

AALTO UNIVERSITY

SPRING 2018

Linear Algebra

ANTTI HANNUKAINEN, NUUTTI HYVÖNEN, VANNI
NOFERINI

May 17, 2021

Contents

Preface	5
Chapter 1. Linear systems	7
1. Solvability of linear systems	7
2. Norms and inner products	12
3. Stability of the solution	21
4. Solving $Ax = b$ on a computer	23
Chapter 2. Eigenvalue problems	25
1. Basic eigenvalue theory	25
2. Hermitian matrices	30
3. Similarity transformations	33
4. Matrix exponential	35
Chapter 3. Least squares problems	41
1. Motivation, definition and equivalence to normal equation	41
2. Projection matrices	43
3. Geometric interpretation of a least squares solution	47
4. QR decomposition	48
5. Singular value decomposition	51

Preface

These lecture notes cover basics of linear algebra from the standpoint of matrix computations: general vector spaces and linear maps are not in the focus of attention, but the primary objects of interest are matrices and (column) vectors that can directly be given as inputs to a computer. The students are expected to have basic knowledge about matrices; for example, the material presented during any of the courses MS-A00XX Matrix Algebra is sufficient. Many of the required fundamental tools are also reviewed in these notes.

The three Chapters [1](#), [2](#) and [3](#) correspond, respectively, to the three main themes of the lectures:

- (1) Solvability and stability of linear systems $A\mathbf{x} = \mathbf{b}$.
- (2) Eigenvalue problems $A\mathbf{v} = \lambda\mathbf{v}$ and their fundamental applications.
- (3) Least squares problems, their geometric interpretation and related matrix decompositions.

For simplicity and notational convenience, Chapters [1](#) and [3](#) consider real vectors and matrices, even though the extension to the complex case would be straightforward. However, Chapter [2](#) touches upon the complex extension that cannot be avoided in the treatment of eigenvalue problems.

Matrices and subspaces are denoted by capital letters (e.g., A or E), vectors by bolded lower case letters (e.g., \mathbf{x} or \mathbf{b}) and scalars by standard lowercase letters (e.g., α or a_{ij}). In particular, the components of a vector $\mathbf{x} \in \mathbb{R}^n$ are x_j , $j = 1, \dots, n$. However, the zero scalar and the zero vector of \mathbb{R}^n are both denoted simply as 0. The Cartesian basis vectors for \mathbb{R}^n are $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ and the identity matrix is denoted as

$$I := [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n] \in \mathbb{R}^{n \times n}.$$

We will frequently use such a notation to write matrices in terms of their column vectors. The positive integers n and m are reserved for spatial dimensions and $i, j, k, l \in \mathbb{N}_0$ are used as generic indices.

CHAPTER 1

Linear systems

Let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. In this chapter, we focus on the following simple problem: Find $\mathbf{x} \in \mathbb{R}^n$ such that

$$(1) \quad A\mathbf{x} = \mathbf{b}.$$

To be more precise, we tackle the following questions:

- (1) Is the problem (1) solvable? Is the solution unique?
- (2) How much does the solution \mathbf{x} change if A or \mathbf{b} is perturbed slightly?
- (3) How accurately can (1) be solved using a computer?

Before starting our analysis by considering the unique solvability of (1), let us make a couple of fundamental definitions that will be indispensable in what follows.

DEFINITION 0.1. *The linear span of the vectors $\mathbf{q}_1, \dots, \mathbf{q}_k \in \mathbb{R}^n$ is the set*

$$\text{span}(\mathbf{q}_1, \dots, \mathbf{q}_k) := \left\{ \mathbf{z} \in \mathbb{R}^n \mid \mathbf{z} = \sum_{j=1}^k \alpha_j \mathbf{q}_j \text{ for some } \boldsymbol{\alpha} \in \mathbb{R}^k \right\}.$$

Moreover, $\mathbf{q}_1, \dots, \mathbf{q}_k \in \mathbb{R}^n$ are called linearly independent if

$$\sum_{j=1}^k \alpha_j \mathbf{q}_j = \mathbf{0} \in \mathbb{R}^n$$

is equivalent to $\boldsymbol{\alpha} = \mathbf{0} \in \mathbb{R}^k$.

If one stacks the vectors $\mathbf{q}_1, \dots, \mathbf{q}_k \in \mathbb{R}^n$ as the columns of a matrix (as we will do frequently),

$$(2) \quad Q = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_k] \in \mathbb{R}^{n \times k},$$

then obviously

$$(3) \quad \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_k) = \left\{ \mathbf{z} \in \mathbb{R}^n \mid \mathbf{z} = Q\boldsymbol{\alpha} \text{ for some } \boldsymbol{\alpha} \in \mathbb{R}^k \right\}.$$

Moreover, $\mathbf{q}_1, \dots, \mathbf{q}_k$ are linearly independent if and only if

$$(4) \quad Q\boldsymbol{\alpha} = \mathbf{0} \iff \boldsymbol{\alpha} = \mathbf{0},$$

i.e., $Q\boldsymbol{\alpha} \in \mathbb{R}^n$ vanishes if and only if $\boldsymbol{\alpha} \in \mathbb{R}^k$ does.

1. Solvability of linear systems

The main aim of this section is to study the existence and uniqueness of a solution to (1). These properties can be characterized in terms of the *nullspace* and the *range* of the matrix A . In the following, we will denote the columns of A by $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n \in \mathbb{R}^m$ (cf. (2)).

DEFINITION 1.1. *Let $A \in \mathbb{R}^{m \times n}$. The range of A is the set*

$$(5) \quad R(A) = \left\{ \mathbf{y} \in \mathbb{R}^m \mid \mathbf{y} = A\mathbf{x} \text{ for some } \mathbf{x} \in \mathbb{R}^n \right\} = \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_n) \subset \mathbb{R}^m$$

and the nullspace of A is the set

$$N(A) = \left\{ \mathbf{x} \in \mathbb{R}^n \mid A\mathbf{x} = \mathbf{0} \right\} \subset \mathbb{R}^n.$$

The sets $R(A)$ and $N(A)$ are closed under addition and scalar multiplication. Indeed, if $\mathbf{x}, \mathbf{y} \in N(A)$, then

$$A(\alpha\mathbf{x} + \beta\mathbf{y}) = \alpha A\mathbf{x} + \beta A\mathbf{y} = 0,$$

meaning that also $\alpha\mathbf{x} + \beta\mathbf{y} \in N(A)$ for any $\alpha, \beta \in \mathbb{R}$. Similarly, if $\mathbf{x}, \mathbf{y} \in R(A)$, i.e. $\mathbf{x} = A\mathbf{z}$ and $\mathbf{y} = A\mathbf{w}$ for some $\mathbf{z}, \mathbf{w} \in \mathbb{R}^n$, then

$$A(\alpha\mathbf{z} + \beta\mathbf{w}) = \alpha A\mathbf{z} + \beta A\mathbf{w} = \alpha\mathbf{x} + \beta\mathbf{y},$$

i.e., $\alpha\mathbf{x} + \beta\mathbf{y} \in R(A)$ for any $\alpha, \beta \in \mathbb{R}$. This means that the subsets $R(A) \subset \mathbb{R}^m$ and $N(A) \subset \mathbb{R}^n$ are actually *subspaces*.

DEFINITION 1.2. Let $E \subset \mathbb{R}^n$ be nonempty and such that for any $\mathbf{x}, \mathbf{y} \in E$ and $\alpha \in \mathbb{R}$

$$\mathbf{x} + \mathbf{y} \in E \quad \text{and} \quad \alpha\mathbf{x} \in E.$$

Then E is called a subspace of \mathbb{R}^n . (In particular, \mathbb{R}^n itself as well the trivial subspace $\{0\}$ are subspaces of \mathbb{R}^n .)

The range of the matrix A induces a condition for the existence of a solution to the equation (1). By definition, for any $\mathbf{b} \in R(A)$, there exists $\mathbf{z} \in \mathbb{R}^n$ such that

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

On the other hand, if (1) has a solution, then obviously $\mathbf{b} \in R(A)$. To sum up, there exists a solution to the linear system (1) if and only if $\mathbf{b} \in R(A) = \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_n)$.

The uniqueness of a solution is studied via the homogeneous problem: Find $\mathbf{y} \in \mathbb{R}^n$ such that

$$A\mathbf{y} = 0 \quad \iff \quad \mathbf{y} \in N(A).$$

If the homogeneous equation only has the trivial solution $\mathbf{y} = 0$, a solution to the linear system (1) is unique (if one exists).

LEMMA 1.1. Assume that (1) has a solution. The solution is unique if and only if $N(A) = \{0\}$ is the trivial subspace.

PROOF. Let $\mathbf{x} \in \mathbb{R}^n$ be a solution of (1) and assume first that $N(A)$ is nontrivial, i.e., there exists $0 \neq \mathbf{z} \in N(A)$. Then,

$$A(\mathbf{x} + \mathbf{z}) = A\mathbf{x} + A\mathbf{z} = A\mathbf{x} = \mathbf{b},$$

meaning that we have constructed a second, distinct solution $\mathbf{x} + \mathbf{z} \neq \mathbf{x}$. This proves the ‘only if’ part of the claim.

In order to prove the ‘if part’, assume that $N(A) = \{0\}$ and let $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^n$ both be solutions to (1), i.e.,

$$A\mathbf{x}_1 = \mathbf{b} \quad \text{and} \quad A\mathbf{x}_2 = \mathbf{b}.$$

Subtracting the two equations gives

$$A(\mathbf{x}_1 - \mathbf{x}_2) = 0.$$

Hence, $\mathbf{x}_1 - \mathbf{x}_2 \in N(A)$, which means by assumption that $\mathbf{x}_1 - \mathbf{x}_2 = 0$, and thus all solutions to (1) must be the same. \square

To summarize the above observations, the problem (1) has a unique solution if and only if

$$(6) \quad \mathbf{b} \in R(A) \quad \text{and} \quad N(A) = \{0\}.$$

Moreover, the proof of Lemma 1.1 indicates that if $\mathbf{x} \in \mathbb{R}^n$ is a solution of (1), then $\mathbf{x} + \mathbf{z}$ is also a solution for any $\mathbf{z} \in N(A)$, which characterizes the possible nonuniqueness related to (1).

It turns out that the combined condition (6) can be true for all $\mathbf{b} \in \mathbb{R}^m$ only if $n = m$, i.e., only if A is a square matrix. This claim can be proved by relating the *dimensions* of $N(A)$ and $R(A)$ to each other and to n . To this end, we need the concept of *basis*.

DEFINITION 1.3. Let $E \subset \mathbb{R}^n$ be a subspace. A basis of E is a set of linearly independent vectors $\{\mathbf{q}_1, \dots, \mathbf{q}_k\} \subset \mathbb{R}^n$ such that

$$E = \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_k).$$

With the notation of the above definition, for each $\mathbf{x} \in E$ there exists a *coordinate vector* $\boldsymbol{\alpha} \in \mathbb{R}^k$ such that

$$(7) \quad \mathbf{x} = \sum_{i=1}^k \mathbf{q}_i \alpha_i \iff \mathbf{x} = Q\boldsymbol{\alpha},$$

where $Q = [\mathbf{q}_1, \dots, \mathbf{q}_k] \in \mathbb{R}^{n \times k}$ is a matrix with the basis vectors as its columns. As the basis vectors are by definition linearly independent, such a coordinate presentation is unique. Indeed, if (7) holds for two coordinate vectors $\boldsymbol{\alpha}, \tilde{\boldsymbol{\alpha}} \in \mathbb{R}^k$, then subtracting the corresponding representations yields

$$Q(\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}) = \mathbf{0} \implies \boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}} = \mathbf{0}$$

by virtue of (4).

Take note that there exists infinitely many different bases for any nontrivial subspace $\{0\} \neq E \subset \mathbb{R}^n$. However, it can be proven that each basis of a subspace E has the same number of basis vectors (the proof is omitted).

DEFINITION 1.4. *Let $E \subset \mathbb{R}^n$ be a subspace and $\{\mathbf{q}_1, \dots, \mathbf{q}_k\} \subset \mathbb{R}^n$ its basis. Then the dimension of E is defined to be*

$$\dim(E) = k = \#\{\mathbf{q}_1, \dots, \mathbf{q}_k\}.$$

For a matrix $A \in \mathbb{R}^{m \times n}$, $\dim(R(A))$ is called the rank of A . (In particular, the dimension of \mathbb{R}^n itself is n since, e.g., the Cartesian basis vectors $\mathbf{e}_1, \dots, \mathbf{e}_n$ form its basis. No other subspace of \mathbb{R}^n has the maximal dimension n .)

In the following example two different bases are introduced for the same subspace.

EXAMPLE 1.1. *Let*

$$A = \begin{bmatrix} 1 & 1 \\ 0 & -1 \\ 2 & 2 \end{bmatrix}.$$

The column vectors of A are clearly linearly independent and so they form a basis for $R(A) \subset \mathbb{R}^3$; see (5). Let

$$\mathbf{w}_1 = \begin{bmatrix} 1 \\ 0 \\ 2 \end{bmatrix}, \quad \mathbf{w}_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$$

be another pair of vectors in \mathbb{R}^3 and denote $W = [\mathbf{w}_1, \mathbf{w}_2] \in \mathbb{R}^{3 \times 2}$.

To study if \mathbf{w}_1 and \mathbf{w}_2 also form a basis for $R(A)$, one has to check two conditions:

- Are the vectors \mathbf{w}_1 and \mathbf{w}_2 linearly independent? (This is obviously true in the considered setting.)
- Can any $\mathbf{y} \in R(A)$ be represented as a linear combination of \mathbf{w}_1 and \mathbf{w}_2 ? (This is not quite obvious.)

The latter question can be answered by first introducing a generic element of $R(A)$, i.e. $\mathbf{y} = A\mathbf{z}$ for some $\mathbf{z} \in \mathbb{R}^2$, and then trying to express it as a linear combination of \mathbf{w}_1 and \mathbf{w}_2 . In other words, one must check if for any $\mathbf{z} \in \mathbb{R}^2$, there exists $\mathbf{x} \in \mathbb{R}^2$ such that

$$W\mathbf{x} = x_1\mathbf{w}_1 + x_2\mathbf{w}_2 = A\mathbf{z},$$

that is,

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 2 & 0 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 1 & 1 \\ 0 & -1 \\ 2 & 2 \end{bmatrix} \mathbf{z}.$$

Gaussian elimination gives

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 1 & 1 \\ 0 & -1 \\ 0 & 0 \end{bmatrix} \mathbf{z},$$

and thus $x_1 = z_1 + z_2$ and $x_2 = -z_2$ provides the needed solution. In consequence, $\{\mathbf{w}_1, \mathbf{w}_2\}$ is a basis of $R(A)$. (Note that we did not need to separately check whether \mathbf{w}_1 and \mathbf{w}_2 belong to $R(A)$ because we knew to begin with that the dimension of $R(A)$ is two.)

We start the actual discussion on the relationship between the dimensions of $R(A)$ and $N(A)$ with another example in \mathbb{R}^3 .

EXAMPLE 1.2. Let

$$A = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 2 & 2 \\ -1 & 1 & 0 \end{bmatrix}.$$

The nullspace of A can be computed by solving the equation

$$(8) \quad A\mathbf{x} = 0.$$

Using Gaussian elimination, it is straightforward to deduce that the solutions to (8) are exactly the scalar multiples of

$$\mathbf{x} = \begin{bmatrix} -1 \\ -1 \\ 1 \end{bmatrix}, \quad \text{i.e.,} \quad N(A) = \{\mathbf{x} \in \mathbb{R}^3 \mid \mathbf{x} = s[-1, -1, 1]^T \text{ for some } s \in \mathbb{R}\}.$$

Hence, the dimension of $N(A)$ is one.

Recall that

$$R(A) = \{\mathbf{y} \in \mathbb{R}^3 \mid \mathbf{y} = A\mathbf{x} \text{ for some } \mathbf{x} \in \mathbb{R}^3\},$$

which is the linear span of the column vectors of A . To deduce the dimension of $R(A)$, we introduce a special basis for \mathbb{R}^3 , namely

$$\mathbf{v}_1 = \begin{bmatrix} -1 \\ -1 \\ 1 \end{bmatrix}, \quad \mathbf{v}_2 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{v}_3 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}.$$

The first vector \mathbf{v}_1 is a basis for $N(A)$, and the main idea is to choose the other two so that the three vectors are linearly independent and thus they together form a basis for \mathbb{R}^3 . Denote $V = [\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3]$. Since the columns of $V \in \mathbb{R}^{3 \times 3}$ form a basis for \mathbb{R}^3 , any vector $\mathbf{x} \in \mathbb{R}^3$ can be given as their linear combination, i.e., as $\mathbf{x} = V\boldsymbol{\alpha}$ for some $\boldsymbol{\alpha} \in \mathbb{R}^3$. Hence, the range of A can be alternatively expressed as

$$R(A) = \{\mathbf{y} \in \mathbb{R}^3 \mid \mathbf{y} = A(V\boldsymbol{\alpha}) \text{ for some } \boldsymbol{\alpha} \in \mathbb{R}^3\}.$$

By construction $\mathbf{v}_1 \in N(A)$, and so

$$A(V\boldsymbol{\alpha}) = (AV)\boldsymbol{\alpha} = [A\mathbf{v}_1, A\mathbf{v}_2, A\mathbf{v}_3]\boldsymbol{\alpha} = [0, A\mathbf{v}_2, A\mathbf{v}_3]\boldsymbol{\alpha} = \alpha_2 A\mathbf{v}_2 + \alpha_3 A\mathbf{v}_3.$$

Hence, $R(A)$ is the linear span of the vectors $A\mathbf{v}_2$ and $A\mathbf{v}_3$, which are linearly independent because they are by construction the (linearly independent) first and second columns of A . In particular, the dimension of $R(A)$ is two.

In Example 1.2,

$$\dim(N(A)) + \dim(R(A)) = 3,$$

which is the dimension of the considered space \mathbb{R}^3 . The basic idea of Example 1.2 is also valid for a general $A \in \mathbb{R}^{m \times n}$, although one needs to use more general arguments to demonstrate that the candidates for the basis vectors of $R(A)$ are linearly independent.

THEOREM 1.1. For any $A \in \mathbb{R}^{m \times n}$,

$$(9) \quad \dim(R(A)) + \dim(N(A)) = n.$$

PROOF. Let $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ be a basis for $N(A)$; in particular, it is assumed that the dimension of $N(A)$ is k . We introduce auxiliary vectors $\mathbf{w}_1, \dots, \mathbf{w}_{n-k}$ such that

$$\{\mathbf{v}_1, \dots, \mathbf{v}_k, \mathbf{w}_1, \dots, \mathbf{w}_{n-k}\}$$

is a basis for the whole \mathbb{R}^n .¹ Define $V \in \mathbb{R}^{n \times k}$ and $W \in \mathbb{R}^{n \times (n-k)}$ via

$$V = [\mathbf{v}_1, \dots, \mathbf{v}_k] \quad \text{and} \quad W = [\mathbf{w}_1, \dots, \mathbf{w}_{n-k}].$$

Because the columns of the composite matrix $[V, W] \in \mathbb{R}^{n \times n}$ form a basis for \mathbb{R}^n , any $\mathbf{x} \in \mathbb{R}^n$ can be written as $\mathbf{x} = [V, W]\boldsymbol{\alpha}$ for some $\boldsymbol{\alpha} \in \mathbb{R}^n$. In particular,

$$R(A) = \{\mathbf{y} \in \mathbb{R}^n \mid \mathbf{y} = A[V, W]\boldsymbol{\alpha} \text{ for some } \boldsymbol{\alpha} \in \mathbb{R}^n\},$$

where

$$\begin{aligned} A[V, W]\boldsymbol{\alpha} &= [A\mathbf{v}_1, \dots, A\mathbf{v}_k, A\mathbf{w}_1, \dots, A\mathbf{w}_{n-k}]\boldsymbol{\alpha} \\ &= [0, \dots, 0, A\mathbf{w}_1, \dots, A\mathbf{w}_{n-k}]\boldsymbol{\alpha} \\ &= \sum_{i=1}^{n-k} \alpha_{i+k} A\mathbf{w}_i \end{aligned}$$

is a linear combination of $n-k$ vectors. Hence, $R(A)$ can be represented as a linear span of $n-k$ vectors, and so its dimension is at most $n-k$.

Completing the proof amounts to showing that the vectors $A\mathbf{w}_1, \dots, A\mathbf{w}_{n-k}$ are linearly independent, i.e., that they form a basis for $R(A)$. Let $\boldsymbol{\alpha} \in \mathbb{R}^{n-k}$ be such that

$$\sum_{i=1}^{n-k} \alpha_i (A\mathbf{w}_i) = [A\mathbf{w}_1, \dots, A\mathbf{w}_{n-k}]\boldsymbol{\alpha} = A\mathbf{W}\boldsymbol{\alpha} = 0,$$

that is, $\mathbf{W}\boldsymbol{\alpha} \in N(A)$. Since the columns of $V \in \mathbb{R}^{n \times k}$ form, by assumption, a basis for $N(A)$, there exists (a unique) $\boldsymbol{\beta} \in \mathbb{R}^k$ such that $\mathbf{W}\boldsymbol{\alpha} = V\boldsymbol{\beta}$, or in an equivalent form,

$$[V, W] \begin{bmatrix} \boldsymbol{\beta} \\ -\boldsymbol{\alpha} \end{bmatrix} = 0.$$

Since the columns of $[V, W] \in \mathbb{R}^{n \times n}$ form a basis for \mathbb{R}^n by our construction, it must hold that $\boldsymbol{\beta} = 0$ and $-\boldsymbol{\alpha} = 0$. In particular, since $\boldsymbol{\alpha}$ vanishes, $A\mathbf{w}_1, \dots, A\mathbf{w}_{n-k}$ are linearly independent. \square

Let us then return to the unique solvability of (1). In order for the ‘unique solvability conditions’ (6) to hold for all $\mathbf{b} \in \mathbb{R}^m$, one must obviously have $R(A) = \mathbb{R}^m$ and $N(A) = \{0\}$. In particular, this yields

$$\dim(R(A)) + \dim(N(A)) = m + 0 = m.$$

By virtue of Theorem 1.1, this is possible only if $m = n$ and thus A is a square matrix. In other words, if $m \neq n$, i.e., A is not square, either (6) does not have any solution for some $\mathbf{b} \in \mathbb{R}^m$, or no solution for (6) is unique, or both.

Let us complete this section by considering the important special case of square matrices, i.e., $m = n$. According to the above considerations, (1) has a solution for all $\mathbf{b} \in \mathbb{R}^n$ if and only if $R(A) = \mathbb{R}^n$, i.e., $\text{rank}(A) = \dim(R(A)) = n$. However, due to Theorem 1.1, this also guarantees that $\dim(N(A)) = 0$, i.e., $N(A) = \{0\}$, and the solution is also unique. Vice versa, the solution to (1) is unique if and only if $\dim(N(A)) = 0$, which also guarantees the existence of a solution by virtue of Theorem 1.1. To summarize, for square matrices the existence and the uniqueness of a solution to (1) are equivalent conditions.

There are also several other conditions that guarantee the unique solvability of (1) for all $\mathbf{b} \in \mathbb{R}^n$ when $m = n$, as indicated by the following corollary.

COROLLARY 1.1. *Let $m = n$ in (1). There exists a unique solution to (1) for any $\mathbf{b} \in \mathbb{R}^n$ if and only if one of the following equivalent conditions holds:*

- (1) $R(A) = \mathbb{R}^n$,
- (2) A has linearly independent columns (or rows),
- (3) $N(A) = \{0\}$,
- (4) $\det(A) \neq 0$, or

¹Such $\mathbf{w}_1, \dots, \mathbf{w}_{n-k}$ could be constructed, e.g., via a variant of the Gram–Schmidt orthogonalization process considered in Chapter 3.

(5) 0 is not an eigenvalue of A .

PROOF. The conditions (1) and (3) were already covered in the above discussion. The condition (2) is equivalent to (1) due to (5). The condition (4) is equivalent to (2) by fundamental properties of the determinant (details are omitted). Finally, having 0 as an eigenvalue is equivalent to having a nontrivial nullspace, i.e., (5) and (3) are equivalent conditions (eigenvalues will be considered in detail in Chapter 2). \square

To sum up, if any of the above conditions (1)–(5) holds for $A \in \mathbb{R}^{n \times n}$, then (1) is uniquely solvable for any $\mathbf{b} \in \mathbb{R}^n$. This unique solution depends linearly on the right-hand side of the equation: Let $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^n$ be the solutions of (1) for the right-hand sides $\mathbf{b}_1, \mathbf{b}_2 \in \mathbb{R}^n$, respectively. Then for any $\alpha, \beta \in \mathbb{R}$,

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

and thus $\alpha\mathbf{x}_1 + \beta\mathbf{x}_2$ is the solution of (1) corresponding to $\mathbf{b} = \alpha\mathbf{b}_1 + \beta\mathbf{b}_2$. Because any linear map can be represented as a matrix (in a given basis), it follows that the solution to (1) can be given as

$$(10) \quad \mathbf{x} = A^{-1}\mathbf{b},$$

for some $A^{-1} \in \mathbb{R}^{n \times n}$, if any of the (equivalent) conditions in Corollary 1.1 holds. The matrix A^{-1} is called the *inverse* of A and if such exists, A is called *invertible*. In particular,

$$A^{-1}A = [A^{-1}\mathbf{a}_1, \dots, A^{-1}\mathbf{a}_n] = [\mathbf{e}_1, \dots, \mathbf{e}_n] = I$$

since the solution to $A\mathbf{x} = \mathbf{a}_j$, $j = 1, \dots, n$, is obviously the j th Cartesian basis vector $\mathbf{e}_j \in \mathbb{R}^n$. One can argue in the same way that also $AA^{-1} = I$, because A is obviously the inverse of A^{-1} (simply consider the symmetry between (1) and (10)).

2. Norms and inner products

To study the stability of the solution to (1) with respect to perturbations, inaccuracies or uncertainties in \mathbf{b} and A , we need tools for measuring the ‘size’ of a vector or a matrix. This section will also introduce the concept of orthogonality via inner products.

2.1. Vector norm. The size of a vector $\mathbf{x} \in \mathbb{R}^n$ is measured by a (vector) norm. The concept of norm is not unique, but there are many useful norms as we will see in what follows.

DEFINITION 2.1. A function $\|\cdot\| : \mathbb{R}^n \rightarrow \mathbb{R}$ is a norm if it satisfies

- (i) $\|\mathbf{x}\| \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$ and $\|\mathbf{x}\| = 0$ if and only if $\mathbf{x} = \mathbf{0}$,
- (ii) $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$,
- (iii) $\|\alpha\mathbf{x}\| = |\alpha|\|\mathbf{x}\|$ for all $\mathbf{x} \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}$.

Definition 2.1 formalizes what one expects from a reasonable measure of length. As an example, consider the space \mathbb{R}^2 and associate to each $\mathbf{x} \in \mathbb{R}^2$ a geometric vector $x_1\mathbf{i} + x_2\mathbf{j}$ as in Figure 1. A natural way to measure the length of such a vector \mathbf{x} is via the Pythagorean theorem, i.e., by introducing the *Euclidean norm* $\|\cdot\|_2 : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined via

$$\|\mathbf{x}\|_2 := (x_1^2 + x_2^2)^{1/2}.$$

This function obviously has the following properties:

- (i) **Positivity:** For all $\mathbf{x} \in \mathbb{R}^2$,

$$\|\mathbf{x}\|_2 \geq 0,$$

with the equality holding only if and only if $\mathbf{x} = \mathbf{0}$.

- (ii) **Scaling:** The distance measured by $\|\cdot\|_2$ is scaling invariant, that is,

$$\|\alpha\mathbf{x}\|_2 = |\alpha|\|\mathbf{x}\|_2$$

for all $\alpha \in \mathbb{R}$.




FIGURE 1. The geometric idea behind measuring the length of a vector in \mathbb{R}^2 .




FIGURE 2. The geometric idea behind the triangle inequality.

- (iii) **Triangle inequality:** The three vectors $\mathbf{x}, \mathbf{y}, \mathbf{x} + \mathbf{y} \in \mathbb{R}^2$ form a triangle in the plane; see Figure 2. For every triangle, the sum of the lengths of two sides is larger than the length of the third one, that is,

$$\|\mathbf{x} + \mathbf{y}\|_2 \leq \|\mathbf{x}\|_2 + \|\mathbf{y}\|_2.$$

Hence, $\|\cdot\|_2 : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a norm in accordance with Definition 2.1. In fact, the geometric properties of the Euclidean norm are the motivation for the general norm of Definition 2.1. Notice also that the Euclidean norm, a.k.a. the 2-norm, generalizes to an arbitrary spatial dimension by defining $\|\cdot\|_2 : \mathbb{R}^n \rightarrow \mathbb{R}$ through

$$(11) \quad \|\mathbf{x}\|_2 = \left(\sum_{i=1}^n |x_i|^2 \right)^{1/2}.$$

The norms used in linear algebra can be divided into generic norms and problem specific norms. Problem specific norms are used when vectors have some special interpretation; they can, e.g., represent coefficients in linear combinations of some elementary functions (cf., e.g., polynomials represented as linear combinations of monomials or a truncated Fourier series). Most theory in linear algebra is given in terms of generic norms that operate on generic vectors that have no special interpretation as such.

The most commonly used norm in \mathbb{R}^n is the Euclidean norm introduced in (11). Other regularly used norms include the family of p -norms (most often the case $p = 1$)

$$(12) \quad \|\mathbf{x}\|_p = \left(\sum_{i=1}^n |x_i|^p \right)^{1/p}, \quad 1 \leq p < \infty.$$

and the ∞ -norm, $\|\mathbf{x}\|_\infty = \max_i |x_i|$. The latter can be obtained as a limit of the p -norms when $p \rightarrow \infty$. The 1 and ∞ -norms are often employed because the related *operator norms* are easy to compute, as we will learn in the following. Proving that $\|\cdot\|_1$ and $\|\cdot\|_\infty$ (or more generally $\|\cdot\|_p$) really are norms is left as an exercise.

EXAMPLE 2.1. (a problem specific norm) Consider a first order polynomial $p(x) = \alpha_0 + \alpha_1 x$ with real coefficients. The size of such a polynomial can be measured by the square integral

$$\int_0^1 p^2(x) dx = \int_0^1 (\alpha_0^2 + 2\alpha_0\alpha_1 x + \alpha_1^2 x^2) dx = \alpha_0^2 + \alpha_0\alpha_1 + \frac{1}{3}\alpha_1^2.$$

On the other hand, any vector in \mathbb{R}^2 can be identified with a first order polynomial via

$$\begin{bmatrix} \alpha_0 \\ \alpha_1 \end{bmatrix} \longleftrightarrow \alpha_0 + \alpha_1 x.$$

Measuring the length of a vector defining the coefficients of a first order polynomial in, say, the 2-norm does not have any immediate interpretation. Instead it is more natural to measure the size of the coefficient vector using the function

$$(13) \quad \left\| \begin{bmatrix} \alpha_0 \\ \alpha_1 \end{bmatrix} \right\| := \left(\alpha_0^2 + \alpha_0 \alpha_1 + \frac{1}{3} \alpha_1^2 \right)^{1/2},$$

which is the square root of the square integral of the first order polynomial $\alpha_0 + \alpha_1 x$ introduced above. We will later prove that (13) really defines a norm, that is, it satisfies the conditions of Definition 2.1.

2.2. Inner product. To begin with, let $n = 2$ and consider two nonzero geometric vectors $x_1 \mathbf{i} + x_2 \mathbf{j}$ and $y_1 \mathbf{i} + y_2 \mathbf{j}$. It follows from the *cosine theorem* that the angle between the two vectors, say θ , satisfies

$$\|\mathbf{x}\|_2^2 + \|\mathbf{y}\|_2^2 - 2\|\mathbf{x}\|_2\|\mathbf{y}\|_2 \cos \theta = \|\mathbf{x} - \mathbf{y}\|_2^2 = \|\mathbf{x}\|_2^2 + \|\mathbf{y}\|_2^2 - 2(x_1 y_1 + x_2 y_2),$$

or, in other words,

$$\cos \theta = \frac{x_1 y_1 + x_2 y_2}{\|\mathbf{x}\|_2 \|\mathbf{y}\|_2}.$$

Motivated by this formula, we define the *dot product* as

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

The dot product introduces the concept of orthogonality in the space \mathbb{R}^2 , that is, when $\mathbf{x} \cdot \mathbf{y} = 0$, then the two vectors \mathbf{x} and \mathbf{y} form a right angle (or one/both of them vanish). We generalize the dot product to an arbitrary spatial dimension by defining

$$\mathbf{x} \cdot \mathbf{y} := \sum_{i=1}^n x_i y_i = \mathbf{x}^T \mathbf{y}$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. In particular, $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ are said to be orthogonal (in the sense of the dot product or the Euclidean inner product) if $\mathbf{x} \cdot \mathbf{y} = 0$, which actually also matches the geometric intuition for $n = 3$.

The basic properties of the dot product motivate the definition of a general *inner product*; it is easy to check that the dot product is an inner product.

DEFINITION 2.2. A mapping $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ is an inner product if the following properties hold:

- (i) *Positive definiteness:* $\langle \mathbf{x}, \mathbf{x} \rangle \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$ and $\langle \mathbf{x}, \mathbf{x} \rangle = 0$ if and only if $\mathbf{x} = \mathbf{0}$.
- (ii) *Symmetry:* $\langle \mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{y}, \mathbf{x} \rangle$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$.
- (iii) *Bilinearity:* $\langle \alpha \mathbf{x} + \beta \mathbf{y}, \mathbf{z} \rangle = \alpha \langle \mathbf{x}, \mathbf{z} \rangle + \beta \langle \mathbf{y}, \mathbf{z} \rangle$ for all $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathbb{R}^n$ and $\alpha, \beta \in \mathbb{R}$.

The conditions (ii) and (iii) together imply that also $\langle \mathbf{z}, \alpha \mathbf{x} + \beta \mathbf{y} \rangle = \alpha \langle \mathbf{z}, \mathbf{x} \rangle + \beta \langle \mathbf{z}, \mathbf{y} \rangle$. Take note that different inner products define different concepts of orthogonality in \mathbb{R}^n via the relation $\langle \mathbf{x}, \mathbf{y} \rangle = 0$.

Before moving on to vector norms that are induced by inner products, let us relate orthogonality and linear independence, starting with a definition.

DEFINITION 2.3. A set of vectors $\{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_k\} \subset \mathbb{R}^n$ is called orthogonal with respect to an inner product $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ if

$$\langle \mathbf{q}_i, \mathbf{q}_j \rangle = 0$$

for all $i \neq j$.

It turns out that an orthogonal set of vectors is always linearly independent, unless one of the vectors is the zero vector.

LEMMA 2.1. If the nonzero vectors $\mathbf{q}_1, \dots, \mathbf{q}_k \in \mathbb{R}^n$ compose a set of orthogonal vectors (with respect to any inner product), then they are linearly independent.

PROOF. Let $\alpha \in \mathbb{R}^k$ be such that

$$(14) \quad \sum_{i=1}^k \alpha_i \mathbf{q}_i = \mathbf{0}.$$

Taking the inner product of the left-hand side of this identity with \mathbf{q}_j , $j = 1, \dots, k$, and using (iii) of Definition 2.2 yields

$$\left\langle \sum_{i=1}^k \alpha_i \mathbf{q}_i, \mathbf{q}_j \right\rangle = \sum_{i=1}^k \alpha_i \langle \mathbf{q}_i, \mathbf{q}_j \rangle = \alpha_j \langle \mathbf{q}_j, \mathbf{q}_j \rangle,$$

where the second step is just the orthogonality assumption. On the other hand, taking the inner product of the right-hand side of (14) with \mathbf{q}_j results in

$$\langle \mathbf{0}, \mathbf{q}_j \rangle = 0 \langle \mathbf{0}, \mathbf{q}_j \rangle = 0$$

due to (iii) of Definition 2.2. Altogether we have deduced that

$$(15) \quad \alpha_j \langle \mathbf{q}_j, \mathbf{q}_j \rangle = \left\langle \sum_{i=1}^k \alpha_i \mathbf{q}_i, \mathbf{q}_j \right\rangle = \langle \mathbf{0}, \mathbf{q}_j \rangle = 0.$$

for any $j = 1, \dots, k$. Since $\langle \mathbf{q}_j, \mathbf{q}_j \rangle > 0$ by (i) of Definition 2.2 and our assumption on the orthogonal vectors, it must hold that $\alpha_j = 0$ for all $j = 1, \dots, k$. Thus the vectors $\mathbf{q}_1, \dots, \mathbf{q}_k$ are linearly independent. \square

A byproduct of Lemma 2.1 is that any set of n nonzero orthogonal vectors $\{\mathbf{q}_1, \dots, \mathbf{q}_n\} \subset \mathbb{R}^n$ forms a basis for the whole of \mathbb{R}^n . Indeed, it can be straightforwardly proven that any set of n linearly independent vectors must be a basis for an n -dimensional (sub)space. Hence, for any $\mathbf{x} \in \mathbb{R}^n$, there exists a representation

$$(16) \quad \mathbf{x} = \sum_{i=1}^n \alpha_i \mathbf{q}_i$$

with some $\alpha \in \mathbb{R}^n$. By taking the inner products of this identity with the orthonormal basis vectors $\mathbf{q}_1, \dots, \mathbf{q}_n$ one at a time, it can be deduced as in the proof of Lemma 2.1 that (cf. (15))

$$(17) \quad \alpha_j = \frac{\langle \mathbf{x}, \mathbf{q}_j \rangle}{\langle \mathbf{q}_j, \mathbf{q}_j \rangle}, \quad j = 1, \dots, n,$$

which provides a means to numerically compute the coefficients in the expansion (16). The (projection) formula (17) for the coefficients of (16) becomes even simpler if $\langle \mathbf{q}_j, \mathbf{q}_j \rangle = 1$ for all $j = 1, \dots, n$. Such a basis is called *orthonormal*, that is, the basis vectors are orthogonal and *normalized*, i.e., of unit length. This statement makes more sense after the definition of norms induced by inner products. The (numerical) construction of orthonormal bases is considered in Chapter 3.

To motivate the connection between inner products and norms, observe that the dot product obviously has a special connection to the Euclidean norm:

$$\|\mathbf{x}\|_2 = (\mathbf{x} \cdot \mathbf{x})^{1/2}$$

for all $\mathbf{x} \in \mathbb{R}^n$. This observation holds more generally: any inner product induces a norm.

LEMMA 2.2. *Let $\langle \cdot, \cdot \rangle$ be an inner product in \mathbb{R}^n . Then the function*

$$(18) \quad \|\mathbf{x}\| := \langle \mathbf{x}, \mathbf{x} \rangle^{1/2}, \quad \mathbf{x} \in \mathbb{R}^n,$$

is a norm. In addition, such a norm satisfies the Cauchy–Schwarz inequality

$$(19) \quad |\langle \mathbf{x}, \mathbf{y} \rangle| \leq \|\mathbf{x}\| \|\mathbf{y}\|$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$.

PROOF. We need to check that $\|\cdot\| : \mathbb{R}^n \rightarrow \mathbb{R}$ satisfies the conditions (i-iii) of Definition 2.1 and that (19) holds.

By the positive definiteness of an inner product,

$$\|\mathbf{x}\| := \langle \mathbf{x}, \mathbf{x} \rangle^{1/2} \geq 0,$$

where the equality holds only if $\mathbf{x} = 0$. This validates (i) of Definition 2.1. Furthermore, due to properties (ii) and (iii) of Definition 2.2,

$$\|\alpha \mathbf{x}\| = \langle \alpha \mathbf{x}, \alpha \mathbf{x} \rangle^{1/2} = (\alpha^2 \langle \mathbf{x}, \mathbf{x} \rangle)^{1/2} = |\alpha| \|\mathbf{x}\|,$$

which demonstrates that (iii) of Definition 2.1 is also satisfied.

Let us next prove the Cauchy–Schwarz inequality (19), which obviously holds if $\mathbf{x} = 0$ or $\mathbf{y} = 0$ since then both sides of (19) vanish due to (i) and (iii) of Definition 2.2 (note that, e.g., $\langle 0, \mathbf{y} \rangle = 0 \langle 0, \mathbf{y} \rangle = 0$). Let thus $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ be nonzero vectors and decompose \mathbf{x} as (cf. (17))

$$\mathbf{x} = \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{y}\|^2} \mathbf{y} + \left(\mathbf{x} - \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{y}\|^2} \mathbf{y} \right).$$

Using (ii) and (iii) of Definition 2.2, we thus obtain that

$$\begin{aligned} \|\mathbf{x}\|^2 &= \frac{\langle \mathbf{x}, \mathbf{y} \rangle^2}{\|\mathbf{y}\|^4} \|\mathbf{y}\|^2 + 2 \left\langle \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{y}\|^2} \mathbf{y}, \mathbf{x} - \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{y}\|^2} \mathbf{y} \right\rangle + \left\| \mathbf{x} - \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{y}\|^2} \mathbf{y} \right\|^2 \\ &= \frac{\langle \mathbf{x}, \mathbf{y} \rangle^2}{\|\mathbf{y}\|^2} + 2 \left(\frac{\langle \mathbf{x}, \mathbf{y} \rangle^2}{\|\mathbf{y}\|^2} - \frac{\langle \mathbf{x}, \mathbf{y} \rangle^2}{\|\mathbf{y}\|^4} \|\mathbf{y}\|^2 \right) + \left\| \mathbf{x} - \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{y}\|^2} \mathbf{y} \right\|^2 \\ &= \frac{\langle \mathbf{x}, \mathbf{y} \rangle^2}{\|\mathbf{y}\|^2} + \left\| \mathbf{x} - \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{y}\|^2} \mathbf{y} \right\|^2. \end{aligned}$$

Dropping the second (positive) term on the right-hand side and multiplying the ensuing inequality by $\|\mathbf{y}\|^2$ results in

$$\|\mathbf{x}\|^2 \|\mathbf{y}\|^2 \geq \langle \mathbf{x}, \mathbf{y} \rangle^2 \quad \text{or} \quad |\langle \mathbf{x}, \mathbf{y} \rangle| \leq \|\mathbf{x}\| \|\mathbf{y}\|,$$

which is the Cauchy–Schwarz inequality.

Finally, the triangle inequality, i.e. (ii) of Definition 2.1, follows by expanding,

$$\|\mathbf{x} + \mathbf{y}\|^2 = \|\mathbf{x}\|^2 + 2 \langle \mathbf{x}, \mathbf{y} \rangle + \|\mathbf{y}\|^2 \leq \|\mathbf{x}\|^2 + 2 \|\mathbf{x}\| \|\mathbf{y}\| + \|\mathbf{y}\|^2 = (\|\mathbf{x}\| + \|\mathbf{y}\|)^2$$

where the inequality is a trivial consequence of (19). \square

We complete this section by pointing out a one-to-one correspondence between positive definite matrices and inner products.

DEFINITION 2.4. A matrix $A \in \mathbb{R}^{n \times n}$ is called positive semidefinite if

$$(20) \quad \mathbf{x}^T A \mathbf{x} = \mathbf{x} \cdot A \mathbf{x} \geq 0$$

for all $\mathbf{x} \in \mathbb{R}^n$. If the second equality in (20) holds only for $\mathbf{x} = 0$, then A is called positive definite. (Nota bene: Sometimes symmetry of A is included in the definition of positive definiteness.)

Take note that any positive definite matrix $A \in \mathbb{R}^{n \times n}$ is invertible: If $\mathbf{x} \in N(A)$, i.e., $A \mathbf{x} = 0$, then

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

where the latter equality holds if and only if $\mathbf{x} = 0$. Hence, $N(A) = \{0\}$ and A is invertible due to the condition (3) in Corollary 1.1. However, most invertible matrices are not positive definite.

The following lemma demonstrates that any inner product in \mathbb{R}^n can be characterized with the help of the Euclidean inner product and a suitable symmetric positive definite matrix. In turn, any symmetric positive definite matrix defines an inner product.

LEMMA 2.3. For any inner product $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$, there exists a symmetric positive definite matrix $A \in \mathbb{R}^{n \times n}$ such that

$$(21) \quad \langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{y}^T A \mathbf{x}.$$

for all $\mathbf{x} \in \mathbb{R}^n$. On the other hand, the formula (21) defines an inner product for any symmetric positive definite $A \in \mathbb{R}^{n \times n}$.

PROOF. We start with the second part of the claim. Let $A \in \mathbb{R}^{n \times n}$ be an arbitrary symmetric positive definite matrix and set

$$\langle \mathbf{x}, \mathbf{y} \rangle_A := \mathbf{y}^T A \mathbf{x}.$$

We need to check that $\langle \cdot, \cdot \rangle_A : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ is an inner product, that is, we need to validate the conditions (i–iii) in Definition 2.2. (i) As A is positive definite,

$$\langle \mathbf{x}, \mathbf{x} \rangle_A = \mathbf{x}^T A \mathbf{x} \geq 0 \quad \text{for all } \mathbf{x} \in \mathbb{R}^n,$$

where the equality holds if and only if $\mathbf{x} = \mathbf{0}$. (ii) The symmetry requirement for an inner product is satisfied due to the assumed symmetry of A and the symmetry of the Euclidean inner product:

$$\langle \mathbf{x}, \mathbf{y} \rangle_A = \mathbf{y}^T A \mathbf{x} = \mathbf{y}^T A^T \mathbf{x} = (A \mathbf{y})^T \mathbf{x} = \mathbf{x}^T A \mathbf{y} = \langle \mathbf{y}, \mathbf{x} \rangle_A.$$

(iii) The bilinearity follows from the (bi)linearity of matrix product (homework).

To complete the proof, let $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ be an arbitrary inner product. Our aim is to construct a symmetric positive definite matrix $A \in \mathbb{R}^{n \times n}$ such that

$$\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{y}^T A \mathbf{x}.$$

Expanding $\mathbf{x} = \sum_{i=1}^n x_i \mathbf{e}_i$ and $\mathbf{y} = \sum_{j=1}^n y_j \mathbf{e}_j$ in the Cartesian basis, it follows that

$$\langle \mathbf{x}, \mathbf{y} \rangle = \left\langle \sum_{i=1}^n x_i \mathbf{e}_i, \sum_{j=1}^n y_j \mathbf{e}_j \right\rangle = \sum_{i,j=1}^n x_i y_j \langle \mathbf{e}_i, \mathbf{e}_j \rangle,$$

where we repeatedly employed (iii) of Definition 2.2. We now define the matrix $A \in \mathbb{R}^{n \times n}$ elementwise as $a_{ji} = \langle \mathbf{e}_i, \mathbf{e}_j \rangle$, $i, j = 1, \dots, n$. Clearly, $a_{ij} = a_{ji}$, i.e., A is symmetric. Moreover,

$$\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i,j=1}^n x_i y_j a_{ji} = \sum_{j=1}^n y_j \sum_{i=1}^n a_{ji} x_i = \mathbf{y}^T A \mathbf{x}$$

by the definition of matrix-vector product. Finally, the constructed A is positive definite since

$$\mathbf{x}^T A \mathbf{x} = \langle \mathbf{x}, \mathbf{x} \rangle > 0 \quad \text{for all } 0 \neq \mathbf{x} \in \mathbb{R}^n$$

because $\langle \cdot, \cdot \rangle$ was assumed to be an inner product; see (i) of Definition 2.2. \square

2.3. Matrix norm. In addition to measuring the length of a vector \mathbf{x} , we also need to be able to measure the ‘size’ of a given matrix $A \in \mathbb{R}^{m \times n}$. Let us begin the discussion with an example.

EXAMPLE 2.2. Let

$$A = \begin{bmatrix} 1 & 0.2 \\ 0.1 & 0.8 \end{bmatrix}.$$

The matrix A associates every vector $\mathbf{x} \in \mathbb{R}^2$ with another one $A \mathbf{x} \in \mathbb{R}^2$. Such a mapping can be visualized (at least) in two ways:

- (1) Draw some pixel image into the plane (e.g., a clown). Map the coordinates of the pixel corners as in Figure 3. Coloring the skewed pixels with the same color as their preimages results in a deformed image that describes the action of A . See Figure 3.
- (2) Due to the linearity of matrix multiplication,

$$A \mathbf{x} = \|\mathbf{x}\|_2 A \left(\frac{\mathbf{x}}{\|\mathbf{x}\|_2} \right).$$

Hence, all information in the matrix A is actually contained in the set

$$A(S) = \{A \mathbf{v} \mid \|\mathbf{v}\|_2 = 1\}$$




FIGURE 3. Two pixelwise discretizations of a clown and their images under the action of the matrix A from Example 2.2.




FIGURE 4. The image of the Euclidean unit sphere under the action of the matrix A from Example 2.2. The red circle is the unit sphere and the blue ellipse its image.

since any vector can be given as a vector of unit length times a suitable scalar. The set $A(S)$ is compared with its preimage, i.e. the unit circle S (or, more generally, the unit sphere), in Figure 4.

The second way to visualize the multiplication by a matrix suggest a way to measure its size.

DEFINITION 2.5. Let $\|\cdot\|$ be some vector norm. The corresponding operator norm for a matrix $A \in \mathbb{R}^{m \times n}$ is defined as

$$\|A\|_{\text{op}} := \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|A\mathbf{x}\|}{\|\mathbf{x}\|} = \max_{\|\mathbf{x}\|=1} \|A\mathbf{x}\|.$$

The definition of an operator norm depends on the considered vector norm. In other words, different vector norms induce different operator norms. The operator norm measures the maximal stretching of the unit sphere under the action of a matrix A ; note that even the unit sphere itself, i.e.,

$$S := \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| = 1\}$$

depends on the investigated vector norm. In Example 2.2, where the Euclidean norm was employed, the operator norm is simply the semi-major axis of the blue ellipse in Figure 4, i.e., $\|A\|_{\text{op}} \approx 1.08$.

In linear algebra, it is customary to drop the subscript from $\|\cdot\|_{\text{op}}$ and denote both a vector norm and the corresponding operator norm in the same way. We follow this convention during these lectures. In particular, the operator norms induced by 1, ∞ and 2-norms are denoted by $\|\cdot\|_1$, $\|\cdot\|_\infty$ and $\|\cdot\|_2$, respectively, or simply by $\|\cdot\|$ if the considered vector norm is clear from the context.

All operator norms share two useful properties, as indicated by the following lemma.

LEMMA 2.4. Any operator norm satisfies,

$$(22) \quad \|AB\| \leq \|A\|\|B\|, \quad A \in \mathbb{R}^{l \times m}, B \in \mathbb{R}^{m \times n}$$

and

$$(23) \quad \|A\mathbf{x}\| \leq \|A\|\|\mathbf{x}\|, \quad A \in \mathbb{R}^{m \times n}, \mathbf{x} \in \mathbb{R}^n.$$

PROOF. Homework. \square

Computing the value of an operator norm requires one to find the maximum of $\|A\mathbf{x}\|\|\mathbf{x}\|^{-1}$ over \mathbb{R}^n or equivalently that of $\|A\mathbf{x}\|$ over the unit sphere corresponding to the examined vector norm. For the operator norms induced by the vector norms $\|\cdot\|_1$, $\|\cdot\|_2$ and $\|\cdot\|_\infty$, this maximum can be found explicitly. This is one of the reasons why these three norms are used as generic norms in linear algebra.

The 2-norm of a matrix can be computed using the following lemma that utilizes eigenvalues of a symmetric matrix. Recall that an eigenvalue of a matrix $A \in \mathbb{R}^{n \times n}$ is such $\lambda \in \mathbb{C}$ that

$$(24) \quad A\mathbf{x} = \lambda\mathbf{x}$$

for some $0 \neq \mathbf{x} \in \mathbb{C}^n$ that is called an eigenvector corresponding to λ .² Eigenvalues and eigenvectors will be studied in detail in Chapter 2. Here, we only need the following result that should be familiar from the basic matrix algebra course: Any symmetric matrix $B \in \mathbb{R}^{n \times n}$ has n orthonormal real eigenvectors (with respect to the dot product/Euclidean inner product), that is,

$$B\mathbf{q}_j = \lambda_j\mathbf{q}_j, \quad j = 1, \dots, n,$$

with $\|\mathbf{q}_j\|_2 = 1$ for all $j = 1, \dots, n$ and $\mathbf{q}_j \cdot \mathbf{q}_k = \mathbf{q}_j^T \mathbf{q}_k = 0$ for all $j \neq k$. Moreover, B can be decomposed as

$$(25) \quad B = Q\Lambda Q^T$$

where $Q = [\mathbf{q}_1, \dots, \mathbf{q}_n] \in \mathbb{R}^{n \times n}$ has the orthonormal eigenvectors as its columns and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n) \in \mathbb{R}^{n \times n}$ is a diagonal matrix composed of the real eigenvalues (some of which may be same). In particular, $Q^T Q = I$ because the columns of Q are orthonormal, i.e., Q is an orthogonal matrix.

LEMMA 2.5. For any $A \in \mathbb{R}^{m \times n}$,

$$\|A\|_2 = (\lambda_{\max}(A^T A))^{1/2},$$

where $\lambda_{\max}(A^T A)$ is the largest eigenvalue of the symmetric matrix $A^T A$.

PROOF. By the monotonicity of the second power,

$$\|A\|_2^2 := \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|A\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2} = \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \frac{\mathbf{x}^T A^T A \mathbf{x}}{\|\mathbf{x}\|_2^2} = \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \frac{\mathbf{x}^T Q\Lambda Q^T \mathbf{x}}{\|\mathbf{x}\|_2^2},$$

where $Q\Lambda Q^T$ is an eigendecomposition of the symmetric matrix $A^T A \in \mathbb{R}^{n \times n}$, with an orthogonal $Q \in \mathbb{R}^{n \times n}$ and a diagonal $\Lambda \in \mathbb{R}^{n \times n}$ as in (25). Defining $\mathbf{y} = Q^T \mathbf{x}$ and noting that $\mathbf{x} = Q\mathbf{y}$ runs through the whole of \mathbb{R}^n when \mathbf{y} does so, we get

$$(26) \quad \|A\|_2^2 = \max_{0 \neq \mathbf{y} \in \mathbb{R}^n} \frac{\mathbf{y}^T \Lambda \mathbf{y}}{\|Q\mathbf{y}\|_2^2} = \max_{0 \neq \mathbf{y} \in \mathbb{R}^n} \frac{\sum_{i=1}^n \lambda_i \mathbf{y}_i^2}{\mathbf{y}^T Q^T Q \mathbf{y}} = \max_{0 \neq \mathbf{y} \in \mathbb{R}^n} \frac{\sum_{i=1}^n \lambda_i \mathbf{y}_i^2}{\|\mathbf{y}\|_2^2},$$

where λ_i , $i = 1, \dots, n$, are the real eigenvalues of $A^T A \in \mathbb{R}^{n \times n}$ and we also employed the identity $Q^T Q = I$.

Notice that all λ_i , $i = 1, \dots, n$, are nonnegative (because $A^T A$ is positive semidefinite):

$$0 \leq \|A\mathbf{q}_i\|_2^2 = \mathbf{q}_i^T A^T A \mathbf{q}_i = \lambda_i \mathbf{q}_i^T \mathbf{q}_i = \lambda_i,$$

where we used the orthonormality of the (eigen)columns of Q (cf. (25)). A straightforward estimate thus gives

$$\|A\|_2^2 \leq \max_{0 \neq \mathbf{y} \in \mathbb{R}^n} \frac{\lambda_{\max}(A^T A) \sum_{i=1}^n \mathbf{y}_i^2}{\|\mathbf{y}\|_2^2} = \lambda_{\max}(A^T A) \max_{0 \neq \mathbf{y} \in \mathbb{R}^n} \frac{\|\mathbf{y}\|_2^2}{\|\mathbf{y}\|_2^2} = \lambda_{\max}(A^T A).$$

²In particular, a real matrix may have complex eigenvalues and eigenvectors.

Letting $1 \leq k \leq n$ be such that $\lambda_k = \lambda_{\max}(A^T A)$ and choosing $\mathbf{y} = \mathbf{e}_k$ in (26) to be the k th Cartesian basis vector, we finally get

$$\|A\|_2^2 \geq \lambda_{\max}(A^T A),$$

which completes the proof as we have altogether established that $\lambda_{\max}(A^T A) \leq \|A\|_2^2 \leq \lambda_{\max}(A^T A)$. \square

COROLLARY 2.1. *Lemma 2.5 remains valid if $\lambda_{\max}(A^T A)$ is replaced by $\lambda_{\max}(AA^T)$.*

PROOF. It can be shown that the $\min\{n, m\}$ largest eigenvalues of the symmetric matrices $A^T A \in \mathbb{R}^{n \times n}$ and $AA^T \in \mathbb{R}^{m \times m}$ are the same³ (when counted according to their algebraic multiplicity); a proof will be presented in connection to singular values in Chapter 3. The assertion is an immediate consequence of the aforementioned fact. \square

It can also be shown that the operator norms of $A \in \mathbb{R}^{m \times n}$ corresponding to the 1 and ∞ vector norms can be computed simply as

$$\|A\|_1 = \max_{1 \leq j \leq n} \sum_{i=1}^m |a_{ij}| \quad \text{and} \quad \|A\|_\infty = \max_{1 \leq i \leq m} \sum_{j=1}^n |a_{ij}|,$$

that is, as the maximal absolute column and row sums of A , respectively.

The operator norms form a subclass of all *matrix norms*. In other words, every matrix norm is not an operator norm, but all operator norms are matrix norms as indicated by the lemma succeeding the following definition.

DEFINITION 2.6. *A function $\|\cdot\| : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is a matrix norm if*

- (i) $\|A\| \geq 0$ and $\|A\| = 0$ if and only if $A = 0$ is the zero matrix,
- (ii) $\|A + B\| \leq \|A\| + \|B\|$,
- (iii) $\|\alpha A\| = |\alpha| \|A\|$

for all $A, B \in \mathbb{R}^{m \times n}$ and $\alpha \in \mathbb{R}$.

LEMMA 2.6. *Let $\|\cdot\|$ denote a vector norm and also the corresponding operator norm. Then $\|\cdot\|$ is also a matrix norm in accordance with Definition 2.6.*

PROOF. Let us verify the conditions (i–iii) of Definition 2.6 one at a time.

(i) Obviously,

$$\|A\| := \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|A\mathbf{x}\|}{\|\mathbf{x}\|} \geq 0.$$

Moreover, if any element of $A \in \mathbb{R}^{m \times n}$ is nonzero, it is easy to find $\mathbf{x} \in \mathbb{R}^n$ such that $A\mathbf{x}$ is also nonzero. Hence, $\|A\| > 0$ if A is not the zero matrix.

(ii) Due to the definition of the operator norm and the linearity of matrix multiplication,

$$\|A + B\| = \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|A\mathbf{x} + B\mathbf{x}\|}{\|\mathbf{x}\|} \leq \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \left(\frac{\|A\mathbf{x}\|}{\|\mathbf{x}\|} + \frac{\|B\mathbf{x}\|}{\|\mathbf{x}\|} \right) \leq \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|A\mathbf{x}\|}{\|\mathbf{x}\|} + \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|B\mathbf{x}\|}{\|\mathbf{x}\|},$$

where the second step is the triangle inequality for the vector norm and the last step corresponds to a simple property of the max function: “The maximum of a sum is less than the sum of the maxima of the summands”.

(iii) Finally, it holds that

$$\|\alpha A\| = \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|\alpha A\mathbf{x}\|}{\|\mathbf{x}\|} = \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \frac{|\alpha| \|A\mathbf{x}\|}{\|\mathbf{x}\|} = |\alpha| \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|A\mathbf{x}\|}{\|\mathbf{x}\|} = |\alpha| \|A\|$$

by the scaling property of the vector norm. \square

³This is not very difficult to believe: $AA^T \mathbf{v} = \lambda \mathbf{v} \Rightarrow A^T A(A^T \mathbf{v}) = \lambda(A^T \mathbf{v})$ and $A^T A \mathbf{v} = \lambda \mathbf{v} \Rightarrow AA^T(A\mathbf{v}) = \lambda(A\mathbf{v})$.

As indicated by Definition 2.6, a matrix norm can simply be defined as a function from the space of matrices $\mathbb{R}^{m \times n}$ to nonnegative real numbers — as long as it satisfies the conditions of Definition 2.6. Examples of such directly defined matrix norms are the Frobenius norm,

$$\|A\|_F := \left(\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2 \right)^{1/2},$$

and the max-norm,

$$\|A\|_{\max} = \max_{1 \leq i \leq m} \max_{1 \leq j \leq n} |a_{ij}|.$$

It is straightforward to verify that both of these really are matrix norms in accordance with Definition 2.6. (The proofs are analogous to showing that the 2 and ∞ vector norms are norms.)

3. Stability of the solution

In this section, we consider the equation (1) in the case that $m = n$ and the matrix A is invertible, i.e., $A \in \mathbb{R}^{n \times n}$ has an inverse $A^{-1} \in \mathbb{R}^{n \times n}$. Recall that the existence of an inverse A^{-1} is equivalent to the unique solvability of (1) for all $\mathbf{b} \in \mathbb{R}^n$, and Corollary 1.1 lists conditions that guarantee the latter. In this section, we employ the norms introduced in the previous section to study how perturbations of A and \mathbf{b} in (1) affect the corresponding solution.

Let \mathbf{x} be the solution to the linear system (1), $\delta A \in \mathbb{R}^{n \times n}$ and $\delta \mathbf{b} \in \mathbb{R}^n$. The perturbations δA and $\delta \mathbf{b}$ model (additive) uncertainties and inaccuracies in A and \mathbf{b} , respectively. The general intuition should thus be that δA and $\delta \mathbf{b}$ are small(ish) compared to A and \mathbf{b} , respectively. To be more precise, we consider the perturbed linear equation

$$(27) \quad (A + \delta A)\tilde{\mathbf{x}} = \mathbf{b} + \delta \mathbf{b}.$$

and study the difference between the solution $\mathbf{x} \in \mathbb{R}^n$ of (1) and the ‘perturbed solution’ $\tilde{\mathbf{x}} \in \mathbb{R}^n$ of (27).

To make such considerations meaningful, we need to first pose a condition under which the unique solvability of (27) follows from that of (1). The topological interpretation of the following theorem is that the invertible matrices form an open subset in the space of all matrices equipped with some operator topology. In layman’s terms, if a matrix is invertible, then so are all nearby matrices.

THEOREM 3.1. *Let $\|\cdot\|$ be an operator norm. If $A \in \mathbb{R}^{n \times n}$ is invertible and $\delta A \in \mathbb{R}^{n \times n}$ satisfies*

$$(28) \quad \|\delta A\| < \frac{1}{\|A^{-1}\|},$$

then $A + \delta A$ is also invertible.

PROOF. Our aim is to prove that $A + \delta A$ satisfies the condition (3) of Corollary 1.1, from which its invertibility follows. The reverse triangle inequality (homework) for $0 \neq \mathbf{z} \in \mathbb{R}^n$ gives,

$$(29) \quad \frac{\|(A + \delta A)\mathbf{z}\|}{\|\mathbf{z}\|} \geq \frac{\|A\mathbf{z}\| - \|\delta A\mathbf{z}\|}{\|\mathbf{z}\|} \geq \min_{0 \neq \mathbf{z} \in \mathbb{R}^n} \frac{\|A\mathbf{z}\|}{\|\mathbf{z}\|} - \max_{0 \neq \mathbf{z} \in \mathbb{R}^n} \frac{\|\delta A\mathbf{z}\|}{\|\mathbf{z}\|}.$$

We make the change of variables $\mathbf{y} = A\mathbf{z}$, i.e. $\mathbf{z} = A^{-1}\mathbf{y}$, in the first term on the right-hand side of (29), which leads to

$$\min_{0 \neq \mathbf{z} \in \mathbb{R}^n} \frac{\|A\mathbf{z}\|}{\|\mathbf{z}\|} = \min_{0 \neq \mathbf{y} \in \mathbb{R}^n} \frac{\|\mathbf{y}\|}{\|A^{-1}\mathbf{y}\|} = \left(\max_{0 \neq \mathbf{y} \in \mathbb{R}^n} \frac{\|A^{-1}\mathbf{y}\|}{\|\mathbf{y}\|} \right)^{-1} = \|A^{-1}\|^{-1}$$

due to the monotonicity of the inverse power on the positive real axis \mathbb{R}_+ . Substituting this back in (29) and recalling the definition of an operator norm, yields

$$\frac{\|(A + \delta A)\mathbf{z}\|}{\|\mathbf{z}\|} \geq \frac{1}{\|A^{-1}\|} - \|\delta A\| > 0,$$

where the strict inequality is our assumption. Equivalently,

$$\|(A + \delta A)\mathbf{z}\| \geq \left(\frac{1}{\|A^{-1}\|} - \|\delta A\| \right) \|\mathbf{z}\| > 0$$

for all $0 \neq \mathbf{z} \in \mathbb{R}^n$. Hence, the nullspace $N(A + \delta A)$ cannot contain any nonzero vectors. \square

The following corollary indicates that the norm of the inverse of $A + \delta A$ can, in fact, be controlled by the norms of the inverse A^{-1} and the perturbation δA .

COROLLARY 3.1. *Under the assumptions of Theorem 3.1, it holds that*

$$\|(A + \delta A)^{-1}\| \leq \frac{\|A^{-1}\|}{1 - \|\delta A\|\|A^{-1}\|}$$

for any operator norm $\|\cdot\|$.

PROOF. Using (22), we first of all obtain

$$\|(A + \delta A)^{-1}\| = \|(I + \delta AA^{-1})A^{-1}\| = \|A^{-1}(I + \delta AA^{-1})^{-1}\| \leq \|A^{-1}\| \|(I + \delta AA^{-1})^{-1}\|.$$

By the definition of an operator norm,

$$(30) \quad \|(I + \delta AA^{-1})^{-1}\| = \max_{0 \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|(I + \delta AA^{-1})^{-1}\mathbf{x}\|}{\|\mathbf{x}\|} = \max_{0 \neq \mathbf{y} \in \mathbb{R}^n} \frac{\|\mathbf{y}\|}{\|(I + \delta AA^{-1})\mathbf{y}\|},$$

where the second step corresponds to the change of variables $\mathbf{x} = (I + \delta AA^{-1})\mathbf{y}$. Furthermore, the denominator on the right-hand side of (30) can be estimated with the help of the reverse triangle inequality as follows:

$$(31) \quad \|(I + \delta AA^{-1})\mathbf{y}\| \geq \|\mathbf{y}\| - \|\delta AA^{-1}\mathbf{y}\| \geq (1 - \|\delta A\|\|A^{-1}\|)\|\mathbf{y}\| > 0,$$

where we also twice used (23) as well as our assumption on the size of $\|A^{-1}\|$. Plugging (30) and (31) in turns in the first estimate of this proof, we finally get

$$\begin{aligned} \|(A + \delta A)^{-1}\| &\leq \|A^{-1}\| \max_{0 \neq \mathbf{y} \in \mathbb{R}^n} \frac{\|\mathbf{y}\|}{\|(I + \delta AA^{-1})\mathbf{y}\|} \leq \|A^{-1}\| \max_{0 \neq \mathbf{y} \in \mathbb{R}^n} \frac{\|\mathbf{y}\|}{(1 - \|\delta A\|\|A^{-1}\|)\|\mathbf{y}\|} \\ &= \frac{\|A^{-1}\|}{1 - \|\delta A\|\|A^{-1}\|}, \end{aligned}$$

which completes the proof. \square

By subtracting the equations (1) and (27), we obtain a linear system that determines the error $\tilde{\mathbf{x}} - \mathbf{x}$, that is,

$$(32) \quad (A + \delta A)(\tilde{\mathbf{x}} - \mathbf{x}) = \delta \mathbf{b} - \delta A \mathbf{x}.$$

The following theorem relates the (relative) error to the relative sizes of the perturbations, i.e. $\|\delta \mathbf{b}\|/\|\mathbf{b}\|$ and $\|\delta A\|/\|A\|$, and the condition number of A , defined as

$$\kappa(A) := \|A\|\|A^{-1}\|.$$

Take note that the condition number of a matrix depends on the considered (operator) norm.

THEOREM 3.2. *Suppose the assumptions of Theorem 3.1 are valid. Then it holds that*

$$(33) \quad \frac{\|\tilde{\mathbf{x}} - \mathbf{x}\|}{\|\mathbf{x}\|} \leq \frac{\kappa(A)}{1 - \frac{\|\delta A\|}{\|A\|}\kappa(A)} \left(\frac{\|\delta \mathbf{b}\|}{\|\mathbf{b}\|} + \frac{\|\delta A\|}{\|A\|} \right)$$

for any $\|\cdot\|$ operator norm and the corresponding condition number $\kappa(A)$.

PROOF. According to Theorem 3.1, the matrix $A + \delta A \in \mathbb{R}^{n \times n}$ is invertible and the error $\tilde{\mathbf{x}} - \mathbf{x}$ can thus be solved from (32) as

$$\tilde{\mathbf{x}} - \mathbf{x} = (A + \delta A)^{-1}(\delta \mathbf{b} - \delta A \mathbf{x}).$$

In particular,

$$\|\tilde{\mathbf{x}} - \mathbf{x}\| \leq \|(A + \delta A)^{-1}\| \|\delta \mathbf{b} - \delta A \mathbf{x}\| \leq \|(A + \delta A)^{-1}\| (\|\delta \mathbf{b}\| + \|\delta A \mathbf{x}\|).$$

due to basic properties of an operator norm and the triangle inequality. The stability result of Corollary 3.1 yields

$$\|\tilde{\mathbf{x}} - \mathbf{x}\| \leq \frac{\|A^{-1}\|}{1 - \|\delta A\| \|A^{-1}\|} (\|\delta \mathbf{b}\| + \|\delta A \mathbf{x}\|).$$

Dividing by $\|\mathbf{x}\|$ and using the estimate $\|\delta A \mathbf{x}\| \leq \|\delta A\| \|\mathbf{x}\|$ gives

$$(34) \quad \frac{\|\tilde{\mathbf{x}} - \mathbf{x}\|}{\|\mathbf{x}\|} \leq \frac{\|A^{-1}\|}{1 - \|\delta A\| \|A^{-1}\|} \left(\frac{\|\delta \mathbf{b}\|}{\|\mathbf{x}\|} + \|\delta A\| \right) = \frac{\|A\| \|A^{-1}\|}{1 - \frac{\|\delta A\|}{\|A\|} \|A^{-1}\| \|A\|} \left(\frac{\|\delta \mathbf{b}\|}{\|A\| \|\mathbf{x}\|} + \frac{\|\delta A\|}{\|A\|} \right),$$

where the latter step is mere algebraic manipulation. Since $\|\mathbf{b}\| = \|A \mathbf{x}\| \leq \|A\| \|\mathbf{x}\|$, we finally obtain

$$(35) \quad \frac{\|\tilde{\mathbf{x}} - \mathbf{x}\|}{\|\mathbf{x}\|} \leq \frac{\|A^{-1}\| \|A\|}{1 - \frac{\|\delta A\|}{\|A\|} \|A^{-1}\| \|A\|} \left(\frac{\|\delta \mathbf{b}\|}{\|\mathbf{b}\|} + \frac{\|\delta A\|}{\|A\|} \right).$$

Substituting the definition of the condition number $\kappa(A)$ completes the proof. \square

4. Solving $Ax = b$ on a computer

Computers operate using finite precision arithmetic: numbers are not stored exactly nor are operations performed exactly. A simple model for error in arithmetic operations is

$$(36) \quad f(a \odot b) = (1 + \delta)(a \odot b),$$

where $a, b \in \mathbb{R}$, \odot is one of the elementary operations, i.e. $\odot = +, -, *, /$, and $|\delta| \leq u$, where u is the machine unit or the machine epsilon. The result of a given operation in finite precision is denoted by $f(a \odot b)$. The size of the machine epsilon depends on how numbers are represented, but it is typically of the size $u = 2.22 \cdot 10^{-16}$.

As an example, due to the finite precision, a numerically computed LU decomposition⁴ produced, say, by MATLAB is not exact, but only an approximation of an exact one. The triangular matrices $L \in \mathbb{R}^{n \times n}$ and $U \in \mathbb{R}^{n \times n}$ given by the computer do not exactly satisfy $A = LU$ for a given $A \in \mathbb{R}^{n \times n}$, but instead

$$LU = A + \delta A.$$

Hence, the utilization of the LU decomposition does not solve the equation (1), but a related linear system.

$$(A + \delta A)\tilde{\mathbf{x}} = \mathbf{b}.$$

An interesting and practical question is: How close is \mathbf{x} to $\tilde{\mathbf{x}}$?

The answer can be given in two steps,

- Estimate the size of δA . For the LU decomposition, it can be shown that

$$\|\delta A\| \leq nu \|L\| \|U\|$$

for the Frobenius, 1 and ∞ norms. In most cases, the ratio $(\|L\| \|U\|) / \|A\|$ is of moderate size, but one can also construct examples for which it tends to infinity. Fortunately, such examples are rarely encountered in practice.

- Use Theorem 3.2 to obtain an estimate for the relative change in the solution:

$$\frac{\|\tilde{\mathbf{x}} - \mathbf{x}\|}{\|\mathbf{x}\|} \leq \frac{\kappa(A)}{1 - \frac{\|\delta A\|}{\|A\|} \kappa(A)} \frac{\|\delta A\|}{\|A\|}.$$

Observe that the relative error $\|\delta A\| / \|A\|$ is amplified by the condition number $\kappa(A)$ both in the numerator and the denominator.

⁴Recall that the LU decomposition is a way of factoring a matrix A as a product of a lower L and an upper U triangular matrix. It is commonly used by, e.g., mathematical software to solve systems of linear equations.

The heuristic of MATLAB is to assume that $\|\delta A\|/\|A\|$ behaves like the machine epsilon and to issue a warning if $\kappa(A)$ is sufficiently large.

To summarize, if $\kappa(A) \geq 10^{16}$, the given linear system cannot be solved on a computer using (mere) double precision arithmetic. Even small floating point errors are amplified by $\kappa(A)$ so that they become significant. To conclude: **All theoretically solvable linear systems cannot be solved on a computer!**

4.1. The bad example: Hilbert matrix. A classical example of a linear system that cannot be solved accurately on a computer arises from approximation theory. Consider approximating a given function f on the interval $(0, 1)$ by the n th degree polynomials P^n in the sense of least squares. In other words, one is interested in the solution of the problem

$$(37) \quad \min_{p \in P^n} \frac{1}{2} \int_0^1 (f - p)^2 dx.$$

This minimization problem is quadratic and can, in principle, be solved rather easily.

The naive approach is to represent the n th degree polynomials in the monomial basis $\{x^{j-1}\}_{j=1}^{n+1}$ as

$$p(x) := \sum_{j=1}^{n+1} \alpha_j x^{j-1},$$

which means that each polynomial in P^n has a one-to-one correspondence with a coefficient vector $\alpha \in \mathbb{R}^{n+1}$. Analogously, the original problem (37) is equivalent to a certain minimization problem over \mathbb{R}^{n+1} :

$$(38) \quad \min_{\alpha \in \mathbb{R}^{n+1}} \frac{1}{2} \int_0^1 \left(f - \sum_{j=1}^{n+1} \alpha_j x^{j-1} \right)^2 dx.$$

The solution of (38) can be characterized, e.g., by taking the gradient with respect to the coefficient vector and equating it to zero. For each $i = 1, \dots, n+1$, this corresponds to

$$\begin{aligned} \frac{\partial}{\partial \alpha_i} \frac{1}{2} \int_0^1 \left(f - \sum_{j=1}^{n+1} \alpha_j x^{j-1} \right)^2 dx &= \int_0^1 \left(\sum_{j=1}^{n+1} \alpha_j x^{j-1} - f \right) x^{i-1} dx \\ &= \sum_{j=1}^{n+1} \alpha_j \int_0^1 x^{i+j-2} dx - \int_0^1 f x^{i-1} dx = 0. \end{aligned}$$

These $n+1$ linear equations can be presented in a matrix form: Find $\alpha \in \mathbb{R}^{n+1}$ such that

$$(39) \quad H\alpha = \mathbf{b},$$

where $H_{ij} = \int_0^1 x^{i+j-2} dx$ and $b_i = \int_0^1 f x^{i-1} dx$ for $i, j = 1, \dots, n+1$. The integrals defining $H \in \mathbb{R}^{(n+1) \times (n+1)}$ can be computed by hand to obtain

$$H_{ij} = \frac{1}{i+j-1}, \quad i, j = 1, \dots, n+1.$$

The condition number of this so-called *Hilbert matrix* grows extremely rapidly as a function of the dimension n . As an example, for $n = 15$ the condition number is already so large that the linear system (39) cannot be accurately solved using a computer.

Regardless of the above explained difficulties, the original polynomial approximation problem (37) can be solved efficiently and accurately. The trick is to use some other basis polynomials in place of the monomials. For example, the use of Legendre polynomials leads to significantly improved stability.

Eigenvalue problems

This second chapter discusses eigenvalues of matrices and their applications. The main application considered in these notes is the definition of matrix valued functions such as the matrix power or matrix exponential. Such functions are used, e.g., to solve recursive equations or to study the behavior of a system of ordinary differential equations close to an equilibrium point. In addition, eigenvalues and eigenvectors serve as a (computational) tool in many problems arising, e.g., from mathematics, statistics, physics or computer sciences.

We will discuss the following themes:

- (1) Basic eigenvalue theory with emphasis on Hermitian matrices,
- (2) Similarity transformations,
- (3) Matrix exponential and its applications.

The eigenvalues and eigenvectors of a *square* matrix $A \in \mathbb{C}^{n \times n}$ are the solutions of the following problem: Find $(\lambda, \mathbf{v}) \in \mathbb{C} \times (\mathbb{C}^n \setminus \{0\})$ such that

$$(40) \quad A\mathbf{v} = \lambda\mathbf{v}.$$

It should be emphasized that an eigenvector can never be the zero vector. When considering eigenvalues, we need to employ complex numbers \mathbb{C} and the spaces of complex valued vectors \mathbb{C}^n and matrices $\mathbb{C}^{n \times n}$ as even real matrices may have complex eigenvalues and vectors.

Almost all theory we have discussed thus far during these lectures directly applies to vectors in \mathbb{C}^n . The only notable exception is the definition of an inner product: The symmetry, i.e. (ii) of Definition 2.2 in Chapter 1, is replaced by the *conjugate symmetry* condition

$$\langle \mathbf{x}, \mathbf{y} \rangle = \overline{\langle \mathbf{y}, \mathbf{x} \rangle}.$$

This also means that the bilinearity condition (iii) of Definition 2.2 turns into *sesquilinearity*, i.e., the linearity with respect to the second variable turns into *antilinearity*: $\langle \mathbf{x}, \alpha\mathbf{y} \rangle = \overline{\alpha}\langle \mathbf{x}, \mathbf{y} \rangle$. The Euclidean inner product and norm are defined in \mathbb{C}^n as

$$\langle \mathbf{x}, \mathbf{y} \rangle := \mathbf{x} \cdot \overline{\mathbf{y}} = \mathbf{y}^* \mathbf{x} \quad \text{and} \quad \|\mathbf{x}\|_2 := (\mathbf{x}^* \mathbf{x})^{1/2},$$

where the ‘overline’ marks componentwise complex conjugation and $*$ denotes the conjugate transpose,

$$\mathbf{x}^* = \overline{\mathbf{x}}^T.$$

In particular, if $\mathbf{x} = \mathbf{a} + i\mathbf{b}$ for some $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, then

$$\overline{\mathbf{x}} = \mathbf{a} - i\mathbf{b}.$$

1. Basic eigenvalue theory

The equation (40) can be reformulated as

$$(A - \lambda I)\mathbf{v} = 0.$$

In other words, if $\lambda \in \mathbb{C}$ is an eigenvalue of A , then the matrix $A - \lambda I$ has a nontrivial nullspace consisting of the corresponding eigenvectors. The eigenvalues can be determined by using the connection between the nullspace and the determinant indicated by Corollary 1.1 in Chapter 1: $N(A - \lambda)$ is nontrivial if and only if

$$\det(A - \lambda I) = 0.$$

When an eigenvalue has been found, the related eigenvector(s) are computed by finding a basis for $N(A - \lambda I)$. The equation (40) defines the eigenvalues $\lambda \in \mathbb{C}$ uniquely, but there exists infinitely many eigenvectors for each eigenvalue: the eigenvectors corresponding to an eigenvalue λ form a subspace of \mathbb{C}^n called the eigenspace,

$$E_\lambda := N(A - \lambda I).$$

The dimension of E_λ can be any integer between 1 and n and it is called the *geometric multiplicity* $\mu_G(\lambda)$ of the eigenvalue λ .

EXAMPLE 1.1. *Let*

$$A = \begin{bmatrix} 3 & -1 \\ 1 & 1 \end{bmatrix}.$$

The eigenvalues of A are the solutions to the equation

$$\det(A - \lambda I) = \det \begin{bmatrix} 3 - \lambda & -1 \\ 1 & 1 - \lambda \end{bmatrix} = (3 - \lambda)(1 - \lambda) + 1 = \lambda^2 - 4\lambda + 4 = (\lambda - 2)^2 = 0.$$

Hence, there is only one eigenvalue $\lambda = 2$ that is a double root of the characteristic polynomial. The eigenspace corresponding to the eigenvalue $\lambda = 2$ is $E_2 = N(A - 2I)$. Because

$$A - 2I = \begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix},$$

it is easy to deduce that $E_2 = \text{span}\{[1, 1]^T\}$ is a one-dimensional subspace of \mathbb{R}^2 .

DEFINITION 1.1. *The characteristic polynomial of a matrix $A \in \mathbb{C}^{n \times n}$ is defined as*

$$p_A(\lambda) = \det(A - \lambda I),$$

i.e., as the polynomial whose roots define the eigenvalues.

Recall the subdeterminant rule that can be employed in analyzing determinants or computing them by hand.¹ The rule can be applied either column-wise,

$$\det A = (-1)^{j-1} \sum_{i=1}^n (-1)^{i-1} a_{ij} \det [A]_{i,j},$$

or row-wise,

$$\det A = (-1)^{j-1} \sum_{i=1}^n (-1)^{i-1} a_{ji} \det [A]_{j,i}.$$

In the above formulas, j is the index of the column or the row with respect to which the determinant is expanded. Moreover, $[A]_{i,j} \in \mathbb{C}^{(n-1) \times (n-1)}$ is the square matrix that is obtained by removing the i th row and the j th column from the matrix A . A couple of examples are in order.

EXAMPLE 1.2. *Let*

$$A = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}$$

Then

$$[A]_{1,1} = \begin{bmatrix} a_{22} & a_{23} \\ a_{32} & a_{33} \end{bmatrix}, \quad [A]_{1,2} = \begin{bmatrix} a_{21} & a_{23} \\ a_{31} & a_{33} \end{bmatrix}, \quad [A]_{1,3} = \begin{bmatrix} a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix}.$$

¹Mathematical software such as MATLAB do *not* numerically evaluate determinants using the subdeterminant rule. Typically, they first form a suitable factorization of the examined matrix involving only triangular matrices, e.g., the LU decomposition. Then two properties of the determinant are used: (i) The determinant of a product is the product of the determinants and (ii) the determinant of a triangular matrix is the product of its diagonal elements.

EXAMPLE 1.3. For

$$A = \begin{bmatrix} 1 & 2 & 0 & 0 \\ 3 & 4 & 0 & 5 \\ 0 & 6 & 7 & 0 \\ 8 & 0 & 0 & 10 \end{bmatrix},$$

the subdeterminant rule with respect to the first row gives

$$\det A = \det \begin{bmatrix} 4 & 0 & 5 \\ 6 & 7 & 0 \\ 0 & 0 & 10 \end{bmatrix} - 2 \det \begin{bmatrix} 3 & 0 & 5 \\ 0 & 7 & 0 \\ 8 & 0 & 10 \end{bmatrix}.$$

Applying the subdeterminant rule next to $[A]_{1,1}$, we obtain

$$\det \begin{bmatrix} 4 & 0 & 5 \\ 6 & 7 & 0 \\ 0 & 0 & 10 \end{bmatrix} = 4 \det \begin{bmatrix} 7 & 0 \\ 0 & 10 \end{bmatrix} + 5 \det \begin{bmatrix} 6 & 7 \\ 0 & 0 \end{bmatrix} = 280.$$

The determinant of $[A]_{1,2}$ is computed recursively according to the same procedure. This gives $\det [A]_{1,2} = -70$ and altogether $\det A = 420$.

With the help of the subdeterminant rule, it is straightforward to prove that p_A is always a polynomial of degree n .

LEMMA 1.1. If $A \in \mathbb{C}^{n \times n}$, then $p_A(\lambda) = \det(A - \lambda I)$ is a polynomial of degree n . Moreover, the coefficient of the n th order term is $(-1)^n$.

PROOF. The proof is based on induction with respect to the dimension of a square matrix.

Base case: For a general $A \in \mathbb{C}^{2 \times 2}$,

$$\det(A - \lambda I) = \det \begin{bmatrix} a_{11} - \lambda & a_{12} \\ a_{21} & a_{22} - \lambda \end{bmatrix} = \lambda^2 - (a_{11} + a_{22})\lambda + (a_{11}a_{22} - a_{12}a_{21})$$

is clearly a second order polynomial and the coefficient of the term λ^2 is $1 = (-1)^2$.

Induction assumption: For any $B \in \mathbb{C}^{k \times k}$, $\det(B - \lambda I)$ is a polynomial of degree k and the coefficient of the k th order term is $(-1)^k$.

Induction step: Let $A \in \mathbb{C}^{(k+1) \times (k+1)}$ be arbitrary. The subdeterminant rule yields

$$(41) \quad \det(A - \lambda I) = (a_{11} - \lambda) \det([A - \lambda I]_{1,1}) + \sum_{i=2}^k (-1)^{i-1} a_{i1} \det([A - \lambda I]_{i,1}).$$

By the induction assumption, each subdeterminant appearing in (41) is a polynomial of degree k having $(-1)^k$ as the coefficient of the k th order term. Hence, the sum term on the right-hand side of (41) amounts to a polynomial of degree k in λ , and the first term on the right-hand side of (41) is a polynomial of degree $k+1$ in λ having $(-1)^{k+1}$ as the coefficient of the $(k+1)$ th order term. Thus, $\det(A - \lambda I)$ is altogether a polynomial of degree $k+1$ having $(-1)^{k+1}$ as the coefficient of the $(k+1)$ th order term, which completes the proof. \square

As the eigenvalues of a matrix A are the roots of its characteristic polynomial p_A , the eigenvalues can be studied by investigating the characteristic polynomial.

LEMMA 1.2. A matrix $A \in \mathbb{C}^{n \times n}$ has at most n distinct eigenvalues.

PROOF. The eigenvalues of $A \in \mathbb{C}^{n \times n}$ are the roots of the characteristic polynomial p_A . Because the polynomial p_A is exactly of the order n , it has at most n distinct roots (and altogether n complex roots if counted according to their multiplicity). \square

Eigenvalues can be multiple roots of the characteristic polynomial, as in Example 1.1. Assume $A \in \mathbb{C}^{n \times n}$ has the distinct eigenvalues $\lambda_1, \dots, \lambda_k$, where $1 \leq k \leq n$. By virtue of the fundamental theorem of algebra, the characteristic polynomial can be factored as

$$(42) \quad p_A(\lambda) = (-1)^n \prod_{i=1}^k (\lambda - \lambda_i)^{\mu_A(\lambda_i)},$$

where $\sum_{i=1}^k \mu_A(\lambda_i) = n$ since $p_A(\lambda)$ is of degree n . The positive integer power $\mu_A(\lambda_i)$ is the *algebraic multiplicity* of the eigenvalue λ_i , i.e., its multiplicity as a root of the characteristic polynomial. Recall that the dimension of the eigenspace $E_{\lambda_i} = N(A - \lambda_i I)$ was dubbed the *geometric multiplicity* of the eigenvalue λ_i and denoted by $\mu_G(\lambda_i)$. It can be shown that the geometric multiplicity $\mu_G(\lambda_i)$ is at most the same as the algebraic multiplicity $\mu_A(\lambda_i)$ for all $i = 1, \dots, k$ (the proof is omitted). However, if $\mu_A(\lambda_i) = 1$, then also $\mu_G(\lambda_i) = \mu_A(\lambda_i) = 1$ because each eigenvalue has at least one-dimensional eigenspace.

The eigenvectors related to different eigenvalues are linearly independent.

THEOREM 1.1. *Any eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_k \in \mathbb{C}^n$ corresponding to the distinct eigenvalues $\lambda_1, \dots, \lambda_k \in \mathbb{C}$ of $A \in \mathbb{C}^{n \times n}$ are linearly independent.*

PROOF. The proof is based on induction with respect to the number of considered eigenvalues.

Base case (i.e., any pair of eigenvectors are linearly independent): Without loss of generality, we may only consider the first two eigenvalues λ_1 and λ_2 ; the proof for any other pair would be analogous. Moreover, we may assume that $\lambda_1 \neq 0$ because one of the two eigenvalues must be nonzero, and we can rename the eigenvalues if necessary.

We argue by contradiction: Assume that the corresponding eigenvectors \mathbf{v}_1 and \mathbf{v}_2 are linearly dependent, that is, there exists $0 \neq \alpha \in \mathbb{C}$ such that

$$\mathbf{v}_1 = \alpha \mathbf{v}_2.$$

Multiplying this equation with A and using the definition of eigenvectors gives

$$\lambda_1 \mathbf{v}_1 = \alpha \lambda_2 \mathbf{v}_2, \quad \text{i.e.,} \quad \mathbf{v}_1 = \frac{\lambda_2}{\lambda_1} \alpha \mathbf{v}_2.$$

Subtracting the above identities leads to

$$\left(1 - \frac{\lambda_2}{\lambda_1}\right) \alpha \mathbf{v}_2 = 0.$$

As $\alpha \neq 0$ by assumption and $\mathbf{v}_2 \neq 0$ by definition, it must hold that $\lambda_2/\lambda_1 = 1$, which is a contradiction because λ_1 and λ_2 are distinct. Hence, \mathbf{v}_1 and \mathbf{v}_2 are linearly independent.

Induction assumption: Assume that any j vectors in the set $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ are linearly independent ($1 \leq j \leq k-1$).

Induction step: Consider any $j+1$ vectors in $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$. After renumbering, we may assume that they are $\mathbf{v}_1, \dots, \mathbf{v}_{j+1}$.

Let $\boldsymbol{\alpha} \in \mathbb{C}^{j+1}$ be such that

$$(43) \quad \sum_{i=1}^{j+1} \alpha_i \mathbf{v}_i = 0.$$

Multiplying (43) by A gives

$$\sum_{i=1}^{j+1} \alpha_i \lambda_i \mathbf{v}_i = 0.$$

Without loss of generality, we may assume that $\lambda_1 \neq 0$; at most one of the distinct eigenvalues is zero and we may rename the eigenvalues and eigenvectors if need be. Hence,

$$\alpha_1 \mathbf{v}_1 + \sum_{i=2}^{j+1} \alpha_i \frac{\lambda_i}{\lambda_1} \mathbf{v}_i = 0.$$

Subtracting this from (43), we get

$$\sum_{i=2}^{j+1} \alpha_i \left(1 - \frac{\lambda_i}{\lambda_1}\right) \mathbf{v}_i = 0.$$

Since the j eigenvectors $\mathbf{v}_2, \dots, \mathbf{v}_{j+1}$ are linearly independent by the induction assumption and $\lambda_i/\lambda_1 \neq 1$ since the considered eigenvalues are distinct, it must hold that $\alpha_i = 0$ for all $i = 2, \dots, j+1$. Using this information, say, in (43), it follows that also $\alpha_1 = 0$, i.e., $\boldsymbol{\alpha} = 0$, and thus $\mathbf{v}_1, \dots, \mathbf{v}_{j+1}$ are linearly independent. \square

If the geometric and the algebraic multiplicity for each eigenvalue of $A \in \mathbb{R}^{n \times n}$ are the same, Theorem 1.1 implies that there exists n linearly independent eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ for A : If $\mu_G(\lambda) = \mu_A(\lambda)$ for all eigenvalues, one can find $\mu_A(\lambda)$ linearly independent eigenvectors from each eigenspace E_λ of A because $\mu_G(\lambda) = \mu_A(\lambda)$ is its dimension. Since Theorem 1.1 guarantees that eigenvectors corresponding to distinct eigenvalues are automatically linearly independent and the algebraic multiplicities sum to n (cf. (42)), one can indeed altogether find n linearly independent eigenvectors. On the other, if $\mu_G(\lambda) < \mu_A(\lambda)$ for any eigenvalue λ , there does not exist a *full set* of linearly independent eigenvectors. On the positive side, if A has n distinct eigenvalues, then $\mu_G(\lambda) = \mu_A(\lambda) = 1$ for all eigenvalues and there are n linearly independent eigenvectors.

Assume then that $A \in \mathbb{R}^{n \times n}$ has n linearly independent eigenvectors and stack them as the columns of a matrix:

$$V := [\mathbf{v}_1, \dots, \mathbf{v}_n] \in \mathbb{C}^{n \times n}.$$

The matrix product AV gives

$$(44) \quad AV = [A\mathbf{v}_1, \dots, A\mathbf{v}_n] = [\lambda_1 \mathbf{v}_1, \dots, \lambda_n \mathbf{v}_n],$$

where the eigenvalues are repeated according to their algebraic/geometric multiplicity. If $\Lambda \in \mathbb{C}^{n \times n}$ is the diagonal matrix carrying the eigenvalues of A in the corresponding order, i.e.,

$$\Lambda = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_n \end{bmatrix},$$

then (44) can be rewritten as

$$AV = V\Lambda.$$

As V has linearly independent columns, it is invertible (see (2) in Corollary 1.1) and we obtain

$$(45) \quad A = V\Lambda V^{-1}.$$

This is an eigendecomposition for the matrix A . If the algebraic and geometric multiplicities are not the same for all eigenvalues, such a decomposition does not exist. If a decomposition of the type (45) exists, A is called *diagonalizable*. Finally, note that a decomposition of the form (45) is never unique: one can change the order of the eigenvalues, the length of the eigenvectors, or even pick different bases for the eigenspaces.

REMARK 1.1. *The decomposition (45) has the following intuitive interpretation: The linear map represented by $A \in \mathbb{R}^{n \times n}$ in the standard/Cartesian basis has a diagonal representation $\Lambda \in \mathbb{R}^{n \times n}$ in the eigenbasis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$. To understand this, let $\boldsymbol{\alpha} \in \mathbb{R}^n$ be the coordinates of an arbitrary $\mathbf{x} \in \mathbb{R}^n$ in the basis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$, that is,*

$$\mathbf{x} = V\boldsymbol{\alpha} \iff \boldsymbol{\alpha} = V^{-1}\mathbf{x}.$$

Hence,

$$A\mathbf{x} = V\Lambda V^{-1}\mathbf{x} = V(\Lambda\boldsymbol{\alpha}),$$

which just means that the coordinates of the image vector $A\mathbf{x} \in \mathbb{R}^n$ in the eigenbasis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ are $\Lambda\boldsymbol{\alpha} \in \mathbb{R}^n$. In other words, the action of the considered linear map on the coordinates of \mathbf{x} in the basis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ is realized by a diagonal matrix Λ , i.e., the coordinates change according to the rule $\boldsymbol{\alpha} \mapsto \Lambda\boldsymbol{\alpha}$. This is why a matrix with a representation of the form (45) is called diagonalizable.

2. Hermitian matrices

Hermitian matrices are encountered in numerous applications. They have also many nice properties that make them ‘easily approachable’. We start our treatment with a couple of definitions that will be put to use immediately afterwards.

DEFINITION 2.1. A matrix $A \in \mathbb{C}^{n \times n}$ is Hermitian if $A^* = A$. In particular, a real matrix is Hermitian if and only if it is symmetric, i.e. $A^T = A$.

DEFINITION 2.2. A matrix $U \in \mathbb{C}^{n \times n}$ is unitary if $U^* = U^{-1}$. In particular, a real matrix is unitary if and only if it is orthogonal, i.e. $U^T = U^{-1}$.

By definition, the diagonal elements of a Hermitian matrix must be real. Observe also that the columns (or the rows) of a unitary matrix form an orthonormal basis for \mathbb{C}^n , that is, they are of unit Euclidean length and mutually orthogonal with respect to the Euclidean inner product. Indeed, if $U \in \mathbb{C}^{n \times n}$ is unitary and we write $U = [\mathbf{u}_1, \dots, \mathbf{u}_n]$, then

$$(46) \quad U^*U = \begin{bmatrix} \mathbf{u}_1^*\mathbf{u}_1 & \mathbf{u}_1^*\mathbf{u}_2 & \cdots & \mathbf{u}_1^*\mathbf{u}_n \\ \mathbf{u}_2^*\mathbf{u}_1 & \mathbf{u}_2^*\mathbf{u}_2 & \cdots & \mathbf{u}_2^*\mathbf{u}_n \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{u}_n^*\mathbf{u}_1 & \mathbf{u}_n^*\mathbf{u}_2 & \cdots & \mathbf{u}_n^*\mathbf{u}_n \end{bmatrix} = I,$$

meaning that the off-diagonal inner products vanish and the squared norms on the diagonal equal one.

Our first aim is to show that any Hermitian matrix $A \in \mathbb{C}^{n \times n}$ is unitarily diagonalizable, that is, there exists a (nonunique) decomposition

$$(47) \quad A = Q\Lambda Q^*,$$

where $Q \in \mathbb{C}^{n \times n}$ is unitary and $\Lambda \in \mathbb{R}^{n \times n}$ a diagonal matrix with *real* entries. This property is closely related to the fact that the eigenvalues of a Hermitian matrix are real and eigenvectors corresponding to different eigenvalues are orthogonal.

LEMMA 2.1. The eigenvalues of a Hermitian $A \in \mathbb{C}^{n \times n}$ are real.

PROOF. Let λ be an arbitrary eigenvalue of A and let \mathbf{v} be a corresponding eigenvector. In particular,

$$\mathbf{v}^*A\mathbf{v} = \mathbf{v}^*(\lambda\mathbf{v}) = \lambda\|\mathbf{v}\|_2^2.$$

On the other hand, as $A^* = A$,

$$\mathbf{v}^*A\mathbf{v} = \mathbf{v}^*A^*\mathbf{v} = (A\mathbf{v})^*\mathbf{v} = (\lambda\mathbf{v})^*\mathbf{v} = \bar{\lambda}\|\mathbf{v}\|_2^2.$$

Subtracting these identities leads to

$$(\lambda - \bar{\lambda})\|\mathbf{v}\|_2^2 = 2\operatorname{Im}\lambda\|\mathbf{v}\|_2^2 = 0.$$

Since $\mathbf{v} \neq 0$ by definition, it must hold that $\operatorname{Im}\lambda = 0$, i.e., λ is real. \square

LEMMA 2.2. Let $\mathbf{q}_1, \mathbf{q}_2 \in \mathbb{C}^n$ be eigenvectors corresponding to distinct eigenvalues $\lambda_1, \lambda_2 \in \mathbb{R}$ of a Hermitian matrix $A \in \mathbb{C}^{n \times n}$, respectively. Then \mathbf{q}_1 and \mathbf{q}_2 are orthogonal in the sense that $\mathbf{q}_1^*\mathbf{q}_2 = 0$.

PROOF. As the eigenvalues of a Hermitian matrix are real, it holds that

$$\lambda_1 \mathbf{q}_1^* \mathbf{q}_2 = (\lambda_1 \mathbf{q}_1)^* \mathbf{q}_2 = (A\mathbf{q}_1)^* \mathbf{q}_2 = \mathbf{q}_1^* A^* \mathbf{q}_2 = \mathbf{q}_1^* A \mathbf{q}_2.$$

On the other hand,

$$\lambda_2 \mathbf{q}_1^* \mathbf{q}_2 = \mathbf{q}_1^* (\lambda_2 \mathbf{q}_2) = \mathbf{q}_1^* A \mathbf{q}_2.$$

Subtracting these equalities gives

$$(\lambda_2 - \lambda_1) \mathbf{q}_1^* \mathbf{q}_2 = 0,$$

which proves the claim since $\lambda_2 \neq \lambda_1$ by assumption. \square

Now we have the necessary tools to establish the unitary decomposition (47).

THEOREM 2.1. *For any Hermitian matrix $A \in \mathbb{C}^{n \times n}$, there exists a unitary $Q \in \mathbb{C}^{n \times n}$ and a diagonal $\Lambda \in \mathbb{R}^{n \times n}$ such that the decomposition (47) is valid.*

PROOF. The proof is based on induction with respect to the dimension n of a Hermitian matrix.

Base case ($n = 1$): A Hermitian 1×1 matrix is a single real number, say, $a \in \mathbb{R}$ that can be decomposed as $a = 1 a 1$, i.e., $\Lambda = a$ itself and $Q = 1$.

Induction assumption: The claim is true for $n = k$.

Induction step: Let $A \in \mathbb{C}^{(k+1) \times (k+1)}$ be Hermitian and note that it has at least one eigenvalue $\lambda_1 \in \mathbb{R}$ with a corresponding eigenvector \mathbf{q}_1 satisfying $\|\mathbf{q}_1\|_2 = 1$ (take any eigenvector for λ and divide it by its 2-norm). Introduce k auxiliary unit vectors $\mathbf{q}_2, \dots, \mathbf{q}_{k+1} \in \mathbb{C}^{k+1}$ that are mutually orthogonal as well as orthogonal to \mathbf{q}_1 ; such can be found, e.g., via the Gram–Schmidt orthogonalization process that will be discussed in Chapter 3. We denote

$$\tilde{Q} = [\mathbf{q}_2, \dots, \mathbf{q}_{k+1}] \in \mathbb{C}^{(k+1) \times k} \quad \text{and} \quad Q = [\mathbf{q}_1, \tilde{Q}] \in \mathbb{C}^{(k+1) \times (k+1)},$$

and note that Q is unitary (see (46)).

Consider

$$\begin{aligned} Q^* A Q &= [\mathbf{q}_1, \tilde{Q}]^* A [\mathbf{q}_1, \tilde{Q}] = \begin{bmatrix} \mathbf{q}_1^* \\ \tilde{Q}^* \end{bmatrix} [A\mathbf{q}_1, A\tilde{Q}] \\ &= \begin{bmatrix} \mathbf{q}_1^* A \mathbf{q}_1 & \mathbf{q}_1^* A \tilde{Q} \\ \tilde{Q}^* A \mathbf{q}_1 & \tilde{Q}^* A \tilde{Q} \end{bmatrix} = \begin{bmatrix} \lambda_1 \mathbf{q}_1^* \mathbf{q}_1 & (\lambda_1 \mathbf{q}_1)^* \tilde{Q} \\ \lambda_1 \tilde{Q}^* \mathbf{q}_1 & \tilde{Q}^* A \tilde{Q} \end{bmatrix} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \tilde{Q}^* A \tilde{Q} \end{bmatrix}, \end{aligned}$$

where the last step follows from the fact that \mathbf{q}_1 is of unit length and orthogonal to the columns of \tilde{Q} . Since the matrix $\tilde{Q}^* A \tilde{Q} \in \mathbb{C}^{k \times k}$ is obviously Hermitian, by the induction assumption there exists a unitary $Q_1 \in \mathbb{C}^{k \times k}$ and a diagonal matrix $\Lambda_1 \in \mathbb{R}^{k \times k}$ such that

$$\tilde{Q}^* A \tilde{Q} = Q_1^* \Lambda_1 Q_1.$$

In consequence,

$$Q^* A Q = \begin{bmatrix} \lambda_1 & 0 \\ 0 & Q_1^* \Lambda_1 Q_1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & Q_1^* \end{bmatrix} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \Lambda_1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & Q_1 \end{bmatrix}.$$

Due to the unitarity of Q , it follows that

$$(48) \quad A = Q \begin{bmatrix} 1 & 0 \\ 0 & Q_1^* \end{bmatrix} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \Lambda_1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & Q_1 \end{bmatrix} Q^* = \left(Q \begin{bmatrix} 1 & 0 \\ 0 & Q_1^* \end{bmatrix} \right) \begin{bmatrix} \lambda_1 & 0 \\ 0 & \Lambda_1 \end{bmatrix} \left(Q \begin{bmatrix} 1 & 0 \\ 0 & Q_1^* \end{bmatrix} \right)^*$$

This completes the proof: (48) is a unitary decomposition of the form (47) because

$$\left(Q \begin{bmatrix} 1 & 0 \\ 0 & Q_1^* \end{bmatrix} \right)^* \left(Q \begin{bmatrix} 1 & 0 \\ 0 & Q_1^* \end{bmatrix} \right) = \begin{bmatrix} 1 & 0 \\ 0 & Q_1 \end{bmatrix} Q^* Q \begin{bmatrix} 1 & 0 \\ 0 & Q_1^* \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & Q_1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & Q_1^* \end{bmatrix} = I,$$

where we used the unitarity of both Q and Q_1 . \square

The decomposition (47) can alternatively be given as

$$AQ = Q\Lambda \quad \iff \quad [A\mathbf{q}_1, \dots, A\mathbf{q}_n] = [\lambda_1\mathbf{q}_1, \dots, \lambda_n\mathbf{q}_n],$$

which demonstrates that the columns of the unitary matrix Q are orthonormal eigenvectors for A and the diagonal elements of Λ , i.e. $\lambda_1, \dots, \lambda_n$, are the corresponding eigenvalues (repeated according to their algebraic/geometric multiplicity). In particular, the orthogonal eigenvectors $\mathbf{q}_1, \dots, \mathbf{q}_n$ form an eigenbasis for the whole of $\mathbb{C}^{n \times n}$ (cf. Lemma 2.1).

As we already know, the eigenvalues of a matrix A are the roots of the corresponding characteristic polynomial. For a Hermitian matrix, the eigenvalues can also be characterized by an optimization problem related to the *Rayleigh quotient*.

DEFINITION 2.3. For a Hermitian matrix $A \in \mathbb{C}^{n \times n}$, the corresponding Rayleigh quotient is defined as

$$R(A, \mathbf{x}) = \frac{\mathbf{x}^* A \mathbf{x}}{\mathbf{x}^* \mathbf{x}} = \frac{\mathbf{x}^* A \mathbf{x}}{\|\mathbf{x}\|_2^2} \in \mathbb{R}.$$

(The Rayleigh quotient is real-valued because the numerator $\mathbf{x}^* A \mathbf{x} = \mathbf{x}^* A^* \mathbf{x} = (\mathbf{x}^* A \mathbf{x})^* = \overline{\mathbf{x}^* A \mathbf{x}}$ is real as it equals its complex conjugate.)

Let the (real) eigenvalues of a Hermitian $A \in \mathbb{C}^{n \times n}$ be arranged in increasing order, i.e.,

$$\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n,$$

where the eigenvalues are repeated according to their algebraic (or geometric) multiplicity. It is easy to see that the minimal and maximal values of the Rayleigh quotient are the smallest and largest eigenvalues of A , respectively.

THEOREM 2.2. For a Hermitian matrix $A \in \mathbb{C}^{n \times n}$, it holds that

$$\lambda_1 = \min_{0 \neq \mathbf{x} \in \mathbb{C}^n} R(A, \mathbf{x}) \quad \text{and} \quad \lambda_n = \max_{0 \neq \mathbf{x} \in \mathbb{C}^n} R(A, \mathbf{x}),$$

where λ_1 is the smallest eigenvalue of A and λ_n the largest.

PROOF. Expand an arbitrary $\mathbf{x} \in \mathbb{C}^n$ in an orthonormal eigenbasis of A ,

$$\mathbf{x} = \sum_{i=1}^n \alpha_i \mathbf{q}_i, \quad \boldsymbol{\alpha} \in \mathbb{C}^n.$$

It follows that (homework),

$$\mathbf{x}^* A \mathbf{x} = \sum_{i=1}^n |\alpha_i|^2 \lambda_i \quad \text{and} \quad \mathbf{x}^* \mathbf{x} = \sum_{i=1}^n |\alpha_i|^2,$$

and the Rayleigh quotient can thus be estimated from below as

$$R(A, \mathbf{x}) = \frac{\mathbf{x}^* A \mathbf{x}}{\mathbf{x}^* \mathbf{x}} = \frac{\sum_{i=1}^n |\alpha_i|^2 \lambda_i}{\sum_{i=1}^n |\alpha_i|^2} \geq \frac{\lambda_1 \sum_{i=1}^n |\alpha_i|^2}{\sum_{i=1}^n |\alpha_i|^2} = \lambda_1.$$

On the other hand, choosing $\alpha_1 = 1$ and $\alpha_i = 0$ for $i = 2, \dots, n$ directly gives

$$R(A, \mathbf{x}) = \frac{\sum_{i=1}^n |\alpha_i|^2 \lambda_i}{\sum_{i=1}^n |\alpha_i|^2} = \lambda_1.$$

This proves the first part of the claim. The second part follows via an analogous argument. \square

The other eigenvalues of a Hermitian matrix A can be characterized via

$$\lambda_k = \min_{\dim U = k} \max_{0 \neq \mathbf{x} \in U} \frac{\mathbf{x}^* A \mathbf{x}}{\mathbf{x}^* \mathbf{x}},$$

where the outer minimization is over subspaces of dimension k . This is known as the Courant–Fisher min-max theorem.

3. Similarity transformations

If a matrix $A \in \mathbb{C}^{n \times n}$ does not have n linearly independent eigenvectors, it is not diagonalizable. However, one can still look for some other decomposition of the form

$$(49) \quad A = SBS^{-1},$$

where B is a simpler matrix than the original A in some suitable sense. For example, if B^n is easier to evaluate than A^n , then the decomposition (49) is useful when defining matrix valued functions via power series expansion because

$$A^n = SB^nS^{-1}.$$

A decomposition of the form (49) merits a separate definition.

DEFINITION 3.1. *The matrices $A \in \mathbb{C}^{n \times n}$ and $B \in \mathbb{C}^{n \times n}$ are called similar if there exists an invertible $S \in \mathbb{C}^{n \times n}$ such that (49) is valid.*

In the spirit of Remark 1.1, two matrices are similar if they represent the same linear mapping in different basis. As an example, if A represents a certain linear map in the Cartesian basis, then B of (49) gives the same linear map in the basis defined by the columns of S . In particular, a diagonalizable matrix is represented by a diagonal matrix in an eigenbasis — and the eigenbasis can be chosen to be orthonormal if the matrix in question is Hermitian.

LEMMA 3.1. *If $A, B \in \mathbb{C}^{n \times n}$ are similar, then the associated characteristic polynomials are the same, i.e., $p_A = p_B$.*

PROOF. Assume that $A = SBS^{-1}$. Using basic properties of the determinant, we deduce

$$\begin{aligned} p_A(\lambda) &= \det(A - \lambda I) = \det(SBS^{-1} - \lambda SS^{-1}) = \det(S(B - \lambda I)S^{-1}) \\ &= \det(S) \det(B - \lambda I) \det(S^{-1}) = \det(S) \det(B - \lambda I) \det(S)^{-1} \\ &= \det(B - \lambda I) = p_B(\lambda), \end{aligned}$$

which proves the claim. □

In what follows, we will introduce two decompositions of the form (49) that are valid for any square matrix $A \in \mathbb{C}^{n \times n}$, independently of whether A has a full set of linearly independent eigenvectors or not. We will start with the Schur decomposition and then briefly discuss the Jordan normal form. Both of these are theoretical tools that are rarely numerically computed in practice. We will omit the existence proof for the Jordan normal form to keep the presentation compact.

The Schur decomposition states that any square matrix is unitarily similar to an upper (or lower) triangular matrix. In other words, for each square matrix there exists an orthonormal basis in which the corresponding linear mapping is represented by an upper triangular matrix (cf. Remark 1.1).

THEOREM 3.1. *For any $A \in \mathbb{C}^{n \times n}$, there exist a unitary $Q \in \mathbb{C}^{n \times n}$ and an upper triangular $T \in \mathbb{C}^{n \times n}$ such that*

$$(50) \quad A = QTQ^*.$$

PROOF. The proof is based on induction with respect to the dimension n .

Base case ($n = 1$): A general 1×1 matrix is a single complex number, say, $a \in \mathbb{C}$ that can be decomposed as $a = 1 a 1$, i.e., $T = a$ itself and $Q = 1$.

Induction assumption: The claim is true for $n = k$.

Induction step: We begin as in the proof of Theorem 2.1: Let $A \in \mathbb{C}^{(k+1) \times (k+1)}$ and note that it has at least one eigenvalue $\lambda_1 \in \mathbb{C}$ with a corresponding eigenvector \mathbf{q}_1 satisfying $\|\mathbf{q}_1\|_2 = 1$.

Introduce k auxiliary unit vectors $\mathbf{q}_2, \dots, \mathbf{q}_{k+1} \in \mathbb{C}^{k+1}$ that are mutually orthogonal as well as orthogonal to \mathbf{q}_1 . We denote

$$\tilde{Q} = [\mathbf{q}_2, \dots, \mathbf{q}_{k+1}] \in \mathbb{C}^{(k+1) \times k} \quad \text{and} \quad Q = [\mathbf{q}_1, \tilde{Q}] \in \mathbb{C}^{(k+1) \times (k+1)},$$

and note that Q is unitary. Exactly as in the proof of Theorem 2.1, we get

$$Q^* A Q = \begin{bmatrix} \mathbf{q}_1^* A \mathbf{q}_1 & \mathbf{q}_1^* A \tilde{Q} \\ \tilde{Q}^* A \mathbf{q}_1 & \tilde{Q}^* A \tilde{Q} \end{bmatrix} = \begin{bmatrix} \lambda_1 \mathbf{q}_1^* \mathbf{q}_1 & \mathbf{q}_1^* A \tilde{Q} \\ \lambda_1 \tilde{Q}^* \mathbf{q}_1 & \tilde{Q}^* A \tilde{Q} \end{bmatrix} = \begin{bmatrix} \lambda_1 & \mathbf{q}_1^* A \tilde{Q} \\ 0 & \tilde{Q}^* A \tilde{Q} \end{bmatrix},$$

where we used the fact that \mathbf{q}_1 is a unit vector and orthogonal to the columns of \tilde{Q} by construction (, but we could not get rid of the top right block as A is not assumed to be Hermitian).

By the induction assumption, for $\tilde{Q}^* A \tilde{Q} \in \mathbb{C}^{k \times k}$ there exists an upper triangular $T_1 \in \mathbb{C}^{k \times k}$ and a unitary $Q_1 \in \mathbb{C}^{k \times k}$ such that

$$\tilde{Q}^* A \tilde{Q} = Q_1 T_1 Q_1^*.$$

Hence,

$$Q^* A Q = \begin{bmatrix} \lambda_1 & \mathbf{q}_1^* A \tilde{Q} \\ 0 & Q_1 T_1 Q_1^* \end{bmatrix} = \begin{bmatrix} \lambda_1 & \mathbf{q}_1^* A \tilde{Q} (Q_1 Q_1^*) \\ 0 & Q_1 T_1 Q_1^* \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & Q_1 \end{bmatrix} \begin{bmatrix} \lambda_1 & \mathbf{q}_1^* A \tilde{Q} Q_1 \\ 0 & T_1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & Q_1^* \end{bmatrix}$$

and by the unitarity of Q ,

(51)

$$A = Q \begin{bmatrix} 1 & 0 \\ 0 & Q_1 \end{bmatrix} \begin{bmatrix} \lambda_1 & \mathbf{q}_1^* A \tilde{Q} Q_1 \\ 0 & T_1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & Q_1^* \end{bmatrix} Q^* = \left(Q \begin{bmatrix} 1 & 0 \\ 0 & Q_1 \end{bmatrix} \right) \begin{bmatrix} \lambda_1 & \mathbf{q}_1^* A \tilde{Q} Q_1 \\ 0 & T_1 \end{bmatrix} \left(Q \begin{bmatrix} 1 & 0 \\ 0 & Q_1 \end{bmatrix} \right)^*.$$

It is easy check that (51) is a decomposition of the required form (50). Indeed, as $T_1 \in \mathbb{C}^{k \times k}$ is upper triangular, so is the midmost matrix in (51). Moreover,

$$\left(Q \begin{bmatrix} 1 & 0 \\ 0 & Q_1 \end{bmatrix} \right)^* \left(Q \begin{bmatrix} 1 & 0 \\ 0 & Q_1 \end{bmatrix} \right) = I,$$

which corresponds exactly to the last formula in the proof of Theorem 2.1. \square

Theorem 3.1 states that any square matrix A is similar to an upper triangular T . Since $p_A(\lambda) = p_T(\lambda) = \det(T - \lambda I)$ by Lemma 3.1 and the determinant of the upper triangular matrix $T - \lambda I$ is the product of its diagonal elements, the eigenvalues of A are on the diagonal of T repeated according to their algebraic multiplicities (cf. (42)).

To complete this section, we introduce the *Jordan normal form* that indicates any square matrix $A \in \mathbb{C}^{n \times n}$ is similar to an ‘almost diagonal’ matrix, i.e., to a *Jordan matrix* that is an upper triangular, block diagonal matrix of the form

$$(52) \quad J = \begin{bmatrix} J_{n_1}(\lambda_1) & & & \\ & J_{n_2}(\lambda_2) & & \\ & & \ddots & \\ & & & J_{n_l}(\lambda_l) \end{bmatrix} \in \mathbb{C}^{n \times n}.$$

Here $J_{n_j}(\lambda_j) \in \mathbb{C}^{n_j \times n_j}$ is a *Jordan block* related to an eigenvalue λ_j ,

$$J_{n_j}(\lambda_j) = \begin{bmatrix} \lambda_j & 1 & & \\ & \lambda_j & 1 & \\ & & \ddots & \ddots \\ & & & \lambda_j & 1 \\ & & & & \lambda_j \end{bmatrix}.$$

Take note that $n_1 + n_2 + \dots + n_l = n$ and J has nonzero elements only on its main diagonal and on its first superdiagonal.

Each eigenvalue is repeated according to its geometric multiplicity in J of (52). On the other hand, by summing the sizes n_j of the Jordan blocks corresponding to a particular eigenvalue, one gets the algebraic multiplicity of that eigenvalue. In consequence, if the algebraic and the geometric multiplicity are the same for each eigenvalue, all Jordan blocks in (52) are of the size 1×1 and J is,

in fact, a diagonal matrix carrying the eigenvalues repeated according to their algebraic/geometric multiplicities. To be slightly more precise, the total number of nonzero elements, i.e. ones, on the first superdiagonal of J is

$$\sum_{j=1}^k (\mu_A(\lambda_j) - \mu_G(\lambda_j)),$$

where $\lambda_1, \dots, \lambda_k$ are the distinct eigenvalues of A .

The following, slightly vague theorem ends this section.

THEOREM 3.2. *Any $A \in \mathbb{C}^{n \times n}$ has a decomposition*

$$(53) \quad A = PJP^{-1}$$

where $J \in \mathbb{C}^{n \times n}$ is as in (52) and the invertible matrix $P \in \mathbb{C}^{n \times n}$ has eigenvectors and so-called generalized eigenvectors of A as its columns.

4. Matrix exponential

Similarity transformations can be used to calculate matrix valued functions such as the square root or the exponential of a matrix. In this section, we concentrate on the matrix exponential and its application to qualitative analysis of systems of ordinary differential equations.

For a scalar argument $t \in \mathbb{R}$, the exponential function can be defined, e.g., via its power series expansion:

$$(54) \quad e^t = \sum_{j=0}^{\infty} \frac{t^j}{j!}.$$

According to basic calculus courses, the series in (54) converges for any $t \in \mathbb{R}$ (and the convergence is, in fact, uniform on any bounded subset of \mathbb{R}). In particular, it can be shown that one is allowed to differentiate the series termwise:

$$\frac{d}{dt} e^t = \sum_{j=0}^{\infty} \frac{d}{dt} \frac{t^j}{j!} = \sum_{j=1}^{\infty} j \frac{t^{(j-1)}}{j!} = \sum_{j=1}^{\infty} \frac{t^{(j-1)}}{(j-1)!} = \sum_{j=0}^{\infty} \frac{t^j}{j} = e^t.$$

In exactly the same manner, one also obtains that

$$\frac{d}{dt} e^{at} = ae^{at}$$

for any constant $a \in \mathbb{C}$. As a consequence, the solution² to the initial value problem

$$(55) \quad x'(t) = ax(t) \quad \text{for } t > 0, \quad x(0) = x_0 \in \mathbb{C},$$

is obviously $x(t) = x_0 e^{at} = e^{at} x_0$. An interesting and practically relevant question is how this argumentation can be generalized if the scalar-valued function $x : \mathbb{R}_+ \rightarrow \mathbb{C}$ is replaced in (55) by a vector-valued version $\mathbf{x} : \mathbb{R}_+ \rightarrow \mathbb{C}^n$ and the scalar coefficient $a \in \mathbb{C}$ is replaced by a matrix $A \in \mathbb{C}^{n \times n}$.

To begin with, notice that we can *formally* define e^{tA} even if $A \in \mathbb{C}^{n \times n}$:

$$(56) \quad e^{tA} := \sum_{j=0}^{\infty} \frac{(tA)^j}{j!} = \sum_{j=0}^{\infty} \frac{t^j A^j}{j!},$$

where A^0 is defined to be the identity matrix $I \in \mathbb{R}^{n \times n}$. The right-hand side of (56) makes algebraic sense since A is a square matrix. In particular, it seems that e^{tA} is also a $n \times n$ matrix for any $t \in \mathbb{R}$. However, it is not quite obvious that the series on the right-hand side of (56) converges (in any reasonable sense), and so it is also not obvious if (56) can be taken as the definition of the *matrix exponential* function. To this end, let us define

$$S_N(t) := \sum_{j=0}^N \frac{t^j A^j}{j!}$$

²In fact, we do not prove on this course that (55) and its vectorial counterpart (58) are uniquely solvable — although they are.

to be the N th partial sum in the infinite series on the right-hand side of (56). For example on the course *Euclidean spaces*, it is proven that in any finite-dimensional normed vector space a sequence converges if and only if it is a Cauchy sequence. In our case, this means that $\lim_{N \rightarrow \infty} S_N(t)$ exists,³ i.e., (56) makes sense, if and only if

$$(57) \quad \|S_N(t) - S_M(t)\| \rightarrow 0 \quad \text{as } N, M \rightarrow \infty$$

for some matrix norm $\|\cdot\|$ on $\mathbb{C}^{n \times n}$.⁴

Let us then prove (57) for an arbitrary *operator norm*; recall that operator norms form a subclass of matrix norms. Without loss of generality we may assume that $N > M$ (the case $M = N$ is trivial and if $M > N$, we can swap the names of the indices). By repetitively applying the triangle inequality and other basic properties of operator norms, we obtain

$$\begin{aligned} \|S_N(t) - S_M(t)\| &= \left\| \sum_{j=M+1}^N \frac{t^j A^j}{j!} \right\| \leq \sum_{j=M+1}^N \left\| \frac{t^j A^j}{j!} \right\| \leq \sum_{j=M+1}^N \frac{(|t| \|A\|)^j}{j!} \\ &= \sum_{k=0}^{N-(M+1)} \frac{(|t| \|A\|)^{k+(M+1)}}{(k+(M+1))!} \leq \frac{(|t| \|A\|)^{M+1}}{(M+1)!} \sum_{k=0}^{N-(M+1)} \frac{(|t| \|A\|)^k}{k!} \\ &\leq \frac{(|t| \|A\|)^{M+1}}{(M+1)!} \sum_{k=0}^{\infty} \frac{(|t| \|A\|)^k}{k!} = \frac{(|t| \|A\|)^{M+1}}{(M+1)!} e^{|t| \|A\|}, \end{aligned}$$

which converges to zero as M (and N) tend to infinity since the factorial in the denominator grows considerably faster than the power in the numerator. Hence, S_N is a Cauchy sequence, and thus it converges for any (fixed) $t \in \mathbb{R}$ and $A \in \mathbb{C}^{n \times n}$. In particular, it makes sense to use (56) as the definition of the matrix exponential function.

Let us then consider the vectorial counterpart of (55), that is, a system of ordinary differential equations of the form

$$(58) \quad \mathbf{x}'(t) = A\mathbf{x}(t) \quad \text{for } t > 0, \quad \mathbf{x}(0) = \mathbf{x}_0,$$

where $A \in \mathbb{C}^{n \times n}$ is the coefficient matrix, $\mathbf{x} : \mathbb{R}_+ \rightarrow \mathbb{C}^n$ is the solution, and $\mathbf{x}_0 \in \mathbb{C}^n$ carries the initial values. Analogously to the case of a single differential equation in (55), the solution to this problem is

$$\mathbf{x}(t) = e^{tA} \mathbf{x}_0$$

as we will demonstrate next. First of all, the initial condition is satisfied:

$$e^{0A} \mathbf{x}_0 = I \mathbf{x}_0 = \mathbf{x}_0$$

because all but the first term in the series on the right-hand side of (56) vanish when $t = 0$.⁵ As in the scalar-valued case, it can moreover be argued that the series on the right-hand side of (56) can be differentiated term by term, which yields

$$\begin{aligned} \frac{d}{dt}(e^{tA} \mathbf{x}_0) &= \left(\sum_{j=0}^{\infty} \frac{d}{dt} \frac{t^j A^j}{j!} \right) \mathbf{x}_0 = \left(\sum_{j=1}^{\infty} j \frac{t^{j-1} A^j}{j!} \right) \mathbf{x}_0 = \left(\sum_{j=1}^{\infty} A \frac{t^{j-1} A^{j-1}}{(j-1)!} \right) \mathbf{x}_0 \\ &= \left(A \sum_{j=0}^{\infty} \frac{t^j A^j}{j!} \right) \mathbf{x}_0 = A(e^{tA} \mathbf{x}_0) \end{aligned}$$

as desired.

³Observe that the space of $n \times n$ matrices, i.e. $\mathbb{C}^{n \times n}$, is a vector space (consider scalar multiplication and addition of matrices) and it is also finite-dimensional (any n^2 matrices that have only one nonzero element at mutually different locations form a basis for $\mathbb{C}^{n \times n}$).

⁴In fact, it does not matter which matrix norm is used since all norms on a finite-dimensional vector space define the same topology (cf. the third exercise sheet).

⁵It is agreed that $(0A)^0 = I$, which can be reasoned by continuity: if t tends to zero from either side, then $(tA)^0 = t^0 A^0 = I$ converges to the identity matrix. (This same convention is actually used already in the scalar-valued definition (54): $e^0 = 0^0 := 1$.)

Similarity transformations can be utilized to write matrix exponentials in more accessible forms; when analyzing this, we drop the ‘time variable’ to simplify the notation and because the matrix exponential e^A also has other applications than systems of ordinary differential equations. For a diagonalizable matrix one simply exploits its similarity with a diagonal matrix, and for the other matrices the standard help is the Jordan normal form. Both of these approaches are theoretical tools that are usually not implemented numerically in practice; there are other numerical algorithms for computing the matrix exponential of a given square matrix.

Let us start with a diagonalizable $A \in \mathbb{C}^{n \times n}$ such that

$$A = X\Lambda X^{-1},$$

where $X \in \mathbb{C}^{n \times n}$ is invertible and the diagonal matrix $\Lambda \in \mathbb{C}^{n \times n}$ has the eigenvalues of A , i.e. $\lambda_1, \dots, \lambda_n$, on its diagonal (repeated according to their algebraic multiplicity). Clearly,

$$A^j = X\Lambda^j X^{-1} = X \begin{bmatrix} \lambda_1^j & & & \\ & \lambda_2^j & & \\ & & \ddots & \\ & & & \lambda_n^j \end{bmatrix} X^{-1}.$$

It easily follows that

$$(59) \quad e^A = \sum_{j=0}^{\infty} \frac{A^j}{j!} = X \left(\sum_{j=0}^{\infty} \frac{\Lambda^j}{j!} \right) X^{-1} = X e^\Lambda X^{-1},$$

which further reduces to

$$e^A = X \begin{bmatrix} \sum_{j=0}^{\infty} \frac{\lambda_1^j}{j!} & & & \\ & \sum_{j=0}^{\infty} \frac{\lambda_2^j}{j!} & & \\ & & \ddots & \\ & & & \sum_{j=0}^{\infty} \frac{\lambda_n^j}{j!} \end{bmatrix} X^{-1} = X \begin{bmatrix} e^{\lambda_1} & & & \\ & e^{\lambda_2} & & \\ & & \ddots & \\ & & & e^{\lambda_n} \end{bmatrix} X^{-1}.$$

This means that calculating the matrix exponential of a diagonalizable A essentially amounts to exponentiating its eigenvalues (as well as computing its eigenvalues and eigenbasis to begin with).

EXAMPLE 4.1. *The matrix*

$$A := \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} -1 & 0 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix}^{-1}$$

is obviously diagonalizable (as (i) we have already diagonalized it and (ii) it is Hermitian/symmetric). According to the above considerations, the matrix exponential of A can be written as

$$e^A = \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} e^{-1} & 0 \\ 0 & e^3 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix}^{-1}.$$

The exponential of a non-diagonalizable $A \in \mathbb{C}^{n \times n}$ can be transformed into a more tractable form by resorting to the Jordan normal form (53): If A , P and J are as in Theorem 3.2, then

$$e^A = P e^J P^{-1},$$

which can be deduced in exactly the same way as (59). Furthermore, it can be argued rather straightforwardly that

$$e^J = \begin{bmatrix} e^{J_{n_1}(\lambda_1)} & & & \\ & e^{J_{n_2}(\lambda_2)} & & \\ & & \ddots & \\ & & & e^{J_{n_l}(\lambda_l)} \end{bmatrix},$$

where the Jordan blocks $J_{n_j}(\lambda_j)$, $j = 1, \dots, l$, are as in (52). The exponentials of the Jordan blocks, i.e. $e^{J_{n_j}(\lambda_j)} \in \mathbb{C}^{n_j \times n_j}$, $j = 1, \dots, l$, can still be written down explicitly with the help of the eigenvalues of A and factorials, but we save our strength by not doing this here.

4.1. Qualitative analysis of differential equations. The matrix exponential is typically not numerically computed, but it can, e.g., be used to obtain qualitative understanding about the behavior of a solution to a system of differential equations close to so-called *equilibrium points*. Consider the rather general (autonomous) system of ordinary differential equations

$$(60) \quad \mathbf{x}'(t) = \mathbf{F}(\mathbf{x}(t)) \quad \text{for } t > 0, \quad \mathbf{x}(0) = \mathbf{x}_0 \in \mathbb{R}^n,$$

where $\mathbf{x} : \mathbb{R}_+ \rightarrow \mathbb{R}^n$ is the solution and $\mathbf{F} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a nice enough function (e.g., twice continuously differentiable). Observe, in particular, that for a real-valued A the system (58) is a special case of (60). An equilibrium point $\tilde{\mathbf{x}} \in \mathbb{R}^n$ of (60) is such that

$$\mathbf{F}(\tilde{\mathbf{x}}) = 0.$$

If a solution to (60) ends up at a certain time $\tilde{t} \geq 0$ to an equilibrium point, i.e. $\mathbf{x}(\tilde{t}) = \tilde{\mathbf{x}}$, it will stay there forever, i.e. $\mathbf{x}(t) = \tilde{\mathbf{x}}$ for all $t \geq \tilde{t}$, as the right-hand side of (60), and thus also the derivative of the solution in question, vanishes at the equilibrium.

Close to an equilibrium point $\tilde{\mathbf{x}}$, the right hand side of (60) can be linearized as

$$\mathbf{F}(\mathbf{x}) \approx \mathbf{F}(\tilde{\mathbf{x}}) + D\mathbf{F}(\tilde{\mathbf{x}})(\mathbf{x} - \tilde{\mathbf{x}}) = D\mathbf{F}(\tilde{\mathbf{x}})(\mathbf{x} - \tilde{\mathbf{x}}),$$

where $D\mathbf{F}(\tilde{\mathbf{x}}) \in \mathbb{R}^{n \times n}$ is the Jacobian matrix of \mathbf{F} evaluated at $\tilde{\mathbf{x}}$. By plugging this approximation in (60), one obtains that

$$\mathbf{x}'(t) \approx D\mathbf{F}(\tilde{\mathbf{x}})(\mathbf{x}(t) - \tilde{\mathbf{x}}) \quad \iff \quad \frac{d}{dt}(\mathbf{x}(t) - \tilde{\mathbf{x}}) \approx D\mathbf{F}(\tilde{\mathbf{x}})(\mathbf{x}(t) - \tilde{\mathbf{x}}),$$

or

$$(61) \quad \mathbf{y}'(t) \approx D\mathbf{F}(\tilde{\mathbf{x}})\mathbf{y}$$

after introducing the spatially shifted solution $\mathbf{y}(t) := \mathbf{x}(t) - \tilde{\mathbf{x}}$. The approximation in (61) is good if one only studies the solution of (60) close to the equilibrium point $\tilde{\mathbf{x}}$, that is, solutions of (61) close to the origin. In this case, a solution $\mathbf{y}(t)$ to (61) is essentially a matrix exponential function defined by the Jacobian matrix $D\mathbf{F}(\tilde{\mathbf{x}})$; cf. (58) and its solution formula with $A = D\mathbf{F}(\tilde{\mathbf{x}})$. Thus, the eigenvalues of the Jacobian matrix determine the local behavior of the system (60) close to the equilibrium point.

EXAMPLE 4.2. *As an example, consider a Lotka–Volterra model*

$$(62) \quad \mathbf{x}'(t) = \mathbf{F}(\mathbf{x}(t)) := \begin{bmatrix} \frac{2}{3}x_1(t) - \frac{4}{3}x_1(t)x_2(t) \\ x_1(t)x_2(t) - x_2(t) \end{bmatrix} \quad \text{for } t > 0, \quad \mathbf{x}(0) = \mathbf{x}_0.$$

This model describes the interaction between predator and prey populations (with certain parameter choices). The unknowns x_1 and x_2 are the number of predators and the number of preys, respectively. A rough understanding about the behavior of the solutions to the Lotka–Volterra system can be obtained by visualizing the vector field \mathbf{F} defining the right-hand side of (62): $\mathbf{F}(\mathbf{z})$ gives a tangent to the trajectory of any solution $\mathbf{x}(t)$ of (62) passing through $\mathbf{z} \in \mathbb{R}^2$; see Figure 1.

Our interest lies with the behavior of the system close to equilibrium points $\tilde{\mathbf{x}} \in \mathbb{R}^2$, where $\mathbf{F}(\tilde{\mathbf{x}}) = 0$. The studied Lotka–Volterra model has two equilibrium points, namely $\tilde{\mathbf{x}}_1 = [1, 0.5]^T$ and $\tilde{\mathbf{x}}_2 = [0, 0]^T$. Around these points, the behavior of the system can be approximated by the corresponding linearized systems, that is,

$$\mathbf{y}'_j(t) = D\mathbf{F}(\tilde{\mathbf{x}}_j)\mathbf{y}_j(t), \quad j = 1, 2,$$

where $\mathbf{y}_j(t) = \mathbf{x}(t) - \tilde{\mathbf{x}}_j$, $j = 1, 2$. For (62),

$$D\mathbf{F}(\mathbf{x}) = \begin{bmatrix} \frac{2}{3} - \frac{4}{3}x_2 & -\frac{4}{3}x_1 \\ x_2 & x_1 - 1 \end{bmatrix},$$

and so

$$D\mathbf{F}(\tilde{\mathbf{x}}_1) = \begin{bmatrix} 0 & -\frac{4}{3} \\ \frac{1}{2} & 0 \end{bmatrix} \quad \text{and} \quad D\mathbf{F}(\tilde{\mathbf{x}}_2) = \begin{bmatrix} \frac{2}{3} & 0 \\ 0 & -1 \end{bmatrix}.$$

Both linearized systems can now be exactly solved and the behavior of their solutions studied using the matrix exponential.




FIGURE 1. The arrows visualize the vector field $\mathbf{F} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ for the Lotka–Volterra model of Example 4.2. A trajectory starting from the black circle is visualized by red stars. The system has two equilibrium points, $(1, 0.5)$ and $(0, 0)$.

Least squares problems

The last theme of these lecture notes are least squares problems that arise, e.g., from the need to fit the parameters of a linear mathematical model to given measurements. To minimize the effect of measurement errors on the estimated parameters, the measurement is typically repeated several times. This leads to an overdetermined problem that includes more observations than model parameters. (In this chapter, we return to the setting where the considered vectors and matrices are real-valued.)

1. Motivation, definition and equivalence to normal equation

We start with two examples:

EXAMPLE 1.1. *A simple example of a least squares problem is the determination of a spring constant from several measurements. Assume that we have at our disposal n weights m_i , $i = 1, \dots, n$, and can measure the elongation x_i of the spring when the weight m_i is attached to it. In addition, assume that the gravity g is known and we are dealing with an ideal spring, meaning that the elongation of the spring under the mass m_i behaves as*

$$x_i = kgm_i, \quad i = 1, \dots, n,$$

where $k \in \mathbb{R}$ is the spring constant. By repeating the measurement with all available weights, one obtains a (trivial) linear system for k :

$$\begin{bmatrix} gm_1 \\ gm_2 \\ \vdots \\ gm_n \end{bmatrix} k = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}.$$

Due to measurement errors, this system does not typically have a solution. Instead, one determines (an estimate for) the parameter k by minimizing the 2-norm of the residual: Let

$$A = \begin{bmatrix} gm_1 \\ gm_2 \\ \vdots \\ gm_n \end{bmatrix} \in \mathbb{R}^{n \times 1} \quad \text{and} \quad \mathbf{b} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} \in \mathbb{R}^n,$$

and seek k as a solution to

$$\min_{k \in \mathbb{R}} \|Ak - \mathbf{b}\|_2^2.$$

It follows easily that

$$\|Ak - \mathbf{b}\|_2^2 = (Ak - \mathbf{b})^T (Ak - \mathbf{b}) = (A^T A)k^2 - 2(A^T \mathbf{b})k + \|\mathbf{b}\|_2^2,$$

which is an equation for an upwards opening parabola in the variable k . The spring constant can thus be solved by computing the zero of the derivative with respect to k , which gives $k = (A^T A)^{-1} A^T \mathbf{b}$.

EXAMPLE 1.2. *As a second example, let us consider finding an ellipsoid that is the best fit to a given set of points $\mathbf{x}^{(i)} \in \mathbb{R}^2$, $i = 1, \dots, n$. There are several ways to define what “best” means and to compute the fit, but for simplicity we have opted to follow an approach that is based on the*




FIGURE 1. A small cylindrical object and two ellipses fitted to its cross section.

so-called conic section presentation of an ellipse. In this presentation, an ellipse is defined as a set

$$(63) \quad \{(x_1, x_2) \in \mathbb{R}^2 \mid x_1^2 + Ax_2^2 + Bx_1x_2 + Cx_1 + Dx_2 + E = 0, \quad A, B, C, D, E \in \mathbb{R}, \quad B^2 - 4A < 0\},$$

which makes fitting an ellipse to a set of points relatively easy, but some work would be needed to find out the semi-axes and center-point of the ellipse.

To be more precise, the task is to define the coefficients A, B, C, D and E so that the data points $\mathbf{x}^{(i)} \in \mathbb{R}^2$, $i = 1, \dots, n$, satisfy the equation defining the set (63) ‘as well as possible’. The simplest way to formulate such an optimality condition is to look for A, B, C, D and E that minimize the sum of the squared discrepancies

$$(64) \quad \sum_{i=1}^n \left((x_1^{(i)})^2 + A(x_2^{(i)})^2 + Bx_1^{(i)}x_2^{(i)} + Cx_1^{(i)} + Dx_2^{(i)} + E \right)^2.$$

Let us collect the unknown coefficients A, B, C, D and E into a single vector $\boldsymbol{\alpha} = [A, B, C, D, E]^T \in \mathbb{R}^5$ and set

$$A = \begin{bmatrix} (x_2^{(1)})^2 & x_1^{(1)}x_2^{(1)} & x_1^{(1)} & x_2^{(1)} & 1 \\ (x_2^{(2)})^2 & x_1^{(2)}x_2^{(2)} & x_1^{(2)} & x_2^{(2)} & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ (x_2^{(n)})^2 & x_1^{(n)}x_2^{(n)} & x_1^{(n)} & x_2^{(n)} & 1 \end{bmatrix} \quad \text{and} \quad \mathbf{b} = - \begin{bmatrix} (x_1^{(1)})^2 \\ (x_1^{(2)})^2 \\ \vdots \\ (x_1^{(n)})^2 \end{bmatrix}.$$

With this notation, the minimization of (64) can be rewritten as follows: find the minimizer $\boldsymbol{\alpha} \in \mathbb{R}^5$ of

$$\|A\boldsymbol{\alpha} - \mathbf{b}\|_2^2,$$

which is of the general least squares form that we will study in more detail in what follows.

Motivated by the preceding two examples, we define a *least squares* (LSQ) solution of our original linear system (1) as follows:

DEFINITION 1.1 (LSQ problem). Let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. A vector $\mathbf{x} \in \mathbb{R}^n$ is called a least squares solution of the equation (1) if it minimizes the squared discrepancy

$$(65) \quad \|A\mathbf{x} - \mathbf{b}\|_2^2$$

measured in the 2-norm.

It turns out that minimizing (65) is equivalent to solving the so-called *normal equation*,

$$(66) \quad A^T A \mathbf{x} = A^T \mathbf{b}.$$

This is not very surprising since (66) characterizes the points where the gradient of the squared sum (65) vanishes (homework). If $N(A) = \{0\}$, then also $N(A^T A) = \{0\}$ ¹, meaning that $A^T A$ is invertible and (66) has a unique solution. However, it is not quite obvious that (66) always has at least one solution; will prove such an existence result later in this chapter.

THEOREM 1.1. *Assume that (66) has at least one solution (as it always has!). Then $\mathbf{x} \in \mathbb{R}^n$ is a minimizer of (65) if and only if it solves (66).*

PROOF. Let $\mathbf{x} \in \mathbb{R}^n$ be a solution to (66) and write an arbitrary $\tilde{\mathbf{x}} \in \mathbb{R}^n$ in the form $\tilde{\mathbf{x}} = \mathbf{x} + \mathbf{z}$, i.e., set $\mathbf{z} = \tilde{\mathbf{x}} - \mathbf{x}$. Then, it holds that

$$\begin{aligned} \|A\tilde{\mathbf{x}} - \mathbf{b}\|_2^2 &= \|A(\mathbf{x} + \mathbf{z}) - \mathbf{b}\|_2^2 = \|(A\mathbf{x} - \mathbf{b}) + A\mathbf{z}\|_2^2 = ((A\mathbf{x} - \mathbf{b}) + A\mathbf{z})^T ((A\mathbf{x} - \mathbf{b}) + A\mathbf{z}) \\ &= \|A\mathbf{x} - \mathbf{b}\|_2^2 + (A\mathbf{x} - \mathbf{b})^T A\mathbf{z} + (A\mathbf{z})^T (A\mathbf{x} - \mathbf{b}) + \|A\mathbf{z}\|_2^2 \\ &= \|A\mathbf{x} - \mathbf{b}\|_2^2 + 2\mathbf{z}^T (A^T A\mathbf{x} - A^T \mathbf{b}) + \|A\mathbf{z}\|_2^2 \\ (67) \quad &= \|A\mathbf{x} - \mathbf{b}\|_2^2 + \|A\mathbf{z}\|_2^2 \geq \|A\mathbf{x} - \mathbf{b}\|_2^2, \end{aligned}$$

where the penultimate step follows from \mathbf{x} being a solution of the normal equation. In other words, \mathbf{x} is a solution of the least squares problem because $\|A\mathbf{x} - \mathbf{b}\|_2 \leq \|A\tilde{\mathbf{x}} - \mathbf{b}\|_2$ for an arbitrary $\tilde{\mathbf{x}} \in \mathbb{R}^n$.

Let $\tilde{\mathbf{x}}$ now also be a least squares solution, which means that equality holds in (67). This implies that $\mathbf{z} = \tilde{\mathbf{x}} - \mathbf{x}$ belongs to $N(A)$. Hence,

$$A^T A\tilde{\mathbf{x}} = A^T A(\mathbf{x} + \mathbf{z}) = A^T A\mathbf{x} + A^T A\mathbf{z} = A^T \mathbf{b},$$

i.e., $\tilde{\mathbf{x}}$ is also a solution to the normal equation (66). \square

The next goal is to enhance our geometric understanding about the normal equation (66): What is the intuitive reason for it being equivalent to the least squares problem (65)? To this end, we need to introduce the concept of a projection.

2. Projection matrices

Projection matrices are related to decompositions of a given vector $\mathbf{x} \in \mathbb{R}^n$ into two parts. To this end, let \mathcal{V} and \mathcal{W} be subspaces of \mathbb{R}^n and assume that

$$(68) \quad \mathcal{V} \cap \mathcal{W} = \{0\} \quad \text{and} \quad \mathcal{V} + \mathcal{W} = \mathbb{R}^n,$$

where

$$\mathcal{V} + \mathcal{W} = \{\mathbf{v} + \mathbf{w} \mid \mathbf{v} \in \mathcal{V} \text{ and } \mathbf{w} \in \mathcal{W}\}.$$

This means that \mathbb{R}^n is the *direct sum* of \mathcal{V} and \mathcal{W} , which is expressed as $\mathbb{R}^n = \mathcal{V} \oplus \mathcal{W}$.

LEMMA 2.1. *If $\mathbb{R}^n = \mathcal{V} \oplus \mathcal{W}$, i.e. (68) is valid, then any \mathbf{x} can be decomposed uniquely as*

$$\mathbf{x} = \mathbf{v} + \mathbf{w},$$

where $\mathbf{v} \in \mathcal{V}$ and $\mathbf{w} \in \mathcal{W}$.

PROOF. Due to (68), it is obvious that any $\mathbf{x} \in \mathbb{R}^n$ has the required decomposition, and so we only need to worry about its uniqueness. Let $\mathbf{v}, \tilde{\mathbf{v}} \in \mathcal{V}$ and $\mathbf{w}, \tilde{\mathbf{w}} \in \mathcal{W}$ be such that

$$\mathbf{v} + \mathbf{w} = \mathbf{x} = \tilde{\mathbf{v}} + \tilde{\mathbf{w}}.$$

Hence,

$$\mathcal{V} \ni \mathbf{v} - \tilde{\mathbf{v}} = \tilde{\mathbf{w}} - \mathbf{w} \in \mathcal{W},$$

and by the first condition in (68), it must hold that

$$\mathbf{v} - \tilde{\mathbf{v}} = \tilde{\mathbf{w}} - \mathbf{w} = 0,$$

as the zero vector is the only common element of \mathcal{V} and \mathcal{W} . In other words, $\mathbf{v} = \tilde{\mathbf{v}}$ and $\mathbf{w} = \tilde{\mathbf{w}}$, and thus the decomposition must be unique. \square

¹Obviously $N(A) \subset N(A^T A)$, but also $N(A^T A) \subset N(A)$: if $\mathbf{z} \in N(A^T A)$, then $0 = \mathbf{z}^T (A^T A\mathbf{z}) = \|A\mathbf{z}\|_2^2$, i.e., $\mathbf{z} \in N(A)$.

The projections, or projection matrices $P_{\mathcal{V}}, P_{\mathcal{W}} \in \mathbb{R}^{n \times n}$ associated to a decomposition $\mathbb{R}^n = \mathcal{V} \oplus \mathcal{W}$ are defined via

$$P_{\mathcal{V}}\mathbf{x} = \mathbf{v} \quad \text{and} \quad P_{\mathcal{W}}\mathbf{x} = \mathbf{w},$$

where $\mathbf{v} \in \mathcal{V}$ and $\mathbf{w} \in \mathcal{W}$ are the unique vectors such that $\mathbf{x} = \mathbf{v} + \mathbf{w}$. It would be straightforward to prove that such matrices exist, i.e., that the functions sending \mathbf{x} to its unique components in \mathcal{V} and \mathcal{W} are linear mappings. However, we take a more constructive path and find explicit representations for $P_{\mathcal{V}}$ and $P_{\mathcal{W}}$.

Let $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ and $\{\mathbf{w}_1, \dots, \mathbf{w}_l\}$ be basis for \mathcal{V} and \mathcal{W} , that is, these sets of vectors are linearly independent and

$$\mathcal{V} = \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\} \quad \text{and} \quad \mathcal{W} = \text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_l\}.$$

According to (68), any vector in \mathbb{R}^n can be presented as a joint linear combination of $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ and $\{\mathbf{w}_1, \dots, \mathbf{w}_l\}$. Let us prove that the joint set of vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_k, \mathbf{w}_1, \dots, \mathbf{w}_l\}$ is also linearly independent: Let $\boldsymbol{\alpha} \in \mathbb{R}^{k+l}$ be such that

$$\sum_{i=1}^k \alpha_i \mathbf{v}_i + \sum_{j=1}^l \alpha_{k+j} \mathbf{w}_j = 0 \quad \iff \quad \mathcal{V} \ni \sum_{i=1}^k \alpha_i \mathbf{v}_i = - \sum_{j=1}^l \alpha_{k+j} \mathbf{w}_j \in \mathcal{W},$$

where both sides must vanish by virtue of (68). Because $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ and $\{\mathbf{w}_1, \dots, \mathbf{w}_l\}$ are separately linearly independent, this means that the whole vector $\boldsymbol{\alpha} \in \mathbb{R}^{k+l}$ must vanish. Hence, $\{\mathbf{v}_1, \dots, \mathbf{v}_k, \mathbf{w}_1, \dots, \mathbf{w}_l\}$ is a set of linearly independent vectors that span \mathbb{R}^n , i.e., they form a basis for \mathbb{R}^n . In particular, their total number $k+l$ must equal the dimension of \mathbb{R}^n ; in the following, we write $l = n - k$.

We stack the two bases as columns of the matrices $W \in \mathbb{R}^{n \times k}$ and $V \in \mathbb{R}^{n \times (n-k)}$, respectively:

$$V = [\mathbf{v}_1, \dots, \mathbf{v}_k] \quad \text{and} \quad W = [\mathbf{w}_1, \dots, \mathbf{w}_{n-k}].$$

With this notation, the two subspaces can be expressed as $\mathcal{V} = R(V)$ and $\mathcal{W} = R(W)$; see (5). The decomposition of \mathbb{R}^n related to \mathcal{V} and \mathcal{W} is thus

$$\mathbf{x} = \mathbf{v} + \mathbf{w} = V\mathbf{z}_V + W\mathbf{z}_W = [V, W] \begin{bmatrix} \mathbf{z}_V \\ \mathbf{z}_W \end{bmatrix}$$

for some ‘coordinate vectors’ $\mathbf{z}_V \in \mathbb{R}^k$ and $\mathbf{z}_W \in \mathbb{R}^{n-k}$. Because the columns of $[W, V] \in \mathbb{R}^{n \times n}$ are a (linearly independent) basis for \mathbb{R}^n , the matrix $[W, V]$ is invertible. Hence,

$$\begin{bmatrix} \mathbf{z}_V \\ \mathbf{z}_W \end{bmatrix} = [V, W]^{-1} \mathbf{x}.$$

On the other hand,

$$\mathbf{z}_V = [I, 0] \begin{bmatrix} \mathbf{z}_V \\ \mathbf{z}_W \end{bmatrix} \quad \text{and} \quad \mathbf{z}_W = [0, I] \begin{bmatrix} \mathbf{z}_V \\ \mathbf{z}_W \end{bmatrix},$$

where $I \in \mathbb{R}^{k \times k}$, $0 \in \mathbb{R}^{k \times (n-k)}$ in the first formula and $I \in \mathbb{R}^{(n-k) \times (n-k)}$, $0 \in \mathbb{R}^{(n-k) \times k}$ in the second one. Altogether, we have thus deduced that

$$(69) \quad P_{\mathcal{V}}\mathbf{x} = V\mathbf{z}_V = V[I, 0] \begin{bmatrix} \mathbf{z}_V \\ \mathbf{z}_W \end{bmatrix} = [V, 0][V, W]^{-1} \mathbf{x},$$

where the second 0 belongs to $\mathbb{R}^{n \times (n-k)}$. Similarly,

$$(70) \quad P_{\mathcal{W}}\mathbf{x} = W\mathbf{z}_W = W[0, I] \begin{bmatrix} \mathbf{z}_V \\ \mathbf{z}_W \end{bmatrix} = [0, W][V, W]^{-1} \mathbf{x},$$

where the latter 0 belongs to $\mathbb{R}^{n \times k}$. In other words, $P_{\mathcal{V}} = [V, 0][V, W]^{-1}$ and $P_{\mathcal{W}} = [0, W][V, W]^{-1}$, where the zeros denote zero matrices of appropriate sizes.

By construction, the matrices $P_{\mathcal{V}}$ and $P_{\mathcal{W}}$ must satisfy

$$P_{\mathcal{V}}^2 = P_{\mathcal{V}} \quad \text{and} \quad P_{\mathcal{W}}^2 = P_{\mathcal{W}}.$$

Indeed, e.g., $P_{\mathcal{V}}\mathbf{x}$ is already an element of \mathcal{V} , and so its unique decomposition of the form (68) must be $P_{\mathcal{V}}\mathbf{x} = P_{\mathcal{V}}\mathbf{x} + 0$, i.e., $P_{\mathcal{V}}^2\mathbf{x} = P_{\mathcal{V}}(P_{\mathcal{V}}\mathbf{x}) = P_{\mathcal{V}}\mathbf{x}$ for all $\mathbf{x} \in \mathbb{R}^n$. This could also be verified by a direct calculation based on (69). In other words, the projection matrices $P_{\mathcal{V}}$ and $P_{\mathcal{W}}$ are

identity maps on their respective ranges. This is a fundamental property that can be taken as the definition of a projection (matrix).

DEFINITION 2.1. A matrix $P \in \mathbb{R}^{n \times n}$ is a projection if it satisfies

$$(71) \quad P^2 = P.$$

Each projection matrix is related to a decomposition of the space \mathbb{R}^n into two parts. Let $P \in \mathbb{R}^{n \times n}$ be a projection and consider the decomposition

$$(72) \quad \mathbf{x} = P\mathbf{x} + (I - P)\mathbf{x},$$

that is, $\mathcal{V} = R(P)$ and $\mathcal{W} = R(I - P)$. Due to the property $P = P^2$ it is easy to see that these two subspaces satisfy the conditions (68): The second one is already proven by (72). On the other hand, if \mathbf{x} belongs to both $R(P)$ and $R(I - P)$, it can be given as $\mathbf{x} = (I - P)\mathbf{z}$ for some \mathbf{z} , but as P is a projection, it also holds that $P\mathbf{x} = \mathbf{x}$. Putting these together,

$$\mathbf{x} = P\mathbf{x} = P(I - P)\mathbf{z} = (P - P^2)\mathbf{z} = (P - P)\mathbf{z} = \mathbf{0},$$

i.e., the zero vector is the only common element of $R(P)$ and $R(I - P)$. Note that $I - P$ is also a projection:

$$(I - P)(I - P) = I - P - P + P^2 = I - 2P + P = I - P,$$

where we once again utilized the identity $P^2 = P$.

A special projection is related to subspaces \mathcal{V} and \mathcal{W} that are orthogonal.

DEFINITION 2.2. Let \mathcal{V} and \mathcal{W} be subspaces of \mathbb{R}^n and $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ be some inner product. The subspaces \mathcal{V} and \mathcal{W} are orthogonal with respect to $\langle \cdot, \cdot \rangle$ if

$$\langle \mathbf{v}, \mathbf{w} \rangle = 0$$

for all $\mathbf{v} \in \mathcal{V}$ and $\mathbf{w} \in \mathcal{W}$.

We also need the definition for the orthogonal complement of a subspace \mathcal{V} .

DEFINITION 2.3. Let \mathcal{V} be subspace of \mathbb{R}^n and $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ some inner product. The orthogonal complement of \mathcal{V} is the subspace

$$(73) \quad \mathcal{V}^\perp := \{\mathbf{w} \in \mathbb{R}^n \mid \langle \mathbf{w}, \mathbf{v} \rangle = 0 \text{ for all } \mathbf{v} \in \mathcal{V}\}.$$

(Proving that \mathcal{V}^\perp is actually a subspace is left as an exercise.)

It holds that

$$(74) \quad \mathcal{V} \cap \mathcal{V}^\perp = \{0\} \quad \text{and} \quad \mathcal{V} + \mathcal{V}^\perp = \mathbb{R}^n,$$

i.e., $\mathcal{V} \oplus \mathcal{V}^\perp = \mathbb{R}^n$. Indeed, if $\mathbf{v} \in \mathcal{V}$ also belongs to \mathcal{V}^\perp , then by definition

$$\langle \mathbf{v}, \mathbf{v} \rangle = 0, \quad \text{i.e.,} \quad \mathbf{v} = \mathbf{0}.$$

On the other hand, if $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ is an orthonormal basis for \mathcal{V} ,² i.e., the vectors are orthogonal (cf. Definition 2.3) and of unit length in the norm defined by the considered inner product, then any $\mathbf{x} \in \mathbb{R}^n$ can be written as

$$(75) \quad \mathbf{x} = \sum_{i=1}^k \langle \mathbf{x}, \mathbf{v}_i \rangle \mathbf{v}_i + \left(\mathbf{x} - \sum_{i=1}^k \langle \mathbf{x}, \mathbf{v}_i \rangle \mathbf{v}_i \right) =: \mathbf{v}_\mathbf{x} + \mathbf{w}_\mathbf{x},$$

which gives \mathbf{x} as a sum of a vector in \mathcal{V} and another in \mathcal{V}^\perp : Clearly, $\mathbf{v}_\mathbf{x} \in \mathcal{V}$ and proving that $\mathbf{w}_\mathbf{x} \in \mathcal{V}^\perp$ is left as a homework (write an arbitrary element of \mathcal{V} in the orthonormal basis, take its inner product with $\mathbf{w}_\mathbf{x}$ and employ the orthonormality of the basis vectors). In fact, $\mathbf{v}_\mathbf{x}$ is the *orthogonal* projection of \mathbf{x} onto \mathcal{V} and $\mathbf{w}_\mathbf{x}$ the *orthogonal* projection of \mathbf{x} onto \mathcal{V}^\perp , as will be defined in the following.

For the rest of this chapter, we only consider orthogonal complements with respect to the Euclidean inner product. A projection matrix related to an orthogonal decomposition $\mathbb{R}^n = \mathcal{V} \oplus \mathcal{V}^\perp$ is called an *orthogonal projection* matrix. Let $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ and $\{\mathbf{v}_1^\perp, \dots, \mathbf{v}_{n-k}^\perp\}$ be (not

²Such a basis can be constructed by the Gram–Schmidt orthogonalization process as we will learn later.

necessarily orthonormal) bases for \mathcal{V} and \mathcal{V}^\perp , respectively, and define the matrices $V \in \mathbb{R}^{n \times k}$ and $V^\perp \in \mathbb{R}^{n \times (n-k)}$ in the standard manner, i.e.,

$$V = [\mathbf{v}_1, \dots, \mathbf{v}_k] \quad \text{and} \quad V^\perp = [\mathbf{v}_1^\perp, \dots, \mathbf{v}_{n-k}^\perp].$$

Note once again that $\mathcal{V} = R(V) = \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_k)$ and $\mathcal{V}^\perp = R(V^\perp) = \text{span}(\mathbf{v}_1^\perp, \dots, \mathbf{v}_{n-k}^\perp)$. By definition (73), the columns of V are orthogonal to those of V^\perp (in the Euclidean inner product), meaning that

$$(76) \quad V^T V^\perp = 0.$$

On the other hand, the decomposition of \mathbb{R}^n related to \mathcal{V} and \mathcal{V}^\perp is

$$(77) \quad \mathbf{x} = V \mathbf{z}_V + V^\perp \mathbf{z}_{V^\perp}$$

for some yet-to-be-defined $\mathbf{z}_V \in \mathbb{R}^k$ and $\mathbf{z}_{V^\perp} \in \mathbb{R}^{n-k}$. Multiplying from the left by V^T and using the orthogonality property (76), we get

$$V^T \mathbf{x} = V^T V \mathbf{z}_V.$$

Because $N(V^T V) = N(V)$ is trivial as the columns of $V \in \mathbb{R}^{n \times k}$ are linearly independent, $V^T V \in \mathbb{R}^{k \times k}$ is invertible and \mathbf{z}_V can be solved as

$$\mathbf{z}_V = (V^T V)^{-1} V^T \mathbf{x}.$$

The projection matrix related to the decomposition $\mathbb{R}^n = \mathcal{V} \oplus \mathcal{V}^\perp$ is thus

$$(78) \quad P_V = V(V^T V)^{-1} V^T \in \mathbb{R}^{n \times n},$$

see (77). In particular, the orthogonal projection does not require one to form a basis for \mathcal{V}^\perp : the auxiliary basis $\{\mathbf{v}_1^\perp, \dots, \mathbf{v}_{n-k}^\perp\}$ of V^\perp does not anymore appear in (78). Moreover, if $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ are orthonormal, then V is orthogonal/unitary, i.e. $V^T V = I$, and so

$$(79) \quad P_V \mathbf{x} = V I^{-1} V^T \mathbf{x} = V \begin{bmatrix} \mathbf{v}_1^T \mathbf{x} \\ \vdots \\ \mathbf{v}_k^T \mathbf{x} \end{bmatrix} = \sum_{i=1}^k (\mathbf{v}_i^T \mathbf{x}) \mathbf{v}_i,$$

which corresponds to the first term in (75).

An orthonormal projection satisfies some useful identities. First of all, since inversion and transposition commute, we have

$$P_V^T = (V(V^T V)^{-1} V^T)^T = (V^T)^T ((V^T V)^T)^{-1} V^T = V(V^T V)^{-1} V^T = P_V,$$

and hence also,

$$(80) \quad P_V^T (I - P_V) = P_V (I - P_V) = P_V - P_V^2 = P_V - P_V = 0.$$

After transposition, this yields

$$(81) \quad (I - P_V)^T P_V = 0.$$

In fact, $I - P_V$ is the orthogonal projection onto \mathcal{V}^\perp (homework).

EXAMPLE 2.1. Let $\mathbf{a}, \mathbf{x} \in \mathbb{R}^2$. Consider a decomposition of the vector \mathbf{x} into two components, one parallel and the other orthogonal to \mathbf{a} . Such a decomposition can be constructed as in (75):

$$\mathbf{x} = \frac{\mathbf{a}^T \mathbf{x}}{\|\mathbf{a}\|_2} \mathbf{a} + \left(\mathbf{x} - \frac{\mathbf{a}^T \mathbf{x}}{\|\mathbf{a}\|_2} \mathbf{a} \right) = \frac{1}{\|\mathbf{a}\|_2^2} \mathbf{a} \mathbf{a}^T \mathbf{x} + \left(I - \frac{1}{\|\mathbf{a}\|_2^2} \mathbf{a} \mathbf{a}^T \right) \mathbf{x},$$

where the division by $\|\mathbf{a}\|_2^2$ corresponds to normalization of \mathbf{a} , i.e., $\mathbf{a}/\|\mathbf{a}\|_2$ is a unit vector. The two components of this decomposition are orthogonal (see Figure 2) and the related projection matrices are

$$(82) \quad P = \frac{1}{\|\mathbf{a}\|_2^2} \mathbf{a} \mathbf{a}^T.$$

and $I - P$.




FIGURE 2. Orthogonal projection of \mathbf{a} onto \mathbf{b} .




FIGURE 3. One possible oblique projection of \mathbf{a} onto \mathbf{b} .

EXAMPLE 2.2. Let $\mathbf{a}, \mathbf{c}, \mathbf{x} \in \mathbb{R}^2$ be such that $[\mathbf{a}, \mathbf{c}] \in \mathbb{R}^{2 \times 2}$ is invertible, i.e., \mathbf{a} and \mathbf{c} are linearly independent. Consider the decomposition of \mathbf{x} in the basis defined by \mathbf{a} and \mathbf{c} :

$$(83) \quad \mathbf{x} = \alpha \mathbf{a} + \beta \mathbf{c}.$$

The coefficients α and β can be solved from the linear system

$$[\mathbf{a}, \mathbf{c}] \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \mathbf{x},$$

and the parameter α can then be extracted from the solution via

$$\alpha = [1, 0][\mathbf{a}, \mathbf{c}]^{-1} \mathbf{x},$$

where $[1, 0] \in \mathbb{R}^{1 \times 2}$. Thus, one of the two projection matrices related to this decomposition is

$$P = [\mathbf{a}, 0][\mathbf{a}, \mathbf{c}]^{-1},$$

where $[\mathbf{a}, 0] \in \mathbb{R}^{2 \times 2}$ (cf. (69)).

3. Geometric interpretation of a least squares solution

With the help of projection matrices, we can now understand the geometry behind the least squares problem (65) and the related normal equation (66). Let $P \in \mathbb{R}^m$ be an orthogonal projection from \mathbb{R}^m to $R(A)$. As P keeps elements in $R(A)$ fixed, the least squares functional can be rewritten as

$$\begin{aligned} \|A\mathbf{x} - \mathbf{b}\|_2^2 &= \|PA\mathbf{x} - \mathbf{b}\|_2^2 = \|P(A\mathbf{x} - \mathbf{b}) - (I - P)\mathbf{b}\|_2^2 \\ &= \|PA\mathbf{x} - P\mathbf{b}\|_2^2 - (P(A\mathbf{x} - \mathbf{b}))^T (I - P)\mathbf{b} - ((I - P)\mathbf{b})^T (P(A\mathbf{x} - \mathbf{b})) + \|(I - P)\mathbf{b}\|_2^2 \\ (84) \quad &= \|A\mathbf{x} - P\mathbf{b}\|_2^2 + \|(I - P)\mathbf{b}\|_2^2, \end{aligned}$$

where the last step follows from (80) or/and (81). Because the second term on the right-hand side of (84) does not depend on \mathbf{x} , to minimize the least squares functional, \mathbf{x} should be chosen such that

$$(85) \quad A\mathbf{x} = P\mathbf{b},$$

which has a (not-necessarily-unique) solution because $P\mathbf{b} \in R(A)$ by definition. In other words, a minimizer of (65) is such \mathbf{x} that $A\mathbf{x}$ equals the orthogonal projection of \mathbf{b} onto $R(A)$.

If $N(A) = \{0\}$, i.e., $A\boldsymbol{\alpha} = 0$ if and only if $\boldsymbol{\alpha} = 0$, the columns of A are a (linearly independent) basis for $R(A)$. Thus, the definition of an orthogonal projection matrix (78) gives $P = A(A^T A)^{-1} A^T$ and (85) turns into

$$(86) \quad A\mathbf{x} = A(A^T A)^{-1} A^T \mathbf{b}.$$

In particular, the unique solution of the normal equation (66), i.e.,

$$\mathbf{x} = (A^T A)^{-1} A^T \mathbf{b}$$

clearly satisfies (86) — as it should because it is the unique minimizer of (65) if $N(A) = N(A^T A) = \{0\}$ according to Theorem 1.1.

4. QR decomposition

The normal equation approach to solving the least squares problem has two drawbacks: First of all, if A does not have a trivial nullspace, the normal equation is not uniquely solvable. Secondly, even if $N(A) = \{0\}$, the conditioning of the normal equation is much worse than that of the original problem (1): For example if $A \in \mathbb{R}^{n \times n}$, the condition numbers corresponding to the 2-norm satisfy

$$\kappa_2(A^T A) = \kappa_2(A)^2,$$

which follows from Lemma 2.5 of Chapter 1 and the fourth exercise sheet after noticing that $A^T A$ is symmetric. According to the discussion in Section 1, if the condition number is squared, solving a linear system becomes much more sensitive to, e.g., floating point errors. In some extreme cases, resorting to the normal equation may even lead to loss of information, as demonstrated by the following example.

EXAMPLE 4.1. *Let*

$$A = \begin{bmatrix} 1 & 1 \\ \epsilon & 0 \\ 0 & \epsilon \end{bmatrix} \quad \text{so that} \quad A^T = \begin{bmatrix} 1 & \epsilon & 0 \\ 1 & 0 & \epsilon \end{bmatrix}.$$

Hence,

$$A^T A = \begin{bmatrix} 1 + \epsilon^2 & 1 \\ 1 & 1 + \epsilon^2 \end{bmatrix}.$$

In floating point arithmetics $\text{fl}(1 + \epsilon^2) = 1$ when ϵ is sufficiently small; note that for $\epsilon \ll 1$, the square ϵ^2 is much smaller than ϵ itself. When this happens,

$$\text{fl}(A^T A) = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix},$$

i.e., data is lost and, in this case, the normal equation is even no longer uniquely solvable. In Matlab, this happens if $\epsilon < 10^{-9}$.

To avoid the dramatic decrease in stability due to the normal equation approach and the related floating point errors, the least squares problem is typically solved using either the QR or the singular value (SVD) decomposition for the matrix A . Both methods can be understood as ways to compute an orthonormal basis for the subspace $R(A)$.

The simplest way to compute the QR decomposition is via the Gram–Schmidt orthogonalization process that is based on the following Lemma. In what follows, “orthogonality” refers to orthogonality in the sense of the Euclidean inner product.

LEMMA 4.1. *Let $\{\mathbf{q}_1, \dots, \mathbf{q}_k\} \subset \mathbb{R}^m$, $k < m$, be a set of orthonormal vectors, i.e., $\|\mathbf{q}_i\|_2 = 1$, $i = 1, \dots, k$, and*

$$\mathbf{q}_i^T \mathbf{q}_j = 0 \quad \text{for } i \neq j.$$

In addition, assume that $\mathbf{a} \in \mathbb{R}^m$ does not belong to $\text{span}(\mathbf{q}_1, \dots, \mathbf{q}_k)$ and define

$$(87) \quad \mathbf{q}_{k+1} = \frac{\tilde{\mathbf{q}}_{k+1}}{\|\tilde{\mathbf{q}}_{k+1}\|_2}, \quad \text{where} \quad \tilde{\mathbf{q}}_{k+1} = \mathbf{a} - \sum_{i=1}^k (\mathbf{a}^T \mathbf{q}_i) \mathbf{q}_i.$$

Then $\{\mathbf{q}_1, \dots, \mathbf{q}_{k+1}\} \subset \mathbb{R}^n$ is a set of orthonormal vectors and

$$(88) \quad \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_{k+1}) = \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_k, \mathbf{a}).$$

PROOF. To begin with note that $\tilde{\mathbf{q}}_{k+1}$ defined in (87) is a nonzero vector because $\mathbf{a} \notin \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_k)$, i.e., \mathbf{a} cannot be given as a linear combination of $\mathbf{q}_1, \dots, \mathbf{q}_k$.

To prove the orthonormality of the set $\{\mathbf{q}_1, \dots, \mathbf{q}_{k+1}\}$ it is enough to prove that \mathbf{q}_{k+1} is of unit Euclidean length and orthogonal to $\mathbf{q}_1, \dots, \mathbf{q}_k$. From the first equation in (87), it is obvious that $\|\mathbf{q}_{k+1}\|_2 = 1$. Moreover, since \mathbf{q}_{k+1} and $\tilde{\mathbf{q}}_{k+1}$ are parallel, it is actually enough to show that $\tilde{\mathbf{q}}_{k+1}$ is orthogonal to $\mathbf{q}_1, \dots, \mathbf{q}_k$: for any $j = 1, \dots, k$, we have

$$\tilde{\mathbf{q}}_{k+1}^T \mathbf{q}_j = \left(\mathbf{a} - \sum_{i=1}^k (\mathbf{a}^T \mathbf{q}_i) \mathbf{q}_i \right)^T \mathbf{q}_j = \mathbf{a}^T \mathbf{q}_j - \sum_{i=1}^k (\mathbf{a}^T \mathbf{q}_i) (\mathbf{q}_i^T \mathbf{q}_j) = \mathbf{a}^T \mathbf{q}_j - \mathbf{a}^T \mathbf{q}_j = 0$$

due to the orthonormality of $\{\mathbf{q}_1, \dots, \mathbf{q}_k\}$.

Although (88) follows straightforwardly from the definition of linear span, let us anyway carefully prove it for the sake of completeness. Assume first that $\mathbf{x} \in \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_{k+1})$, i.e.,

$$\mathbf{x} = \sum_{i=1}^{k+1} \alpha_i \mathbf{q}_i = \sum_{i=1}^k \alpha_i \mathbf{q}_i + \alpha_{k+1} \mathbf{q}_{k+1}$$

for some $\alpha \in \mathbb{R}^{k+1}$. Note that (87) can be rewritten in the form

$$\mathbf{q}_{k+1} = \frac{1}{\|\tilde{\mathbf{q}}_{k+1}\|_2} \left(\mathbf{a} - \sum_{i=1}^k (\mathbf{a}^T \mathbf{q}_i) \mathbf{q}_i \right).$$

Hence,

$$\mathbf{x} = \sum_{i=1}^k \alpha_i \mathbf{q}_i + \frac{\alpha_{k+1}}{\|\tilde{\mathbf{q}}_{k+1}\|_2} \left(\mathbf{a} - \sum_{i=1}^k (\mathbf{a}^T \mathbf{q}_i) \mathbf{q}_i \right) = \sum_{i=1}^k \left(\alpha_i - \frac{\alpha_{k+1} \mathbf{a}^T \mathbf{q}_i}{\|\tilde{\mathbf{q}}_{k+1}\|_2} \right) \mathbf{q}_i + \frac{\alpha_{k+1}}{\|\tilde{\mathbf{q}}_{k+1}\|_2} \mathbf{a},$$

which is obviously in $\text{span}(\mathbf{q}_1, \dots, \mathbf{q}_k, \mathbf{a})$, meaning that $\text{span}(\mathbf{q}_1, \dots, \mathbf{q}_{k+1}) \subset \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_k, \mathbf{a})$.

On the other hand, if $\mathbf{x} \in \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_k, \mathbf{a})$, then for some $\alpha \in \mathbb{R}^{k+1}$,

$$\begin{aligned} \mathbf{x} &= \sum_{i=1}^k \alpha_i \mathbf{q}_i + \alpha_{k+1} \mathbf{a} = \sum_{i=1}^k \alpha_i \mathbf{q}_i + \alpha_{k+1} \left(\|\tilde{\mathbf{q}}_{k+1}\|_2 \mathbf{q}_{k+1} + \sum_{i=1}^k (\mathbf{a}^T \mathbf{q}_i) \mathbf{q}_i \right) \\ &= \sum_{i=1}^k (\alpha_i + \alpha_{k+1} \mathbf{a}^T \mathbf{q}_i) \mathbf{q}_i + \alpha_{k+1} \|\tilde{\mathbf{q}}_{k+1}\|_2 \mathbf{q}_{k+1}, \end{aligned}$$

which clearly belongs to $\text{span}(\mathbf{q}_1, \dots, \mathbf{q}_{k+1})$. Hence, also $\text{span}(\mathbf{q}_1, \dots, \mathbf{q}_k, \mathbf{a}) \subset \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_{k+1})$, which completes the proof. \square

The intuitive idea of (87) is that one first subtracts from \mathbf{a} its projections onto the one-dimensional subspaces defined by $\mathbf{q}_1, \dots, \mathbf{q}_k$, leaving only the component of \mathbf{a} orthogonal to $\text{span}(\mathbf{q}_1, \dots, \mathbf{q}_k)$, and then this component is normalized. In fact, one can write the second equation of (87) in the form

$$\tilde{\mathbf{q}}_{k+1} = (I - P_k) \mathbf{a},$$

where $P_k \in \mathbb{R}^{m \times m}$ is the orthogonal projection matrix onto the subspace $\text{span}(\mathbf{q}_1, \dots, \mathbf{q}_k)$; see (79).

Using Lemma 4.1, it is straightforward to compute an orthonormal basis for the subspace

$$R(A) = \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_n) \subset \mathbb{R}^m,$$

assuming the columns $\mathbf{a}_1, \dots, \mathbf{a}_n$ of the matrix $A \in \mathbb{R}^{m \times n}$ are linearly independent, i.e., assuming $N(A) = \{0\}$. Indeed, such basis $\{\mathbf{q}_1, \dots, \mathbf{q}_n\}$ can be recursively obtained via

$$\mathbf{q}_j = \frac{\tilde{\mathbf{q}}_j}{\|\tilde{\mathbf{q}}_j\|_2}, \quad \text{where} \quad \tilde{\mathbf{q}}_j = \mathbf{a}_j - \sum_{i=1}^{j-1} (\mathbf{a}_j^T \mathbf{q}_i) \mathbf{q}_i, \quad \text{for } j = 1, \dots, n.$$

In other words, one first defines \mathbf{q}_1 by simply normalizing \mathbf{a}_1 , then one computes a unit vector \mathbf{q}_2 that is orthogonal to \mathbf{q}_1 and satisfies $\text{span}(\mathbf{q}_1, \mathbf{q}_2) = \text{span}(\mathbf{q}_1, \mathbf{a}_2) = \text{span}(\mathbf{a}_1, \mathbf{a}_2)$, then one continues by computing a unit vector \mathbf{q}_3 that is orthogonal to both \mathbf{q}_1 and \mathbf{q}_2 and satisfies

$$\text{span}(\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3) = \text{span}(\mathbf{q}_1, \mathbf{q}_2, \mathbf{a}_3) = \text{span}(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3),$$

and so on until \mathbf{q}_n is computed and it holds that $\text{span}(\mathbf{q}_1, \dots, \mathbf{q}_n) = \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_n) = R(A)$.

Take note that one can get the original columns of A back via

$$(89) \quad \mathbf{a}_j = \|\tilde{\mathbf{q}}_j\|_2 \mathbf{q}_j + \sum_{i=1}^{j-1} (\mathbf{a}_j^T \mathbf{q}_i) \mathbf{q}_i, \quad j = 1, \dots, n,$$

which demonstrates that, for any $j = 1, \dots, n$, the j th column \mathbf{a}_j of A can be given as a linear combination of $\mathbf{q}_1, \dots, \mathbf{q}_j$, i.e., of (only) the first j orthonormal basis vectors of $R(A)$ produced by the Gram–Schmidt process. Defining in the standard manner $Q = [\mathbf{q}_1, \dots, \mathbf{q}_n] \in \mathbb{R}^{m \times n}$ and collecting the coefficients in the linear combinations of (89) as columns of an upper triangular matrix $R \in \mathbb{R}^{n \times n}$, the equations (89) can be written neatly in a matrix form

$$A = QR.$$

To be more precise, R can be given elementwise as

$$R_{i,j} = \begin{cases} \mathbf{a}_j^T \mathbf{q}_i & \text{if } i < j, \\ \|\tilde{\mathbf{q}}_i\|_2 & \text{if } i = j, \\ 0 & \text{if } i > j. \end{cases}$$

Note also that $Q^T Q = I \in \mathbb{R}^{n \times n}$ because the columns of Q are orthonormal.

Once a QR decomposition has been computed for A ,³ the least squares problem (65) can be formulated *equivalently* as the equation

$$R\mathbf{x} = Q^T \mathbf{b}.$$

There are (at least) two ways to deduce this. First of all, since $N(A^T A) = N(A) = \{0\}$ by assumption, the unique solution of the least squares problem (65) is the unique solution of the normal equation (66). On the other hand, since

$$A^T = R^T Q^T \quad \text{and} \quad A^T A = R^T Q^T Q R = R^T I R = R^T R,$$

the normal equation can be written as

$$R^T R \mathbf{x} = R^T Q^T \mathbf{b} \quad \iff \quad R \mathbf{x} = Q^T \mathbf{b},$$

where the last step follows by multiplying from the left by the inverse of R^T .⁴

The second way is based on using the reformulation (85) of the least squares problem and writing the orthogonal projection matrix $P \in \mathbb{R}^{m \times m}$ onto $R(A)$ with the help of the orthonormal basis $\{\mathbf{q}_1, \dots, \mathbf{q}_n\}$, that is,

$$P = Q(Q^T Q)^{-1} Q^T = Q I^{-1} Q^T = Q Q^T$$

according to (78). Hence, (85) reduces to

$$\mathbf{A}\mathbf{x} = QR\mathbf{x} = QQ^T \mathbf{b} \quad \iff \quad R\mathbf{x} = Q^T \mathbf{b},$$

where “ \Rightarrow ” follows by multiplying from left with Q^T and “ \Leftarrow ” by multiplying from left with Q .

There are two ways to numerically implement the Gram–Schmidt process, or the computation of a QR decomposition, namely classical and modified. In floating point arithmetics these two procedures have very different numerical stability properties. The more stable *modified Gram–Schmidt* procedure can be implemented in Matlab as follows:

³Observe that when introducing the QR decomposition, we assumed that the columns of $A \in \mathbb{R}^{m \times n}$ are linearly independent, i.e., $N(A) = \{0\}$, and so everything that follows is conditional on this assumption. In particular, this excludes the case $n > m$.

⁴Observe that both R and R^T are invertible: As triangular square matrices, their eigenvalues are their diagonal elements $\|\tilde{\mathbf{q}}_i\|_2 > 0$, $i = 1, \dots, n$. Hence, zero is not an eigenvalue for either of these square matrices, and thus they are both invertible.

```

function [Q,R] = my_gsmith(A)

Q = [];
for i=1:size(A,2)
    q = A(:,i);

    for k=1:size(Q,2)
        R(k,i) = q'*Q(:,k);
        q = q - R(k,i)*Q(:,k);
    end
    R(i,i) = norm(q);
    Q(:,i) = q/R(i,i);
end

```

And the *classical Gram-Schmidt* procedure as follows:

```

function [Q,R] = my_c_gsmith(A)

Q = [];
for i=1:size(A,2)
    q = A(:,i);

    for k=1:size(Q,2)
        R(k,i) = q'*Q(:,k);
    end

    for k=1:size(Q,2)
        q = q - R(k,i)*Q(:,k);
    end
    R(i,i) = norm(q);
    Q(:,i) = q/R(i,i);
end

```

In exact arithmetics, both these algorithms return $Q \in \mathbb{R}^{m \times n}$ such that $Q^T Q = I$ and $R(Q) = R(A)$ as well as an upper triangular $R \in \mathbb{R}^{n \times n}$ such that $A = QR$. However, the two implementations of the Gram–Schmidt process have very different numerical stability properties. The exact definition of numerical stability is in this case rather complicated. One can, e.g., measure the orthogonality of Q , i.e. $\|I - Q^T Q\|$, or the accuracy of the decomposition, i.e. $\|A - QR\|$, in some suitable matrix norm. The loss of orthogonality is more prone to stability issues than the accuracy of the decomposition. For the modified Gram–Schmidt, one can prove numerical stability in both of these measures, whereas the classical Gram–Schmidt is not numerically stable.

5. Singular value decomposition

The *singular value decomposition* (SVD) is a general matrix factorization that can be formed for *any* matrix. For a symmetric (or Hermitian) and positive definite matrix, the SVD coincides with the unitary diagonalization, or the eigenvalue decomposition; more generally, the eigenvalue and singular value decompositions are almost equivalent for any symmetric matrix, meaning that they can be obtained from one another by simply changing signs of certain matrix elements.

The intuitive idea behind the singular value decomposition is the following: Any matrix $A \in \mathbb{R}^{m \times n}$ can be represented as a diagonal matrix $\Sigma \in \mathbb{R}^{m \times n}$ between a certain *orthonormal* basis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ for \mathbb{R}^n and another *orthonormal* basis $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ for \mathbb{R}^m .⁵ In other words, $A\mathbf{x}$ can be evaluated by first computing the coordinates of \mathbf{x} in the orthonormal basis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$, then scaling these coordinates by certain non-negative scalars, and finally interpreting the scaled coordinates as the coordinates of $A\mathbf{x}$ in the basis $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$. If $m > n$, the coordinates/projections of $A\mathbf{x}$ corresponding to $\mathbf{u}_{n+1}, \dots, \mathbf{u}_m$ are automatically set to zero; if

⁵Here and in what follows, “orthonormal” is to be understood in the sense of the Euclidean inner product.

$n > m$, the coordinates/projections of \mathbf{x} corresponding to $\mathbf{v}_{m+1}, \dots, \mathbf{v}_n$ do not affect $A\mathbf{x}$ in any way.

Forming the SVD corresponds to finding the aforementioned orthonormal basis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ and $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$, i.e., the right-hand and left-hand *singular vectors*, as well as the corresponding ‘scaling scalars’

$$\sigma_1 \geq \sigma_2 \geq \sigma_{\min\{m,n\}} \geq 0$$

that are called the *singular values*, are arranged in decreasing order, and are the diagonal elements of $\Sigma \in \mathbb{R}^{m \times n}$. We denote by $p \leq \min\{m, n\}$ the largest index for which the corresponding singular value is positive, i.e., $\sigma_p > 0$ but $\sigma_{p+1} = 0$ (or $p = \min\{m, n\}$). Note that, according to the above explanation, all information in A is actually encoded in $\{\mathbf{v}_1, \dots, \mathbf{v}_p\}$, $\{\mathbf{u}_1, \dots, \mathbf{u}_p\}$ and $\sigma_1, \dots, \sigma_p$ because the components of \mathbf{x} in the directions $\mathbf{v}_{p+1}, \dots, \mathbf{v}_n$ are either scaled by zero singular values or altogether ignored when evaluating $A\mathbf{x}$. It will turn out that $p = \text{rank}(A) = \text{dim}(R(A))$.

Let us formulate the above ‘construction’ in a matrix form. A SVD for $A \in \mathbb{R}^{m \times n}$ is

$$(90) \quad A = U\Sigma V^T,$$

where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ are orthogonal (or unitary) matrices, i.e., their columns are the aforementioned orthonormal basis for $\mathbb{R}^{m \times m}$ and $\mathbb{R}^{n \times n}$, respectively, and $\Sigma \in \mathbb{R}^{m \times n}$ carries the singular values in the decreasing order on its diagonal. To be more precise, when $n > m$, Σ takes the form

$$\Sigma = \begin{bmatrix} \sigma_1 & & & & & \\ & \sigma_2 & & & & \\ & & \ddots & & & \\ & & & \ddots & & \\ & & & & \sigma_m & \\ & & & & & 0 & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \end{bmatrix},$$

and in the case $m > n$, Σ has the form

$$\Sigma = \begin{bmatrix} \sigma_1 & & & & & & & \\ & \sigma_2 & & & & & & \\ & & \ddots & & & & & \\ & & & \ddots & & & & \\ & & & & \sigma_n & & & \\ & & & & & 0 & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \\ & & & & & & & \end{bmatrix},$$

that is, there are extra zeros either at the right edge ($n > m$) or at the bottom ($m > n$) of Σ . The SVD is not unique: the matrices U and V can be chosen in different ways.

Let us verify that the decomposition (90) is concordant with the verbal explanation of the SVD presented in the beginning of this section: For any $\mathbf{x} \in \mathbb{R}^n$ we have

$$(91) \quad A\mathbf{x} = U\Sigma V^T \mathbf{x} = U\Sigma \begin{bmatrix} \mathbf{v}_1^T \mathbf{x} \\ \vdots \\ \mathbf{v}_n^T \mathbf{x} \end{bmatrix} = U \begin{bmatrix} \sigma_1 \mathbf{v}_1^T \mathbf{x} \\ \vdots \\ \sigma_{\min\{m,n\}} \mathbf{v}_{\min\{m,n\}}^T \mathbf{x} \\ 0 \end{bmatrix} = \sum_{i=1}^{\min\{m,n\}} \sigma_i (\mathbf{v}_i^T \mathbf{x}) \mathbf{u}_i = \sum_{i=1}^p \sigma_i \mathbf{u}_i (\mathbf{v}_i^T \mathbf{x}),$$

where in the third to last step $0 \in \mathbb{R}^{m-n}$ if $m > n$ and otherwise it should be ignored. As $\mathbf{v}_1^T \mathbf{x}, \dots, \mathbf{v}_n^T \mathbf{x}$ are the coordinates/projections of $\mathbf{x} \in \mathbb{R}^n$ in the orthonormal basis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$, the decomposition (90) is exactly of the desired form. In particular, the right-hand side of (90) demonstrates that A is completely determined by the first p singular values and singular vectors.

We still need to prove that any matrix really has a SVD. We start with a lemma that is related to the eigenvalues and eigenvectors of $A^T A$, i.e., entities that will turn out to be closely related to the SVD of A itself.

LEMMA 5.1. *Let $A \in \mathbb{R}^{m \times n}$. Let $\{\mathbf{v}_1, \dots, \mathbf{v}_n\} \subset \mathbb{R}^n$ be an orthonormal set of eigenvectors for $A^T A \in \mathbb{R}^{n \times n}$ arranged so that the corresponding eigenvalues, $\lambda_1, \dots, \lambda_n$, form a non-increasing*

sequence. Let p denote the number of those $\mathbf{v}_1, \dots, \mathbf{v}_n$ that correspond to positive eigenvalues. Then, $A\mathbf{v}_i = 0$ for $i = p + 1, \dots, n$, and

$$(92) \quad \mathbf{u}_i := \frac{1}{\sqrt{\lambda_i}} A\mathbf{v}_i, \quad i = 1, \dots, p,$$

form an orthonormal set in \mathbb{R}^m . In particular, $0 \leq p \leq \min\{m, n\}$.

PROOF. To begin with, note that there really exists an orthonormal eigenbasis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\} \subset \mathbb{R}^n$ for $A^T A \in \mathbb{R}^{n \times n}$ because it is symmetric (or Hermitian); see Theorem 2.1 of Chapter 2 and the succeeding discussion. Moreover, all eigenvalues of $A^T A$ are non-negative since $A^T A$ is positive semi-definite:

$$0 \leq \|A\mathbf{v}_i\|_2^2 = \mathbf{v}_i^T A^T A \mathbf{v}_i = \lambda_i \|\mathbf{v}_i\|_2^2 = \lambda_i$$

for any $i = 1, \dots, n$. This formula also proves that $A\mathbf{v}_i = 0$ if and only if $\lambda_i = 0$, i.e., if and only if $i = p + 1, \dots, n$; in particular, none of the vectors $\mathbf{u}_1, \dots, \mathbf{u}_p$ defined by (92) is a zero vector.

Let us then take a closer look at the (left-hand singular) vectors defined by (92). It holds that

$$\mathbf{u}_i^T \mathbf{u}_j = \frac{1}{\sqrt{\lambda_i \lambda_j}} \mathbf{v}_i^T A^T A \mathbf{v}_j = \frac{\lambda_j}{\sqrt{\lambda_i \lambda_j}} \mathbf{v}_i^T \mathbf{v}_j$$

which vanishes when $i \neq j$ due to the orthonormality of $\{\mathbf{v}_1, \dots, \mathbf{v}_p\}$. On the other hand,

$$\|\mathbf{u}_i\|_2^2 = \mathbf{u}_i^T \mathbf{u}_i = \frac{\lambda_i}{\sqrt{\lambda_i \lambda_i}} \mathbf{v}_i^T \mathbf{v}_i = \|\mathbf{v}_i\|_2^2 = 1,$$

which completes the proof of the orthonormality of $\{\mathbf{u}_1, \dots, \mathbf{u}_p\}$.

Finally, because $\{\mathbf{v}_1, \dots, \mathbf{v}_p\}$ is an orthonormal set in \mathbb{R}^n and $\{\mathbf{u}_1, \dots, \mathbf{u}_p\}$ is an orthonormal set in \mathbb{R}^m , it must hold that $p \leq n$ and $p \leq m$, i.e., $p \leq \min\{m, n\}$. \square

As a side note, observe that all positive eigenvalues of $A^T A \in \mathbb{R}^{n \times n}$ are also eigenvalues of $AA^T \in \mathbb{R}^{m \times m}$ with $\mathbf{u}_1, \dots, \mathbf{u}_p$ being the corresponding orthonormal eigenvectors:

$$AA^T \mathbf{u}_i = \frac{1}{\sqrt{\lambda_i}} A(A^T A \mathbf{v}_i) = \lambda_i \left(\frac{1}{\sqrt{\lambda_i}} A\mathbf{v}_i \right) = \lambda_i \mathbf{u}_i$$

for $i = 1, \dots, p$.

THEOREM 5.1. Let $A \in \mathbb{R}^{m \times n}$. Then there exists an integer $0 \leq p \leq \min\{m, n\}$, positive coefficients $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p > 0$ and sets of orthonormal vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_p\} \subset \mathbb{R}^n$ and $\{\mathbf{u}_1, \dots, \mathbf{u}_p\} \subset \mathbb{R}^m$ such that

$$A = \sum_{i=1}^p \sigma_i \mathbf{u}_i \mathbf{v}_i^T.$$

PROOF. Let $\{\mathbf{v}_1, \dots, \mathbf{v}_n\} \subset \mathbb{R}^n$ be an orthonormal eigenbasis for $A^T A \in \mathbb{R}^n$ organized in the same way as in Lemma 5.1. Let $\mathbf{x} \in \mathbb{R}^n$ be arbitrary and write it in the eigenbasis of $A^T A \in \mathbb{R}^n$ as

$$\mathbf{x} = \sum_{i=1}^n (\mathbf{v}_i^T \mathbf{x}) \mathbf{v}_i.$$

Because $A\mathbf{v}_i = 0$ for $i = p + 1, \dots, n$ by Lemma 5.1, we have

$$A\mathbf{x} = \sum_{i=1}^n (\mathbf{v}_i^T \mathbf{x}) A\mathbf{v}_i = \sum_{i=1}^p (\mathbf{v}_i^T \mathbf{x}) A\mathbf{v}_i = \sum_{i=1}^p (\mathbf{v}_i^T \mathbf{x}) \sqrt{\lambda_i} \mathbf{u}_i = \sum_{i=1}^p \sqrt{\lambda_i} \mathbf{u}_i \mathbf{v}_i^T \mathbf{x}$$

where the orthonormal vectors $\mathbf{u}_1, \dots, \mathbf{u}_p$ are defined by (92). After defining $\sigma_i = \sqrt{\lambda_i}$, $i = 1, \dots, p$, the proof is complete. \square

Let $0 \leq p \leq \min\{m, n\}$, $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$, $\{\mathbf{u}_1, \dots, \mathbf{u}_p\}$ and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p > 0$ be as defined in Lemma 5.1 and Theorem 5.1. Based on Theorem 5.1, the matrix form of SVD, i.e. (90), is obtained by defining $\sigma_{p+1}, \dots, \sigma_{\min\{m, n\}} = 0$ and introducing the missing orthonormal columns of U , i.e., $\mathbf{u}_{p+1}, \dots, \mathbf{u}_m$, via, e.g., the Gram–Schmidt orthogonalization process. It is easy to check

that with these choices (91) is valid for any $\mathbf{x} \in \mathbb{R}^n$ by virtue of Theorem 5.1, and thus we have proved the existence of (90) for any $A \in \mathbb{R}^{m \times n}$.

Although the ‘extra singular vectors’ $\mathbf{u}_{p+1}, \dots, \mathbf{u}_m$ and $\mathbf{v}_{p+1}, \dots, \mathbf{v}_n$ are not actually needed in constructing A (cf. Theorem 5.1), they can be used to determine $R(A)$, $N(A)$ and their orthogonal complements.

THEOREM 5.2. *Let $A \in \mathbb{R}^{m \times n}$, consider its SVD given by (90) and denote (still) by p the number of positive singular values of A . Then,*

- (1) $R(A) = \text{span}(\mathbf{u}_1, \dots, \mathbf{u}_p)$,
- (2) $N(A) = \text{span}(\mathbf{v}_{p+1}, \dots, \mathbf{v}_n)$,
- (3) $R(A)^\perp = \text{span}(\mathbf{u}_{p+1}, \dots, \mathbf{u}_m)$,
- (4) $N(A)^\perp = \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_p)$.

In particular, $\text{rank}(A) = \dim(R(A)) = p$.

PROOF. According to Theorem 5.1,

$$R(A) = \left\{ \mathbf{y} \in \mathbb{R}^m \mid \mathbf{y} = \sum_{i=1}^p \sigma_i (\mathbf{v}_i^T \mathbf{x}) \mathbf{u}_i \text{ for some } \mathbf{x} \in \mathbb{R}^n \right\}$$

which clearly is a subset of $\text{span}(\mathbf{u}_1, \dots, \mathbf{u}_p)$. On the other hand, an arbitrary element in $\text{span}(\mathbf{u}_1, \dots, \mathbf{u}_p)$, say,

$$\mathbf{y} = \sum_{i=1}^p \alpha_i \mathbf{u}_i, \quad \boldsymbol{\alpha} \in \mathbb{R}^p,$$

can be written as

$$A \left(\sum_{j=1}^p \frac{\alpha_j}{\sigma_j} \mathbf{v}_j \right) = \sum_{i=1}^p \sigma_i \mathbf{u}_i \left(\sum_{j=1}^p \frac{\alpha_j}{\sigma_j} \mathbf{v}_i^T \mathbf{v}_j \right) = \sum_{i=1}^p \sigma_i \frac{\alpha_i}{\sigma_i} \mathbf{u}_i = \mathbf{y}$$

due to the orthonormality of $\{\mathbf{v}_1, \dots, \mathbf{v}_p\}$. Hence also $\text{span}(\mathbf{u}_1, \dots, \mathbf{u}_p) \subset R(A)$, and altogether $R(A) = \text{span}(\mathbf{u}_1, \dots, \mathbf{u}_p)$. This also proves that $\text{rank}(A) = p$. Moreover, since the $m - p$ vectors $\mathbf{u}_{p+1}, \dots, \mathbf{u}_m$ are mutually orthonormal as well as orthogonal to $R(A)$, they must form a basis for the $m - p$ dimensional orthogonal complement $R(A)^\perp$.

Let then $\mathbf{x} \in N(A)$ be arbitrary and write it in the orthonormal basis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ of \mathbb{R}^n as

$$\mathbf{x} = \sum_{j=1}^n \alpha_j \mathbf{v}_j, \quad \boldsymbol{\alpha} \in \mathbb{R}^n.$$

By Theorem 5.1,

$$0 = A\mathbf{x} = A \left(\sum_{j=1}^n \alpha_j \mathbf{v}_j \right) = \sum_{i=1}^p \sigma_i \mathbf{u}_i \left(\sum_{j=1}^n \alpha_j \mathbf{v}_i^T \mathbf{v}_j \right) = \sum_{i=1}^p \sigma_i \alpha_i \mathbf{u}_i$$

due to the orthonormality of $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$. Since $\sigma_1, \dots, \sigma_p$ are positive and $\mathbf{u}_1, \dots, \mathbf{u}_p$ are linearly independent, it must hold that $\alpha_1 = \dots = \alpha_p = 0$, meaning that

$$\mathbf{x} = \sum_{j=p+1}^n \alpha_j \mathbf{v}_j \in \text{span}(\mathbf{v}_{p+1}, \dots, \mathbf{v}_n).$$

In other words, $N(A) \subset \text{span}(\mathbf{v}_{p+1}, \dots, \mathbf{v}_n)$. On the other hand, via an analogous calculation it follows easily that any $\mathbf{x} \in \text{span}(\mathbf{v}_{p+1}, \dots, \mathbf{v}_n)$ belongs to $N(A)$, and so $N(A) = \text{span}(\mathbf{v}_{p+1}, \dots, \mathbf{v}_n)$. Finally, the remaining orthonormal vectors $\mathbf{v}_1, \dots, \mathbf{v}_p$ must form a basis for the $n - (n - p) = p$ dimensional orthogonal complement $N(A)^\perp$. □

We complete these lecture notes by demonstrating the SVD can be used for solving least squares problems even if $N(A) = 0$.

THEOREM 5.3. *Let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. Consider the SVD of A given by (90) and denote (still) by p the number of positive singular values of A . Then, the vector*

$$(93) \quad \mathbf{x}^\dagger := \sum_{j=1}^p \frac{1}{\sigma_j} (\mathbf{u}_j^T \mathbf{b}) \mathbf{v}_j$$

is a solution to the normal equation (66).

PROOF. The proof is simply a straightforward computation based on the representations

$$A = \sum_{i=1}^p \sigma_i \mathbf{u}_i \mathbf{v}_i^T \quad \text{and} \quad A^T = \left(\sum_{j=1}^p \sigma_j \mathbf{u}_j \mathbf{v}_j^T \right)^T = \sum_{j=1}^p \sigma_j (\mathbf{u}_j \mathbf{v}_j^T)^T = \sum_{j=1}^p \sigma_j \mathbf{v}_j \mathbf{u}_j^T$$

provided by Theorem 5.1. Indeed,

$$\begin{aligned} A\mathbf{x}^\dagger &= \sum_{i=1}^p \sigma_i \mathbf{u}_i \mathbf{v}_i^T \left(\sum_{j=1}^p \frac{1}{\sigma_j} (\mathbf{u}_j^T \mathbf{b}) \mathbf{v}_j \right) = \sum_{i=1}^p \sigma_i \mathbf{u}_i \left(\sum_{j=1}^p \frac{1}{\sigma_j} (\mathbf{u}_j^T \mathbf{b}) (\mathbf{v}_i^T \mathbf{v}_j) \right) \\ &= \sum_{i=1}^p \frac{\sigma_i}{\sigma_i} (\mathbf{u}_i^T \mathbf{b}) \mathbf{u}_i = \sum_{i=1}^p (\mathbf{u}_i^T \mathbf{b}) \mathbf{u}_i, \end{aligned}$$

where we once again used the orthonormality of $\{\mathbf{v}_1, \dots, \mathbf{v}_p\}$. Analogously,

$$A^T(A\mathbf{x}^\dagger) = \sum_{j=1}^p \sigma_j \mathbf{v}_j \mathbf{u}_j^T \left(\sum_{i=1}^p (\mathbf{u}_i^T \mathbf{b}) \mathbf{u}_i \right) = \sum_{j=1}^p \sigma_j \mathbf{v}_j \left(\sum_{i=1}^p (\mathbf{u}_i^T \mathbf{b}) (\mathbf{u}_j^T \mathbf{u}_i) \right) = \sum_{j=1}^p \sigma_j (\mathbf{u}_j^T \mathbf{b}) \mathbf{v}_j,$$

which equals

$$A^T \mathbf{b} = \sum_{j=1}^p \sigma_j \mathbf{v}_j \mathbf{u}_j^T \mathbf{b} = \sum_{j=1}^p \sigma_j (\mathbf{u}_j^T \mathbf{b}) \mathbf{v}_j,$$

and so the proof is complete. \square

COROLLARY 5.1. *Let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. Consider the SVD of A given by (90) and denote (still) by p the number of positive singular values of A . All solutions of the least squares problem (65) are of the form*

$$(94) \quad \mathbf{x} = \mathbf{x}^\dagger + \mathbf{z}$$

for some $\mathbf{z} \in N(A)$.

PROOF. Since we have (finally!) found a solution \mathbf{x}^\dagger to the normal equation (66) independently of whether $N(A) = \{0\}$, by virtue of Theorem 1.1 we know that \mathbf{x} is a solution of the least squares problem (65) if and only if it solves the normal equation (66). In consequence, it is enough to prove that all solutions of the normal equation (66) are of the form (94).

Let \mathbf{x} be of the form (94). Obviously,

$$A^T A \mathbf{x} = A^T A \mathbf{x}^\dagger + A^T A \mathbf{z} = A^T \mathbf{b} + 0 = A^T \mathbf{b},$$

and so \mathbf{x} solves the normal equation (66). On the other hand, if \mathbf{x} is an arbitrary solution of the normal equation, we may define $\mathbf{z} = \mathbf{x} - \mathbf{x}^\dagger$, i.e., $\mathbf{x} = \mathbf{x}^\dagger + \mathbf{z}$. It follows that

$$A^T A \mathbf{z} = A^T A \mathbf{x} - A^T A \mathbf{x}^\dagger = A^T \mathbf{b} - A^T \mathbf{b} = 0.$$

Hence, $\mathbf{z} \in N(A^T A) = N(A)$ and so \mathbf{x} is of the form (94). This completes the proof. \square

According to Theorems 5.2 and 5.3, the special least squares solution \mathbf{x}^\dagger belongs to $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_p) = N(A)^\perp$. Hence, any (other) solution (94) to the least squares problem (65) satisfies

$$\|\mathbf{x}\|_2^2 = \|\mathbf{x}^\dagger + \mathbf{z}\|_2^2 = \|\mathbf{x}^\dagger\|_2^2 + 2\mathbf{z}^T \mathbf{x}^\dagger + \|\mathbf{z}\|_2^2 = \|\mathbf{x}^\dagger\|_2^2 + \|\mathbf{z}\|_2^2 \geq \|\mathbf{x}^\dagger\|_2^2,$$

and so \mathbf{x}^\dagger is the solution to (65) with the smallest Euclidean norm, i.e. the *unique* so-called *minimum norm solution* to (65).

