

Linear Algebra Done Wrong

Sergei Treil

DEPARTMENT OF MATHEMATICS, BROWN UNIVERSITY

Copyright © Sergei Treil, 2003

Preface

Contents

Preface	iii
Chapter 1. Basic Notions	1
§1. Vector spaces	1
§2. Linear combinations, bases.	4
§3. Linear Transformations. Matrix–vector multiplication	9
§4. Composition of linear transformations and matrix multiplication.	13
§5. Invertible transformations and matrices. Isomorphisms	16
§6. Subspaces.	22
Chapter 2. Systems of linear equations	25
§1. Different faces of linear systems.	25
§2. Solution of a linear system. Echelon and reduced echelon forms	26
§3. Analyzing the pivots.	32
§4. Dimension. Finite-dimensional spaces.	36
§5. General solution of a linear system.	37
§6. Fundamental subspaces of a matrix. Rank.	39
§7. Representation of a linear transformation in arbitrary bases. Change of coordinates formula.	46
Chapter 3. Determinants	51
§1. Introduction.	51
§2. What properties determinant should have.	52
§3. Constructing the determinant.	54

§4. Formal definition. Existence and uniqueness of the determinant.	61
§5. Cofactor expansion.	63
§6. Minors and rank.	67
Chapter 4. Spectral theory (eigenvalues and eigenvectors)	69
§1. Main definitions	70
§2. Diagonalization.	72
Chapter 5. Inner product spaces	79
§1. Inner product in $\mathbb{R}^n$ and $\mathbb{C}^n$. Inner product spaces.	79
§2. Orthogonality	84
§3. Least square solution. Formula for the orthogonal projection	85
§4. Fundamental subspaces revisited.	90
§5. Isometries and unitary operators. Unitary and orthogonal matrices.	93
Chapter 6. Structure of operators in inner product spaces.	97
§1. Schur (upper triangular) representation of an operator.	97
§2. Spectral theorem for self-adjoint and normal operators.	99
§3. Polar and singular value decompositions.	102
§4. Structure of orthogonal matrices	109

Basic Notions

1. Vector spaces

A vector space V is a collection of objects, called vectors, along with two operations, addition of vectors and multiplication by a number (scalar), such that the following properties hold:

The first 4 properties deal with the addition:

1. Commutativity: $\mathbf{v} + \mathbf{w} = \mathbf{w} + \mathbf{v}$ for all $\mathbf{v}, \mathbf{w} \in V$;
2. Associativity: $(\mathbf{u} + \mathbf{v}) + \mathbf{w} = \mathbf{u} + (\mathbf{v} + \mathbf{w})$ for all $\mathbf{u}, \mathbf{v}, \mathbf{w} \in V$;
3. Zero vector: there exists a special vector, denoted by $\mathbf{0}$ such that $\mathbf{v} + \mathbf{0} = \mathbf{v}$ for all $\mathbf{v} \in V$;
4. Additive inverse: For every vector $\mathbf{v} \in V$ there exists a vector $\mathbf{w} \in V$ such that $\mathbf{v} + \mathbf{w} = \mathbf{0}$. Such additive inverse is usually denoted as $-\mathbf{v}$;

The next two properties concern multiplication:

5. Multiplicative identity: $1\mathbf{v} = \mathbf{v}$ for all $\mathbf{v} \in V$;
6. Multiplicative associativity: $(\alpha\beta)\mathbf{v} = \alpha(\beta\mathbf{v})$ for all $\mathbf{v} \in V$ and all scalars α, β ;

And finally, two distributive properties, which connect multiplication and addition:

7. $\alpha(\mathbf{u} + \mathbf{v}) = \alpha\mathbf{u} + \alpha\mathbf{v}$ for all $\mathbf{u}, \mathbf{v} \in V$ and all scalars α ;
8. $(\alpha + \beta)\mathbf{v} = \alpha\mathbf{v} + \beta\mathbf{v}$ for all $\mathbf{v} \in V$ and all scalars α, β .

Remark. It is not hard to show that zero vector $\mathbf{0}$ is unique. It is also easy to show that given $\mathbf{v} \in V$ the inverse vector $-\mathbf{v}$ is unique.

In fact, properties 2 and 3 can be deduced from the properties 5 and 8: they imply that $\mathbf{0} = 0\mathbf{v}$ for any $\mathbf{v} \in V$, and that $-\mathbf{v} = (-1)\mathbf{v}$.

If scalars are usual real numbers, we call the space V a *real* vector space. If scalars are complex numbers, i.e. if we can multiply by complex numbers, we call the space V a *complex* vector space.

Note, that any complex vector space is a real vector space as well (if we can multiply by complex numbers, we can multiply by real numbers), but not the other way around.

It is also possible to consider a situation when the scalars are elements of an arbitrary field $\mathbb{F}$. In this case we say that V is a vector space over the field $\mathbb{F}$. Although many of the construction work for general fields, in this text we consider only real and complex vector spaces, i.e. $\mathbb{F}$ is always either $\mathbb{R}$ or $\mathbb{C}$.

1.1. Examples.

Example. The space $\mathbb{R}^n$ consists of all columns of size n ,

$$\mathbf{v} = \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{pmatrix}$$

whose entries are real numbers. Addition and multiplication are defined entrywise, i.e.

$$\alpha \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{pmatrix} = \begin{pmatrix} \alpha v_1 \\ \alpha v_2 \\ \vdots \\ \alpha v_n \end{pmatrix}, \quad \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{pmatrix} + \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_n \end{pmatrix} = \begin{pmatrix} v_1 + w_1 \\ v_2 + w_2 \\ \vdots \\ v_n + w_n \end{pmatrix}$$

Example. The space $\mathbb{C}^n$ also consists of columns of size n , only the entries now are complex numbers. Addition and multiplication are defined exactly as in the case of $\mathbb{R}^n$, the only difference is that we can now multiply vectors by *complex* numbers, i.e. $\mathbb{C}^n$ is a *complex* vector space.

Example. The space $M_{m \times n}$ of $m \times n$ matrices: the multiplication and addition are defined entrywise. If we allow only real entries (and so only multiplication only by reals), then we have a real vector space; if we allow complex entries and multiplication by complex numbers, we deal with a complex space.

Example. The space $\mathbb{P}_n$ of polynomials of degree at most n , consists of all polynomials of form

$$p(t) = a_0 + a_1 t + a_2 t^2 + \dots + a_n t^n,$$

where t is the independent variable. In the case of real coefficients a_k we have a real space, complex coefficients give us a complex space.

Question: What are zero vectors in each of the above examples?

more to follow...

2. Linear combinations, bases.

Let V be a vector space, and let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p \in V$ be a collection of vectors. A *linear combination* of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p$ is a sum of form

$$\alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \dots + \alpha_p \mathbf{v}_p = \sum_{k=1}^p \alpha_k \mathbf{v}_k.$$

Definition 2.1. A system of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n \in V$ is called a *basis* (for the vector space V) if any vector $\mathbf{v} \in V$ admits a *unique* representation as a linear combination

$$\mathbf{v} = \alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \dots + \alpha_n \mathbf{v}_n = \sum_{k=1}^n \alpha_k \mathbf{v}_k.$$

Another way to say is that the equation $x_1 \mathbf{v}_1 + x_2 \mathbf{v}_2 + \dots + x_n \mathbf{v}_n = \mathbf{v}$ (with unknowns x_k) has a unique solution for arbitrary right side $\mathbf{v}$.

Before discussing any properties of bases¹, let us give few examples, showing that such objects exist, and it makes sense to study them.

Example 2.2. In the first example the space V is $\mathbb{R}^n$. Consider vectors

$$\mathbf{e}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad \mathbf{e}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad \mathbf{e}_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \\ \vdots \\ 0 \end{pmatrix}, \dots, \quad \mathbf{e}_n = \begin{pmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 1 \end{pmatrix},$$

(the vector $\mathbf{e}_k$ has all entries 0 except the entry number k , which is 1). The system of vectors $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ is a basis in $\mathbb{R}^n$. Indeed, any vector

$$\mathbf{v} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} \in \mathbb{R}^n$$

can be represented as the linear combination

$$\mathbf{v} = x_1 \mathbf{e}_1 + x_2 \mathbf{e}_2 + \dots + x_n \mathbf{e}_n = \sum_{k=1}^n x_k \mathbf{e}_k$$

and this representation is unique. The system $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n \in \mathbb{R}^n$ is called *the standard basis* in $\mathbb{R}^n$

¹the plural for the “basis” is *bases*, the same as the plural for “base”

Example 2.3. In this example the space is the space $\mathbb{P}_n$ of the polynomials of degree at most n . Consider vectors (polynomials) $\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n \in \mathbb{P}_n$ defined by

$$\mathbf{e}_0 := 1, \quad \mathbf{e}_1 := t, \quad \mathbf{e}_2 := t^2, \quad \mathbf{e}_3 := t^3, \quad \dots, \quad \mathbf{e}_n := t^n.$$

Clearly, any polynomial $p, p(t) = a_0 + a_1t + a_2t^2 + \dots + a_nt^n$ admits unique representation

$$p = a_0\mathbf{e}_0 + a_1\mathbf{e}_1 + \dots + a_n\mathbf{e}_n.$$

So the system $\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n \in \mathbb{P}_n$ is a basis in $\mathbb{P}_n$. We will call it the standard basis in $\mathbb{P}_n$.

Remark 2.4. If a vector space V has a basis $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$, then any vector $\mathbf{v}$ is uniquely defined by its coefficients in the decomposition $\mathbf{v} = \sum_{k=1}^n \alpha_k \mathbf{v}_k$. So, if we stack the coefficients α_k in a column, we can operate with them as with column vectors, i.e. as with elements of $\mathbb{R}^n$.

The definition of a basis says that any vector admits a unique representation as a linear combination. This statement is in fact two statements, namely that the representation exists and that it is unique. Let us analyze these two statements separately.

If we only take existence we get the following notion

Definition 2.5. A system of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p \in V$ is called a generating system (also a *spanning system*, and also a *complete system*) in V any vector $\mathbf{v} \in V$ admits representation as a linear combination

$$\mathbf{v} = \alpha_1\mathbf{v}_1 + \alpha_2\mathbf{v}_2 + \dots + \alpha_p\mathbf{v}_p = \sum_{k=1}^p \alpha_k \mathbf{v}_k.$$

The only difference with the definition of a basis is that we do not assume that the representation above is unique.

The words *generating*, *spanning* and *complete* here are synonyms. I personally prefer the term *complete*, because of my operator theory background. Generating and spanning are more often used in linear algebra textbooks.

Clearly, any basis is a generating (complete) system. Also, if we have a basis, say $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$, and we add to it several vectors, say $\mathbf{v}_{n+1}, \dots, \mathbf{v}_p$, then the new system will be a generating (complete) system. Indeed, we can represent any vector as a linear combination of the vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$, and just ignore the new ones (by putting corresponding coefficients $\alpha_k = 0$).

Now, let us turn our attention to the uniqueness.

Definition. A linear combination $\alpha_1\mathbf{v}_1 + \alpha_2\mathbf{v}_2 + \dots + \alpha_p\mathbf{v}_p$ is called *trivial* if $\alpha_k = 0 \forall k$.

A trivial linear combination is always (for all choices of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p$) equal $\mathbf{0}$, and that is probably the reason for the name.

Definition. A system of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p \in V$ is called *linearly independent* if only the trivial linear combination ($\sum_{k=1}^p \alpha_k \mathbf{v}_k$ with $\alpha_k = \mathbf{0} \ \forall k$) of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p$ equals $\mathbf{0}$.

In other words, the system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p$ is linearly independent iff the equation $x_1 \mathbf{v}_1 + x_2 \mathbf{v}_2 + \dots + x_p \mathbf{v}_p = \mathbf{0}$ (with unknowns x_k) has only trivial solution $x_1 = x_2 = \dots = x_p = 0$.

If a system is not linearly independent, it is called *linearly dependent*. By negating the definition of linear independence, we get the following

Definition. A system of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p$ is called *linearly dependent* if $\mathbf{0}$ can be represented as a nontrivial linear combination, $\mathbf{0} = \sum_{k=1}^p \alpha_k \mathbf{v}_k$. Non-trivial here means that at least one of the coefficient α_k is non-zero. This can be (and usually is) written as $\sum_{k=1}^p |\alpha_k| \neq 0$.

So, restating the definition we can say, that a system is linearly dependent if and only if there exist scalars $\alpha_1, \alpha_2, \dots, \alpha_p$, $\sum_{k=1}^p |\alpha_k| \neq 0$ such that

$$\sum_{k=1}^p \alpha_k \mathbf{v}_k = \mathbf{0}.$$

An alternative definition (in terms of equations) is that a system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p$ is linearly dependent iff the equation

$$x_1 \mathbf{v}_1 + x_2 \mathbf{v}_2 + \dots + x_p \mathbf{v}_p = \mathbf{0}$$

(with unknowns x_k) has a non-trivial solution. Non-trivial here again means that at least one of x_k is different from 0, and it can be written as $\sum_{k=1}^p |x_k| \neq 0$.

The following proposition gives an alternative description of linearly dependent systems.

Proposition 2.6. *A system of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p \in V$ is linearly dependent if and only if one of the vectors $\mathbf{v}_k$ can be represented as a linear combination of the other vectors,*

$$(2.1) \quad \mathbf{v}_k = \sum_{\substack{j=1 \\ j \neq k}}^p \beta_j \mathbf{v}_j.$$

Proof. Suppose the system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p$ is linearly dependent. Then there exist scalars α_k , $\sum_{k=1}^p |\alpha_k| \neq 0$ such that

$$\alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \dots + \alpha_p \mathbf{v}_p = \mathbf{0}.$$

Let k be the index such that $\alpha_k \neq 0$. Then, moving all terms except $\alpha_k \mathbf{v}_k$ to the right side we get

$$\alpha_k \mathbf{v}_k = - \sum_{\substack{j=1 \\ j \neq k}}^p \alpha_j \mathbf{v}_j.$$

Dividing both sides by α_k we get (2.1) with $\beta_j = -\alpha_j/\alpha_k$.

On the other hand, if (2.1) holds, $\mathbf{0}$ can be represented as a non-trivial linear combination

$$\mathbf{v}_k - \sum_{\substack{j=1 \\ j \neq k}}^p \beta_j \mathbf{v}_j = \mathbf{0}.$$

□

Obviously, any basis is a linearly independent system. Indeed, if a system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ is a basis, $\mathbf{0}$ admits a unique representation

$$\mathbf{0} = \alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \dots + \alpha_n \mathbf{v}_n = \sum_{k=1}^n \alpha_k \mathbf{v}_k.$$

Since the trivial linear combination always gives $\mathbf{0}$, the trivial linear combination must be the *only one* giving $\mathbf{0}$.

So, as we already discussed, if a system is a basis it is a complete (generating) and linearly independent system. The following proposition shows that the converse implication is also true.

Proposition 2.7. *A system of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n \in V$ is a basis if and only if it is linearly independent and complete (generating).*

Proof. We already know that a basis is always linearly independent and complete, so in one direction the proposition is already proved.

Let us prove the other direction. Suppose a system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ is linearly independent and complete. Take an arbitrary vector $\mathbf{v} \in V$. Since the system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ is linearly complete (generating), $\mathbf{v}$ can be represented as

$$\mathbf{v} = \alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \dots + \alpha_n \mathbf{v}_n = \sum_{k=1}^n \alpha_k \mathbf{v}_k.$$

We only need to show that this representation is unique.

Suppose $\mathbf{v}$ admits another representation

$$\mathbf{v} = \sum_{k=1}^n \tilde{\alpha}_k \mathbf{v}_k.$$

In many textbooks a basis is defined as a complete and linearly independent system. By Proposition 2.7 this definition is equivalent to ours

Then

$$\sum_{k=1}^n (\alpha_k - \tilde{\alpha}_k) \mathbf{v}_k = \sum_{k=1}^n \alpha_k \mathbf{v}_k - \sum_{k=1}^n \tilde{\alpha}_k \mathbf{v}_k = \mathbf{v} - \mathbf{v} = \mathbf{0}.$$

Since the system is linearly independent, $\alpha_k - \tilde{\alpha}_k = 0 \ \forall k$, and thus the representation $\mathbf{v} = \alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \dots + \alpha_n \mathbf{v}_n$ is unique. $\square$

Remark. In many textbooks a basis is defined as a complete and linearly independent system (by Proposition 2.7 this definition is equivalent to ours). Although this definition is more common, than one presented in this text, I prefer the later. It emphasizes the main property of a basis, namely that any vector admits a unique representation as a linear combination.

Proposition 2.8. *Any (finite) generating system contains a basis.*

Proof. Suppose $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p \in V$ is a generating (complete) set. If it is linearly independent, it is a basis, and we are done.

Suppose it is not linearly independent, i.e. it is linearly dependent. Then there exists a vector $\mathbf{v}_k$ which can be represented as a linear combination of the vectors $\mathbf{v}_j, j \neq k$.

Since $\mathbf{v}_k$ can be represented as a linear combination of vectors $\mathbf{v}_j, j \neq k$, any linear combination of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p$ can be represented as a linear combination of the same vectors without $\mathbf{v}_k$ (i.e. the vectors $\mathbf{v}_j, 1 \leq j \leq p, j \neq k$). So, if we delete the vector $\mathbf{v}_k$, the new system will still be a complete one.

If the new system is linearly independent, we are done. If not, we repeat the procedure.

Repeating this procedure finitely many times we arrive to a linearly independent and complete system, because otherwise we delete all vectors and end up with an empty set.

So, any finite complete (generating) set contains a complete linearly independent subset, i.e. a basis. $\square$

3. Linear Transformations. Matrix–vector multiplication

A *transformation* T from a set X to a set Y is a rule that for each argument (input) $x \in X$ assigns a value (output) $y = T(x) \in Y$.

The set X is called *domain* of T , and the set Y is called *codomain* of T .

We write $T : X \rightarrow Y$ to say that T is a transformation with the domain X and the codomain Y .

The words “transformation”, “transformation”, “mapping”, “map”, “operator”, “function” all denote the same object.

Definition. Let V, W be vector spaces. A transformation $T : V \rightarrow W$ is called linear if

1. $T(\mathbf{u} + \mathbf{v}) = T(\mathbf{u}) + T(\mathbf{v}) \quad \forall \mathbf{u}, \mathbf{v} \in V$;
2. $T(\alpha \mathbf{v}) = \alpha T(\mathbf{v})$ for all $\mathbf{v} \in V$ and for all scalars α .

Properties 1 and 2 together are equivalent to the following one:

$$T(\alpha \mathbf{u} + \beta \mathbf{v}) = \alpha T(\mathbf{u}) + \beta T(\mathbf{v}) \quad \text{for all } \mathbf{u}, \mathbf{v} \in V \quad \text{and for all scalars } \alpha, \beta.$$

3.1. Examples. You dealt with linear transformation before, may be without even suspecting it, as the example below shows.

Example. Differentiation: Let $V = \mathbb{P}_n$ (the set of polynomials of degree at most n), $W = \mathbb{P}_{n-1}$, and let $T : \mathbb{P}_n \rightarrow \mathbb{P}_{n-1}$ be the differentiation operator,

$$T(p) := p' \quad \forall p \in \mathbb{P}_n.$$

Since $(f + g)' = f' + g'$ and $(\alpha f)' = \alpha f'$, this is a linear transformation.

Example. Rotation: in this example $V = W = \mathbb{R}^2$ (the usual coordinate plane), and a transformation $T_\gamma : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ takes a vector in $\mathbb{R}^2$ and rotates it counterclockwise by γ radians. Since T_γ rotates the plane as a whole, it rotates as a whole the parallelogram used to define a sum of two vectors (parallelogram law). Therefore the property 1 of linear transformation holds. It is also easy to see that the property 2 is also true.

Example. Reflection: in this example again $V = W = \mathbb{R}^2$, and the transformation $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the reflection in the first coordinate axis, see the fig. It can also be shown geometrically, that this transformation is linear, but we will use another way to show that.

Namely, it is easy to write a formula for T ,

$$T\left(\begin{pmatrix} x_1 \\ x_2 \end{pmatrix}\right) = \begin{pmatrix} x_1 \\ -x_2 \end{pmatrix}$$

and from this formula it is easy to check that the transformation is linear.

Example. Let us investigate transformations $T : \mathbb{R} \rightarrow \mathbb{R}$. Any such transformation is given by the formula

$$T(x) = ax \quad \text{where } a = T(1).$$

Indeed,

$$T(x) = T(x \times 1) = xT(1) = xa = ax.$$

So, any linear transformation of $\mathbb{R}$ is just a multiplication by a constant.

3.2. Linear transformations $\mathbb{R}^n \rightarrow \mathbb{R}^m$. Matrix-column multiplication. It turns out that a linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ also can be represented as a multiplication, but not by a number, but by a matrix.

Let us see how. Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear transformation. What information do we need to compute $T(\mathbf{x})$ for all vectors $\mathbf{x} \in \mathbb{R}^n$? My claim is that it is sufficient how T acts on the standard basis $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ of $\mathbb{R}^n$. Namely, it is sufficient to know n vectors in $\mathbb{R}^m$ (i.e. of size m),

$$\mathbf{a}_1 = T(\mathbf{e}_1), \quad \mathbf{a}_2 := T(\mathbf{e}_2), \quad \dots, \quad \mathbf{a}_n := T(\mathbf{e}_n).$$

Indeed, let

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix}.$$

Then $\mathbf{x} = x_1\mathbf{e}_1 + x_2\mathbf{e}_2 + \dots + x_n\mathbf{e}_n = \sum_{k=1}^n x_k\mathbf{e}_k$

$$T(\mathbf{x}) = T\left(\sum_{k=1}^n x_k\mathbf{e}_k\right) = \sum_{k=1}^n T(x_k\mathbf{e}_k) = \sum_{k=1}^n x_k T(\mathbf{e}_k) = \sum_{k=1}^n x_k\mathbf{a}_k.$$

So, if we join the vectors (columns) $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n$ together in a matrix $A = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n]$ ($\mathbf{a}_k$ being the k th column of A $k = 1, 2, \dots, n$), this matrix contains all the information about T .

Let us show how one should define the product of a matrix and a vector (column) to represent the transformation T as a product, $T(\mathbf{x}) = A\mathbf{x}$. Let

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,n} \\ a_{2,1} & a_{2,2} & \dots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \dots & a_{m,n} \end{pmatrix}$$

Recall, that the column number k of A is the vector $\mathbf{a}_k$, i.e.

$$\mathbf{a}_k = \begin{pmatrix} a_{1,k} \\ a_{2,k} \\ \vdots \\ a_{m,k} \end{pmatrix}$$

Then if we want $A\mathbf{x} = T(\mathbf{x})$ we get

$$A\mathbf{x} = \sum_{k=1}^n x_k \mathbf{a}_k = x_1 \begin{pmatrix} a_{1,1} \\ a_{2,1} \\ \vdots \\ a_{m,1} \end{pmatrix} + x_2 \begin{pmatrix} a_{1,2} \\ a_{2,2} \\ \vdots \\ a_{m,2} \end{pmatrix} + \dots + x_n \begin{pmatrix} a_{1,n} \\ a_{2,n} \\ \vdots \\ a_{m,n} \end{pmatrix}$$

So, the matrix–vector multiplication should be performed by the following *column by coordinate* rule:

multiply each column of the matrix by the corresponding coordinate of the vector.

Example.

$$\begin{pmatrix} 1 & 2 & 3 \\ 3 & 2 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} = 1 \begin{pmatrix} 1 \\ 3 \end{pmatrix} + 2 \begin{pmatrix} 2 \\ 2 \end{pmatrix} + 3 \begin{pmatrix} 3 \\ 1 \end{pmatrix} = \begin{pmatrix} 15 \\ 10 \end{pmatrix}$$

The “column by coordinate” rule is very well adapted for parallel computing. It will be also very important in different theoretical constructions later.

However, when doing computations manually, it is more convenient to compute the result one entry at a time. This can be expressed as the following *row by column* rule

To get the entry number k of the result, one need to multiply row number k of the matrix by the vector, that is, if $A\mathbf{x} = \mathbf{y}$, then $y_k = \sum_{j=1}^n a_{k,j}x_j$, $k = 1, 2, \dots, m$;

here x_j and y_k are coordinates of the vectors $\mathbf{x}$ and $\mathbf{y}$ respectively, and $a_{j,k}$ are the entries of the matrix A .

Example.

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} = \begin{pmatrix} 1 \cdot 1 + 2 \cdot 2 + 3 \cdot 3 \\ 4 \cdot 1 + 5 \cdot 2 + 6 \cdot 3 \end{pmatrix} = \begin{pmatrix} 15 \\ 32 \end{pmatrix}$$

3.3. Conclusions.

- To get the matrix of a linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ one needs to join the vectors $\mathbf{a}_k = T\mathbf{e}_k$ (where $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ is the standard basis in $\mathbb{R}^n$) into a matrix: k th column of the matrix is $\mathbf{a}_k$, $k = 1, 2, \dots, n$.
- If the matrix A of the linear transformation T is known, then $T(\mathbf{x})$ can be found by the matrix–vector multiplication, $T(\mathbf{x}) = A\mathbf{x}$. To

perform matrix–vector multiplication one can use either “column by coordinate” or “row by column” rule.

The latter seems more appropriate for manual computations. The former is well adapted for parallel computers, and will be used in different theoretical constructions.

For a linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$, its matrix is usually denoted as $[T]$. However, very often people do not distinguish between a linear transformation and its matrix, and use the same symbol for both. When it does not lead to a confusion, we will also use the same symbol for a transformation and its matrix.

The notation $T\mathbf{v}$ is often used instead of $T(\mathbf{v})$.

Since a linear transformation is essentially a multiplication, the notation $T\mathbf{v}$ is often used instead of $T(\mathbf{v})$. We will also use this notation. Note that usual order of algebraic operations apply, i.e. $T\mathbf{v} + \mathbf{u}$ means $T(\mathbf{v}) + \mathbf{u}$, not $T(\mathbf{v} + \mathbf{u})$.

In the matrix vector multiplication using the “row by column” rule be sure that you have the same number of entries in the row and in the column. The entries in the row and in the column should end simultaneously: if not, the multiplication is not defined.

Remark. In the matrix–vector multiplication $A\mathbf{x}$ the number of columns of the matrix A must coincide with the size of the vector $\mathbf{x}$, i.e. a vector in $\mathbb{R}^n$ can only be multiplied by an $m \times n$ matrix.

It makes sense, since an $m \times n$ matrix defines a linear transformation $\mathbb{R}^n \rightarrow \mathbb{R}^m$, so vector $\mathbf{x}$ must belong to $\mathbb{R}^n$.

The easiest way to remember this is to remember that if performing multiplication you run out of some elements faster, then the multiplication is not defined. For example, if using the “row by column” rule you run out of row entries, but still have some unused entries in the vector, the multiplication is not defined. It is also not defined if you run out of vector’s entries, but still have unused entries in the column.

4. Composition of linear transformations and matrix multiplication.

4.1. Definition of the matrix multiplication. Knowing matrix–vector multiplication, one can easily guess what is the natural way to define the product AB of two matrices: Let us multiply by A each column of B (matrix–vector multiplication) and join the resulting column-vectors into a matrix. Formally,

if $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_r$ are the columns of B , then $A\mathbf{b}_1, A\mathbf{b}_2, \dots, A\mathbf{b}_r$ are the columns of the matrix AB .

Recalling the *row by column* rule for the matrix–vector multiplication we get the following *row by column rule for the matrices*

the entry $(AB)_{j,k}$ (the entry in the row j and column k) of the product AB is defined by

$$(AB)_{j,k} = (\text{row } \#j \text{ of } A) \cdot (\text{column } \#k \text{ of } B)$$

Formally it can be rewritten as

$$(AB)_{j,k} = \sum_l a_{j,l} b_{l,k},$$

if $a_{j,k}$ and $b_{j,k}$ are entries of the matrices A and B respectively.

I intentionally did not speak about sizes of the matrices A and B , but if we recall the row by column rule for the matrix–vector multiplication, we can see that in order for the multiplication to be defined, the size of a row of A should be equal to the size of a column of B .

In other words the product AB is defined if and only if A is an $m \times n$ and B is $n \times r$ matrix.

4.2. Motivation: composition of linear transformations. One can ask yourself here: Why are we using such complicated rule of multiplication? Why don't we just multiply matrices entrywise?

And the answer is, that the multiplication, as it is defined above, arises naturally from the composition of linear transformations.

Suppose we have two linear transformations, $T_1 : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $T_2 : \mathbb{R}^r \rightarrow \mathbb{R}^n$. Define the composition $T = T_1 \circ T_2$ of the transformations T_1, T_2 as

$$T(\mathbf{x}) = T_1(T_2(\mathbf{x})) \quad \forall \mathbf{x} \in \mathbb{R}^r.$$

Note that $T_1(\mathbf{x}) \in \mathbb{R}^m$. Since $T_1 : \mathbb{R}^n \rightarrow \mathbb{R}^m$ the expression $T_1(T_2(\mathbf{x}))$ is well defined and the result belongs to $\mathbb{R}^m$. So, $T : \mathbb{R}^r \rightarrow \mathbb{R}^m$.

It is easy to show that T is a linear transformation (exercise), so it is defined by an $m \times r$ matrix. How one can find this matrix, knowing matrices of T_1 and T_2 ?

We will usually identify a linear transformation and its matrix, but in the next few paragraphs we will distinguish them

Let A be the matrix of T_1 and B be the matrix of T_2 . As we discussed in the previous section, the columns of T are vectors $T(\mathbf{e}_1), T(\mathbf{e}_2), \dots, T(\mathbf{e}_r)$, where $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_r$ is the standard basis in $\mathbb{R}^r$. For $k = 1, 2, \dots, r$ we have

$$T(\mathbf{e}_k) = T_1(T_2(\mathbf{e}_k)) = T_1(B\mathbf{e}_k) = T_1(\mathbf{b}_k) = A\mathbf{b}_k$$

(operators T_2 and T_1 are simply the multiplication by B and A respectively).

So, the columns of the matrix of T are $A\mathbf{b}_1, A\mathbf{b}_2, \dots, A\mathbf{b}_r$, and that is exactly how the matrix AB was defined!

Let us return to identifying again a linear transformation with its matrix. Since the matrix multiplication agrees with the composition, we can (and will) write T_1T_2 instead of $T_1 \circ T_2$ and $T_1T_2\mathbf{x}$ instead of $T_1(T_2(\mathbf{x}))$.

Note: order of transformations!

Note that in the composition T_1T_2 the transformation T_2 is applied first! The way to remember this is to see that in $T_1T_2\mathbf{x}$ the transformation T_2 meets $\mathbf{x}$ first.

Remark. There is another way of checking the dimensions of matrices in a product, different from the row by column rule: for a composition T_1T_2 to be defined it is necessary that $T_2\mathbf{x}$ belongs to the domain of T_1 . If T_2 acts from some space, say $\mathbb{R}^r$ to $\mathbb{R}^n$, then T_1 must act from $\mathbb{R}^n$ to some space, say $\mathbb{R}^m$. So, in order for T_1T_2 to be defined the matrices of T_1 and T_2 should be of sizes $m \times n$ and $n \times r$ respectively—the same condition as obtained from the *row by column* rule.

Example. Let $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the reflection in the line $x_1 = 3x_2$. It is a linear transformation, so let us find its matrix. To find the matrix, we need to compute $T\mathbf{e}_1$ and $T\mathbf{e}_2$. However, the direct computation of $T\mathbf{e}_1$ and $T\mathbf{e}_2$ involves significantly more trigonometry than a sane person is willing to remember.

An easier way to find the matrix of T is to represent it as a composition of simple linear transformation. Namely, let γ be the angle between the x_1 axis and the line $x_1 = 3x_2$, and let T_0 be the reflection in the x_1 -axis. Then to get the reflection T we can first rotate the plane by the angle $-\gamma$, moving the line $x_1 = 3x_2$ to the x_1 -axis, then reflect everything in the x_1 axis, and then rotate the plane by γ , taking everything back. Formally it can be written as

$$T = R_\gamma T_0 R_{-\gamma}$$

(note the order of terms!), where R_γ is the rotation by γ . The matrix of T_0 is easy to compute,

$$T_0 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$

the rotation matrices are known

$$R_\gamma = \begin{pmatrix} \cos \gamma & -\sin \gamma \\ \sin \gamma & \cos \gamma \end{pmatrix},$$

$$R_{-\gamma} = \begin{pmatrix} \cos(-\gamma) & -\sin(-\gamma) \\ \sin(-\gamma) & \cos(-\gamma) \end{pmatrix} = \begin{pmatrix} \cos \gamma & \sin \gamma \\ -\sin \gamma & \cos \gamma \end{pmatrix}$$

To compute $\sin \gamma$ and $\cos \gamma$ take a vector in the line $x_1 = 3x_2$, say a vector $(3, 1)^T$. Then

$$\cos \gamma = \frac{\text{first coordinate}}{\text{length}} = \frac{3}{\sqrt{3^2 + 1^2}} = \frac{3}{\sqrt{10}}$$

and similarly

$$\sin \gamma = \frac{\text{second coordinate}}{\text{length}} = \frac{1}{\sqrt{3^2 + 1^2}} = \frac{1}{\sqrt{10}}$$

Gathering everything together we get

$$\begin{aligned} T &= R_\gamma T_0 R_{-\gamma} = \frac{1}{\sqrt{10}} \begin{pmatrix} 3 & -1 \\ 1 & 3 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \frac{1}{\sqrt{10}} \begin{pmatrix} 3 & 1 \\ -1 & 3 \end{pmatrix} \\ &= \frac{1}{10} \begin{pmatrix} 3 & -1 \\ 1 & 3 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} 3 & 1 \\ -1 & 3 \end{pmatrix} \end{aligned}$$

It remains only to perform matrix multiplication here to get the final result. $\square$

4.3. Properties of matrix multiplication. The matrix multiplication enjoys a lot of properties, familiar to us from the high school algebra:

1. Associativity: $A(BC) = (AB)C$, provided that either left or right side is well defined;
2. Distributivity: $A(B + C) = AB + AC$, $(A + B)C = AC + BC$, provided either left or right side of each equation is well defined;
3. One can take scalar multiplies out: $A(\alpha B) = \alpha AB$.

These properties are easy to prove. One should prove the corresponding properties for linear transformations, and they almost trivially follow from the definitions. And the properties of linear transformations imply the properties for the matrix multiplication.

The new twist here is that the commutativity fails:

matrix multiplication is non-commutative, i.e. generally for matrices $AB \neq BA$.

One can see easily it would be unreasonable to expect the commutativity of matrix multiplication. Indeed, let A and B be matrices of sizes $m \times n$ and $n \times r$ respectively. Then the product AB is well defined, but if $m \neq r$, BA is not defined.

Even when both products are well defined, for example, when A and B are $n \times n$ (square) matrices, the multiplication is still non-commutative. If we just pick the matrices A and B at random, the chances are that $AB \neq BA$: we have to be very lucky to get $AB = BA$.

4.4. Transposed matrices and multiplication. Given a matrix A , the *transposed* matrix A^T is defined by transforming the rows of A into the columns. For example

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

So, the columns of A^T are the rows of A and vice versa, the rows of A^T are the columns of A .

The formal definition is as follows: $(A^T)_{j,k} = (A)_{k,j}$ meaning that the entry of A^T in the row number j and column number k equals the entry of A in the row number k and column number j .

The transposed matrix has a very nice interpretation in terms of linear transformations, namely it gives the so-called *adjoint* transformation. We will study this in detail later, but for now transposition will be just a useful formal operation.

One of the first uses is that we can write a column vector $\mathbf{x} \in \mathbb{R}^n$ as $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$. If we put the column vertically, it will use significantly more space.

A simple analysis of the row by columns rule shows that

$$(AB)^T = B^T A^T,$$

i.e. when you take the transpose of the product, you change the order of the terms.

5. Invertible transformations and matrices. Isomorphisms

5.1. Identity transformation and identity matrix. Among all linear transformations, there is a special one, the identity transformation (operator) I , $I\mathbf{x} = \mathbf{x}$, $\forall \mathbf{x}$.

To be precise, there are infinitely many identity transformations: for any vector space V , there is the identity transformation $I = I_V : V \rightarrow V$, $I_V \mathbf{x} = \mathbf{x}$, $\forall \mathbf{x} \in V$. However, when it is does not lead to the confusion

we will use the same symbol I for all identity operators (transformations). We will use the notation I_V only we want to emphasize in what space the transformation is acting.

Clearly, if $I : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the identity transformation in $\mathbb{R}^n$, its matrix is an $n \times n$ matrix

$$I = I_n = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix}$$

(1 on the main diagonal and 0 everywhere else). When we want to emphasize the size of the matrix, we use the notation I_n ; otherwise we just use I .

Clearly, for an arbitrary linear transformation A , the equalities

$$AI = A, \quad IA = A$$

hold (whenever the product is defined).

5.2. Invertible transformations.

Definition. Let $A : V \rightarrow W$ be a linear transformation. We say that the transformation A is *left invertible* if there exist a transformation $B : W \rightarrow V$ such that

$$BA = I \quad (I = I_V \text{ here}).$$

The transformation A is called *right invertible* if there exists a linear transformation $C : W \rightarrow V$ such that

$$AC = I \quad (\text{here } I = I_W).$$

The transformations B and C are called *left* and *right inverses* of A . Note, that we did not assume the uniqueness of B or C here, and generally left and right inverses are not unique.

Definition. A linear transformation $A : V \rightarrow W$ is called *invertible* if it is both right and left invertible.

Theorem 5.1. *If a linear transformation $A : V \rightarrow W$ is invertible, then its left and right inverses B and C are unique and coincide.*

Corollary. *A transformation $A : V \rightarrow W$ is invertible if and only if there exists a unique linear transformation (denoted A^{-1}), $A^{-1} : W \rightarrow V$ such that*

$$A^{-1}A = I_V, \quad AA^{-1} = I_W.$$

The transformation A^{-1} is called the *inverse* of A .

Often, the symbol E is used in Linear Algebra textbooks for the identity matrix. I prefer I , since it is used in the operator theory

very often this property is used as the definition of an invertible transformation

Proof of Theorem 5.1. Let $BA = I$ and $AC = I$. Then

$$BAC = B(AC) = BI = B.$$

On the other hand

$$BAC = (BA)C = IC = C,$$

and therefore $B = C$.

Suppose for some transformation B_1 we have $B_1A = I$. Repeating the above reasoning with B_1 instead of B we get $B_1 = C$. Therefore the left inverse B is unique. The uniqueness of C is proved similarly. $\square$

Definition. A matrix is called *invertible* (resp. *left invertible*, *right invertible*) if the corresponding linear transformation is invertible (resp. left invertible, right invertible).

Theorem 5.1 asserts that a matrix A is invertible if there exists a unique matrix A^{-1} such that $A^{-1}A = I$, $AA^{-1} = I$. The matrix A^{-1} is called (surprise) the *inverse* of A .

Examples.

1. The identity transformation (matrix) is invertible, $I^{-1} = I$;
2. The rotation R_γ

$$R_\gamma = \begin{pmatrix} \cos \gamma & -\sin \gamma \\ \sin \gamma & \cos \gamma \end{pmatrix}$$

is invertible, and the inverse is given by $(R_\gamma)^{-1} = R_{-\gamma}$. This equality is clear from the geometric description of R_γ , and it also can be checked by the matrix multiplication;

3. The column $(1, 1)^T$ is left invertible but not right invertible. One of the possible left inverses in the row $(1/2, 1/2)$.

To show that this matrix is not right invertible, we just notice that there are more than one left inverse. **Exercise:** describe all left inverses of this matrix.

4. The row $(1, 1)$ is right invertible, but not left invertible. The column $(1/2, 1/2)^T$ is a possible right inverse.

Remark 5.2. An invertible matrix *must* be square ($n \times n$). Moreover, if a square matrix A has either left or right inverse, it is invertible. So it is sufficient to check only one of the identities $AA^{-1} = I$, $A^{-1}A = I$.

An invertible matrix must be square. (To be proved later)

This fact will be proved later. Until we prove this fact, we will not use it. I presented it here only to stop trying wrong directions.

5.2.1. Properties of the inverse transformation.

Inverse of a product:
 $(AB)^{-1} = B^{-1}A^{-1}$.
 Note the change of
 order

Theorem 5.3 (Inverse of the product). *If linear transformations A and B are invertible (and such that the product AB is defined), then the product AB is invertible and*

$$(AB)^{-1} = B^{-1}A^{-1}$$

(note the change of the order!)

Proof. Direct computation shows:

$$(AB)(B^{-1}A^{-1}) = A(BB^{-1})A^{-1} = AIA^{-1} = AA^{-1} = I$$

and similarly

$$(B^{-1}A^{-1})(AB) = B^{-1}(A^{-1}A)B = B^{-1}IB = B^{-1}B = I$$

□

Remark 5.4. The invertibility of the product AB does not imply the invertibility of the factors A and B (can you think of an example?). However, if one of the factors (either A or B) and the product AB are invertible, then the second factor is also invertible.

We leave the proof of this fact as an exercise.

Theorem 5.5 (Inverse of A^T). *If a matrix A is invertible, then A^T is also invertible and*

$$(A^T)^{-1} = (A^{-1})^T$$

Proof. Using $(AB)^T = B^T A^T$ we get

$$(A^{-1})^T A^T = (AA^{-1})^T = I^T = I,$$

and similarly

$$A^T (A^{-1})^T = (A^{-1}A)^T = I^T = I.$$

□

And finally, if A is invertible, then A^{-1} is also invertible, $(A^{-1})^{-1} = A$. So, let us summarize main properties of the inverse:

1. If A is invertible, then A^{-1} is also invertible, $(A^{-1})^{-1} = A$;
2. If A and B are invertible and the product AB is defined, then AB is invertible and $(AB)^{-1} = B^{-1}A^{-1}$.
3. If A is invertible, then A^T is also invertible and $(A^T)^{-1} = (A^{-1})^T$.

5.3. Isomorphism. Isomorphic spaces. An invertible linear transformation $A : V \rightarrow W$ is called an *isomorphism*. We did not introduce anything new here, it is just another shorter name for the object we already studied.

Two vector spaces V and W are called *isomorphic* (denoted $V \cong W$) if there is an isomorphism $A : V \rightarrow W$.

Isomorphic spaces can be considered as different representation of the same space, meaning that all properties, constructions involving vector space operations are preserved under isomorphism.

The theorem below illustrates this statement.

Theorem 5.6. *Let $A : V \rightarrow W$ be an isomorphism, and let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ be a basis in V . Then the system $A\mathbf{v}_1, A\mathbf{v}_2, \dots, A\mathbf{v}_n$ is a basis in W .*

We leave the proof of the theorem as an exercise.

Remark. In the above theorem one can replace “basis” by “linearly independent”, or “generating”, or “linearly dependent”—all these properties are preserved under isomorphisms.

Remark. If A is an isomorphism, then so is A^{-1} . Therefore in the above theorem we can state that $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ is a basis if and only if $A\mathbf{v}_1, A\mathbf{v}_2, \dots, A\mathbf{v}_n$ is a basis.

The inverse to the Theorem 5.6 is also true

Theorem 5.7. *Let $A : V \rightarrow W$ be a linear map, and let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ and $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n$ be bases in V and W respectively. If $A\mathbf{v}_k = \mathbf{w}_k$, $k = 1, 2, \dots, n$, then A is an isomorphism.*

Proof. Define the inverse transformation A^{-1} by $A^{-1}\mathbf{w}_k = \mathbf{v}_k$, $k = 1, 2, \dots, n$ (as we know, a linear transformation is defined by its values on a basis). $\square$

Examples.

1. Let $A : \mathbb{R}^{n+1} \rightarrow \mathbb{P}_n$ ($\mathbb{P}_n$ is the set of polynomials $\sum_{k=0}^n a_k t^k$ of degree at most n) is defined by

$$A\mathbf{e}_1 = 1, A\mathbf{e}_2 = t, \dots, A\mathbf{e}_n = t^{n-1}, A\mathbf{e}_{n+1} = t^n$$

By Theorem 5.7 A is an isomorphism, so $\mathbb{P}_n \cong \mathbb{R}^{n+1}$.

2. Let V be a (real) vector space with a basis $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$. Define transformation $A : \mathbb{R}^n \rightarrow V$ by

$$A\mathbf{e}_k = \mathbf{v}_k, \quad k = 1, 2, \dots, n,$$

where $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ is the standard basis in $\mathbb{R}^n$. Again by Theorem 5.7 A is an isomorphism, so $V \cong \mathbb{R}^n$.

Any real vector space with a basis is isomorphic to $\mathbb{R}^n$

3. $M_{2 \times 3} \cong \mathbb{R}^6$;
4. More generally, $M_{m \times n} \cong \mathbb{R}^{m \cdot n}$

5.4. Invertibility and equations.

Theorem 5.8. *Let $A : V \rightarrow W$ be a linear transformation. Then A is invertible if and only if for any right side $\mathbf{b} \in W$ the equation* Doesn't this remind you of a basis?

$$A\mathbf{x} = \mathbf{b}$$

has a unique solution $\mathbf{x} \in V$.

Proof. Suppose A is invertible. Then $\mathbf{x} = A^{-1}\mathbf{b}$ solves the equation $A\mathbf{x} = \mathbf{b}$. To show that the solution is unique, suppose that for some other vector $\mathbf{x}_1 \in V$

$$A\mathbf{x}_1 = \mathbf{b}$$

Multiplying this identity by A^{-1} from the left we get

$$A^{-1}A\mathbf{x} = A^{-1}\mathbf{b},$$

and therefore $\mathbf{x}_1 = A^{-1}\mathbf{b} = \mathbf{x}$. Note, that both identities, $AA^{-1} = I$ and $A^{-1}A = I$ were used here.

Let us now suppose that the equation $A\mathbf{x} = \mathbf{b}$ has a unique solution $\mathbf{x}$ for any $\mathbf{b} \in W$. Let us use symbol $\mathbf{y}$ instead of $\mathbf{b}$. We know that given $\mathbf{y} \in W$ the equation

$$A\mathbf{x} = \mathbf{y}$$

has a unique solution $\mathbf{x} \in V$. Let us call this solution $B(\mathbf{y})$.

Let us check that B is a linear transformation. We need to show that $B(\alpha\mathbf{y}_1 + \beta\mathbf{y}_2) = \alpha B(\mathbf{y}_1) + \beta B(\mathbf{y}_2)$. Let $\mathbf{x}_k := B(\mathbf{y}_k)$, $k = 1, 2$, i.e. $A\mathbf{x}_k = \mathbf{y}_k$, $k = 1, 2$. Then

$$A(\alpha\mathbf{x}_1 + \beta\mathbf{x}_2) = \alpha A\mathbf{x}_1 + \beta A\mathbf{x}_2 = \alpha\mathbf{y}_1 + \beta\mathbf{y}_2,$$

which means

$$B(\alpha\mathbf{y}_1 + \beta\mathbf{y}_2) = \alpha B(\mathbf{y}_1) + \beta B(\mathbf{y}_2).$$

□

Recalling the definition of a basis we get the following corollary of Theorem 5.7.

Corollary 5.9. *An $m \times n$ matrix is invertible if and only if its columns form a basis in $\mathbb{R}^m$.*

6. Subspaces.

A *subspace* of a vector space V is a subset $V_0 \subset V$ of V which is closed under the vector addition and multiplication by scalars, i.e.

1. If $\mathbf{v} \in V_0$ then $\alpha\mathbf{v} \in V_0$ for all scalars α ;
2. For any $\mathbf{u}, \mathbf{v} \in V_0$ the sum $\mathbf{u} + \mathbf{v} \in V_0$;

Again, the conditions 1 and 2 can be replaced by the following one:

$$\alpha\mathbf{u} + \beta\mathbf{v} \in V_0 \quad \text{for all } \mathbf{u}, \mathbf{v} \in V_0, \text{ and for all scalars } \alpha, \beta.$$

Note, that a subspace $V_0 \subset V$ with the operations (vector addition and multiplication by scalars) inherited from V is a vector space. Indeed, because all operations are inherited from the vector space V they must satisfy all eight axioms of the vector space. The only thing that could possibly go wrong, is that the result of some operation does not belong to V_0 . But the definition of subspace prohibit this!

Now let us consider some examples:

1. *Trivial* subspaces of a space V , namely V itself and $\{\mathbf{0}\}$ (the subspace consisting only of zero vector). Note, that the empty set $\emptyset$ is not a vector space, since it does not contain a zero vector, so it is not a subspace.

With each linear transformation $A : V \rightarrow W$ we can associate the following two subspaces:

2. The *null space*, or *kernel* of A , which is denoted as $\text{Null } A$ or $\text{Ker } A$ and consists of all vectors $\mathbf{v} \in V$ such that $A\mathbf{v} = \mathbf{0}$
3. The range $\text{Ran } A$ is defined as the set of all vectors $\mathbf{w} \in W$ which can be represented as $\mathbf{w} = A\mathbf{v}$ for some $\mathbf{v} \in V$.

If A is a matrix, i.e. $A : \mathbb{R}^m \rightarrow \mathbb{R}^n$, then recalling *column by coordinate* rule of the matrix–vector multiplication, we can see that any vector $\mathbf{w} \in \text{Ran } A$ can be represented as a linear combination of columns of the matrix A . That explains why the term column space (and notation $\text{Col } A$) is often used for the range of the matrix. So, for a matrix A , the notation $\text{Col } A$ is often used instead of $\text{Ran } A$.

And now the last example.

4. Given a system of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r \in V$ its *linear span* (sometimes called simply *span*) $\mathcal{L}\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r\}$ is the collection of all vectors $\mathbf{v} \in V$ that can be represented as a linear combination $\mathbf{v} = \alpha_1\mathbf{v}_1 + \alpha_2\mathbf{v}_2 + \dots + \alpha_r\mathbf{v}_r$ of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r$. The notation $\text{span}\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r\}$ is also used instead of $\mathcal{L}\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r\}$

It is easy to check that in all of these examples we indeed have subspaces. We leave this as an exercise for the reader. Some of the statements will be proved later in the text.

Systems of linear equations

1. Different faces of linear systems.

There exist several points of view on what a system of linear equations, or in short a *linear system* is. The first, naïve one is, that it is simply a collection of m linear equations with n unknowns $x_1, x_2, \dots, x_n$,

$$\begin{cases} a_{1,1}x_1 + a_{1,2}x_2 + \dots + a_{1,n}x_n = b_1 \\ a_{2,1}x_1 + a_{2,2}x_2 + \dots + a_{2,n}x_n = b_2 \\ \dots \\ a_{m,1}x_1 + a_{m,2}x_2 + \dots + a_{m,n}x_n = b_m \end{cases}$$

To solve the system is to find *all* n -tuples of numbers $x_1, x_2, \dots, x_n$ which satisfy all m equations simultaneously.

If we denote $\mathbf{x} := (x_1, x_2, \dots, x_n)^T \in \mathbb{R}^n$, $\mathbf{b} = (b_1, b_2, \dots, b_m)^T \in \mathbb{R}^m$, and

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,n} \\ a_{2,1} & a_{2,2} & \dots & a_{2,n} \\ \vdots & \vdots & & \vdots \\ a_{m,1} & a_{m,2} & \dots & a_{m,n} \end{pmatrix},$$

then the above linear system can be written in the *matrix form* (as a *matrix-vector equation*)

$$A\mathbf{x} = \mathbf{b}.$$

To solve the above equation is to find all vectors $\mathbf{x} \in \mathbb{R}^n$ satisfying $A\mathbf{x} = \mathbf{b}$.

And finally, recalling the “column by coordinate” rule of the matrix-vector multiplication, we can write the system as a *vector equation*

$$x_1 \mathbf{a}_1 + x_2 \mathbf{a}_2 + \dots + x_n \mathbf{a}_n = \mathbf{b},$$

where $\mathbf{a}_k$ is the k th column of the matrix A , $\mathbf{a}_k = (a_{1,k}, a_{2,k}, \dots, a_{m,k})^T$, $k = 1, 2, \dots, n$.

Note, that these three examples are essentially just different representations of the same mathematical object.

Before explaining how to solve a linear system, let us notice that it does not matter how we call the unknowns, x_k , y_k or something else. So, all the information necessary to solve the system is contained in the matrix A , which is called *the coefficient matrix* of the system and in the vector (right side) $\mathbf{b}$. Hence, the all information we need is contained in the following matrix

$$\left(\begin{array}{cccc|c} a_{1,1} & a_{1,2} & \dots & a_{1,n} & b_1 \\ a_{2,1} & a_{2,2} & \dots & a_{2,n} & b_2 \\ \vdots & \vdots & & \vdots & \vdots \\ a_{m,1} & a_{m,2} & \dots & a_{m,n} & b_m \end{array} \right)$$

which is obtained by attaching the column $\mathbf{b}$ to the matrix A . This matrix is called the *augmented matrix* of the system. We will usually put the vertical line separating A and $\mathbf{b}$ to distinguish between the augmented matrix and the coefficient matrix.

2. Solution of a linear system. Echelon and reduced echelon forms

Linear system are solved by the *Gauss–Jordan elimination* (which is sometimes called *em row reduction*). By performing operations on rows of the augmented matrix of the system (i.e. on the equations), we reduce it to a simple form, the so-called *echelon form*. And when the system is in the *echelon form*, one can easily write the solution.

2.1. Row operations. There are three types of row operations we use:

1. Row exchange: interchange two rows of the matrix;
2. Scaling: multiply a row by a non-zero scalar a ;
3. Row replacement: replace a row $\# k$ by its sum with a constant multiple of a row $\# j$; all other rows remain intact;

It is clear that the operations 1 and 2 do not change the set of solutions of the system; they essentially do not change the system.

As for the operation 3, one can easily see that it does not lose solutions. Namely, let a “new” system be obtained from an “old” one by a row operation of type 3. Then any solution of the “old” system is a solution of the “new” one.

To see that we do not gain anything extra, i.e. that any solution of the “new” system is also a solution of the “old” one, we just notice that row operation of type 3 are *reversible*, i.e. the “old” system also can be obtained from the “new” one by applying a row operation of type 3 (can you say which one?)

2.1.1. *Row operations and multiplication by elementary matrices.* There is another, more “advanced” explanation why the above row operations are legal. Namely, every row operation is equivalent to the multiplication of the matrix from the left by one of the special elementary matrices.

Namely, the multiplication by the matrix

$$\begin{matrix} & & j & & k & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \\ j & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \\ k & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \end{matrix} \begin{pmatrix} 1 & & & & & & & & \\ & \ddots & & & & & & & \\ & & 1 & & & & & & \\ \dots & \dots & \dots & 0 & \dots & \dots & \dots & 1 & \\ & & & & 1 & & & & \\ & & & & & \ddots & & & \\ & & & & & & 1 & & \\ \dots & \dots & \dots & 1 & \dots & \dots & \dots & 0 & \\ & & & & & & & & 1 & \ddots & \\ & & & & & & & & & & 1 \end{pmatrix}$$

just interchanges the rows number j and number k . Multiplication by the matrix

$$\begin{matrix} & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ k & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \end{matrix} \begin{pmatrix} 1 & & & & & & & & \\ & \ddots & & & & & & & \\ & & 1 & 0 & & & & & \\ \dots & \dots & 0 & a & 0 & & & & \\ & & & 0 & 1 & & & & \\ & & & & & \ddots & 0 & & \\ & & & & & & 0 & 1 & \end{pmatrix}$$

multiplies the row number k by a . And finally, the multiplication by the matrix

$$j \begin{pmatrix} 1 & & \vdots & & \vdots & & \\ & \ddots & \vdots & & \vdots & & 0 \\ \dots & \dots & 1 & \dots & 0 & & \\ & & \vdots & \ddots & \vdots & & \\ k \dots & \dots & a & \dots & 1 & & \\ 0 & & & & & \ddots & \\ & & & & & & 1 \end{pmatrix}$$

A way to describe (or to remember) these elementary matrices: they are obtained from I by applying the corresponding row operation to it

adds to the row $\#k$ row $\#j$ multiplied by a , and leaves all other rows intact.

To see, that the multiplication by these matrices works as advertised, one can just see how the multiplications act on vectors (columns).

Note that all these matrices are invertible (compare with reversibility of row operations). The inverse of the first matrix is the matrix itself. To get the inverse of the second one, one just replaces a by $1/a$. And finally, the inverse of the third matrix is obtained by replacing a by $-a$. To see that the inverses are indeed obtained this way, one again can simply check how they act on columns.

So, performing a row operation on the augmented matrix of the system $A\mathbf{x} = \mathbf{b}$ is equivalent to the multiplication of the system (from the left) by a special invertible matrix E . Clearly, any solution of the equation

$$A\mathbf{x} = \mathbf{b}$$

is also a solution of

$$EA\mathbf{x} = E\mathbf{b}.$$

Multiplying this equation (from the left) by E^{-1} we get that any its solution is the solution of the equation

$$E^{-1}EA\mathbf{x} = E^{-1}E\mathbf{b},$$

which is the original equation $A\mathbf{x} = \mathbf{b}$. So, a row operation does not change the solution set of a system.

2.2. Row reduction. And now how the row reduction is carried out. The main step consists of 3 substeps:

1. Find the leftmost non-zero column of the matrix;
2. Make sure, by applying row operations of type 2, if necessary, that the first (the upper) entry of this column is non-zero.

This entry will be called the *pivot entry* or simply the *pivot*;

3. “Kill” (i.e. make them 0) all non-zero entries below the pivot by adding (subtracting) an appropriate multiple of the first row from the rows number $2, 3, \dots, m$.

We apply the main step to a matrix, then we leave the first row alone and apply the main step to rows $2, \dots, m$, then to rows $3, \dots, m$, etc.

The point to remember is that after we subtract a multiple of a row from all rows below it (step 3), we leave it alone and do not change it in any way, not even interchange it with another row.

After applying the main step finitely many times (at most m), we get what is called the *echelon form* of the matrix.

2.2.1. *An example of row reduction.* Let us consider the following linear system:

$$\begin{cases} x_1 + 2x_2 + 3x_3 = 1 \\ 3x_1 + 2x_2 + x_3 = 7 \\ 2x_1 + x_2 + 2x_3 = 1 \end{cases}$$

The augmented matrix of the system is

$$\left(\begin{array}{ccc|c} 1 & 2 & 3 & 1 \\ 3 & 2 & 1 & 7 \\ 2 & 1 & 2 & 1 \end{array} \right)$$

Subtracting $3 \cdot \text{Row\#1}$ from the second row, and subtracting $2 \cdot \text{Row\#1}$ from the third one we get:

$$\left(\begin{array}{ccc|c} 1 & 2 & 3 & 1 \\ 3 & 2 & 1 & 7 \\ 2 & 1 & 2 & 1 \end{array} \right) \begin{array}{l} -3R_1 \\ -2R_1 \end{array} \sim \left(\begin{array}{ccc|c} 1 & 2 & 3 & 1 \\ 0 & -4 & -8 & 4 \\ 0 & -3 & -4 & -1 \end{array} \right)$$

Multiplying the second row by $-1/4$ we get

$$\left(\begin{array}{ccc|c} 1 & 2 & 3 & 1 \\ 0 & 1 & 2 & -1 \\ 0 & -3 & -4 & -1 \end{array} \right)$$

Adding $3 \cdot \text{Row\#2}$ to the third row we obtain

$$\left(\begin{array}{ccc|c} 1 & 2 & 3 & 1 \\ 0 & 1 & 2 & -1 \\ 0 & -3 & -4 & -1 \end{array} \right) -3R_2 \sim \left(\begin{array}{ccc|c} 1 & 2 & 3 & 1 \\ 0 & 1 & 2 & -1 \\ 0 & 0 & 2 & -4 \end{array} \right)$$

Now we can use the so called *back substitution* to solve the system. Namely, from the last row (equation) we get $x_3 = -2$. Then from the second equation we get

$$x_2 = -1 - 2x_3 = -1 - 2(-2) = 3,$$

and finally, from the first row (equation)

$$x_1 = 1 - 2x_2 - 3x_3 = 1 - 6 + 6 = 1.$$

So, the solution is

$$\begin{cases} x_1 = 1 \\ x_2 = 3, \\ x_3 = -2, \end{cases}$$

or in the vector form

$$\mathbf{x} = \begin{pmatrix} 1 \\ 3 \\ -2 \end{pmatrix}$$

or $\mathbf{x} = (1, 3, -2)^T$. We can check the solution by multiplying $A\mathbf{x}$, where A is the coefficient matrix.

Instead of using back substitution, we can do row reduction from down to top, killing all the entries above the main diagonal of the coefficient matrix: we start by multiplying the last row by 1/2, and the rest is pretty self-explanatory:

$$\begin{aligned} \left(\begin{array}{ccc|c} 1 & 2 & 3 & 1 \\ 0 & 1 & 2 & -1 \\ 0 & 0 & 1 & -2 \end{array} \right) & \begin{array}{l} -3R_3 \\ -2R_3 \end{array} \sim \left(\begin{array}{ccc|c} 1 & 2 & 0 & 7 \\ 0 & 1 & 0 & 3 \\ 0 & 0 & 1 & -2 \end{array} \right) \begin{array}{l} -2R_2 \\ \end{array} \\ & \sim \left(\begin{array}{ccc|c} 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 3 \\ 0 & 0 & 1 & -2 \end{array} \right) \end{aligned}$$

and we just read the solution $\mathbf{x} = (1, 3, -2)^T$ off the augmented matrix.

We leave it as an exercise to the reader to formulate the algorithm for the backward phase of the row reduction.

2.3. Echelon form. A matrix is in *echelon form* if it satisfies the following two conditions:

1. All zero rows (i.e. the rows with all entries equal 0), if any, are below all non-zero entries.

For a non-zero row, let us call the leftmost non-zero entry the *leading entry*. Then the second property of the echelon form can be formulated as follows:

2. For any non-zero row its leading entry is strictly to the right of the leading entry in the previous row.

Pivots: leading
(rightmost non-zero
entries) in a row

The leading entry in each row in echelon form is also called *pivot entry*, or simply *pivot*, because these entries are exactly the *pivots* we used in the row reduction.

A particular case of the echelon form is the so-called *triangular form*. We got this form in our example above. In this form the coefficient matrix is square ($n \times n$), all its entries on the main diagonal are non-zero, and all the

entries below the main diagonal are zero. The right side, i.e. the rightmost column of the augmented matrix can be arbitrary.

After the backward phase of the row reduction, we get what the so-called *reduced echelon form* of the matrix: coefficient matrix equal I , as in the above example, is a particular case of the reduced echelon form.

The general definition is as follows: we say that a matrix is in the *reduced echelon form*, if it is in the echelon form and

3. All pivot entries are equal 1;
4. All entries above the pivots are 0. Note, that all entries below the pivots are also 0 because of the echelon form.

To get reduced echelon form from echelon form, we work from the bottom to the top and from the right to the left, using row replacement to kill all entries above the pivots.

An example of the reduced echelon form is the system with the coefficient matrix equal I . In this case, one just reads the solution from the reduced echelon form. In general case, one can also easily read the solution from the reduced echelon form. For example, let the reduced echelon form of the system (augmented matrix) be

$$\left(\begin{array}{ccccc|c} \boxed{1} & 2 & 0 & 0 & 0 & 1 \\ 0 & 0 & \boxed{1} & 5 & 0 & 2 \\ 0 & 0 & 0 & 0 & \boxed{1} & 3 \end{array} \right);$$

here we boxed the pivots. The idea is to move the variables, corresponding to the columns without pivot (the so-called *free variables*) to the right side. Then we can just write the solution.

$$\begin{cases} x_1 = 1 - 2x_2 \\ x_2 \text{ is free} \\ x_3 = 2 - 5x_2 \\ x_4 \text{ is free} \\ x_5 = 3 \end{cases}$$

or in the vector form

$$\mathbf{x} = \begin{pmatrix} 1 - 2x_2 \\ x_2 \\ 1 - 5x_2 \\ x_4 \\ 3 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \\ 3 \end{pmatrix} + x_2 \begin{pmatrix} -2 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + x_4 \begin{pmatrix} 0 \\ 0 \\ -5 \\ 1 \\ 0 \end{pmatrix}$$

One can also find the solution from the echelon form by using back substitution: the idea is to work from bottom to top, moving all free variables to the right side.

3. Analyzing the pivots.

All questions about existence of a solution and its uniqueness can be answered by analyzing pivots in the echelon (reduced echelon) form of the augmented matrix of the system. First of all, let us ask yourself a question: when the equation $A\mathbf{x} = \mathbf{b}$ is *inconsistent*, i.e. when it does not have a solution? The answer follows immediately, if one just thinks about it:

a system is inconsistent (does not have a solution) if and only if there is a pivot in the last row of an echelon form of the *augmented* matrix, i.e. iff an echelon form of the augmented matrix has a row $(0 \ 0 \ \dots \ 0 \mid b)$, $b \neq 0$ in it.

Indeed, such a row corresponds to the equation $0x_1 + 0x_2 + \dots + 0x_n = b \neq 0$ that does not have a solution. And if we don't have such a row, we just make the reduced echelon form and then read the solution off it.

Now, three more statements. Note, that they both deal with *coefficient matrix*, not with the augmented matrix of the system.

1. A solution (if it exists) is unique iff there are no free variables, that is if and only if the echelon form of the coefficient matrix has a pivot in every column;
2. Equation $A\mathbf{x} = \mathbf{b}$ is consistent for all right sides $\mathbf{b}$ if and only if the echelon form of the coefficient matrix has a pivot in every row.
3. Equation $A\mathbf{x} = \mathbf{b}$ has a *unique solution* for any right side $\mathbf{b}$ if and only if echelon form of the coefficient matrix A has a pivot in every column and every row.

The first statement is trivial, because free variables are responsible for all non-uniqueness. I should only emphasize that this statement *does not tell anything* about the existence.

The second statement is a tiny bit more complicated. If we have a pivot in every row of the coefficient matrix, we cannot have the pivot in the last column of the *augmented* matrix, so the system is always consistent, no matter what the right side $\mathbf{b}$ is.

Let us show that if we have a zero row in the echelon form of the coefficient matrix A , then we can pick a right side $\mathbf{b}$ such that the system $A\mathbf{x} = \mathbf{b}$ is not consistent. Let A_e be echelon form of the coefficient matrix A . Then

$$A_e = EA,$$

where E is the product of elementary matrices, corresponding to the row operations, $E = E_N, \dots, E_2, E_1$. If A_e has a zero row, then the last row is also zero. Therefore, if we put $\mathbf{b}_e = (0, \dots, 0, 1)^T$ (all entries are 0, except

the last one), then the equation

$$A_e \mathbf{x} = \mathbf{b}_e$$

does not have a solution. Multiplying this equation by E^{-1} from the left, an recalling that $E^{-1}A_e = A$, we get that the equation

$$A\mathbf{x} = E^{-1}\mathbf{b}_e$$

does not have a solution.

Finally, statement 3 immediately follows from statements 1 and 2. $\square$

From the above analysis of pivots we get several very important corollaries. The main observation we use is

In echelon form, any row and any column have no more than 1 pivot in it (it can have 0 pivots)

3.1. Corollaries about linear independence and bases. Dimension.

Questions as to when a system of vectors in $\mathbb{R}^n$ is a basis, a linearly independent or a spanning system, can be easily answered by the row reduction.

Proposition 3.1. *Let we have a system of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m \in \mathbb{R}^n$, and let $A = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m]$ be an $n \times m$ matrix with columns $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m$. Then*

1. *The system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m$ is linearly independent iff echelon form of A has a pivot in every column;*
2. *The system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m$ is complete in $\mathbb{R}^n$ (spanning, generating) iff echelon form of A has a pivot in every row;*
3. *The system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m$ is a basis in $\mathbb{R}^n$ iff echelon form of A has a pivot in every column and in every row.*

Proof. The system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m \in \mathbb{R}^m$ is linearly independent if and only if the equation

$$x_1\mathbf{v}_1 + x_2\mathbf{v}_2 + \dots + x_m\mathbf{v}_m = \mathbf{0}$$

has the unique (trivial) solution $x_1 = x_2 = \dots = x_m = 0$, or equivalently, the equation $A\mathbf{x} = \mathbf{0}$ has unique solution $\mathbf{x} = \mathbf{0}$. By statement 1 above, it happens if and only if there is a pivot in every column of the matrix.

Similarly, the system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m \in \mathbb{R}^m$ is complete in $\mathbb{R}^n$ if and only if the equation

$$x_1\mathbf{v}_1 + x_2\mathbf{v}_2 + \dots + x_m\mathbf{v}_m = \mathbf{b}$$

has a solution for any right side $\mathbf{b} \in \mathbb{R}^n$. By statement 2 above, it happens if and only if there is a pivot in every column in echelon form of the matrix.

And finally, the system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m \in \mathbb{R}^n$ is a basis in $\mathbb{R}^n$ if and only if the equation

$$x_1\mathbf{v}_1 + x_2\mathbf{v}_2 + \dots + x_m\mathbf{v}_m = \mathbf{b}$$

has unique solution for any right side $\mathbf{b} \in \mathbb{R}^n$. By statement 3 this happens if and only if there is a pivot in every column and in every row of echelon form of A . $\square$

Proposition 3.2. *Any linearly independent system of vectors in $\mathbb{R}^n$ cannot have more than n vectors in it.*

Proof. Let a system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m \in \mathbb{R}^n$ be linearly independent, and let $A = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m]$ be the $n \times m$ matrix with columns $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m$. By Proposition 3.1 echelon form of A must have a pivot in every column, which is impossible if $m > n$ (number of pivots cannot be more than number of rows). $\square$

Proposition 3.3. *Any two bases in a vector space V have the same number of vectors in them.*

Proof. Let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ and $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_m$ be two different bases in V . Without loss of generality we can assume that $n \leq m$. Consider an isomorphism $A : \mathbb{R}^n \rightarrow V$ defined by

$$A\mathbf{e}_k = \mathbf{v}_k, \quad k = 1, 2, \dots, n,$$

where $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ is the standard basis in $\mathbb{R}^n$.

Since A^{-1} is also an isomorphism, the system

$$A^{-1}\mathbf{w}_1, A^{-1}\mathbf{w}_2, \dots, A^{-1}\mathbf{w}_m$$

is a basis (see Theorem 5.6 in Chapter 1). So it is linearly independent, and by Proposition 3.2, $m \leq n$. Together with the assumption $n \leq m$ this implies that $m = n$. $\square$

The statement below is a particular case of the above proposition.

Proposition 3.4. *Any basis in $\mathbb{R}^n$ must have exactly n vectors in it.*

Proof. This fact follows immediately from the previous proposition, but there is also a direct proof. Let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m$ be a basis in $\mathbb{R}^n$ and let A be the $n \times m$ matrix with columns $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m$. The fact that the system is a basis, means that the equation

$$A\mathbf{x} = \mathbf{b}$$

has a unique solution for any (all possible) right side $\mathbf{b}$. The existence means that there is a pivot in every row (of a reduced echelon form of the matrix),

hence the number of pivots is exactly n . The uniqueness mean that there is pivot in every column of the coefficient matrix (its echelon form), so

$$m = \text{number of columns} = \text{number of pivots} = n$$

□

Proposition 3.5. *Any spanning (generating) set in $\mathbb{R}^n$ must have at least n vectors.*

Proof. Let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m$ be a complete system in $\mathbb{R}^n$, and let A be $n \times m$ matrix with columns $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m$. Statement 2 of Proposition 3.1 implies that echelon form of A has a pivot in every row. Since number of pivots cannot exceed the number of rows, $n \leq m$. □

3.2. Corollaries about invertible matrices.

Proposition 3.6. *An invertible matrix must be square ($n \times n$).*

Proof. We know that an $m \times n$ matrix is invertible if and only if its columns form a basis in $\mathbb{R}^m$ (see Corollary 5.9 in Section 5.4, Chapter 1). Proposition 3.4 above implies that $m = n$. □

Proposition 3.7. *If a square ($n \times n$) matrix is left invertible, or if it is right right invertible, then it is invertible. In other words, to check the invertibility of a square matrix A it is sufficient to check only one of the conditions $AA^{-1} = I$, $A^{-1}A = I$.*

Note, that this proposition apply only to square matrices!

Proof. We know that matrix A is invertible if and only if the equation $A\mathbf{x} = \mathbf{b}$ has a unique solution for any right side $\mathbf{b}$. This happens if and only if echelon form of the matrix A has pivots in every row and in every column.

If a matrix A is left invertible, the equation $A\mathbf{x} = \mathbf{0}$ has unique solution $\mathbf{x} = \mathbf{0}$. Indeed, if B is a left inverse of A (i.e. $BA = I$), and $\mathbf{x}$ satisfies

$$A\mathbf{x} = \mathbf{0},$$

then multiplying this identity by B from the left we get $\mathbf{x} = \mathbf{0}$, so the solution is unique. Therefore, echelon form of A has pivots in every row. If the matrix A is square ($n \times n$), the echelon form also has pivots in every column, so the matrix is invertible.

If a matrix A is right invertible, and C is its right inverse ($AC = I$), then for $\mathbf{x} = C\mathbf{b}$, $\mathbf{b} \in \mathbb{R}^n$

$$A\mathbf{x} = AC\mathbf{b} = I\mathbf{b} = \mathbf{b}.$$

Therefore, for any right side $\mathbf{b}$ the equation $A\mathbf{x} = \mathbf{b}$ has a solution $\mathbf{x} = C\mathbf{b}$. Thus, echelon form of A has pivots in every row. If A is square, it also has a pivot in every column, so A is invertible. $\square$

4. Dimension. Finite-dimensional spaces.

Definition. The dimension $\dim V$ of a vector space V is the number of vectors in a basis.

For a vector space consisting only of zero vector $\mathbf{0}$ we put $\dim V = 0$. If V does not have a (finite) basis, we put $\dim V = \infty$.

If $\dim V$ is finite, we call the space V *finite-dimensional*; otherwise we call it *infinite-dimensional*.

Proposition 3.3 asserts that the dimension is well defined, i.e. that it does not depend on the choice of a basis.

Proposition 2.8 from Chapter 1 states that any finite spanning system in a vector space V contains a basis. This immediately implies the following

Proposition 4.1. *A vector space V is finite-dimensional if and only if it has a finite spanning system.*

Suppose, that we have a system of vectors in a finite-dimensional vector space, and we want to check if it is a basis (or if it is linearly independent, or if it is complete)? Probably the simplest way is to use an isomorphism $A : V \rightarrow \mathbb{R}^n$, $n = \dim V$ to move the problem to $\mathbb{R}^n$, where all such questions can be answered by row reduction (studying pivots).

Note, that if $\dim V = n$, then there always exists an isomorphism $A : V \rightarrow \mathbb{R}^n$. Indeed, if $\dim V = n$ then there exists a basis $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n \in V$, and one can define an isomorphism $A : V \rightarrow \mathbb{R}^n$ by

$$A\mathbf{v}_k = \mathbf{e}_k, \quad k = 1, 2, \dots, n.$$

As an example, let us give the following two corollaries of the above Propositions 3.2, 3.5:

Proposition 4.2. *Any linearly independent system in a finite-dimensional vector space V cannot have more than $\dim V$ vectors in it.*

Proof. Let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m \in V$ be a linearly independent system, and let $A : V \rightarrow \mathbb{R}^n$ be an isomorphism. Then $A\mathbf{v}_1, A\mathbf{v}_2, \dots, A\mathbf{v}_m$ is a linearly independent system in $\mathbb{R}^n$, and by Proposition 3.2 $m \leq n$. $\square$

Proposition 4.3. *Any complete independent system in a finite-dimensional vector space V must have at least $\dim V$ vectors in it.*

Proof. Let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m \in V$ be a complete system in V , and let $A : V \rightarrow \mathbb{R}^n$ be an isomorphism. Then $A\mathbf{v}_1, A\mathbf{v}_2, \dots, A\mathbf{v}_m$ is a complete system in $\mathbb{R}^n$, and by Proposition 3.5 $m \geq n$. $\square$

5. General solution of a linear system.

In this short section we discuss the structure of the general solution (i.e. of the solution set) of a linear system.

We call a system $A\mathbf{x} = \mathbf{b}$ homogeneous, if the right side, $\mathbf{b} = \mathbf{0}$, i.e. a homogeneous system is a system of form $A\mathbf{x} = \mathbf{0}$.

With each system

$$A\mathbf{x} = \mathbf{b}$$

we can associate a homogeneous system just by putting $\mathbf{b} = \mathbf{0}$.

Theorem 5.1 (General solution of a linear equation). *Let a vector $\mathbf{x}_1$ satisfy the equation $A\mathbf{x} = \mathbf{b}$, and let H be the set of all solutions of the associated homogeneous system*

$$A\mathbf{x} = \mathbf{0}.$$

Then the set

$$\{\mathbf{x} = \mathbf{x}_1 + \mathbf{x}_h : \mathbf{x}_h \in H\}$$

is the set of all solutions of the equation $A\mathbf{x} = \mathbf{b}$.

In other words, this theorem can be stated as

$$\boxed{\begin{array}{l} \text{General solution} \\ \text{of } A\mathbf{x} = \mathbf{b} \end{array}} = \boxed{\begin{array}{l} \text{A particular solu-} \\ \text{tion of } A\mathbf{x} = \mathbf{b} \end{array}} + \boxed{\begin{array}{l} \text{General solution} \\ \text{of } A\mathbf{x} = \mathbf{0} \end{array}}.$$

Proof. Fix a vector $\mathbf{x}_1$ satisfying $A\mathbf{x}_1 = \mathbf{b}$. Let a vector $\mathbf{x}_h$ satisfy $A\mathbf{x}_h = \mathbf{0}$. Then for $\mathbf{x} = \mathbf{x}_1 + \mathbf{x}_h$ we have

$$A\mathbf{x} = A(\mathbf{x}_1 + \mathbf{x}_h) = A\mathbf{x}_1 + A\mathbf{x}_h = \mathbf{b} + \mathbf{0} = \mathbf{b},$$

so any $\mathbf{x}$ of form

$$\mathbf{x} = \mathbf{x}_1 + \mathbf{x}_h, \quad \mathbf{x}_h \in H$$

is a solution of $A\mathbf{x} = \mathbf{b}$.

let now $\mathbf{x}$ be satisfy $A\mathbf{x} = \mathbf{b}$. Then for $\mathbf{x}_h := \mathbf{x} - \mathbf{x}_1$ we get

$$A\mathbf{x}_h = A(\mathbf{x} - \mathbf{x}_1) = A\mathbf{x} - A\mathbf{x}_1 = \mathbf{b} - \mathbf{b} = \mathbf{0},$$

so $\mathbf{x}_h \in H$. Therefore *any* solution $\mathbf{x}$ of $A\mathbf{x} = \mathbf{b}$ can be represented as $\mathbf{x} = \mathbf{x}_1 + \mathbf{x}_h$ with some $\mathbf{x}_h \in H$. $\square$

The power of this theorem is in its generality. It applies to all linear equations, we do not have to assume here that vector spaces are finite-dimensional. You will meet this theorem in differential equations, integral equations, partial differential equations, etc. Besides showing the structure of the solution set, this theorem allows to separate investigation of uniqueness from the study of existence. Namely, to study uniqueness, we only need to analyze uniqueness of the homogeneous equation $A\mathbf{x} = \mathbf{0}$, which always have a solution.

There is an immediate application in this course: this theorem allows us to check a solution of a system $A\mathbf{x} = \mathbf{b}$. For example, consider a system

$$\begin{pmatrix} 2 & 3 & 1 & 4 & -9 \\ 1 & 1 & 1 & 1 & -3 \\ 1 & 1 & 1 & 2 & -5 \\ 2 & 2 & 2 & 3 & -8 \end{pmatrix} \mathbf{x} = \begin{pmatrix} 17 \\ 6 \\ 8 \\ 14 \end{pmatrix}$$

Performing row reduction one can find the solution of this system

$$(5.1) \quad \mathbf{x} = \begin{pmatrix} 3 \\ 1 \\ 0 \\ 2 \\ 0 \end{pmatrix} + x_3 \begin{pmatrix} -2 \\ 1 \\ 1 \\ 0 \\ 0 \end{pmatrix} + x_5 \begin{pmatrix} 2 \\ -1 \\ 0 \\ 2 \\ 1 \end{pmatrix}, \quad x_3, x_5 \in \mathbb{R}$$

The parameters x_3, x_5 can be denoted here by any other letters, t and s , for example; we keeping notation x_3 and x_5 here only to remind that they came from the corresponding free variables.

Now, let us suppose, that we are just given this solution, and we want to check whether it is correct or not. Of course, we can repeat the row operations, but this is too time consuming. And moreover, if the solution was obtained by some nonstandard method, it can look differently from what we get from the row reduction. For example the formula

$$(5.2) \quad \mathbf{x} = \begin{pmatrix} 3 \\ 1 \\ 0 \\ 2 \\ 0 \end{pmatrix} + s \begin{pmatrix} -2 \\ 1 \\ 1 \\ 0 \\ 0 \end{pmatrix} + t \begin{pmatrix} 0 \\ 0 \\ 1 \\ 2 \\ 1 \end{pmatrix}, \quad s, t \in \mathbb{R}$$

gives the same set as (5.1) (can you say why?); here we just replaced the last vector in (5.1) by its sum with the second one. So, this formula is different from the solution we got from the row reduction, but it is nevertheless correct.

The simplest way to check that (5.1) and (5.2) give us correct solutions, is to check that the first vector $(3, 1, 0, 2, 0)^T$ satisfies the equation $A\mathbf{x} = \mathbf{b}$,

and that the other two (the ones with the parameters x_3 and x_5 or s and t in front of them) should satisfy the associated homogeneous equation $A\mathbf{x} = \mathbf{0}$.

If this checks out, we will be assured, that any vector $\mathbf{x}$ defined by (5.1) or (5.2) is indeed a solution.

Note, that this method of checking the solution does not guarantee that (5.1) (or (5.2)) gives us all the solutions. For example, if we just somehow miss the term with x_2 , the above method of checking will still work fine.

So, how can we guarantee, that we did not miss any free variable, and there should not be extra term in (5.1)?

What comes to mind, is to count the pivots again. In this example, if one does row operations, the number of pivots is 3. So indeed, there should be 2 free variables, and it looks like we did not miss anything in (5.1).

To be able to *prove* this, we will need new notions of fundamental subspaces and of rank of a matrix. I should also mention, that in this particular example, one even does not have to perform row operations to check that there are only 2 free variables, and that formulas (5.1) and (5.2) both give correct general solution.

6. Fundamental subspaces of a matrix. Rank.

As we discussed above in Section 6 of Chapter 1, with any linear transformation $A : V \rightarrow W$ we can associate two subspaces, namely, its kernel, or null space

$$\text{Ker } A = \text{Null } A := \{\mathbf{v} \in V : A\mathbf{v} = \mathbf{0}\} \subset V,$$

and its range

$$\text{Ran } A = \{\mathbf{w} \in W : \mathbf{w} = A\mathbf{v} \text{ for some } \mathbf{v} \in V\} \subset W.$$

In other words, the kernel $\text{Ker } A$ is the solution set of the homogeneous equation $A\mathbf{x} = \mathbf{0}$, and the range $\text{Ran } A$ is exactly the set of all right sides $\mathbf{b} \in W$ for which the equation $A\mathbf{x} = \mathbf{b}$ has a solution.

If A is an $m \times n$ matrix, i.e. a mapping from $\mathbb{R}^n$ to $\mathbb{R}^m$, then it follows from the “column by coordinate” rule of the matrix multiplication that any vector $\mathbf{w} \in \text{Ran } A$ can be represented as a linear combination of columns of A . This explains the name *column space* (notation $\text{Col } A$), which is often used instead of $\text{Ran } A$.

If A is a matrix, then in addition to $\text{Ran } A$ and $\text{Ker } A$ one can also consider the range and kernel for the transposed matrix A^T . Often the term *row space* is used for $\text{Ran } A^T$ and the term *left null space* is used for $\text{Ker } A^T$ (but usually no special notation is introduced).

The four subspaces $\text{Ran } A$, $\text{Ker } A$, $\text{Ran } A^T$, $\text{Ker } A^T$ are called the *fundamental subspaces* of the matrix A . In this section we will study important relations between the dimensions of the four fundamental subspaces.

We will need the following definition, which is one of the fundamental notions of Linear Algebra

Definition. Given a linear transformation (matrix) A its rank $\text{rank } A$ is the dimension of the range of A

$$\text{rank } A := \dim \text{Ran } A.$$

6.1. Computing fundamental subspaces and rank. To compute fundamental subspaces and rank of a matrix, one needs to do echelon reduction. Namely, let A be the matrix, and A_e be its echelon form

1. The pivot columns of the *original* matrix A (i.e. the columns where after row operations we will have pivots in the echelon form) give us a basis (one of many possible) in $\text{Ran } A$.
2. The pivot *rows* of the echelon form A_e give us a basis in the row space. Of course, it is possible just to transpose the matrix, and then do row operations. But if we already got echelon form of A , say to compute $\text{Ran } A$, then we get $\text{Ran } A^T$ for free.
3. To find a basis in the null space $\text{Ker } A$ one needs to solve the homogeneous equation $A\mathbf{x} = \mathbf{0}$: the details will be seen from the example below.

Example. Consider a matrix

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

Performing row operations we get the echelon form

$$\begin{pmatrix} \boxed{1} & 1 & 2 & 2 & 1 \\ 0 & 0 & \boxed{-3} & -3 & -1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

(the pivots are boxed here). So, the columns 1 and 3 of the *original matrix*, i.e. the columns

$$\begin{pmatrix} 1 \\ 2 \\ 3 \\ 1 \end{pmatrix}, \quad \begin{pmatrix} 2 \\ 2 \\ 3 \\ -1 \end{pmatrix}$$

give us a basis in $\text{Ran } A$. We also got a basis in the row space $\text{Ran } A^T$ for free: the first and second row of the *echelon form* of A , i.e. the vectors

$$\begin{pmatrix} 1 \\ 1 \\ 2 \\ 2 \\ 1 \end{pmatrix}, \quad \begin{pmatrix} 0 \\ 0 \\ -3 \\ -3 \\ -1 \end{pmatrix}$$

(we put the vectors vertically here. The question of whether to put vectors here vertically as columns, or horizontally as rows is really a matter of convention. Our reason for putting them vertically is that although we call $\text{Ran } A^T$ the *row space* we define it as a column space of A^T)

To compute the basis in the null space $\text{Ker } A$ we need to solve the equation $A\mathbf{x} = \mathbf{0}$. Compute the *reduced* echelon form of A , which in this example is

$$\begin{pmatrix} \boxed{1} & 1 & 0 & 0 & 1/3 \\ 0 & 0 & \boxed{1} & 1 & 1/3 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

Note, that when solving the homogeneous equation $A\mathbf{x} = \mathbf{0}$, it is not necessary to write the whole augmented matrix, it is sufficient to work with the coefficient matrix. Indeed, in this case the last column of the augmented matrix is the column of zeroes, which does not change under row operations. So, we can just keep this column in mind, without actually writing it. Keeping this last zero column in mind, we can read the solution off the reduced echelon form above:

$$\begin{cases} x_1 = -x_2 - \frac{1}{3}x_5, \\ x_2 \text{ is free,} \\ x_3 = -x_4 - \frac{1}{3}x_5 \\ x_4 \text{ is free,} \\ x_5 \text{ is free,} \end{cases}$$

or, in the vector form

$$(6.1) \quad \mathbf{x} = \begin{pmatrix} -x_2 - \frac{1}{3}x_5 \\ x_2 \\ -x_4 - \frac{1}{3}x_5 \\ x_4 \\ x_5 \end{pmatrix} = x_2 \begin{pmatrix} -1 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + x_4 \begin{pmatrix} 0 \\ 0 \\ -1 \\ 1 \\ 0 \end{pmatrix} + x_5 \begin{pmatrix} -1/3 \\ 0 \\ -1/3 \\ 0 \\ 1 \end{pmatrix}$$

The vectors at each free variable, i.e. in our case the vectors

$$\begin{pmatrix} -1 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} 0 \\ 0 \\ -1 \\ 1 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} -1/3 \\ 0 \\ -1/3 \\ 0 \\ 1 \end{pmatrix}$$

form a basis in $\text{Ker } A$.

Unfortunately, there is no shortcut for finding a basis in $\text{Ker } A^T$: one just need honestly solve the equation $A^T \mathbf{x} = \mathbf{0}$. Knowledge of the echelon form of A does not help here.

6.2. Explanation of the computing bases in the fundamental subspaces. So, why do the above methods indeed give us bases in the fundamental subspaces?

6.2.1. *The null space* $\text{Ker } A$. The case of the null space $\text{Ker } A$ is probably the simplest one: since we solved the equation $A\mathbf{x} = \mathbf{0}$, i.e. found all the solutions, then any vector in $\text{Ker } A$ is a linear combination of the vectors we obtained. Thus, the vectors we obtained form a spanning system in $\text{Ker } A$. To see that the system is linearly independent, let us multiply each vector by the corresponding free variable and add everything, see (6.1). Then for each free variable x_k the entry number k of the resulting vector is exactly x_k , see (6.1) again, so the only way this vector (the linear combination) can be $\mathbf{0}$ is when all free variables are 0.

6.2.2. *The column space* $\text{Ran } A$. Let us now explain why the method for finding a basis in the column space $\text{Ran } A$ works. First of all notice, that the pivot columns of the *reduced echelon* form A_{re} of A form a basis in $\text{Ran } A_{\text{re}}$ (not in the column space of the original matrix, but of its reduced echelon form!). Since row operations are just left multiplications by invertible matrices, they do not change linear independence. Therefore, the pivot columns of the *original* matrix A are linearly independent.

Let us now show that the pivot columns of A span the column space of A . Let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r$ be the pivot columns of A , and let $\mathbf{v}$ be an arbitrary column of A . We want to show that $\mathbf{v}$ can be represented as a linear combination of the pivot columns $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r$,

$$\mathbf{v} = \alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \dots + \alpha_r \mathbf{v}_r.$$

the reduced echelon form A_{re} is obtained from A by the left multiplication

$$A_{\text{re}} = EA,$$

where E is a product of elementary matrices, so E is an invertible matrix. The vectors $E\mathbf{v}_1, E\mathbf{v}_2, \dots, E\mathbf{v}_r$ are the pivot columns of A_{re} , and the column

$\mathbf{v}$ of A is transformed to the column $E\mathbf{v}$ of A_{re} . Since the pivot columns of A_{re} form a basis in $\text{Ran } A_{\text{re}}$, vector $E\mathbf{v}$ can be represented as a linear combination

$$E\mathbf{v} = \alpha_1 E\mathbf{v}_1 + \alpha_2 E\mathbf{v}_2 + \dots + \alpha_r E\mathbf{v}_r.$$

Multiplying this equality by E^{-1} from the left we get the representation

$$\mathbf{v} = \alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \dots + \alpha_r \mathbf{v}_r,$$

so indeed pivot columns of A span $\text{Ran } A$.

6.2.3. *The row space* $\text{Ran } A^T$. It is easy to see that the pivot rows of the echelon form A_e of A are linearly independent. Indeed, let $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_r$ are the transposed (since we agreed always to put vectors vertically) pivot rows of A_e . Suppose

$$\alpha_1 \mathbf{w}_1 + \alpha_2 \mathbf{w}_2 + \dots + \alpha_r \mathbf{w}_r = \mathbf{0}.$$

Consider the first non-zero entry of $\mathbf{w}_1$. Since for all other vectors $\mathbf{w}_2, \mathbf{w}_3, \dots, \mathbf{w}_r$ the corresponding entries equal 0 (by the definition of echelon form), we can conclude that $\alpha_1 = 0$. So we can just ignore the first term in the sum.

Consider now the first non-zero entry of $\mathbf{w}_2$. The corresponding entries of the vectors $\mathbf{w}_3, \dots, \mathbf{w}_r$ are 0, so $\alpha_2 = 0$. Repeating this procedure, we get that $\alpha_k = 0 \ \forall k = 1, 2, \dots, r$.

To see that vectors $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_r$ span the row space, one can notice that *row operations do not change the row space*. This can be obtained directly from analyzing row operations, but we present here a more formal way to demonstrate this fact.

For a transformation A and a set X let us denote by $A(X)$ the set of all elements y which can be represented as $y = A(x)$, $x \in X$,

$$A(X) := \{y = A(x) : x \in X\}.$$

If A is an $m \times n$ matrix, and A_e is its echelon form, A_e is obtained from A by left multiplication

$$A_e = EA,$$

where E is an $m \times m$ invertible matrix (the product of the corresponding elementary matrices). Then

$$\text{Ran } A_e^T = \text{Ran}(A^T E^T) = A^T(\text{Ran } E^T) = A^T(\mathbb{R}^m) = \text{Ran } A^T,$$

so indeed $\text{Ran } A^T = \text{Ran } A_e^T$.

6.3. The Rank Theorem. Dimensions of fundamental subspaces.

There are some applications, involving finding bases in the column space or in the null space of a matrix. For example, as it was shown above, solution of a homogeneous equation amounts to finding a basis in the null space $\text{Ker } A$. Finding a basis in a column space means simply is simply extracting a basis from a spanning set, by removing unnecessary vectors (columns).

However, the most important application of the above methods of computing bases of fundamental subspaces is the relations between their dimensions.

Theorem 6.1 (The rank Theorem). *For a matrix A*

$$\text{rank } A = \text{rank } A^T.$$

This theorem is often stated as follows:

The *column rank* of a matrix coincides with its *row rank*.

The proof of this theorem is trivial, since dimensions of both $\text{Ran } A$ and $\text{Ran } A^T$ equal to the number of pivots in the echelon form of A .

The following theorem gives us important relations between dimensions of the fundamental spaces. It is often also called the Rank Theorem

Theorem 6.2. *Let A be an $m \times n$ matrix, i.e. a linear transformation from $\mathbb{R}^n$ to $\mathbb{R}^m$. Then*

1. $\dim \text{Ker } A + \dim \text{Ran } A = \dim \text{Ker } A + \text{rank } A = n$ (*dimension of the domain of A*);
2. $\dim \text{Ker } A^T + \dim \text{Ran } A^T = \dim \text{Ker } A^T + \text{rank } A^T = \dim \text{Ker } A^T + \text{rank } A = m$ (*dimension of the codomain of A*);

Proof. The proof, modulo the above algorithms of finding bases in the fundamental subspaces, is almost trivial. The first statement is simply the fact that the number of free variables ($\dim \text{Ker } A$) plus the number of basic variables (i.e. the number of pivots, i.e. $\text{rank } A$) adds up to the number of columns (i.e. to n).

The second statement, if one takes into account that $\text{rank } A = \text{rank } A^T$ is simply the first statement applied to A^T . $\square$

As one of the application of the above theorem, let us recall the example from Section 5. There we considered a system

$$\begin{pmatrix} 2 & 3 & 1 & 4 & -9 \\ 1 & 1 & 1 & 1 & -3 \\ 1 & 1 & 1 & 2 & -5 \\ 2 & 2 & 2 & 3 & -8 \end{pmatrix} \mathbf{x} = \begin{pmatrix} 17 \\ 6 \\ 8 \\ 14 \end{pmatrix}$$

and we claimed that its general solution given by

$$\mathbf{x} = \begin{pmatrix} 3 \\ 1 \\ 0 \\ 2 \\ 0 \end{pmatrix} + x_3 \begin{pmatrix} -2 \\ 1 \\ 1 \\ 0 \\ 0 \end{pmatrix} + x_5 \begin{pmatrix} 2 \\ -1 \\ 0 \\ 2 \\ 1 \end{pmatrix}, \quad x_3, x_5 \in \mathbb{R}$$

or by

$$\mathbf{x} = \begin{pmatrix} 3 \\ 1 \\ 0 \\ 2 \\ 0 \end{pmatrix} + s \begin{pmatrix} -2 \\ 1 \\ 1 \\ 0 \\ 0 \end{pmatrix} + t \begin{pmatrix} 0 \\ 0 \\ 1 \\ 2 \\ 1 \end{pmatrix}, \quad s, t \in \mathbb{R}$$

We checked in Section 5 that a vector $\mathbf{x}$ given by either formula is indeed a solution of the equation. But how can we guarantee that any of the formulas describe *all* solutions.

First of all we know, that in either formula, the last 2 vectors (the ones multiplied by the parameters) belong to $\text{Ker } A$. It is easy to see, that in either case both vectors are linearly independent (two vectors are linearly dependent if and only if one is a multiple of another).

Now, let us count dimensions: interchanging first and second rows and performing first round of row reduction

$$\begin{array}{l} -2R_1 \\ -R_1 \\ -2R_1 \end{array} \begin{pmatrix} 1 & 1 & 1 & 1 & -3 \\ 2 & 3 & 1 & 4 & -9 \\ 1 & 1 & 1 & 2 & -5 \\ 2 & 2 & 2 & 3 & -8 \end{pmatrix} \sim \begin{pmatrix} 1 & 1 & 1 & 1 & -3 \\ 0 & 1 & -1 & 2 & -3 \\ 0 & 0 & 0 & 1 & -2 \\ 0 & 0 & 0 & 1 & -2 \end{pmatrix}$$

we see that there are three pivots already, so $\text{rank } A \geq 3$. (Actually, we already can see that the rank is 3, but it is enough just to have the estimate here). By Theorem 6.2 $\text{rank } A + \dim \text{Ker } A = 5$, hence $\dim \text{Ker } A \leq 2$, and therefore there cannot be more than 2 linearly independent vectors in $\text{Ker } A$. Therefore, last 2 vectors in either formula form a basis in $\text{Ker } A$, so either formula give all solutions of the equation.

An important corollary of the rank theorem, is the following theorem connecting existence and uniqueness for linear equations.

Theorem 6.3. *Let A be an $m \times n$ matrix. Then the equation*

$$A\mathbf{x} = \mathbf{b}$$

has a solution for every $\mathbf{b} \in \mathbb{R}^m$ if and only if the dual equation

$$A^T \mathbf{x} = \mathbf{0}$$

has unique (only trivial) solution. (Note, that in the second equation we have A^T , not A).

Proof. The proof follows immediately from Theorem 6.2 by counting the dimensions. We leave details as an exercise to the reader. $\square$

There is a very nice geometric interpretation of the second rank theorem (Theorem 6.2). Namely, the statement 1 of the theorem says, that if a transformation $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ has trivial kernel ($\text{Ker } A = \{\mathbf{0}\}$), then the dimensions of the domain $\mathbb{R}^n$ and of the range $\text{Ran } A$ coincide: if the kernel is non-trivial, then the transformation “kills” $\dim \text{Ker } A$ dimensions, so $\dim \text{Ran } A = n - \dim \text{Ker } A$.

7. Representation of a linear transformation in arbitrary bases. Change of coordinates formula.

The material we have learned about linear transformations and their matrices can be easily extended to transformations in abstract vector spaces with finite bases. In this section we will distinguish between a linear transformation T and its matrix, the reason being that we consider different bases, so one linear transformation can have different matrix representation.

7.1. Coordinate vector. Let V be a vector space with a basis $\mathcal{B} := \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n\}$. Any vector $\mathbf{v} \in V$ admits a unique representation as a linear combination

$$\mathbf{v} = x_1 \mathbf{b}_1 + x_2 \mathbf{b}_2 + \dots + x_n \mathbf{b}_n = \sum_{k=1}^n x_k \mathbf{b}_k.$$

Numbers $x_1, x_2, \dots, x_n$ are called *coordinates* of the vector $\mathbf{v}$ in the basis $\mathcal{B}$. It is convenient to join these coordinates into the so-called *coordinate vector* of $\mathbf{v}$ relative to the basis $\mathcal{B}$, which is the column vector

$$[\mathbf{v}]_{\mathcal{B}} := \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} \in \mathbb{R}^n.$$

Note that the mapping

$$\mathbf{v} \mapsto [\mathbf{v}]_{\mathcal{B}}$$

is an isomorphism between V and $\mathbb{R}^n$.

7.2. Matrix of a linear transformation. Let $T : V \rightarrow W$ be a linear transformation, and let $\mathcal{A} = \{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n\}$, $\mathcal{B} := \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_m\}$ be bases in V and W respectively.

A matrix of the transformation T in (or with respect to) the bases $\mathcal{A}$ and $\mathcal{B}$ is an $m \times n$ matrix, denoted by $[T]_{\mathcal{B}\mathcal{A}}$, which relates the coordinate vectors $[T\mathbf{v}]_{\mathcal{B}}$ and $[\mathbf{v}]_{\mathcal{A}}$,

$$[T\mathbf{v}]_{\mathcal{B}} = [T]_{\mathcal{B}\mathcal{A}}[\mathbf{v}]_{\mathcal{A}};$$

notice the balance of symbols $\mathcal{A}$ and $\mathcal{B}$ here: that is the reason we put the first basis $\mathcal{A}$ into the second position.

The matrix $[T]_{\mathcal{B}\mathcal{A}}$ is easy to find: its k th column is just the coordinate vector $[T\mathbf{a}_k]_{\mathcal{B}}$ (compare with finding the matrix of a linear transformation from $\mathbb{R}^n$ to $\mathbb{R}^m$).

As in the case of standard bases, composition of linear transformations is equivalent to multiplication of their matrices: one only has to be a bit more careful about bases. Namely, let $T_1 : X \rightarrow Y$ and $T_2 : Y \rightarrow Z$ be linear transformation, and let $\mathcal{A}, \mathcal{B}$ and $\mathcal{C}$ are bases in X, Y and Z respectively. The for the composition $T = T_2T_1$,

$$T : X \rightarrow Z, \quad T\mathbf{x} := T_2(T_1(\mathbf{x}))$$

we have

$$(7.1) \quad [T]_{\mathcal{C}\mathcal{A}} = [T_2T_1]_{\mathcal{C}\mathcal{A}} = [T_1]_{\mathcal{C}\mathcal{B}}[T_2]_{\mathcal{B}\mathcal{A}}$$

(notice again the balance of indices here).

The proof here goes exactly as in the case of $\mathbb{R}^n$ spaces with standard bases, so we do not repeat it here. Another possibility is to transfer everything to the spaces $\mathbb{R}^n$ via coordinate isomorphisms $\mathbf{v} \mapsto [\mathbf{v}]_{\mathcal{B}}$. And then one does not need any proof, everything follows from the results about matrix multiplication.

7.3. Change of coordinate matrix. Let we have two bases $\mathcal{A} = \{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n\}$ and $\mathcal{B} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n\}$ in a vector space V . Consider the identity transformation $I = I_V$ and its matrix $[I]_{\mathcal{B}\mathcal{A}}$ in these bases. By the definition

$$[\mathbf{v}]_{\mathcal{B}} = [I]_{\mathcal{B}\mathcal{A}}[\mathbf{v}]_{\mathcal{A}}, \quad \forall \mathbf{v} \in V,$$

i.e. for any vector $\mathbf{v} \in V$ the matrix $[I]_{\mathcal{B}\mathcal{A}}$ transforms its coordinates in the basis $\mathcal{A}$ into coordinates in the basis $\mathcal{B}$. The matrix $[I]_{\mathcal{B}\mathcal{A}}$ is often called the *change of coordinates* (from the basis $\mathcal{A}$ to the basis $\mathcal{B}$) matrix.

The matrix $[I]_{\mathcal{B}\mathcal{A}}$ is easy to compute: according to the general rule of finding matrix of a linear transformation, its k th column is the coordinate representation $[\mathbf{a}_k]_{\mathcal{B}}$ of k th element of the basis $\mathcal{A}$

Note that

$$[I]_{\mathcal{A}\mathcal{B}} = ([I]_{\mathcal{B}\mathcal{A}})^{-1},$$

(follows immediately from the multiplication of matrices rule (7.1)), so any change of coordinate matrix is always invertible.

7.3.1. An example: change of coordinates from the standard basis. Let our space V be $\mathbb{R}^n$, and let we have a basis $\mathcal{B} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n\}$ there. We also have the standard basis $\mathcal{S} = \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\}$ there. The change of coordinates matrix $[I]_{\mathcal{S}\mathcal{B}}$ is easy to compute:

$$[I]_{\mathcal{S}\mathcal{B}} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n] =: B,$$

i.e. it is just the matrix B whose k th column is the vector (column) $\mathbf{v}_k$. And in the other direction

$$[I]_{\mathcal{A}\mathcal{B}} = ([I]_{\mathcal{S}\mathcal{B}})^{-1} = B^{-1}.$$

For example, consider a basis

$$\mathcal{B} = \left\{ \begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ 1 \end{pmatrix} \right\}$$

in $\mathbb{R}^2$, and let $\mathcal{S}$ denote the standard basis there. Then

$$[I]_{\mathcal{S}\mathcal{B}} = \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix} =: B$$

and

$$[I]_{\mathcal{B}\mathcal{S}} = [I]_{\mathcal{S}\mathcal{B}}^{-1} = B^{-1} = \frac{1}{3} \begin{pmatrix} -1 & 2 \\ 2 & -1 \end{pmatrix}$$

(we know how to compute inverses, and it is also easy to check that the above matrix is indeed the inverse of B)

7.3.2. An example: going through the standard basis. Let in the space of polynomials of degree at most 1 we have bases

$$\mathcal{A} = \{1, 1+x\}, \quad \text{and} \quad \mathcal{B} = \{1+2x, 1-2x\},$$

and we want to find the change of coordinate matrix $[I]_{\mathcal{B}\mathcal{A}}$.

Of course, we can always take vectors from the basis $\mathcal{A}$ and try to decompose them in the basis $\mathcal{B}$: it involves solving linear systems, and we can do that.

However, I think the following way is simpler. In $\mathbb{P}_1$ we also have the standard basis $\mathcal{S} = \{1, x\}$, and for this basis

$$[I]_{\mathcal{S}\mathcal{A}} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} =: A, \quad [I]_{\mathcal{S}\mathcal{B}} = \begin{pmatrix} 1 & 1 \\ 2 & -2 \end{pmatrix} =: B,$$

and taking the inverses

$$[I]_{\mathcal{A}\mathcal{S}} = A^{-1} = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}, \quad [I]_{\mathcal{B}\mathcal{S}} = B^{-1} = \frac{1}{4} \begin{pmatrix} 2 & 1 \\ 2 & -1 \end{pmatrix}.$$

Then

$$[I]_{\mathcal{B}\mathcal{A}} = [I]_{\mathcal{B}\mathcal{S}}[I]_{\mathcal{S}\mathcal{A}} = B^{-1}A = \frac{1}{4} \begin{pmatrix} 2 & 1 \\ 2 & -1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

and

$$[I]_{\mathcal{A}\mathcal{B}} = [I]_{\mathcal{A}\mathcal{S}}[I]_{\mathcal{S}\mathcal{B}} = A^{-1}B = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 2 & -2 \end{pmatrix}$$

Noice the balace of indices here

7.4. Matrix of a transformation and change of coordinates. Let $T : V \rightarrow W$ be a linear transformation, and let $\mathcal{A}, \tilde{\mathcal{A}}$ be two bases in V and let $\mathcal{B}, \tilde{\mathcal{B}}$ be two bases in W . Suppose we know the matrix $[T]_{\mathcal{B}\mathcal{A}}$, and we would like to find the matrix representation with respect to new bases $\tilde{\mathcal{A}}, \tilde{\mathcal{B}}$, i.e. the matrix $[T]_{\tilde{\mathcal{B}}\tilde{\mathcal{A}}}$. The rule is very simple:

to get the matrix in the “new” bases one has to surround the matrix in the “old” bases by change of coordinates matrctes.

I did not mention here what change of coordinate matrix should go where, because we don’t have any choice if we follow the balance of indices rule. Namely, matrix representation of a linear transformation changes according to the formula

$$[T]_{\tilde{\mathcal{B}}\tilde{\mathcal{A}}} = [I]_{\tilde{\mathcal{B}}\mathcal{B}}[T]_{\mathcal{B}\mathcal{A}}[I]_{\mathcal{A}\tilde{\mathcal{A}}}$$

Notice the balance of indices.

The proof can be done just by analyzing what each of the matrices does.

7.5. Case of one basis: similar matrices. Let V be a vector space and let $\mathcal{A} = \{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n\}$ be a basis in V . Consider a linear transformation $T : V \rightarrow V$ and let $[T]_{\mathcal{A}\mathcal{A}}$ be its matrix in this basis (we use the same basis for “inputs” and “outputs”)

The case when we use the same basis for “inputs” and “outputs” is very important (because in this case we can multiply a matrix by itself), so let us study this case a bit more carefully. Notice, that very often in this case the shorter notation $[T]_{\mathcal{A}}$ is used instead of $[T]_{\mathcal{A}\mathcal{A}}$. However, the two index notation $[T]_{\mathcal{A}\mathcal{A}}$ is better adapted to the balance of indices rule, so I recommend using it (or at least always keep it in mind) when doing change of coordinates.

Let $\mathcal{B} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n\}$ be another basis in V . By the change of coordinate rule above

$$[T]_{\mathcal{B}\mathcal{B}} = [I]_{\mathcal{B}\mathcal{A}}[T]_{\mathcal{A}\mathcal{A}}[I]_{\mathcal{A}\mathcal{B}}$$

$[T]_{\mathcal{A}}$ is often used instead of $[T]_{\mathcal{A}\mathcal{A}}$. It is shorter, but two index notation is better adapted to the balance of indices rule.

Recalling that

$$[I]_{\mathcal{B}\mathcal{A}} = [I]_{\mathcal{A}\mathcal{B}}^{-1}$$

and denoting $Q := [I]_{\mathcal{A}\mathcal{B}}$, we can rewrite the above formula as

$$[T]_{\mathcal{B}\mathcal{B}} = Q^{-1}[T]_{\mathcal{A}\mathcal{A}}Q.$$

This gives a motivation for the following definition

Definition 7.1. We say that a matrix A is similar to a matrix B if there exists an invertible matrix Q such that $A = Q^{-1}BQ$.

Since an invertible matrix must be square, it follows from counting dimensions, that similar matrices A and B has to be square and of the same size. If A is similar to B , i.e. if $A = Q^{-1}BQ$, then

$$B = QAQ^{-1} = (Q^{-1})^{-1}A(Q^{-1})$$

(since Q^{-1} is invertible), therefore B is similar to A . So, we can just say that A and B are *similar*.

The above reasoning shows, that it does not matter where to put Q and where Q^{-1} : one can use the formula $A = QBQ^{-1}$ in the definition of similarity.

The above discussion shows, that one can treat similar matrices as different matrix representation of the same linear operator (transformation).

Determinants

1. Introduction.

The reader probably already met determinants in calculus or algebra, at least the determinants of 2×2 and 3×3 matrices. For a 2×2 matrix

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

the determinant is simply $ad - bc$; the determinant of a 3×3 matrix can be found by the “Star of David” rule.

In this section we would like to introduce determinants for $n \times n$ matrices. I don’t want just to give a formal definition: I want to give some motivation first, then derive some properties the determinant should have, and then to show that if we want to have these properties, then we do not any choice, and arrive to several equivalent definitions of the determinant.

It is more convenient to start not with the determinant of a matrix, but with determinant of a system of vectors: there is no difference, since we always can join vectors together (say as columns) to form a matrix.

Let we have n vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ in $\mathbb{R}^n$ (notice that the number of vectors coincide with dimension), and we want to find the n -dimensional volume of the parallelepiped determined by these vectors.

The parallelepiped determined by the vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ can be defined as the collection of all vectors $\mathbf{v} \in \mathbb{R}^n$ that can be represented as

$$\mathbf{v} = t_1 \mathbf{v}_1 + t_2 \mathbf{v}_2 + \dots + t_n \mathbf{v}_n, \quad 0 \leq t_k \leq 1 \quad \forall k = 1, 2, \dots, n.$$

It can be easily visualized when $n = 2$ (parallelogram) and $n = 3$ (parallelepiped). So, what is the n -dimensional volume?

If $n = 2$ it is area; if $n = 3$ it is indeed the volume. In dimension 1 it is just the length.

And finally some notation. For a system of vectors (columns) $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ we will denote its determinant (that we are going to construct) as $D(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n)$. If we join these vectors in a matrix A (column number k of A is $\mathbf{v}_k$), then we will use the notation $\det A$,

$$\det A = D(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n)$$

Also, for a matrix

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,n} \\ a_{2,1} & a_{2,2} & \dots & a_{2,n} \\ a_{n,1} & a_{n,2} & \dots & a_{n,n} \end{pmatrix}$$

its determinant is often is denoted by

$$\begin{vmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,n} \\ a_{2,1} & a_{2,2} & \dots & a_{2,n} \\ a_{n,1} & a_{n,2} & \dots & a_{n,n} \end{vmatrix}.$$

2. What properties determinant should have.

We know, that for dimensions 2 and 3 “volume” of a parallelepiped is determined by the *base times height* rule: if we pick one vector, then height is the distance from this vector to the subspace spanned by the remaining vectors, and the base is the $(n-1)$ -dimensional volume of the parallelepiped determined by the remaining vectors.

Now let us generalize this idea to higher dimensions. For a moment we do not care about how exactly to determine height and base: we will show, that if we assume that base and height satisfy some natural properties, then we do not have any choice, and the volume (determinant) is uniquely defined.

2.1. Linearity in each argument. First of all, if we multiply vector $\mathbf{v}_1$ by a positive number a , then the height (i.e. the distance to the linear span $\mathcal{L}(\mathbf{v}_2, \dots, \mathbf{v}_n)$) is multiplied by a . If we admit negative heights, then this property holds for all scalars a , and so for the determinant $D(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n)$ of the system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ should satisfy

$$D(\alpha \mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n) = \alpha D(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n).$$

Of, course, there is nothing special about vector $\mathbf{v}_1$, so for any index k

$$(2.1) \quad D(\mathbf{v}_1, \dots, \underset{k}{\alpha \mathbf{v}_k}, \dots, \mathbf{v}_n) = \alpha D(\mathbf{v}_1, \dots, \underset{k}{\mathbf{v}_k}, \dots, \mathbf{v}_n)$$

To get the next property, let us notice that if we add 2 vectors, then the “height” of the result should be equal the sum of the “heights” of summands,

i.e. that

$$(2.2) \quad D(\mathbf{v}_1, \dots, \underbrace{\mathbf{u}_k + \mathbf{v}_k}_k, \dots, \mathbf{v}_n) = D(\mathbf{v}_1, \dots, \underbrace{\mathbf{u}_k}_k, \dots, \mathbf{v}_n) + D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_k}_k, \dots, \mathbf{v}_n)$$

In other words, the above two properties say that the determinant of n vectors is *linear in each argument (vector)*, meaning that if we fix $n - 1$ vectors and interpret the remaining vector as a variable (argument), we get a linear function.

Remark. We already know that *linearity* is a very nice property, that helps in many situations. So, admitting negative heights (and therefore negative volumes) is a very small parice to pay to get linearity: we can always put absolute value afterwards.

In fact, by admitting negative heights, we did not sacrifice anything! Contrary, we even gained something, because sign of the determinant contains some information about the system of vectors (orientation).

2.2. Preservation under “column replacement”. The next property also seems natural. Namely, if we take a vector, and add to it a multiple of another vector, the “height” does not change, so

$$(2.3) \quad D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_j + \alpha \mathbf{v}_k}_j, \dots, \underbrace{\mathbf{v}_k}_k, \dots, \mathbf{v}_n) = D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_j}_j, \dots, \underbrace{\mathbf{v}_k}_k, \dots, \mathbf{v}_n)$$

In other words, if we apply the *column operation* of the third type, the determinant does not change.

Remark. Although it is not essential here, let us notice that the second part of linearity (property (2.2)) is not independent: it can be deduced from properties (2.1) and (2.3).

We leave the proof as an exercise for the reader.

2.3. Antisymmetry. The next property the determinant should have, is that if we interchange 2 vectors, the determinant changes sign:

$$(2.4) \quad D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_k}_j, \dots, \underbrace{\mathbf{v}_j}_k, \dots, \mathbf{v}_n) = -D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_j}_j, \dots, \underbrace{\mathbf{v}_k}_k, \dots, \mathbf{v}_n).$$

At the first sight this property does not look natural, but it can be deduced from the previous ones. Namely, applying property (2.3) three

Functions of several variables that change sign when one interchanges any two arguments are called *antisymmetric*.

times, and then using (2.1) we get

$$\begin{aligned}
D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_j}_{j}, \dots, \underbrace{\mathbf{v}_k}_{k}, \dots, \mathbf{v}_n) &= \\
&= D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_j}_{j}, \dots, \underbrace{\mathbf{v}_k - \mathbf{v}_j}_{k}, \dots, \mathbf{v}_n) \\
&= D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_j + (\mathbf{v}_k - \mathbf{v}_j)}_j, \dots, \underbrace{\mathbf{v}_k - \mathbf{v}_j}_k, \dots, \mathbf{v}_n) \\
&= D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_k}_{j}, \dots, \underbrace{\mathbf{v}_k - \mathbf{v}_j}_k, \dots, \mathbf{v}_n) \\
&= D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_k}_{j}, \dots, \underbrace{(\mathbf{v}_k - \mathbf{v}_j) - \mathbf{v}_k}_k, \dots, \mathbf{v}_n) \\
&= D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_k}_{j}, \dots, \underbrace{-\mathbf{v}_j}_k, \dots, \mathbf{v}_n) \\
&= -D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_k}_{j}, \dots, \underbrace{\mathbf{v}_j}_k, \dots, \mathbf{v}_n).
\end{aligned}$$

2.4. Normalization. And the last property is the easiest one: for the standard basis $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ in $\mathbb{R}^n$ the corresponding parallelepiped is the n -dimensional unit cube, so

$$(2.5) \quad D(\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n) = 1.$$

In matrix notation this can be written as

$$\det(I) = 1$$

3. Constructing the determinant.

The plan of the game is now as follows: using the properties that as we decided in Section 2 the determinant should have, we derive other properties of the determinant, some of them highly non-trivial. We will show how to use these properties to compute the determinant using our old friend—row reduction.

And later, in Section 4 we will show that the determinant, i.e. a function with the desired properties exists and unique: we have to be sure that the object we are computing and studying exists.

3.1. Basic properties. We will use the following basic properties of the determinant:

1. Determinant is linear in each column, i.e. in vector notation for every index k

$$D(\mathbf{v}_1, \dots, \underbrace{\alpha \mathbf{u}_k + \beta \mathbf{v}_k}_k, \dots, \mathbf{v}_n) = \alpha D(\mathbf{v}_1, \dots, \mathbf{u}_k, \dots, \mathbf{v}_n) + \beta D(\mathbf{v}_1, \dots, \mathbf{v}_k, \dots, \mathbf{v}_n)$$

for all scalars α, β .

2. Determinant is *antisymmetric*, i.e. if one interchanges two columns, the determinant changes sign.
3. Normalization property: $\det I = 1$.

All these properties were discussed above in Section 2. The first property is just the (2.1) and (2.2) combined. The second one is (2.4), and the last one is the normalization property (2.5). Note, that we did not use property (2.3): it can be deduced from the above three. These three properties completely define determinant!

3.2. Properties of determinant deduced from the basic properties.

Proposition 3.1. *For a square matrix A the following statements hold:*

1. *If A has a zero column, then $\det A = 0$.*
2. *If A has two equal columns, then $\det A = 0$;*
3. *If one column of A is a multiple of another, then $\det A = 0$;*
4. *If columns of A are linearly dependent, i.e. if the matrix is not invertible, then $\det A = 0$.*

Proof. Statement 1 follows immediately from linearity. If we multiply the zero column by zero, we do not change the matrix and its determinant. But by the property 1 above, we should get 0.

The fact that determinant is antisymmetric, implies statement 2. Indeed, if we interchange two equal columns, we change nothing, so determinant remains the same. On the other hand, interchanging two columns changes sign of determinant, so

$$\det A = -\det A,$$

which is possible only if $\det A = 0$.

Statement 3 is immediate corollary of statement 2 and linearity.

To prove the last statement, let us first suppose that the first vector $\mathbf{v}_1$ is a linear combination of the other vectors,

$$\mathbf{v}_1 = \alpha_2 \mathbf{v}_2 + \alpha_3 \mathbf{v}_3 + \dots + \alpha_n \mathbf{v}_n = \sum_{k=2}^n \alpha_k \mathbf{v}_k.$$

Then by linearity we have (in vector notation)

$$\begin{aligned} D(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n) &= D\left(\left(\sum_{k=2}^n \alpha_k \mathbf{v}_k\right), \mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_n\right) \\ &= \sum_{k=2}^n \alpha_k D(\mathbf{v}_k, \mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_n) \end{aligned}$$

and each determinant in the sum is zero because of two equal columns.

Let us now consider general case, i.e. let us assume that the system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ is linearly dependent. Then one of the vectors, say $\mathbf{v}_k$ can be represented as a linear combination of the others. Interchanging this vector with $\mathbf{v}_1$ we arrive to the situation we just treated, so

$$D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_k}_{\mathbf{v}_1}, \dots, \mathbf{v}_n) = -D(\mathbf{v}_k, \dots, \underbrace{\mathbf{v}_1}_{\mathbf{v}_k}, \dots, \mathbf{v}_n) = -0 = 0,$$

so the determinant in this case is also 0. $\square$

The next proposition generalizes property (2.3). As we already have said above, this property can be deduced from the three “basic” properties of the determinant, we are using in this section.

Note, that adding to a column a multiple of itself is prohibited here. We can only add multiples of the other columns.

Proposition 3.2. *The determinant does not change if we add to a column a linear combination of the other columns (leaving the other columns intact). In particular, the determinant is preserved under “column replacement” (column operation of third type).*

Proof. Fix a vector $\mathbf{v}_k$, and let $\mathbf{u}$ be a linear combination of the other vectors,

$$\mathbf{u} = \sum_{j \neq k} \alpha_j \mathbf{v}_j.$$

Then by linearity

$$D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_k + \mathbf{u}}_{\mathbf{v}_k}, \dots, \mathbf{v}_n) = D(\mathbf{v}_1, \dots, \underbrace{\mathbf{v}_k}_{\mathbf{v}_k}, \dots, \mathbf{v}_n) + D(\mathbf{v}_1, \dots, \underbrace{\mathbf{u}}_{\mathbf{0}}, \dots, \mathbf{v}_n),$$

and by Proposition 3.2 the last term is zero. $\square$

3.3. Determinants of diagonal and triangular matrices. Now we are ready to compute determinant for some important special classes of matrices. The first class is the so-called *diagonal* matrices. Let us recall that a square matrix $A = \{a_{j,k}\}_{j,k=1}^n$ is called *diagonal* if all entries *off the main diagonal* are zero, i.e. if $a_{j,k} = 0$ for all $j \neq k$. We will often use the notation

$\text{diag}(a_1, a_2, \dots, a_n)$ for the diagonal matrix

$$\begin{pmatrix} a_1 & 0 & \dots & 0 \\ 0 & a_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & a_n \end{pmatrix}$$

Since a diagonal matrix $\text{diag}(a_1, a_2, \dots, a_n)$ can be obtained from the identity matrix I by multiplying column number k by a_k ,

Determinant of a diagonal matrix equal the product of the diagonal entries,

$$\det(\text{diag}(a_1, a_2, \dots, a_n)) = a_1 a_2 \dots a_n.$$

The next important class is the class of so-called *triangular* matrices. A square matrix $A = \{a_{j,k}\}_{j,k=1}^n$ is called *upper triangular* if all entries *below* the main diagonal are 0, i.e. if $a_{j,k} = 0$ for all $k < j$. A square matrix is called *lower triangular* if all entries *above* the main are 0, i.e. if $a_{j,k} = 0$ for all $j < k$. We call a matrix *triangular*, if it is either lower or upper triangular matrix.

It is easy to see that

Determinant of a *triangular* matrix equals to the product of the diagonal entries,

$$\det A = a_{1,1} a_{2,2} \dots a_{n,n}.$$

Indeed, if a triangular matrix has zero on the main diagonal, it is not invertible (can be easily checked by row operation) and therefore both sides equal zero. If all diagonal entries are non-zero, then using column replacement (column operations of third type) one can transform the matrix into a diagonal one with the same diagonal entries: For upper triangular matrix one should first subtract appropriate multiples of the first column from the columns number $2, 3, \dots, n$, “killing” all entries in the first row, then subtract appropriate multiples of the second column from columns number $3, \dots, n$, and so on.

To treat the case of lower triangular matrices one has to do “column reduction” from the left to the right, i.e. first subtract appropriate multiples of the last column from columns number $n-1, \dots, 2, 1$, and so on.

3.4. Computing the determinant. Now we know how to compute determinants, using their properties: one just need to do column reduction (i.e. row reduction for A^T) keeping track of column operations changing the determinant. Fortunately, the most often used operation—row replacement, i.e. operation of third type—does not change the determinant. So we

only need to keep track of interchanging of columns and of multiplication of column by a scalar.

If echelon form of A^T does not have pivots in every column (and row), then A is not invertible, so $\det A = 0$. If A is invertible, we arrive to triangular matrix, and $\det A$ is the product of diagonal entries times the correction from column interchanges and multiplications.

The above algorithm implies that $\det A$ can be zero only if a matrix A is not invertible. Combining this with the last statement of Proposition 3.1 we get

Proposition 3.3. *$\det A = 0$ if and only if A is invertible. Equivalent statement: $\det A \neq 0$ if and only if A is invertible.*

Note, that although we now know how to compute determinants, the determinant is still not defined. One can ask: why don't we define it as the result we get from the above algorithm? The problem is that formally this result is not well defined: that means we did not prove that different sequences of column operations yield same answer.

3.5. Determinants of a transpose and of a product. Determinants of elementary matrices.

In this section we prove two important theorems

Theorem 3.4 (Determinant of a transpose). *For a square matrix A ,*

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

This theorem implies that for all statement about columns we discussed above, the corresponding statements about rows are also true. In particular, determinants behave under *row operations* the same way they behave under *column operations*. So, we can use row operations to compute determinants.

Theorem 3.5 (Determinant of a product). *For $n \times n$ matrices A and B*

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

In other words

Determinant of a product equals product of determinants.

To prove both theorem we need the following lemma

Lemma 3.6. *For a square matrix A and an elementary matrix E (of the same size)*

$$\det(AE) = (\det A)(\det E)$$

Proof. The proof can be done just by direct checking: determinants of special matrices are easy to compute; right multiplication by an elementary matrix is a column operation, and effect of column operations on the determinant is well known.

This can look like a lucky coincidence, that the determinants of elementary matrices agree with the corresponding column operations, but it is not a coincidence at all.

Namely, for a column operation the corresponding elementary matrix can be obtained from the identity matrix I by this column operation. So, its determinant is 1 (determinant of I) times the effect of the column operation.

And that is all! It may be hard to realize at first, but above paragraph is a *complete and rigorous* proof of the lemma! $\square$

Applying N times Lemma 3.6 we get the following corollary.

Corollary 3.7. *For any matrix A and any sequence of elementary matrices $E_1, E_2, \dots, E_N$ (all matrices are $n \times n$)*

$$\det(AE_1E_2 \dots E_N) = (\det A)(\det E_1)(\det E_2) \dots (\det E_N)$$

Lemma 3.8. *Any invertible matrix is a product of elementary matrices.*

Proof. We know that any invertible matrix is row equivalent to the identity matrix, which is its reduced echelon form. So

$$I = E_N H_{N-1} \dots E_2 E_1 A,$$

and therefore any invertible matrix can be represented as a product of elementary matrices,

$$A = E_1^{-1} E_2^{-1} \dots E_{N-1}^{-1} I = E_1^{-1} E_2^{-1} \dots E_{N-1}^{-1}$$

(inverse of an elementary matrix is an elementary matrix). $\square$

Proof of Theorem 3.4. First of all, it can be easily checked, that for an elementary matrix E we have $\det E = \det(E^T)$. Notice, that it is sufficient to prove theorem only for invertible matrices A , since if A is not invertible then A^T is also not invertible, and both determinants are zero.

By Lemma 3.8 matrix A can be represented as a product of elementary matrices,

$$A = E_1 E_2 \dots E_N,$$

and by Corollary 3.7 the determinant of A is the product of determinants of the elementary matrices. Since taking the transpose just transposes each elementary matrix and reverses their order, Corollary 3.7 implies that $\det A = \det A^T$. $\square$

Proof of Theorem 3.5. Let us first suppose that the matrix B is invertible. Then Lemma 3.8 implies that B can be represented as a product of elementary matrices

$$B = E_1 E_2 \dots E_N,$$

and so by Corollary 3.7

$$\det(AB) = (\det A)[(\det E_1)(\det E_2) \dots (\det E_N)] = (\det A)(\det B).$$

If B is not invertible, then the product AB is also not invertible, and the theorem just says that $0 = 0$.

To check that the product $AB = C$ is not invertible, let us assume that it is invertible. Then multiplying the identity $AB = C$ by C^{-1} from the left, we get $C^{-1}AB = I$, so $C^{-1}A$ is a left inverse of B . So B is left invertible, and since it is square, it is invertible. We got a contradiction. $\square$

3.6. Summary of properties of determinant. First of all, let us say once more, that *determinant is defined only for square matrices!* Since we now know that $\det A = \det(A^T)$, the statements that we knew about columns, true for rows too.

1. Determinant is linear in each row (column) when the other rows (columns) are fixed.
2. If one interchanges two rows (columns) of a matrix A , the determinant changes sign.
3. For a triangular (in particular, for a diagonal) matrix its determinant is the product of the diagonal entries. In particular, $\det I = 1$.
4. If a matrix A has a zero row (or column), $\det A = 0$.
5. If a matrix A has two equal rows (columns), $\det A = 0$.
6. If one of the rows (columns) of A is a linear combination of the other rows (columns), i.e. if the matrix is not invertible, then $\det A = 0$;
More generally
7. $\det A = 0$ if and only if A is not invertible, or equivalently
8. $\det A \neq 0$ if and only if A is invertible.
9. $\det A$ does not change if we add to a row (column) a linear combination of the other rows (columns). In particular, the determinant is preserved under the row (column) replacement, i.e. under the row (column) operation of the third kind.
10. $\det A^T = \det A$.
11. $\det(AB) = (\det A)(\det B)$.
And finally,
12. If A is an $n \times n$ matrix, then $\det(aA) = a^n \det A$.

The last property follows from the linearity of the determinant, if we recall that to multiply a matrix A by a we have to multiply by a each row. And each multiplication multiplies the determinant by a .

4. Formal definition. Existence and uniqueness of the determinant.

In this section we arrive to the formal definition of the determinant. We show that a function, satisfying the basic properties 1, 2, 3 from Section 3 exists, and moreover, such function is unique, i.e. we do not have any choice in constructing the determinant.

Consider an $n \times n$ matrix $A = \{a_{j,k}\}_{j,k=1}^n$, and let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ be its columns, i.e.

$$\mathbf{v}_k = \begin{pmatrix} a_{1,k} \\ a_{2,k} \\ \vdots \\ a_{n,k} \end{pmatrix} = a_{1,k}\mathbf{e}_1 + a_{2,k}\mathbf{e}_2 + \dots + a_{n,k}\mathbf{e}_n = \sum_{j=1}^n a_{j,k}\mathbf{e}_j.$$

Using linearity of the determinant we expand it in the first column $\mathbf{v}_1$:

$$(4.1) \quad D(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n) = D\left(\sum_{j=1}^n a_{j,1}\mathbf{e}_j, \mathbf{v}_2, \dots, \mathbf{v}_n\right) = \sum_{j=1}^n a_{j,1}D(\mathbf{e}_j, \mathbf{v}_2, \dots, \mathbf{v}_n).$$

Then we expand it in the second column, then in the third, and so on. We get

$$D(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n) = \sum_{j_1=1}^n \sum_{j_2=1}^n \dots \sum_{j_n=1}^n a_{j_1,1}a_{j_2,2} \dots a_{j_n,n} D(\mathbf{e}_{j_1}, \mathbf{e}_{j_2}, \dots, \mathbf{e}_{j_n}).$$

Notice, that we have to use a different index of summation for each column: we call them $j_1, j_2, \dots, j_n$; the index j_1 here is the same as the index j in (4.1).

It is a huge sum, it contains n^n terms. Fortunately, some of the terms are zero. Namely, if any 2 of the indices $j_1, j_2, \dots, j_n$ coincide, the determinant $D(\mathbf{e}_{j_1}, \mathbf{e}_{j_2}, \dots, \mathbf{e}_{j_n})$ is zero, because there are two equal rows here.

So, let us rewrite the sum, omitting all non-zero terms. The most convenient way to do that is using the notion of a *permutation*. A permutation of an ordered set $\{1, 2, \dots, n\}$ is a rearrangement of its elements. A convenient way to represent a permutation is by using a function

$$\sigma : \{1, 2, \dots, n\} \rightarrow \{1, 2, \dots, n\},$$

where $\sigma(k)$ gives new position of the element number k .

Such function σ has to be one-to-one (different values for different arguments) and onto (takes all values from the codomain). Such functions (one-to-one and onto) are called *bijections*, and they give one-to-one correspondence between two sets.

Although it is not directly relevant here, let us notice, that it is well-known in combinatorics, that the number of different perturbations of the set $\{1, 2, \dots, n\}$ is exactly $n!$. The set of all permutations of the set $\{1, 2, \dots, n\}$ will be denoted $\text{Perm}(n)$.

Using the notion of a permutation, we can rewrite the determinant as

$$D(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n) = \sum_{\sigma \in \text{Perm}(n)} a_{\sigma(1),1} a_{\sigma(2),2} \dots a_{\sigma(n),n} D(\mathbf{e}_{\sigma(1)}, \mathbf{e}_{\sigma(2)}, \dots, \mathbf{e}_{\sigma(n)}).$$

The matrix with columns $\mathbf{e}_{\sigma(1)}, \mathbf{e}_{\sigma(2)}, \dots, \mathbf{e}_{\sigma(n)}$ can be obtained from the identity matrix by finitely many column interchanges, so the determinant

$$D(\mathbf{e}_{\sigma(1)}, \mathbf{e}_{\sigma(2)}, \dots, \mathbf{e}_{\sigma(n)})$$

is 1 or -1 depending on the number of column interchanges.

To formalize that, we define *sign* σ of a permutation σ to be 1 if even number of interchanges is necessary to rearrange the n -tuple $1, 2, \dots, n$ into $\sigma(1), \sigma(2), \dots, \sigma(n)$, and $\text{sign}(\sigma) = -1$ if the number of interchanges is odd.

It is a well-known fact from the combinatorics, that the signum of permutation is well defined, i.e. that although there are infinitely many ways to get the n -tuple $\sigma(1), \sigma(2), \dots, \sigma(n)$ from $1, 2, \dots, n$, the number of interchanges is either always odd or always even.

One of the ways to show that is to count the number K of pairs j, k , $j < k$ such that $\sigma(j) > \sigma(k)$, and see if the number is even or odd. We call the permutation *odd* if K is odd and *even* if K is even. Then define *signum* of σ to be $(-1)^K$. We want to show that signum and sign coincide, so sign is well defined.

If $\sigma(k) = k \forall k$, then the number of such pairs is 0, so *signum* of such *identity* permutation is 1. Note also, that any elementary transpose, which interchange two neighbors, changes the signum of a permutation, because it changes (increases or decreases) the number of the pairs exactly by 1. So, to get from a permutation to another one always need an even number of elementary transposes if the permutation have the same signum, and an odd number if the signums are different.

Finally, any interchange of two entries can be achieved by an odd number of elementary transposes. This implies that *signum* changes under an

interchange of two entries. So, to get from $1, 2, \dots, n$ to an even permutation (positive signum) one always need even number of interchanges, and odd number of interchanges is needed to get an odd permutation (negative signum). That means *signum* and *sign* coincide, and so *sign* is well defined.

So, if we want determinant to satisfy basic properties 1–3 from Section 3, we must define it as

$$(4.2) \quad \det A = \sum_{\sigma \in \text{Perm}(n)} a_{\sigma(1),1} a_{\sigma(2),2} \dots a_{\sigma(n),n} \text{sign}(\sigma),$$

where the sum is taken over all permutations of the set $\{1, 2, \dots, n\}$.

If we define the determinant this way, it is easy to check that it satisfies the basic properties 1–3 from Section 3. Indeed, it is linear in each column, because for each column every term (product) in the sum contains exactly one entry from this column.

Interchanging two columns of A just adds an extra interchange to the perturbation, so right side in (4.2) changes sign. And finally, for the identity matrix I , the right side of (4.2) is 1 (it has one non-zero term).

5. Cofactor expansion.

For an $n \times n$ matrix $A = \{a_{j,k}\}_{j,k=1}^n$ let $A_{j,k}$ denotes the $(n-1) \times (n-1)$ matrix obtained from A by crossing out row number j and column number k .

Theorem 5.1 (Cofactor expansion of determinant). *Let A be an $n \times n$ matrix. For each j , $1 \leq j \leq n$, determinant of A can be expanded in the row number j as*

$$\begin{aligned} \det A &= a_{j,1}(-1)^{j+1} \det A_{j,1} + a_{j,2}(-1)^{j+2} \det A_{j,2} + \dots + a_{j,n}(-1)^{j+n} \det A_{j,n} \\ &= \sum_{k=1}^n a_{j,k}(-1)^{j+k} \det A_{j,k}. \end{aligned}$$

Similarly, for each k , $1 \leq k \leq n$, the determinant can be expanded in the column number k ,

$$\det A = \sum_{j=1}^n a_{j,k}(-1)^{j+k} \det A_{j,k}$$

Proof. Let us prove formula for the expansion in the row number 1: the formula for expansion in the row number k then can be obtained from it by interchanging rows number 1 and k . And since $\det A = \det A^T$, column expansion follows automatically.

Let us first consider a special case, when the first row has one non zero term $a_{1,1}$. Performing column operations on columns $2, 3, \dots, n$ we transform A to the lower triangular form. The determinant of A then can be computed as

$$\boxed{\begin{array}{l} \text{the product of diagonal} \\ \text{entries of the triangular} \\ \text{matrix} \end{array}} \times \boxed{\begin{array}{l} \text{correcting factor from} \\ \text{the column operations} \end{array}}.$$

But the product of all diagonal entries except the first one (i.e. without $a_{1,1}$) times the correcting factor is exactly $\det A_{1,1}$, so in this particular case $\det A = a_{1,1} \det A_{1,1}$.

Let us now consider the case when all entries in the first row except $a_{1,2}$ are zeroes. This case can be reduced to the previous one by interchanging columns number 1 and 2, and therefore in this case $\det A = (-1) \det A_{1,2}$.

The case when $a_{1,3}$ is the only non-zero entry in the first row, can be reduced to the previous one by interchanging rows 2 and 3, so in this case $\det A = a_{1,3} \det A_{1,3}$.

Repeating this procedure we get that in the case when $a_{1,k}$ is the only non-zero entry in the first row $\det A = (-1)^{1+k} \det A_{1,k}$.

In the general case, linearity of the determinant implies that

$$\det A = \det A^{(1)} + \det A^{(2)} + \dots + \det A^{(n)} = \sum_{k=1}^n \det A^{(k)}$$

where the matrix $A^{(k)}$ is obtained from A by replacing all entries in the first row except $a_{1,k}$ by 0. As we just discussed above

$$\det A^{(k)} = (-1)^{1+k} \det A_{1,k},$$

so

$$\det A = \sum_{k=1}^n (-1)^{1+k} \det A_{1,k}.$$

□

The numbers

$$C_{j,k} = (-1)^{j+k} \det A_{j,k}$$

are called *cofactors*. Using this notation formula for expansion of the determinant in the row number j can be rewritten as

$$\det A = a_{j,1}C_{j,1} + a_{j,2}C_{j,2} + \dots + a_{j,n}C_{j,n} = \sum_{k=1}^n a_{j,k}C_{j,k}.$$

Similarly, expansion in the row number k can be written as

$$\det A = a_{1,k}C_{1,k} + a_{2,k}C_{2,k} + \dots + a_{n,k}C_{n,k} = \sum_{j=1}^n a_{j,k}C_{j,k}$$

Remark. Very often the cofactor expansion formula is used as the definition of determinant. It is not difficult to show that the quantity given by this formula satisfies the basic properties of the determinant: the normalization property is trivial, the proof of antisymmetry is easy. However, the proof of linearity is a bit tedious (although not too difficult).

Very often the cofactor expansion formula is used as the definition of determinant.

Remark. Although it looks very nice, the cofactor expansion formula is not suitable for computing determinant of matrices bigger than 3×3 .

As one can count it requires $n!$ multiplications, and $n!$ grows very rapidly. For example, cofactor expansion of a 20×20 matrix require $20! \approx 2.4!10^{18}$ multiplications: it would take a computer performing a *billion* mutiplications per second over 77 years to perform the multiplications.

On the other hand, computing the determinant of an $n \times n$ matrix using *row reduction* requires $(n^3 + 2n - 3)/3$ multiplications (and about the same number of additions). It would take a computer performing a *million* operations per second (very slow, by today's standards) a fraction of a second to compute the determinant of a 100×100 matrix by row reduction.

It only can be practical to apply the cofactor expansion formula in higher dimensions if a row (or a column) has a lot of zero entries.

However, the cofactor exansion formula is of great theoretical importance, as the next section shows.

5.1. Cofactor formula for the inverse matrix. The matrix $C = \{C_{j,k}\}_{j,k=1}^n$ whose entries are cofactors of a given matrix A is called the *cofactor matrix* of A .

Theorem 5.2. *Let A be an invertible matrix and let C be its cofactor matrix. Then*

$$A^{-1} = \frac{1}{\det A} C^T.$$

Proof. Let us find the product AC^T . The diagonal entry number j is obtained by multiplying j th row of A by j th column of A (i.e. j th row of C), so

$$(AC^T)_{j,j} = a_{j,1}C_{j,1} + a_{j,2}C_{j,2} + \dots + a_{j,n}C_{j,n} = \det A,$$

by the cofactor expansion formula.

To get off diagonal terms we need to multiply j th row of A by k th column of C^T , $j \neq k$, to get

$$a_{j,1}C_{k,1} + a_{j,2}C_{k,2} + \dots + a_{j,n}C_{k,n}.$$

It follows from the cofactor expansions formula (expanding in k th row) that this is the determinant of the matrix obtained from A by replacing row number k by the row number j (and leaving all other rows as they were). But the columns j and k of this matrix coincide, so the determinant is 0. So, all off-diagonal entries of AC^T are zeroes (and all diagonal ones equal $\det A$), thus

$$AC^T = \det A.$$

That means that the matrix $\frac{1}{\det A} C^T$ is a right inverse of A , and since A is square, it is the inverse. $\square$

Recalling that for an invertible matrix A the equation $A\mathbf{x} = \mathbf{b}$ has a unique solution

$$\mathbf{x} = A^{-1}\mathbf{b} = \frac{1}{\det A} C^T \mathbf{b},$$

we get the following corollary of the above theorem.

Corollary 5.3 (Cramer's rule). *For an invertible matrix A the entry number k of the solution of the equation $A\mathbf{x} = \mathbf{b}$ is given by the formula*

$$x_k = \frac{\det B_k}{\det A},$$

where the matrix B_k is obtained from A by replacing column number k of A by the vector $\mathbf{b}$.

5.2. Some applications of the cofactor formula for the inverse.

Example (Inverting 2×2 matrices). The cofactor formula really shines when one needs to invert a 2×2 matrix

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

The cofactors are just entries (1×1 matrices), the cofactor matrix is

$$\begin{pmatrix} d & -b \\ -c & a \end{pmatrix},$$

so the inverse matrix A^{-1} is given by the formula

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

While the cofactor formula for the inverse does not look practical or dimensions higher than 3, it has a great theoretical value, as the examples below illustrate.

Example (Matrix with integer inverse). Suppose that we want to construct a matrix A with integer entries, such that its inverse also have integer entries (inverting such matrix would make a nice homework problem: no messing with fractions). If $\det A = 1$ and its entries are integer, the cofactor formula for inverses implies that A^{-1} also have integer entries.

Note, that it is easy to construct an integer matrix A with $\det A = 1$: one should start with a triangular matrix with 1 on the main diagonal, and then apply several row or column replacements (operations of the third type) to make the matrix look generic.

Example (Inverse of a polynomial matrix). Another example is to consider a *polynomial matrix* $A(x)$, i.e. a matrix whose entries are not numbers but polynomials $a_{j,k}(x)$ of the variable x . If $\det A(x) \equiv 1$, then the inverse matrix $A^{-1}(x)$ is also a polynomial matrix.

If $\det A(x) = p(x) \neq 0$, it follows from the cofactor expansion that $p(x)$ is a polynomial, so $A^{-1}(x)$ has rational entries: moreover, $p(x)$ is a multiple of each denominator.

6. Minors and rank.

For a matrix A let us consider its $k \times k$ *submatrix*, obtained by taking k rows and k columns. The determinant of this matrix is called a *minor* of order k . Note, that an $m \times n$ matrix has $\binom{m}{k} \cdot \binom{n}{k}$ different $k \times k$ submatrices, and so it has $\binom{m}{k} \cdot \binom{n}{k}$ minors of order k .

Theorem 6.1. *For a non-zero matrix A its rank equals to the maximal integer k such that there exists a non-zero minor of order k .*

Proof. Let us first show, that if $k > \text{rank } A$ then all minors of order k are 0. Indeed, since the dimension of the column space $\text{Ran } A$ is $\text{rank } A < k$, any k columns of A are linearly dependent. Therefore, for any $k \times k$ submatrix of A its column are linearly dependent, and so all minors of order k are 0.

To complete the proof we need to show that there exists a non-zero minor of order $k = \text{rank } A$. There can be many such minors, but probably the easiest way to get such a minor is to take pivot rows and pivot column (i.e. rows and columns of the original matrix, containing a pivot). This $k \times k$ submatrix has the same pivots as the original matrix, so it is invertible (pivot in every column and every row) and its determinant is non-zero. $\square$

This theorem does not look very useful, because it is much easier to perform row reduction than to compute all minors. However, it is of great theoretical importance, as the following corollary shows.

Corollary 6.2. *Let $A = A(x)$ be an $m \times n$ polynomial matrix (i.e. a matrix whose entries are polynomials of x). Then $\text{rank } A(x)$ is constant everywhere, except maybe finitely many points.*

Proof. Let r be the largest integer such that $\text{rank } A(x) = r$ for some x . To show that such r exists, we first try $r = \min\{m, n\}$. If there exists x such that $\text{rank } A(x) = r$, we found r . If not, we replace r by $r - 1$ and try again. After finitely many steps we either stop or hit 0. So, r exists.

Let x_0 be a point such that $\text{rank } A(x_0) = r$, and let M be a minor of order k such that $M(x_0) \neq 0$. Since $M(x)$ is the determinant of a $k \times k$ polynomial matrix, $M(x)$ is a polynomial. Since $M(x_0) \neq 0$, it is not identically zero, so it can be zero only at finitely many points. So, everywhere except maybe finitely many points $\text{rank } A(x) \geq r$. But by the definition of r , $\text{rank } A(x) \leq r$ for all x .

□

Spectral theory (eigenvalues and eigenvectors)

The spectral theory is the main tool that helps us to understand the structure of a linear operator. In this chapter we consider only operators acting from a vector space to itself (or, equivalently, $n \times n$ matrices). If we have such a linear transformation $A : V \rightarrow V$, we can multiply it to itself, take any power of it, any polynomial.

The main idea of the spectral theory is to split the operator in the simple blocks to analyze each block separately.

To explain the main idea, let us consider *difference equations*. Many processes can be described by the equations of the following type

$$\mathbf{x}_{n+1} = A\mathbf{x}_n, \quad n = 0, 1, 2, \dots,$$

where $A : V \rightarrow V$ is a linear transformation, and $\mathbf{x}_n$ is the state of the system at the time n . Given the initial state $\mathbf{x}_0$ we would like to know the state $\mathbf{x}_n$ at the time n , analyze the long time behaviour of $\mathbf{x}_n$, etc.¹

At the first look the problem looks trivial: the solution $\mathbf{x}_n$ is given by the formula $\mathbf{x}_n = A^n \mathbf{x}_0$. But what if n is huge: thousands, millions? Or what if we want to analyze the behaviour of $\mathbf{x}_n$ as $n \rightarrow \infty$?

¹The difference equations are discrete time analogues of the *differential equation* $\mathbf{x}'(t) = A\mathbf{x}(t)$. To solve the differential equation, one needs to compute $e^{tA} := \sum_{k=0}^{\infty} t^k A^k / k!$, and spectral theory also helps in doing this.

And here the idea of *eigenvalues* and *eigenvectors* comes handy: suppose that $A\mathbf{x}_0 = \lambda\mathbf{x}_0$, where λ is some scalar. Then $A^2\mathbf{x}_0 = \lambda^2\mathbf{x}_0$, $A^3\mathbf{x}_0 = \lambda^3\mathbf{x}_0$, $\dots$, $A^n\mathbf{x}_0 = \lambda^n\mathbf{x}_0$, so the behaviour of the solution is very well understood

In this section we will consider only operators in finitedimensional spaces. The spectral theory in infinitely many dimensions is significantly more complicated, and most of result presented here fails in infinitedimensional setting.

1. Main definitions

1.1. Eigenvalues, eigenvectors, spectrum. A scalar λ is called an *eigenvalue* of an operator $A : V \rightarrow V$ if there exists a *non-zero* vector $\mathbf{v} \in V$ such that

$$A\mathbf{v} = \lambda\mathbf{v}.$$

The vector $\mathbf{v}$ is called the *eigenvector* of A (corresponding to the eigenvalue λ).

If we know that λ is an eigenvalue, the eigenvectors are easy to find: one just has to solve the equation $A\mathbf{x} = \lambda\mathbf{x}$, or, equivalently

$$(A - \lambda I)\mathbf{x} = \mathbf{0}.$$

So, finding *all* eigenvectors, corresponding to an eigenvalue λ is simply finding the nullspace of $A - \lambda I$. The nullspace $\text{Ker}(A - \lambda I)$, i.e. the set of all eigenvectors and $\mathbf{0}$ vector, is called the *eigenspace*.

The set of all eigenvalues of an operator A is called *spectrum* of A , and is usually denoted $\sigma(A)$.

1.2. Finding eigenvalues: characteristic polynomials. A scalar λ is an eigenvalue if and only if the nullspace $\text{Ker}(A - \lambda I)$ is non-trivial (so the equation $(A - \lambda I)\mathbf{x} = \mathbf{0}$ has a non-trivial solution).

Let A acts on $\mathbb{R}^n$ (i.e. $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$). Since the matrix of A is square, $A - \lambda I$ has a non-trivial nullspace if and only if it is not invertible. And we know that a square matrix is not invertible if and only if its determinant is 0. Therefore

$$\boxed{\lambda \in \sigma(A), \text{ i.e. } \lambda \text{ is an eigenvalue of } A \iff \det(A - \lambda I) = 0}$$

If A is an $n \times n$ matrix, the determinant $\det(A - \lambda I)$ is a polynomial of degree n of the variable λ . This polynomial is called the *characteristic polynomial* of A . So, to find all eigenvalues of A one just needs to compute the characteristic polynomial and find all its roots.

This method of finding spectrum of an operator is not very practical in higher dimensions. Finding roots of a polynomial of high degree can be a very difficult problem, and it is impossible to solve the equation of

degree higher than 4 in radicals. So, in higher dimensions different numerical methods of finding eigenvalues and eigenvectors are used.

1.3. Characteristic polynomial of an operator. So we know how to find spectrum of a matrix. But how to find eigenvalues of an operator acting in an abstract vector space? The recipe is simple:

Take an arbitrary basis, and compute eigenvalues of the matrix of the operator in this basis

But how do we know that the result does not depend on a choice of the basis?

There can be several possible explanations. One is based on the notion of *similar matrices*. Let us recall that square matrices A and B are called similar if there exist an invertible matrix S such that

$$A = SBS^{-1}.$$

Note, that determinants of similar matrices coincide. Indeed

$$\det A = \det(SBS^{-1}) = \det S \det B \det S^{-1} = \det B$$

because $\det S^{-1} = 1/\det S$. Note that if $A = SBS^{-1}$ then

$$A - \lambda I = SBS^{-1} - \lambda SIS^{-1} = S(BS^{-1} - \lambda IS^{-1}) = S(B - \lambda I)S^{-1},$$

so the matrices $A - \lambda I$ and $B - \lambda I$ are similar. Therefore

$$\det(A - \lambda I) = \det(B - \lambda I),$$

i.e.

characteristic polynomials of similar matrices coincide.

If $T : V \rightarrow V$ is a linear transformation, and $\mathcal{A}$ and $\mathcal{B}$ are two bases in V , then

$$[T]_{\mathcal{A}\mathcal{A}} = [I]_{\mathcal{A}\mathcal{B}}[T]_{\mathcal{B}\mathcal{B}}[I]_{\mathcal{B}\mathcal{A}}$$

and since $[I]_{\mathcal{B}\mathcal{A}} = ([I]_{\mathcal{A}\mathcal{B}})^{-1}$ the matrices $[T]_{\mathcal{A}\mathcal{A}}$ and $[T]_{\mathcal{B}\mathcal{B}}$ are similar.

In other words, matrices of a linear transformation in different bases are similar.

Therefore, we can define the characteristic polynomial of an operator as the characteristic polynomial of its matrix in some basis. As we just have discussed above, the result does not depend on the choice of the basis, so characteristic polynomial of an operator is well defined.

2. Diagonalization.

Suppose an operator (matrix) A has a basis $\mathcal{B} = \mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ of eigenvectors, and let $\lambda_1, \lambda_2, \dots, \lambda_n$ be the corresponding eigenvalues. Then the matrix of A in this basis is the diagonal matrix with $\lambda_1, \lambda_2, \dots, \lambda_n$ on the diagonal

$$[A]_{\mathcal{B}\mathcal{B}} = \text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_n\} = \begin{pmatrix} \lambda_1 & & 0 \\ & \lambda_2 & \\ 0 & & \ddots \\ & & & \lambda_n \end{pmatrix}.$$

Therefore, it is easy to find an N th power of the operator A . Namely, its matrix in the basis $\mathcal{B}$ is

$$[A^N]_{\mathcal{B}\mathcal{B}} = \text{diag}\{\lambda_1^N, \lambda_2^N, \dots, \lambda_n^N\} = \begin{pmatrix} \lambda_1^N & & 0 \\ & \lambda_2^N & \\ 0 & & \ddots \\ & & & \lambda_n^N \end{pmatrix}.$$

Moreover, functions of the operator are also very easy to compute: for example the operator (matrix) exponent e^{tA} is defined as $e^{tA} = I + tA + \frac{t^2 A^2}{2!} + \frac{t^3 A^3}{3!} + \dots = \sum_{k=0}^{\infty} \frac{t^k A^k}{k!}$, and its matrix in the basis $\mathcal{B}$ is

$$[e^{tA}]_{\mathcal{B}\mathcal{B}} = \text{diag}\{e^{\lambda_1 t}, e^{\lambda_2 t}, \dots, e^{\lambda_n t}\} = \begin{pmatrix} e^{\lambda_1 t} & & 0 \\ & e^{\lambda_2 t} & \\ 0 & & \ddots \\ & & & e^{\lambda_n t} \end{pmatrix}.$$

To find the matrices in the standard basis $\mathcal{S}$, we need to recall that the change of coordinate matrix $[I]_{\mathcal{S}\mathcal{B}}$ is the matrix with columns $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$. Let us call this matrix S , then

$$A = [A]_{\mathcal{S}\mathcal{S}} = S \begin{pmatrix} \lambda_1 & & 0 \\ & \lambda_2 & \\ 0 & & \ddots \\ & & & \lambda_n \end{pmatrix} S^{-1} = SDS^{-1},$$

where we use D for the diagonal matrix in the middle.

Similarly

$$A^N = SD^N S^{-1} = S \begin{pmatrix} \lambda_1^N & & & 0 \\ & \lambda_2^N & & \\ & & \ddots & \\ 0 & & & \lambda_n^N \end{pmatrix} S^{-1}$$

and similarly for e^{tA} .

Another way of thinking about powers (or other functions) of diagonalizable operators is to see that if operator A can be represented as $A = SDS^{-1}$, then

$$A^N = \underbrace{(SDS^{-1})(SDS^{-1}) \dots (SDS^{-1})}_{N \text{ times}} = SD^N S^{-1}$$

and it is easy to compute the N th power of a diagonal matrix.

The following theorem is almost trivial.

Theorem 2.1. *A matrix A admits a representation $A = SDS^{-1}$, where D is a diagonal matrix if and only if there exists a basis of eigenvectors of A .*

Proof. We already discussed above that if there is a basis of eigenvectors, then the matrix admits the representation $A = SDS^{-1}$, where $S = [I]_{\mathcal{S}\mathcal{B}}$ is the change of coordinate matrix from coordinates in the basis $\mathcal{B}$ to the standard coordinates.

On the other hand if the matrix admits the representation $A = SDS^{-1}$ with a diagonal matrix D , then columns of S are eigenvectors of A (column number k corresponds to the k th diagonal entry of D). Since S is invertible, its columns form a basis. $\square$

Theorem 2.2. *Let $\lambda_1, \lambda_2, \dots, \lambda_r$ be distinct eigenvalues of A , and let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r$ be the corresponding eigenvectors. Then vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r$ are linearly independent.*

Proof. We will use induction on r . The case $r = 1$ is trivial, because by the definition an eigenvector is non-zero. And a system consisting of one non-zero vector is linearly independent.

Suppose that the statement of the theorem is true for $r - 1$. Suppose there exists a non-trivial linear combination

$$(2.1) \quad c_1 \mathbf{v}_1 + c_2 \mathbf{v}_2 + \dots + c_r \mathbf{v}_r = \sum_{k=1}^r c_k \mathbf{v}_k = \mathbf{0}.$$

Applying $A - \lambda_r I$ to (2.1) and using the fact that $(A - \lambda_r I)\mathbf{v}_r = \mathbf{0}$ we get

$$\sum_{k=1}^{r-1} c_k(\lambda_k - \lambda_r)\mathbf{v}_k = \mathbf{0}.$$

By the induction hypothesis vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{r-1}$ are linearly independent, so $c_k(\lambda_k - \lambda_r) = 0$ for $k = 1, 2, \dots, r-1$. Since $\lambda_k \neq \lambda_r$ we can conclude that $c_k = 0$ for $k < r$. Then it follows from (2.1) that $c_r = 0$, i.e. we have the trivial linear combination. $\square$

Corollary 2.3. *If an operator $A : V \rightarrow V$ has exactly $n = \dim V$ distinct eigenvalues, then it is diagonalizable.*

Proof. For each eigenvalue λ_k let $\mathbf{v}_k$ be a corresponding eigenvector (just pick one eigenvector for each eigenvalue). By Theorem 2.2 the system $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ is linearly independent, and since it consists of exactly $n = \dim V$ vectors it is a basis. $\square$

2.1. Bases of subspaces (AKA direct sums of subspaces). Let $V_1, V_2, \dots, V_p$ be subspaces of a vector space V . We say that the system of subspaces is a basis in V if any vector $\mathbf{v} \in V$ admits a unique representation as a sum

$$\mathbf{v} = \mathbf{v}_1 + \mathbf{v}_2 + \dots + \mathbf{v}_p = \sum_{k=1}^p \mathbf{v}_k, \quad \mathbf{v}_k \in V_k.$$

We also say, that a system of subspaces $V_1, V_2, \dots, V_p$ is linearly independent if the equation

$$\mathbf{v}_1 + \mathbf{v}_2 + \dots + \mathbf{v}_p = \mathbf{0}, \quad \mathbf{v}_k \in V_k$$

has only trivial solution ($\mathbf{v}_k = \mathbf{0} \forall k = 1, 2, \dots, p$).

Another way to phrase that is to say that a system of subspaces $V_1, V_2, \dots, V_p$ is linearly independent if and only if any system of non-zero vectors $\mathbf{v}_k$, where $\mathbf{v}_k \in V_k$, is linearly independent.

Remark 2.4. From the above definition one can immediately see that Theorem 2.1 states in fact that the system of eigenspaces E_k of an operator A

$$E_k := \text{Ker}(A - \lambda_k I), \quad \lambda_k \in \sigma(A),$$

is linearly independent.

There is a simple example of a basis of subspaces. Let V be a vector space with a basis $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$. Split the set of indices $1, 2, \dots, n$ into p subsets $\Lambda_1, \Lambda_2, \dots, \Lambda_p$, and define subspaces $V_k := \text{span}\{\mathbf{v}_j : j \in \Lambda_k\}$. Clearly subspaces V_k form a basis of V .

The following theorem shows that in finite-dimensional case it is essentially the only possible example of a basis of subspaces.

Theorem 2.5. *Let $V_1, V_2, \dots, V_p$ be a basis of subspaces, and let in each subspace V_k we have a basis (of vectors) $\mathcal{B}_k$ ². Then the union $\cup_k \mathcal{B}_k$ of these bases is a basis in V .*

To prove the theorem we need the following lemma

Lemma 2.6. *Let $V_1, V_2, \dots, V_p$ be a linearly independent family of subspaces, and let in each subspace V_k we have a linearly independent system $\mathcal{B}_k$ of vectors³. Then the union $\mathcal{B} := \cup_k \mathcal{B}_k$ is a linearly independent system.*

Proof. The proof of the lemma is almost trivial, if one thinks a bit about it. The main difficulty in writing the proof is a choice of a appropriate notation. Instead of using two indices (one for the number k and the other for the number of a vector in $\mathcal{B}_k$, let us use “flat” notation.

Namely, let n be the number of vectors in $\mathcal{B} := \cup_k \mathcal{B}_k$. Let us split the set of indices $1, 2, \dots, n$ into p sets $\Lambda_1, \Lambda_2, \dots, \Lambda_p$, such that the number of indices in Λ_k coincides with the number of vectors in $\mathcal{B}_k$, and index vectors in $\mathcal{B}_k$ by indices in Λ_k . In this notation the set $\mathcal{B}_k$ consists of vectors $\mathbf{b}_j : j \in \Lambda_k$.

Suppose we have a non-trivial linear combination

$$(2.2) \quad c_1 \mathbf{b}_1 + c_2 \mathbf{b}_2 + \dots + c_n \mathbf{b}_n = \sum_{j=1}^n c_j \mathbf{b}_j = \mathbf{0}.$$

Denote

$$\mathbf{v}_k = \sum_{j \in \Lambda_k} c_j \mathbf{b}_j.$$

Then (2.2) can be rewritten as

$$\mathbf{v}_1 + \mathbf{v}_2 + \dots + \mathbf{v}_p = \mathbf{0}.$$

Since $\mathbf{v}_k \in V_k$ and the system of subspaces V_k is linearly independent, $\mathbf{v}_k = \mathbf{0} \forall k$. That means that for every k

$$\sum_{j \in \Lambda_k} c_j \mathbf{b}_j = \mathbf{0},$$

and since the system of vectors $\mathbf{b}_j : j \in \Lambda_k$ (i.e. the system $\mathcal{B}_k$) are linearly independent, we have $c_j = 0$ for all $j \in \Lambda_k$. Since it is true for all Λ_k , we can conclude that $c_j = 0$ for all j . $\square$

²We do not list here vectors in $\mathcal{B}_k$, one just should keep in mind that each $\mathcal{B}_k$ consists of finitely many vectors in V_k

³again here we do not name each vector in $\mathcal{B}_k$ individually, we just keep in mind that each set $\mathcal{B}_k$ consists of finitely many vectors.

Proof of Theorem 2.5. To prove the theorem we will use the same notation as in the proof of Lemma 2.6, i.e. the system $\mathcal{B}_k$ consists of vectors $\mathbf{b}_j$, $j \in \Lambda_k$.

Lemma 2.6 asserts that the system of vectors $\mathbf{b}_j$, $j = 1, 2, \dots, n$ is linearly independent, so it only remains to show that the system is complete.

Since the system of subspaces $V_1, V_2, \dots, V_p$ is a basis, any vector $\mathbf{v} \in V$ can be represented as

$$\mathbf{v} = \mathbf{v}_1 p_1 + \mathbf{v}_2 p_2 + \dots + \mathbf{v}_p p_p = \sum_{k=1}^p \mathbf{v}_k, \quad \mathbf{v}_k \in V_k.$$

Since the vectors $\mathbf{b}_j$, $j \in \Lambda_k$ form a basis in V_k , vectors $\mathbf{v}_k$ can be represented as

$$\mathbf{v}_k = \sum_{j \in \Lambda_k} c_j \mathbf{b}_j,$$

and therefore $\mathbf{v} = \sum_{j=1}^n c_j \mathbf{b}_j$. $\square$

2.2. Criterion of diagonalizability. First of all let us mention a simple necessary condition. Since for a diagonal matrix $D = \text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_n\}$

$$\det(D - \lambda I) = (\lambda_1 - \lambda)(\lambda_2 - \lambda) \dots (\lambda_n - \lambda),$$

we can see that if an operator A is diagonalizable, its characteristic polynomial splits into the product of monomials. Note, that any polynomial can be decomposed into the product of monomials, if we allow complex coefficients (i.e. complex eigenvalues).

In what follows we always assume that the characteristic polynomial splits into the product of monomials, either by working in a complex vector space, or simply assuming that A has exactly $n = \dim V$ eigenvalues (counting multiplicity).

Theorem 2.7. *An operator $A : V \rightarrow V$ is diagonalizable if and only if for each eigenvalue λ the dimension of the eigenspace $\text{Ker}(A - \lambda I)$ (i.e. the geometric multiplicity) coincides with the algebraic multiplicity of λ .*

Proof. First of all let us note, that for a diagonal matrix algebraic and geometric multiplicities of eigenvalues coincide, and therefore the same holds for the diagonalizable operators.

Let us now prove another implication. Let $\lambda_1, \lambda_2, \dots, \lambda_p$ be eigenvalues of A , and let $E_k := \text{Ker}(A - \lambda_k I)$ be the corresponding eigenspaces. According to Remark 2.4, the subspaces E_k , $k = 1, 2, \dots, p$ are linearly independent.

Let $\mathcal{B}_k$ be a basis in E_k . By Lemma 2.6 the system $\mathcal{B} = \cup_k \mathcal{B}_k$ is a linearly independent system of vectors.

We know that each $\mathcal{B}_k$ consists of $\dim E_k (= \text{multiplicity of } \lambda_k)$ vectors. So the number of vectors in $\mathcal{B}$ equal to the sum of multiplicities of eigenvectors λ_k , which is exactly $n = \dim V$. So, we have a linearly independent system of $\dim V$ eigenvectors, which means it is a basis. $\square$

2.3. Some example.

2.3.1. *Real eigenvalues.* Consider matrix

$$A = \begin{pmatrix} 1 & 2 \\ 8 & 1 \end{pmatrix}.$$

Its characteristic polynomial equals to

$$\begin{vmatrix} 1-\lambda & 2 \\ 8 & 1-\lambda \end{vmatrix} = (1-\lambda)^2 - 16$$

and its roots (eigenvalues) are $\lambda = 5$ and $\lambda = -3$. For the eigenvalue $\lambda = 5$

$$A - 5I = \begin{pmatrix} 1-5 & 2 \\ 8 & 1-5 \end{pmatrix} = \begin{pmatrix} -4 & 2 \\ 8 & -4 \end{pmatrix}$$

A basis in its nullspace consists of one vector $(1, 2)^T$, so this is the corresponding eigenvector.

Similarly, for $\lambda = -3$

$$A - \lambda I = A + 3I = \begin{pmatrix} 4 & 2 \\ 8 & 4 \end{pmatrix}$$

and the eigenspace $\text{Ker}(A + 3I)$ is spanned by the vector $(1, -2)^T$. The matrix A can be diagonalized as

$$A = \begin{pmatrix} 1 & 2 \\ 8 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 2 & -2 \end{pmatrix} \begin{pmatrix} 5 & 0 \\ 0 & -3 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 2 & -2 \end{pmatrix}^{-1}$$

2.3.2. *Complex eigenvalues.* Consider matrix

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

Its characteristic polynomial is

$$\begin{vmatrix} 1-\lambda & 2 \\ -2 & 1-\lambda \end{vmatrix} = (1-\lambda)^2 + 2^2$$

and the eigenvalues (roots of the characteristic polynomial) are $\lambda = 1 \pm 2i$. For $\lambda = 1 + 2i$

$$A - \lambda I = \begin{pmatrix} -2i & 2 \\ -2 & -2i \end{pmatrix}$$

This matrix has rank 1, so the eigenspace $\text{Ker}(A - \lambda I)$ is spanned by one vector, for example by $(1, i)^T$.

Since the matrix A is real, we do not need to compute an eigenvector for $\lambda = 1 - 2i$: we can get it for free by taking the complex conjugate of the above eigenvector. So, for $\lambda = 1 - 2i$ a corresponding eigenvector is $(1, -i)^T$, and so matrix A can be diagonalized as

$$A = \begin{pmatrix} 1 & 1 \\ i & -i \end{pmatrix} \begin{pmatrix} 1+2i & 0 \\ 0 & 1-2i \end{pmatrix} \begin{pmatrix} 1 & 1 \\ i & -i \end{pmatrix}^{-1}$$

2.3.3. *A non-diagonalizable matrix.* Consider matrix

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

Its characteristic polynomial is

$$\begin{vmatrix} 1-\lambda & 1 \\ 0 & 1-\lambda \end{vmatrix} = (1-\lambda)^2,$$

so A has an eigenvalue 1 of multiplicity 2. But the eigenspace $\text{Ker}(A - 1I)$ is one dimensional (spanned by the vector $(1, 0)^T$). Therefore, although the characteristic polynomial of A is a product of monomials, A cannot be diagonalized.

Inner product spaces

1. Inner product in $\mathbb{R}^n$ and $\mathbb{C}^n$. Inner product spaces.

1.1. Inner product and norm in $\mathbb{R}^n$. In dimensions 2 and 3, we defined the length of a vector $\mathbf{x}$ (i.e. the distance from its endpoint to the origin) by the Pythagorean rule, for example in $\mathbb{R}^2$ the length of the vector is defined as

$$\|\mathbf{x}\| = \sqrt{x_1^2 + x_2^2 + x_3^2}.$$

It is natural to generalize this formula for all n , to define the *norm* of the vector $\mathbf{x} \in \mathbb{R}^n$ as

$$\|\mathbf{x}\| = \sqrt{x_1^2 + x_2^2 + \dots + x_n^2}$$

The word *norm* is used as a fancy replacement to the word length.

The *dot product* in $\mathbb{R}^3$ was defined as $\mathbf{x} \cdot \mathbf{y} = x_1y_1 + x_2y_2 + x_3y_3$, where $\mathbf{x} = (x_1, x_2, x_3)^T$ and $\mathbf{y} = (y_1, y_2, y_3)^T$.

Similarly, in $\mathbb{R}^n$ one can define the *inner product* $(\mathbf{x}, \mathbf{y})$ of two vectors $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$, $\mathbf{y} = (y_1, y_2, \dots, y_n)^T$ by

$$(\mathbf{x}, \mathbf{y}) := x_1y_1 + x_2y_2 + \dots + x_ny_n = \mathbf{y}^T \mathbf{x},$$

so $\|\mathbf{x}\| = \sqrt{(\mathbf{x}, \mathbf{x})}$.

Note, that $\mathbf{y}^T \mathbf{x} = \mathbf{x}^T \mathbf{y}$, and we use the notation $\mathbf{y}^T \mathbf{x}$ only to be consistent.

1.2. Inner product and norm in $\mathbb{C}^n$. Let us now define norm and inner product for $\mathbb{C}^n$. As we have seen before, the complex space $\mathbb{C}^n$ is most natural space from the point of view of spectral theory: even if one starts from a matrix with real coefficients (or operator on a real vectors space), the eigenvalues can be complex, and one needs to work in a complex space.

While the notation $\mathbf{x} \cdot \mathbf{y}$ and term “dot product” is often used for the inner product, for the reasons which will be clear later, we prefer the notation $(\mathbf{x}, \mathbf{y})$

For a complex number $z = x + iy$, we have $|z|^2 = x^2 + y^2 = z\bar{z}$. If $\mathbf{z} \in \mathbb{C}^n$ is given by

$$\mathbf{z} = \begin{pmatrix} z_1 \\ z_2 \\ \vdots \\ z_n \end{pmatrix} = \begin{pmatrix} x_1 + iy_1 \\ x_2 + iy_2 \\ \vdots \\ x_n + iy_n \end{pmatrix},$$

it is natural to define its norm $\|\mathbf{z}\|$ by

$$\|\mathbf{z}\|^2 = \sum_{k=1}^n (x_k^2 + y_k^2) = \sum_{k=1}^n |z_k|^2.$$

Let us try to define an inner product on $\mathbb{C}^n$ such that $\|\mathbf{z}\|^2 = (\mathbf{z}, \mathbf{z})$. One of the choices is to define $(\mathbf{z}, \mathbf{w})$ by

$$(\mathbf{z}, \mathbf{w}) = z_1 \bar{w}_1 + z_2 \bar{w}_2 + \dots + z_n \bar{w}_n = \sum_{k=1}^n z_k \bar{w}_k,$$

and that will be our definition of the inner product in $\mathbb{C}^n$.

To simplify the notation, let us introduce a new notion. For a matrix A let us define its *Hermitian adjoint*, or simply *adjoint* A^* by $A^* = \overline{A}^T$, meaning that we take the transpose of the matrix, and then take the complex conjugate of each entry. Note, that for a real matrix A , $A^* = A^T$.

Using the notion of A^* , one can write the inner product in $\mathbb{C}^n$ as

$$(\mathbf{z}, \mathbf{w}) = \mathbf{w}^* \mathbf{z}.$$

Remark. It is easy to see that one can define a different inner product in $\mathbb{C}^n$ such that $\|\mathbf{z}\|^2 = (\mathbf{z}, \mathbf{z})$, namely the inner product given by

$$(\mathbf{z}, \mathbf{w})_1 = \bar{z}_1 w_1 + \bar{z}_2 w_2 + \dots + \bar{z}_n w_n = \mathbf{z}^* \mathbf{w}.$$

We did not specify, what properties we want the inner product to satisfy, but $\mathbf{z}^* \mathbf{w}$ and $\mathbf{w}^* \mathbf{z}$ are the only reasonable choices giving $\|\mathbf{z}\|^2 = (\mathbf{z}, \mathbf{z})$.

Note, that the above two choices of the inner product are essentially equivalent: the only difference between them is notational, because $(\mathbf{z}, \mathbf{w})_1 = (\mathbf{w}, \mathbf{z})$.

While the second choice of the inner product looks more natural, the first one, $(\mathbf{z}, \mathbf{w}) = \mathbf{w}^* \mathbf{z}$ is more widely used, so we will use it as well.

1.3. Inner product spaces. The inner product we defined for $\mathbb{R}^n$ and $\mathbb{C}^n$ satisfies the following properties:

1. (Conjugate) symmetry: $(\mathbf{x}, \mathbf{y}) = \overline{(\mathbf{y}, \mathbf{x})}$; note, that for a real space, this property is just symmetry, $(\mathbf{x}, \mathbf{y}) = (\mathbf{y}, \mathbf{x})$;
2. Linearity: $(\alpha \mathbf{x} + \beta \mathbf{y}, \mathbf{z}) = \alpha(\mathbf{x}, \mathbf{z}) + \beta(\mathbf{y}, \mathbf{z})$ for all vector $\mathbf{x}, \mathbf{y}, \mathbf{z}$ and all scalars α, β ;
3. Non-negativity: $(\mathbf{x}, \mathbf{x}) \geq 0 \forall \mathbf{x}$;

4. Non-degeneracy: $(\mathbf{x}, \mathbf{x}) = 0$ if and only if $\mathbf{x} = \mathbf{0}$.

Let V be a (complex or real) vector space. An *inner product* on V is a function, that assign to each pair of vectors $\mathbf{x}, \mathbf{y}$ a scalar, denoted by $(\mathbf{x}, \mathbf{y})$ such that the properties 1–4 from the previous section are satisfied.

Note that for a real space V we assume that $(\mathbf{x}, \mathbf{y})$ is always real, and for a complex space the inner product $(\mathbf{x}, \mathbf{y})$ can be complex.

A space V together with an inner product on it is called an *inner product space*. Given an inner product space, one defines the norm on it by

$$\|\mathbf{x}\| = \sqrt{(\mathbf{x}, \mathbf{x})}.$$

1.3.1. Examples.

Example 1.1. Let V be $\mathbb{R}^n$ or $\mathbb{C}^n$. We already have an inner product $(\mathbf{x}, \mathbf{y}) = \mathbf{y}^* \mathbf{x} = \sum_{k=1}^n x_k \overline{y_k}$ defined above.

This inner product is called the *standard* inner product in $\mathbb{R}^n$ or $\mathbb{C}^n$

We will use symbol $\mathbb{F}$ to denote both $\mathbb{C}$ and $\mathbb{R}$. When we have some statement about space $\mathbb{F}^n$, it means the statement is true for both $\mathbb{R}^n$ and $\mathbb{C}^n$.

Example 1.2. Let V be the space $\mathbb{P}_n$ of polynomials of degree at most n . Define the inner product by

$$(f, g) = \int_{-1}^1 f(t) \overline{g(t)} dt.$$

It is easy to check, that the above properties 1–4 are satisfied.

This definition works both for complex and real cases. In the real case we only allow polynomials with real coefficients, and we do not need the complex conjugate here.

Let us recall, that for a square matrix A , its *trace* is defined as the sum of the diagonal entries,

$$\text{trace } A := \sum_{k=1}^n a_{k,k}.$$

Example 1.3. For the space $M_{m \times n}$ of $m \times n$ matrices let us define the inner product by

$$(A, B) = \text{trace}(B^* A).$$

Again, it is easy to check that the properties 1–4 are satisfied, i.e. that we indeed defined an inner product.

1.4. Properties of inner product. The statements we get in this section are true for any inner product space, not only for $\mathbb{F}^n$. To prove them we use only properties 1–4 of the inner product.

First of all let us notice, that properties 1 and 2 imply that

$$2'. (\mathbf{x}, \alpha\mathbf{y} + \beta\mathbf{z}) = \overline{\alpha}(\mathbf{x}, \mathbf{y}) + \overline{\beta}(\mathbf{x}, \mathbf{z}).$$

Indeed,

$$\begin{aligned} (\mathbf{x}, \alpha\mathbf{y} + \beta\mathbf{z}) &= \overline{(\alpha\mathbf{y} + \beta\mathbf{z}, \mathbf{x})} = \overline{\alpha(\mathbf{y}, \mathbf{x}) + \beta(\mathbf{z}, \mathbf{x})} = \\ &= \overline{\alpha(\mathbf{y}, \mathbf{x})} + \overline{\beta(\mathbf{z}, \mathbf{x})} = \overline{\alpha}(\mathbf{x}, \mathbf{y}) + \overline{\beta}(\mathbf{x}, \mathbf{z}) \end{aligned}$$

Note also that the property 2 implies that for all vectors $\mathbf{x}$

$$(\mathbf{0}, \mathbf{x}) = (\mathbf{x}, \mathbf{0}) = 0.$$

Lemma 1.4. *Let $\mathbf{x}$ be a vector in an inner product space V . Then $\mathbf{x} = \mathbf{0}$ if and only if*

$$(1.1) \quad (\mathbf{x}, \mathbf{y}) = 0 \quad \forall \mathbf{y} \in V.$$

Proof. Since $(\mathbf{0}, \mathbf{y}) = 0$ we only need to show that (1.1) implies $\mathbf{x} = \mathbf{0}$. Putting $\mathbf{y} = \mathbf{x}$ in (1.1) we get $(\mathbf{x}, \mathbf{x}) = 0$, so $\mathbf{x} = \mathbf{0}$. $\square$

Applying the above lemma to the difference $\mathbf{x} - \mathbf{y}$ we get the following

Corollary 1.5. *Let $\mathbf{x}, \mathbf{y}$ be vectors in an inner product space V . The equality $\mathbf{x} = \mathbf{y}$ holds if and only if*

$$(\mathbf{x}, \mathbf{z}) = (\mathbf{y}, \mathbf{z}) \quad \forall \mathbf{z} \in V.$$

The following corollary is very simple, but will be used a lot

Corollary 1.6. *Suppose two operators $A, B : X \rightarrow Y$ satisfy*

$$(A\mathbf{x}, \mathbf{y}) = (B\mathbf{x}, \mathbf{y}) \quad \forall \mathbf{x} \in X, \forall \mathbf{y} \in Y.$$

Then $A = B$

Proof. By the previous corollary (fix $\mathbf{x}$ and take all possible $\mathbf{y}$ s) we get $A\mathbf{x} = B\mathbf{x}$. Since this is true for all $\mathbf{x} \in X$, the transformations A and B coincide. $\square$

The following property relates the norm and the inner product.

Theorem 1.7 (Cauchy–Schwarz inequality).

$$|(\mathbf{x}, \mathbf{y})| \leq \|\mathbf{x}\| \cdot \|\mathbf{y}\|.$$

Proof. The proof we are going to present, is not the shortest one, but it gives a lot for the understanding.

Let us consider the real case first. If $\mathbf{y} = \mathbf{0}$, the statement is trivial, so we can assume that $\mathbf{y} \neq \mathbf{0}$. By the properties of an inner product, for all scalar t

$$0 \leq \|\mathbf{x} - t\mathbf{y}\|^2 = (\mathbf{x} - t\mathbf{y}, \mathbf{x} - t\mathbf{y}) = \|\mathbf{x}\|^2 - 2t(\mathbf{x}, \mathbf{y}) + t^2\|\mathbf{y}\|^2.$$

In particular, this inequality should hold for $t = \frac{(\mathbf{x}, \mathbf{y})}{\|\mathbf{y}\|^2}$ ¹, and for this point the inequality becomes

$$0 \leq \|\mathbf{x}\|^2 - 2\frac{(\mathbf{x}, \mathbf{y})^2}{\|\mathbf{y}\|^2} + \frac{(\mathbf{x}, \mathbf{y})^2}{\|\mathbf{y}\|^2} = \|\mathbf{x}\|^2 - \frac{(\mathbf{x}, \mathbf{y})^2}{\|\mathbf{y}\|^2},$$

which is exactly the inequality we need to prove.

There are several possible ways to treat the complex case. One is to replace $\mathbf{x}$ by $\alpha\mathbf{x}$, where α is a complex constant, $|\alpha| = 1$ such that $(\alpha\mathbf{x}, \mathbf{y})$ is real, and then repeat the proof for the real case.

The other possibility is again to consider

$$\begin{aligned} 0 \leq \|\mathbf{x} - t\mathbf{y}\|^2 &= (\mathbf{x} - t\mathbf{y}, \mathbf{x} - t\mathbf{y}) = (\mathbf{x}, \mathbf{x} - t\mathbf{y}) - t(\mathbf{y}, \mathbf{x} - t\mathbf{y}) \\ &= \|\mathbf{x}\|^2 - t(\mathbf{x}, \mathbf{y}) - \bar{t}(\mathbf{y}, \mathbf{x}) + |t|^2\|\mathbf{y}\|^2. \end{aligned}$$

Substituting $t = \frac{(\mathbf{x}, \mathbf{y})}{\|\mathbf{y}\|^2}$ into this inequality, we get

$$0 \leq \|\mathbf{x}\|^2 - \frac{|(\mathbf{x}, \mathbf{y})|^2}{\|\mathbf{y}\|^2}$$

which is the inequality we need.

Note, that the above paragraph is in fact a complete formal proof of the theorem. The reasoning before that was only to explain why do we need to pick this particular value of t . $\square$

The following *polarization identities*, which allow to reconstruct the inner product from the norm:

Lemma 1.8 (Polarization identities). *For $\mathbf{x}, \mathbf{y} \in V$*

$$(\mathbf{x}, \mathbf{y}) = \frac{1}{4} (\|\mathbf{x} + \mathbf{y}\|^2 - \|\mathbf{x} - \mathbf{y}\|^2)$$

if V is a real inner product space, and

$$(\mathbf{x}, \mathbf{y}) = \frac{1}{4} \sum_{\alpha=\pm 1, \pm i} \alpha \|\mathbf{x} + \alpha\mathbf{y}\|^2$$

if V is a complex space.

¹That is the point where the above quadratic polynomial has minimum: it can be computed, for example by taking the derivative in t and equating it to 0

The lemma is proved by direct computation. We leave the proof as an exercise for the reader.

2. Orthogonality

3. Least square solution. Formula for the orthogonal projection

As it was discussed before in Chapter 2, the equation

$$A\mathbf{x} = \mathbf{b}$$

has a solution if and only if $\mathbf{b} \in \text{Ran } A$. But what to do if we want to solve an equation that does not have a solution?

This seems to be a silly question, because if there is no solution, then there is no solution. But situations when we want to solve an equation that does not have a solution can appear naturally, for example, if we obtained the equation from an experiment. If we do not have any errors, the right side $\mathbf{b}$ belongs to the column space $\text{Ran } A$, and equation is consistent. But in real life it is impossible to avoid errors in measurements, so it is possible that an equation that in theory should be consistent, does not have a solution. So, what one can do in this situation?

3.1. Least square solution. The simplest idea is to write down the error

$$\|A\mathbf{x} - \mathbf{b}\|$$

and try to find $\mathbf{x}$ minimizing it. If we can find $\mathbf{x}$ such that the error is 0, the system is consistent and we have exact solution. Otherwise, we get the so-called *least square* solution.

The term *least square* arises from the fact that minimizing $\|A\mathbf{x} - \mathbf{b}\|$ is equivalent to minimizing

$$\|A\mathbf{x} - \mathbf{b}\|^2 = \sum_{k=1}^m |(A\mathbf{x})_k - b_k|^2 = \sum_{k=1}^m \left| \sum_{j=1}^n A_{k,j}x_j - b_k \right|^2$$

i.e. to minimizing the sum of squares of linear functions.

There are several ways to find the least square solution. If we are in $\mathbb{R}^n$, and everything is real, we can forget about absolute values. Then we can just take partial derivatives with respect to x_j and find the where all of them are 0, which gives us minimum.

3.1.1. Geometric approach. However, there is a simpler way of finding the minimum. Namely, if we take all possible vectors $\mathbf{x}$, then $A\mathbf{x}$ gives us all possible vectors in $\text{Ran } A$, so minimum of $\|A\mathbf{x} - \mathbf{b}\|$ is exactly the distance from $\mathbf{b}$ to $\text{Ran } A$. Therefore the value of $\|A\mathbf{x} - \mathbf{b}\|$ is minimal if and only if $A\mathbf{x} = P_{\text{Ran } A}\mathbf{b}$, where $P_{\text{Ran } A}$ stands for the orthogonal projection onto the column space $\text{Ran } A$.

So, to find the least square solution we simply need to solve the equation

$$A\mathbf{x} = P_{\text{Ran } A}\mathbf{b}.$$

If we know an orthogonal basis $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ in $\text{Ran } A$, we can find vector $P_{\text{Ran } A} \mathbf{b}$ by the formula

$$P_{\text{Ran } A} \mathbf{b} = \sum_{k=1}^n \frac{(\mathbf{b}, \mathbf{v}_k)}{\|\mathbf{v}_k\|^2} \mathbf{v}_k.$$

If we only know a basis in $\text{Ran } A$, we need to use the Gram–Schmidt orthogonalization to obtain an orthogonal basis from it.

So, theoretically, the problem is solved, but the solution is not very simple: it involves Gram–Schmidt orthogonalization, which can be computationally intensive. Fortunately, there exists a simpler solution.

3.1.2. Normal equation. Namely, $A\mathbf{x}$ is the orthogonal projection $P_{\text{Ran } A} \mathbf{b}$ if and only if $\mathbf{b} - A\mathbf{x} \perp \text{Ran } A$ ($A\mathbf{x} \in \text{Ran } A$ for all $\mathbf{x}$).

If $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n$ are columns of A , then the condition $A\mathbf{x} \perp \text{Ran } A$ can be rewritten as

$$\mathbf{b} - A\mathbf{x} \perp \mathbf{a}_k, \quad \forall k = 1, 2, \dots, n.$$

That means

$$0 = (\mathbf{b} - A\mathbf{x}, \mathbf{a}_k) = \mathbf{a}_k^* (\mathbf{b} - A\mathbf{x}) \quad \forall k = 1, 2, \dots, n.$$

Joining rows $\mathbf{a}_k^*$ together we get that these equations are equivalent to

$$A^* (\mathbf{b} - A\mathbf{x}) = \mathbf{0},$$

which in turn is equivalent to the so-called *normal equation*

$$A^* A \mathbf{x} = A^* \mathbf{b}.$$

A solution of this equation gives us the least square solution of $A\mathbf{x} = \mathbf{b}$.

Note, that the least square solution is unique if and only if $A^* A$ is invertible.

3.2. Formula for the orthogonal projection. As we already discussed above, if $\mathbf{x}$ is a solution of the *normal equation* $A^* A \mathbf{x} = A^* \mathbf{b}$ (i.e. a least square solution of $A\mathbf{x} = \mathbf{b}$), then $A\mathbf{x} = P_{\text{Ran } A} \mathbf{b}$. So, to find the orthogonal projection of $\mathbf{b}$ onto the column space $\text{Ran } A$ we need to solve the normal equation $A^* A \mathbf{x} = A^* \mathbf{b}$, and then multiply the solution by A .

If operator $A^* A$ is invertible, the solution of the normal equation $A^* A \mathbf{x} = A^* \mathbf{b}$ is given by $\mathbf{x} = (A^* A)^{-1} A^* \mathbf{b}$, so the orthogonal projection $P_{\text{Ran } A} \mathbf{b}$ can be computed as

$$P_{\text{Ran } A} \mathbf{b} = A(A^* A)^{-1} A^* \mathbf{b}.$$

Since this is true for all $\mathbf{b}$,

$$P_{\text{Ran } A} = A(A^* A)^{-1} A^*$$

is the formula for the matrix of the orthogonal projection onto $\text{Ran } A$.

The following theorem implies that if an $m \times n$ matrix A has rank n , then A^*A is invertible.

Theorem 3.1. *For an $m \times n$ matrix A*

$$\text{Ker } A = \text{Ker}(A^*A).$$

Indeed, according to this theorem, if rank A is n , then $\text{Ker } A = \{\mathbf{0}\}$. Therefore $\text{Ker}(A^*A) = \{\mathbf{0}\}$ (pivot in every column), and since A^*A is square, it is invertible.

We leave the proof of the theorem as an exercise. To prove the equality $\text{Ker } A = \text{Ker}(A^*A)$ one needs to prove two inclusions $\text{Ker}(A^*A) \subset \text{Ker } A$ and $\text{Ker } A \subset \text{Ker}(A^*A)$. One of the inclusion is trivial, for the other one use the fact that

$$\|A\mathbf{x}\|^2 = (A\mathbf{x}, A\mathbf{x}) = (A^*A\mathbf{x}, \mathbf{x}).$$

3.3. An example: line fitting. Let us introduce few examples where the least square solution appears naturally. Suppose that we know that two quantities x and y are related by the law $y = a + bx$. The coefficients a and b are unknown, and we would like to find them from an experiment data.

Suppose we run the experiment n times, and we got n pairs (x_k, y_k) , $k = 1, 2, \dots, n$. Ideally, all the points (x_k, y_k) should be on a straight line, but because of errors in measurements, it usually does not happen: the point are usually close to some line, but not exactly on it. And that is where the least square solution helps!

Ideally, the coefficients a and b should satisfy the equations

$$a + bx_k = y_k, \quad k = 1, 2, \dots, n$$

(note that here, x_k and y_k are some fixed numbers, and the unknowns are a and b). If it is possible to find such a and b we are lucky: if not, the standard thing to do, is to minimize the total quadratic error

$$\sum_{k=1}^n |a + bx_k - y_k|^2.$$

But minimizing this error is exactly finding the least square solution of the system

$$\begin{pmatrix} 1 & x_1 \\ 1 & x_2 \\ \vdots & \vdots \\ 1 & x_n \end{pmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix}$$

(let us remind, that x_k y_k are some given numbers, and the unknowns are a and b).

3.3.1. *An example.* Suppose our data (x_k, y_k) consist of pairs

$$(-2, 4), (-1, 2), (0, 1), (2, 1), (3, 1).$$

Then we need to find the least square solution of

$$\begin{pmatrix} 1 & -2 \\ 1 & -1 \\ 1 & 0 \\ 1 & 2 \\ 1 & 3 \end{pmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{pmatrix} 4 \\ 2 \\ 1 \\ 1 \\ 1 \end{pmatrix}$$

Then

$$A^*A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ -2 & -1 & 0 & 2 & 3 \end{pmatrix} \begin{pmatrix} 1 & -2 \\ 1 & -1 \\ 1 & 0 \\ 1 & 2 \\ 1 & 3 \end{pmatrix} = \begin{pmatrix} 5 & 2 \\ 2 & 18 \end{pmatrix}$$

and

$$A^*\mathbf{b} = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ -2 & -1 & 0 & 2 & 3 \end{pmatrix} \begin{pmatrix} 4 \\ 2 \\ 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 9 \\ -5 \end{pmatrix}$$

so the normal equation $A^*A\mathbf{x} = A^*\mathbf{b}$ is rewritten as

$$\begin{pmatrix} 5 & 2 \\ 2 & 18 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} 9 \\ -5 \end{pmatrix}.$$

The solution of this equation is

$$a = 2, b = -1/2,$$

so the best fitting straight line is

$$y = 2 - 1/2x.$$

3.4. Other examples: curves and planes. The least square method is not limited to the line fitting: it can be applied to more general curves, as well as to surfaces in higher dimensions.

The only constrain here is that the parameters we want to find be involved linearly. The general algorithm is as follows:

1. Find the equations that your data should satisfy if there is exact fit;
2. Write these equations as a linear system, where unknowns are the parameters you want to find. Note, that the system need not to be consistent (and usually is not);
3. Find the least square solution of the system.

3.4.1. *An example: curve fitting.* For example, suppose we know that the relation between x and y is given by the quadratic law $y = a + bx + cx^2$, so we want to fit a parabola $y = a + bx + cx^2$ to the data. Then our unknowns a, b, c should satisfy the equations

$$a + bx_k + cx_k^2 = y_k, \quad k = 1, 2, \dots, n$$

or, in matrix form

$$\begin{pmatrix} 1 & x_1 & x_1^2 \\ 1 & x_2 & x_2^2 \\ \vdots & \vdots & \vdots \\ 1 & x_n & x_n^2 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix}$$

For example, for the data from the previous example we need to find the least square solution of

$$\begin{pmatrix} 1 & -2 & 4 \\ 1 & -1 & 1 \\ 1 & 0 & 0 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} 4 \\ 2 \\ 1 \\ 1 \\ 1 \end{pmatrix}.$$

Then

$$A^*A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ -2 & -1 & 0 & 2 & 3 \\ 4 & 1 & 0 & 4 & 9 \end{pmatrix} \begin{pmatrix} 1 & -2 & 4 \\ 1 & -1 & 1 \\ 1 & 0 & 0 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \end{pmatrix} = \begin{pmatrix} 5 & 2 & 18 \\ 2 & 18 & 26 \\ 18 & 26 & 114 \end{pmatrix}$$

and

$$A^*\mathbf{b} = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ -2 & -1 & 0 & 2 & 3 \\ 4 & 1 & 0 & 4 & 9 \end{pmatrix} \begin{pmatrix} 4 \\ 2 \\ 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 9 \\ -5 \\ 31 \end{pmatrix}.$$

Therefore the normal equation $A^*A\mathbf{x} = A^*\mathbf{b}$ is

$$\begin{pmatrix} 5 & 2 & 18 \\ 2 & 18 & 26 \\ 18 & 26 & 114 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} 9 \\ -5 \\ 31 \end{pmatrix}$$

which has the unique solution

$$a = 86/77, \quad b = -62/77, \quad c = 43/154.$$

Therefore,

$$y = 86/77 - 62x/77 + 43x^2/154$$

is the best fitting parabola.

3.4.2. *Plane fitting.* As another example, let us fit a plane $z = a + bx + cy$ to the data

$$(x_k, y_k, z_k) \in \mathbb{R}^3, \quad k = 1, 2, \dots, n.$$

The equations we should have in the case of exact fit are

$$a + bx_k + cy_k = z_k, \quad k = 1, 2, \dots, n,$$

or, in the matrix form

$$\begin{pmatrix} 1 & x_1 & y_1 \\ 1 & x_2 & y_2 \\ \vdots & \vdots & \vdots \\ 1 & x_n & y_n \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} z_1 \\ z_2 \\ \vdots \\ z_n \end{pmatrix}.$$

So, to find the best fitting plane, we need to find the best square solution of this system (unknowns are a, b, c).

4. Fundamental subspaces revisited.

4.1. Adjoint matrices and adjoint operators. Let us recall that for an $m \times n$ matrix A its *Hermitian adjoint* (or simply *adjoint*) A^* is defined by $A^* := \overline{A^T}$. In other words, matrix A^* is obtained from the transposed matrix A^T by taking complex conjugate of each entry.

The following identity is the main property of adjoint matrix:

$$(A\mathbf{x}, \mathbf{y}) = (\mathbf{x}, A^*\mathbf{y}) \quad \forall \mathbf{x}, \mathbf{y} \in V.$$

Before proving this identity, let us introduce some useful formulas. Let us recall that for transposed matrices we have the identity $(AB)^T = B^T A^T$. Since for complex numbers z and w we have $\overline{\overline{z}w} = \overline{z}\overline{w}$, the identity

$$(AB)^* = B^* A^*$$

holds for the adjoint.

Also, since $(A^T)^T = A$ and $\overline{\overline{z}} = z$,

$$(A^*)^* = A.$$

Now, we are ready to prove the main identity:

$$(A\mathbf{x}, \mathbf{y}) = \mathbf{y}^* A\mathbf{x} = (A^*\mathbf{y})^* \mathbf{x} = (\mathbf{x}, A^*\mathbf{y});$$

the first and the last equalities here follow from the definition of inner product in $\mathbb{F}^n$, and the middle one follows from the fact that

$$(A^*\mathbf{x})^* = \mathbf{x}^* (A^*)^* = \mathbf{x}^* A.$$

4.1.1. *Uniqueness of the adjoint.* The above main identity $(A\mathbf{x}, \mathbf{y}) = (\mathbf{x}, A^*\mathbf{y})$ is often used as the definition of the adjoint operator. Let us first notice that the adjoint operator is unique: if a matrix B satisfies

$$(A\mathbf{x}, \mathbf{y}) = (\mathbf{x}, B\mathbf{y}) \quad \forall \mathbf{x}, \mathbf{y},$$

then $B = A^*$. Indeed, by the definition of A^*

$$(\mathbf{x}, A^*\mathbf{y}) = (\mathbf{x}, B\mathbf{y}) \quad \forall \mathbf{x}$$

and therefore by Corollary 1.5 $A^*\mathbf{y} = B\mathbf{y}$. Since it is true for all $\mathbf{y}$, the linear transformations, and therefore the matrices A^* and B coincide.

4.1.2. *Adjoint transformation in abstract setting.* The above main identity $(A\mathbf{x}, \mathbf{y}) = (\mathbf{x}, A^*\mathbf{y})$ can be used to define the adjoint operator in abstract setting, where $A : V \rightarrow W$ is an operator acting from one inner product space to another. Namely, we define $A^* : W \rightarrow V$ be the operator satisfying

$$(A\mathbf{x}, \mathbf{y}) = (\mathbf{x}, A^*\mathbf{y}) \quad \forall \mathbf{x} \in V, \forall \mathbf{y} \in W.$$

Why does such an operator exists? We can simply construct it: consider orthonormal bases $\mathcal{A} = \mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ in V and $\mathcal{B} = \mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_m$ in W . If $[A]_{\mathcal{B}\mathcal{A}}$ is the matrix of A with respect to these bases, we define the operator A^* by defining its matrix $[A^*]_{\mathcal{A}\mathcal{B}}$ as

$$[A^*]_{\mathcal{A}\mathcal{B}} = ([A]_{\mathcal{B}\mathcal{A}})^*.$$

We leave the proof that this indeed gives the adjoint operator as an exercise for the reader.

Note, that the reasoning in the above Sect. 4.1.1 implies that the adjoint operator is unique.

4.1.3. *Useful formulas.* Below we present the properties of the adjoint operators (matrices) we will use a lot. We leave the proofs as an exercise for the reader.

1. $(A + B)^* = A^* + B^*$;
2. $(\alpha A)^* = \bar{\alpha} A^*$;
3. $(AB)^* = B^* A^*$;
4. $(A^*)^* = A$.

4.2. Relation between fundamental subspaces.

Theorem 4.1. *Let $A : V \rightarrow W$ be an operator acting from one inner product space to another. Then*

1. $\text{Ker } A^* = (\text{Ran } A)^\perp$;
2. $\text{Ker } A = (\text{Ran } A^*)^\perp$;
3. $\text{Ran } A = (\text{Ker } A^*)^\perp$;

4. $\text{Ran } A^* = (\text{Ker } A)^\perp$;
5. $(\mathbf{y}, A\mathbf{x}) = (A^*\mathbf{y}, \mathbf{x})$.

Proof. First of all, let us notice, that since for a subspace E we have $(E^\perp)^\perp = E$, the statements 1 and 3 are equivalent. Similarly, for the same reason, the statements 2 and 4 are equivalent as well. And finally, the statement 2 is exactly the statement 1 applied to the operator A^* (here we use the fact that $(A^*)^* = A$).

So, to prove the theorem we only need to prove the statement 1.

We will present 2 proofs of this statement: a “matrix” proof, and an “invariant”, or “coordinate-free” one.

In the “matrix” proof, we assume that A is an $m \times n$ matrix, i.e. that $A : \mathbb{F}^n \rightarrow \mathbb{F}^m$. The general case can be always reduced to this one by picking orthonormal bases in V and W , and considering the matrix of A in this bases.

Let $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n$ be the columns of A . Note, that $\mathbf{x} \in (\text{Ran } A)^\perp$ if and only if $\mathbf{x} \perp \mathbf{a}_k$ (i.e. $(\mathbf{x}, \mathbf{a}_k) = 0$) $\forall k = 1, 2, \dots, n$.

By the definition of the inner product in $\mathbb{F}^n$, that means

$$0 = (\mathbf{x}, \mathbf{a}_k) = \mathbf{a}_k^* \mathbf{x} \quad \forall k = 1, 2, \dots, n.$$

Since $\mathbf{a}_k^*$ is the row number k of A^* , the above n equalities are equivalent to the equation

$$A^* \mathbf{x} = \mathbf{0}.$$

So, we proved that $\mathbf{x} \in (\text{Ran } A)^\perp$ if and only if $A^* \mathbf{x} = \mathbf{0}$, and that is exactly the statement 1.

Now, let us present the “coordinate-free” proof. The inclusion $\mathbf{x} \in (\text{Ran } A)^\perp$ means that $\mathbf{x}$ is orthogonal to all vectors of the form $A\mathbf{y}$, i.e. that

$$(\mathbf{x}, A\mathbf{y}) = 0 \quad \forall \mathbf{y}.$$

Since $(\mathbf{x}, A\mathbf{y}) = (A^* \mathbf{x}, \mathbf{y})$, the last identity is equivalent to

$$(A^* \mathbf{x}, \mathbf{y}) = 0 \quad \forall \mathbf{y},$$

and by Lemma 1.4 this happens if and only if $A^* \mathbf{x} = \mathbf{0}$. So we proved that $\mathbf{x} \in (\text{Ran } A)^\perp$ if and only if $A^* \mathbf{x} = \mathbf{0}$, which is exactly the statement 1 of the theorem. $\square$

The above theorem makes the structure of the operator A and the geometry of fundamental subspaces much more transparent. It follows from this theorem that the operator A can be represented as a composition of orthogonal projection onto $\text{Ran } A^*$ and an isomorphism from $\text{Ran } A^*$ to $\text{Ran } A$.

5. Isometries and unitary operators. Unitary and orthogonal matrices.

5.1. Main definitions.

Definition. An operator $U : X \rightarrow Y$ is called an *isometry*, if it preserves the norm,

$$\|U\mathbf{x}\| = \|\mathbf{x}\| \quad \forall \mathbf{x} \in X.$$

The following theorem shows that an isometry preserves the inner product

Theorem 5.1. *An operator $U : X \rightarrow Y$ is an isometry if and only if it preserves the inner product, i.e. if and only if*

$$(\mathbf{x}, \mathbf{y}) = (U\mathbf{x}, U\mathbf{y}) \quad \forall \mathbf{x}, \mathbf{y} \in X.$$

Proof. The proof uses polarization identities (Lemma 1.8). For example, if X is a complex space

$$\begin{aligned} (U\mathbf{x}, U\mathbf{y}) &= \frac{1}{4} \sum_{\alpha=\pm 1, \pm i} \alpha \|U\mathbf{x} + \alpha U\mathbf{y}\|^2 \\ &= \frac{1}{4} \sum_{\alpha=\pm 1, \pm i} \alpha \|U(\mathbf{x} + \alpha\mathbf{y})\|^2 \\ &= \frac{1}{4} \sum_{\alpha=\pm 1, \pm i} \alpha \|\mathbf{x} + \alpha\mathbf{y}\|^2 = (\mathbf{x}, \mathbf{y}). \end{aligned}$$

Similarly, for a real space X

$$\begin{aligned} (U\mathbf{x}, U\mathbf{y}) &= \frac{1}{4} (\|U\mathbf{x} + U\mathbf{y}\|^2 - \|U\mathbf{x} - U\mathbf{y}\|^2) \\ &= \frac{1}{4} (\|U(\mathbf{x} + \mathbf{y})\|^2 - \|U(\mathbf{x} - \mathbf{y})\|^2) \\ &= \frac{1}{4} (\|\mathbf{x} + \mathbf{y}\|^2 - \|\mathbf{x} - \mathbf{y}\|^2) = (\mathbf{x}, \mathbf{y}). \end{aligned}$$

□

Lemma 5.2. *An operator $U : X \rightarrow Y$ is an isometry if and only if $U^*U = I$.*

Proof. By the definitions of the isometry and of the adjoint operator

$$(\mathbf{x}, \mathbf{y}) = (U\mathbf{x}, U\mathbf{y}) = (U^*U\mathbf{x}, \mathbf{y}) \quad \forall \mathbf{x}, \mathbf{y} \in X.$$

Therefore, if $U^*U = I$, we have $(\mathbf{x}, \mathbf{y}) = (U\mathbf{x}, U\mathbf{y})$ and therefore U is an isometry.

On the other hand, if U is an isometry, then for all $\mathbf{x} \in X$

$$(U^*U\mathbf{x}, \mathbf{y}) = (\mathbf{x}, \mathbf{y}) \quad \forall \mathbf{y} \in X,$$

and therefore by Corollary 1.5 $U^*U\mathbf{x} = \mathbf{x}$. Since it is true for all $\mathbf{x} \in X$, we have $U^*U = I$. $\square$

The above lemma implies that an isometry is always left invertible (U^* being a left inverse).

Definition. An isometry $U : X \rightarrow Y$ is called a *unitary operator* if it is invertible.

Proposition 5.3. *An isometry $U : X \rightarrow Y$ is a unitary operator if and only if $\dim X = \dim Y$.*

Proof. Since U is an isometry, it is left invertible, and since $\dim X = \dim Y$, it is invertible (a left invertible square matrix is invertible).

On the other hand, if $U : X \rightarrow Y$ is invertible, $\dim X = \dim Y$ (only square matrices are invertible, isomorphic spaces have equal dimensions). $\square$

A square matrix U is called *unitary* if $U^*U = I$, i.e. a unitary matrix is a matrix of a unitary operator acting in $\mathbb{F}^n$.

A unitary matrix with real entries is called an *orthogonal matrix*. In other words, an orthogonal matrix is a matrix of a unitary operator acting in the real space $\mathbb{R}^n$.

Few properties of unitary operators:

1. For a unitary transformation U , $U^{-1} = U^*$;
2. If U is unitary, $U^* = U^{-1}$ is also unitary;
3. If U is a isometry, and $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ is an orthonormal basis, then $U\mathbf{v}_1, U\mathbf{v}_2, \dots, U\mathbf{v}_n$ is an orthonormal system. Moreover, if U is unitary, $U\mathbf{v}_1, U\mathbf{v}_2, \dots, U\mathbf{v}_n$ is an orthonormal basis.
4. A product of unitary operators is a unitary operator as well.

5.2. Examples. First of all, let us notice, that

a matrix U is an isometry if and only if its columns form an orthonormal system.

This statement can be checked directly by computing the product U^*U .

It is easy to check that the columns of the rotation matrix

$$\begin{pmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{pmatrix}$$

are orthogonal to each other, and that each column has norm 1. Therefore, the rotation matrix is an isometry, and since it is square, it is unitary. Since all entries of the rotation matrix are real, it is an orthogonal matrix.

The next example is more abstract. Let X and Y be inner product spaces, $\dim X = \dim Y = n$, and let $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$ and $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n$ be orthonormal bases in X and Y respectively. Define an operator $U : X \rightarrow Y$ by

$$U\mathbf{x}_k = \mathbf{y}_k, \quad k = 1, 2, \dots, n.$$

Since for a vector $\mathbf{x} = c_1\mathbf{x}_1 + c_2\mathbf{x}_2 + \dots + c_n\mathbf{x}_n$

$$\|\mathbf{x}\|^2 = |c_1|^2 + |c_2|^2 + \dots + |c_n|^2$$

and

$$\|U\mathbf{x}\|^2 = \|U(\sum_{k=1}^n c_k\mathbf{x}_k)\|^2 = \|\sum_{k=1}^n c_k\mathbf{y}_k\|^2 = \sum_{k=1}^n |c_k|^2,$$

one can conclude that $\|U\mathbf{x}\| = \|\mathbf{x}\|$ for all $\mathbf{x} \in X$, so U is a unitary operator.

5.3. Properties of unitary operators.

Proposition 5.4. *Let U be a unitary matrix. Then*

1. $|\det U| = 1$. In particular, for an orthogonal matrix $\det U = \pm 1$;
2. If λ is an eigenvalue of U , then $|\lambda| = 1$

Remark. Note, that for an orthogonal matrix, an eigenvalue (unlike the determinant) does not have to be real. Our old friend, the rotation matrix gives an example.

Proof of Proposition 5.4. Let $\det U = z$. Then

$$|z|^2 = \bar{z}z = \det(U^*U) = \det I = 1,$$

so $|\det U| = |z| = 1$. Statement 1 is proved.

To prove statement 2 let us notice that if $U\mathbf{x} = \lambda\mathbf{x}$ then

$$\|U\mathbf{x}\| = \|\lambda\mathbf{x}\| = |\lambda| \cdot \|\mathbf{x}\|,$$

so $|\lambda| = 1$. □

5.4. Unitary equivalent operators.

Definition. Operators (matrices) A and B are called *unitarily equivalent* if there exists a unitary operator U such that $A = UBU^*$.

Since for a unitary U we have $U^{-1} = U^*$, any two unitary equivalent matrices are similar as well.

Converse is not true, it is easy to construct a pair of similar matrices, which are not unitarily equivalent.

The following proposition gives a way to construct a counterexample

Proposition 5.5. *A matrix A is unitarily equivalent to a diagonal one if and only if it has an orthogonal (orthonormal) basis of eigenvectors.*

Proof. Let $A = UBU^*$ and let $A\mathbf{x} = \lambda\mathbf{x}$. Then $BU\mathbf{x} = UAU^*U\mathbf{x} = UA\mathbf{x} = U(\lambda\mathbf{x}) = \lambda U\mathbf{x}$, i.e. $U\mathbf{x}$ is an eigenvector of B .

So, let A be unitarily equivalent to a diagonal matrix D , i.e. let $D = UAU^*$. The vectors $\mathbf{e}_k$ of the standard basis are eigenvectors of D , so the vectors $U\mathbf{e}_k$ are eigenvectors of A . Since U is unitary, the system $U\mathbf{e}_1, U\mathbf{e}_2, \dots, U\mathbf{e}_n$ is an orthonormal basis.

Let now A has an orthogonal basis $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$ of eigenvectors. Dividing each vector $\mathbf{u}_k$ by its norm if necessary, we can always assume that the system $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$ is an *orthonormal* basis. Let D be the matrix of A in the basis $\mathcal{B} = \mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$. Clearly, D is a diagonal matrix.

Denote by U the matrix with columns $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$. Since the columns form an orthonormal basis, U is unitary. The standard change of coordinate formula implies

$$A = [A]_{SS} = [I]_{SB}[A]_{BB}[I]_{BS} = UDU^{-1}$$

and since U is unitary, $A = UDU^*$. □

Structure of operators in inner product spaces.

In this chapter we again assuming that all spaces are finite-dimensional

1. Schur (upper triangular) representation of an operator.

Theorem 1.1. *Let $A : X \rightarrow X$ be an operator acting in a complex inner product space. There exists an orthonormal basis $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$ in X such that the matrix of A in this basis is upper triangular.*

In other words, any $n \times n$ matrix A can be represented as $T = UTU^$, where U is a unitary, and T is an upper triangular matrix.*

Proof. We prove the theorem using the induction in $\dim X$. If $\dim X = 1$ the theorem is trivial, since any 1×1 matrix is upper triangular.

Suppose we proved that the theorem is true if $\dim X = n - 1$, and we want to prove it for $\dim X = n$.

Let λ_1 be an eigenvalue of A , and let $\mathbf{u}_1$, $\|\mathbf{u}_1\| = 1$ be a corresponding eigenvector, $A\mathbf{u}_1 = \lambda_1\mathbf{u}_1$. Denote $E = \mathbf{u}_1^\perp$, and let $\mathbf{v}_2, \dots, \mathbf{v}_n$ be some orthonormal basis in E (clearly, $\dim E = \dim X - 1 = n - 1$), so $\mathbf{u}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$

is an orthonormal basis in X . In this basis the matrix of A has the form

$$(1.1) \quad \left(\begin{array}{c|c} \lambda_1 & * \\ \hline 0 & \\ \vdots & \\ 0 & A_1 \end{array} \right);$$

here all entries below λ_1 are zeroes, and $*$ means that we do not care what entries are in the first row right of λ_1 .

We do care enough about the lower left $(n-1) \times (n-1)$ block, to give it name: we denote it as A_1 .

Note, that A_1 defines a linear transformation in E , and since $\dim E = n-1$, the induction hypothesis implies that there exists an orthonormal basis (let us denote it as $\mathbf{u}_2, \dots, \mathbf{u}_n$) in which the matrix of A_1 is upper triangular.

So, matrix of A in the orthonormal basis $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$ has the form (1.1), where matrix A_1 is upper triangular. Therefore, the matrix of A in this basis is upper triangular as well. $\square$

Remark. Note, that the subspace $E = \mathbf{u}_1^\perp$ introduced in the proof is not invariant under A , i.e. the inclusion $AE \subset E$ does not necessarily holds. That means that A_1 is not a part of A , it is some operator constructed from A .

Note also, that $AE \subset E$ if and only if all entries denoted by $*$ (i.e. all entries in the first row, except λ_1) are zero.

Remark. Note, that even if we start from a real matrix A , the matrices U and T can have complex entries. The rotation matrix

$$\begin{pmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{pmatrix}, \quad \alpha \neq k\pi, k \in \mathbb{Z}$$

is not unitarily equivalent (not even similar) to a real upper triangular matrix. Indeed, eigenvalues of this matrix are complex, and the eigenvalues of an upper triangular matrix are its diagonal entries.

The following theorem is a real-valued version of Theorem 1.1

Theorem 1.2. *Let $A : X \rightarrow X$ be an operator acting in a real inner product space. Suppose that all eigenvalues of A are real. Then there exists an orthonormal basis $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$ in X such that the matrix of A in this basis is upper triangular.*

In other words, any real $n \times n$ matrix A can be represented as $T = UTU^ = UTU^T$, where U is an orthogonal, and T is a real upper triangular matrices.*

Proof. To prove the theorem we just need to analyze the proof of Theorem 1.1. Let us assume (we can always do that without loss of generality, that the operator (matrix) A acts in $\mathbb{R}^n$).

Suppose, the theorem is true for $(n-1) \times (n-1)$ matrices. As in the proof of Theorem 1.1 let λ_1 be a real eigenvalue of A , $\mathbf{u}_1 \in \mathbb{R}^n$, $\|\mathbf{u}_1\| = 1$ be a corresponding eigenvector, and let $\mathbf{v}_2, \dots, \mathbf{v}_n$ be on orthonormal system (in $\mathbb{R}^n$) such that $\mathbf{u}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ is an orthonormal basis in $\mathbb{R}^n$.

The matrix of A in this basis has form (1.1), where A_1 is some real matrix.

If we can prove that matrix A_1 has only real eigenvalues, then we are done. Indeed, then by induction hypothesis there exists an orthonormal basis $\mathbf{u}_2, \dots, \mathbf{u}_n$ in $E = \mathbf{u}_1^\perp$ such that the matrix of A_1 in this basis is upper triangular, so the matrix of A in the basis $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$ is also upper triangular.

To show that A_1 has only real eigenvalues, let us notice that

$$\det(A - \lambda I) = (\lambda_1 - \lambda) \det(A_1 - \lambda)$$

(take the cofactor expansion in the first row, for example), and so any eigenvalue of A_1 is also an eigenvalue of A . But A has only real eigenvalues! $\square$

2. Spectral theorem for self-adjoint and normal operators.

In this section we deal with matrices (operators) which are unitarily equivalent to diagonal matrices.

Let us recall that an operator is called self-adjoint if $A = A^*$.

Theorem 2.1. *Let $A = A^*$ be a self-adjoint operator in an inner product space X (the space can be complex or real). Then all eigenvalues of A are real, and there exists an orthonormal basis of eigenvectors.*

This theorem can be restated in matrix form as follows

Theorem 2.2. *Let $A = A^*$ be a self-adjoint (and therefore square) matrix. Then A can be represented as*

$$A = UDU^*,$$

where U is a unitary matrix and D is a diagonal matrix with real entries.

Moreover, if the matrix A is real, matrix U can be chosen to be orthogonal.

Proof. To prove Theorems 2.1 and 2.2 let us first apply Theorem 1.1 (Theorem 1.2 if X is a real space) to find an orthonormal basis in X such that the matrix of A in this basis is upper triangular. And now let us ask ourselves the question: what upper triangular matrices are self-adjoint?

The answer is immediate: an upper triangular matrix is self-adjoint if and only if it is a diagonal matrix with real entries. Theorem 2.1 (and so Theorem 2.2) is proved. $\square$

Let us give an independent proof to the fact that eigenvalues of a self-adjoint operators are real. Let $A = A^*$ and $A\mathbf{x} = \lambda\mathbf{x}$, $\mathbf{x} \neq \mathbf{0}$. Then

$$(A\mathbf{x}, \mathbf{x}) = (\lambda\mathbf{x}, \mathbf{x}) = \lambda(\mathbf{x}, \mathbf{x}) = \lambda\|\mathbf{x}\|^2.$$

On the other hand,

$$(A\mathbf{x}, \mathbf{x}) = (\mathbf{x}, A^*\mathbf{x}) = (\mathbf{x}, A\mathbf{x}) = (\mathbf{x}, \lambda\mathbf{x}) = \bar{\lambda}(\mathbf{x}, \mathbf{x}) = \bar{\lambda}\|\mathbf{x}\|^2,$$

so $\lambda\|\mathbf{x}\|^2 = \bar{\lambda}\|\mathbf{x}\|^2$. Since $\|\mathbf{x}\| \neq 0$ ($\mathbf{x} \neq \mathbf{0}$), we can conclude $\lambda = \bar{\lambda}$, so λ is real.

It also follows from Theorem 2.1 that eigenspaces of a self-adjoint operator are orthogonal. Let us give an alternative proof of this result.

Proposition 2.3. *Let $A = A^*$ be a self-adjoint operator, and let $\mathbf{u}, \mathbf{v}$ be its eigenvectors, $A\mathbf{u} = \lambda\mathbf{u}$, $A\mathbf{v} = \mu\mathbf{v}$. Then, if $\lambda \neq \mu$, the eigenvectors $\mathbf{u}$ and $\mathbf{v}$ are orthogonal.*

Proof. This proposition follows from the spectral theorem (Theorem 1.1), but here we are giving a direct proof. Namely,

$$(A\mathbf{u}, \mathbf{v}) = (\lambda\mathbf{u}, \mathbf{v}) = \lambda(\mathbf{u}, \mathbf{v}).$$

On the other hand

$$(A\mathbf{u}, \mathbf{v}) = (\mathbf{u}, A^*\mathbf{v}) = (\mathbf{u}, A\mathbf{v}) = (\mathbf{u}, \mu\mathbf{v}) = \bar{\mu}(\mathbf{u}, \mathbf{v}) = \mu(\mathbf{u}, \mathbf{v})$$

(the last equality holds because eigenvalues of a self-adjoint operator are real), so $\lambda(\mathbf{u}, \mathbf{v}) = \mu(\mathbf{u}, \mathbf{v})$. If $\lambda \neq \mu$ it is possible only if $(\mathbf{u}, \mathbf{v}) = 0$. $\square$

Now let us try to find what matrices are unitarily equivalent to a diagonal one. It is easy to check that for a diagonal matrix D

$$D^*D = DD^*.$$

Therefore $A^*A = AA^*$ if the matrix of A in some orthonormal basis is diagonal.

Definition. An operator (matrix) N is called *normal* if $N^*N = NN^*$.

Clearly, any self-adjoint operator ($A^*A = AA^*$) is normal. Also, any unitary operator $U : X \rightarrow X$ is normal since $U^*U = UU^* = I$.

Note, that a normal operator is an operator acting in one space, not from one space to another. So, if U is a unitary operator acting from one space to another, we cannot say that U is unitary.

Theorem 2.4. *Any normal operator N in a complex vector space has an orthonormal basis of eigenvectors.*

*In other words, any matrix N satisfying $N^*N = NN^*$ can be represented as*

$$N = UDU^*,$$

where U is a unitary matrix, and D is a diagonal one.

Remark. Note, that in the above theorem even if N is a real matrix, we did not claim that matrices U and D are real. Moreover, it can be easily shown, that if D is real, N must be self-adjoint.

Proof of Theorem 2.4. To prove Theorem 2.4 we apply Theorem 1.1 to get an orthonormal basis, such that the matrix of N in this basis is upper triangular. To complete the proof of the theorem we only need to show that an upper triangular normal matrix must be diagonal.

We will prove this using induction in the dimension of matrix. The case of 1×1 matrix is trivial, since any 1×1 matrix is diagonal.

Suppose we have proved that any $(n-1) \times (n-1)$ upper triangular normal matrix is diagonal, and we want to prove it for $n \times n$ matrices. Let N be $n \times n$ upper triangular normal matrix. We can write it as

$$N = \left(\begin{array}{c|ccc} a_{1,1} & a_{1,2} & \dots & a_{1,n} \\ 0 & & & \\ \vdots & & N_1 & \\ 0 & & & \end{array} \right)$$

where N_1 is an upper triangular $(n-1) \times (n-1)$ matrix.

Let us compare upper left entries (first row first column) of N^*N and NN^* . Direct computation shows that that

$$(N^*N)_{1,1} = \overline{a_{1,1}}a_{1,1} = |a_{1,1}|^2$$

and

$$(NN^*)_{1,1} = |a_{1,1}|^2 + |a_{1,2}|^2 + \dots + |a_{1,n}|^2.$$

So, $(N^*N)_{1,1} = (NN^*)_{1,1}$ if and only if $a_{1,2} = \dots = a_{1,n} = 0$. Therefore, the matrix N has the form

$$N = \left(\begin{array}{c|ccc} a_{1,1} & 0 & \dots & 0 \\ 0 & & & \\ \vdots & & N_1 & \\ 0 & & & \end{array} \right)$$

It follows from the above representation that

$$N^*N = \left(\begin{array}{c|ccc} |a_{1,1}|^2 & 0 & \dots & 0 \\ \hline 0 & & & \\ \vdots & & N_1^*N_1 & \\ 0 & & & \end{array} \right), \quad NN^* = \left(\begin{array}{c|ccc} |a_{1,1}|^2 & 0 & \dots & 0 \\ \hline 0 & & & \\ \vdots & & N_1N_1^* & \\ 0 & & & \end{array} \right)$$

so $N_1^*N_1 = N_1N_1^*$. That means the matrix N_1 is also normal, and by the induction hypothesis it is diagonal. So the matrix N is also diagonal. $\square$

The following proposition gives a very useful characterization of normal operators.

Proposition 2.5. *An operator $N : X \rightarrow X$ is normal if and only if*

$$\|N\mathbf{x}\| = \|N^*\mathbf{x}\| \quad \forall \mathbf{x} \in X.$$

Proof. Let N be normal, $N^*N = NN^*$. Then

$$\begin{aligned} \|N\mathbf{x}\|^2 &= (N\mathbf{x}, N\mathbf{x}) = (N^*N\mathbf{x}, \mathbf{x}) = (NN^*\mathbf{x}, \mathbf{x}) = (N^*\mathbf{x}, N^*\mathbf{x}) = \|N^*\mathbf{x}\|^2 \\ \text{so } \|N\mathbf{x}\| &= \|N^*\mathbf{x}\|. \end{aligned}$$

Let now

$$\|N\mathbf{x}\| = \|N^*\mathbf{x}\| \quad \forall \mathbf{x} \in X.$$

Polarization identities (Lemma 1.8 in Chapter 5) imply that for all $\mathbf{x}, \mathbf{y} \in X$

$$\begin{aligned} (N^*N\mathbf{x}, \mathbf{y}) &= (N\mathbf{x}, N\mathbf{y}) = \sum_{\alpha=\pm 1, \pm i} \alpha \|N\mathbf{x} + \alpha N\mathbf{y}\|^2 \\ &= \sum_{\alpha=\pm 1, \pm i} \alpha \|N(\mathbf{x} + \alpha\mathbf{y})\|^2 \\ &= \sum_{\alpha=\pm 1, \pm i} \alpha \|N^*(\mathbf{x} + \alpha\mathbf{y})\|^2 \\ &= \sum_{\alpha=\pm 1, \pm i} \alpha \|N^*\mathbf{x} + \alpha N^*\mathbf{y}\|^2 \\ &= (N^*\mathbf{x}, N^*\mathbf{y}) = (NN^*\mathbf{x}, \mathbf{y}) \end{aligned}$$

and therefore (see Corollary 1.6) $N^*N = NN^*$. $\square$

3. Polar and singular value decompositions.

3.1. Positive definite operators. Square roots.

Definition. A self adjoint operator $A : X \rightarrow X$ is called *positive definite* if

$$(A\mathbf{x}, \mathbf{x}) > 0 \quad \forall \mathbf{x} \neq \mathbf{0},$$

and it is called *positive semidefinite* if

$$(A\mathbf{x}, \mathbf{x}) \geq 0 \quad \forall \mathbf{x} \in X.$$

We will use the notation $A > 0$ for positive definite operators, and $A \geq 0$ for positive semi-definite.

The following theorem describes positive definite and semidefinite operators.

Theorem 3.1. *Let $A = A^*$. Then*

1. $A > 0$ if and only if all eigenvalues of A are positive.
2. $A \geq 0$ if and only if all eigenvalues of A are non-negative.

Proof. Pick an orthonormal basis such that matrix of A in this basis is diagonal (see Theorem 2.1). To finish the proof it remains to notice that a diagonal matrix is positive definite (positive semidefinite) if and only if all its diagonal entries are positive (non-negative). $\square$

Corollary 3.2. *Let $A = A^* \geq 0$ be a positive semidefinite operator. There exists a unique positive semidefinite operator B such that $B^2 = A$*

Such B is called (positive) square root of A and is denoted as $\sqrt{A}$ or $A^{1/2}$.

Proof. Let us prove that $\sqrt{A}$ exists. Let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ be an orthonormal basis of eigenvectors of A , and let $\lambda_1, \lambda_2, \dots, \lambda_n$ be the corresponding eigenvalues. Note, that since $A \geq 0$, all $\lambda_k \geq 0$.

In the basis $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ the matrix of A is a diagonal matrix $\text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_n\}$ with entries $\lambda_1, \lambda_2, \dots, \lambda_n$ on the diagonal. Define the matrix of B in the same basis as $\text{diag}\{\sqrt{\lambda_1}, \sqrt{\lambda_2}, \dots, \sqrt{\lambda_n}\}$.

Clearly, $B = B^* \geq 0$ and $B^2 = A$.

To prove that such B is unique, let us suppose that there exists an operator $C = C^* \geq 0$ such that $C^2 = A$. Let $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$ be an orthonormal basis of eigenvectors of C , and let $\mu_1, \mu_2, \dots, \mu_n$ be the corresponding eigenvalues (note that $\mu_k \geq 0 \forall k$). The matrix of C in the basis $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$ is a diagonal one $\text{diag}\{\mu_1, \mu_2, \dots, \mu_n\}$, and therefore the matrix of $A = C^2$ in the same basis is $\text{diag}\{\mu_1^2, \mu_2^2, \dots, \mu_n^2\}$. This implies that any eigenvalue λ of A is of form μ_k^2 , and, moreover, if $A\mathbf{x} = \lambda\mathbf{x}$, then $C\mathbf{x} = \sqrt{\lambda}\mathbf{x}$.

Therefore in the basis $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ above, the matrix of C has the diagonal form $\text{diag}\{\sqrt{\lambda_1}, \sqrt{\lambda_2}, \dots, \sqrt{\lambda_n}\}$, i.e. $B = C$. $\square$

3.2. Modulus of an operator. Singular values. Consider an operator $A : X \rightarrow Y$. Its *Hermitian square* A^*A is a positive semidefinite operator acting in X . Indeed,

$$(A^*A)^* = A^*A^{**} = A^*A$$

and

$$(A^*Ax, x) = (Ax, Ax) = \|Ax\|^2 \geq 0 \quad \forall x \in X.$$

Therefore, there exists a (unique) positive-semidefinite square root $R = \sqrt{A^*A}$. This operator R is called the *modulus* of the operator A , and is often denoted as $|A|$.

The modulus of A shows how “big” the operator A is:

Proposition 3.3. *For a linear operator $A : X \rightarrow Y$*

$$\||A|x\| = \|Ax\| \quad \forall x \in X.$$

Proof. For any $x \in X$

$$\begin{aligned} \||A|x\|^2 &= (|A|x, |A|x) = (|A|^*|A|x, x) = (|A|^2x, x) \\ &= (A^*Ax, x) = (Ax, Ax) = \|Ax\|^2 \end{aligned}$$

□

Corollary 3.4.

$$\text{Ker } A = \text{Ker } |A| = (\text{Ran } |A|)^\perp.$$

Proof. The first equality follows immediately from Proposition 3.3, the second one follows from the identity $\text{Ker } T = (\text{Ran } T^*)^\perp$ ($|A|$ is self-adjoint). □

Theorem 3.5 (Polar decomposition of an operator). *Let $A : X \rightarrow X$ be an operator (square matrix). Then A can be represented as*

$$A = U|A|,$$

where U is a unitary operator.

Remark. The unitary operator U is generally not unique. As one will see from the proof of the theorem, U is unique if and only if A is invertible.

Remark. The polar decomposition $A = U|A|$ also holds for operators $A : X \rightarrow Y$ acting from one space to another. But in this case we can only guarantee that U is an *isometry* from $\text{Ran } |A| = (\text{Ker } A)^\perp$ to Y .

If $\dim X \leq \dim Y$ this isometry can be extended to the isometry from the whole X to Y (if $\dim X = \dim Y$ this will be a unitary operator).

Proof of Theorem 3.5. Consider a vector $x \in \text{Ran } |A|$. Then vector x can be represented as $x = |A|v$ for some vector $v \in X$.

Define $U_0x := Av$. By Proposition 3.3

$$\|U_0x\| = \|Av\| = \||A|v\| = \|x\|$$

so it looks like U is an isometry from $\text{Ran } |A|$ to X .

But first we need to prove that U_0 is well defined. Let v_1 be another vector such that $x = |A|v_1$. But $x = |A|v = |A|v_1$ means that $v - v_1 \in$

$\text{Ker } |A| = \text{Ker } A$ (cf Corollary 3.4), so $A\mathbf{v} = A\mathbf{v}_1$, meaning that $U_0\mathbf{x}$ is well defined.

By the construction $A = U_0|A|$. We leave as an exercise for the reader to check that U_0 is a linear transformation.

To extend U_0 to a unitary operator U , let us find some unitary transformation $U_1 : \text{Ker } A \rightarrow (\text{Ran } A)^\perp = \text{Ker } A^*$. It is always possible to do, since for square matrices $\dim \text{Ker } A = \dim \text{Ker } A^*$ (the Rank Theorem).

It is easy to check that $U = U_0 + U_1$ is a unitary operator, and that $A = U|A|$. $\square$

3.3. Singular values. Singular value decomposition.

Definition. Eigenvalues of $|A|$ are called the *singular values* of A . In other words, if $\lambda_1, \lambda_2, \dots, \lambda_n$ are eigenvalues of A^*A , then $\sqrt{\lambda_1}, \sqrt{\lambda_2}, \dots, \sqrt{\lambda_n}$ are singular values of A .

Consider an operator $A : X \rightarrow Y$, and let $\sigma_1, \sigma_2, \dots, \sigma_r$ be the non-zero singular values of A , counting multiplicities. That means $\sigma_1^2, \sigma_2^2, \dots, \sigma_r^2$ are eigenvalues of A^*A , and let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r$ be an orthonormal basis of eigenvectors of A^*A , $A^*A\mathbf{v}_k = \sigma_k^2\mathbf{v}_k$.

Proposition 3.6. *The system*

$$\mathbf{w}_k := \frac{1}{\sigma_k} A\mathbf{v}_k, \quad k = 1, 2, \dots, r$$

is an orthonormal system.

Proof.

$$(A\mathbf{v}_j, A\mathbf{v}_k) = (A^*A\mathbf{v}_j, \mathbf{v}_k) = (\sigma_j^2\mathbf{v}_j, \mathbf{v}_k) = \sigma_j^2(\mathbf{v}_j, \mathbf{v}_k) = \begin{cases} 0, & j \neq k \\ \sigma_j^2, & j = k \end{cases}$$

since $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r$ is an orthonormal system. $\square$

In the notation of the above proposition, the operator A can be represented as

$$(3.1) \quad A = \sum_{k=1}^r \sigma_k \mathbf{w}_k \mathbf{v}_k^*,$$

or, equivalently

$$(3.2) \quad A\mathbf{x} = \sum_{k=1}^r \sigma_k (\mathbf{x}, \mathbf{v}_k) \mathbf{w}_k.$$

Indeed, let $\mathbf{v}_{r+1}, \dots, \mathbf{v}_n$ be an orthonormal basis in $\text{Ker } A = \text{Ker } |A| = \text{Ker } A^*A$, so $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ is an orthonormal basis in X . Then

$$\sum_{k=1}^r \sigma_k \mathbf{w}_k \mathbf{v}_k^* \mathbf{v}_j = \sigma_j \mathbf{w}_j \mathbf{v}_j^* \mathbf{v}_j = \sigma_j \mathbf{w}_j = A \mathbf{v}_j \quad \text{if } j = 1, 2, \dots, r,$$

and

$$\sum_{k=1}^r \sigma_k \mathbf{w}_k \mathbf{v}_k^* \mathbf{v}_j = \mathbf{0} = A \mathbf{v}_j \quad \text{for } j > r.$$

So the operators in the left and right sides of (3.1) coincide on the basis $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$, so they are equal.

Definition. The above decomposition (3.1) (or (3.2)) is called the *singular value decomposition* of the operator A

Remark. Singular value decomposition is not unique. Why?

Lemma 3.7. *Let A can be represented as*

$$A = \sum_{k=1}^r \sigma_k \mathbf{w}_k \mathbf{v}_k^*$$

where $\sigma_k > 0$ and $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r, \mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_r$ are some orthonormal systems.

Then this representation gives a singular value decomposition of A .

Proof. We only need to show that $\mathbf{v}_k$ are eigenvalues of A^*A , $A^*A \mathbf{v}_k = \sigma_k^2 \mathbf{v}_k$. Since $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r$ is an orthonormal system,

$$\mathbf{w}_k^* \mathbf{w}_j = (\mathbf{w}_j, \mathbf{w}_k) = \delta_{k,j} := \begin{cases} 0, & j \neq k \\ 1, & j = k, \end{cases}$$

and therefore

$$A^*A = \sum_{k=1}^r \sigma_k^2 \mathbf{v}_k \mathbf{v}_k^*.$$

Since $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r$ is an orthonormal system

$$A^*A \mathbf{v}_j = \sum_{k=1}^r \sigma_k^2 \mathbf{v}_k \mathbf{v}_k^* \mathbf{v}_j = \sigma_j^2 \mathbf{v}_j$$

thus $\mathbf{v}_k$ are eigenvectors of A^*A . □

Corollary 3.8. *Let*

$$A = \sum_{k=1}^r \sigma_k \mathbf{w}_k \mathbf{v}_k^*$$

be a singular value decomposition of A . Then

$$A^* = \sum_{k=1}^r \sigma_k \mathbf{v}_k \mathbf{w}_k^*$$

is a singular value decomposition of A^*

3.4. Matrix representation of the singular value decomposition.

The singular value decomposition can be written in a nice matrix form. It is especially easy to do if the operator A is invertible. In this case $\dim X = \dim Y = n$, and the operator A has n non-zero singular values (counting multiplicities), so the singular value decomposition has the form

$$A = \sum_{k=1}^n \sigma_k \mathbf{w}_k \mathbf{v}_k^*$$

where $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ and $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n$ are orthonormal bases in X and Y respectively. It can be rewritten as

$$A = W \Sigma V^*,$$

where $\Sigma = \text{diag}\{\sigma_1, \sigma_2, \dots, \sigma_n\}$ and V and W are unitary matrices with columns $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ and $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n$ respectively.

Such representation can be written even if A is not invertible. Let us first consider the case $\dim X = \dim Y = n$, and let $\sigma_1, \sigma_2, \dots, \sigma_r$, $r < n$ be non-zero singular values of A . Let

$$A = \sum_{k=1}^r \sigma_k \mathbf{w}_k \mathbf{v}_k^*$$

be a singular value decomposition of A . To represent A as $W \Sigma V$ let us complete the systems $\{\mathbf{v}_k\}_{k=1}^r$, $\{\mathbf{w}_k\}_{k=1}^r$ to orthonormal bases. Namely, let $\mathbf{v}_{r+1}, \dots, \mathbf{v}_n$ and $\mathbf{w}_{r+1}, \dots, \mathbf{w}_n$ be an orthonormal bases in $\text{Ker } A = \text{Ker } |A|$ and $(\text{Ran } A)^\perp$ respectively. Then $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ and $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n$ are orthonormal bases in X and Y respectively and A can be represented as

$$A = W \Sigma V^*,$$

where Σ is $n \times n$ diagonal matrix $\text{diag}\{\sigma_1, \dots, \sigma_r, 0, \dots, 0\}$, and V , W are $n \times n$ unitary matrices with columns $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ and $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n$ respectively.

3.4.1. From singular value decomposition to the polar decomposition. Note, that if we know the singular value decomposition $A = W \Sigma V^*$ of a square matrix A , we can write a polar decomposition of A :

$$A = W \Sigma V^* = (W V^*)(V \Sigma V^*) = U |A|$$

so $|A| = V \Sigma V^*$ and $U = W V^*$.

3.4.2. *General matrix form of the singular value decomposition.* In the general case when $\dim X = n$, $\dim Y = m$ (i.e. A is an $m \times n$ matrix), the above representation $A = W\Sigma V^*$ is also possible. Namely, if

$$A = \sum_{k=1}^r \sigma_k \mathbf{w}_k \mathbf{v}_k^*$$

is a singular value decomposition of A , we need to complete the systems $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r$ and $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_r$ to orthonormal bases in X and Y respectively. Then A can be represented as

$$A = W\Sigma V^*,$$

where $V \in M_{n \times n}$ and $W \in M_{m \times m}$ are unitary matrices with columns $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ and $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_m$ respectively, and Σ is a “diagonal” $m \times n$ matrix

$$\Sigma_{j,k} = \begin{cases} \sigma_k & j = k \leq r : \\ 0 & \text{otherwise.} \end{cases}$$

In other words, to get the matrix Σ one has to take the diagonal matrix $\text{diag}\{\sigma_1, \sigma_2, \dots, \sigma_r\}$ and make it to an $m \times n$ matrix by adding extra zeroes “southeast”.

3.5. Examples.

3.6. Concluding remarks.

4. Structure of orthogonal matrices