

Imperial College
London

LECTURE NOTES

IMPERIAL COLLEGE LONDON

DEPARTMENT OF COMPUTING

Mathematical Methods

Instructors:

Marc Deisenroth, Mahdi Cheraghchi

Version: November 30, 2016

Contents

1	Course Overview	3
2	Sequences	6
2.1	The Convergence Definition	6
2.2	Illustration of Convergence	7
2.3	Common Converging Sequences	8
2.4	Combinations of Sequences	9
2.5	Sandwich Theorem	9
2.6	Ratio Tests for Sequences	11
2.7	Proof of Ratio Tests	12
2.8	Useful Techniques for Manipulating Absolute Values	13
2.9	Properties of Real Numbers	14
3	Series	15
3.1	Geometric Series	16
3.2	Harmonic Series	16
3.3	Series of Inverse Squares	17
3.4	Common Series and Convergence	18
3.5	Convergence Tests	19
3.6	Absolute Convergence	22
3.7	Power Series and the Radius of Convergence	23
3.8	Proofs of Ratio Tests	24
4	Power Series	27
4.1	Basics of Power Series	27
4.2	Maclaurin Series	27
4.3	Taylor Series	31
4.4	Taylor Series Error Term	33
4.5	Deriving the Cauchy Error Term	34
4.6	Power Series Solution of ODEs	34
5	Linear Algebra	36
5.1	Linear Equation Systems	38
5.2	Groups	39
5.2.1	Definitions	39
5.2.2	Examples	40

5.3	Matrices	40
5.3.1	Matrix Multiplication	40
5.3.2	Inverse and Transpose	42
5.3.3	Multiplication by a Scalar	43
5.3.4	Compact Representations of Linear Equation Systems	43
5.4	Gaussian Elimination	43
5.4.1	Example: Solving a Simple Linear Equation System	44
5.4.2	Elementary Transformations	45
5.4.3	The Minus-1 Trick for Solving Homogeneous Equation Systems	50
5.4.4	Applications of Gaussian Elimination in Linear Algebra	51
5.5	Vector Spaces	53
5.5.1	Examples	54
5.5.2	Generating Set and Vector Subspaces	55
5.6	Linear (In)Dependence	56
5.7	Basis and Dimension	60
5.7.1	Rank	64
5.8	Intersection of Subspaces	65
5.9	Linear Mappings	67
5.9.1	Image and Kernel (Null Space)	68
5.9.2	Matrix Representation of Linear Mappings	70
5.9.3	Basis Change	73
5.10	Determinants	77
5.11	Eigenvalues	79
5.11.1	Characteristic Polynomial	81
5.11.2	Example: Eigenspace Computation	81
5.11.3	Eigenvalues in Practical Applications	82
5.12	Diagonalization	84
5.12.1	Applications	88
5.12.2	Cayley-Hamilton Theorem	88
5.13	Scalar Products	89
5.13.1	Lengths, Distances, Orthogonality	90
5.13.2	Applications	91
5.14	Projections	92
5.14.1	Projection onto a Line	93
5.14.2	Projection onto General Subspaces	95
5.14.3	Applications	97
5.15	Affine Subspaces	98
5.15.1	Intersection of Affine Subspaces	99
5.16	Affine Mappings	104
5.17	Rotations	104
5.17.1	Rotations in the Plane	105
5.17.2	Rotations in Three Dimensions	105
5.17.3	Rotations in n Dimensions	107
5.17.4	Properties of Rotations	107

Preface

These lecture notes are intended for undergraduate students of computer science at Imperial College London and cover basic mathematical concepts that are required in other courses. The lecture notes are based on Jeremy Bradley's lecture from 2014/15, but the course has been partially restructured since then.

Many people have helped improving these lecture notes by providing comments, suggestions and corrections. In particular, we want to thank Romain Barnoud, Ruhi Choudhury and Tony Field for valuable feedback and additions.

Marc Deisenroth and Mahdi Cheraghchi
October 2016

Chapter 1

Course Overview

This introductory course covers background that is essential to many other courses in the Department of Computing. We will approximately split this course into two parts: Analysis and Linear Algebra. In Analysis, we will cover topics, including power series and convergence of sequences and series. The part on Linear Algebra will cover linear equation systems (and their solution), linear mappings and their matrix representations, vector spaces, eigenvalues, affine mappings, and applications, e.g., projections and rotations.

There are many courses in the department that require an understanding of mathematical concepts. These courses range from optimization, operations research and complexity to machine learning, robotics, graphics and computer vision.

The department offers additional courses that cover other and more advanced mathematical concepts: Probability and Statistics (2nd year), Computational Techniques (2nd year), Mathematics for Inference and Machine Learning (4th year).

Figures 1.1–1.2 give an overview of the relevance of this course to other courses in the Department of Computing. Figure 1.1 details the topics from this course that are relevant in other courses, whereas Figure 1.2 shows which courses in Years 2–4 rely on material covered in this introductory course.

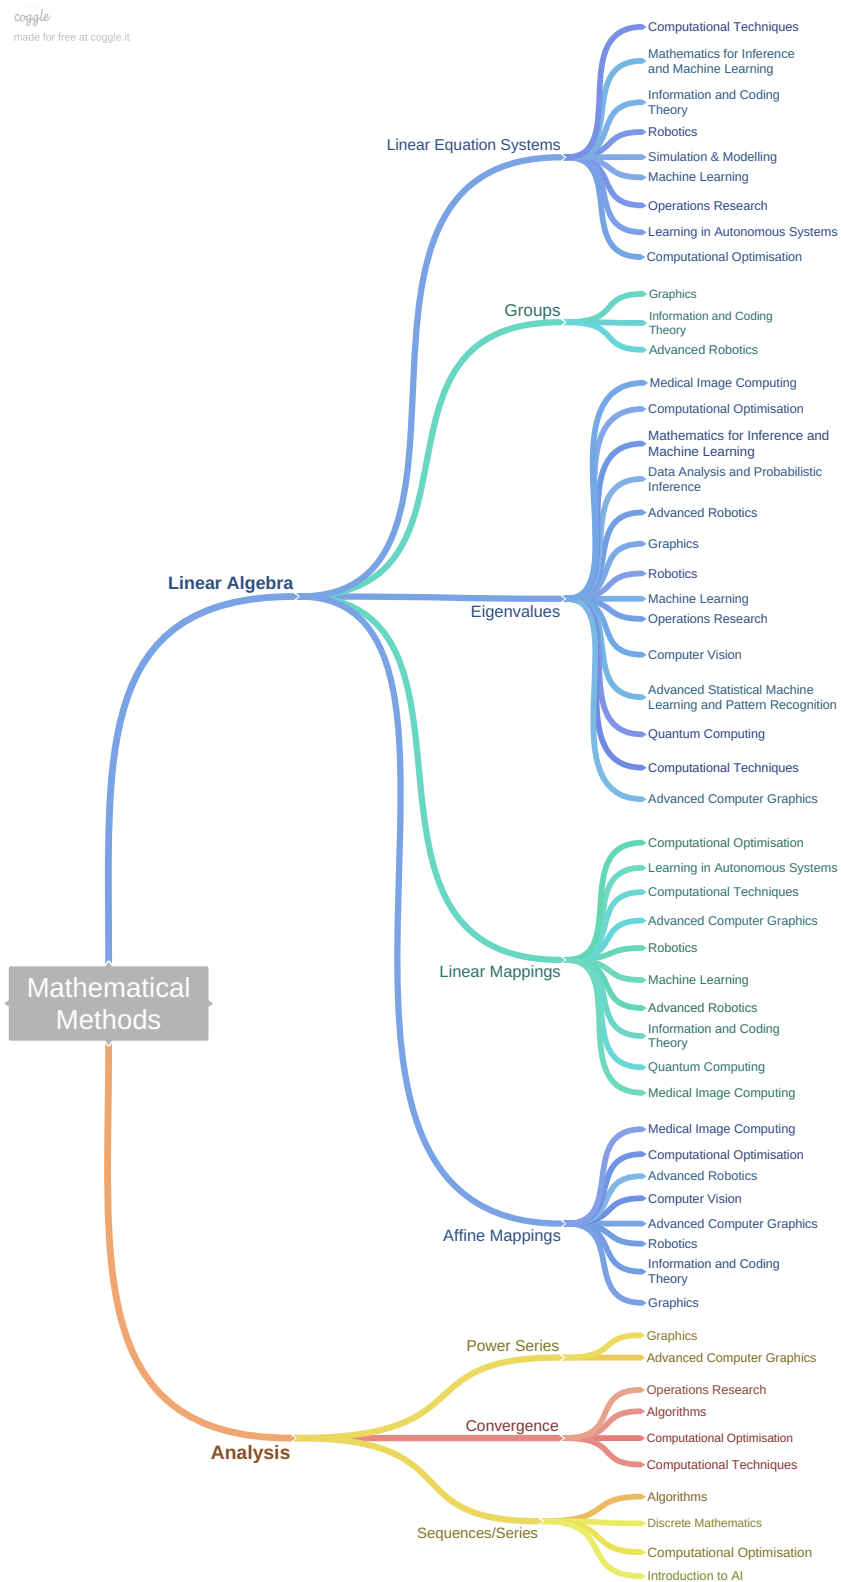


Figure 1.1: Overview of the relevance of topics covered in this course to other courses in the Department of Computing.

coggle
made for free at coggle.it

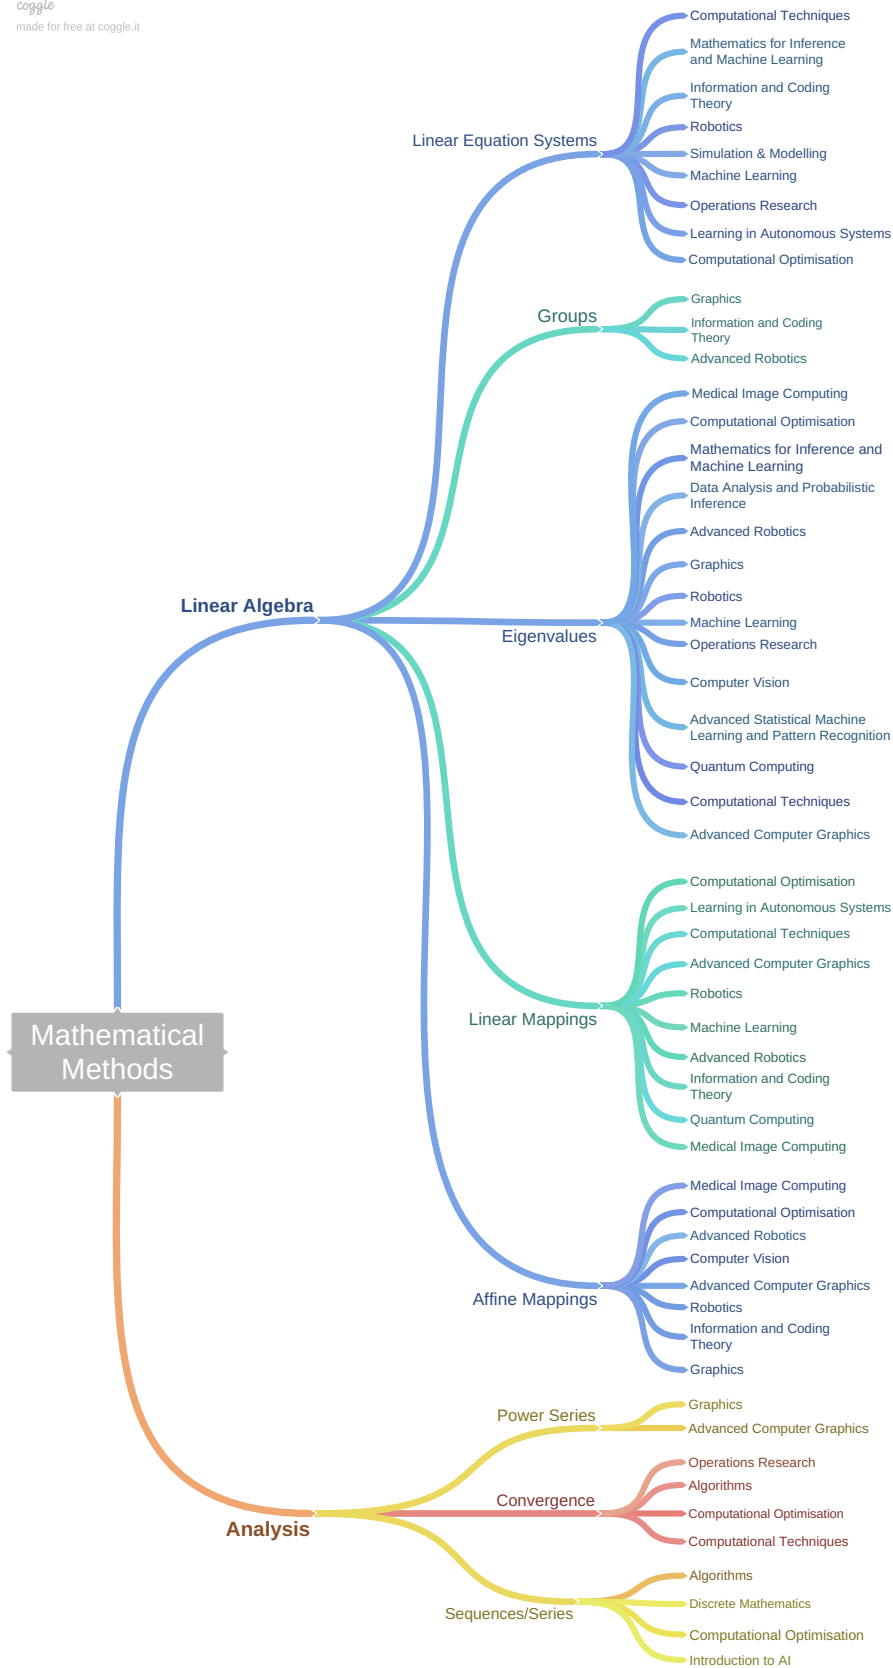


Figure 1.2: Overview of courses (ordered by year) in the Department of Computing that rely on topics covered in this course.

Chapter 2

Sequences

2.1 The Convergence Definition

Given a sequence of real numbers a_1, a_2, \dots , we would like to have a precise definition of what it means for this sequence to converge to a limit l . You may already have a good intuitive understanding of what it means for $a_n \rightarrow l$ as $n \rightarrow \infty$ or $\lim_{n \rightarrow \infty} a_n = l$. Below is the mathematical definition that formally represents the concept.

Definition A sequence a_n for $n \geq 1$ converges to a limit $l \in \mathbb{R}$, written $a_n \rightarrow l$ as $n \rightarrow \infty$, if and only if the following statement can be shown to be true:

For all $\epsilon > 0$, we can always find a positive integer N , such that, for all $n > N$:

$$|a_n - l| < \epsilon$$

This statement is best understood by taking the last part first and working backwards.

The limit inequality $|a_n - l| < \epsilon$ is called the *limit inequality* and says that the distance from the n th term in the sequence to the limit should be less than ϵ .

For all $n > N$... There should be some point N after which all the terms a_n in the sequence are within ϵ of the limit.

For all $\epsilon > 0$... No matter what value of ϵ we pick, however tiny, we will still be able to find some value of N such that all the terms to the right of a_N are within ϵ of the limit.

How do we prove convergence So that is the definition, but it also tells us how to prove that a sequence converges. We need to find the N for any given value of ϵ . If we can find this quantity then we have a rigorous proof of convergence for a given sequence. We have an example below for $a_n = 1/n$.

2.2 Illustration of Convergence

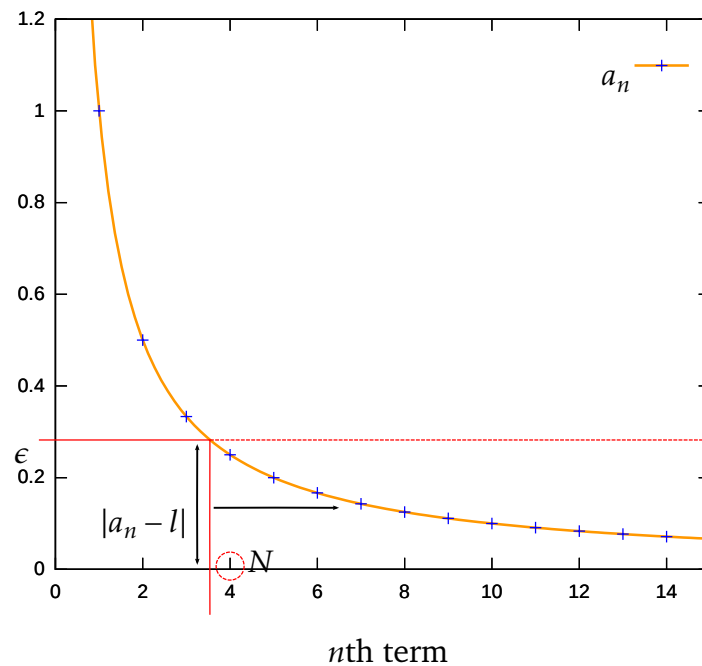


Figure 2.1: The convergence of the sequence $a_n = 1/n$ for $n \geq 1$ tending to 0

A demonstration of the convergence of $a_n = 1/n$ for $n \geq 1$ is shown in Figure 2.1. The *game* is to find a value of N for whatever value of ϵ is chosen. N represents the point after which all a_n will be within ϵ of the limit. Clearly the smaller the value of ϵ chosen, the further to the right in Figure 2.1 we will have to go to achieve this, thus the larger the value of N we will have to choose.

Applying the limit inequality In the example in Figure 2.1 we have taken a particular value of ϵ to demonstrate the concepts involved. Taking $\epsilon = 0.28$ (picked arbitrarily), we apply the $|a_n - l| < \epsilon$ inequality, to get:

$$\begin{aligned} |1/n| &< 0.28 \\ 1/n &< 0.28 \\ n &> 1/0.28 \\ n &> 3.57 \end{aligned}$$

This states that for $n > 3.57$ the continuous function $1/n$ will be less than 0.28 from the limit. However, we are dealing with a sequence a_n which is only defined at integer points on this function. Hence we need to find the N th term, a_N , which is the first point in the sequence to also be less than this value of ϵ . In this case we can see that the next term above 3.57, that is $N = 4$, satisfies this requirement.

For all $n > N$ We are further required to ensure that all other points to the right of a_N in the sequence a_n for $n > N$ are also within 0.28 of the limit. We know this to be

true because the condition on n we originally obtained was $n > 3.57$, so we can be sure that $a_N, a_{N+1}, a_{N+2}, \dots$ are all closer to the limit than 0.28.

For all $\epsilon > 0$ The final point to note is that it is not sufficient to find N for just a single value of ϵ . We need to find N for every value of $\epsilon > 0$. Since N will vary with ϵ as mentioned above, we are clearly required to find a function $N(\epsilon)$.

In the case of $a_n = 1/n$, this is straightforward. We apply the limit inequality in general and derive a condition for n .

$$\frac{1}{n} < \epsilon \quad \Leftrightarrow \quad n > \frac{1}{\epsilon}$$

We are looking for a greater-than inequality. So long as we get one, we can select the next largest integer value; we use the ceiling function to do this, giving:

$$N(\epsilon) = \left\lceil \frac{1}{\epsilon} \right\rceil$$

Now whatever value of ϵ we are given, we can find a value of N using this function and we are done.

2.3 Common Converging Sequences

It is very useful to have an intuition about certain common sequences and whether they converge. Below is a list of useful results, which can be proved directly. I would recommend you get practice at direct proof on some of the sequences below.

1. $a_n = \frac{1}{n} \rightarrow 0$, also $a_n = \frac{1}{n^2} \rightarrow 0, a_n = \frac{1}{\sqrt{n}} \rightarrow 0$
2. In general, $a_n = \frac{1}{n^c} \rightarrow 0$ for some positive real constant $c > 0$
3. $a_n = \frac{1}{2^n} \rightarrow 0$, also $a_n = \frac{1}{3^n} \rightarrow 0, a_n = \frac{1}{e^n} \rightarrow 0$
4. In general, $a_n = \frac{1}{c^n} \rightarrow 0$ for some real constant $|c| > 1$;
or equivalently $a_n = c^n \rightarrow 0$ for some real constant $|c| < 1$
5. $a_n = \frac{1}{n!} \rightarrow 0$ (Hard to prove directly, easier to use a ratio test in Section 2.6).
6. $a_n = \frac{1}{\log n} \rightarrow 0$ for $n > 1$

2.4 Combinations of Sequences

Given sequences a_n and b_n which converge to limits a and b respectively, then

1. $\lim_{n \rightarrow \infty} (\lambda a_n) = \lambda a$ for some real constant λ
2. $\lim_{n \rightarrow \infty} (a_n + b_n) = a + b$
3. $\lim_{n \rightarrow \infty} (a_n - b_n) = a - b$
4. $\lim_{n \rightarrow \infty} (a_n b_n) = ab$
5. $\lim_{n \rightarrow \infty} \frac{a_n}{b_n} = \frac{a}{b}$ provided that $b \neq 0$

We proved rule (2) in lectures, you can find proofs of the other results in “Analysis: an introduction to proof” by Stephen Lay.

Example These sequence combination results are useful for being able to understand convergence properties of combined sequences. As an example, if we take the sequence:

$$a_n = \frac{3n^2 + 2n}{6n^2 + 7n + 1}$$

We can transform this sequence into a combination of sequences that we know how to deal with, by dividing numerator and denominator by n^2 .

$$\begin{aligned} a_n &= \frac{3 + \frac{2}{n}}{6 + \frac{7}{n} + \frac{1}{n^2}} \\ &= \frac{b_n}{c_n} \end{aligned}$$

where $b_n = 3 + \frac{2}{n}$ and $c_n = 6 + \frac{7}{n} + \frac{1}{n^2}$. We know from rule (5) above that if $b_n \rightarrow b$ and $c_n \rightarrow c$, then $\frac{b_n}{c_n} \rightarrow \frac{b}{c}$. This means that we can investigate the convergence of b_n and c_n separately.

Similarly we can rewrite b_n as $3 + d_n$, for $d_n = \frac{2}{n}$. We know that $d_n = \frac{2}{n} \rightarrow 0$ as it is one of the common sequence convergence results, thus $b_n \rightarrow 3$ by rule (2) above. By a similar argument, we can see that $c_n \rightarrow 6$ using composition rule (2) and common results.

Finally, we get the result that $a_n \rightarrow \frac{1}{2}$.

2.5 Sandwich Theorem

The Sandwich theorem is an alternative way, and often a simpler way, of proving that a sequence converges to a limit. In order to use the Sandwich theorem, you need to have two sequences (one of which can be constant) which bound the sequence

you wish to reason about. The two bounding sequences should be given as known results, or proved to converge prior to use. The bounding sequences should be easier to reason about than the sequence to be reasoned about.

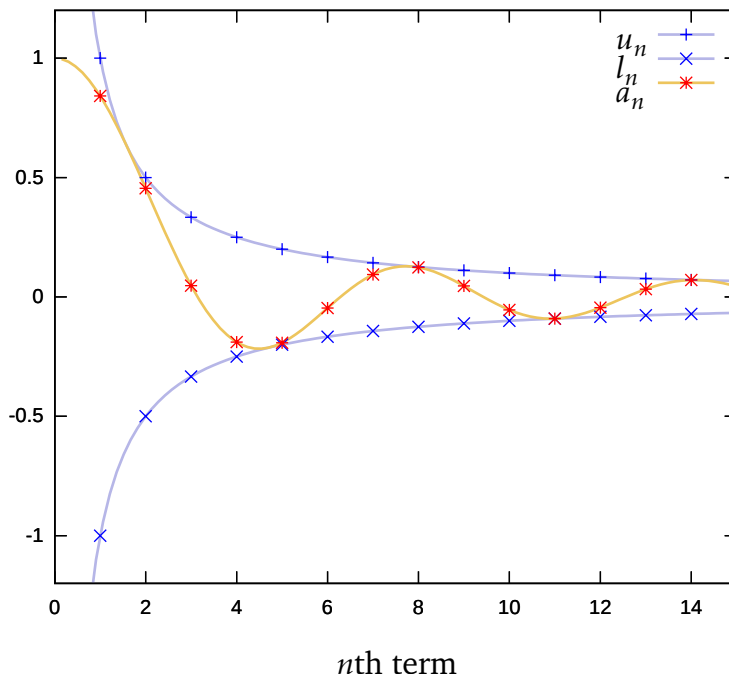


Figure 2.2: Demonstration of the Sandwich theorem. The sequence a_n is sandwiched between two bounding sequences u_n and l_n which converge to the same limit.

Sandwich Theorem Given sequences u_n and l_n for $n \geq 1$, where both $u_n \rightarrow l$ and $l_n \rightarrow l$ as $n \rightarrow \infty$.

If we can show that a third sequence a_n satisfies:

$$l_n \leq a_n \leq u_n \text{ for all } n \geq N \text{ for some } N \in \mathbb{N}$$

Then $a_n \rightarrow l$ as $n \rightarrow \infty$.

Demonstration We provide an illustration of the sandwich theorem in Figure 2.2. Here, we have a sequence $a_n = \frac{\sin n}{n}$ which oscillates around its limit of 0. We construct two sandwiching sequences $u_n = 1/n$ and $l_n = -1/n$. Since we know (or can easily prove) that u_n and l_n both tend to 0, we only need to show that $l_n \leq a_n$ and $a_n \leq u_n$ from some point $n > N$. In this case we can show it for all $n > 0$.

In this instance, we can prove this in one line and refer to the property of the sin function.

$$-\frac{1}{n} \leq \frac{\sin n}{n} \leq \frac{1}{n} \quad \Rightarrow \quad -1 \leq \sin n \leq 1$$

which is true for all n by the range of $\sin n$.

Proof of sandwich theorem We know that l_n and u_n tend to the same limit l . We use the definition of convergence for both these sequences.

Given some $\epsilon > 0$, we can find an N_1 and N_2 such that for all $n > N_1$, $|l_n - l| < \epsilon$ and for all $n > N_2$, $|u_n - l| < \epsilon$. If we want both sequences to be simultaneously less than ϵ from the limit l for a single value of N' , then we have to pick $N' = \max(N_1, N_2)$.

So for $N' = \max(N_1, N_2)$, for all $n > N'$, we know that both:

$$|l_n - l| < \epsilon \text{ and } |u_n - l| < \epsilon$$

Therefore, removing the modulus sign and using the fact that $l_n < u_n$:

$$l - \epsilon < l_n < u_n < l + \epsilon$$

We also have proved in the course of applying the sandwich theorem that $l_n \leq a_n \leq u_n$ for all $n > N$. Hence, for all $n > \max(N, N')$:

$$\begin{aligned} l - \epsilon &< l_n \leq a_n \leq u_n < l + \epsilon \\ l - \epsilon &< a_n < l + \epsilon \\ -\epsilon &< a_n - l < \epsilon \\ \Rightarrow |a_n - l| &< \epsilon \end{aligned}$$

This allows us to say that a_n converges to l as a result.

2.6 Ratio Tests for Sequences

A third technique (and a very simple one) for proving sequence convergence is called the ratio test. So-called as it compares the ratio of consecutive terms in the sequence. The test verifies whether a sequence converges to 0 or diverges. If a sequence b_n has a non-zero limit, l , then the ratio test can still be applied to a modified sequence $a_n = b_n - l$, which will converge to 0 (if it passes the test!).

Ratio convergence test If $\left| \frac{a_{n+1}}{a_n} \right| \leq c < 1$ for some $c \in \mathbb{R}$ and for all sufficiently large n (i.e., for all $n \geq N$ for some integer N), then $a_n \rightarrow 0$ as $n \rightarrow \infty$.

Ratio divergence test If $\left| \frac{a_{n+1}}{a_n} \right| \geq c > 1$ for some $c \in \mathbb{R}$ and for all sufficiently large n (i.e., for all $n \geq N$ for some integer N), then a_n diverges as $n \rightarrow \infty$.

Example of convergence ratio test As a simple example of the ratio test in use we can look at $a_n = 3^{-n}$, then $\left| \frac{a_{n+1}}{a_n} \right| = \left| \frac{3^{-(n+1)}}{3^{-n}} \right| = \frac{1}{3} < 1$. So we can say that $a_n = 3^{-n}$ converges to 0 as $n \rightarrow \infty$.

Limit ratio test The limit ratio test is a direct corollary of the standard ratio tests. Instead of looking at the ratio of consecutive terms in the sequence, we look at the limit of that ratio. This is useful when the standard test is not conclusive and leaves terms involving n in the ratio expression. To use this test, we compute:

$$r = \lim_{n \rightarrow \infty} \left| \frac{a_{n+1}}{a_n} \right|$$

If $r < 1$ the a_n sequence converges to 0, and if $r > 1$, it diverges

Example of limit ratio test An example of the limit ratio test, we consider the sequence $a_n = \frac{1}{n!}$. Taking $\left| \frac{a_{n+1}}{a_n} \right| = \left| \frac{1/(n+1)!}{1/n!} \right| = \frac{1}{n+1}$. This is not a conclusive result by the normal ratio test (although clearly less than 1, there is no constant that sits between it and 1). So we consider the limit $r = \lim_{n \rightarrow \infty} \frac{1}{n+1} = 0$. Hence a_n converges to 0.

2.7 Proof of Ratio Tests

For completeness, I will include proofs of the ratio tests above. If you are finding the material fairly straightforward then you may enjoy looking at these. I do not expect you to be able to reproduce these arguments in an exam although I do expect you to be able to apply the tests in question.

Proof of convergence ratio test In the case where $\left| \frac{a_{n+1}}{a_n} \right| \leq c < 1$ for $n \geq N$, we can take the left hand inequality and multiply by $|a_n|$ to get $|a_{n+1}| \leq c|a_n|$ for $n \geq N$. Also $|a_n| \leq c|a_{n-1}|$ and so on. In fact we can create a chain of inequalities down to the point $n = N$, where the property is first true:

$$|a_n| \leq c|a_{n-1}| \leq c^2|a_{n-2}| \cdots \leq c^{n-N}|a_N|$$

Thus:

$$\begin{aligned} |a_n| &\leq c^{n-N}|a_N| \\ &\leq kc^n \text{ where } k = \frac{|a_N|}{c^N} \text{ is a constant} \end{aligned}$$

As $c < 1$, we know that a new sequence $b_n = kc^n$ converges to 0 (standard result (4) from Section 2.3) and we have bounded the sequence a_n between b_n and $-b_n$. Thus by the sandwich theorem, a_n must also converge to 0.

A very similar proof can be constructed to show the divergence of a_n should $\left| \frac{a_{n+1}}{a_n} \right| \geq c > 1$ for $n \geq N$.

Proof of limit ratio test Justifying that this works demonstrates an interesting (and new) use of the limit definition. We define a new sequence, $b_n = \left| \frac{a_{n+1}}{a_n} \right|$, and we can consider the limit of this new sequence, r .

As with all sequences, we can say that for all $\epsilon > 0$, we can find an N such that for all $n > N$, $|b_n - r| < \epsilon$.

As before we can expand the modulus to give:

$$r - \epsilon < b_n < r + \epsilon \quad (*)$$

Now we take the definition of the limit literally and pick a carefully chosen value for $\epsilon > 0$ in both cases.

In the case where $r < 1$ (remember that this corresponds to a sequence that we wish to show converges), we choose $\epsilon = \frac{1-r}{2} > 0$. Equation (*) becomes:

$$\begin{aligned} r - \frac{1-r}{2} < b_n < r + \frac{1-r}{2} \\ \frac{3r-1}{2} < b_n < \frac{r+1}{2} \end{aligned}$$

Taking the right hand side, we see that $b_n < \frac{r+1}{2}$ and since $r < 1$, we can show that $\frac{r+1}{2} < 1$ also. If we take $c = \frac{r+1}{2} < 1$, then by the original ratio test, we can say that the original sequence a_n must converge.

In the case where $r > 1$ (remember that this corresponds to a sequence that we wish to show diverges), we choose $\epsilon = \frac{r-1}{2} > 0$. Equation (*) becomes:

$$\begin{aligned} r - \frac{r-1}{2} < b_n < r + \frac{r-1}{2} \\ \frac{r+1}{2} < b_n < \frac{3r-1}{2} \end{aligned}$$

In this case, we take the left hand side of the inequality, $b_n > \frac{r+1}{2}$ and since $r > 1$, we can show that $\frac{r+1}{2} > 1$ also. If we now take $c = \frac{r+1}{2} > 1$, then by the original ratio test, we can say that the original sequence a_n must diverge.

2.8 Useful Techniques for Manipulating Absolute Values

Many of the reasoning techniques behind limits require the ability to manipulate the modulus or absolute values of expressions. There are some properties it is useful to know. These will come in handy for other topics, including series and radius of convergence reasoning.

A commonly used decomposition of the absolute value operator is:

$$|x| < a \quad \Leftrightarrow \quad -a < x < a$$

Where we are combining absolute values of expressions, the following are useful, where x, y as real numbers:

1. $|xy| = |x| \times |y|$

2. $\left| \frac{x}{y} \right| = \frac{|x|}{|y|}$
3. $|x + y| \leq |x| + |y|$, the triangle inequality
4. $|x - y| \geq ||x| - |y||$

2.9 Properties of Real Numbers

Now we have looked at limits of sequences of real numbers, we will see how this relates to some fundamental properties of sets of real numbers, which will lead us to describe the fundamental axiom of analysis.

We will cover some definitions of basic concepts for sets of real numbers: If S is a set of real numbers.

1. u is an upper bound on S if $u \geq s$ for any $s \in S$
2. l is a lower bound on S if $l \leq s$ for any $s \in S$
3. we say that S is bounded above if S has an upper bound
4. we say that S is bounded below if S has a lower bound
5. we say that S is bounded if S has an upper and lower bound

A set S can have many upper or lower bounds (or none), so:

1. we call $\sup(S)$ the least upper bound of S
2. we call $\inf(S)$ the greatest lower bound of S

By convention, if a set S has no upper bound, e.g., $S = \{x \mid x > 0\}$, then $\sup(S) = \infty$. Similarly if S has no lower bound then $\inf(S) = -\infty$.

The fundamental axiom of analysis states that every increasing sequence of real numbers that is bounded above, must converge.

Chapter 3

Series

Series are sums of sequences. Whether they converge or not is an important factor in determining whether iterative numerical algorithms terminate.

An infinite series is a summation of the form:

$$S = \sum_{n=1}^{\infty} a_n$$

for some real sequence a_n for $n \geq 1$.

Partial sums One easy way of determining the convergence of series is to construct the partial sum – the sum up to the n th term of the series:

$$S_n = \sum_{i=1}^n a_i$$

S_n itself is a sequence and can be treated as such. A series converges or diverges if and only if its sequence of partial sums converges or diverges.

Bounded above When $a_i > 0$ for all i , then the partial sum S_n is an increasing sequence, that is:

$$S_1 < S_2 < S_3 < \dots$$

If you can show that the sequence is bounded above, then as with other increasing sequences, it must converge to some limit.

Series terms tend to 0 For a series $S = \sum_{n=1}^{\infty} a_n$ to have any chance of converging, the terms a_n have to converge to 0 as $n \rightarrow \infty$. If a_n diverges or tends to a non-zero limit then the series S will diverge. Note that $a_n \rightarrow 0$ is not sufficient to ensure that S itself converges (otherwise we would not need the rest of this section).

The tail of a series It is important to understand that, as with sequences, the convergence behaviour of a series lies in the tail and we do not really care about transient values of early terms such as a_1, a_2, a_3 . If $S = \sum_{n=1}^{\infty} a_n$ converges or diverges then so will $\sum_{n=10}^{\infty} a_n$ or $\sum_{n=1000}^{\infty} a_n$ or in general $\sum_{n=N}^{\infty} a_n$ for any natural number N . This explains why, in comparison tests, we need only worry about later terms in the series from a certain point N , rather than the whole series.

3.1 Geometric Series

This is an example of the use of partial sums to determine the convergence of a series. A common and well-known example is the *geometric series*

$$G = \sum_{n=1}^{\infty} x^n$$

If G does exist, $G = x + \sum_{n=2}^{\infty} x^n = x + x \sum_{n=1}^{\infty} x^n = x + xG$ so:

$$G = \frac{x}{1-x}$$

How can we determine when G does exist? When the series and thus the sequence of partial sums converges. In this case, we can find an expression of the n th partial sum of the geometric series. Similarly, n th partial sum $G_n = x + \sum_{i=2}^n x^i = x + x \sum_{i=1}^{n-1} x^i = x + x(G_n - x^n)$, so:

$$G_n = \frac{x - x^{n+1}}{1-x}$$

G_n is a sequence, so we can determine if G_n converges or diverges in the limit, taking care to take into account behaviour that may depend on x . The x^{n+1} term is the only varying term in the sequence, so from our knowledge of sequences we can say that this and thus G_n will converge if $|x| < 1$.

$$G_n \rightarrow \frac{x}{1-x} \text{ as } n \rightarrow \infty \text{ for } |x| < 1$$

by rules of convergence for sequences.

Similarly, for $|x| > 1$, G_n diverges as $n \rightarrow \infty$. For $x = 1$, $G_n = n$ which also diverges.

3.2 Harmonic Series

This is a very important result as it is used to show that other series diverge by using the comparison test, described later.

The harmonic series is given by

$$S = \sum_{n=1}^{\infty} \frac{1}{n} = 1 + \frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \frac{1}{5} + \dots$$

It is an important result to show that it diverges to ∞ .

By grouping terms in the series, we can see this intuitively:

$$S = 1 + \frac{1}{2} + \underbrace{\left(\frac{1}{3} + \frac{1}{4}\right)}_{> \frac{1}{4} + \frac{1}{4}} + \underbrace{\left(\frac{1}{5} + \frac{1}{6} + \frac{1}{7} + \frac{1}{8}\right)}_{> \frac{1}{8} + \frac{1}{8} + \frac{1}{8} + \frac{1}{8}} + \cdots > 1 + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \cdots$$

So considering partial sums, we get that

$$S_n = \sum_{i=1}^{2^n} \frac{1}{i} > 1 + \frac{n}{2}$$

Clearly $1 + \frac{n}{2}$ diverges and since S_n is greater than this sequence, it must also diverge as $n \rightarrow \infty$ and so does S .

3.3 Series of Inverse Squares

Similarly to the harmonic series, this result is important since it is used to show that other series converge by comparison.

In order to understand the convergence of $S = \sum_{n=1}^{\infty} \frac{1}{n^2}$, the series of inverse squares, we must first consider the limit of the series:

$$T = \sum_{i=1}^{\infty} \frac{1}{i(i+1)}$$

Using partial fractions, we can rewrite $\frac{1}{i(i+1)}$ as $\frac{1}{i} - \frac{1}{i+1}$. The the n th partial sum of T can therefore be written as:

$$T_n = \sum_{i=1}^n \left(\frac{1}{i} - \frac{1}{i+1} \right) = 1 - \frac{1}{n+1} \quad (3.1)$$

(To see why the right hand side is true, try writing out the first and last few terms of the sum and observing which cancel). From this, we see that T_n converges giving $T = 1$.

We will use the partial sum result to gives us a handle on being able to understand the convergence or divergence of:

$$S = \sum_{n=1}^{\infty} \frac{1}{n^2}$$

We start by considering terms of the partial sum $S_n = \sum_{i=1}^n \frac{1}{i^2}$, and we notice that:

$$\frac{1}{i(i+1)} < \frac{1}{i^2} < \frac{1}{(i-1)i} \text{ for } i \geq 2$$

We sum from $i = 2$ (to avoid a 0 denominator on the right hand side) to n to get:

$$\frac{1}{2 \cdot 3} + \frac{1}{3 \cdot 4} + \frac{1}{4 \cdot 5} + \cdots + \frac{1}{n(n+1)} < \sum_{i=2}^n \frac{1}{i^2} < \frac{1}{2 \cdot 1} + \frac{1}{3 \cdot 2} + \frac{1}{4 \cdot 3} + \cdots + \frac{1}{n(n-1)}$$

We note that the left hand and right hand series above differ only in the first and last terms, so can be rewritten as:

$$\left(\sum_{i=1}^n \frac{1}{i(i+1)} \right) - \frac{1}{2} < \sum_{i=2}^n \frac{1}{i^2} < \sum_{i=1}^{n-1} \frac{1}{i(i+1)}$$

Now we add 1 across the whole inequality to allow us to sum $\frac{1}{i^2}$ from 1 and get the required partial sum S_n as our middle expression:

$$\frac{1}{2} + \sum_{i=1}^n \frac{1}{i(i+1)} < S_n < 1 + \sum_{i=1}^{n-1} \frac{1}{i(i+1)}$$

This effectively bounds the sequence of partial sums, S_n . We use our result of (3.1) to get:

$$\frac{3}{2} - \frac{1}{n+1} < S_n < 2 - \frac{1}{n}$$

We see that the upper bound on S_n (and thus also S) is 2 and the lower bound is $\frac{3}{2}$. Since the partial sum sequence is increasing, the existence of an upper bound proves that the series converges. The value of the bound itself must lie between $\frac{3}{2}$ and 2.

3.4 Common Series and Convergence

As with sequences, it is extremely useful to have a good knowledge of whether some common series converge or not. These will be very useful in comparison tests for both convergence and divergence arguments.

Diverging series It should be noted that if $S = \sum_{n=1}^{\infty} a_n$ and the sequence a_n does not itself tend to 0 as $n \rightarrow \infty$, then S will certainly diverge.

1. Harmonic series: $S = \sum_{n=1}^{\infty} \frac{1}{n}$ diverges.

2. Harmonic primes: $S = \sum_{p: \text{prime}} \frac{1}{p}$ diverges.

3. Geometric series: $S = \sum_{n=1}^{\infty} x^n$ diverges for $|x| \geq 1$.

Converging series The limit for the inverse squares series is given for interest only, the essential thing is to know that the series converges.

1. Geometric series: $S = \sum_{n=1}^{\infty} x^n$ converges to $\frac{x}{1-x}$ so long as $|x| < 1$.
2. Inverse squares series: $S = \sum_{n=1}^{\infty} \frac{1}{n^2}$ converges to $\frac{\pi^2}{6}$.
3. $\frac{1}{n^c}$ series: $S = \sum_{n=1}^{\infty} \frac{1}{n^c}$ converges for $c > 1$.

3.5 Convergence Tests

In this section we focus on series with non-negative a_n terms. In these series the partial sums are increasing and hence either they converge, if the partial sums are bounded, or they diverge, if they are not.

In the following tests, we take:

1. a_i is a non-negative term in the series $\sum_{i=1}^{\infty} a_i$ that we wish to reason about
2. $\sum_{i=1}^{\infty} c_i$ is a series that we have already established converges with sum c
3. $\sum_{i=1}^{\infty} d_i$ is a series that we have already established diverges

Comparison test Let $\lambda > 0$ and $N \in \mathbb{N}$. Then

1. if $a_i \leq \lambda c_i$ for all $i > N$, then $\sum_{i=1}^{\infty} a_i$ converges
2. if $a_i \geq \lambda d_i$ for all $i > N$, then $\sum_{i=1}^{\infty} a_i$ diverges

Limit comparison test Sometimes the following form of comparison test is easier, as it does not involve having to find a constant λ to make the first form of the test work:

1. if $\lim_{i \rightarrow \infty} \frac{a_i}{c_i}$ exists, then $\sum_{i=1}^{\infty} a_i$ converges
2. if $\lim_{i \rightarrow \infty} \frac{d_i}{a_i}$ exists, then $\sum_{i=1}^{\infty} a_i$ diverges

D'Alembert's ratio test The D'Alembert ratio test is a very useful test for quickly discerning the convergence or divergence of a series. It exists in many subtly different forms.

From some point $N \in \mathbb{N}$ in the series, for all $i \geq N$:

1. If $\frac{a_{i+1}}{a_i} \geq 1$, then $\sum_{i=1}^{\infty} a_i$ diverges
2. If there exists a k such that $\frac{a_{i+1}}{a_i} \leq k < 1$, then $\sum_{i=1}^{\infty} a_i$ converges

D'Alembert's limit ratio test From this definition we can obtain (Section 3.8) the easier-to-use and more often quoted version of the test which looks at the limit of the ratio of successive terms. This version states:

1. if $\lim_{i \rightarrow \infty} \frac{a_{i+1}}{a_i} > 1$ then $\sum_{i=1}^{\infty} a_i$ diverges
2. if $\lim_{i \rightarrow \infty} \frac{a_{i+1}}{a_i} = 1$ then $\sum_{i=1}^{\infty} a_i$ may converge or diverge
3. if $\lim_{i \rightarrow \infty} \frac{a_{i+1}}{a_i} < 1$ then $\sum_{i=1}^{\infty} a_i$ converges

There are some simple examples of series, $\sum_{n=1}^{\infty} a_n$, which converge or diverge but give a limit value of 1 when applying the limit ratio test. $a_n = 1$ for all n diverges but gives a limit of 1 in the above test. By contrast, $a_n = \frac{1}{n^2}$ for $n \geq 1$ converges but also gives a limit ratio of 1.

Integral test This test is based on the idea¹ that for a sequence $S = \sum_{n=1}^{\infty} a_n$ where $a_n = f(n)$ is a decreasing function, e.g., $\frac{1}{n}$ as $n \geq 1$:

$$\sum_{n=1}^{\infty} a_{n+1} < \int_1^{\infty} f(x) dx < \sum_{n=1}^{\infty} a_n$$

In practice, you evaluate the integral and if it diverges you know that the series diverges by the right hand inequality. If the integral converges, you know the series does also; this time by the left hand inequality.

Formally stated: suppose that $f(x)$ is a continuous, positive and decreasing function on the interval $[N, \infty)$ and that $a_n = f(n)$ then:

1. If $\int_N^{\infty} f(x) dx$ converges, so does $\sum_{n=N}^{\infty} a_n$

¹If you have come across Riemann Sums before, we are essentially using $\sum_{i=1}^{\infty} a_n$ and $\sum_{i=1}^{\infty} a_{n+1}$ as left and right Riemann sums which bound the integral $\int_1^{\infty} f(x) dx$ where $a_n = f(n)$.

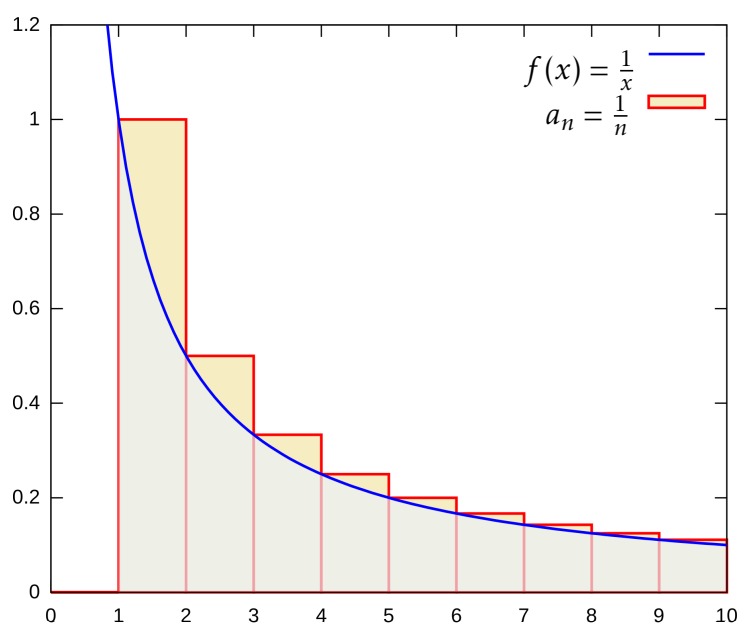


Figure 3.1: A divergence test: the series $\sum_{n=1}^{\infty} \frac{1}{n}$ as compared to $\int_1^{\infty} \frac{1}{x} dx$, showing that the series is bounded below by the integral.

2. If $\int_N^{\infty} f(x) dx$ diverges, so does $\sum_{n=N}^{\infty} a_n$

As usual we do not have to concern ourselves with the first $N - 1$ terms as long as they are finite. If we can show that the tail converges or diverges, i.e., a_n for $n \geq N$, then the whole series will follow suit.

Example: Integral test divergence We can apply the integral test to show that $\sum_{n=1}^{\infty} \frac{1}{n}$ diverges. Figure 3.1 shows $\sum_{n=1}^{\infty} a_n$ where $a_n = \frac{1}{n}$ as a sum of $1 \times a_n$ rectangular areas for $n \geq 1$. When displayed like this, and only because $f(x) = \frac{1}{x}$ is a decreasing function, we can see that the series $\sum_{n=1}^{\infty} a_n$ is strictly greater than the corresponding integral $\int_1^{\infty} f(x) dx$. If we evaluate this integral:

$$\begin{aligned} \int_1^{\infty} f(x) dx &= \lim_{b \rightarrow \infty} \int_1^b \frac{1}{x} dx \\ &= \lim_{b \rightarrow \infty} [\ln x]_1^b \\ &= \lim_{b \rightarrow \infty} (\ln b - \ln 1) \end{aligned}$$

we get a diverging result. Hence the original series $\sum_{n=1}^{\infty} a_n$ must also diverge.

Example: Integral test convergence We can apply the integral test to show that $\sum_{n=1}^{\infty} \frac{1}{n^2}$ converges. Figure 3.2 shows $\sum_{n=1}^{\infty} a_{n+1}$ where $a_n = \frac{1}{n^2}$ as a sum of $1 \times a_{n+1}$

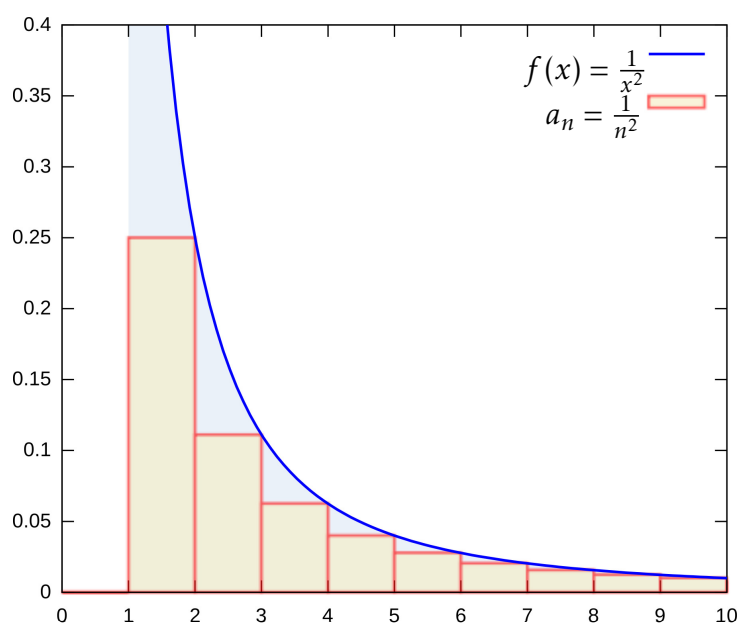


Figure 3.2: A convergence test: the series $\sum_{n=2}^{\infty} \frac{1}{n^2}$ as compared to $\int_1^{\infty} \frac{1}{x^2} dx$, showing that the series is bounded above by the integral.

rectangular areas for $n \geq 1$. When displayed like this, and again only because $f(x) = \frac{1}{x^2}$ is a decreasing function, we can see that the series $\sum_{n=1}^{\infty} a_{n+1}$ is strictly less than the corresponding integral $\int_1^{\infty} f(x) dx$. If we evaluate this integral:

$$\begin{aligned} \int_1^{\infty} f(x) dx &= \lim_{b \rightarrow \infty} \int_1^b \frac{1}{x^2} dx \\ &= \lim_{b \rightarrow \infty} \left[-\frac{1}{x} \right]_1^b \\ &= \lim_{b \rightarrow \infty} \left(1 - \frac{1}{b} \right) \\ &= 1 \end{aligned}$$

Since we get a converging result in the integral, we know that the series $\sum_{n=1}^{\infty} a_{n+1}$ must also converge. Hence $\sum_{n=1}^{\infty} \frac{1}{n^2}$ also converges as it is only one (finite) term different.

3.6 Absolute Convergence

Thus far we have only dealt with series involving positive a_n terms. The difficulty with series that involve a mixture of positive and negative terms, is that the convergence behaviour can depend on the order in which you add the terms up.

Consider $S = 1 - 1 + 1 - 1 + 1 - 1 \dots$, the partial sums of this series oscillate between 1 and 0. If we take two positive terms (from later in the same series) for every

negative term, i.e., $S = 1 + 1 - 1 + 1 + 1 - 1 + 1 + 1 - 1$, we get a series that diverges to $+\infty$. However both series have the same terms and the same number of -1 and $+1$ terms.

Absolute convergence In cases of series with a mixture of positive and negative terms, $S = \sum_{n=1}^{\infty} a_n$, it makes sense to look at the series of absolute values of the terms of a series, that is:

$$S' = \sum_{n=1}^{\infty} |a_n|$$

If this series, S' , converges then we can be sure that the original series S will also converge. S is said to converge absolutely.

A series that converges absolutely also converges, however the same is not true the other way round. A simple example of this is the alternating harmonic series, $S = \sum_{n=1}^{\infty} (-1)^{n-1} \frac{1}{n}$. This series converges to $\ln 2$ (when summed in this order). However the series involving the absolute values of the same terms is the normal harmonic series, $S' = \sum_{n=1}^{\infty} \frac{1}{n}$, which we know diverges.

Tests for absolute convergence Where we have a series with a mixture of positive and negative terms, a_n , we can test for absolute convergence by applying any of the above tests to the absolute value of the terms $|a_n|$. For example the limit ratio test becomes: for $N \in \mathbb{N}$,

1. if $\lim_{n \rightarrow \infty} \left| \frac{a_{n+1}}{a_n} \right| > 1$ then $\sum_{n=1}^{\infty} a_n$ diverges
2. if $\lim_{n \rightarrow \infty} \left| \frac{a_{n+1}}{a_n} \right| = 1$ then $\sum_{n=1}^{\infty} a_n$ may converge or diverge
3. if $\lim_{n \rightarrow \infty} \left| \frac{a_{n+1}}{a_n} \right| < 1$ then $\sum_{n=1}^{\infty} a_n$ converges absolutely (and thus also converges)

3.7 Power Series and the Radius of Convergence

As a glimpse forward to the next topic, we will look at how series convergence affects power series. As you may know, it is possible to represent high-order functions as power series, that is:

$$f(x) = \sum_{n=0}^{\infty} a_n x^n$$

However, such series may only converge for certain ranges of input parameter x . We have already seen this with the geometric series, which represents a power series expansion of the function, $f(x) = \frac{x}{1-x}$ (from Section 3.1) as:

$$f(x) = \sum_{n=1}^{\infty} x^n$$

However this limit, $\frac{x}{1-x}$, and therefore also the power series expansion are only equivalent if $|x| < 1$ as we have seen.

The radius of convergence represents the size of x set for which a power series converges. If a power series converges for $|x| < R$, then the radius of convergence is said to be R . Sometimes, we consider series which only converge within some distance R of a point a , in this case, we might get the following result $|x - a| < R$. In this case also the radius of convergence is R . If the series converges for all values of x , then the radius is given as infinite.

Example: calculating the radius of convergence Let us take a power series:

$$S = \sum_{n=1}^{\infty} n^2 x^n$$

We can apply the D'Alembert ratio test to this as with any other series, however we have to leave the x variable free, and we have to cope with possible negative x values.

We apply the version of the D'Alembert ratio test which tests for absolute convergence (and thus also convergence) since, depending on the value of x , we may have negative terms in the series.

$$\begin{aligned} \lim_{n \rightarrow \infty} \left| \frac{a_{n+1}}{a_n} \right| &= \lim_{n \rightarrow \infty} \left| \frac{(n+1)^2 x^{n+1}}{n^2 x^n} \right| \\ &= \lim_{n \rightarrow \infty} \left| x \left(1 + \frac{1}{n} \right)^2 \right| \\ &= |x| \end{aligned}$$

If we look at the ratio test, it says that if the result is < 1 it will converge. So our general convergence condition on this series, is that $|x| < 1$. Therefore the radius of convergence is 1.

3.8 Proofs of Ratio Tests

As with the ratio test proofs for converging sequences, I do not expect you to come up with these arguments in an exam. However some of you will be interested in how these tests work.

Proof of D'Alembert ratio test We take the two cases separately:

Case $\frac{a_{i+1}}{a_i} \geq 1$ for all $i \geq N$: Taking $a_N = C$ and the non-increasing case of $\frac{a_{i+1}}{a_i} = 1$, we get a partial sum from the N th term of:

$$\sum_{i=N}^{n+N-1} a_i = Cn$$

This diverges as it is proportional to n . A comparison test with this constant term series and the case where $\frac{a_{i+1}}{a_i} > 1$ shows that the case of the ratio being greater than 1 also diverges.

Case $\frac{a_{i+1}}{a_i} \leq k < 1$: We can see that this series converges by performing a term-by-term comparison with a geometric series with ratio $x = k$. From $i \geq N$, we get:

$$\begin{aligned} a_{N+1} &\leq k a_N \\ a_{N+2} &\leq k a_{N+1} \leq k^2 a_N \\ a_{N+3} &\leq k a_{N+2} \leq k^2 a_{N+1} \leq k^3 a_N \\ &\vdots \\ a_{N+m} &\leq k^m a_N \end{aligned}$$

We can rewrite this last as $a_{N+m} \leq \lambda k^{N+m}$ where $\lambda = \frac{a_N}{k^N}$ is a constant. Letting $n = N + m$, allows us to write that $a_n \leq \lambda k^n$ for $n \geq N$. Since we know that $k < 1$ we are sure that the geometric series $b_n = k^n$ converges. This is precisely what we require under the conditions of the comparison test, to show that $\sum_{n=1}^{\infty} a_n$ converges.

Proof of D'Alembert limit ratio test Taking the two cases in turn:

Case $\lim_{i \rightarrow \infty} \frac{a_{i+1}}{a_i} = l > 1$: We are looking to show that $\sum_{i=1}^{\infty} a_i$ diverges. Using the definition of the limit, we know that from some point N for all $i > N$:

$$\left| \frac{a_{i+1}}{a_i} - l \right| < \epsilon$$

$$l - \epsilon < \frac{a_{i+1}}{a_i} < l + \epsilon$$

Since we are trying to show that this series diverges, we would like to use the left hand inequality to show that the ratio $\frac{a_{i+1}}{a_i} > 1$ from some point onwards, for example $i > N$. To do this we can pick $\epsilon = l - 1$, recalling that $l > 1$, so this is allowed as $\epsilon > 0$. We now get that for all $i > N$:

$$1 < \frac{a_{i+1}}{a_i}$$

and so by the non-limit D'Alembert ratio test, the series $\sum_{i=1}^{\infty} a_i$ diverges.

Case $\lim_{i \rightarrow \infty} \frac{a_{i+1}}{a_i} = l < 1$: We are looking to show that $\sum_{i=1}^{\infty} a_i$ converges in this case. Again using the limit definition we get for all $i > N$:

$$\left| \frac{a_{i+1}}{a_i} - l \right| < \epsilon$$

$$l - \epsilon < \frac{a_{i+1}}{a_i} < l + \epsilon$$

We are aiming to show that $\frac{a_{i+1}}{a_i} \leq k < 1$ for some k , so that we can avail ourselves of the second D'Alembert ratio test case. So we are looking to use the right hand inequality on this occasion. This means that we will need to set $k = l + \epsilon$ but we still need to guarantee that $k < 1$ (and obviously $l < k$ also as we need to set $k = l + \epsilon$ where $\epsilon > 0$). We can do this easily by situating k between l and 1 with $k = \frac{l+1}{2}$. This sets $\epsilon = \frac{1-l}{2}$ and importantly $\epsilon > 0$ since $l < 1$. Going forward from setting ϵ (and because we have fixed it to work), we get for all $i > N$:

$$\frac{a_{i+1}}{a_i} < k < 1$$

So, by the non-limit D'Alembert ratio test, the series $\sum_{i=1}^{\infty} a_i$ converges.

Chapter 4

Power Series

4.1 Basics of Power Series

We can represent a function $f(x)$ by a series expansion in x

$$f(x) = \sum_{i=0}^{\infty} a_i x^i$$

for coefficients $a_i \in \mathbb{R}$ and $i = 1, 2, \dots$. This is called a *power series* because it is a series of powers of the argument x .

For example:

$$f(x) = (1+x)^2 = 1 + 2x + x^2$$

has $a_0 = 1, a_1 = 2, a_2 = 1$ and $a_i = 0$ for $i > 2$. However in general the series may be infinite provided it converges. We will use ratio-of-convergence arguments to check this.

4.2 Maclaurin Series

For a function $f(x)$:

$$f(x) = a_0 + a_1 x + a_2 x^2 + a_3 x^3 + \dots$$

Suppose the value of the function $f(x)$ is known at $x = 0$. Then we have an easy way of finding the first coefficient of our power series, a_0 . We can set $x = 0$, to get: $f(0) = a_0$.

Suppose the derivatives of the function $f(x)$ are known at $x = 0$, that is:

$$f'(0), f''(0), \dots, f^{(n)}(0)$$

where $f^{(n)}(a)$ represents the function $f(x)$ differentiated n times and then taking the derivative expression with $x = a$. Now if we differentiate $f(x)$ to get rid of the constant term:

$$f'(x) = a_1 + 2.a_2 x + 3.a_3 x^2 + 4.a_4 x^3 + \dots$$

and again set $x = 0$, we get $a_1 = f'(0)$.

Differentiate again to get rid of the constant term:

$$f''(x) \equiv f^{(2)}(x) = 2.1.a_2 + 3.2.a_3x + 4.3.a_4x^2 + \dots$$

Set $x = 0$ and repeat the process, we get:

$$a_2 = \frac{f^{(2)}(0)}{2!}$$

In general, for the n th term, we get:

$$a_n = \frac{f^{(n)}(0)}{n!} \text{ for } n \geq 0$$

Maclaurin series Suppose $f(x)$ is differentiable infinitely many times and that it has a power series representation (*series expansion*) of the form

$$f(x) = \sum_{i=0}^{\infty} a_i x^i$$

as above.

Differentiating n times gives

$$f^{(n)}(x) = \sum_{i=n}^{\infty} a_i i(i-1)\dots(i-n+1)x^{i-n}$$

Setting $x = 0$, we have $f^{(n)}(0) = n!a_n$ because all terms but the first have x as a factor. Hence we obtain Maclaurin's series:

$$f(x) = \sum_{n=0}^{\infty} f^{(n)}(0) \frac{x^n}{n!}$$

As we have seen in the Series notes, it is important to check the radius of convergence (set of valid values for x) that ensures that this series converges.

The addition of each successive term in the Maclaurin series creates a closer approximation to the original function around the point $x = 0$, as we will demonstrate.

Example Taking $f(x) = (1+x)^3$. We repeat the process of differentiating repeatedly and setting x to zero in each derivative expression:

1. $f(0) = 1$ so $a_0 = 1$
2. $f'(x) = 3(1+x)^2$ so $f'(0) = 3$ and $a_1 = \frac{3}{1!} = 3$
3. $f''(x) = 3.2(1+x)$ so $f''(0) = 6$ and $a_2 = \frac{6}{2!} = 3$
4. $f'''(x) = 3.2.1$ so $f'''(0) = 6$ and $a_3 = \frac{6}{3!} = 1$

Higher derivatives are all 0 and so (as we know), so we end up with:

$$(1+x)^3 = 1 + 3x + 3x^2 + x^3$$

In this case, since the series expansion is finite we know that the Maclaurin series will be accurate (and certainly converge) for all x , therefore the radius of convergence is infinite.

We now consider the partial sums of the Maclaurin expansion $f_n(x) = \sum_{i=0}^n \frac{f^{(i)}(x)}{i!} x^i$. For $f(x) = (1+x)^3$ these would be:

$$f_0(x) = 1$$

$$f_1(x) = 1 + 3x$$

$$f_2(x) = 1 + 3x + 3x^2$$

$$f_3(x) = 1 + 3x + 3x^2 + x^3$$

We can plot these along with $f(x)$ in Figure 4.1, and we see that each partial sum is a successively better approximation to $f(x)$ around the point $x = 0$.

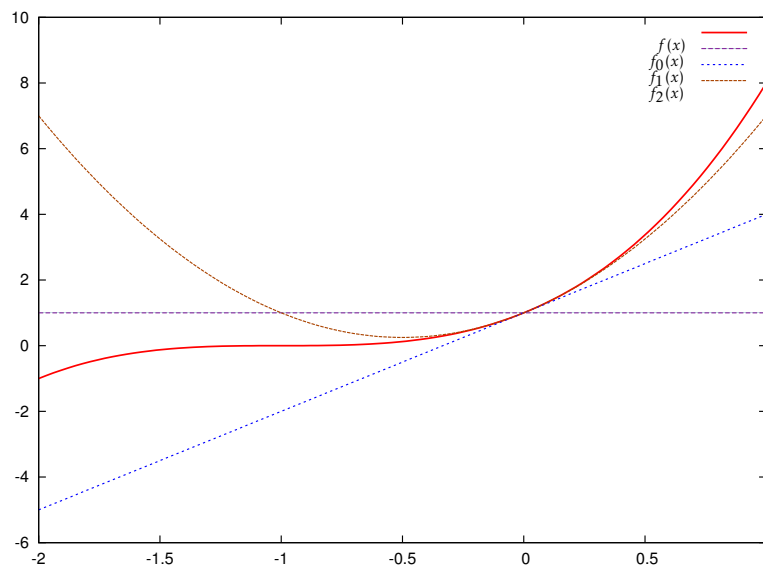


Figure 4.1: Successive approximations from the Maclaurin series expansion of $f(x) = (1+x)^3$ where $f_0(x) = 1$, $f_1(x) = 1 + 3x$ and $f_2(x) = 1 + 3x + 3x^2$.

Example 2: $f(x) = (1-x)^{-1}$ We probably know that the power series for this function is the geometric series in x , in which all $a_i = 1$.

1. $f(0) = 1$, so far so good!
2. $f'(x) = -(1-x)^{-2}(-1) = (1-x)^{-2}$ so $f'(0) = 1$
3. $f''(x) = -2!(1-x)^{-3}(-1) = 2!(1-x)^{-3}$ so $f''(0) = 2!$

Differentiating repeatedly,

$$\begin{aligned} f^{(n)}(x) &= (-1)(-2)\dots(-n)(1-x)^{-(n+1)}(-1)^n \\ &= n!(1-x)^{-(n+1)} \end{aligned}$$

so

$$a_n = \frac{f^{(n)}(0)}{n!} = \frac{n!(1)^{-(n+1)}}{n!} = 1$$

Thus $(1-x)^{-1} = \sum_{i=0}^{\infty} x^i$ unsurprisingly since this is the sum of an infinite geometric series in x .

We can check convergence by using the absolute convergence version of D'Alembert's ratio test, where $b_n = x^n$ are the series terms:

$$\begin{aligned} \lim_{n \rightarrow \infty} \left| \frac{b_{n+1}}{b_n} \right| &= \lim_{n \rightarrow \infty} \left| \frac{x^{n+1}}{x^n} \right| \\ &= |x| \end{aligned}$$

Thus convergence is given by $|x| < 1$ and a radius of convergence of 1 for this power series.

Example 3: $f(x) = \ln(1+x)$

1. $f(0) = 0$ because $\ln 1 = 0$ so no constant term
2. $f'(x) = (1+x)^{-1}$ so $f'(0) = 1$
3. $f''(x) = (-1)(1+x)^{-2}$ so $f''(0) = -1$

Giving an n th derivative of:

$$f^{(n)}(x) = \frac{(-1)^{n-1}(n-1)!}{(1+x)^n} \text{ and thus } a_n = \frac{f^{(n)}(0)}{n!} = \frac{(-1)^{n-1}}{n}$$

Therefore series for $f(x) = \ln(1+x)$ is:

$$\ln(1+x) = \frac{x}{1} - \frac{x^2}{2} + \frac{x^3}{3} - \frac{x^4}{4} + \dots$$

Taking series terms of $b_n = \frac{(-1)^{n-1}x^n}{n}$, we get a ratio test of $\lim_{n \rightarrow \infty} \left| \frac{b_{n+1}}{b_n} \right| = |x|$ and a convergence condition of $|x| < 1$.

Indeed we see that if we were to set $x = -1$, we would get $\ln 0 = -\infty$ and a corresponding power series of:

$$\ln 0 = - \sum_{n=1}^{\infty} \frac{1}{n}$$

which we already know diverges.

On the other hand setting $x = 1$, gives the alternating harmonic series and a nice convergence result:

$$\ln 2 = \sum_{n=1}^{\infty} \frac{(-1)^{n-1}}{n}$$

which goes to show that you cannot be sure whether a ratio test of 1 will give a converging or diverging series.

4.3 Taylor Series

A more general result is:

$$\begin{aligned} f(x) &= f(a) + \frac{f^{(1)}(a)}{1!}(x-a) + \frac{f^{(2)}(a)}{2!}(x-a)^2 + \dots \\ &= \sum_{n=0}^{\infty} \frac{f^{(n)}(a)}{n!}(x-a)^n \end{aligned}$$

allows us to create power series of a function around a particular point $x = a$ that we may be interested in. This allows us to create successively better approximations to the function $f(x)$ about $x = a$ directly rather than having to use huge numbers of terms from the Maclaurin series to achieve the same result.

Setting $a = 0$ gives us an expansion around zero and we recover the Maclaurin series as a special case.

Example For a series like $f(x) = \ln x$, we would not be able to create a Maclaurin series since $\ln 0$ is not defined (is singular). So this is an example where an expansion around another point is required. We show a Taylor expansion about $x = 2$:

1. $f(x) = \ln x$
2. $f'(x) = \frac{1}{x}$
3. $f''(x) = -\frac{1}{x^2}$
4. $f'''(x) = 2!(-1)^2 \frac{1}{x^3}$
5. $f^{(n)} = \frac{(-1)^{n-1}(n-1)!}{x^n}$ for $n > 0$

This gives us a Taylor series of:

$$f(x) = \ln 2 + \sum_{n=1}^{\infty} \frac{(-1)^{n-1}}{n2^n}(x-2)^n$$

We can now show how the Taylor series approximates the function $f(x)$ around the point $x = a$ in the same way as the Maclaurin series does around the point $x = 0$.

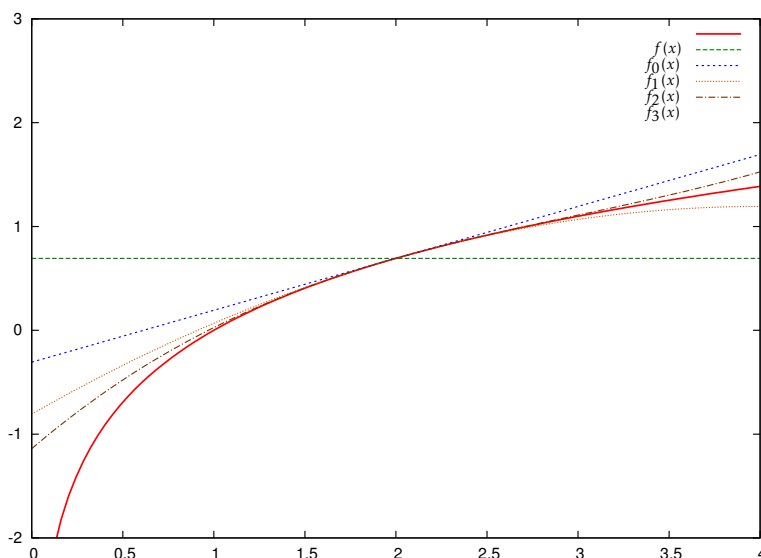


Figure 4.2: Successive approximations from the Taylor series expansion of $f(x) = \ln x$ around the point $x = 2$.

Figure 4.2 shows the first four partial sums $f_n(x) = \sum_{i=0}^n \frac{f^{(i)}(a)}{i!} (x-a)^i$ from the Taylor series for $f(x) = \ln x$ around the point $x = 2$, where:

$$\begin{aligned} f_0(x) &= \ln 2 \\ f_1(x) &= \ln 2 + \frac{x-2}{2} \\ f_2(x) &= \ln 2 + \frac{x-2}{2} - \frac{(x-2)^2}{8} \\ f_3(x) &= \ln 2 + \frac{x-2}{2} - \frac{(x-2)^2}{8} + \frac{(x-2)^3}{24} \end{aligned}$$

As for Maclaurin series we still need to be aware of the radius of convergence for this series. The techniques for calculating it are identical for Taylor series.

Taking the absolute ratio test for terms $b_n = \frac{(-1)^{n-1}}{n2^n} (x-2)^n$, where we can ignore the $\ln 2$ term since we are looking in the limit of $n \rightarrow \infty$:

$$\begin{aligned} \lim_{n \rightarrow \infty} \left| \frac{b_{n+1}}{b_n} \right| &= \lim_{n \rightarrow \infty} \left| \frac{(x-2)^{n+1}/(n+1)2^{n+1}}{(x-2)^n/n2^n} \right| \\ &= \lim_{n \rightarrow \infty} \frac{|x-2|}{2} \frac{n}{n+1} \\ &= \frac{|x-2|}{2} \end{aligned}$$

This gives us a convergence condition of $|x-2| < 2$ and a radius of convergence of 2 for this power series. In general the radius of convergence is limited by the nearest singularity (such as $x = 0$ in this case) for a real series or pole if we extend complex series to the complex plane (we may or may not have time to investigate poles!).

4.4 Taylor Series Error Term

In computing, numerical and scientific applications it is usual that we truncate the series at the k th term so that the computation terminates. In these circumstances it is important to know how close we are to the actual value of $f(x)$ (in the worst case) and are able to approximate the error that we have introduced by not summing the entire series.

The infinite series gets translated into a finite one:

$$\begin{aligned} f(x) &= \sum_{n=0}^{\infty} \frac{f^{(n)}(a)}{n!} (x-a)^n \\ &= \sum_{n=0}^k \frac{f^{(n)}(a)}{n!} (x-a)^n + \frac{f^{(k+1)}(c)}{(k+1)!} (x-a)^{k+1} \end{aligned} \quad (4.1)$$

where c is a constant that lies between x and a . The term $\frac{f^{(k+1)}(c)}{(k+1)!} (x-a)^{k+1}$ is known as the Lagrange error term and it replaces the tail of the infinite series from the $(k+1)$ th term onwards. Although it is precise there is no easy way to find c . So in practice the bound $a < c < x$ or $x < c < a$ is used to generate a worst-case error for the term.

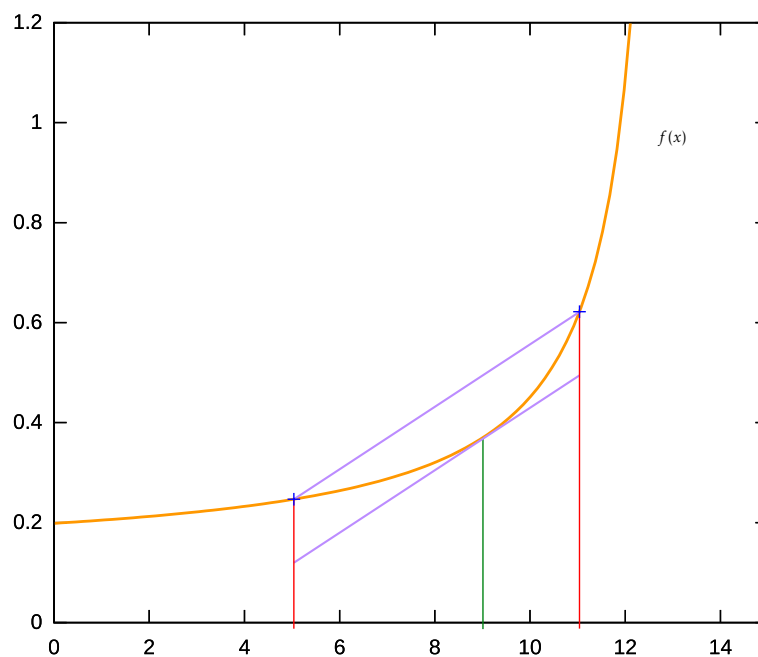


Figure 4.3: Mean Value Theorem

Figure 4.3 shows a continuous and differentiable function $f(x)$. The mean value theorem states that between two points x and a , there is a point c such that $f'(c)$ is equal to the gradient between the points $(a, f(a))$ and $(x, f(x))$.

Given that we can compute the gradient between those two points, we get:

$$\begin{aligned} f'(c) &= \frac{f(x) - f(a)}{x - a} \\ f(x) &= f(a) + (x - a)f'(c) \end{aligned} \quad (4.2)$$

for some c between x and a . Taylor's theorem can be thought of as a general version of Mean Value Theorem but expressed in terms of the $(k + 1)$ th derivative instead of the 1st.

In fact if you set $k = 0$ in (4.1), you get the mean value theorem of (4.2).

4.5 Deriving the Cauchy Error Term

Let us start by defining the set of partial sums to the Taylor series. Here we are going to fix x and vary the offset variable.

$$F_k(t) = \sum_{n=0}^k \frac{f^{(n)}(t)}{n!} (x - t)^n$$

Now $F_k(a)$ is the standard form of the n th partial sum and $F_k(x) = f(x)$ for all k . Thus:

$$\begin{aligned} F_k(x) - F_k(a) &= f(x) - \sum_{n=0}^k \frac{f^{(n)}(a)}{n!} (x - a)^n \\ &= R_k(x) \end{aligned}$$

4.6 Power Series Solution of ODEs

Consider the differential equation

$$\frac{dy}{dx} = ky$$

for constant k , given that $y = 1$ when $x = 0$.

Try the [series solution](#)

$$y = \sum_{i=0}^{\infty} a_i x^i$$

Find the coefficients a_i by differentiating term by term, to obtain the [identity](#), for $i \geq 0$:

Matching coefficients

$$\sum_{i=1}^{\infty} a_i i x^{i-1} \equiv \sum_{i=0}^{\infty} k a_i x^i \equiv \sum_{i=1}^{\infty} k a_{i-1} x^{i-1}$$

Comparing coefficients of x^{i-1} for $i \geq 1$

$$ia_i = ka_{i-1} \quad \text{hence}$$

$$a_i = \frac{k}{i}a_{i-1} = \frac{k}{i} \cdot \frac{k}{i-1}a_{i-2} = \dots = \frac{k^i}{i!}a_0$$

When $x = 0, y = a_0$ so $a_0 = 1$ by the boundary condition. Thus

$$y = \sum_{i=0}^{\infty} \frac{(kx)^i}{i!} = e^{kx}$$

Chapter 5

Linear Algebra

This chapter is largely based on the lecture notes and books by Drumm and Weil (2001); Strang (2003); Hogben (2013); Liesen and Mehrmann (2015) as well as Pavel Grinfeld’s Linear Algebra series¹. Another excellent source is Gilbert Strang’s Linear Algebra course at MIT².

Linear algebra is the study of vectors. The vectors we know from school are called “geometric vectors”. In general, however, vectors are special objects that can be added together and multiplied by scalars to produce another object of the same kind. Any object that satisfies these two properties can be considered a vector. Here are some examples of such vector objects:

1. Geometric vectors. This example of a vector may be familiar from High School. Geometric vectors are directed segments, which can be drawn, see Fig. 5.1. Two vectors \vec{x}, \vec{y} can be added, such that $\vec{x} + \vec{y} = \vec{z}$ is another geometric vector. Furthermore, $\lambda \vec{x}, \lambda \in \mathbb{R}$ is also a geometric vector. In fact, it is the original vector scaled by λ . Therefore, geometric vectors are instances of the vector concepts introduced above.
2. Polynomials are also vectors: Two polynomials can be added together, which results in another polynomial; and they can be multiplied by a scalar $\lambda \in \mathbb{R}$, and the result is a polynomial as well. Therefore, polynomial are (rather unusual) instances of vectors. Note that polynomials are very different from geometric vectors. While geometric vectors are concrete “drawings”, polynomials are abstract concepts. However, they are both vectors.
3. \mathbb{R}^n is a set of numbers, and its elements are n -tuples. \mathbb{R}^n is even more abstract than polynomials, and the most general concept we consider in this course. For example,

$$\mathbf{a} = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \in \mathbb{R}^3 \quad (5.1)$$

¹<http://tinyurl.com/nahclwm>

²<http://tinyurl.com/29p5q8j>

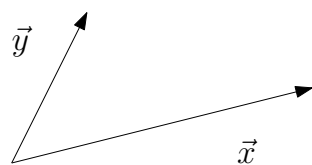


Figure 5.1: Example of two geometric vectors in two dimensions.

is an example of a triplet of numbers. Adding two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ component-wise results in another vector: $\mathbf{a} + \mathbf{b} = \mathbf{c} \in \mathbb{R}^n$. Moreover, multiplying $\mathbf{a} \in \mathbb{R}^n$ by $\lambda \in \mathbb{R}$ results in a scaled vector $\lambda \mathbf{a} \in \mathbb{R}^n$.

4. Audio signals are vectors, too. Audio signals are represented as a series of numbers. We can add audio signals together, and their sum is a new audio signals. If we scale an audio signal, we also obtain an audio signal. Therefore, audio signals are a type of vector, too.

Linear algebra focuses on the similarities between these vector concepts: We can add them together and multiply them by scalars. We will largely focus on the third kind of vectors since most algorithms in linear algebra are formulated in \mathbb{R}^n . There is a 1:1 correspondence between any kind of vector and \mathbb{R}^n . By studying \mathbb{R}^n , we implicitly study all other vectors. Although \mathbb{R}^n is rather abstract, it is most useful.

Practical Applications of Linear Algebra

Linear algebra centers around solving linear equation systems and is at the core of many computer science applications. Here is a selection:³

- Ranking of web pages (web search)
- Linear programming (optimization)
- Error correcting codes (e.g., in DVDs)
- Decomposition of sounds into different sources
- Projections, rotations, scaling (computer graphics)
- Data visualization
- En/Decryption algorithms (cryptography)
- State estimation and optimal control (e.g., in robotics and dynamical systems)
- Numerics: determine whether a computation is numerically stable (e.g., in large-scale data-analytics systems, optimization, machine learning)
- Find out whether a function is convex (e.g., in Optimization)

³More details can be found on Jeremy Kun's blog: <http://tinyurl.com/olkbkct>

5.1 Linear Equation Systems

Systems of linear equations play a central part of linear algebra. Many problems can be formulated as systems of linear equations, and linear algebra gives us the tools for solving them.

Example

A company produces products N_1, \dots, N_n for which resources R_1, \dots, R_m are required. To produce a unit of product N_j , a_{ij} units of resource R_i are needed, where $i = 1, \dots, m$ and $j = 1, \dots, n$.

The objective is to find an optimal production plan, i.e., a plan how many units x_j of product N_j should be produced if a total of b_i units of resource R_i are available and (ideally) no resources are left over.

If we produce x_1, \dots, x_n units of the corresponding products, we need a total of

$$a_{i1}x_1 + \dots + a_{in}x_n \quad (5.2)$$

many units of resource R_i . The desired optimal production plan $(x_1, \dots, x_n) \in \mathbb{R}^n$, therefore, has to satisfy the following system of equations:

$$\begin{aligned} a_{11}x_1 + \dots + a_{1n}x_n &= b_1 \\ \vdots &= \vdots, \\ a_{m1}x_1 + \dots + a_{mn}x_n &= b_m \end{aligned} \quad (5.3)$$

where $a_{ij} \in \mathbb{R}$ and $b_i \in \mathbb{R}$.

Equation (5.3) is the general form of a **linear equation system**, and x_1, \dots, x_n are the **unknowns** of this linear equation system. Every n -tuple $(x_1, \dots, x_n) \in \mathbb{R}^n$ that satisfies (5.3) is a **solution** of the linear equation system.

Example

The linear equation system

$$\begin{aligned} x_1 + x_2 + x_3 &= 3 & (1) \\ x_1 - x_2 + 2x_3 &= 2 & (2) \\ 2x_1 + 3x_3 &= 1 & (3) \end{aligned} \quad (5.4)$$

has **no solution**: Adding the first two equations yields $(1)+(2) = 2x_1 + 3x_3 = 5$, which contradicts the third equation (3).

Let us have a look at the linear equation system

$$\begin{aligned} x_1 + x_2 + x_3 &= 3 & (1) \\ x_1 - x_2 + 2x_3 &= 2 & (2) \\ x_2 + x_3 &= 2 & (3) \end{aligned} \quad (5.5)$$

From the first and third equation it follows that $x_1 = 1$. From $(1)+(2)$ we get $2+3x_3 = 5$, i.e., $x_3 = 1$. From (3), we then get that $x_2 = 1$. Therefore, $(1, 1, 1)$ is the only possible and **unique solution** (verify by plugging in).

As a third example, we consider

$$\begin{aligned} x_1 + x_2 + x_3 &= 3 & (1) \\ x_1 - x_2 + 2x_3 &= 2 & (2) \\ 2x_1 + 3x_3 &= 5 & (3) \end{aligned} \quad (5.6)$$

Since (1)+(2)=(3), we can omit the third equation (redundancy). From (1) and (2), we get $2x_1 = 5 - 3x_3$ and $2x_2 = 1 + x_3$. We define $x_3 = a \in \mathbb{R}$ as a free variable, such that any triplet

$$\left(\frac{5}{2} - \frac{3}{2}a, \frac{1}{2} + \frac{1}{2}a, a \right), \quad a \in \mathbb{R} \quad (5.7)$$

is a solution to the linear equation system, i.e., we obtain a solution set that contains **infinitely many** solutions.

In general, for a real-valued linear equation system we obtain either no, exactly one or infinitely many solutions.

For a systematic approach to solving linear equation systems, we will introduce a useful compact notation. We will write the linear equation system from (5.3) in the following form:

$$x_1 \begin{bmatrix} a_{11} \\ \vdots \\ a_{m1} \end{bmatrix} + x_2 \begin{bmatrix} a_{12} \\ \vdots \\ a_{m2} \end{bmatrix} + \cdots + x_n \begin{bmatrix} a_{1n} \\ \vdots \\ a_{mn} \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_m \end{bmatrix} \Leftrightarrow \begin{bmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & & \vdots \\ a_{m1} & \cdots & a_{mn} \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_m \end{bmatrix}. \quad (5.8)$$

In order to work with these **matrices**, we need to have a close look at the underlying algebraic structures and define computation rules.

5.2 Groups

Groups play an important role in computer science. Besides providing a fundamental framework for operations on sets, they are heavily used in cryptography, coding theory and graphics.

5.2.1 Definitions

Consider a set G and an operation $\otimes : G \rightarrow G$ defined on G . For example, \otimes could be $+, \cdot$ defined on $\mathbb{R}, \mathbb{N}, \mathbb{Z}$ or \cup, \cap, \setminus defined on $\mathcal{P}(B)$, the power set of B .

Then (G, \otimes) is called a **group** if

- **Closure** of G under \otimes : $\forall x, y \in G : x \otimes y \in G$
- **Associativity**: $\forall x, y, z \in G : (x \otimes y) \otimes z = x \otimes (y \otimes z)$
- **Neutral element**: $\exists e \in G \forall x \in G : x \otimes e = x$ and $e \otimes x = x$

- **Inverse element:** $\forall x \in G \exists y \in G : x \otimes y = e$ and $y \otimes x = e$. We often write x^{-1} to denote the inverse element of x .⁴

If additionally $\forall x, y \in G : x \otimes y = y \otimes x$ then (G, \otimes) is **Abelian group** (commutative).

5.2.2 Examples

$(\mathbb{Z}, +)$ is a group, whereas $(\mathbb{N}_0, +)$ ⁵ is not: Although $(\mathbb{N}_0, +)$ possesses a neutral element (0), the inverse elements are missing.

(\mathbb{Z}, \cdot) is not a group: Although (\mathbb{Z}, \cdot) contains a neutral element (1), the inverse elements for any $z \in \mathbb{Z}, z \neq \pm 1$, are missing.

(\mathbb{R}, \cdot) is not a group since 0 does not possess an inverse element. However, $(\mathbb{R} \setminus \{0\})$ is Abelian.

$(\mathbb{R}^n, +), (\mathbb{Z}^n, +), n \in \mathbb{N}$ are Abelian if $+$ is defined componentwise, i.e.,

$$(x_1, \dots, x_n) + (y_1, \dots, y_n) = (x_1 + y_1, \dots, x_n + y_n). \quad (5.9)$$

Then, $e = (0, \dots, 0)$ is the neutral element and $(x_1, \dots, x_n)^{-1} := (-x_1, \dots, -x_n)$ is the inverse element.

5.3 Matrices

Definition 1 (Matrix)

With $m, n \in \mathbb{N}$ a real-valued (m, n) **matrix** A is an $m \cdot n$ -tuple of elements $a_{ij}, i = 1, \dots, m, j = 1, \dots, n$, which is ordered according to a rectangular scheme consisting of m rows and n columns:

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix}, \quad a_{ij} \in \mathbb{R}. \quad (5.10)$$

$(1, n)$ -matrices are called **rows**, $(m, 1)$ -matrices are called **columns**. These special matrices are also called **row/column vectors**.

$\mathbb{R}^{m \times n}$ is the set of all real-valued (m, n) -matrices. $A \in \mathbb{R}^{m \times n}$ can be equivalently represented as $A \in \mathbb{R}^{mn}$. Therefore, $(\mathbb{R}^{m \times n}, +)$ is Abelian group (with componentwise addition as defined in (5.9)).

5.3.1 Matrix Multiplication

For $A \in \mathbb{R}^{m \times n}, B \in \mathbb{R}^{n \times k}$ (note the size of the matrices!) the elements c_{ij} of the product $C = AB \in \mathbb{R}^{m \times k}$ are defined as

$$c_{ij} = \sum_{l=1}^n a_{il} b_{lj}, \quad i = 1, \dots, m, \quad j = 1, \dots, k. \quad (5.11)$$

⁴The inverse element is defined with respect to the operation \otimes and does not necessarily mean $\frac{1}{x}$.

⁵ $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$

This means, to compute element c_{ij} we multiply the elements of the i th row of A with the j th column of B ⁶ and sum them up.⁷

Remark 1

Matrices can only be multiplied if their “neighboring” dimensions match. For instance, an $n \times k$ -matrix A can be multiplied with a $k \times m$ -matrix B , but only from the left side:

$$\underbrace{A}_{n \times k} \underbrace{B}_{k \times m} = \underbrace{C}_{n \times m} \quad (5.12)$$

The product BA is not defined if $m \neq n$ since the neighboring dimensions do not match.

Remark 2

Note that matrix multiplication is **not** defined as an element-wise operation on matrix elements, i.e., $c_{ij} \neq a_{ij}b_{ij}$ (even if the size of A, B was chosen appropriately).⁸

Example

For $A = \begin{bmatrix} 1 & 2 & 3 \\ 3 & 2 & 1 \end{bmatrix} \in \mathbb{R}^{2 \times 3}$, $B = \begin{bmatrix} 0 & 2 \\ 1 & -1 \\ 0 & 1 \end{bmatrix} \in \mathbb{R}^{3 \times 2}$, we obtain

$$AB = \begin{bmatrix} 1 & 2 & 3 \\ 3 & 2 & 1 \end{bmatrix} \begin{bmatrix} 0 & 2 \\ 1 & -1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 2 & 3 \\ 2 & 5 \end{bmatrix} \in \mathbb{R}^{2 \times 2}, \quad (5.13)$$

$$BA = \begin{bmatrix} 0 & 2 \\ 1 & -1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 & 3 \\ 3 & 2 & 1 \end{bmatrix} = \begin{bmatrix} 6 & 4 & 2 \\ -2 & 0 & 2 \\ 3 & 2 & 1 \end{bmatrix} \in \mathbb{R}^{3 \times 3}. \quad (5.14)$$

From this example, we can already see that matrix multiplication is not commutative, i.e., $AB \neq BA$.

Definition 2 (Identity Matrix)

In $\mathbb{R}^{n \times n}$, we define the **identity matrix** as

$$I_n = \begin{bmatrix} 1 & 0 & \cdots & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & 1 & 0 \\ 0 & \cdots & \cdots & 0 & 1 \end{bmatrix} \in \mathbb{R}^{n \times n}. \quad (5.15)$$

With this, $A \cdot I_n = A = I_n A$ for all $A \in \mathbb{R}^{n \times n}$. Therefore, the identity matrix is the neutral element with respect to matrix multiplication “ \cdot ” in $(\mathbb{R}^{n \times n}, \cdot)$.⁹

⁶They are both of length k , such that we can compute $a_{il}b_{lj}$ for $l = 1, \dots, k$.

⁷Later, we will call this the **scalar product** or **dot product** of the corresponding row and column.

⁸This kind of element-wise multiplication appears often in computer science where we multiply (multi-dimensional) arrays with each other.

⁹If $A \in \mathbb{R}^{m \times n}$ then I_n is only a right neutral element, such that $AI_n = A$. The corresponding left-neutral element would be I_m since $I_m A = A$.

Properties

- Associativity: $\forall A \in \mathbb{R}^{m \times n}, B \in \mathbb{R}^{n \times p}, C \in \mathbb{R}^{p \times q} : (AB)C = A(BC)$
- Distributivity: $\forall A_1, A_2 \in \mathbb{R}^{m \times n}, B \in \mathbb{R}^{n \times p} : (A_1 + A_2)B = A_1B + A_2B$
 $A(B + C) = AB + AC$
- $\forall A \in \mathbb{R}^{m \times n} : I_m A = A I_n = A$. Note that $I_m \neq I_n$ for $m \neq n$.

5.3.2 Inverse and Transpose**Definition 3 (Inverse)**

For a square matrix¹⁰ $A \in \mathbb{R}^{n \times n}$ a matrix $B \in \mathbb{R}^{n \times n}$ with $AB = I_n = BA$ is called **inverse** and denoted by A^{-1} .

Not every matrix A possesses an inverse A^{-1} . If this inverse does exist, A is called **regular/invertible**, otherwise **singular**. We will discuss these properties much more later on in the course.

Remark 3

The set of regular (invertible) matrices $A \in \mathbb{R}^{n \times n}$ is a group with respect to matrix multiplication as defined in (5.11) and is called **general linear group** $GL(n, \mathbb{R})$.

Definition 4 (Transpose)

For $A \in \mathbb{R}^{m \times n}$ the matrix $B \in \mathbb{R}^{n \times m}$ with $b_{ij} = a_{ji}$ is called the **transpose** of A . We write $B = A^\top$.

For a square matrix A^\top is the matrix we obtain when we “mirror” A on its main diagonal.¹¹ In general, A^\top can be obtained by writing the columns of A as the rows of A^\top .

Remark 4

- $AA^{-1} = I = A^{-1}A$
- $(AB)^{-1} = B^{-1}A^{-1}$
- $(A^\top)^\top = A$
- $(A + B)^\top = A^\top + B^\top$
- $(AB)^\top = B^\top A^\top$
- If A is invertible, $(A^{-1})^\top = (A^\top)^{-1}$
- Note: $(A + B)^{-1} \neq A^{-1} + B^{-1}$. Example: in the scalar case $\frac{1}{2+4} = \frac{1}{6} \neq \frac{1}{2} + \frac{1}{4}$.

A is **symmetric** if $A = A^\top$. Note that this can only hold for (n, n) -matrices (**quadratic matrices**). The sum of symmetric matrices is symmetric, but this does not hold for the product in general (although it is always defined). A counterexample is

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}. \quad (5.16)$$

¹⁰The number columns equals the number of rows.

¹¹The main diagonal (sometimes principal diagonal, primary diagonal, leading diagonal, or major diagonal) of a matrix A is the collection of entries A_{ij} where $i = j$.

5.3.3 Multiplication by a Scalar

Let $A \in \mathbb{R}^{m \times n}$ and $\lambda \in \mathbb{R}$. Then $\lambda A = K$, $K_{ij} = \lambda a_{ij}$. Practically, λ scales each element of A . For $\lambda, \psi \in \mathbb{R}$ it holds:

- Distributivity:
 $(\lambda + \psi)C = \lambda C + \psi C$, $C \in \mathbb{R}^{m \times n}$
 $\lambda(B + C) = \lambda B + \lambda C$, $B, C \in \mathbb{R}^{m \times n}$
- Associativity:
 $(\lambda\psi)C = \lambda(\psi C)$, $C \in \mathbb{R}^{m \times n}$
 $\lambda(BC) = (\lambda B)C = B(\lambda C)$, $B \in \mathbb{R}^{m \times n}, C \in \mathbb{R}^{n \times k}$.
 Note that this allows us to move scalar values around.
- $(\lambda C)^T = C^T \lambda^T = C^T \lambda = \lambda C^T$ since $\lambda = \lambda^T$ for all $\lambda \in \mathbb{R}$.

5.3.4 Compact Representations of Linear Equation Systems

If we consider a linear equation system

$$\begin{aligned} 2x_1 + 3x_2 + 5x_3 &= 1 \\ 4x_1 - 2x_2 - 7x_3 &= 8 \\ 9x_1 + 5x_2 - 3x_3 &= 2 \end{aligned}$$

and use the rules for matrix multiplication, we can write this equation system in a more compact form as

$$\begin{bmatrix} 2 & 3 & 5 \\ 4 & -2 & -7 \\ 9 & 5 & -3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 1 \\ 8 \\ 2 \end{bmatrix}. \quad (5.17)$$

Note that x_1 scales the first column, x_2 the second one, and x_3 the third one. Generally, linear equation systems can be compactly represented in their matrix form as $Ax = b$, see (5.3), and the product Ax is a (linear) combination of the columns of A .¹²

5.4 Solving Linear Equation Systems via Gaussian Elimination

In (5.3), we have introduced the general form of an equation system, i.e.,

$$\begin{aligned} a_{11}x_1 + \cdots + a_{1n}x_n &= b_1 \\ &\vdots \quad \quad \quad \vdots \\ a_{m1}x_1 + \cdots + a_{mn}x_n &= b_m \end{aligned} \quad (5.18)$$

¹²We will discuss linear combinations in Section 5.5.

where $a_{ij} \in \mathbb{R}$ and $b_i \in \mathbb{R}$ are known constants and x_j are unknowns, $i = 1, \dots, m$, $j = 1, \dots, n$. Thus far, we have introduced matrices as a compact way of formulating linear equation systems, i.e., such that we can write $A\mathbf{x} = \mathbf{b}$, see (5.8). Moreover, we defined basic matrix operations, such as addition and multiplication of matrices. In the following, we will introduce a constructive and systematic way of solving linear equation systems. Before doing this, we introduce the **augmented matrix** $[A|\mathbf{b}]$ of the linear equation system $A\mathbf{x} = \mathbf{b}$. This augmented matrix will turn out to be useful when solving linear equation systems.

5.4.1 Example: Solving a Simple Linear Equation System

Now we are turning towards solving linear equation systems. Before doing this in a systematic way using Gaussian elimination, let us have a look at an example. Consider the following linear equation system:

$$\begin{bmatrix} 1 & 0 & 8 & -4 \\ 0 & 1 & 2 & 12 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} 42 \\ 8 \end{bmatrix}. \quad (5.19)$$

This equation system is in a particularly easy form, where the first two columns consist of a 1 and a 0.¹³ Remember that we want to find scalars x_1, \dots, x_4 , such that $\sum_{i=1}^4 x_i \mathbf{c}_i = \mathbf{b}$, where we define \mathbf{c}_i to be the i th column of the matrix and \mathbf{b} the right-hand-side of (5.19). A solution to the problem in (5.19) can be found immediately by taking 42 times the first column and 8 times the second column, i.e.,

$$\mathbf{b} = \begin{bmatrix} 42 \\ 8 \end{bmatrix} = 42 \begin{bmatrix} 1 \\ 0 \end{bmatrix} + 8 \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \quad (5.20)$$

Therefore, one solution vector is $[42, 8, 0, 0]^\top$. This solution is called a **particular solution** or **special solution**. However, this is not the only solution of this linear equation system. To capture all the other solutions, we need to be creative of generating $\mathbf{0}$ in a non-trivial way using the columns of the matrix: Adding a couple of $\mathbf{0}$ s to our special solution does not change the special solution. To do so, we express the third column using the first two columns (which are of this very simple form):

$$\begin{bmatrix} 8 \\ 2 \end{bmatrix} = 8 \begin{bmatrix} 1 \\ 0 \end{bmatrix} + 2 \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad (5.21)$$

such that $\mathbf{0} = 8\mathbf{c}_1 + 2\mathbf{c}_2 - 1\mathbf{c}_3 + 0\mathbf{c}_4$. In fact, any scaling of this solution produces the $\mathbf{0}$ vector:

$$\lambda_1 \begin{bmatrix} 8 \\ 2 \\ -1 \\ 0 \end{bmatrix} = \mathbf{0}, \quad \lambda_1 \in \mathbb{R}. \quad (5.22)$$

¹³Later, we will say that this matrix is in reduced row echelon form.

Following the same line of reasoning, we express the fourth column of the matrix in (5.19) using the first two columns and generate another set of non-trivial versions of $\mathbf{0}$ as

$$\lambda_2 \begin{bmatrix} -4 \\ 12 \\ 0 \\ -1 \end{bmatrix} = \mathbf{0}, \quad \lambda_2 \in \mathbb{R}. \quad (5.23)$$

Putting everything together, we obtain all solutions of the linear equation system in (5.19), which is called the **general solution**, as

$$\begin{bmatrix} 42 \\ 8 \\ 0 \\ 0 \end{bmatrix} + \lambda_1 \begin{bmatrix} 8 \\ 2 \\ -1 \\ 0 \end{bmatrix} + \lambda_2 \begin{bmatrix} -4 \\ 12 \\ 0 \\ -1 \end{bmatrix}. \quad (5.24)$$

Remark 5

- *The general approach we followed consisted of the following three steps:*
 1. *Find a particular solution to $A\mathbf{x} = \mathbf{b}$*
 2. *Find all solutions to $A\mathbf{x} = \mathbf{0}$*
 3. *Combine the solutions from 1. and 2. to the general solution.*
- *Neither the general nor the particular solution is unique.*

The linear equation system in the example above was easy to solve because the matrix in (5.19) has this particularly convenient form, which allowed us to find the particular and the general solution by inspection. However, general equation systems are not of this simple form. Fortunately, there exists a constructive way of transforming any linear equation system into this particularly simple form: **Gaussian elimination**.

The rest of this section will introduce Gaussian elimination, which will allow us to solve all kinds of linear equation systems by first bringing them into a simple form and then applying the three steps to the simple form that we just discussed in the context of the example in (5.19), see Remark 5.

5.4.2 Elementary Transformations

Key to solving linear equation systems are **elementary transformations** that keep the solution set the same¹⁴, but that transform the equation system into a simpler form:

- Exchange of two equations (or: rows in the matrix representing the equation system)

¹⁴Therefore, the original and the modified equation system are **equivalent**.

- Multiplication of an equation (row) with a constant $\lambda \in \mathbb{R} \setminus \{0\}$
- Addition an equation (row) to another equation (row)

Example

We want to find the solutions of the following system of equations:

$$\begin{aligned} -2x_1 + 4x_2 - 2x_3 - x_4 + 4x_5 &= -3 \\ 4x_1 - 8x_2 + 3x_3 - 3x_4 + x_5 &= 2 \\ x_1 - 2x_2 + x_3 - x_4 + x_5 &= 0 \\ x_1 - 2x_2 - 3x_4 + 4x_5 &= a \end{aligned}, \quad a \in \mathbb{R} \quad (5.25)$$

We start by converting this system of equations into the compact matrix notation $Ax = b$. We no longer mention the variables x explicitly and build the **augmented matrix**

$$\left[\begin{array}{ccccc|c} -2 & 4 & -2 & -1 & 4 & -3 \\ 4 & -8 & 3 & -3 & 1 & 2 \\ 1 & -2 & 1 & -1 & 1 & 0 \\ 1 & -2 & 0 & -3 & 4 & a \end{array} \right] \begin{array}{l} \text{Swap with } R_3 \\ \\ \text{Swap with } R_1 \\ \end{array}$$

where we used the vertical line to separate the left-hand-side from the right-hand-side in (5.25). Swapping rows 1 and 3 leads to

$$\left[\begin{array}{ccccc|c} 1 & -2 & 1 & -1 & 1 & 0 \\ 4 & -8 & 3 & -3 & 1 & 2 \\ -2 & 4 & -2 & -1 & 4 & -3 \\ 1 & -2 & 0 & -3 & 4 & a \end{array} \right] \begin{array}{l} \\ -4R_1 \\ +2R_1 \\ -R_1 \end{array}$$

When we now apply the indicated transformations (e.g., subtract Row 1 4 times from Row 2), we obtain

$$\begin{aligned} & \left[\begin{array}{ccccc|c} 1 & -2 & 1 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 & -3 & 2 \\ 0 & 0 & 0 & -3 & 6 & -3 \\ 0 & 0 & -1 & -2 & 3 & a \end{array} \right] \begin{array}{l} \\ \\ -R_2 \\ \end{array} \quad \rightsquigarrow \quad \left[\begin{array}{ccccc|c} 1 & -2 & 1 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 & -3 & 2 \\ 0 & 0 & 0 & -3 & 6 & -3 \\ 0 & 0 & 0 & -3 & 6 & a-2 \end{array} \right] \begin{array}{l} \\ \\ \\ -R_3 \end{array} \\ & \rightsquigarrow \left[\begin{array}{ccccc|c} 1 & -2 & 1 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 & -3 & 2 \\ 0 & 0 & 0 & -3 & 6 & -3 \\ 0 & 0 & 0 & 0 & 0 & a+1 \end{array} \right] \begin{array}{l} \\ \cdot(-1) \\ \cdot(-\frac{1}{3}) \\ \end{array} \\ & \rightsquigarrow \left[\begin{array}{ccccc|c} 1 & -2 & 1 & -1 & 1 & 0 \\ 0 & 0 & 1 & -1 & 3 & -2 \\ 0 & 0 & 0 & 1 & -2 & 1 \\ 0 & 0 & 0 & 0 & 0 & a+1 \end{array} \right] \end{aligned}$$

This (augmented) matrix is in a convenient form, the **row-echelon form (REF)**. Reverting this compact notation back into the explicit notation with the variables we

seek, we obtain

$$\begin{array}{rccccrcr} x_1 & - & 2x_2 & + & x_3 & - & x_4 & + & x_5 & = & 0 \\ & & & & x_3 & - & x_4 & + & 3x_5 & = & -2 \\ & & & & & & x_4 & - & 2x_5 & = & 1 \\ & & & & & & & & 0 & = & a+1 \end{array} \quad (5.26)$$

Only for $a = -1$, this equation system can be solved. A **particular solution** is given by

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} = \begin{bmatrix} 2 \\ 0 \\ -1 \\ 1 \\ 0 \end{bmatrix} \quad (5.27)$$

and the **general solution**, which captures the set of all possible solutions, is given as

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} = \begin{bmatrix} 2 \\ 0 \\ -1 \\ 1 \\ 0 \end{bmatrix} + \lambda_1 \begin{bmatrix} 2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \lambda_2 \begin{bmatrix} 2 \\ 0 \\ -1 \\ 2 \\ 1 \end{bmatrix}, \quad \lambda_1, \lambda_2 \in \mathbb{R} \quad (5.28)$$

Remark 6 (Pivots and Staircase Structure)

The leading coefficient of a row (the first nonzero number from the left) is called **pivot** and is always strictly to the right of the leading coefficient of the row above it. This ensures that an equation system in row echelon form always has a “staircase” structure.

Definition 5 (Row-Echelon Form)

A matrix is in **row-echelon form** (REF) if

- All nonzero rows (rows with at least one nonzero element) are above any rows of all zeros (all zero rows, if any, belong at the bottom of the matrix), and
- The leading coefficient (the first nonzero number from the left, also called the pivot) of a nonzero row is always strictly to the right of the leading coefficient of the row above it¹⁵.

Remark 7 (Basic and Free Variables)

The variables corresponding to the pivots in the row-echelon form are called **basic variables**, the other variables are **free variables**. For example, in (5.26), x_1, x_3, x_4 are basic variables, whereas x_2, x_5 are free variables.

¹⁵In some literature, it is required that the leading coefficient is 1.

Remark 8 (Obtaining a Particular Solution)

The row echelon form makes our lives easier when we need to determine a particular solution. To do this, we express the right-hand side of the equation system using the pivot columns, such that $\mathbf{b} = \sum_{i=1}^P \lambda_i \mathbf{p}_i$, where \mathbf{p}_i , $i = 1, \dots, P$, are the pivot columns. The λ_i are determined easiest if we start with the most-right pivot column and work our way to the left.

In the above example, we would try to find $\lambda_1, \lambda_2, \lambda_3$ such that

$$\lambda_1 \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \lambda_2 \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} + \lambda_3 \begin{bmatrix} -1 \\ -1 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ -2 \\ 1 \\ 0 \end{bmatrix} \quad (5.29)$$

From here, we find relatively directly that $\lambda_3 = 1, \lambda_2 = -1, \lambda_1 = 2$. When we put everything together, we must not forget the non-pivot columns for which we set the coefficients implicitly to 0. Therefore, we get the particular solution

$$\mathbf{x} = \begin{bmatrix} 2 \\ 0 \\ -1 \\ 0 \end{bmatrix}. \quad (5.30)$$

Example

In the following, we will go through solving a linear equation system in matrix form. Consider the problem of finding $\mathbf{x} = [x_1, x_2, x_3]^T$, such that $\mathbf{A}\mathbf{x} = \mathbf{b}$, where

$$\mathbf{A} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 4 \\ 6 \\ 8 \end{bmatrix} \quad (5.31)$$

First, we write down the augmented matrix $[\mathbf{A} | \mathbf{b}]$, which is given by

$$\left[\begin{array}{ccc|c} 1 & 2 & 3 & 4 \\ 4 & 5 & 6 & 6 \\ 7 & 8 & 9 & 8 \end{array} \right],$$

which we now transform into row echelon form using the elementary row operations:

$$\begin{aligned} & \left[\begin{array}{ccc|c} 1 & 2 & 3 & 4 \\ 4 & 5 & 6 & 6 \\ 7 & 8 & 9 & 8 \end{array} \right] \begin{array}{l} -4R_1 \\ -7R_1 \end{array} \quad \rightsquigarrow \quad \left[\begin{array}{ccc|c} 1 & 2 & 3 & 4 \\ 0 & -3 & -6 & -10 \\ 0 & -6 & -12 & -20 \end{array} \right] \begin{array}{l} \cdot(-\frac{1}{3}) \\ -2R_2 \end{array} \\ & \rightsquigarrow \left[\begin{array}{ccc|c} 1 & 2 & 3 & 4 \\ 0 & 1 & 2 & \frac{10}{3} \\ 0 & 0 & 0 & 0 \end{array} \right] \end{aligned}$$

From the row echelon form, we see that x_3 is a free variable. To find a particular solution, we can set x_3 to any real number. For convenience, we choose $x_3 = 0$, but any other number would have worked. With $x_3 = 0$, we obtain a particular solution

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} -\frac{8}{3} \\ \frac{10}{3} \\ 0 \end{bmatrix}. \quad (5.32)$$

To find the general solution, we combine the particular solution with the solution of the homogeneous equation system $A\mathbf{x} = \mathbf{0}$. There are two ways of getting there: the matrix view and the equation system view. Looking at it from the matrix perspective, we need to express the third column of the row-echelon form in terms of the first two columns. This can be done by seeing that

$$\begin{bmatrix} 3 \\ 2 \\ 0 \end{bmatrix} = -\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + 2\begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix} \Leftrightarrow -\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + 2\begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix} - \begin{bmatrix} 3 \\ 2 \\ 0 \end{bmatrix} = \mathbf{0}. \quad (5.33)$$

We now take the coefficients $[-1, 2, -1]^\top$ of these columns that are a non-trivial representation of $\mathbf{0}$ as the solution (and any multiple of it) to the homogeneous equation system $A\mathbf{x} = \mathbf{0}$.

An alternative and equivalent way is to remember that we wanted to solve a linear equation system, we find the solution to the homogeneous equation system by expressing x_3 in terms of x_1, x_2 . From the row echelon form, we see that $x_2 + 2x_3 = 0 \Rightarrow x_3 = -\frac{1}{2}x_2$. With this, we now look at the first set of equations and obtain $x_1 + 2x_2 + 3x_3 = 0 \Rightarrow x_1 - x_3 = 0 \Rightarrow x_3 = x_1$.

Independent of whether we use the matrix or equation system view, we arrive at the general solution

$$\begin{bmatrix} -\frac{8}{3} \\ \frac{10}{3} \\ 0 \end{bmatrix} + \lambda \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}, \lambda \in \mathbb{R}. \quad (5.34)$$

Remark 9 (Reduced Row Echelon Form)

An equation system is in **reduced row echelon form**¹⁶ if

- It is in row echelon form
- Every pivot must be 1
- The pivot is the only non-zero entry in its column.

The reduced row echelon form will play an important role later in this section (Section 5.4.3) because it allows us to determine the general solution of a linear equation system in a straightforward way.

¹⁶also: **row-reduced echelon form** or **row canonical form**

Example (Reduced Row Echelon Form)

Verify that the following matrix is in reduced row echelon form (the pivots are colored in red):

$$A = \begin{bmatrix} \color{red}{1} & 3 & 0 & 0 & 3 \\ 0 & 0 & \color{red}{1} & 0 & 9 \\ 0 & 0 & 0 & \color{red}{1} & -4 \end{bmatrix} \quad (5.35)$$

The key idea for finding the solutions of $A\mathbf{x} = \mathbf{0}$ is to look at the *non-pivot columns*, which we will need to express as a (linear) combination of the pivot columns. The reduced row echelon form makes this relatively straightforward, and we express the non-pivot columns in terms of sums and multiples of the pivot columns that are on their left: The second column is three times the first column (we can ignore the pivot columns on the right of the second column). Therefore, to obtain $\mathbf{0}$, we need to subtract the second column from three times the first column. Now, we look at the fifth column, which is our second non-pivot column. The fifth column is given by 3 times the first pivot column, 9 times the second pivot column, and -4 times the third pivot column. We need to keep track of the indices of the pivot columns and translate this into 3 times the first column, 9 times the third pivot column (which is our second pivot column) and -4 times the fourth column (which is the third pivot column). Then we need to subtract the fifth column to obtain $\mathbf{0}$ —in the end, we are still solving a homogeneous equation system.

To summarize, all solutions of $A\mathbf{x} = \mathbf{0}, \mathbf{x} \in \mathbb{R}^5$ are given by

$$\lambda_1 \begin{bmatrix} 3 \\ -1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \lambda_2 \begin{bmatrix} 3 \\ 0 \\ 9 \\ -4 \\ -1 \end{bmatrix}, \quad \lambda_1, \lambda_2 \in \mathbb{R}. \quad (5.36)$$

5.4.3 The Minus-1 Trick for Solving Homogeneous Equation Systems

In the following, we introduce a practical trick for reading out the solutions \mathbf{x} of a homogeneous linear equation system $A\mathbf{x} = \mathbf{0}$, where $A \in \mathbb{R}^{k \times n}, \mathbf{x} \in \mathbb{R}^n$.

To start, we assume that A is in reduced row echelon form without any rows that just contain zeros (e.g., after applying Gaussian elimination), i.e.,

$$A = \begin{bmatrix} 0 & \cdots & 0 & \color{red}{1} & * & \cdots & * & 0 & * & \cdots & * & 0 & * & \cdots & * \\ \vdots & & \vdots & 0 & 0 & \cdots & 0 & \color{red}{1} & * & \cdots & * & \vdots & \vdots & & \vdots \\ \vdots & & \vdots & \vdots & \vdots & & \vdots & 0 & * & \ddots & * & 0 & \vdots & & \vdots \\ \vdots & & \vdots & \vdots & \vdots & & \vdots & \vdots & * & \ddots & * & 0 & \vdots & & \vdots \\ 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & * & \cdots & * & \color{red}{1} & * & \cdots & * \end{bmatrix} \quad (5.37)$$

Note that the columns j_1, \dots, j_k with the pivots (marked red) are the standard unit vectors $e_1, \dots, e_k \in \mathbb{R}^k$.

We now extend this matrix to an $n \times n$ -matrix \tilde{A} by adding $n - k$ rows of the form

$$\begin{bmatrix} 0 & \cdots & 0 & -1 & 0 & \cdots & 0 \end{bmatrix}, \quad (5.38)$$

such that the diagonal of the augmented matrix \tilde{A} contains only 1 or -1 . Then, the columns of \tilde{A} , which contain the -1 as pivots are solutions of the homogeneous equation system $Ax = 0$.¹⁷ To be more precise, these columns form a basis (Section 5.7) of the solution space of $Ax = 0$, which we will later call the **kernel** or **null space** (Section 5.9.1).

Example

Let us revisit the matrix in (5.35), which is already in reduced row echelon form:

$$A = \begin{bmatrix} 1 & 3 & 0 & 0 & 3 \\ 0 & 0 & 1 & 0 & 9 \\ 0 & 0 & 0 & 1 & -4 \end{bmatrix}. \quad (5.39)$$

We now augment this matrix to a 5×5 matrix by adding rows of the form (5.38) at the places where the pivots on the diagonal are missing and obtain

$$\tilde{A} = \begin{bmatrix} 1 & 3 & 0 & 0 & 3 \\ 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 9 \\ 0 & 0 & 0 & 1 & -4 \\ 0 & 0 & 0 & 0 & -1 \end{bmatrix} \quad (5.40)$$

From this form, we can immediately read out the solutions of $Ax = 0$ by taking the columns of \tilde{A} , which contain -1 on the diagonal:

$$\lambda_1 \begin{bmatrix} 3 \\ -1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \lambda_2 \begin{bmatrix} 3 \\ 0 \\ 9 \\ -4 \\ -1 \end{bmatrix}, \quad \lambda_1, \lambda_2 \in \mathbb{R}, \quad (5.41)$$

which is identical to the solution in (5.36) that we obtained by “insight”.

5.4.4 Applications of Gaussian Elimination in Linear Algebra

Gaussian elimination can also be used to find the rank of a matrix (Chapter 5.7), to calculate the determinant of a matrix (Chapter 5.10), the null space, and the inverse

¹⁷The proof of this trick is out of the scope of this course.

of an invertible square matrix. Because of its relevance to central concepts in Linear Algebra, Gaussian elimination is the most important algorithm we will cover. Gaussian elimination is an important prerequisite for CO343 Operations Research. The initial part of the course discusses in depth systems of linear equations with more columns than rows, and the related notions of basic solutions and non-basic solutions.

Calculating the Inverse

To compute the inverse A^{-1} of $A \in \mathbb{R}^{n \times n}$, we need to satisfy $AA^{-1} = I_n$. We can write this down as a set of simultaneous linear equations $AX = I_n$, where we solve for $X = [x_1 | \dots | x_n]$. We use the augmented matrix notation for a compact representation of this set of linear equation systems and obtain

$$[A|I_n] \rightsquigarrow \dots \rightsquigarrow [I_n|A^{-1}]. \quad (5.42)$$

This means that if we bring the augmented equation system into reduced row echelon form, we can read off the inverse on the right-hand side of the equation system.

Example

For $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$, we determine its inverse by solving the following linear equation system:

$$\left[\begin{array}{cc|cc} 1 & 2 & 1 & 0 \\ 3 & 4 & 0 & 1 \end{array} \right]$$

We bring this system now into reduced row echelon form

$$\begin{aligned} & \left[\begin{array}{cc|cc} 1 & 2 & 1 & 0 \\ 3 & 4 & 0 & 1 \end{array} \right] \xrightarrow{-3R_1} \left[\begin{array}{cc|cc} 1 & 2 & 1 & 0 \\ 0 & -2 & -3 & 1 \end{array} \right] \xrightarrow{+R_2 \cdot (-\frac{1}{2})} \\ & \rightsquigarrow \left[\begin{array}{cc|cc} 1 & 0 & -2 & 1 \\ 0 & 1 & \frac{3}{2} & -\frac{1}{2} \end{array} \right] \end{aligned}$$

The right-hand side of this augmented equation system contains the inverse

$$A^{-1} = \begin{bmatrix} -2 & 1 \\ \frac{3}{2} & -\frac{1}{2} \end{bmatrix}. \quad (5.43)$$

Example

To determine the inverse of

$$A = \begin{bmatrix} 1 & 0 & 2 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 2 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} \quad (5.44)$$

we write down the augmented matrix

$$\left[\begin{array}{cccc|cccc} 1 & 0 & 2 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 2 & 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 \end{array} \right]$$

and transform it into reduced row echelon form

$$\left[\begin{array}{cccc|cccc} 1 & 0 & 0 & 0 & -1 & 2 & -2 & 2 \\ 0 & 1 & 0 & 0 & 1 & -1 & 2 & -2 \\ 0 & 0 & 1 & 0 & 1 & -1 & 1 & -1 \\ 0 & 0 & 0 & 1 & -1 & 0 & -1 & 2 \end{array} \right],$$

such that the desired inverse is given as its right-hand side:

$$A^{-1} = \begin{bmatrix} -1 & 2 & -2 & 2 \\ 1 & -1 & 2 & -2 \\ 1 & -1 & 1 & -1 \\ -1 & 0 & -1 & 2 \end{bmatrix}. \quad (5.45)$$

Remark 10

You may have encountered a way of computing the inverse of a matrix using co-factors and/or cross-products. This approach only works in three dimensions and is not used in practice.

5.5 Vector Spaces

When we discussed group theory, we were looking at sets G and inner operations on G , i.e., mappings $G \times G \rightarrow G$. In the following, we will consider sets that in addition to an inner operation $+$ also contain an outer operation \cdot , the multiplication by a scalar $\lambda \in \mathbb{R}$.

Definition 6 (Vector space)

A real-valued **vector space** (also called an \mathbb{R} -vector space) is a set V with two operations

$$+ : V \times V \rightarrow V \quad (5.46)$$

$$\cdot : \mathbb{R} \times V \rightarrow V \quad (5.47)$$

where

1. $(V, +)$ is an Abelian group
2. Distributivity:

$$(a) \lambda \cdot (\mathbf{x} + \mathbf{y}) = \lambda \cdot \mathbf{x} + \lambda \cdot \mathbf{y} \quad \forall \lambda \in \mathbb{R}, \mathbf{x}, \mathbf{y} \in V$$

$$(b) (\lambda + \psi) \cdot \mathbf{x} = \lambda \cdot \mathbf{x} + \psi \cdot \mathbf{x} \quad \forall \lambda, \psi \in \mathbb{R}, \mathbf{x} \in V$$

$$3. \text{ Associativity (outer operation): } \lambda \cdot (\psi \cdot \mathbf{x}) = (\lambda\psi) \cdot \mathbf{x} \quad \forall \lambda, \psi \in \mathbb{R}, \mathbf{x} \in V$$

$$4. \text{ Neutral element with respect to the outer operation: } 1 \cdot \mathbf{x} = \mathbf{x}, \quad \forall \mathbf{x} \in V$$

The elements $\mathbf{x} \in V$ are called **vectors**. The neutral element of $(V, +)$ is the zero vector $\mathbf{0} = [0, \dots, 0]^\top$, and the inner operation $+$ is called **vector addition**. The elements $\lambda \in \mathbb{R}$ are called **scalars** and the outer operation \cdot is a **multiplication by scalars**.¹⁸

Remark 11

When we started the course, we defined vectors as special objects that can be added together and multiplied by scalars to yield another element of the same kind (see p. 36). Examples were geometric vectors, polynomials and \mathbb{R}^n . Definition 6 gives now the corresponding formal definition and applies to all kinds of vectors. We will continue focusing on vectors as elements of \mathbb{R}^n because it is the most general formulation, and most algorithms are formulated in \mathbb{R}^n .

Remark 12

Note that a “vector multiplication” $\mathbf{a}\mathbf{b}$, $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, is not defined. Theoretically, we could define it in two ways: (a) We could define an element-wise multiplication, such that $\mathbf{c} = \mathbf{a} \cdot \mathbf{b}$ with $c_j = a_j b_j$. This “array multiplication” is common to many programming languages but makes mathematically only limited sense; (b) By treating vectors as $n \times 1$ matrices (which we usually do), we can use the matrix multiplication as defined in (5.11). However, then the dimensions of the vectors do not match. Only the following multiplications for vectors are defined: $\mathbf{a}\mathbf{b}^\top$ (outer product), $\mathbf{a}^\top \mathbf{b}$ (inner/scalar product).

5.5.1 Examples

- $V = \mathbb{R}^n, n \in \mathbb{N}$ is a vector space with operations defined as follows:
 - Addition: $\mathbf{x} + \mathbf{y} = (x_1, \dots, x_n) + (y_1, \dots, y_n) = (x_1 + y_1, \dots, x_n + y_n)$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$
 - Multiplication by scalars: $\lambda \mathbf{x} = \lambda(x_1, \dots, x_n) = (\lambda x_1, \dots, \lambda x_n)$ for all $\lambda \in \mathbb{R}, \mathbf{x} \in \mathbb{R}^n$
- $V = \mathbb{R}^{m \times n}, m, n \in \mathbb{N}$ is a vector space with

$$\text{– Addition: } \mathbf{A} + \mathbf{B} = \begin{bmatrix} a_{11} + b_{11} & \cdots & a_{1n} + b_{1n} \\ \vdots & & \vdots \\ a_{m1} + b_{m1} & \cdots & a_{mn} + b_{mn} \end{bmatrix} \text{ is defined elementwise for}$$

all $\mathbf{A}, \mathbf{B} \in V$

¹⁸Note: A **scalar product** is something different, and we will get to this in Section 5.13.

– Multiplication by scalars: $\lambda A = \begin{bmatrix} \lambda a_{11} & \cdots & \lambda a_{1n} \\ \vdots & & \vdots \\ \lambda a_{m1} & \cdots & \lambda a_{mn} \end{bmatrix}$ as defined in Section 5.3.

Remember that $\mathbb{R}^{m \times n}$ is equivalent to \mathbb{R}^{mn} .

- $V = \mathbb{C}$, with the standard definition of addition of complex numbers.

Remark 13 (Notation)

The three vector spaces $\mathbb{R}^n, \mathbb{R}^{n \times 1}, \mathbb{R}^{1 \times n}$ are only different with respect to the way of writing. In the following, we will not make a distinction between \mathbb{R}^n and $\mathbb{R}^{n \times 1}$, which allows us to write n -tuples as **column vectors**

$$\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}. \quad (5.48)$$

This will simplify the notation regarding vector space operations. However, we will distinguish between $\mathbb{R}^{n \times 1}$ and $\mathbb{R}^{1 \times n}$ (the **row vectors**) to avoid confusion with matrix multiplication. By default we write \mathbf{x} to denote a column vector, and a row vector is denoted by \mathbf{x}^\top , the **transpose** of \mathbf{x} .

5.5.2 Generating Set and Vector Subspaces

Definition 7 (Linear Combination)

Consider a vector space V and a finite number of vectors $\mathbf{x}_1, \dots, \mathbf{x}_k \in V$. Then, every vector $\mathbf{v} \in V$ of the form

$$\mathbf{v} = \lambda_1 \mathbf{x}_1 + \cdots + \lambda_k \mathbf{x}_k = \sum_{i=1}^k \lambda_i \mathbf{x}_i \in V \quad (5.49)$$

with $\lambda_1, \dots, \lambda_k \in \mathbb{R}$ is a **linear combination** of the vectors $\mathbf{x}_1, \dots, \mathbf{x}_k$.

Definition 8 (Generating Set/Span)

Consider an \mathbb{R} -vector space V and $A = \{\mathbf{x}_1, \dots, \mathbf{x}_k\} \subset V$. If every vector $\mathbf{v} \in V$ can be expressed as a linear combination of $\mathbf{x}_1, \dots, \mathbf{x}_k$, A is called a **generating set** or **span**, which spans the vector space V . In this case, we write $V = [A]$ or $V = [\mathbf{x}_1, \dots, \mathbf{x}_k]$.

Definition 9 (Vector Subspace)

Let V be an \mathbb{R} -vector space and $U \subset V$, $U \neq \emptyset$. U is called **vector subspace** of V (or **linear subspace**) if U is a vector space with the vector space operations $+$ and \cdot restricted to $U \times U$ and $\mathbb{R} \times U$.

Remark 14

If $U \subset V$ and V is a vector space, then U naturally inherits many properties directly from V because they are true for all $\mathbf{x} \in V$, and in particular for all $\mathbf{x} \in U \subset V$. This includes the Abelian group properties, the distributivity, the associativity and the neutral element. What we still do need to show is

1. $U \neq \emptyset$, in particular: $\mathbf{0} \in U$
2. Closure of U :
 - (a) With respect to the outer operation: $\forall \lambda \in \mathbb{R} \forall \mathbf{x} \in U : \lambda \mathbf{x} \in U$
 - (b) With respect to the inner operation: $\forall \mathbf{x}, \mathbf{y} \in U : \mathbf{x} + \mathbf{y} \in U$.

Examples

- For every vector space V the trivial subspaces are V itself and $\{\mathbf{0}\}$.
- The solution set of a homogeneous linear equation system $A\mathbf{x} = \mathbf{0}$ with n unknowns $\mathbf{x} = [x_1, \dots, x_n]^\top$ is a subspace of \mathbb{R}^n .
- However, the solution of an inhomogeneous equation system $A\mathbf{x} = \mathbf{b}$, $\mathbf{b} \neq \mathbf{0}$ is not a subspace of \mathbb{R}^n .
- The intersection of arbitrarily many subspaces is a subspace itself.
- The intersection of all subspaces $U_i \subset V$ is called **linear hull** of V .

Remark 15

Every subspace $U \subset \mathbb{R}^n$ is the solution space of a homogeneous linear equation system $A\mathbf{x} = \mathbf{0}$.

5.6 Linear (In)Dependence

In Section 5.5, we learned about linear combinations of vectors, see (5.49). The $\mathbf{0}$ vector can always be written as the linear combination of k vectors $\mathbf{x}_1, \dots, \mathbf{x}_k$ because $\mathbf{0} = \sum_{i=1}^k 0\mathbf{x}_i$ is always true. In the following, we are interested in non-trivial linear combinations of a set of vectors to represent $\mathbf{0}$.

Definition 10 (Linear (In)dependence)

Let us consider a vector space V with $k \in \mathbb{N}$ and $\mathbf{x}_1, \dots, \mathbf{x}_k \in V$. If there is a non-trivial linear combination, such that $\mathbf{0} = \sum_{i=1}^k \lambda_i \mathbf{x}_i$ with at least one $\lambda_i \neq 0$, the vectors $\mathbf{x}_1, \dots, \mathbf{x}_k$ are **linearly dependent**. If only the trivial solution exists, i.e., $\lambda_1 = \dots = \lambda_k = 0$ the vectors $\mathbf{x}_1, \dots, \mathbf{x}_k$ are **linearly independent**.

Intuitively, a set of linearly dependent vectors contains some redundancy, whereas linearly independent vectors are all essential. Throughout this chapter, we will formalize this intuition more.

Example From Wikipedia (2015)

A geographic example may help to clarify the concept of linear independence. A person describing the location of a certain place might say, "It is 3 miles North and 4 miles East of here." This is sufficient information to describe the location, because the geographic coordinate system may be considered as a 2-dimensional vector space

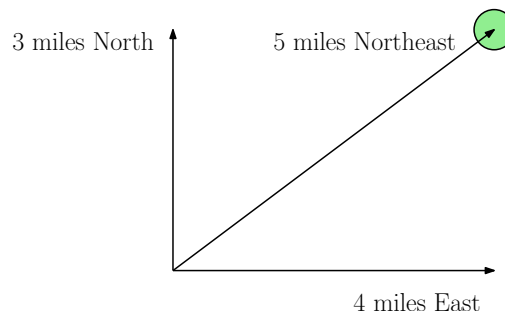


Figure 5.2: Linear dependence of three vectors in a two-dimensional space (plane).

(ignoring altitude and the curvature of the Earth’s surface). The person might add, “The place is 5 miles Northeast of here.” Although this last statement is true, it is not necessary to find this place (see Fig. 5.2 for an illustration).

In this example, the “3 miles North” vector and the “4 miles East” vector are linearly independent. That is to say, the north vector cannot be described in terms of the east vector, and vice versa. The third “5 miles Northeast” vector is a linear combination of the other two vectors, and it makes the set of vectors linearly dependent, that is, one of the three vectors is unnecessary.

Remark 16

The following properties are useful to find out whether vectors are linearly independent.

- *k vectors are either linearly dependent or linearly independent. There is no third option.*
- *If at least one of the vectors $\mathbf{x}_1, \dots, \mathbf{x}_k$ is $\mathbf{0}$ then they are linearly dependent. The same holds if two vectors are identical.*
- *The vectors $\mathbf{x}_1, \dots, \mathbf{x}_k \neq \mathbf{0}$, $k \geq 2$, are linearly dependent if and only if (at least) one of them is a linear combination of the others.*
- *A practical way of checking whether vectors $\mathbf{x}_1, \dots, \mathbf{x}_k \in V$ are linearly independent is to use Gaussian elimination: Write all vectors as columns of a matrix A . Gaussian elimination yields a matrix in (reduced) row echelon form.*
 - *The pivot columns indicate the vectors, which are linearly independent of the previous¹⁹ vectors (note that there is an ordering of vectors when the matrix is built).*
 - *The non-pivot columns can be expressed as linear combinations of the pivot columns on their left. For instance, in*

$$\begin{bmatrix} 1 & 3 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (5.50)$$

¹⁹the vectors on the left

the first and third column are pivot columns. The second column is a non-pivot column because it is 3 times the first column.

If all columns are pivot columns, the column vectors are linearly independent. If there is at least one non-pivot column, the columns are linearly dependent.

Example

Consider \mathbb{R}^4 with

$$\mathbf{x}_1 = \begin{bmatrix} 1 \\ 2 \\ -3 \\ 4 \end{bmatrix}, \quad \mathbf{x}_2 = \begin{bmatrix} 1 \\ 1 \\ 0 \\ 2 \end{bmatrix}, \quad \mathbf{x}_3 = \begin{bmatrix} -1 \\ -2 \\ 1 \\ 1 \end{bmatrix}. \quad (5.51)$$

To check whether they are linearly dependent, we follow the general approach and solve

$$\lambda_1 \mathbf{x}_1 + \lambda_2 \mathbf{x}_2 + \lambda_3 \mathbf{x}_3 = \lambda_1 \begin{bmatrix} 1 \\ 2 \\ -3 \\ 4 \end{bmatrix} + \lambda_2 \begin{bmatrix} 1 \\ 1 \\ 0 \\ 2 \end{bmatrix} + \lambda_3 \begin{bmatrix} -1 \\ -2 \\ 1 \\ 1 \end{bmatrix} = \mathbf{0} \quad (5.52)$$

for $\lambda_1, \dots, \lambda_3$. We write the vectors \mathbf{x}_i , $i = 1, 2, 3$, as the columns of a matrix and apply Gaussian elimination.

$$\begin{array}{ccc} \left[\begin{array}{ccc|c} 1 & 1 & -1 & 0 \\ 2 & 1 & -2 & 0 \\ -3 & 0 & 1 & 0 \\ 4 & 2 & 1 & 0 \end{array} \right] \begin{array}{l} -2R_1 \\ +3R_1 \\ -4R_1 \end{array} & \rightsquigarrow & \left[\begin{array}{ccc|c} 1 & 1 & -1 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 3 & -2 & 0 \\ 0 & -2 & 5 & 0 \end{array} \right] \begin{array}{l} +3R_2 \\ -2R_2 \end{array} \\ \\ \rightsquigarrow \left[\begin{array}{ccc|c} 1 & 1 & -1 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -2 & 0 \\ 0 & 0 & 5 & 0 \end{array} \right] \begin{array}{l} \cdot(-1) \\ \cdot(-\frac{1}{2}) \\ +\frac{5}{2}R_3 \end{array} & \rightsquigarrow & \left[\begin{array}{ccc|c} 1 & 1 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{array} \right] \end{array}$$

Here, every column of the matrix is a pivot column²⁰, i.e., every column is linearly independent of the columns on its left. Therefore, there is no non-trivial solution, and we require $\lambda_1 = 0, \lambda_2 = 0, \lambda_3 = 0$. Hence, the vectors $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3$ are linearly independent.

Remark 17

Consider an \mathbb{R} -vector space V with k linearly independent vectors $\mathbf{b}_1, \dots, \mathbf{b}_k$ and m linear combinations

$$\mathbf{x}_1 = \sum_{i=1}^k \lambda_{i1} \mathbf{b}_i, \quad (5.53)$$

²⁰Note that the matrix is not in reduced row echelon form; it also does not need to be.

$$\vdots \quad (5.54)$$

$$\mathbf{x}_m = \sum_{i=1}^k \lambda_{im} \mathbf{b}_i. \quad (5.55)$$

We want to test whether $\mathbf{x}_1, \dots, \mathbf{x}_m$ are linearly independent. For this purpose, we follow the general approach of testing when $\sum_{j=1}^m \psi_j \mathbf{x}_j = \mathbf{0}$ and obtain

$$\mathbf{0} = \sum_{j=1}^m \psi_j \mathbf{x}_j = \sum_{j=1}^m \psi_j \left(\sum_{i=1}^k \lambda_{ij} \mathbf{b}_i \right) = \sum_{j=1}^m \sum_{i=1}^k \psi_j \lambda_{ij} \mathbf{b}_i = \sum_{i=1}^k \left(\sum_{j=1}^m \psi_j \lambda_{ij} \right) \mathbf{b}_i. \quad (5.56)$$

Therefore, $\mathbf{x}_1, \dots, \mathbf{x}_m$ are linearly independent **if and only if** the column vectors

$$\boldsymbol{\lambda}_1 = \begin{bmatrix} \lambda_{11} \\ \vdots \\ \lambda_{k1} \end{bmatrix} \in \mathbb{R}^k, \quad \dots, \quad \boldsymbol{\lambda}_m = \begin{bmatrix} \lambda_{1m} \\ \vdots \\ \lambda_{km} \end{bmatrix} \in \mathbb{R}^k \quad (5.57)$$

are linearly independent.

Proof 1

Since $\mathbf{b}_1, \dots, \mathbf{b}_k$ are linearly independent it follows that for all $j = 1, \dots, m$ we get $\sum_{i=1}^k \lambda_{ij} \mathbf{b}_i = \mathbf{0}$ if and only if $\lambda_{ij} = 0, i = 1, \dots, k$. Therefore, $\sum_{j=1}^m \psi_j \lambda_{ij} = 0$ for $i = 1, \dots, k$. This implies $\sum_{i=1}^k (\sum_{j=1}^m \psi_j \lambda_{ij}) \mathbf{b}_i = \mathbf{0}$. Hence, $\sum_{j=1}^m \psi_j \mathbf{x}_j = \mathbf{0}$ is equivalent to $\sum_{j=1}^m \psi_j \boldsymbol{\lambda}_j = \mathbf{0}$.

Remark 18

In a vector space V , m linear combinations of k vectors $\mathbf{x}_1, \dots, \mathbf{x}_k$ are linearly dependent if $m > k$.

Example

Consider a set of linearly independent vectors $\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3, \mathbf{b}_4 \in \mathbb{R}^n$ and

$$\begin{aligned} \mathbf{x}_1 &= \mathbf{b}_1 & - & 2\mathbf{b}_2 & + & \mathbf{b}_3 & - & \mathbf{b}_4 \\ \mathbf{x}_2 &= -4\mathbf{b}_1 & - & 2\mathbf{b}_2 & & & + & 4\mathbf{b}_4 \\ \mathbf{x}_3 &= 2\mathbf{b}_1 & + & 3\mathbf{b}_2 & - & \mathbf{b}_3 & - & 3\mathbf{b}_4 \\ \mathbf{x}_4 &= 17\mathbf{b}_1 & - & 10\mathbf{b}_2 & + & 11\mathbf{b}_3 & + & \mathbf{b}_4 \end{aligned} \quad (5.58)$$

Are the vectors $\mathbf{x}_1, \dots, \mathbf{x}_4 \in \mathbb{R}^n$ linearly independent? To answer this question, we investigate whether the column vectors

$$\left\{ \begin{bmatrix} 1 \\ -2 \\ 1 \\ -1 \end{bmatrix}, \begin{bmatrix} -4 \\ -2 \\ 0 \\ 4 \end{bmatrix}, \begin{bmatrix} 2 \\ 3 \\ -1 \\ -3 \end{bmatrix}, \begin{bmatrix} 17 \\ -10 \\ 11 \\ 1 \end{bmatrix} \right\} \quad (5.59)$$

are linearly independent. The reduced row echelon form of the corresponding linear equation system with coefficient matrix

$$\mathbf{A} = \begin{bmatrix} 1 & -4 & 2 & 17 \\ -2 & -2 & 3 & -10 \\ 1 & 0 & -1 & 11 \\ -1 & 4 & -3 & 1 \end{bmatrix} \quad (5.60)$$

is given as

$$\begin{bmatrix} 1 & 0 & 0 & -7 \\ 0 & 1 & 0 & -15 \\ 0 & 0 & 1 & -18 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (5.61)$$

From the reduced row echelon form, we see that the corresponding linear equation system is non-trivially solvable: The last column is not a pivot column, and $x_4 = -7x_1 - 15x_2 - 18x_3$. Therefore, x_1, \dots, x_4 are linearly dependent as x_4 lies in the span of x_1, \dots, x_3 .

5.7 Basis and Dimension

In a vector space V , we are particularly interested in the set of linearly independent vectors A that possesses the property that any vector $v \in V$ can be obtained by a linear combination of vectors in A . We discussed generating sets previously. Now, we will be more specific and characterize the smallest generating set that spans a vector (sub)space.

Definition 11 (Basis)

Consider a real vector space V and $A \subset V$

- A generating set A of V is called **minimal** if there exists no smaller set $\tilde{A} \subset A \subset V$, which spans V .
- Every linearly independent generating set of V is minimal and is called **basis** of V .

Let V be a real vector space and $B \subset V, B \neq \emptyset$. Then, the following statements are equivalent:

- B is basis of V
- B is a minimal generating set
- B is a maximal linearly independent subset of V .

- Every vector $\mathbf{x} \in V$ is a linear combination of vectors from B , and every linear combination is unique, i.e., with

$$\mathbf{x} = \sum_{i=1}^k \lambda_i \mathbf{b}_i = \sum_{i=1}^k \psi_i \mathbf{b}_i \quad (5.62)$$

and $\lambda_i, \psi_i \in \mathbb{R}$, $\mathbf{b}_i \in B$ it follows that $\lambda_i = \psi_i$, $i = 1, \dots, k$.

Example

- In \mathbb{R}^3 , the canonical/standard basis is

$$B = \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\}. \quad (5.63)$$

- Different bases in \mathbb{R}^3 are

$$B_1 = \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right\}, \quad B_2 = \left\{ \begin{bmatrix} 0.53 \\ 0.86 \\ -0.43 \end{bmatrix}, \begin{bmatrix} 1.83 \\ 0.31 \\ 0.34 \end{bmatrix}, \begin{bmatrix} -2.25 \\ -1.30 \\ 3.57 \end{bmatrix} \right\} \quad (5.64)$$

- The set

$$A = \left\{ \begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \end{bmatrix}, \begin{bmatrix} 2 \\ -1 \\ 0 \\ 2 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \\ -4 \end{bmatrix} \right\} \quad (5.65)$$

is linearly independent, but not a generating set (and no basis): For instance, the vector $[1, 0, 0, 0]^\top$ cannot be obtained by a linear combination of elements in A .

Remark 19

- Every vector space V possesses a basis B .
- The examples above show that there can be many bases of a vector space V , i.e., there is no unique basis. However, all bases possess the same number of elements, the **basis vectors**.
- We only consider finite-dimensional vector spaces V . In this case, the **dimension of V** is the number of basis vectors, and we write $\dim(V)$.

- If $U \subset V$ is a subspace of V then $\dim(U) \leq \dim(V)$ and $\dim(U) = \dim(V)$ if and only if $U = V$.

Remark 20

A basis of a subspace $U = [\mathbf{x}_1, \dots, \mathbf{x}_m] \subset \mathbb{R}^n$ can be found by executing the following steps:

1. Write the spanning vectors as columns of a matrix A
2. Apply Gaussian elimination algorithm to A .
3. The spanning vectors associated with the pivot columns form a basis of U .

Example (Determining a Basis)

For a vector subspace $U \subset \mathbb{R}^5$, spanned by the vectors

$$\mathbf{x}_1 = \begin{bmatrix} 1 \\ 2 \\ -1 \\ -1 \\ -1 \end{bmatrix}, \quad \mathbf{x}_2 = \begin{bmatrix} 2 \\ -1 \\ 1 \\ 2 \\ -2 \end{bmatrix}, \quad \mathbf{x}_3 = \begin{bmatrix} 3 \\ -4 \\ 3 \\ 5 \\ -3 \end{bmatrix}, \quad \mathbf{x}_4 = \begin{bmatrix} -1 \\ 8 \\ -5 \\ -6 \\ 1 \end{bmatrix} \in \mathbb{R}^5, \quad (5.66)$$

we are interested in finding out which vectors $\mathbf{x}_1, \dots, \mathbf{x}_4$ are a basis for U . For this, we need to check whether $\mathbf{x}_1, \dots, \mathbf{x}_4$ are linearly independent. Therefore, we need to solve

$$\sum_{i=1}^4 \lambda_i \mathbf{x}_i = \mathbf{0}, \quad (5.67)$$

which leads to a homogeneous equation system with the corresponding matrix

$$[\mathbf{x}_1 | \mathbf{x}_2 | \mathbf{x}_3 | \mathbf{x}_4] = \begin{bmatrix} 1 & 2 & 3 & -1 \\ 2 & -1 & -4 & 8 \\ -1 & 1 & 3 & -5 \\ -1 & 2 & 5 & -6 \\ -1 & -2 & -3 & 1 \end{bmatrix}. \quad (5.68)$$

With the basic transformation of linear equation systems, we obtain

$$\begin{array}{l} \begin{bmatrix} 1 & 2 & 3 & -1 \\ 2 & -1 & -4 & 8 \\ -1 & 1 & 3 & -5 \\ -1 & 2 & 5 & -6 \\ -1 & -2 & -3 & 1 \end{bmatrix} \begin{array}{l} -2R_1 \\ +R_1 \\ +R_1 \\ +R_1 \end{array} \quad \rightsquigarrow \quad \begin{bmatrix} 1 & 2 & 3 & -1 \\ 0 & -5 & -10 & 10 \\ 0 & 3 & 6 & -6 \\ 0 & 4 & 8 & -7 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{array}{l} \cdot(-\frac{1}{5}) \\ \cdot\frac{1}{3} - R_2 \\ -4R_2 \end{array} \\ \\ \rightsquigarrow \begin{bmatrix} 1 & 2 & 3 & -1 \\ 0 & 1 & 2 & -2 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{array}{l} +R_4 \\ +2R_4 \\ \text{swap with } R_3 \end{array} \quad \rightsquigarrow \quad \begin{bmatrix} 1 & 2 & 3 & 0 \\ 0 & 1 & 2 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{array}{l} -2R_2 \end{array} \end{array}$$

$$\rightsquigarrow \begin{bmatrix} 1 & 0 & -1 & 0 \\ 0 & 1 & 2 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

From this reduced-row echelon form we see that x_1, x_2, x_4 belong to the pivot columns, and, therefore, linearly independent (because the linear equation system $\lambda_1 x_1 + \lambda_2 x_2 + \lambda_4 x_4 = \mathbf{0}$ can only be solved with $\lambda_1 = \lambda_2 = \lambda_4 = 0$). Therefore, $\{x_1, x_2, x_4\}$ is a basis of U .

Example (Finding a Simple Basis)

Let us now consider a slightly different problem: Instead of finding out which vectors x_1, \dots, x_4 of the span of U form a basis, we are interested in finding a “simple” basis for U . Here, “simple” means that we are interested in basis vectors with many coordinates equal to 0.

To solve this problem we replace the vectors x_1, \dots, x_4 with suitable linear combinations. **In practice**, we write x_1, \dots, x_4 as *row vectors* in a matrix and perform Gaussian elimination:

$$\begin{array}{l} \begin{bmatrix} 1 & 2 & -1 & -1 & -1 \\ 2 & -1 & 1 & 2 & -2 \\ 3 & -4 & 3 & 5 & -3 \\ -1 & 8 & -6 & -6 & 1 \end{bmatrix} \rightsquigarrow \begin{bmatrix} 1 & 2 & -1 & -1 & -1 \\ 0 & -5 & 3 & 4 & 0 \\ 0 & -10 & 6 & 8 & 0 \\ 0 & 10 & -7 & -7 & 0 \end{bmatrix} \\ \rightsquigarrow \begin{bmatrix} 1 & 2 & -1 & -1 & -1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & -\frac{3}{5} & -\frac{4}{5} & 0 \\ 0 & 0 & -1 & 1 & 0 \end{bmatrix} \rightsquigarrow \begin{bmatrix} 1 & 2 & -1 & -1 & -1 \\ 0 & 1 & 0 & -\frac{7}{5} & 0 \\ 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \\ \rightsquigarrow \begin{bmatrix} 1 & 2 & 0 & -2 & -1 \\ 0 & 1 & 0 & -\frac{7}{5} & 0 \\ 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \rightsquigarrow \begin{bmatrix} 1 & 0 & 0 & -\frac{4}{5} & -1 \\ 0 & 1 & 0 & -\frac{7}{5} & 0 \\ 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \end{array}$$

From the reduced row echelon form, the simple basis vectors are the rows with the leading 1s (the “steps”).

$$U = \left[\underbrace{\begin{bmatrix} 1 \\ 0 \\ 0 \\ -\frac{4}{5} \\ -1 \end{bmatrix}}_{b_1}, \underbrace{\begin{bmatrix} 0 \\ 1 \\ 0 \\ -\frac{7}{5} \\ 0 \end{bmatrix}}_{b_2}, \underbrace{\begin{bmatrix} 0 \\ 0 \\ 1 \\ -1 \\ 0 \end{bmatrix}}_{b_3} \right] \quad (5.69)$$

and $B = \{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}$ is a (simple) basis of U (check linear independence).

5.7.1 Rank

- The number of linearly independent columns of a matrix $A \in \mathbb{R}^{m \times n}$ equals the number of linearly independent rows and is called **rank** of A and is denoted by $\text{rk}(A)$.
- $\text{rk}(A) = \text{rk}(A^\top)$, i.e., the column rank equals the row rank.
- The columns of $A \in \mathbb{R}^{m \times n}$ span a subspace $U \subset \mathbb{R}^m$ with $\dim(U) = \text{rk}(A)$. A basis of U can be found by applying Gaussian elimination to A to identify the pivot columns.
- The rows of $A \in \mathbb{R}^{m \times n}$ span a subspace $W \subset \mathbb{R}^n$ with $\dim(W) = \text{rk}(A)$. A basis of W can be found by applying the Gaussian elimination algorithm to A^\top .
- For all $A \in \mathbb{R}^{n \times n}$ holds: A is regular (invertible) if and only if $\text{rk}(A) = n$.
- For all $A \in \mathbb{R}^{m \times n}$ and all $\mathbf{b} \in \mathbb{R}^m$: The linear equation system $A\mathbf{x} = \mathbf{b}$ can be solved if and only if $\text{rk}(A) = \text{rk}(A|\mathbf{b})$, where $A|\mathbf{b}$ denotes the augmented system.
- For $A \in \mathbb{R}^{m \times n}$ the subspace of solutions for $A\mathbf{x} = \mathbf{0}$ possesses dimension $n - \text{rk}(A)$.²¹
- A matrix $A \in \mathbb{R}^{m \times n}$ has **full rank** if its rank equals the largest possible for a matrix of the same dimensions, which is the lesser of the number of rows and columns, i.e., $\text{rk}(A) = \min(m, n)$. A matrix is said to be **rank deficient** if it does not have full rank.

Example

- $A = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix}$. A possesses two linearly independent rows (and columns).
Therefore, $\text{rk}(A) = 2$.
- $A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 8 & 12 \end{bmatrix}$. We see that the second row is a multiple of the first row, such that the row-echelon form of A is $\begin{bmatrix} 1 & 2 & 3 \\ 0 & 0 & 0 \end{bmatrix}$, and $\text{rk}(A) = 1$.

²¹Later, we will be calling this subspace **kernel** or **nullspace**.

• $A = \begin{bmatrix} 1 & 2 & 1 \\ -2 & -3 & 1 \\ 3 & 5 & 0 \end{bmatrix}$ We use Gaussian elimination to determine the rank:

$$\begin{aligned} & \begin{bmatrix} 1 & 2 & 1 \\ -2 & -3 & 1 \\ 3 & 5 & 0 \end{bmatrix} +R_1 - R_2 \quad \rightsquigarrow \quad \begin{bmatrix} 1 & 2 & 1 \\ -2 & -3 & 1 \\ 0 & 0 & 0 \end{bmatrix} +2R_1 \\ & \rightsquigarrow \begin{bmatrix} 1 & 2 & 1 \\ 0 & -1 & 3 \\ 0 & 0 & 0 \end{bmatrix} \end{aligned}$$

Here, we see that the number of linearly independent rows and columns is 2, such that $\text{rk}(A) = 2$.

5.8 Intersection of Subspaces

In the following, we determine a basis of the intersection $U_1 \cap U_2$ of two subspaces $U_1, U_2 \subset V$. This means, we are interested in finding all $\mathbf{x} \in V$, such that $\mathbf{x} \in U_1$ and $\mathbf{x} \in U_2$.

Consider $U_1 = [\mathbf{b}_1, \dots, \mathbf{b}_k] \subset V$ and $U_2 = [\mathbf{c}_1, \dots, \mathbf{c}_l] \subset V$. We know that and $\mathbf{x} \in U_1$ can be represented as a linear combination $\sum_{i=1}^k \lambda_i \mathbf{b}_i$ of the basis vectors (or spanning vectors) $\mathbf{b}_1, \dots, \mathbf{b}_k$. Equivalently $\mathbf{x} = \sum_{j=1}^l \psi_j \mathbf{c}_j$. Therefore, the approach is to find $\lambda_1, \dots, \lambda_k$ and/or ψ_1, \dots, ψ_l , such that

$$\sum_{i=1}^k \lambda_i \mathbf{b}_i = \mathbf{x} = \sum_{j=1}^l \psi_j \mathbf{c}_j \quad (5.70)$$

$$\Leftrightarrow \sum_{i=1}^k \lambda_i \mathbf{b}_i - \sum_{j=1}^l \psi_j \mathbf{c}_j = \mathbf{0}. \quad (5.71)$$

This condition has the form of a homogeneous linear equation system, which we know how to solve: We write the basis vectors into a matrix

$$A = [\mathbf{b}_1 \quad \dots \quad \mathbf{b}_k \quad -\mathbf{c}_1 \quad \dots \quad -\mathbf{c}_l] \quad (5.72)$$

and solve the linear equation system

$$A \begin{bmatrix} \lambda_1 \\ \vdots \\ \lambda_k \\ \psi_1 \\ \vdots \\ \psi_l \end{bmatrix} = \mathbf{0} \quad (5.73)$$

to find either $\lambda_1, \dots, \lambda_k$ or ψ_1, \dots, ψ_l , which we can then use to determine $U_1 \cap U_2$.

Example

We consider

$$U_1 = \left[\begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \end{bmatrix} \right] \subset \mathbb{R}^4, \quad U_2 = \left[\begin{bmatrix} -1 \\ 1 \\ 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} \right] \subset \mathbb{R}^4. \quad (5.74)$$

To find a basis of $U_1 \cap U_2$, we need to find all $\mathbf{x} \in V$ that can be represented as linear combinations of the basis vectors of U_1 and U_2 , i.e.,

$$\sum_{i=1}^3 \lambda_i \mathbf{b}_i = \mathbf{x} = \sum_{j=1}^2 \psi_j \mathbf{c}_j, \quad (5.75)$$

where \mathbf{b}_i and \mathbf{c}_j are the basis vectors of U_1 and U_2 , respectively. The matrix $A = [\mathbf{b}_1 | \mathbf{b}_2 | \mathbf{b}_3 | -\mathbf{c}_1 | -\mathbf{c}_2]$ from (5.72) is given as

$$A = \begin{bmatrix} 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & -1 & -1 \\ 0 & 1 & 1 & -2 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}. \quad (5.76)$$

By using Gaussian elimination, we determine the corresponding reduced row echelon form

$$\begin{bmatrix} 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & -2 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \quad (5.77)$$

We keep in mind that we are interested in finding $\lambda_1, \lambda_2, \lambda_3 \in \mathbb{R}$ and/or $\psi_1, \psi_2 \in \mathbb{R}$ with

$$\begin{bmatrix} 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & -2 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \\ \psi_1 \\ \psi_2 \end{bmatrix} = \mathbf{0}. \quad (5.78)$$

From here, we can immediately see that $\psi_2 = 0$ and $\psi_1 \in \mathbb{R}$ is a free variable since it corresponds to a non-pivot column, and our solution is

$$U_1 \cap U_2 = \psi_1 \mathbf{c}_1 = \begin{bmatrix} -1 \\ 1 \\ 2 \\ 0 \end{bmatrix}, \quad \psi_1 \in \mathbb{R}. \quad (5.79)$$

Remark 21

Alternatively, we could have used $\lambda_1 = -\psi_1, \lambda_2 = 2\psi_1, \lambda_3 = 0$ and determined the (same) solution via the basis vectors of U_1 as

$$\psi_1 \left(- \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} + 2 \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix} \right) = \psi_1 \begin{bmatrix} -1 \\ 1 \\ 2 \\ 0 \end{bmatrix} = \begin{bmatrix} -1 \\ 1 \\ 2 \\ 0 \end{bmatrix}, \quad \psi_1 \in \mathbb{R}. \quad (5.80)$$

5.9 Linear Mappings

In the following, we will study mappings on vector spaces that preserve their structure. Earlier on in the course, we said that vectors are objects that can be added together and multiplied by a scalar, and the resulting object is still a vector. This property we wish to preserve: Consider two real vector spaces V, W . A mapping $\Phi : V \rightarrow W$ preserves the structure of the vector space if

$$\Phi(\mathbf{x} + \mathbf{y}) = \Phi(\mathbf{x}) + \Phi(\mathbf{y}) \quad (5.81)$$

$$\Phi(\lambda \mathbf{x}) = \lambda \Phi(\mathbf{x}) \quad (5.82)$$

for all $\mathbf{x}, \mathbf{y} \in V$ and $\lambda \in \mathbb{R}$. We can summarize this in the following definition:

Definition 12 (Linear Mapping)

For real vector spaces V, W , a mapping $\Phi : V \rightarrow W$ is called **linear** (or **vector space homomorphism**) if

$$\forall \mathbf{x}, \mathbf{y} \in V \forall \lambda, \psi \in \mathbb{R} : \Phi(\lambda \mathbf{x} + \psi \mathbf{y}) = \lambda \Phi(\mathbf{x}) + \psi \Phi(\mathbf{y}). \quad (5.83)$$

Important special cases:

- **Isomorphism:** $\Phi : V \rightarrow W$ linear and bijective
- **Endomorphism:** $\Phi : V \rightarrow V$ linear
- **Automorphism:** $\Phi : V \rightarrow V$ linear and bijective
- We define $\text{id}_V : V \rightarrow V, \mathbf{x} \mapsto \mathbf{x}$ as the **identity mapping** in V .

Example: Homomorphism

The mapping $\Phi : \mathbb{R}^2 \rightarrow \mathbb{C}, \Phi(\mathbf{x}) = x_1 + ix_2$, is a homomorphism:

$$\begin{aligned} \Phi \left(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} \right) &= (x_1 + y_1) + i(x_2 + y_2) = x_1 + ix_2 + y_1 + iy_2 = \Phi \left(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \right) + \Phi \left(\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} \right) \\ \Phi \left(\lambda \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \right) &= \lambda x_1 + \lambda i x_2 = \lambda(x_1 + ix_2) = \lambda \Phi \left(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \right) \end{aligned} \quad (5.84)$$

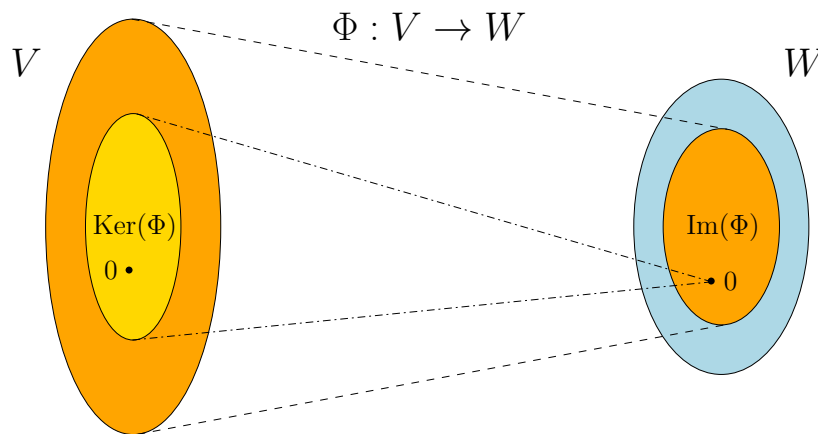


Figure 5.3: Kernel and Image of a linear mapping $\Phi : V \rightarrow W$.

This justifies why complex numbers can be represented as tuples in \mathbb{R}^2 : There is a bijective linear mapping (we only showed linearity, but not the bijection) that converts the elementwise addition of tuples in \mathbb{R}^2 into the set of complex numbers with the corresponding addition.

Theorem 1

Finite-dimensional \mathbb{R} -vector spaces V and W are isomorph if and only if $\dim(V) = \dim(W)$.

5.9.1 Image and Kernel (Null Space)

Definition 13 (Image and Kernel)

For $\Phi : V \rightarrow W$, we define the **kernel/null space**

$$\ker(\Phi) := \Phi^{-1}(\{\mathbf{0}_W\}) = \{v \in V : \Phi(v) = \mathbf{0}_W\} \quad (5.85)$$

and the **image/range**

$$\text{Im}(\Phi) := \Phi(V) = \{w \in W | \exists v \in V : \Phi(v) = w\}. \quad (5.86)$$

More informally, the kernel is the set of vectors in $v \in V$ that Φ maps onto the neutral element $\mathbf{0}_W$ in W . The image is the set of vectors $w \in W$ that can be “reached” by Φ from any vector in V . An illustration is given in Figure 5.3.

Remark 22

Consider a linear mapping $\Phi : V \rightarrow W$, where V, W are vector spaces.

- It always holds that $\Phi(\{\mathbf{0}_V\}) = \mathbf{0}_W$ and, therefore, $\mathbf{0}_V \in \ker(\Phi)$. In particular, the null space is never empty.
- $\text{Im}(\Phi) \subset W$ is a subspace of W , and $\ker(\Phi) \subset V$ is a subspace of V .
- Φ is injective (one-to-one) if and only if $\ker(\Phi) = \{\mathbf{0}\}$

Remark 23 (Null Space and Column Space)

- For $A \in \mathbb{R}^{m \times n}$ the mapping $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^m$, $\mathbf{x} \mapsto A\mathbf{x}$ is linear. For $A = (\mathbf{a}_1 | \dots | \mathbf{a}_n)$ we obtain

$$\text{Im}(\Phi) = \{A\mathbf{x} : \mathbf{x} \in \mathbb{R}^n\} = \{x_1\mathbf{a}_1 + \dots + x_n\mathbf{a}_n : x_1, \dots, x_n \in \mathbb{R}\} = [\mathbf{a}_1, \dots, \mathbf{a}_n] \subset \mathbb{R}^m, \quad (5.87)$$

i.e., the image is the span of the columns of A , also called the **column space**. Therefore, the column space (image) is a subspace of \mathbb{R}^m , where m is the “height” of the matrix.

- The kernel/null space $\ker(\Phi)$ is the general solution to the linear homogeneous equation system $A\mathbf{x} = \mathbf{0}$ and captures all possible linear combinations of the elements in \mathbb{R}^n that produce $\mathbf{0} \in \mathbb{R}^m$.
- The kernel (null space) is a subspace of \mathbb{R}^n , where n is the “width” of the matrix.
- The null space focuses on the relationship among the columns, and we can use it to determine whether/how we can express a column as a linear combination of other columns.
- The purpose of the null space is to determine whether a solution of the linear equation system is unique and, if not, to capture all possible solutions.

Example (Image and Kernel of a Linear Mapping)

The mapping

$$\Phi : \mathbb{R}^4 \rightarrow \mathbb{R}^2, \quad \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} \mapsto \begin{bmatrix} 1 & 2 & -1 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} x_1 + 2x_2 - x_3 \\ x_1 + x_4 \end{bmatrix} \quad (5.88)$$

$$= x_1 \begin{bmatrix} 1 \\ 1 \end{bmatrix} + x_2 \begin{bmatrix} 2 \\ 0 \end{bmatrix} + x_3 \begin{bmatrix} -1 \\ 0 \end{bmatrix} + x_4 \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad (5.89)$$

is linear. To determine $\text{Im}(\Phi)$ we can simply take the span of the columns of the transformation matrix and obtain

$$\text{Im}(\Phi) = \left[\begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 2 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right]. \quad (5.90)$$

To compute the kernel (null space) of Φ , we need to solve $A\mathbf{x} = \mathbf{0}$, i.e., we need to solve a homogeneous equation system. To do this, we use Gaussian elimination to transform A into reduced row echelon form:

$$\begin{bmatrix} 1 & 2 & -1 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix} \rightsquigarrow \begin{bmatrix} 1 & 0 & 0 & 1 \\ 1 & 2 & -1 & 0 \end{bmatrix} -R_1 \cdot \left(\frac{1}{2}\right) \rightsquigarrow \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & -\frac{1}{2} & -\frac{1}{2} \end{bmatrix}$$

This matrix is now in reduced row echelon form, and we can now use the Minus-1 Trick to compute a basis of the kernel (see Section 5.4.3).²² This gives us now the kernel (null space) as

$$\ker(\Phi) = \left[\begin{array}{c} 0 \\ \frac{1}{2} \\ 1 \\ 0 \end{array}, \begin{array}{c} -1 \\ \frac{1}{2} \\ 0 \\ 1 \end{array} \right]. \quad (5.91)$$

Theorem 2 (Rank-Nullity Theorem)

For vector spaces V, W and a linear mapping $\Phi : V \rightarrow W$ it holds that

$$\dim(\ker(\Phi)) + \dim(\text{Im}(\Phi)) = \dim(V) \quad (5.92)$$

Remark 24

Consider \mathbb{R} -vector spaces V, W, X . Then:

- For linear mappings $\Phi : V \rightarrow W$ and $\Psi : W \rightarrow X$ the mapping $\Psi \circ \Phi : V \rightarrow X$ is also linear.
- If $\Phi : V \rightarrow W$ is an isomorphism then $\Phi^{-1} : W \rightarrow V$ is an isomorphism as well.
- If $\Phi : V \rightarrow W, \Psi : V \rightarrow W$ are linear then $\Phi + \Psi$ and $\lambda\Phi, \lambda \in \mathbb{R}$ are linear, too.

5.9.2 Matrix Representation of Linear Mappings

Any n -dimensional \mathbb{R} -vector space is isomorphic to \mathbb{R}^n (Theorem 1). If we define a basis $\{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ of V we can construct an isomorphism concretely. In the following, the order of the basis vectors will be important. Therefore, we write

$$B = (\mathbf{b}_1, \dots, \mathbf{b}_n) \quad (5.93)$$

and call this n -tuple an **ordered basis** of V .

Definition 14 (Coordinates)

Consider an \mathbb{R} -vector space V and an ordered basis $B = (\mathbf{b}_1, \dots, \mathbf{b}_n)$. For $\mathbf{x} \in V$ we obtain a unique representation (linear combination)

$$\mathbf{x} = \alpha_1 \mathbf{b}_1 + \dots + \alpha_n \mathbf{b}_n \quad (5.94)$$

²²Alternatively, we can express the non-pivot columns (columns 3 and 4) as linear combinations of the pivot-columns (columns 1 and 2). The third column \mathbf{a}_3 is equivalent to $-\frac{1}{2}$ times the second column \mathbf{a}_2 . Therefore, $\mathbf{0} = \mathbf{a}_3 + \frac{1}{2}\mathbf{a}_2$. In the same way, we see that $\mathbf{a}_4 = \mathbf{a}_1 - \frac{1}{2}\mathbf{a}_2$ and, therefore, $\mathbf{0} = \mathbf{a}_1 - \frac{1}{2}\mathbf{a}_2 - \mathbf{a}_4$.

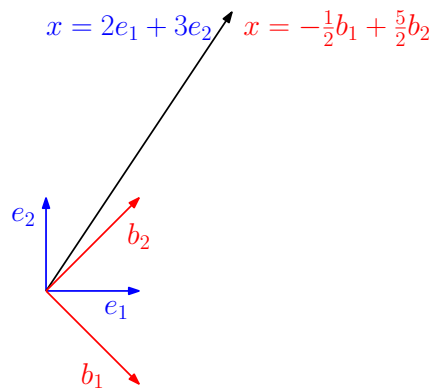


Figure 5.4: The coordinates of the vector x are the coefficients of the linear combination of the basis vectors to represent x .

of x with respect to B . Then $\alpha_1, \dots, \alpha_n$ are the **coordinates** of x with respect to B and the vector

$$\begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix} \in \mathbb{R}^n \quad (5.95)$$

is the **coordinate vector/coordinate representation** of x with respect to B .

Remark 25

Intuitively, the basis vectors can be understood as units (including rather odd units, such as “apples”, “bananas”, “kilograms” or “seconds”). However, let us have a look at a geometric vector $x \in \mathbb{R}^2$ with the coordinates

$$\begin{bmatrix} 2 \\ 3 \end{bmatrix} \quad (5.96)$$

with respect to the standard basis e_1, e_2 in \mathbb{R}^2 . This means, we can write $x = 2e_1 + 3e_2$. However, we do not have to choose the standard basis to represent this vector. For the basis $b_1 = [1, -1]^\top, b_2 = [1, 1]^\top$, we would obtain the coordinates $\frac{1}{2}[-1, 5]^\top$ to represent the same vector (see Fig. 5.4).

In the following, we will be looking at mappings that transform coordinate vectors with respect to one basis into coordinate vectors with respect to a different basis.

Remark 26

For an n -dimensional \mathbb{R} -vector space V and a basis B of V , the mapping $\Phi : \mathbb{R}^n \rightarrow V$, $\Phi(e_i) = b_i$, $i = 1, \dots, n$, is linear (and because of Theorem 1 an isomorphism), where (e_1, \dots, e_n) is the standard basis of \mathbb{R}^n .

Now we are ready to make an explicit connection between matrices and linear mappings between finite-dimensional vector spaces.

Definition 15 (Transformation matrix)

Consider \mathbb{R} -vector spaces V, W with corresponding (ordered) bases $= (\mathbf{b}_1, \dots, \mathbf{b}_n)$ and $C = (\mathbf{c}_1, \dots, \mathbf{c}_m)$. Moreover, we consider a linear mapping $\Phi : V \rightarrow W$. For $j \in \{1, \dots, n\}$

$$\Phi(\mathbf{b}_j) = a_{1j}\mathbf{c}_1 + \dots + a_{mj}\mathbf{c}_m \quad (5.97)$$

is the unique representation of $\Phi(\mathbf{b}_j)$ with respect to C . Then, we call the $m \times n$ -matrix

$$\mathbf{A}_\Phi := ((a_{ij})) \quad (5.98)$$

the **transformation matrix** of Φ (with respect to the ordered bases B of V and C of W).

Remark 27

- The coordinates of $\Phi(\mathbf{b}_j)$ are the j -th column of \mathbf{A}_Φ .
- $\text{rk}(\mathbf{A}_\Phi) = \dim(\text{Im}(\Phi))$
- Consider (finite-dimensional) \mathbb{R} -vector spaces V, W with ordered bases B, C , $\Phi : V \rightarrow W$ linear with transformation matrix \mathbf{A}_Φ . If \hat{x} is the coordinate vector of $x \in V$ and \hat{y} the coordinate vector of $y = \Phi(x) \in W$, then

$$\hat{y} = \mathbf{A}_\Phi \hat{x}. \quad (5.99)$$

This means that the transformation matrix can be used to map coordinates with respect to an ordered basis in V to coordinates with respect to an ordered basis in W .

Example (Transformation Matrix)

Consider a homomorphism $\Phi : V \rightarrow W$ and ordered bases $B = (\mathbf{b}_1, \dots, \mathbf{b}_3)$ of V and $C = (\mathbf{c}_1, \dots, \mathbf{c}_4)$ of W . With

$$\begin{aligned} \Phi(\mathbf{b}_1) &= \mathbf{c}_1 - \mathbf{c}_2 + 3\mathbf{c}_3 - \mathbf{c}_4 \\ \Phi(\mathbf{b}_2) &= 2\mathbf{c}_1 + \mathbf{c}_2 + 7\mathbf{c}_3 + 2\mathbf{c}_4 \\ \Phi(\mathbf{b}_3) &= 3\mathbf{c}_2 + \mathbf{c}_3 + 4\mathbf{c}_4 \end{aligned} \quad (5.100)$$

the transformation matrix \mathbf{A}_Φ with respect to B and C satisfies $\Phi(\mathbf{b}_k) = \sum_{i=1}^4 \alpha_{ik}\mathbf{c}_i$ for $k = 1, \dots, 3$ and is given as

$$\mathbf{A}_\Phi = (\alpha_1 | \alpha_2 | \alpha_3) = \begin{bmatrix} 1 & 2 & 0 \\ -1 & 1 & 3 \\ 3 & 7 & 1 \\ -1 & 2 & 4 \end{bmatrix}, \quad (5.101)$$

where the α_j , $j = 1, 2, 3$, are the coordinate vectors of $\Phi(\mathbf{b}_j)$ with respect to C .

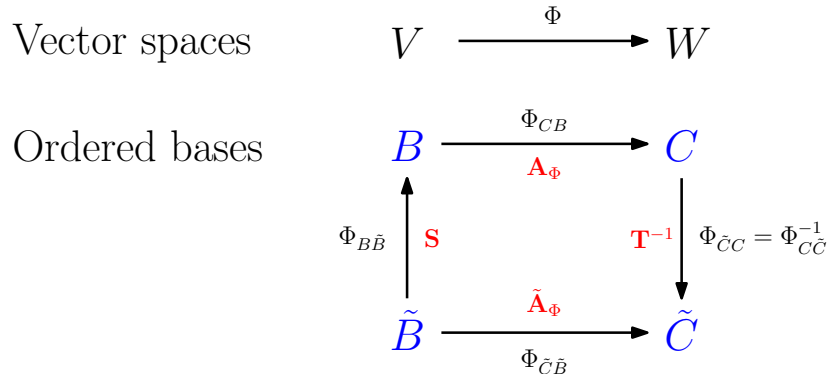


Figure 5.5: For a homomorphism $\Phi : V \rightarrow W$ and ordered bases B, \tilde{B} of V and C, \tilde{C} of W (marked in blue), we can express the mapping $\Phi_{\tilde{B}\tilde{C}}$ with respect to the bases \tilde{B}, \tilde{C} equivalently as a composition of the homomorphisms $\Phi_{\tilde{C}\tilde{B}} = \Phi_{\tilde{C}C} \circ \Phi_{CB} \circ \Phi_{B\tilde{B}}$ with respect to the bases in the subscripts. The corresponding transformation matrices are in red.

5.9.3 Basis Change

In the following, we will have a closer look at how transformation matrices of a linear mapping $\Phi : V \rightarrow W$ change if we change the bases in V and W . Consider

$$B = (\mathbf{b}_1, \dots, \mathbf{b}_n), \quad \tilde{B} = (\tilde{\mathbf{b}}_1, \dots, \tilde{\mathbf{b}}_n) \tag{5.102}$$

ordered bases of V and

$$C = (\mathbf{c}_1, \dots, \mathbf{c}_m), \quad \tilde{C} = (\tilde{\mathbf{c}}_1, \dots, \tilde{\mathbf{c}}_n) \tag{5.103}$$

ordered bases of W . Moreover, \mathbf{A}_Φ is the transformation matrix of the linear mapping $\Phi : V \rightarrow W$ with respect to the bases B and C , and $\tilde{\mathbf{A}}_\Phi$ is the corresponding transformation mapping with respect to \tilde{B} and \tilde{C} . We will now investigate how \mathbf{A} and $\tilde{\mathbf{A}}$ are related, i.e., how/whether we can transform \mathbf{A}_Φ into $\tilde{\mathbf{A}}_\Phi$ if we choose to perform a basis change from B, C to \tilde{B}, \tilde{C} .

We can write the vectors of the new basis \tilde{B} of V as a linear combination of the basis vectors of B , such that

$$\tilde{\mathbf{b}}_j = s_{1j}\mathbf{b}_1 + \dots + s_{nj}\mathbf{b}_n, \quad j = 1, \dots, n. \tag{5.104}$$

Similarly, we write the new basis vectors \tilde{C} of W as a linear combination of the basis vectors of C , which yields

$$\tilde{\mathbf{c}}_k = t_{1k}\mathbf{c}_1 + \dots + t_{mk}\mathbf{c}_m. \tag{5.105}$$

We define $\mathbf{S} = ((s_{ij})) \in \mathbb{R}^{n \times n}$ as the transformation matrix that maps coordinates with respect to \tilde{B} onto coordinates with respect to B and $\mathbf{T} = ((t_{lk})) \in \mathbb{R}^{m \times m}$ as the transformation matrix that maps coordinates with respect to \tilde{C} onto coordinates with respect to C . In particular, the j th column of \mathbf{S} are the coordinate representations of $\tilde{\mathbf{b}}_j$ with respect to B and the j th columns of \mathbf{T} is the coordinate representation of $\tilde{\mathbf{c}}_j$ with respect to C . Note that both \mathbf{S} and \mathbf{T} are regular.

For all $j = 1, \dots, n$, we get

$$\Phi(\tilde{\mathbf{b}}_j) = \sum_{k=1}^m \underbrace{\tilde{a}_{kj} \tilde{\mathbf{c}}_k}_{\in W} = \sum_{k=1}^m \tilde{a}_{kj} \sum_{i=1}^m t_{ik} \mathbf{c}_i = \sum_{i=1}^m \left(\sum_{k=1}^m t_{ik} \tilde{a}_{kj} \right) \mathbf{c}_i \quad (5.106)$$

where we expressed the new basis vectors $\tilde{\mathbf{c}}_k \in W$ as linear combinations of the basis vectors $\mathbf{c}_i \in W$. When we express the $\tilde{\mathbf{b}}_k \in V$ as linear combinations of $\mathbf{b}_i \in V$, we arrive at

$$\Phi(\tilde{\mathbf{b}}_j) = \Phi\left(\sum_{k=1}^n s_{kj} \mathbf{b}_k\right) = \sum_{k=1}^n s_{kj} \Phi(\mathbf{b}_k) = \sum_{k=1}^n s_{kj} \sum_{i=1}^m a_{ik} \mathbf{c}_i = \sum_{i=1}^m \left(\sum_{k=1}^n a_{ik} s_{kj} \right) \mathbf{c}_i \quad (5.107)$$

Comparing (5.106) and (5.107), it follows for all $j = 1, \dots, n$ and $i = 1, \dots, m$ that

$$\sum_{k=1}^m t_{ik} \tilde{a}_{kj} = \sum_{k=1}^n a_{ik} s_{kj} \quad (5.108)$$

and, therefore,

$$T \tilde{\mathbf{A}}_\Phi = \mathbf{A} \mathbf{S}, \quad (5.109)$$

such that

$$\tilde{\mathbf{A}}_\Phi = \mathbf{T}^{-1} \mathbf{A} \mathbf{S}. \quad (5.110)$$

Hence, with a basis change in V (B is replaced with \tilde{B}) and W (C is replaced with \tilde{C}) the transformation matrix \mathbf{A}_Φ of a linear mapping $\Phi : V \rightarrow W$ is replaced by an equivalent matrix $\tilde{\mathbf{A}}_\Phi$ with

$$\tilde{\mathbf{A}}_\Phi = \mathbf{T}^{-1} \mathbf{A}_\Phi \mathbf{S}. \quad (5.111)$$

Figure 5.5 illustrates this relation: Consider a homomorphism $\Phi : V \rightarrow W$ and ordered bases B, \tilde{B} of V and C, \tilde{C} of W . The mapping Φ_{CB} is an instantiation of Φ and maps basis vectors of B onto linear combinations of basis vectors of C . Assuming, we know the transformation matrix \mathbf{A}_Φ of Φ_{CB} with respect to the ordered bases B, C . When we perform a basis change from B to \tilde{B} in V and from C to \tilde{C} in W , we can determine the corresponding transformation matrix $\tilde{\mathbf{A}}_\Phi$ as follows: First, we find a linear mapping $\Phi_{B\tilde{B}} : V \rightarrow V$ that maps coordinates with respect to the new basis \tilde{B} onto the (unique) coordinates with respect to the “old” basis B (in V). Then, we use the transformation matrix \mathbf{A}_Φ of $\Phi_{CB} : V \rightarrow W$ to map these coordinates onto the coordinates with respect to C in W . Finally, we use a linear mapping $\Phi_{\tilde{C}C} : W \rightarrow W$ to map the coordinates with respect to C onto coordinates with respect to \tilde{C} . Therefore, we can express the linear mapping $\Phi_{\tilde{C}\tilde{B}}$ as a composition of linear mappings that involve the “old” basis.²³

$$\Phi_{\tilde{C}\tilde{B}} = \Phi_{\tilde{C}C} \circ \Phi_{CB} \circ \Phi_{B\tilde{B}} = \Phi_{\tilde{C}C}^{-1} \circ \Phi_{CB} \circ \Phi_{B\tilde{B}}. \quad (5.112)$$

²³Why does $\Phi_{\tilde{C}C}^{-1}$ exist?

Definition 16 (Equivalence)

Two matrices $A, \tilde{A} \in \mathbb{R}^{m \times n}$ are **equivalent** if there exist regular matrices $S \in \mathbb{R}^{n \times n}$ and $T \in \mathbb{R}^{m \times m}$, such that $\tilde{A} = T^{-1}AS$.

Definition 17 (Similarity)

Two matrices $A, \tilde{A} \in \mathbb{R}^{n \times n}$ are **similar** if there exists a regular matrix $S \in \mathbb{R}^{n \times n}$ with $\tilde{A} = S^{-1}AS$

Remark 28

Similar matrices are always equivalent. However, equivalent matrices are not necessarily similar.

Remark 29

Consider \mathbb{R} -vector spaces V, W, X . From Remark 24 we already know that for linear mappings $\Phi : V \rightarrow W$ and $\Psi : W \rightarrow X$ the mapping $\Psi \circ \Phi : V \rightarrow X$ is also linear. With transformation matrices A_Φ and A_Ψ of the corresponding mappings, the overall transformation matrix $A_{\Psi \circ \Phi}$ is given by $A_{\Psi \circ \Phi} = A_\Psi A_\Phi$.

In light of this remark, we can look at basis changes from the perspective of composing linear mappings:

- A_Φ is the transformation matrix of a linear mapping $\Phi_{CB} : V \rightarrow W$ with respect to the bases B, C .
- \tilde{A}_Φ is the transformation matrix of a linear mapping $\Phi_{\tilde{C}\tilde{B}} : V \rightarrow W$ with respect to the bases \tilde{B}, \tilde{C} .
- S is the transformation matrix of a linear mapping $\Phi_{\tilde{B}B} : V \rightarrow V$ (automorphism) that represents \tilde{B} in terms of B .
- T is the transformation matrix of a linear mapping $\Phi_{\tilde{C}C} : W \rightarrow W$ (automorphism) that represents \tilde{C} in terms of C .

If we (informally) write down the transformations just in terms of bases then

- $A_\Phi : B \rightarrow C$
- $\tilde{A}_\Phi : \tilde{B} \rightarrow \tilde{C}$
- $S : \tilde{B} \rightarrow B$
- $T : \tilde{C} \rightarrow C$ and $T^{-1} : C \rightarrow \tilde{C}$

and

$$\tilde{B} \rightarrow \tilde{C} = \tilde{B} \rightarrow B \rightarrow C \rightarrow \tilde{C} \quad (5.113)$$

$$\tilde{A}_\Phi = T^{-1}A_\Phi S. \quad (5.114)$$

Note that the execution order in (5.114) is from right to left because vectors are multiplied at the right-hand side.

Example

Consider a linear mapping $\Phi : \mathbb{R}^3 \rightarrow \mathbb{R}^4$ whose transformation matrix is

$$\mathbf{A}_\Phi = \begin{bmatrix} 1 & 2 & 0 \\ -1 & 1 & 3 \\ 3 & 7 & 1 \\ -1 & 2 & 4 \end{bmatrix} \quad (5.115)$$

with respect to the standard bases

$$\mathbf{B} = \left(\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right), \quad \mathbf{C} = \left(\begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right). \quad (5.116)$$

We seek the transformation matrix $\tilde{\mathbf{A}}_\Phi$ of Φ with respect to the new bases

$$\tilde{\mathbf{B}} = \left(\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \right) \in \mathbb{R}^3, \quad \tilde{\mathbf{C}} = \left(\begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right). \quad (5.117)$$

Then,

$$\mathbf{S} = \begin{bmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix}, \quad \mathbf{T} = \begin{bmatrix} 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.118)$$

where the i th column of \mathbf{S} is the coordinate representation of $\tilde{\mathbf{b}}_i$ in terms of the basis vectors of B .²⁴ Similarly, the j th column of \mathbf{T} is the coordinate representation of $\tilde{\mathbf{c}}_j$ in terms of the basis vectors of C .

Therefore, we obtain

$$\tilde{\mathbf{A}}_\Phi = \mathbf{T}^{-1} \mathbf{A}_\Phi \mathbf{S} = \begin{bmatrix} 1/2 & 1/2 & -1/2 & -1/2 \\ 1/2 & -1/2 & 1/2 & -1/2 \\ -1/2 & 1/2 & 1/2 & 1/2 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 3 & 2 & 1 \\ 0 & 4 & 2 \\ 10 & 8 & 4 \\ 1 & 6 & 3 \end{bmatrix} = \begin{bmatrix} -4 & -4 & -2 \\ 6 & 0 & 0 \\ 4 & 8 & 4 \\ 1 & 6 & 3 \end{bmatrix}. \quad (5.119)$$

Soon, we will be able to exploit the concept of a basis change to find a basis with respect to which the transformation matrix of an endomorphism has a particularly simple (diagonal) form.

²⁴Since B is the standard basis, this representation is straightforward to find. For a general basis B we would need to solve a linear equation system to find the λ_i such that $\sum_{i=1}^3 \lambda_i \mathbf{b}_i = \tilde{\mathbf{b}}_j$, $j = 1, \dots, 3$.

5.10 Determinants

Determinants are important concepts in linear algebra. For instance, they indicate whether a matrix can be inverted or we can use them to check for linear independence. A geometric intuition is that the absolute value of the determinant of real vectors is equal to the volume of the parallelepiped spanned by those vectors. Determinants will play a very important role for determining eigenvalues and eigenvectors (Section 5.11).

Determinants are only defined for square matrices $A \in \mathbb{R}^{n \times n}$, and we write $\det(A)$ or $|A|$.

Remark 30

- For $n = 1$, $\det(A) = \det(a_{11}) = a_{11}$
- For $n = 2$,

$$\det(A) = \begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix} = a_{11}a_{22} - a_{12}a_{21} \quad (5.120)$$

- For $n = 3$ (Sarrus rule):

$$\begin{vmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{vmatrix} = a_{11}a_{22}a_{33} + a_{21}a_{32}a_{13} + a_{31}a_{12}a_{23} - a_{31}a_{22}a_{13} - a_{11}a_{32}a_{23} - a_{21}a_{12}a_{33} \quad (5.121)$$

- For an upper/lower triangular matrix A , the determinant is the product of the diagonal elements: $\det(A) = \prod_{i=1}^n a_{ii}$

Remark 31 (Properties of Determinants)

- $\det(AB) = \det(A)\det(B)$
- $\det(A) = 0 \Leftrightarrow A$ is singular (not invertible)
- Alternatively: A is regular $\Leftrightarrow \det(A) \neq 0$.
- $\det(A) = \det(A^\top)$
- If A is regular then $\det(A^{-1}) = 1/\det(A)$
- Similar matrices possess the same determinant. Therefore, for a linear mapping $\Phi : V \rightarrow V$ all transformation matrices A_Φ of Φ have the same determinant.

Theorem 3

For $A \in \mathbb{R}^{n \times n}$:

1. Adding a multiple of a column/row to another one does not change $\det(A)$.
2. Multiplication of a column/row with $\lambda \in \mathbb{R}$ scales $\det(A)$ by λ . In particular, $\det(\lambda A) = \lambda^n \det(A)$.

3. Swapping two rows/columns changes the sign of $\det(A)$.

Because of this theorem, we can use Gaussian elimination to compute $\det(A)$. However, we need to pay attention to swapping the sign when swapping rows.

Example

$$\begin{vmatrix} 2 & 0 & 1 & 2 & 0 \\ 2 & -1 & 0 & 1 & 1 \\ 0 & 1 & 2 & 1 & 2 \\ -2 & 0 & 2 & -1 & 2 \\ 2 & 0 & 0 & 1 & 1 \end{vmatrix} = \begin{vmatrix} 2 & 0 & 1 & 2 & 0 \\ 0 & -1 & -1 & -1 & 1 \\ 0 & 1 & 2 & 1 & 2 \\ 0 & 0 & 3 & 1 & 2 \\ 0 & 0 & -1 & -1 & 1 \end{vmatrix} = \begin{vmatrix} 2 & 0 & 1 & 2 & 0 \\ 0 & -1 & -1 & -1 & 1 \\ 0 & 0 & 1 & 0 & 3 \\ 0 & 0 & 3 & 1 & 2 \\ 0 & 0 & -1 & -1 & 1 \end{vmatrix} \quad (5.122)$$

$$= \begin{vmatrix} 2 & 0 & 1 & 2 & 0 \\ 0 & -1 & -1 & -1 & 1 \\ 0 & 0 & 1 & 0 & 3 \\ 0 & 0 & 0 & 1 & -7 \\ 0 & 0 & 0 & -1 & 4 \end{vmatrix} = \begin{vmatrix} 2 & 0 & 1 & 2 & 0 \\ 0 & -1 & -1 & -1 & 1 \\ 0 & 0 & 1 & 0 & 3 \\ 0 & 0 & 0 & 1 & -7 \\ 0 & 0 & 0 & 0 & -3 \end{vmatrix} = 6 \quad (5.123)$$

We first used Gaussian elimination to bring A into triangular form, and then exploited the fact that the determinant of a triangular matrix is the product of its diagonal elements.

Theorem 4 (Laplace Expansion)

Consider a matrix $A = ((a_{ij})) \in \mathbb{R}^{n \times n}$. We define $A_{i,j}$ to be the matrix that remains if we delete the i th row and the j th column from A . Then, for all $j = 1, \dots, n$:

1. $\det(A) = \sum_{k=1}^n (-1)^{k+j} a_{kj} \det(A_{k,j})$ “Expansion along column j ”
2. $\det(A) = \sum_{k=1}^n (-1)^{k+j} a_{jk} \det(A_{j,k})$ “Expansion along row j ”

The scalar $(-1)^{k+1} \det(A_{k,j})$ is also called a **cofactor**.

Remark 32 (Checkerboard Pattern)

The signs $+/-$ resulting from $(-1)^{k+j}$ follow a checkerboard pattern, e.g., for a 3×3 -matrix it looks as follows:

$$\begin{bmatrix} + & - & + \\ - & + & - \\ + & - & + \end{bmatrix}. \quad (5.124)$$

Example (Expansion along Rows)

In the following, we will go through an example of how we can simplify a determinant by expanding along a row. In particular, we expand along the first row in the following example (the Sarrus rule can be recovered from this):

$$\begin{vmatrix} 1 & -2 & -3 \\ 1 & 0 & 1 \\ 2 & 4 & 4 \end{vmatrix} = 1(-1)^{1+1} \begin{vmatrix} 0 & 1 \\ 4 & 4 \end{vmatrix} + (-2)(-1)^{1+2} \begin{vmatrix} 1 & 1 \\ 2 & 4 \end{vmatrix} + (-3)(-1)^{1+3} \begin{vmatrix} 1 & 0 \\ 2 & 4 \end{vmatrix} \quad (5.125)$$

$$= -4 + 2(4 - 2) - 3 \cdot 4 = -4 + 4 - 12 = -12 \quad (5.126)$$

It is usually advisable to expand along rows/columns with many 0 entries.

Example

Let us re-compute the example in (5.122)

$$\begin{vmatrix} 2 & 0 & 1 & 2 & 0 \\ 2 & -1 & 0 & 1 & 1 \\ 0 & 1 & 2 & 1 & 2 \\ -2 & 0 & 2 & -1 & 2 \\ 2 & 0 & 0 & 1 & 1 \end{vmatrix} = \begin{vmatrix} 2 & 0 & 1 & 2 & 0 \\ 0 & -1 & -1 & -1 & 1 \\ 0 & 1 & 2 & 1 & 2 \\ 0 & 0 & 3 & 1 & 2 \\ 0 & 0 & -1 & -1 & 1 \end{vmatrix} \stackrel{\text{1st col.}}{=} (-1)^{1+1} 2 \cdot \begin{vmatrix} -1 & -1 & -1 & 1 \\ 1 & 2 & 1 & 2 \\ 0 & 3 & 1 & 2 \\ 0 & -1 & -1 & 1 \end{vmatrix} \quad (5.127)$$

If we now subtract the fourth row from the first row and multiply (-2) times the third column to the fourth column we obtain

$$2 \begin{vmatrix} -1 & 0 & 0 & 0 \\ 1 & 2 & 1 & 0 \\ 0 & 3 & 1 & 0 \\ 0 & -1 & -1 & 3 \end{vmatrix} \stackrel{\text{1st row}}{=} -2 \begin{vmatrix} 2 & 1 & 0 \\ 3 & 1 & 0 \\ -1 & -1 & 3 \end{vmatrix} \stackrel{\text{3rd col.}}{=} (-2) \cdot 3(-1)^{3+3} \cdot \begin{vmatrix} 2 & 1 \\ 3 & 1 \end{vmatrix} = 6 \quad (5.128)$$

5.11 Eigenvalues

Definition 18 (Eigenvalue, Eigenvector)

For an \mathbb{R} -vector space V and an endomorphism $\Phi : V \rightarrow V$, the scalar $\lambda \in \mathbb{R}$ is called **eigenvalue** of Φ if there exists an $x \in V \setminus \{0\}$, with

$$\Phi(x) = \lambda x. \quad (5.129)$$

The corresponding vector x is called **eigenvector** of Φ associated with eigenvalue λ .

Definition 19 (Eigenspace and Spectrum)

- The set of all eigenvectors of Φ associated with an eigenvalue λ spans a subspace of V , which is called **eigenspace** of Φ with respect to λ and denoted by E_λ .
- The set of all eigenvalues of Φ is called **spectrum** of Φ .

Remark 33

- Apparently, $E_\lambda = \ker(\Phi - \lambda \text{id}_V)$ since

$$\Phi(x) = \lambda x \Leftrightarrow \Phi(x) - \lambda x = 0 \Leftrightarrow (\Phi - \lambda \text{id}_V)x = 0 \Leftrightarrow x \in \ker(\Phi - \lambda \text{id}_V). \quad (5.130)$$

- A matrix $A \in \mathbb{R}^{n \times n}$ uniquely determines the linear mapping $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$, $x \mapsto Ax$. Therefore, we can also talk of eigenvalues, eigenvectors and eigenspaces of square matrices.
- Similar matrices possess the same eigenvalues
- Eigenvectors are not unique: If x is an eigenvector of Φ with eigenvalue λ , then αx , $\alpha \in \mathbb{R} \setminus \{0\}$, is an eigenvector with the same eigenvalue. Therefore, there exists an infinite number of eigenvectors for every eigenvalue λ .

Theorem 5

Consider an \mathbb{R} -vector space V and a linear map $\Phi : V \rightarrow V$ with pairwise different eigenvalues $\lambda_1, \dots, \lambda_k$ and corresponding eigenvectors x_1, \dots, x_k . Then the eigenvectors x_1, \dots, x_k are linearly independent.

More informally, Theorem 5 states that eigenvectors belonging to different eigenvalues are linearly independent.²⁵

Remark 34

For an n -dimensional vector space V , an endomorphism $\Phi : V \rightarrow V$ (and equivalently the corresponding transformation matrix $A \in \mathbb{R}^{n \times n}$) possesses at most n different eigenvalues.

The following statements are equivalent:

- λ is eigenvalue of $A \in \mathbb{R}^{n \times n}$
- There exists an $x \in \mathbb{R}^n \setminus \{0\}$ with $Ax = \lambda x$ or, equivalently, $(A - \lambda I_n)x = 0$
- $(A - \lambda I_n)x = 0$ can be solved non-trivially, i.e., $x \neq 0$.
- $\text{rk}(A - \lambda I_n) < n$
- $\det(A - \lambda I_n) = 0$

Note that A and A^\top possess the same eigenvalues, but not the same eigenvectors.

Remark 35 (Geometric Interpretation)

Geometrically, an eigenvector corresponding to a real, nonzero eigenvalue points in a direction that is stretched, and the eigenvalue is the factor by which it is stretched. If the eigenvalue is negative, the direction is reversed. In particular, the eigenvector does not change its direction under Φ .²⁶

²⁵Here is a short video that goes through a proof (by contradiction):

<http://tinyurl.com/hhxks1k>

²⁶A short video on the geometric interpretation of eigenvalues is available at

<http://tinyurl.com/p6nhkvf>.

5.11.1 Characteristic Polynomial

In the following, we will discuss how to determine the eigenspaces of an endomorphism Φ .²⁷ For this, we need to introduce the characteristic polynomial first.

Definition 20 (Characteristic Polynomial)

For $\lambda \in \mathbb{R}$ and an endomorphism Φ on \mathbb{R}^n with transformation matrix A

$$p(\lambda) = \det(A - \lambda I) = a_0 + a_1 \lambda + a_2 \lambda^2 + \cdots + a_{n-1} \lambda^{n-1} + (-1)^n \lambda^n, \quad a_0, \dots, a_{n-1} \in \mathbb{R}, \quad (5.131)$$

is the **characteristic polynomial** of A . In particular,

$$a_0 = \det(A), \quad (5.132)$$

$$a_{n-1} = (-1)^{n-1} \text{tr}(A), \quad (5.133)$$

where $\text{tr}(A) = \sum_{i=1}^n a_{ii}$ is the **trace** of A and defined as the sum of the diagonal elements of A .

Theorem 6

$\lambda \in \mathbb{R}$ is eigenvalue of $A \in \mathbb{R}^{n \times n}$ if and only if λ is a root of the characteristic polynomial $p(\lambda)$ of A .

Remark 36

1. If λ is an eigenvalue of $A \in \mathbb{R}^{n \times n}$ then the corresponding eigenspace E_λ is the solution space of the homogeneous linear equation system $(A - \lambda I_n)\mathbf{x} = \mathbf{0}$.
2. Similar matrices possess the same characteristic polynomial.

5.11.2 Example: Eigenspace Computation

$$\bullet A = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$$

1. Characteristic polynomial: $p(\lambda) = |A - \lambda I_2| = \begin{vmatrix} 1 - \lambda & 0 \\ 1 & 1 - \lambda \end{vmatrix} = (1 - \lambda)^2$. Therefore $\lambda = 1$ is the only root of p and, therefore, the only eigenvalue of A
2. To compute the eigenspace for the eigenvalue $\lambda = 1$, we need to compute the null space of $A - I$:

$$(A - 1 \cdot I)\mathbf{x} = \mathbf{0} \Leftrightarrow \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} \mathbf{x} = \mathbf{0} \quad (5.134)$$

$$\Rightarrow E_1 = \left[\begin{bmatrix} 0 \\ 1 \end{bmatrix} \right] \quad (5.135)$$

²⁷It turns out that it is sufficient to work directly with the corresponding transformation mappings $A_\Phi \in \mathbb{R}^{n \times n}$.

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

1. Characteristic polynomial: $p(\lambda) = \det(A - \lambda I) = \lambda^2 + 1$. For $\lambda \in \mathbb{R}$ there exist no eigenvalue of A . However, for $\lambda \in \mathbb{C}$ we find $\lambda_1 = i$, $\lambda_2 = -i$.
2. The corresponding eigenspaces (for $\lambda_i \in \mathbb{C}$) are

$$E_i = \left[\begin{bmatrix} 1 \\ i \end{bmatrix} \right], \quad E_{-i} = \left[\begin{bmatrix} 1 \\ -i \end{bmatrix} \right]. \quad (5.136)$$

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

1. Characteristic polynomial:

$$p(\lambda) = \begin{vmatrix} -\lambda & -1 & 1 & 1 \\ -1 & 1-\lambda & -2 & 3 \\ 2 & -1 & -\lambda & 0 \\ 1 & -1 & 1 & -\lambda \end{vmatrix} = \begin{vmatrix} -\lambda & -1 & 1 & 1 \\ 0 & -\lambda & -1 & 3-\lambda \\ 0 & 1 & -2-\lambda & 2\lambda \\ 1 & -1 & 1 & -\lambda \end{vmatrix} \quad (5.137)$$

$$= \begin{vmatrix} -\lambda & -1-\lambda & 0 & 1 \\ 0 & -\lambda & -1-\lambda & 3-\lambda \\ 0 & 1 & -1-\lambda & 2\lambda \\ 1 & 0 & 0 & -\lambda \end{vmatrix} \quad (5.138)$$

$$= (-\lambda) \begin{vmatrix} -\lambda & -1-\lambda & 3-\lambda \\ 1 & -1-\lambda & 2\lambda \\ 0 & 0 & -\lambda \end{vmatrix} - \begin{vmatrix} -1-\lambda & 0 & 1 \\ -\lambda & -1-\lambda & 3-\lambda \\ 1 & -1-\lambda & 2\lambda \end{vmatrix} \quad (5.139)$$

$$= (-\lambda)^2 \begin{vmatrix} -\lambda & -1-\lambda \\ 1 & -1-\lambda \end{vmatrix} - \begin{vmatrix} -1-\lambda & 0 & 1 \\ -\lambda & -1-\lambda & 3-\lambda \\ 1 & -1-\lambda & 2\lambda \end{vmatrix} \quad (5.140)$$

$$= (1+\lambda)^2(\lambda^2 - 3\lambda + 2) = (1+\lambda)^2(1-\lambda)(2-\lambda) \quad (5.141)$$

Therefore, the eigenvalues of A are $\lambda_1 = -1$, $\lambda_2 = 1$, $\lambda_3 = 2$.

2. The corresponding eigenspaces are the solutions of $(A - \lambda_i I)x = 0$, $i = 1, 2, 3$, and given by

$$E_{-1} = \left[\begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix} \right], \quad E_1 = \left[\begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \right], \quad E_2 = \left[\begin{bmatrix} 1 \\ 0 \\ 1 \\ 1 \end{bmatrix} \right]. \quad (5.142)$$

5.11.3 Eigenvalues in Practical Applications

Eigenvalues and eigenvectors are essential concepts that are widely used.

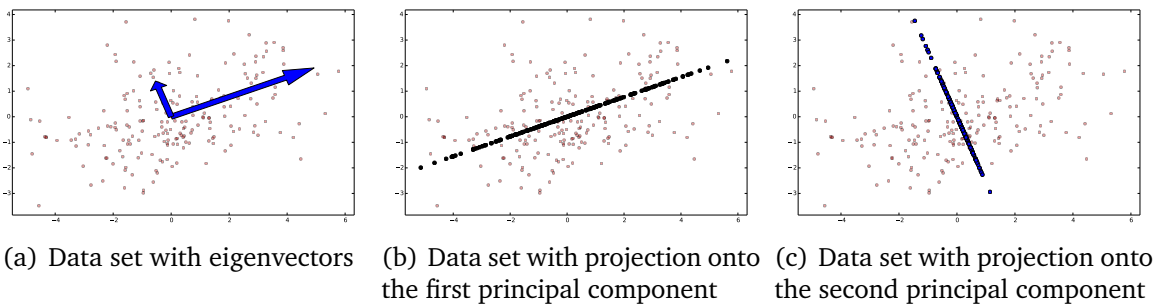


Figure 5.6: (a) Two-dimensional data set (red) with corresponding eigenvectors (blue), scaled by the magnitude of the corresponding eigenvectors. The longer the eigenvector the higher the variability (spread) of the data along this axis. (b) For optimal (linear) dimensionality reduction, we would project the data onto the subspace spanned by the eigenvector associated with the largest eigenvalue. (c) Projection of the data set onto the subspace spanned by the eigenvector associated with the smaller eigenvalue leads to a larger projection error.

- Eigenvectors and eigenvalues are fundamental to **principal component analysis** (PCA²⁸, Hotelling (1933)), which is commonly used for dimensionality reduction in face recognition, data visualization and other machine learning applications. Eigenvalues of a “data matrix” tell us, which dimensions of the high-dimensional data contain a significant signal/information.²⁹ These dimensions are important to keep, whereas dimensions associated with small eigenvalues can be discarded without much loss. The eigenvectors associated with the largest eigenvalues are called **principal components**. Figure 5.6 illustrates this for two-dimensional data.
- Eigenvalues were used by Claude Shannon to determine the theoretical limit to how much information can be transmitted through a communication medium like your telephone line or through the air. This is done by calculating the eigenvectors and eigenvalues of the communication channel (expressed as a matrix), and then waterfilling on the eigenvalues. The eigenvalues are then, in essence, the gains of the fundamental modes of the channel, which themselves are captured by the eigenvectors.
- Google uses the eigenvector corresponding to the maximal eigenvalue of the Google matrix to determine the rank of a page for search. The idea that the PageRank algorithm³⁰ brought up was that the importance of any web page can be judged by looking at the pages that link to it. For this, we write down all websites as a huge directed graph that shows which page links to which.

²⁸also known as Kosambi-KarhunenLoève Transform

²⁹To be more precise, PCA can be considered a method for (1) performing a basis change from the standard basis toward the eigenbasis in \mathbb{R}^d , $d < n$, (2) projecting the data in \mathbb{R}^n onto the subspace spanned by the d eigenvectors corresponding to the d largest eigenvalues (which are called the principal components).

³⁰Developed at Stanford University by Larry Page and Sergey Brin in 1996.

PageRank computes the weight (importance) $x_i \geq 0$ of a website a_i by counting the number of pages pointing to a_i . PageRank also take the importance of the website into account that links to a website to a_i . Then, the navigation behavior of a user can be described by a transition matrix A of this graph that tells us with what (click) probability somebody will end up on a different website. The matrix A has the property that for any initial rank/importance vector x of a website the sequence x, Ax, A^2x, \dots converges to a vector x^* . This vector is called the **PageRank** and satisfies $Ax^* = x^*$, i.e., it is an eigenvector (with corresponding eigenvalue 1) of A .³¹ Liesen and Mehrmann (2015) provide a more detailed description of PageRank with some examples. More details and different perspectives on PageRank can be found at <http://tinyurl.com/83tehpk>.

- Eigenvalues are frequently used to determine **numerical stability**, e.g., when inverting matrices. Since a computer can only represent numbers with a finite precision, we often look at **condition numbers** of matrices, i.e., the ratio $\frac{|\lambda_{\max}|}{|\lambda_{\min}|}$ of the biggest to the smallest eigenvalue. If this ratio exceeds a threshold (e.g., 10^8), matrix inversion may become numerically unstable and lead to inaccurate results.
- Consider a linear time-invariant system that evolves according to $x_{t+1} = Ax_t$, where x is called the “state” and A the “transition matrix”. We are often interested in analyzing the limit-behavior of such a system for $t \rightarrow \infty$. This requires us to compute $\lim_{t \rightarrow \infty} A^t x_0$ for some initial state x_0 . To analyze the limit behavior (convergence to a fixed point or a limit cycle or divergence), we perform a basis change into the eigenbasis, such that $A = S^{-1}DS$, where D is a diagonal matrix with the eigenvalues on its diagonal.³² Then, $\lim_{t \rightarrow \infty} A^t x_0 = (S^{-1}DS)^t x_0 = S^{-1}D^t S x_0$. Therefore, the eigenvalues of A can tell us something about convergence/divergence (e.g., if there is an eigenvalue $\lambda > 1$ of A , the system diverges for $t \rightarrow \infty$).

5.12 Diagonalization

Diagonal matrices are possess a very simple structure and they allow for a very fast computation of determinants and inverses, for instance. In this section, we will have a closer look at how to transform matrices into diagonal form. More specifically, we will look at endomorphisms of finite-dimensional vector spaces, which are similar to a diagonal matrix, i.e., endomorphisms whose transformation matrix attains diagonal structure for a suitable basis. Here, we finally have a practically important application of the basis change we discussed in Section 5.9.3 and eigenvalues form Section 5.11.

³¹When normalizing x^* , such that $\|x^*\| = 1$ we can interpret the entries as probabilities.

³²We will discuss this in Section 5.12.

Definition 21 (Diagonal Form)

A matrix $A \in \mathbb{R}^{n \times n}$ is **diagonalizable** if it is similar³³ to a diagonal matrix

$$\begin{bmatrix} c_1 & 0 & \cdots & \cdots & 0 \\ 0 & c_2 & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & & \vdots \\ 0 & \cdots & 0 & c_{n-1} & 0 \\ 0 & \cdots & \cdots & 0 & c_n \end{bmatrix}. \quad (5.143)$$

Theorem 7

For an endomorphism Φ of an n -dimensional \mathbb{R} -vector space V the following statements are equivalent:

1. Φ is diagonalizable.
2. The transformation matrix A_Φ is diagonalizable.
3. There exists a basis in V consisting of the eigenvectors of Φ .³⁴
4. The sum of the dimensions of the eigenspaces of Φ is n .³⁵

Theorem 8

For an n -dimensional \mathbb{R} -vector space V and a linear mapping $\Phi : V \rightarrow V$ the following holds: Φ is diagonalizable if and only if

1. Its characteristic polynomial $p(\lambda)$ is given in the form

$$p(\lambda) = (-1)^n (\lambda - c_1)^{r_1} \cdots (\lambda - c_k)^{r_k} \quad (5.144)$$

with $r_i \in \mathbb{N}$ and pairwise different roots $c_i \in \mathbb{R}$ and

2. For $i = 1, \dots, k$

$$\dim(\text{Im}(\Phi - c_i \text{id}_V)) = n - r_i \quad (5.145)$$

$$\stackrel{\text{Rank-Nullity}}{\iff} \dim(\ker(\Phi - c_i \text{id}_V)) = r_i. \quad (5.146)$$

In (5.144) we say that the characteristic polynomial decomposes into **linear factors**. The second requirement in (5.145) says that the dimension of the eigenspace E_{c_i} must correspond to the (algebraic) multiplicity r_i of the eigenvalues in the characteristic polynomial, $i = 1, \dots, k$.³⁶ The dimension of the eigenspace E_{c_i} is the dimension of the kernel/null space of $\Phi - c_i \text{id}_V$.

³³Remember: Two matrices A, D are similar if and only if there exists an invertible matrix S , such that $D = S^{-1}AS$.

³⁴Therefore, we need n eigenvectors.

³⁵In particular, an endomorphism Φ of an n -dimensional \mathbb{R} -vector space with n different eigenvalues is diagonalizable.

³⁶More explicitly, the **algebraic multiplicity** of c_i is the number of times r_i it is repeated as a root of the characteristic polynomial.

Theorem 8 holds equivalently if we replace Φ with $A \in \mathbb{R}^{n \times n}$ and id_V with I_n .

If Φ is diagonalizable, we can find a suitable basis such that Φ possesses a transformation matrix of the form

$$A_\Phi = \begin{bmatrix} c_1 & 0 & \cdots & \cdots & \cdots & \cdots & 0 \\ 0 & \ddots & \ddots & & & & \vdots \\ \vdots & \ddots & c_1 & \ddots & & & \vdots \\ \vdots & & \ddots & \ddots & \ddots & & \vdots \\ \vdots & & & \ddots & c_k & \ddots & \vdots \\ \vdots & & & & \ddots & \ddots & 0 \\ 0 & \cdots & \cdots & \cdots & \cdots & 0 & c_k \end{bmatrix}, \quad (5.147)$$

where each eigenvalue c_i appears r_i times (its multiplicity in the characteristic polynomial and the dimension of the corresponding eigenspace) on the diagonal.

Example

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

1. Characteristic polynomial: $p(\lambda) = (1 - \lambda)^2$
2. The dimension of eigenspace is $2 - \text{rk}(A - I) = 1 \neq 2$. Because of Theorem 8 A is not diagonalizable.

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

- Characteristic polynomial: $p(\lambda) = 1 + \lambda^2$.
 - For $\lambda \in \mathbb{R}$ there exist no roots of $p(\lambda)$, and the characteristic polynomial does not decompose into linear factors. Therefore, A is not diagonalizable.
 - If we consider a \mathbb{C} -vector space and $\lambda \in \mathbb{C}$, then $p(\lambda) = (i - \lambda)(-i - \lambda)$ and A has two eigenvalues, the characteristic polynomial decomposes into linear factors, and A is diagonalizable (verify the dimension of the eigenspaces).

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

1. Characteristic polynomial: $p(\lambda) = (1 + \lambda)^2(1 - \lambda)(2 - \lambda)$. The eigenvalues are $c_1 = -1, c_2 = 1, c_3 = 2$ with algebraic multiplicities $r_1 = 2, r_2 = 1, r_3 = 1$, respectively.
2. Dimension of eigenspaces: $\dim(E_{c_1}) = 1 \neq r_1$.

Therefore, A cannot be diagonalized.

$$\bullet A = \begin{bmatrix} 3 & 2 & -1 \\ 2 & 6 & -2 \\ 0 & 0 & 2 \end{bmatrix}.$$

1. Characteristic polynomial: $p(\lambda) = (2 - \lambda)^2(7 - \lambda)$. Therefore, $c_1 = 2, c_2 = 7, r_1 = 2, r_2 = 1$
2. Dimension of eigenspaces: $\dim(A - c_1 I_3) = 2 = r_1, \dim(A - c_2 I_3) = 1 = r_2$

Therefore, A is diagonalizable.

Let us now discuss a practical way of constructing diagonal matrices.

Remark 37

If $A \in \mathbb{R}^{n \times n}$ is diagonalizable and $(\mathbf{b}_1, \dots, \mathbf{b}_n)$ is an ordered basis of eigenvectors of A with $A\mathbf{b}_i = c_i\mathbf{b}_i, i = 1, \dots, n$, then it holds that for the regular matrix $S = (\mathbf{b}_1 | \dots | \mathbf{b}_n)$

$$S^{-1}AS = \begin{bmatrix} c_1 & 0 & \cdots & \cdots & 0 \\ 0 & c_2 & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \cdots & 0 & c_{n-1} & 0 \\ 0 & \cdots & \cdots & 0 & c_n \end{bmatrix}. \quad (5.148)$$

The diagonal matrix in (5.148) is the transformation matrix of $x \mapsto Ax$ with respect to the eigenbasis $(\mathbf{b}_1, \dots, \mathbf{b}_n)$.

Example

Coming back to the above example, where we wanted to determine the diagonal form of

$$A = \begin{bmatrix} 3 & 2 & -1 \\ 2 & 6 & -2 \\ 0 & 0 & 2 \end{bmatrix}. \quad (5.149)$$

From above, we already know that A is diagonalizable. We now determine the eigenbasis of \mathbb{R}^3 that allows us to transform A into a similar matrix in diagonal form via $S^{-1}AS$:

1. The eigenspaces are

$$E_2 = \underbrace{\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}}_{=: \mathbf{b}_1}, \underbrace{\begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix}}_{=: \mathbf{b}_2}, \quad E_7 = \underbrace{\begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix}}_{=: \mathbf{b}_3} \quad (5.150)$$

2. We now collect the eigenvectors in a matrix and obtain

$$S = (\mathbf{b}_1 | \mathbf{b}_2 | \mathbf{b}_3) = \begin{bmatrix} 1 & -2 & 1 \\ 0 & 1 & 2 \\ 1 & 0 & 0 \end{bmatrix} \quad (5.151)$$

such that

$$S^{-1}AS = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 7 \end{bmatrix}. \quad (5.152)$$

Remark 38

The dimension of the eigenspace \mathbb{E}_λ is called **geometric multiplicity** of λ and cannot be greater than the algebraic multiplicity of the corresponding eigenvalue λ .

Remark 39

So far, we computed diagonal matrices as $D = S^{-1}AS$. However, we can equally write $A = SDS^{-1}$. Here, we can interpret the transformation matrix A as follows: S^{-1} performs a basis change from the standard basis into the eigenbasis. Then, D then scales the vector along the axes of the eigenbasis, and S transforms the scaled vectors back into the standard/canonical coordinates.

5.12.1 Applications

Diagonal matrices $D = S^{-1}AS$ exhibit the nice properties that they can be easily raised to a power:

$$A^k = (SDS^{-1})^k = SD^kS^{-1}. \quad (5.153)$$

Computing D^k is easy because we apply this operation individually to any diagonal element. As an example, this allows to compute inverses of D in $\mathcal{O}(n)$ instead of $\mathcal{O}(n^3)$.

A different property of diagonal matrices is that they decouple variables. This is important in probability theory to interpret random variables, e.g., for Gaussian distributions.

With diagonal matrices, it is easier to analyze properties of differential equations, which play an important role in any kind of (linear) dynamical system.

5.12.2 Cayley-Hamilton Theorem

Theorem 9 (Cayley-Hamilton)

Let V be an n -dimensional \mathbb{R} -vector space and $\Phi : V \rightarrow V$ an endomorphism with transformation matrix A_Φ and characteristic polynomial p . Then,

$$p(\Phi) = 0 \quad (5.154)$$

(and equivalently, $p(A_\Phi) = \mathbf{0}$).

Remark 40

- Note that the right hand side of (5.154) is the zero mapping (or the $\mathbf{0}$ -matrix when we use the transformation matrix A_Φ).
- The importance of the Cayley-Hamilton theorem is not the existence of a (non-trivial) polynomial q , such that $q(\Phi) = 0$, but that the characteristic polynomial has this property.

Applications

- Find an expression for A^{-1} in terms of $I, A, A^2, \dots, A^{n-1}$.

Example

$A = \begin{bmatrix} 1 & -1 \\ 2 & 1 \end{bmatrix}$ has the characteristic polynomial $p(\lambda) = \lambda^2 - 2\lambda + 3$. Then Theorem 9 states that $A^2 - 2A + 3I = \mathbf{0}$ and, therefore,

$$-A^2 + 2A = 3I \quad (5.155)$$

$$\Leftrightarrow A \frac{1}{3}(-A + 2I) = I \quad (5.156)$$

$$\Leftrightarrow A^{-1} = \frac{1}{3}(2I - A). \quad (5.157)$$

Generally, for an invertible matrix $A \in \mathbb{R}^{n \times n}$, A^{-1} can be written as an $(n-1)$ -th order polynomial expression in A .

- Find an expression of A^m , $m \geq n$, in terms of $I, A, A^2, \dots, A^{n-1}$

5.13 Scalar Products

Definition 22

Let $\beta : V \times V \rightarrow \mathbb{R}$ be a bilinear mapping (i.e., linear in both arguments).

- β is called **symmetric** if $\beta(x, y) = \beta(y, x)$ for all $x, y \in V$.
- β is called **positive definite** if for all $x \neq \mathbf{0}$: $\beta(x, x) > 0$. $\beta(\mathbf{0}, \mathbf{0}) = 0$.
- A positive definite, symmetric bilinear mapping $\beta : V \times V \rightarrow \mathbb{R}$ is called **scalar product/dot product/inner product** on V . We typically write $\langle x, y \rangle$ instead of $\beta(x, y)$.
- The pair $(V, \langle \cdot, \cdot \rangle)$ is called **Euclidean vector space** or (real) **vector space with scalar product**.

Example

- For $V = \mathbb{R}^n$ we define the **standard scalar product** $\langle x, y \rangle := x^\top y = \sum_{i=1}^n x_i y_i$.

- $V = \mathbb{R}^2$. If we define $\beta(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle := x_1 y_1 - (x_1 y_2 + x_2 y_1) + 2x_2 y_2$ then β is a scalar product but different from the standard scalar product.

In a Euclidean vector space, the scalar product allows us to introduce concepts, such as lengths, distances and orthogonality.

5.13.1 Lengths, Distances, Orthogonality

Definition 23 (Norm)

Consider a Euclidean vector space $(V, \langle \cdot, \cdot \rangle)$. Then $\|\mathbf{x}\| := \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$ is the **length** or **norm** of $\mathbf{x} \in V$. The mapping

$$\|\cdot\| : V \rightarrow \mathbb{R} \quad (5.158)$$

$$\mathbf{x} \mapsto \|\mathbf{x}\| \quad (5.159)$$

is called **norm**.

Example

In geometry, we are often interested in lengths of vectors. We can now use the scalar product to compute them. For instance, in a Euclidean vector space with standard scalar product, if $\mathbf{x} = [1, 2]^\top$ then its norm/length is $\|\mathbf{x}\| = \sqrt{1^2 + 2^2} = \sqrt{5}$

Remark 41

The norm $\|\cdot\|$ possesses the following properties:

1. $\|\mathbf{x}\| \geq 0$ for all $\mathbf{x} \in V$ and $\|\mathbf{x}\| = 0 \Leftrightarrow \mathbf{x} = \mathbf{0}$
2. $\|\lambda \mathbf{x}\| = |\lambda| \cdot \|\mathbf{x}\|$ for all $\mathbf{x} \in V$ and $\lambda \in \mathbb{R}$
3. **Minkowski inequality:** $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$ for all $\mathbf{x}, \mathbf{y} \in V$

Definition 24 (Distance and Metric)

Consider a Euclidean vector space $(V, \langle \cdot, \cdot \rangle)$. Then $d(\mathbf{x}, \mathbf{y}) := \|\mathbf{x} - \mathbf{y}\|$ is called **distance** of $\mathbf{x}, \mathbf{y} \in V$. The mapping

$$d : V \times V \rightarrow \mathbb{R} \quad (5.160)$$

$$(\mathbf{x}, \mathbf{y}) \mapsto d(\mathbf{x}, \mathbf{y}) \quad (5.161)$$

is called **metric**.

A metric d satisfies:

1. d is positive definite, i.e., $d(\mathbf{x}, \mathbf{y}) \geq 0$ for all $\mathbf{x}, \mathbf{y} \in V$ and $d(\mathbf{x}, \mathbf{y}) = 0 \Leftrightarrow \mathbf{x} = \mathbf{y}$

2. d is symmetric, i.e., $d(\mathbf{x}, \mathbf{y}) = d(\mathbf{y}, \mathbf{x})$ for all $\mathbf{x}, \mathbf{y} \in V$.
3. **Triangular inequality:** $d(\mathbf{x}, \mathbf{z}) \leq d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})$.

Definition 25 (Orthogonality)

Vectors \mathbf{x} and \mathbf{y} are **orthogonal** if $\langle \mathbf{x}, \mathbf{y} \rangle = 0$, and we write $\mathbf{x} \perp \mathbf{y}$

Theorem 10

Let $(V, \langle \cdot, \cdot \rangle)$ be a Euclidean vector space and $\mathbf{x}, \mathbf{y}, \mathbf{z} \in V$. Then:

1. **Cauchy-Schwarz inequality:** $|\langle \mathbf{x}, \mathbf{y} \rangle| \leq \|\mathbf{x}\| \|\mathbf{y}\|$
2. **Minkowski inequality:** $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$
3. **Triangular inequality:** $d(\mathbf{x}, \mathbf{z}) \leq d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})$
4. **Parallelogram law:** $\|\mathbf{x} + \mathbf{y}\|^2 + \|\mathbf{x} - \mathbf{y}\|^2 = 2\|\mathbf{x}\|^2 + 2\|\mathbf{y}\|^2$
5. $4\langle \mathbf{x}, \mathbf{y} \rangle = \|\mathbf{x} + \mathbf{y}\|^2 - \|\mathbf{x} - \mathbf{y}\|^2$
6. $\mathbf{x} \perp \mathbf{y} \Leftrightarrow \|\mathbf{x} + \mathbf{y}\|^2 = \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2$

The Cauchy-Schwarz inequality allows us to define angles ω in Euclidean vector spaces between two vectors \mathbf{x}, \mathbf{y} . Assume that $\mathbf{x} \neq \mathbf{0}, \mathbf{y} \neq \mathbf{0}$. Then

$$-1 \leq \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{x}\| \|\mathbf{y}\|} \leq 1 \quad (5.162)$$

Therefore, there exists a unique $\omega \in [0, \pi)$ with

$$\cos \omega = \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{x}\| \|\mathbf{y}\|} \quad (5.163)$$

The number ω is the **angle** between \mathbf{x} and \mathbf{y} .

5.13.2 Applications

Scalar products allow us to compute angles between vectors or distances. A major purpose of scalar products is to determine whether vectors are orthogonal to each other; in this case $\langle \mathbf{x}, \mathbf{y} \rangle = 0$. This will play an important role when we discuss projections in Section 5.14. The scalar product also allows us to determine specific bases of vector (sub)spaces, where each vector is orthogonal to all others (orthogonal bases) using the Gram-Schmidt method³⁷. These bases are important optimization and numerical algorithms for solving linear equation systems. For instance, Krylov subspace methods³⁸, such as Conjugate Gradients or GMRES, minimize residual errors that are orthogonal to each other (Stoer and Burlirsch, 2002).

³⁷not discussed in this course

³⁸The basis for the Krylov subspace is derived from the Cayley-Hamilton theorem, see Section 5.12.2, which allows us to compute the inverse of a matrix in terms of a linear combination of its powers.




Figure 5.7: The shade is the projection of the ball onto a plane (table) with the center being the light source. Adapted from <http://tinyurl.com/ka4t28t>

In machine learning, scalar products are important in the context of kernel methods (Schölkopf and Smola, 2002). Kernel methods exploit the fact that many linear algorithms can be expressed purely by scalar product computations.³⁹ Then, the “kernel trick” allows us to compute these scalar products implicitly in a (potentially infinite-dimensional) feature space, without even knowing this feature space explicitly. This allowed the “non-linearization” of many algorithms used in machine learning, such as kernel-PCA (Schölkopf et al., 1998) for dimensionality reduction. Gaussian processes (Rasmussen and Williams, 2006) also fall into the category of kernel methods and are the current state-of-the-art in probabilistic regression (fitting curves to data points).

5.14 Projections

Projections are an important class of linear transformations (besides rotations and reflections⁴⁰). Projections play an important role in graphics (see e.g., Figure 5.7), coding theory, statistics and machine learning. We often deal with data that is very high-dimensional. However, often only a few dimensions are important. In this case, we can project the original very high-dimensional data into a lower-dimensional feature space⁴¹ and work in this lower-dimensional space to learn more about the data set and extract patterns. For example, machine learning tools such as Principal Component Analysis (PCA) by Hotelling (1933) and Deep Neural Networks (e.g., deep auto-encoders, first applied by Deng et al. (2010)) heavily exploit this idea. In the following, we will focus on linear orthogonal projections.

³⁹Matrix-vector multiplication $Ax = b$ falls into this category since b_i is a scalar product of the i th row of A with x .

⁴⁰Note that translations are not linear. Why?

⁴¹“Feature” is just a commonly word for “data representation”.

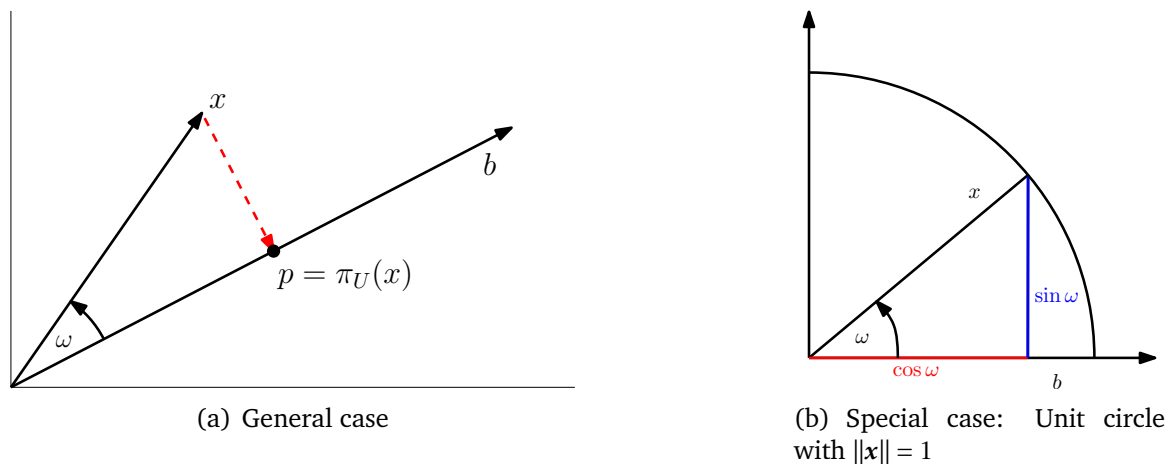


Figure 5.8: Projection $\pi_U(x)$ of x onto a subspace U with basis b .

Definition 26 (Projection)

Let V be an \mathbb{R} -vector space and $W \subset V$ a subspace of V . A linear mapping $\pi : V \rightarrow W$ is called a **projection** if $\pi^2 = \pi \circ \pi = \pi$.

Remark 42 (Projection matrix)

Since linear mappings can be expressed by transformation matrices, the definition above applies equally to a special kind of transformation matrices, the **projection matrices** P_π , which exhibit the property that $P_\pi^2 = P_\pi$.

Since $P_\pi^2 = P_\pi$ it follows that all eigenvalues of P_π are either 0 or 1. The corresponding eigenspaces are the kernel and image of the projection, respectively.⁴²

In the following, we will derive projections from points in the Euclidean vector space $(\mathbb{R}^n, \langle \cdot, \cdot \rangle)$ onto subspaces. We will start with one-dimensional subspaces, the lines. If not mentioned otherwise, we assume the standard scalar product $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^\top \mathbf{y}$.

5.14.1 Projection onto a Line

Assuming we are given a line through the origin with basis vector $\mathbf{b} \in \mathbb{R}^m$.⁴³ This is a one-dimensional subspace $U \subset \mathbb{R}^m$ spanned by \mathbf{b} . When we project $\mathbf{x} \in \mathbb{R}^m$ onto U , we want to find the point $\pi_U(\mathbf{x}) = \mathbf{p} \in U$ that is closest to \mathbf{x} . Let us characterize some properties of the projection $\mathbf{p} = \pi_U(\mathbf{x})$ (Fig. 5.8(a) serves as an illustration):

- \mathbf{p} is closest to \mathbf{x} , where “closest” implies that $\|\mathbf{x} - \mathbf{p}\|$ is minimal. This means, the segment $\mathbf{p} - \mathbf{x}$ from \mathbf{p} to \mathbf{x} is orthogonal to U and, therefore, the basis \mathbf{b} of U . The orthogonality condition yields $\langle \mathbf{p} - \mathbf{x}, \mathbf{b} \rangle = (\mathbf{p} - \mathbf{x})^\top \mathbf{b} = 0$ (when we assume the standard scalar product).
- The projection $\mathbf{p} = \pi_U(\mathbf{x})$ of \mathbf{x} onto U must be a multiple of the basis/direction vector \mathbf{b} .⁴⁴ Thus, $\mathbf{p} = \lambda \mathbf{b}$, for some $\lambda \in \mathbb{R}$.

⁴²A good illustration is given here: <http://tinyurl.com/p5jn5ws>.

⁴³In affine geometry, the space spanned by \mathbf{b} is also called a **direction (space)**.

⁴⁴More generally, a linear combination of the basis vectors of U .

In the following three steps, we determine λ , the projection point $\pi_U(\mathbf{x}) = \mathbf{p} \in U$ and the projection matrix \mathbf{P}_π that maps arbitrary $\mathbf{x} \in \mathbb{R}^m$ onto U .

1. Finding λ . We know that $\mathbf{p} = \lambda\mathbf{b}$. Therefore,

$$\mathbf{x} - \mathbf{p} = \mathbf{x} - \lambda\mathbf{b} \perp \mathbf{b} \quad (5.164)$$

$$\Leftrightarrow \langle \mathbf{x} - \lambda\mathbf{b}, \mathbf{b} \rangle = 0 \quad (5.165)$$

$$\Leftrightarrow \langle \mathbf{x}, \mathbf{b} \rangle - \lambda \langle \mathbf{b}, \mathbf{b} \rangle = 0 \quad (5.166)$$

$$\Leftrightarrow \lambda = \frac{\langle \mathbf{x}, \mathbf{b} \rangle}{\langle \mathbf{b}, \mathbf{b} \rangle} = \frac{\langle \mathbf{x}, \mathbf{b} \rangle}{\|\mathbf{b}\|^2} \quad (5.167)$$

With $\langle \cdot, \cdot \rangle$ being the standard scalar product, we get

$$\lambda = \frac{\mathbf{b}^\top \mathbf{x}}{\mathbf{b}^\top \mathbf{b}} = \frac{\mathbf{b}^\top \mathbf{x}}{\|\mathbf{b}\|^2} \quad (5.168)$$

2. Finding the projection point $\mathbf{p} = \pi_U(\mathbf{x}) \in U$. Since $\mathbf{p} = \lambda\mathbf{b}$ we immediately obtain with (5.168) that

$$\mathbf{p} = \lambda\mathbf{b} = \frac{\mathbf{b}^\top \mathbf{x}}{\|\mathbf{b}\|^2} \mathbf{b}. \quad (5.169)$$

We can also compute the length of \mathbf{p} (i.e., the distance from $\mathbf{0}$) by means of Definition 23:

$$\|\mathbf{p}\| = \|\lambda\mathbf{b}\| = |\lambda| \|\mathbf{b}\| = \frac{|\mathbf{b}^\top \mathbf{x}|}{\|\mathbf{b}\|^2} \|\mathbf{b}\| = |\cos \omega| \|\mathbf{x}\| \|\mathbf{b}\| \frac{\|\mathbf{b}\|}{\|\mathbf{b}\|^2} = \cos \omega \|\mathbf{x}\|. \quad (5.170)$$

Here, ω is the angle between \mathbf{x} and \mathbf{b} . This equation should be familiar from trigonometry: If $\|\mathbf{x}\| = 1$ it lies on the unit circle. Then the projection onto the horizontal axis \mathbf{b} is exactly $\cos \omega$. An illustration is given in Figure 5.8(b)

3. Finding the projection matrix \mathbf{P}_π . We know that a projection is a linear mapping (see Definition 26). Therefore, there exists a projection matrix \mathbf{P}_π , such that $\pi_U(\mathbf{x}) = \mathbf{p} = \mathbf{P}_\pi \mathbf{x}$. With

$$\mathbf{p} = \lambda\mathbf{b} = \mathbf{b} \frac{\mathbf{b}^\top \mathbf{x}}{\|\mathbf{b}\|^2} = \frac{\mathbf{b}\mathbf{b}^\top}{\|\mathbf{b}\|^2} \mathbf{x} \quad (5.171)$$

we immediately see that

$$\mathbf{P}_\pi = \frac{\mathbf{b}\mathbf{b}^\top}{\|\mathbf{b}\|^2}. \quad (5.172)$$

Note that $\mathbf{b}\mathbf{b}^\top$ is a matrix (with rank 1) and $\|\mathbf{b}\|^2 = \langle \mathbf{b}, \mathbf{b} \rangle$ is a scalar.

The projection matrix \mathbf{P}_π projects any vector $\mathbf{x} \in \mathbb{R}^m$ onto the line through the origin with direction \mathbf{b} (equivalently, the subspace U spanned by \mathbf{b}).

Example (Projection onto a Line)

Find the projection matrix P_π onto the line through the origin spanned by $\mathbf{b} = \begin{bmatrix} 1 & 2 & 2 \end{bmatrix}^\top$. \mathbf{b} is a direction and a basis of the one-dimensional subspace (line through origin).

With (5.172), we obtain

$$\mathbf{P}_\pi = \frac{\mathbf{b}\mathbf{b}^\top}{\mathbf{b}^\top\mathbf{b}} = \frac{1}{9} \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix} \begin{bmatrix} 1 & 2 & 2 \end{bmatrix} = \frac{1}{9} \begin{bmatrix} 1 & 2 & 2 \\ 2 & 4 & 4 \\ 2 & 4 & 4 \end{bmatrix}. \quad (5.173)$$

Let us now choose a particular \mathbf{x} and see whether it lies in the subspace spanned by \mathbf{b} . For $\mathbf{x} = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^\top$, the projected point is

$$\mathbf{p} = \mathbf{P}_\pi \mathbf{x} = \frac{1}{9} \begin{bmatrix} 1 & 2 & 2 \\ 2 & 4 & 4 \\ 2 & 4 & 4 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} = \frac{1}{9} \begin{bmatrix} 5 \\ 10 \\ 10 \end{bmatrix} \in \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}. \quad (5.174)$$

Note that the application of \mathbf{P}_π to \mathbf{p} does not change anything, i.e., $\mathbf{P}_\pi \mathbf{p} = \mathbf{p}$.⁴⁵ This is expected because according to Definition 26 we know that a projection matrix \mathbf{P}_π satisfies $\mathbf{P}_\pi^2 \mathbf{x} = \mathbf{P}_\pi \mathbf{x}$.

5.14.2 Projection onto General Subspaces

In the following, we generalize projections to the case where we project vectors $\mathbf{x} \in \mathbb{R}^n$ onto general subspaces $U \subset \mathbb{R}^n$. Assume that $(\mathbf{b}_1, \dots, \mathbf{b}_m)$ is an ordered basis of U .⁴⁶

Projections $\pi_U(\mathbf{x})$ onto U are elements of U . Therefore, they exhibit the property that they can be represented as linear combinations of the basis vectors of U , i.e., a projected point is mapped onto the subspace spanned by the columns of $\mathbf{B} \in \mathbb{R}^{n \times m}$, where $\mathbf{B} = (\mathbf{b}_1 | \dots | \mathbf{b}_m)$. Projections $\pi_U(\mathbf{x}) \in U$ of \mathbf{x} onto the subspace U spanned by $(\mathbf{b}_1, \dots, \mathbf{b}_m)$ can be expressed as linear combinations of the basis vectors of U , such that $\mathbf{p} = \sum_{i=1}^m \lambda_i \mathbf{b}_i$.

As before, we follow a three-step procedure to find \mathbf{p} and the projection matrix \mathbf{P}_π :

1. Find $\lambda_1, \dots, \lambda_m$, such that the linear combination $\mathbf{p} = \sum_{i=1}^m \lambda_i \mathbf{b}_i = \mathbf{B} \boldsymbol{\lambda}$ is closest to $\mathbf{x} \in \mathbb{R}^n$. As in the 1D case, “closest” means “minimum distance”, which implies that the line connecting $\mathbf{p} \in U$ and $\mathbf{x} \in \mathbb{R}^n$ must be orthogonal to all basis vectors of U . Therefore, we obtain the conditions (assuming the standard scalar product)

$$\langle \mathbf{b}_1, \mathbf{x} - \mathbf{p} \rangle = \mathbf{b}_1^\top (\mathbf{x} - \mathbf{p}) = 0 \quad (5.175)$$

⁴⁵ \mathbf{p} is therefore an eigenvector of \mathbf{P}_π , and the corresponding eigenvalue is 1.

⁴⁶If U is given by a set of spanning vectors, make sure you determine a basis $\mathbf{b}_1, \dots, \mathbf{b}_m$ before proceeding.

$$\vdots \quad (5.176)$$

$$\langle \mathbf{b}_m, \mathbf{x} - \mathbf{p} \rangle = \mathbf{b}_m^\top (\mathbf{x} - \mathbf{p}) = 0 \quad (5.177)$$

which, with $\mathbf{p} = \mathbf{B}\boldsymbol{\lambda}$, can be written as

$$\mathbf{b}_1^\top (\mathbf{x} - \mathbf{B}\boldsymbol{\lambda}) = 0 \quad (5.178)$$

$$\vdots \quad (5.179)$$

$$\mathbf{b}_m^\top (\mathbf{x} - \mathbf{B}\boldsymbol{\lambda}) = 0 \quad (5.180)$$

such that we obtain a homogeneous linear equation system

$$\begin{bmatrix} \mathbf{b}_1^\top \\ \vdots \\ \mathbf{b}_m^\top \end{bmatrix} \begin{bmatrix} \mathbf{x} - \mathbf{B}\boldsymbol{\lambda} \end{bmatrix} = \mathbf{0} \quad \Leftrightarrow \quad \mathbf{B}^\top (\mathbf{x} - \mathbf{B}\boldsymbol{\lambda}) = \mathbf{0}. \quad (5.181)$$

Factorizing yields

$$\mathbf{B}^\top (\mathbf{x} - \mathbf{B}\boldsymbol{\lambda}) = \mathbf{0} \quad \Leftrightarrow \quad \mathbf{B}^\top \mathbf{B}\boldsymbol{\lambda} = \mathbf{B}^\top \mathbf{x}, \quad (5.182)$$

and the expression on the right-hand side is called **normal equation**.⁴⁷ Since the vectors $\mathbf{b}_1, \dots, \mathbf{b}_m$ are a basis and, therefore, linearly independent, $\mathbf{B}^\top \mathbf{B}$ is regular and can be inverted. This allows us to solve for the optimal coefficients

$$\boldsymbol{\lambda} = (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{x}. \quad (5.183)$$

The matrix $(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top$ is often called the **pseudo-inverse** of \mathbf{B} , which can be computed for non-square matrices \mathbf{B} . It only requires that $\mathbf{B}^\top \mathbf{B}$ is positive definite. In practical applications (e.g., linear regression), we often add $\epsilon \mathbf{I}$ to $\mathbf{B}^\top \mathbf{B}$ to guarantee positive definiteness or increase numerical stability. This “ridge” can be rigorously derived using Bayesian inference.⁴⁸

2. Find the projection of $\pi_U(\mathbf{x}) = \mathbf{p} \in U$. We already established that $\mathbf{p} = \mathbf{B}\boldsymbol{\lambda}$. Therefore, with (5.183)

$$\mathbf{p} = \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{x}. \quad (5.184)$$

3. Find the projection matrix \mathbf{P}_π . From (5.184) we can immediately see that the projection matrix that solves $\mathbf{P}_\pi \mathbf{x} = \mathbf{p}$ must be

$$\mathbf{P}_\pi = \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top. \quad (5.185)$$

Remark 43

Comparing the solutions for projecting onto a one-dimensional subspace and the general case, we see that the general case includes the 1D case as a special case: If $\dim(U) = 1$ then $\mathbf{B}^\top \mathbf{B}$ is just a scalar and we can rewrite the projection matrix in (5.185) $\mathbf{P}_\pi = \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top$ as $\mathbf{P}_\pi = \frac{\mathbf{B}\mathbf{B}^\top}{\mathbf{B}^\top \mathbf{B}}$, which is exactly the projection matrix in (5.172).

⁴⁷You may see the normal equation again when you take courses on machine learning, state estimation or robotics.

⁴⁸For more details, please see CO-493 or other machine learning courses.

Example

For a subspace $U = \left[\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 2 \end{bmatrix} \right] \subset \mathbb{R}^3$ and $\mathbf{x} = \begin{bmatrix} 6 \\ 0 \\ 0 \end{bmatrix} \in \mathbb{R}^3$ find λ , the projection point \mathbf{p} and the projection matrix \mathbf{P}_π .

First, we see that the generating set of U is a basis (linear independence) and write the basis vectors of U into a matrix $\mathbf{B} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{bmatrix}$.

Second, we compute the matrix $\mathbf{B}^\top \mathbf{B}$ and the vector $\mathbf{B}^\top \mathbf{x}$ as

$$\mathbf{B}^\top \mathbf{B} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{bmatrix} = \begin{bmatrix} 3 & 3 \\ 3 & 5 \end{bmatrix} \quad \mathbf{B}^\top \mathbf{x} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \end{bmatrix} \begin{bmatrix} 6 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 6 \\ 0 \end{bmatrix} \quad (5.186)$$

Third, we solve the normal equation $\mathbf{B}^\top \mathbf{B} \lambda = \mathbf{B}^\top \mathbf{x}$ to find λ :

$$\begin{bmatrix} 3 & 3 \\ 3 & 5 \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} = \begin{bmatrix} 6 \\ 0 \end{bmatrix} \Rightarrow \lambda = \begin{bmatrix} 5 \\ -3 \end{bmatrix} \quad (5.187)$$

Fourth, the projection $\pi_U(\mathbf{x}) = \mathbf{p}$ of \mathbf{x} onto U , i.e., into the column space of \mathbf{B} , can be directly computed via

$$\mathbf{p} = \mathbf{B} \lambda = \begin{bmatrix} 5 \\ 2 \\ -1 \end{bmatrix} \quad (5.188)$$

The corresponding **projection/reconstruction error** is $\|\mathbf{x} - \mathbf{p}\| = \left\| \begin{bmatrix} 1 & -2 & 1 \end{bmatrix}^\top \right\| = \sqrt{6}$

Fifth, the general projection matrix (for and $\mathbf{x} \in \mathbb{R}^3$) is given by

$$\mathbf{P}_\pi = \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top = \frac{1}{6} \begin{bmatrix} 5 & 2 & 1 \\ 2 & 2 & 2 \\ -1 & 2 & 5 \end{bmatrix} \quad (5.189)$$

To verify the results, we can (a) check whether the error vector $\mathbf{p} - \mathbf{x}$ is orthogonal to all basis vectors of U , (b) verify that $\mathbf{P}_\pi = \mathbf{P}_\pi^2$ (see Definition 26).

Remark 44

In vector spaces with non-standard scalar products, we have to pay attention when computing angles and distances, which are defined by means of the scalar product.

5.14.3 Applications

Projections are often used in computer graphics, e.g., to generate shadows, see Figure 5.7. In optimization, orthogonal projections are often used to (iteratively) minimize residual errors. This also has applications in machine learning, e.g., in linear

regression where we want to find a (linear) function that minimizes the residual errors, i.e., the lengths of the orthogonal projections of the data onto the line (Bishop, 2006). PCA (Hotelling, 1933) also uses projections to reduce the dimensionality of high-dimensional data: First, PCA determines an orthogonal basis of the data space. It turns out that this basis is the eigenbasis of the data matrix. The importance of each individual dimension of the data is proportional to the corresponding eigenvalue. Finally, we can select the eigenvectors corresponding to the largest eigenvalues to reduce the dimensionality of the data, and this selection results in the minimal residual error of the data points projected onto the subspace spanned by these **principal components**. Figure 5.6 illustrates the projection onto the first principal component for a two-dimensional data set.

5.15 Affine Subspaces

In the following, we will have a closer look at geometric properties of vector spaces. For this purpose, we define an affine space.

Definition 27 (Affine Subspace)

Let V be a vector space, $\mathbf{x}_0 \in V$ and $U \subset V$ a subspace. Then the subset

$$L = \mathbf{x}_0 + U := \{\mathbf{x}_0 + \mathbf{u} : \mathbf{u} \in U\} = \{\mathbf{v} \in V \mid \exists \mathbf{u} \in U : \mathbf{v} = \mathbf{x}_0 + \mathbf{u}\} \subset V \quad (5.190)$$

is called **affine subspace** of V .⁴⁹ U is called **direction** or **direction space**, and \mathbf{x}_0 is called **support point**.

Note that the definition of an affine subspace excludes $\mathbf{0}$ if the support point $\mathbf{x}_0 \notin U$. Therefore, an affine subspace is not a (linear) subspace (vector subspace) of V for $\mathbf{x}_0 \notin U$.

Examples of affine subspaces are points, lines and planes in \mathbb{R}^2 , which do not (necessarily) go through the origin.

Remark 45

- Consider two affine subspaces $L = \mathbf{x}_0 + U$ and $\tilde{L} = \tilde{\mathbf{x}}_0 + \tilde{U}$ of a vector space V . $L \subset \tilde{L}$ if and only if $U \subset \tilde{U}$ and $\mathbf{x}_0 - \tilde{\mathbf{x}}_0 \in \tilde{U}$.
- Affine subspaces are often described by **parameters**: Consider a k -dimensional affine space $L = \mathbf{x}_0 + U$ of V . If $(\mathbf{b}_1, \dots, \mathbf{b}_k)$ is an (ordered) basis of U , then every element $\mathbf{x} \in L$ can be (uniquely) described as

$$\mathbf{x} = \mathbf{x}_0 + \lambda_1 \mathbf{b}_1 + \dots + \lambda_k \mathbf{b}_k, \quad (5.191)$$

where $\lambda_1, \dots, \lambda_k \in \mathbb{R}$.

This representation is called **parametric equation** of L with directional vectors $\mathbf{b}_1, \dots, \mathbf{b}_k$ and **parameters** $\lambda_1, \dots, \lambda_k$.

Example

⁴⁹ L is also called **linear manifold**.

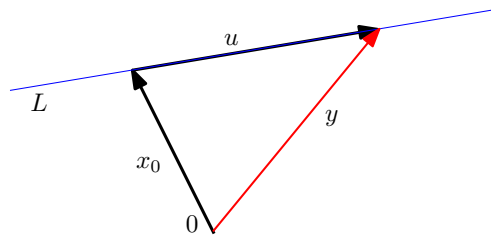


Figure 5.9: Points y on a line lie in an affine subspace L with support point x_0 and direction u .

- One-dimensional affine subspaces are called **lines** and can be written as $y = x_0 + \lambda x_1$, where $\lambda \in \mathbb{R}$, where $U = [x_1] \subset \mathbb{R}^n$ is a one-dimensional subspace of \mathbb{R}^n .
- Two-dimensional affine subspaces of \mathbb{R}^n are called **planes**. The parametric equation for planes is $y = x_0 + \lambda_1 x_1 + \lambda_2 x_2$, where $\lambda_1, \lambda_2 \in \mathbb{R}$ and $[x_1, x_2]$ is a basis of $U \subset \mathbb{R}^n$.
- In an n -dimensional vector space V , the $(n - 1)$ -dimensional affine subspaces are called **hyperplanes**.
- For $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$ the solution of the linear equation system $Ax = b$ is either the empty set or an affine subspace of \mathbb{R}^n of dimension $n - \text{rk}(A)$. In particular, the solution of the linear equation $\lambda_1 x_1 + \dots + \lambda_n x_n = b$, where $(\lambda_1, \dots, \lambda_n) \neq (0, \dots, 0)$, is a hyperplane in \mathbb{R}^n .

- In \mathbb{R}^n every k -dimensional affine subspace is the solution of a linear inhomogeneous equation system $Ax = b$, where $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$ and $\text{rk}(A) = n - k$.⁵⁰

•

5.15.1 Intersection of Affine Subspaces

In the following, we consider the problem of determining the intersection $L_1 \cap L_2$ of two affine subspaces $L_1 = x_1 + U_1$ and $L_2 = x_2 + U_2$ where $U_1, U_2 \subset V$ are subspaces with ordered bases (b_1, \dots, b_k) and (c_1, \dots, c_m) , respectively.

We follow the same approach that we used in the discussion of the intersection of vector subspace (Section 5.8):

$$x \in L_1 \Leftrightarrow \exists \lambda_1, \dots, \lambda_k \in \mathbb{R} : x = x_1 + \sum_{i=1}^k \lambda_i b_i \quad (5.192)$$

⁵⁰Recall that for homogeneous equation systems $Ax = 0$ the solution was a vector subspace (not affine).

$$\mathbf{x} \in L_2 \Leftrightarrow \exists \psi_1, \dots, \psi_m \in \mathbb{R} : \mathbf{x} = \mathbf{x}_2 + \sum_{j=1}^k \psi_j \mathbf{c}_j \quad (5.193)$$

We exploited that since $\mathbf{b}_1, \dots, \mathbf{b}_k$ are a basis of U_1 , every $\mathbf{x} \in L_1$ can (uniquely) be represented as a linear combination of the basis vectors plus the support point \mathbf{x}_1 . The same reasoning applies to $\mathbf{x} \in L_2$ with the basis $\mathbf{c}_1, \dots, \mathbf{c}_m$ of U_2 .

$$\forall \mathbf{x} \in L_1 \cap L_2 : \mathbf{x} \in L_1 \wedge \mathbf{x} \in L_2 \quad (5.194)$$

$$\Leftrightarrow \exists \lambda_1, \dots, \lambda_k, \psi_1, \dots, \psi_m \in \mathbb{R} : \mathbf{x}_1 + \sum_{i=1}^k \lambda_i \mathbf{b}_i = \mathbf{x} = \mathbf{x}_2 + \sum_{j=1}^m \psi_j \mathbf{c}_j \quad (5.195)$$

$$\Leftrightarrow \exists \lambda_1, \dots, \lambda_k, \psi_1, \dots, \psi_m \in \mathbb{R} : \sum_{i=1}^k \lambda_i \mathbf{b}_i - \sum_{j=1}^m \psi_j \mathbf{c}_j = \mathbf{x}_2 - \mathbf{x}_1 \quad (5.196)$$

$$\Leftrightarrow \exists \lambda_1, \dots, \lambda_k, \psi_1, \dots, \psi_m \in \mathbb{R} : \underbrace{\begin{bmatrix} \mathbf{b}_1 & \dots & \mathbf{b}_k & -\mathbf{c}_1 & \dots & -\mathbf{c}_m \end{bmatrix}}_{=:A} \underbrace{\begin{bmatrix} \lambda_1 \\ \vdots \\ \lambda_k \\ \psi_1 \\ \vdots \\ \psi_m \end{bmatrix}}_{=:y} = \underbrace{\mathbf{x}_2 - \mathbf{x}_1}_{=:z} \quad (5.197)$$

Because of this, finding the intersection of two affine spaces reduces to solving an inhomogeneous linear equation system $A\mathbf{y} = \mathbf{z}$ to determine either $\lambda_1, \dots, \lambda_k$ or ψ_1, \dots, ψ_m . When we have these coefficients, $L_1 \cap L_2$ can be described using (5.192) or (5.193).

Example

Consider two linear spaces $L_1, L_2 \subset \mathbb{R}^4$:

$$L_1 = \underbrace{\begin{bmatrix} 2 \\ 0 \\ 0 \\ 1 \end{bmatrix}}_{\mathbf{x}_1} + \underbrace{\left[\begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \end{bmatrix} \right]}_{U_1}, \quad L_2 = \underbrace{\begin{bmatrix} 3 \\ 1 \\ 0 \\ 0 \end{bmatrix}}_{\mathbf{x}_2} + \underbrace{\left[\begin{bmatrix} -1 \\ 1 \\ 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} \right]}_{U_2} \quad (5.198)$$

We want to determine the intersection $L_1 \cap L_2$. We know that holds that $\mathbf{x} \in L_1 \cap L_2 \Leftrightarrow \mathbf{x} \in L_1 \wedge \mathbf{x} \in L_2 \Leftrightarrow \exists \lambda_1, \lambda_2, \lambda_3, \psi_1, \psi_2 \in \mathbb{R}$, such that

$$\underbrace{\begin{bmatrix} 2 \\ 0 \\ 0 \\ 1 \end{bmatrix} + \lambda_1 \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} + \lambda_2 \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix} + \lambda_3 \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \end{bmatrix}}_{\in L_1} = \mathbf{x} = \underbrace{\begin{bmatrix} 3 \\ 1 \\ 0 \\ 0 \end{bmatrix} + \psi_1 \begin{bmatrix} -1 \\ 1 \\ 2 \\ 0 \end{bmatrix} + \psi_2 \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}}_{\in L_2}. \quad (5.199)$$

From this, we obtain the following inhomogeneous equation system with unknowns $\lambda_1, \lambda_2, \lambda_3, \psi_1, \psi_2$:

$$\left[\begin{array}{ccccc|c} 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & -1 & -1 & 1 \\ 0 & 1 & 1 & -2 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & -1 \end{array} \right]$$

Using Gaussian elimination, we quickly obtain the reduced row echelon form

$$\left[\begin{array}{ccccc|c} 1 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & -2 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{array} \right]$$

We exploit the reduced row echelon form to obtain the general solution

$$\begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \\ \psi_1 \\ \psi_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ -1 \\ 0 \\ 1 \end{bmatrix} + \alpha \begin{bmatrix} 1 \\ -2 \\ 0 \\ -1 \\ 0 \end{bmatrix}, \quad \alpha \in \mathbb{R}. \quad (5.200)$$

Note that we have only determined the general solution for the coordinates of the linear combinations that describe the intersection $L_1 \cap L_2$, see (5.199). To obtain the intersection, we use $\lambda_1, \lambda_2, \lambda_3$ from (5.200) in the left-hand-side of (5.199) or, alternatively, ψ_1, ψ_2 in the left-hand-side of (5.199) to obtain the intersection.

Let us do both—we should obtain the same result.

- Let us start with $\lambda_1, \lambda_2, \lambda_3$ to obtain $L_1 \cap L_2$ via L_1 . From (5.199), we know that $\mathbf{x} \in L_1 \cap L_2$ can be written as $\mathbf{x}_1 + \sum_{i=1}^3 \lambda_i \mathbf{b}_i$ where the \mathbf{b}_i are the basis vectors of U_1 . With the λ_i in (5.200), we, therefore, obtain

$$L_1 \cap L_2 = \begin{bmatrix} 2 \\ 0 \\ 0 \\ 1 \end{bmatrix} + (1 + \alpha) \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} + (1 - 2\alpha) \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix} - \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \end{bmatrix} \quad (5.201)$$

$$= \begin{bmatrix} 3 \\ 2 \\ 0 \\ 0 \end{bmatrix} + \alpha \begin{bmatrix} 1 \\ -1 \\ -2 \\ 0 \end{bmatrix} = \begin{bmatrix} 3 \\ 2 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} -1 \\ 1 \\ 2 \\ 0 \end{bmatrix}. \quad (5.202)$$

- Now, we consider ψ_1, ψ_2 to obtain $L_1 \cap L_2$ via L_2 . From (5.200), we obtain $\psi_1 = \alpha \in \mathbb{R}$ is a free parameter, and $\psi_2 = 1$. Plugging this into (5.199), we obtain that any $\mathbf{x} \in L_1 \cap L_2$ satisfies

$$L_1 \cap L_2 = \begin{bmatrix} 3 \\ 1 \\ 0 \\ 0 \end{bmatrix} + \psi_1 \begin{bmatrix} -1 \\ 1 \\ 2 \\ 0 \end{bmatrix} + 1 \cdot \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 3 \\ 2 \\ 0 \\ 0 \end{bmatrix} + \psi_1 \begin{bmatrix} -1 \\ 1 \\ 2 \\ 0 \end{bmatrix} = \begin{bmatrix} 3 \\ 2 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} -1 \\ 1 \\ 2 \\ 0 \end{bmatrix}, \quad (5.203)$$

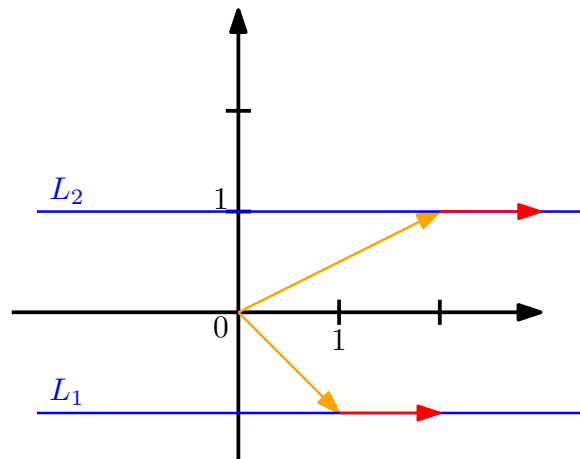


Figure 5.10: Parallel lines. The affine subspaces L_1 and L_2 are parallel with $L_1 \cap L_2 = \emptyset$.

We see that the solutions in (5.202) and (5.203) are identical. Overall, $L_1 \cap L_2$ is a line (1D affine subspace) in \mathbb{R}^4 .

Remark 46

The intersection of a finite set of hyperplanes in \mathbb{R}^n is either empty or an affine subspace. Moreover, every k -dimensional affine subspace of \mathbb{R}^n is the intersection of finitely many hyperplanes.

Definition 28 (Parallelism)

Two affine subspaces L_1 and L_2 are **parallel** ($L_1 \parallel L_2$) if the following holds for the corresponding direction spaces U_1, U_2 : $U_1 \subset U_2$ or $U_2 \subset U_1$.⁵¹

Remark 47

Parallel affine subspaces, which do not contain each other (such that $L_1 \subset L_2$ or $L_2 \subset L_1$), have no points in common, i.e., $L_1 \cap L_2 = \emptyset$.

Example (Parallel affine subspaces)

The affine subspaces

$$L_1 = \begin{bmatrix} 1 \\ -1 \end{bmatrix} + \left[\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right], \quad L_2 = \begin{bmatrix} 2 \\ 1 \end{bmatrix} + \left[\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right] \quad (5.204)$$

are parallel ($U_1 = U_2$ in this case) and $L_1 \cap L_2 = \emptyset$ because the lines are offset as illustrated in Figure 5.10.

We talk about **skew lines** if they are neither parallel nor have an intersection point. Imagine non-parallel lines in 3D that “miss” each other.

⁵¹Note that this definition of parallel allows for $L_1 \subset L_2$ or $L_2 \subset L_1$.

Example

1. Consider two lines $g = \mathbf{x}_1 + U_1, h = \mathbf{x}_2 + U_2$ in \mathbb{R}^2 , where $U_1, U_2 \subset \mathbb{R}$ are one-dimensional subspaces.
 - (a) If $g \cap h \neq \emptyset$:
 - For $\dim(U_1 \cap U_2) = 0$, we get a single point as the intersection.
 - For $\dim(U_1 \cap U_2) = 1$ it follows that $U_1 = U_2$ and $g = h$.
 - (b) If $g \cap h = \emptyset$:
 - For $\dim(U_1 \cap U_2) = 1$ it follows again that $U_1 = U_2$, and $g \parallel h$, $g \neq h$ (because they do not intersect).
 - The case $\dim(U_1 \cap U_2) = 0$ cannot happen in \mathbb{R}^2 .
2. Consider two lines $g = \mathbf{x}_1 + U_1, h = \mathbf{x}_2 + U_2$ in $\mathbb{R}^n, n \geq 3$
 - (a) For $g \cap h \neq \emptyset$ we obtain (as in \mathbb{R}^2) that either g and h intersect in a single point or $g = h$.
 - (b) If $g \cap h = \emptyset$:
 - For $\dim(U_1 \cap U_2) = 0$, we obtain that g and h are skew lines (this cannot happen in \mathbb{R}^2). This means, there exists no plane that contains both g and h .
 - If $\dim(U_1 \cap U_2) = 1$ it follows that $g \parallel h$.
3. Consider two hyper-planes $L_1 = \mathbf{x}_1 + U_1$ and $L_2 = \mathbf{x}_2 + U_2$ in $\mathbb{R}^n, n = 3$, where $U_1, U_2 \subset \mathbb{R}^2$ are two-dimensional subspaces.
 - (a) If $L_1 \cap L_2 \neq \emptyset$ the intersection is an affine subspace. The kind of affine subspace depends on the dimension $\dim(U_1 \cap U_2)$ of the intersection of the corresponding direction spaces.
 - $\dim(U_1 \cap U_2) = 2$: Then $U_1 = U_2$ and, therefore, $L_1 = L_2$.
 - $\dim(U_1 \cap U_2) = 1$: The intersection is a line.
 - $\dim(U_1 \cap U_2) = 0$: Cannot happen in \mathbb{R}^3 .
 - (b) If $L_1 \cap L_2 = \emptyset$, then $\dim(U_1 \cap U_2) = 2$ (no other option possible in \mathbb{R}^3) and $L_1 \parallel L_2$.
4. Consider two planes $L_1 = \mathbf{x}_1 + U_1$ and $L_2 = \mathbf{x}_2 + U_2$ in $\mathbb{R}^n, n = 4$, where $U_1, U_2 \subset \mathbb{R}^2$ are two-dimensional subspaces.
 - (a) For $L_1 \cap L_2 \neq \emptyset$ the additional case is possible that the planes intersect in a point.
 - (b) For $L_1 \cap L_2 = \emptyset$ the additional case is possible that $\dim(U_1 \cap U_2) = 1$. This means that the planes are not parallel, they have no point in common, but there is a line g such that $g \parallel L_1$ and $g \parallel L_2$.
5. For two planes $L_1 = \mathbf{x}_1 + U_1$ and $L_2 = \mathbf{x}_2 + U_2$ in $\mathbb{R}^n, n > 4$, where $U_1, U_2 \subset \mathbb{R}^2$ are two-dimensional subspaces, all kinds of intersections are possible.




Figure 5.11: The robotic arm needs to rotate its joints in order to pick up objects or to place them correctly. Figure taken from (Deisenroth et al., 2015).

5.16 Affine Mappings

Definition 29 (Affine mapping)

For two \mathbb{R} -vector spaces V, W and a linear mapping $\Phi : V \rightarrow W$ and $\mathbf{a} \in W$ the mapping

$$\phi : V \rightarrow W \quad (5.205)$$

$$x \mapsto \mathbf{a} + \Phi(x) \quad (5.206)$$

is an **affine mapping** from V to W . The vector \mathbf{a} is called **translation vector** of ϕ .

- Every affine mapping $\phi : V \rightarrow W$ is also the composition of a linear mapping $\Phi : V \rightarrow W$ and a translation $\tau : W \rightarrow W$ in W , such that $\phi = \tau \circ \Phi$. The mappings Φ and τ are uniquely determined.
- The composition $\phi' \circ \phi$ of affine mappings $\phi : V \rightarrow W, \phi' : W \rightarrow X$ is affine.
- Affine mappings keep the geometric structure invariant, and preserve the dimension and parallelism.

Theorem 11

Let V, W be finite-dimensional \mathbb{R} -vector spaces and $\phi : V \rightarrow W$ an affine mapping. Then:

- If L is an affine subspace of V then $\phi(L)$ is an affine subspace of W and $\dim(\phi(L)) \leq \dim(W)$.
- ϕ preserves parallelism, i.e., for all affine subspaces $L_1, L_2 \subset V$ it follows that if $L_1 \parallel L_2$ in V then $\phi(L_1) \parallel \phi(L_2)$ in W .

5.17 Rotations

An important category of linear mappings⁵² are rotations. A rotation is acting to rotate an object (counterclockwise) by an angle θ about the origin. Important application areas of rotations include computer graphics and robotics. For example, in

⁵²more precisely: endomorphisms

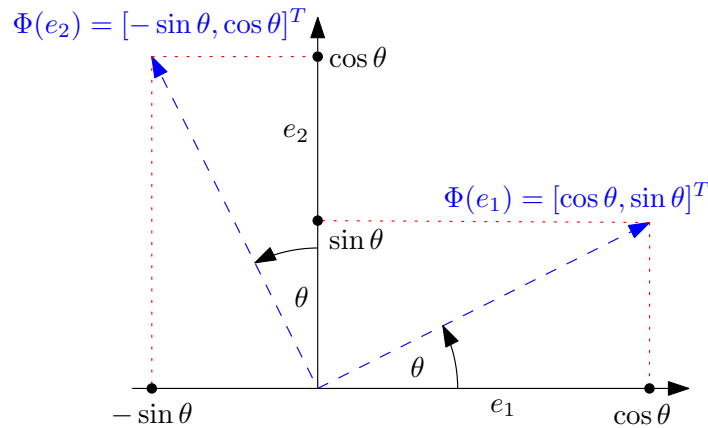


Figure 5.12: Rotation of the standard basis in \mathbb{R}^2 by an angle θ .

robotics, it is often important to know how to rotate the joints of a robotic arm in order to pick up or place an object, see Figure 5.11.

5.17.1 Rotations in the Plane

Consider the standard basis in \mathbb{R}^2 $\left\{e_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, e_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}\right\}$, which defines the standard coordinate system in \mathbb{R}^2 . Assume we want to rotate this coordinate system by an angle θ as illustrated in Figure 5.12. Note that the rotated vectors are still linearly independent and, therefore, are a basis of \mathbb{R}^2 . This means that the rotation performs a basis change.

Since rotations Φ are linear mappings, we can express them by a **rotation matrix** $R(\theta)$. Trigonometry allows us to determine the coordinates of the rotated axes with respect to the standard basis in \mathbb{R}^2 . We obtain

$$\Phi(e_1) = \begin{bmatrix} \cos \theta \\ \sin \theta \end{bmatrix}, \quad \Phi(e_2) = \begin{bmatrix} -\sin \theta \\ \cos \theta \end{bmatrix}. \quad (5.207)$$

Therefore, the rotation matrix that performs the basis change into the rotated coordinates $R(\theta)$ is given as

$$R(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \quad (5.208)$$

Generally, in \mathbb{R}^2 rotations of any vector $x \in \mathbb{R}^2$ by an angle θ are given by

$$x' = \Phi(x) = R(\theta)x = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} x_1 \cos \theta - x_2 \sin \theta \\ x_1 \sin \theta + x_2 \cos \theta \end{bmatrix}. \quad (5.209)$$

5.17.2 Rotations in Three Dimensions

In three dimensions, we have to define what “counterclockwise” means: We will go by the convention that a “counterclockwise” (planar) rotation about an axis refers

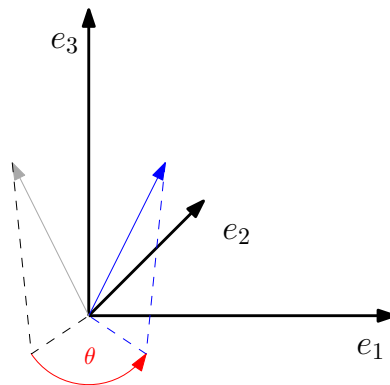


Figure 5.13: Rotation of a general vector (gray) in \mathbb{R}^3 by an angle θ about the e_3 -axis. The rotated vector is shown in blue.

to a rotation about an axis when we look at the axis “from the end toward the origin”. In \mathbb{R}^3 there are three (planar) rotations about three standard basis vectors (see Figure 5.13):

- Counterclockwise rotation about the e_1 -axis

$$R_1(\theta) = [\Phi(e_1) \quad \Phi(e_2) \quad \Phi(e_3)] = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{bmatrix} \quad (5.210)$$

Here, the e_1 coordinate is fixed, and the counterclockwise rotation is performed in the e_2e_3 plane.

- Counterclockwise rotation about the e_2 -axis

$$R_2(\theta) = \begin{bmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{bmatrix} \quad (5.211)$$

If we rotate the e_1e_3 plane about the e_2 axis, we need to look at the e_2 axis from its “tip” toward the origin.

- Counterclockwise rotation about the e_3 -axis

$$R_3(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (5.212)$$

Figure 5.13 to illustrates this.

5.17.3 Rotations in n Dimensions

Definition 30 (Givens Rotation)

Let V be an n -dimensional Euclidean vector space and $\Phi : V \rightarrow V$ an automorphism with transformation matrix

$$\mathbf{R}_{ij}(\theta) := \begin{bmatrix} \mathbf{I}_{i-1} & \mathbf{0} & \cdots & \cdots & \mathbf{0} \\ \mathbf{0} & \cos \theta & \mathbf{0} & -\sin \theta & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_{j-i} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \sin \theta & \mathbf{0} & \cos \theta & \mathbf{0} \\ \mathbf{0} & \cdots & \cdots & \mathbf{0} & \mathbf{I}_{n-j} \end{bmatrix} \in \mathbb{R}^{n \times n}, \quad \theta \in \mathbb{R}, \quad (5.213)$$

for $1 \leq i < j \leq n$. Then $\mathbf{R}_{ij}(\theta)$ is called a **Givens rotation**. Essentially, $\mathbf{R}_{ij}(\theta)$ is the identity matrix \mathbf{I}_n with

$$r_{ii} = \cos \theta, \quad r_{ij} = -\sin \theta, \quad r_{ji} = \sin \theta, \quad r_{jj} = \cos \theta. \quad (5.214)$$

In two dimensions (i.e., $n = 2$), we obtain (5.208) as a special case.

5.17.4 Properties of Rotations

Rotations exhibit some useful properties:

- The composition of rotations can be expressed as a single rotation where the angles are summed up:

$$\mathbf{R}(\phi)\mathbf{R}(\theta) = \mathbf{R}(\phi + \theta) \quad (5.215)$$

- Rotations preserve lengths, i.e., $\|\mathbf{x}\| = \|\mathbf{R}(\phi)\mathbf{x}\|$. The original vector and the rotated vector have the same length. This also implies that $\det(\mathbf{R}(\phi)) = 1$
- Rotations preserve distances, i.e., $\|\mathbf{x} - \mathbf{y}\| = \|\mathbf{R}_\theta(\mathbf{x}) - \mathbf{R}_\theta(\mathbf{y})\|$. In other words, rotations leave the distance between any two points unchanged after the transformation
- Rotations in three (or more) dimensions are generally not commutative. Therefore, the order in which rotations are applied is important, even if they rotate about the same point. Only in two dimensions vector rotations are commutative, such that $\mathbf{R}(\phi)\mathbf{R}(\theta) = \mathbf{R}(\theta)\mathbf{R}(\phi)$ for all $\phi, \theta \in [0, 2\pi)$, and form an Abelian group (with multiplication) if they rotate about the same point (e.g., the origin).
- Rotations have no real eigenvalues, except we rotate about $n\pi$ where $n \in \mathbb{Z}$.

Bibliography

- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Information Science and Statistics. Springer-Verlag. pages 98
- Deisenroth, M. P., Fox, D., and Rasmussen, C. E. (2015). Gaussian Processes for Data-Efficient Learning in Robotics and Control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):408–423. pages 104
- Deng, L., Seltzer, M. L., Yu, D., Acero, A., Mohamed, A.-r., and Hinton, G. E. (2010). Binary Coding of Speech Spectrograms using a Deep Auto-Encoder. In *Interspeech*, pages 1692–1695. pages 92
- Drumm, V. and Weil, W. (2001). *Lineare Algebra und Analytische Geometrie*. Lecture Notes, Universität Karlsruhe. pages 36
- Hogben, L., editor (2013). *Handbook of Linear Algebra*. Discrete Mathematics and Its Applications. Chapman and Hall, 2nd edition. pages 36
- Hotelling, H. (1933). Analysis of a Complex of Statistical Variables into Principal Components. *Journal of Educational Psychology*, 24:417–441. pages 83, 92, 98
- Liesen, J. and Mehrmann, V. (2015). *Linear Algebra*. Springer Undergraduate Mathematics Series. Springer. pages 36, 84
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. The MIT Press, Cambridge, MA, USA. pages 92
- Schölkopf, B. and Smola, A. J. (2002). *Learning with Kernels—Support Vector Machines, Regularization, Optimization, and Beyond*. Adaptive Computation and Machine Learning. The MIT Press, Cambridge, MA, USA. pages 92
- Schölkopf, B., Smola, A. J., and Müller, K.-R. (1998). Nonlinear Component Analysis as a Kernel Eigenvalue Problem. *Neural Computation*, 10(5):1299–1319. pages 92
- Stoer, J. and Burlirsch, R. (2002). *Introduction to Numerical Analysis*. Springer. pages 91
- Strang, G. (2003). *Introduction to Linear Algebra*. Wellesley-Cambridge Press, 3rd edition. pages 36

Wikipedia (2015). Linear Dependence. https://en.wikipedia.org/wiki/Linear_independence.
pages 56