


Hans Bruun Nielsen

LINEAR ALGEBRA
for
IT & HEALTH



September 2010

DTU Informatics – IMM
Richard Petersens Plads
DTU – building 321
DK-2800 Kgs. Lyngby

Phone: +45 4525 3351
www.imm.dtu.dk

Preface

This lecture note has been written for the teaching in the course *Modelling and Programming*. The aim has been to give an easy, but sufficient introduction to the subject of linear algebra, leaving out the parts that are not strictly necessary.

I wish to thank Jeppe Revall Frisvad for his careful reading of the note and many useful suggestions. Also, I appreciate the fruitful discussions with Peter Dalgaard and Rasmus Reinhold Paulsen.

Hans Bruun Nielsen

The front page illustration is discussed in Example 1.14 on page 8.

Contents

Preface	iii
1. Introduction	1
1.1. Vectors and matrices	1
1.2. Basic vector arithmetic	3
1.3. Basic matrix arithmetic	5
1.4. Linear mapping	8
1.5. Basis and coordinates	10
2. Linear Systems of Equations	15
2.1. Simple systems	17
2.2. Gaussian elimination	18
2.3. Complete solution	22
2.4. Matrix inverse	25
2.5. LU factorization	28
2.6. SPD systems	30
3. Eigenvalues and Eigenvectors	35
3.1. Properties of eigensolutions	35
3.2. Applications of eigensolutions	39
Differential equations	39
Simplification of quadratics	42
Google's PageRank	44
4. Linear Least Squares Problems	47
4.1. Overdetermined systems	48
4.2. QR factorization	52
4.3. Singular value decomposition	53
Literature	57
Notation	58
List of Danish words	59
Index	60

1. Introduction

1.1. Vectors and matrices

A *vector* $\mathbf{x} \in \mathbb{R}^n$ (also called an n -vector) is a list of n elements, x_1, x_2, \dots, x_n , all of which are real numbers: $x_i \in \mathbb{R}$.

Example 1.1. In the usual Cartesian system a *geometric vector* \vec{v} can be identified by the coordinates (v_1, v_2) , cf the figure. This pair of coordinates is a vector in \mathbb{R}^2 .




Figure 1.1. Geometric vector.

Conversely, one may visualize a vector $\mathbf{x} \in \mathbb{R}^2$ by a geometric vector \vec{x} in the plane, having the coordinates (x_1, x_2) .

It may help intuition to generalize this idea: One may think of an n -vector as an “arrow” in the n -dimensional space. ■

Example 1.2. During the development of a new drug you need to perform clinical experiments to investigate how fast the drug is excreted from the human body. This investigation may consist in measuring the drug concentration in the kidneys at certain times after taking the drug. The results may be represented by two vectors, \mathbf{t} and \mathbf{y} , where y_i is the concentration at time t_i .




Figure 1.2. Concentration y as function of time t .

The figure is a MATLAB plot of the results of such an experiment. Time is given in hours after the drug was taken. It seems that there was an 8 hour break during the experiment.

We shall return to this problem several times, first in Example 1.5. In Chapter 3 we discuss a mathematical model for the behaviour shown in Figure 1.2. ■

The *zero vector* $\mathbf{0}$ has all its elements equal to zero. We say that \mathbf{x} is nonzero if it has at least one element different from zero.

A *matrix* $\mathbf{A} \in \mathbb{R}^{m \times n}$ (also called an $m \times n$ matrix) is a list of $m \cdot n$ real numbers, organized in m rows, each of which has n elements. Alternatively we can think of \mathbf{A} as organized in n columns, with m elements in each. The element in position (i, j) (ie the j^{th} element in the i^{th} row) is denoted a_{ij} or $(\mathbf{A})_{ij}$. The *diagonal* of \mathbf{A} consists of the elements a_{ii} , $i = 1, 2, \dots, \min\{m, n\}$.

Example 1.3. Consider the 2×3 matrix

$$\mathbf{A} = \begin{pmatrix} 1.2 & 3.4 & 5.6 \\ 7.8 & 9.1 & -2.3 \end{pmatrix}.$$

We have, for instance, $a_{12} = 3.4$ and $a_{23} = -2.3$. The diagonal elements are $a_{11} = 1.2$ and $a_{22} = 9.1$. ■

Note that $\mathbb{R}^{1 \times n}$ and $\mathbb{R}^{m \times 1}$ contain matrices with only one row and one column, respectively. Such degenerated matrices are equivalent to vectors, and we talk about a *row vector* if $\mathbf{x} \in \mathbb{R}^{1 \times n}$ and a *column vector* if $\mathbf{y} \in \mathbb{R}^{m \times 1}$. If nothing else is specified, $\mathbf{x} \in \mathbb{R}^n$ is considered to be a column vector.

The elements in the entire i^{th} row of the the matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ is a row vector, which we denote $\mathbf{A}_{i,:}$. Similarly, $\mathbf{A}_{:,j} \in \mathbb{R}^{m \times 1}$ consists of the elements in the j^{th} column of \mathbf{A} .

Example 1.4. For the matrix from Example 1.3, we see that

$$\mathbf{A}_{1,:} = (1.2 \quad 3.4 \quad 5.6), \quad \mathbf{A}_{:,3} = \begin{pmatrix} 5.6 \\ -2.3 \end{pmatrix}. \quad \blacksquare$$

Example 1.5. The results from the clinical experiment in Example 1.2 may be stored in a matrix \mathbf{A} , with the times and the measured concentrations in the first and second column, respectively: $\mathbf{A}_{:,1} = \mathbf{t}$, $\mathbf{A}_{:,2} = \mathbf{y}$.

If the experiment involves 4 patients that are measured simultaneously, then the results may be stored in a matrix \mathbf{A} with 5 columns, where $\mathbf{A}_{:,j+1}$ holds the y_i -values for the j^{th} person. ■

A *square matrix* has the same number of rows and columns, $m = n$. A *diagonal matrix* is a square matrix where all the elements outside the diagonal are zero. The *identity matrix* \mathbf{I} is a diagonal matrix with all diagonal elements equal to 1. A *lower triangular matrix* is a square matrix where all the elements above the diagonal are zero. Similarly, all the elements below the diagonal are zero in an *upper triangular matrix*.

Example 1.6. In the case $m = n = 3$ we have

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

As examples of a diagonal matrix and a lower and an upper triangular matrix we give

$$\mathbf{D} = \begin{pmatrix} 2 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 3 \end{pmatrix}, \quad \mathbf{L} = \begin{pmatrix} 2 & 0 & 0 \\ 5 & -1 & 0 \\ 2 & 7 & 3 \end{pmatrix}, \quad \mathbf{U} = \begin{pmatrix} 2 & 3 & -4 \\ 0 & -1 & 7 \\ 0 & 0 & 3 \end{pmatrix}.$$

For the sake of better visual impression we often omit the zeros outside the diagonal in a diagonal matrix and above/below the diagonal in a lower/upper triangular matrix, for instance

$$\mathbf{I} = \begin{pmatrix} 1 & & \\ & 1 & \\ & & 1 \end{pmatrix}, \quad \mathbf{U} = \begin{pmatrix} 2 & 3 & -4 \\ & -1 & 7 \\ & & 3 \end{pmatrix}. \quad \blacksquare$$

Example 1.7. MATLAB is an acronym for MATrix LABoratory. The basic storage element is a matrix. Let \mathbf{A} be a MATLAB variable, then the command `[m,n] = size(A)` returns m , the number of rows in \mathbf{A} , and n , the number of columns in \mathbf{A} .

If $m = 0$ or $n = 0$, then \mathbf{A} is empty.

If $m = n = 1$, then \mathbf{A} is a scalar.

If $m = 1$ and $n > 1$, then \mathbf{A} is a row vector.

If $m > 1$ and $n = 1$, then \mathbf{A} is a column vector.

If $m > 1$ and $n > 1$, then \mathbf{A} is a proper matrix. ■

1.2. Basic vector arithmetic

1° *Transpose.* The transpose of the vector \mathbf{x} is denoted \mathbf{x}^T . If \mathbf{x} is a column vector, then \mathbf{x}^T is the row vector with the same elements, and vice versa.

2° *Multiplication by a scalar.* Let $\mathbf{x} \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}$. Then

$$\mathbf{z} = \alpha \mathbf{x}$$

is the vector with elements $z_i = \alpha x_i$, $i = 1, 2, \dots, n$.

3° *Addition.* Let $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{y} \in \mathbb{R}^m$, and consider the sum

$$\mathbf{z} = \mathbf{x} + \mathbf{y}.$$

This is meaningful only if \mathbf{x} and \mathbf{y} have the same number of elements, ie $m = n$. If one distinguishes between row and column vectors, the two vectors must also be of the same type, ie both of them must be either row vectors or column vectors. When these conditions are satisfied, \mathbf{z} is a vector of the same type with elements $z_i = x_i + y_i$, $i = 1, 2, \dots, n$.

4° *Multiplication of two vectors.* Let $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{y} \in \mathbb{R}^m$ be column vectors. They can be multiplied in three different ways:

4a° *Inner product* (also called *scalar product* or *dot product*): This is defined only if the two vectors have the same number of elements. The inner product of two n -vectors \mathbf{x} and \mathbf{y} is

$$\mathbf{x}^T \mathbf{y} = x_1 y_1 + x_2 y_2 + \dots + x_n y_n. \quad (1.1)$$

Thus, the inner product is a real number, a *scalar*. It follows immediately from the definition that swapping the two vectors does not change their scalar product:

$$\mathbf{y}^T \mathbf{x} = \mathbf{x}^T \mathbf{y} .$$

The vectors \mathbf{x} and \mathbf{y} are said to be *orthogonal* if their scalar product is zero, $\mathbf{x}^T \mathbf{y} = 0$.

4b° The *outer product* of the column vectors \mathbf{x} and \mathbf{y} is written

$$\mathbf{x} \mathbf{y}^T .$$

This is a matrix with the number of rows given by the number of elements in \mathbf{x} , n , and the number of columns is given by the number of elements in \mathbf{y} , m . The elements in the matrix are

$$(\mathbf{x} \mathbf{y}^T)_{ij} = x_i y_j \quad \begin{cases} i = 1, 2, \dots, m \\ j = 1, 2, \dots, n \end{cases} .$$

In general $\mathbf{y} \mathbf{x}^T \neq \mathbf{x} \mathbf{y}^T$. The outer product may also be expressed row-wise or column-wise:

$$\begin{aligned} (\mathbf{x} \mathbf{y}^T)_{i,:} &= x_i \mathbf{y}^T, & i = 1, 2, \dots, n, \\ (\mathbf{x} \mathbf{y}^T)_{:,j} &= y_j \mathbf{x}, & j = 1, 2, \dots, m. \end{aligned}$$

4c° *Element-wise multiplication*. In some cases it is relevant to compute the vector defined by element-wise multiplication of two vectors with the same number of elements, $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. We use the notation

$$\mathbf{z} = \mathbf{x} \otimes \mathbf{y} .$$

This is also a vector with n elements, given by $z_i = x_i y_i$, $i = 1, 2, \dots, n$.

Example 1.8. Let

$$\mathbf{x} = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} 4 \\ -5 \\ 6 \end{pmatrix} .$$

Then

$$2\mathbf{x}^T = (2 \ 4 \ 6), \quad \mathbf{x}^T + \mathbf{y}^T = (5 \ -3 \ 9),$$

$$\mathbf{x}^T \mathbf{y} = 12, \quad \mathbf{y} \mathbf{x}^T = \begin{pmatrix} 4 & 8 & 12 \\ -5 & -10 & -15 \\ 6 & 12 & 18 \end{pmatrix}, \quad \mathbf{x} \otimes \mathbf{y} = \begin{pmatrix} 4 \\ -10 \\ 18 \end{pmatrix} .$$

In MATLAB the inner, outer, and element-wise products of the vectors \mathbf{x} and \mathbf{y} are obtained by the commands

$$\text{ip} = \mathbf{x}' * \mathbf{y}, \quad \text{op} = \mathbf{x} * \mathbf{y}', \quad \text{ep} = \mathbf{x} .* \mathbf{y} \quad \blacksquare$$

5° *Norm*. In many applications we are interested in being able to talk about the “length” of a vector, and we introduce the norm for that purpose. The scalar product of an n -vector \mathbf{x} and itself is

$$\mathbf{x}^T \mathbf{x} = x_1^2 + x_2^2 + \dots + x_n^2 .$$

This is a real, nonnegative number, and we define the *norm* of \mathbf{x} , $\|\mathbf{x}\|$, as

$$\|\mathbf{x}\| = \sqrt{\mathbf{x}^T \mathbf{x}} . \quad (1.2)$$

It is fairly easy to verify that $\|\mathbf{x}\|$ satisfies the following so-called *norm conditions*

$$5a^\circ \quad \|\mathbf{x}\| \geq 0 \quad \text{for all } \mathbf{x} ,$$

$$5b^\circ \quad \|\mathbf{x}\| = 0 \quad \Leftrightarrow \quad \mathbf{x} = \mathbf{0} ,$$

$$5c^\circ \quad \|\alpha \mathbf{x}\| = |\alpha| \cdot \|\mathbf{x}\| \quad \text{for all } \alpha \in \mathbb{R} ,$$

$$5d^\circ \quad \|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\| \quad \text{for all } \mathbf{y} \in \mathbb{R}^n .$$

Example 1.9. Let \vec{x} and \vec{y} be geometric vectors with Cartesian coordinates (x_1, x_2) and (y_1, y_2) . Their *dot product* is

$$\vec{x} \cdot \vec{y} = x_1 y_1 + x_2 y_2 ,$$

and the length of \vec{x} is

$$|\vec{x}| = \sqrt{x_1^2 + x_2^2} .$$

If $\mathbf{x}, \mathbf{y} \in \mathbb{R}^2$ are the column vectors with the coordinates, we see that $\mathbf{x}^T \mathbf{y} = \vec{x} \cdot \vec{y}$ and $\|\mathbf{x}\| = |\vec{x}|$.

The norm $\|\mathbf{x}\|$ generalizes this concept of “length” to vectors with any number of elements.

The condition 5d^o is known as the *triangle inequality*. The reason is that if a triangle in the plane has two sides given by \vec{x} and \vec{y} , then the length of the third side is $|\vec{x} + \vec{y}|$, and this cannot be larger than the sum of the lengths of the two other sides. ■

1.3. Basic matrix arithmetic

1^o *Transpose.* The transpose of the matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ is the matrix $\mathbf{A}^T \in \mathbb{R}^{n \times m}$. It is obtained from \mathbf{A} by interchanging rows and columns: $(\mathbf{A}^T)_{ij} = (\mathbf{A})_{ji}$.

2^o *Multiplication by a scalar.* Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\alpha \in \mathbb{R}$. Then

$$\mathbf{C} = \alpha \mathbf{A}$$

is the matrix with elements $c_{ij} = \alpha a_{ij}$, $i = 1, 2, \dots, m$; $j = 1, 2, \dots, n$.

3^o *Addition of two matrices.* This makes sense only, if the two matrices are of the same type, ie if they have the same number of rows and the same number of columns. If $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{m \times n}$, then

$$\mathbf{C} = \mathbf{A} + \mathbf{B}$$

is the matrix with $c_{ij} = a_{ij} + b_{ij}$, $i = 1, 2, \dots, m$; $j = 1, 2, \dots, n$.

4^o *Matrix-vector multiplication.* Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{x} \in \mathbb{R}^{n \times 1}$, and consider the product

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

The result is defined by the scalar products

$$y_i = \mathbf{A}_{i,:} \mathbf{x} \quad i = 1, 2, \dots, m,$$

so the matrix–vector product is

$$\mathbf{A} \mathbf{x} = \begin{pmatrix} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n \\ \vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n \end{pmatrix}. \quad (1.3)$$

Example 1.10. With the matrix from Example 1.3 we get

$$\mathbf{A} = \begin{pmatrix} 1.2 & 3.4 & 5.6 \\ 7.8 & 9.1 & -2.3 \end{pmatrix}, \quad \mathbf{A}^T = \begin{pmatrix} 1.2 & 7.8 \\ 3.4 & 9.1 \\ 5.6 & -2.3 \end{pmatrix},$$

$$\mathbf{A} \begin{pmatrix} 2 \\ -2 \\ 3 \end{pmatrix} = \begin{pmatrix} 1.2 \cdot 2 + 3.4 \cdot (-2) + 5.6 \cdot 3 \\ 7.8 \cdot 2 + 9.1 \cdot (-2) - 2.3 \cdot 3 \end{pmatrix} = \begin{pmatrix} 12.4 \\ -9.5 \end{pmatrix},$$

$$\mathbf{A}^T \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} 1.2 - 7.8 \\ 3.4 - 9.1 \\ 5.6 + 2.3 \end{pmatrix} = \begin{pmatrix} 6.6 \\ 5.7 \\ -7.9 \end{pmatrix}.$$

If \mathbf{D} is a diagonal matrix, it follows immediately from Eq. (1.3) that the vector $\mathbf{y} = \mathbf{D}\mathbf{x}$ has the elements $y_i = d_{ii}x_i$. In particular, $\mathbf{I}\mathbf{x} = \mathbf{x}$. ■

Example 1.11. Consider the parabola

$$p(t) = c_1 t^2 + c_2 t + c_3,$$

where c_1 , c_2 and c_3 are known coefficients. From basic vector arithmetic it follows that we can express $p(t)$ by

$$p(t) = \begin{pmatrix} t^2 & t & 1 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ c_3 \end{pmatrix}.$$

Now, let $y_i = p(t_i)$, $i = 1, 2, 3$, with $t_i = i$. Then it follows that

$$\begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix} = \begin{pmatrix} t_1^2 & t_1 & 1 \\ t_2^2 & t_2 & 1 \\ t_3^2 & t_3 & 1 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ c_3 \end{pmatrix} = \begin{pmatrix} 1 & 1 & 1 \\ 4 & 2 & 1 \\ 9 & 3 & 1 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ c_3 \end{pmatrix}.$$

The figure illustrates the problem for a specific choice of the coefficients.




Figure 1.3. $p(t) = 0.5t^2 - 1.6t + 1.9$.

We return to this problem in Examples 2.1 and 2.6. ■

5° *Matrix–matrix multiplication.* Let $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{B} \in \mathbb{R}^{n \times q}$, and consider the product

$$\mathbf{C} = \mathbf{A}\mathbf{B}.$$

This is defined by

$$(\mathbf{C})_{ij} = \mathbf{A}_{i,:} \mathbf{B}_{:,j}, \quad \begin{cases} i = 1, 2, \dots, m \\ j = 1, 2, \dots, q \end{cases}. \quad (1.4)$$

In words: the (i, j) th element in \mathbf{C} is the inner product of the i th row in \mathbf{A} and the j th column in \mathbf{B} . The product exists only if these two vectors have the same number of elements, i.e. the number of columns in \mathbf{A} must be equal to the number of rows in \mathbf{B} , $n = p$. The result can be expressed column-wise:

$$\mathbf{C}_{:,j} = \mathbf{A}\mathbf{B}_{:,j}, \quad j = 1, 2, \dots, q, \quad (1.5)$$

i.e. the j th column in \mathbf{C} is the matrix–vector product of \mathbf{A} and the j th column in \mathbf{B} .

We shall sometimes use the following identity,

$$(\mathbf{A}\mathbf{B})^T = \mathbf{B}^T \mathbf{A}^T. \quad (1.6)$$

We already saw this relation in the case where the two matrices degenerate to a row and column matrix, respectively, see Eq. (1.1). For general matrices it is easy to prove the identity by using the definition of the matrix–matrix product.

Example 1.12. If both \mathbf{A} and \mathbf{B} are square and have the same number of elements, then both $\mathbf{A}\mathbf{B}$ and $\mathbf{B}\mathbf{A}$ are defined; for instance

$$\mathbf{A} = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} -5 & 6 \\ 7 & 8 \end{pmatrix}$$

gives

$$\mathbf{A}\mathbf{B} = \begin{pmatrix} 9 & 22 \\ 13 & 50 \end{pmatrix}, \quad \mathbf{B}\mathbf{A} = \begin{pmatrix} 13 & 14 \\ 31 & 46 \end{pmatrix}.$$

This illustrates, in general $\mathbf{A}\mathbf{B} \neq \mathbf{B}\mathbf{A}$. To illustrate Eq. (1.6) we see that

$$\mathbf{A}^T \mathbf{B}^T = \begin{pmatrix} 1 & 3 \\ 2 & 7 \end{pmatrix} \begin{pmatrix} -5 & 7 \\ 6 & 8 \end{pmatrix} = \begin{pmatrix} 13 & 31 \\ 14 & 46 \end{pmatrix},$$

which is equal to $(\mathbf{B}\mathbf{A})^T$.

With the rectangular matrix from Example 1.3,

$$\mathbf{M} = \begin{pmatrix} 1.2 & 3.4 & 5.6 \\ 7.8 & 9.1 & -2.3 \end{pmatrix},$$

the matrix products $\mathbf{A}\mathbf{M}$ and $\mathbf{M}^T \mathbf{A}^T$ are defined, while $\mathbf{M}\mathbf{A}$ is not defined.

$$\mathbf{A}\mathbf{M} = \begin{pmatrix} 16.8 & 21.6 & 1 \\ 34.8 & 46.6 & 7.6 \end{pmatrix}, \quad \mathbf{M}^T \mathbf{A}^T = \begin{pmatrix} 16.8 & 34.8 \\ 21.6 & 46.6 \\ 1 & 7.6 \end{pmatrix}.$$

Again the result agrees with Eq. (1.6). ■

Exercise 1.13. Let \mathbf{A} be an $n \times n$ matrix and let \mathbf{D} be an $n \times n$ diagonal matrix. Show that $\mathbf{B} = \mathbf{D}\mathbf{A}$ has the rows

$$\mathbf{B}_{i,:} = d_{ii}\mathbf{A}_{i,:}, \quad i = 1, 2, \dots, n .$$

Hint: Use Eq. (1.5) and Example 1.10.

Next, show that $\mathbf{C} = \mathbf{AD}$ has the columns

$$\mathbf{C}_{:,j} = d_{jj}\mathbf{A}_{:,j}, \quad j = 1, 2, \dots, n .$$

Hint: Use Eq. (1.6) and the result for \mathbf{DA} . ■

1.4. Linear mapping

Let \mathbf{A} be an $m \times n$ matrix and let \mathbf{x} be an n -vector. The m -vector $\mathbf{y} = \mathbf{Ax}$ is said to be obtained by a *linear mapping* from \mathbb{R}^n into \mathbb{R}^m , cf Figure 1.4 below. We say




Figure 1.4. Linear mapping, $\mathbf{y} = \mathbf{Ax}$.

that \mathbf{y} is the *image* of \mathbf{x} and that \mathbf{A} is the *mapping matrix*.

Example 1.14. Linear mappings are widely used in mathematical models. In *CT scanning* or *MR scanning*, for instance, you get a vector \mathbf{b} of measurements, which do not directly make sense. However, you can construct a matrix \mathbf{A} such that for a given distribution of mass density \mathbf{d} in the scanned object, you would measure $\mathbf{y} = \mathbf{Ad}$. The problem is to find the distribution \mathbf{x} such that $\mathbf{Ax} = \mathbf{b}$. On the front page we show such a picture with \mathbf{b} obtained by MR scanning.

The solution of the problem $\mathbf{Ax} = \mathbf{b}$ is the subject of the next chapter. First, however, we need to discuss some properties of linear mappings. ■

Using the definition (1.3) of the matrix–vector product and the rules for vector arithmetic, we see that the relation

$$\mathbf{A}(\alpha\mathbf{u} + \beta\mathbf{v}) = \alpha(\mathbf{Au}) + \beta(\mathbf{Av}) \quad (1.7)$$

is true for all choices of n -vectors \mathbf{u} , \mathbf{v} , and scalars α, β . This relation is known as the *linearity condition*.

In the special case where $\alpha = \beta = 0$, it follows that

$$\mathbf{A}\mathbf{0}_{[n]} = \mathbf{0}_{[m]}, \quad (1.8)$$

where the indices $[n]$ and $[m]$ are used to indicate that the zero vectors are in \mathbb{R}^n and \mathbb{R}^m , respectively. In words: a linear mapping maps the zero vector into the zero vector.

Example 1.15. Let

$$\mathbf{A} = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}, \quad \mathbf{z} = \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}, \quad \mathbf{w} = \begin{pmatrix} 1 \\ -2 \\ 1 \end{pmatrix}.$$

Using Eq. (1.3) we get

$$\mathbf{A}\mathbf{z} = \begin{pmatrix} 8 \\ 20 \\ 32 \end{pmatrix}, \quad \mathbf{A}\mathbf{w} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}. \quad \blacksquare$$

This example shows that for some linear mappings there are nonzero vectors \mathbf{x} for which $\mathbf{A}\mathbf{x} = \mathbf{0}$. In order to get a better understanding of this we look at the definition (1.3) of the *matrix–vector product*,

$$\mathbf{A}\mathbf{x} = \begin{pmatrix} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n \\ \vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n \end{pmatrix}.$$

From the rules for vector arithmetic it follows that an equivalent formulation of the right-hand side is

$$\mathbf{y} = \mathbf{A}\mathbf{x} = x_1\mathbf{A}_{:,1} + x_2\mathbf{A}_{:,2} + \cdots + x_n\mathbf{A}_{:,n}. \quad (1.9)$$

We say that \mathbf{y} is a *linear combination* of the vectors $\mathbf{A}_{:,1}, \mathbf{A}_{:,2}, \dots, \mathbf{A}_{:,n}$.

If $\mathbf{x} = \mathbf{0}$, we naturally get $\mathbf{y} = \mathbf{0}$, and if $\mathbf{y} \neq \mathbf{0}$ for all nonzero \mathbf{x} (ie a vector with at least one nonzero element), then the vectors $\mathbf{A}_{:,1}, \mathbf{A}_{:,2}, \dots, \mathbf{A}_{:,n}$ (the columns in \mathbf{A}) are said to be *linearly independent*. If there exists a nonzero \mathbf{x} such that the image $\mathbf{y} = \mathbf{0}$, then the columns in \mathbf{A} are said to be *linearly dependent*. Assume, for instance, that $\mathbf{A}\mathbf{x} = \mathbf{0}$ for a vector \mathbf{x} with $x_1 \neq 0$. Then Eq. (1.9) implies that

$$\mathbf{A}_{:,1} = -\frac{x_2}{x_1}\mathbf{A}_{:,2} - \cdots - \frac{x_n}{x_1}\mathbf{A}_{:,n}.$$

This means that the first column in \mathbf{A} is a linear combination of the other columns. In general, if the columns in \mathbf{A} are linearly dependent, then one or more of the columns can be expressed as a linear combination of the other columns. The selection of which column to express in terms of the others is normally not unique.

Example 1.16. If, for instance, $\mathbf{A}_{:,k} = \mathbf{0}$, then the columns in \mathbf{A} are linearly dependent: we may take $x_k = 1$ and all other $x_j = 0$.

Two nonzero vectors \mathbf{u} and \mathbf{v} are linearly dependent if and only if they are proportional, ie there is a real number β such that $\mathbf{u} = \beta\mathbf{v}$.

For the matrix from Example 1.15,

$$\mathbf{A} = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix},$$

we can select two columns in three different ways (neglecting the ordering):

$$\begin{pmatrix} 1 & 2 \\ 4 & 5 \\ 7 & 8 \end{pmatrix}, \quad \begin{pmatrix} 1 & 3 \\ 4 & 6 \\ 7 & 9 \end{pmatrix}, \quad \begin{pmatrix} 2 & 3 \\ 5 & 6 \\ 8 & 9 \end{pmatrix}.$$

In all three cases the two columns are linearly independent.

In Example 1.15 we saw that

$$\mathbf{A}_{:,1} - 2\mathbf{A}_{:,2} + \mathbf{A}_{:,3} = \mathbf{0}.$$

This means that

$$\mathbf{A}_{:,1} = 2\mathbf{A}_{:,2} + \mathbf{A}_{:,3}, \quad \mathbf{A}_{:,2} = \frac{1}{2}\mathbf{A}_{:,1} + \frac{1}{2}\mathbf{A}_{:,3}, \quad \mathbf{A}_{:,3} = 2\mathbf{A}_{:,2} - \mathbf{A}_{:,1}.$$

These relations are easily verified. ■

A linear mapping from \mathbb{R}^n into \mathbb{R}^m is sometimes confused with an *affine mapping* from \mathbb{R}^n into \mathbb{R}^m . In this case the image \mathbf{y} is given by

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{b}, \tag{1.10}$$

where the m -vector \mathbf{b} may be zero or nonzero. Obviously, the linearity condition Eq. (1.7) is not satisfied by an affine mapping with a nonzero \mathbf{b} .


1.5. Basis and Coordinates

Example 1.17. Let \vec{b}_1 and \vec{b}_2 be two geometric vectors in the plane. If they are linearly independent, ie they are not parallel, then we can write any \vec{x} in the plane as a linear combination of \vec{b}_1 and \vec{b}_2 ,

$$\vec{x} = \tilde{x}_1 \vec{b}_1 + \tilde{x}_2 \vec{b}_2.$$

The figure illustrates such a linear combination.

Figure 1.5. $\vec{x} = -1.3\vec{b}_1 + 1.5\vec{b}_2$.



The observation in this example generalizes. If we know n linearly independent vectors in \mathbb{R}^n , $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n$, then any n -vector \mathbf{x} can be written as a linear combination of these vectors:

$$\mathbf{x} = \tilde{x}_1 \mathbf{b}_1 + \tilde{x}_2 \mathbf{b}_2 + \dots + \tilde{x}_n \mathbf{b}_n. \tag{1.11}$$

The vectors $\{\mathbf{b}_j\}$ are said to form a *basis* for \mathbb{R}^n , and the $\{\tilde{x}_j\}$ are the *coordinates* of \mathbf{x} with respect to this basis.

The *usual basis* in \mathbb{R}^n is the vectors $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$, where $\mathbf{e}_j = \mathbf{I}_{:,j}$, ie its j^{th} element is 1 and all the other elements are zero. The coordinates of a vector \mathbf{x} with respect to this basis are simply the elements in \mathbf{x} .

Example 1.18. The usual basis in \mathbb{R}^3 is

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

The vector with coordinates 2, 3 and 4 is

$$\mathbf{x} = 2 \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + 3 \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} + 4 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 2 \\ 3 \\ 4 \end{pmatrix}. \quad \blacksquare$$

By comparison with Eq. (1.9) we see that Eq. (1.11) is equivalent to

$$\mathbf{x} = \mathbf{B} \tilde{\mathbf{x}}, \quad (1.12)$$

where $\tilde{\mathbf{x}}$ is the vector of coordinates with respect to the basis $\{\mathbf{b}_j\}$, and the matrix \mathbf{B} has \mathbf{b}_j as its j^{th} column, $\mathbf{B}_{:,j} = \mathbf{b}_j$.

Thus, given the coordinates, the vector $\mathbf{x} \in \mathbb{R}^n$ is obtained by a linear mapping of $\tilde{\mathbf{x}} \in \mathbb{R}^n$ with the mapping matrix \mathbf{B} . The *inverse problem* is: given \mathbf{x} , find $\tilde{\mathbf{x}}$. How to do that in the general case is the subject of the next chapter. In special cases, however, the inverse problem is simple:

If the basis vectors are *orthogonal*, ie

$$\mathbf{b}_i^T \mathbf{b}_j = 0 \quad \text{for all } i \neq j,$$

then it follows from Eq. (1.11) that

$$\begin{aligned} \mathbf{b}_k^T \mathbf{x} &= \tilde{x}_1 (\mathbf{b}_k^T \mathbf{b}_1) + \cdots + \tilde{x}_k (\mathbf{b}_k^T \mathbf{b}_k) + \cdots + \tilde{x}_n (\mathbf{b}_k^T \mathbf{b}_n) \\ &= \tilde{x}_k (\mathbf{b}_k^T \mathbf{b}_k), \end{aligned}$$

showing that

$$\tilde{x}_k = \frac{\mathbf{b}_k^T \mathbf{x}}{\mathbf{b}_k^T \mathbf{b}_k}, \quad k = 1, 2, \dots, n.$$

The computation is even easier if the basis is *orthonormal*. This means that the basis vectors are not only orthogonal, but also each of them has norm $\|\mathbf{b}_j\| = 1$, cf Eq. (1.2),

$$\mathbf{b}_i^T \mathbf{b}_j = \begin{cases} 0 & \text{if } i \neq j, \\ 1 & \text{if } i = j. \end{cases} \quad (1.13)$$

In this case the computation simplifies to

$$\tilde{x}_k = \mathbf{b}_k^T \mathbf{x}, \quad k = 1, 2, \dots, n. \quad (1.14)$$

Example 1.19. It is simple to verify that the usual basis is orthonormal. Another example of an orthonormal basis in \mathbb{R}^2 is

$$\mathbf{b}_1 = \begin{pmatrix} \cos v \\ \sin v \end{pmatrix}, \quad \mathbf{b}_2 = \begin{pmatrix} -\sin v \\ \cos v \end{pmatrix},$$

for any angle v . We easily see that the two vectors are nonzero and not proportional, so they form a basis in \mathbb{R}^2 for any choice of v . Next,

$$\begin{aligned}\mathbf{b}_1^T \mathbf{b}_2 &= \cos v \cdot \sin v - \sin v \cdot \cos v = 0, \\ \mathbf{b}_1^T \mathbf{b}_1 &= \mathbf{b}_2^T \mathbf{b}_2 = \cos^2 v + \sin^2 v = 1.\end{aligned}$$

Thus, the two vectors satisfy Eq. (1.13).

The matrix

$$\mathbf{B} = \begin{pmatrix} \cos v & -\sin v \\ \sin v & \cos v \end{pmatrix}$$

is known as a *rotation matrix*, cf the next example. ■

Example 1.20. We shall demonstrate that a change of basis may simplify a problem: We want to find the points $P: (x_1, x_2)$ that satisfy the quadratic equation

$$73x_1^2 - 72x_1x_2 + 52x_2^2 = 100.$$

The vectors

$$\mathbf{b}_1 = \begin{pmatrix} 0.6 \\ 0.8 \end{pmatrix}, \quad \mathbf{b}_2 = \begin{pmatrix} -0.8 \\ 0.6 \end{pmatrix}$$

form an orthonormal basis in \mathbb{R}^2 , and we introduce the coordinates $(\tilde{x}_1, \tilde{x}_2)$ with respect to this basis:

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 0.6 & -0.8 \\ 0.8 & 0.6 \end{pmatrix} \begin{pmatrix} \tilde{x}_1 \\ \tilde{x}_2 \end{pmatrix}.$$

In accordance with this we use $0.6\tilde{x}_1 - 0.8\tilde{x}_2$ and $0.8\tilde{x}_1 + 0.6\tilde{x}_2$ to replace x_1 and x_2 , respectively, in the quadratic equation:

$$\begin{aligned}100 &= 73(0.6\tilde{x}_1 - 0.8\tilde{x}_2)^2 \\ &\quad - 72(0.6\tilde{x}_1 - 0.8\tilde{x}_2)(0.8\tilde{x}_1 + 0.6\tilde{x}_2) \\ &\quad + 52(0.8\tilde{x}_1 + 0.6\tilde{x}_2)^2 \\ &= 73(0.36\tilde{x}_1^2 - 0.96\tilde{x}_1\tilde{x}_2 + 0.64\tilde{x}_2^2) \\ &\quad - 72(0.48\tilde{x}_1^2 - 0.28\tilde{x}_1\tilde{x}_2 - 0.48\tilde{x}_2^2) \\ &\quad + 52(0.64\tilde{x}_1^2 + 0.96\tilde{x}_1\tilde{x}_2 - 0.36\tilde{x}_2^2) \\ &= 25\tilde{x}_1^2 + 100\tilde{x}_2^2.\end{aligned}$$

So, we got rid of the the term involving the product of the two elements in the vector. After dividing by 100 on both sides of this equation we get

$$\left(\frac{\tilde{x}_1}{2}\right)^2 + \left(\frac{\tilde{x}_2}{1}\right)^2 = 1.$$

This is the equation (with respect to the basis \mathbf{b}_1 and \mathbf{b}_2) for points on an ellipse with half axes 2 and 1. In order to draw this curve we exploit that a point on the ellipse can be expressed by

$$\begin{pmatrix} \tilde{x}_1 \\ \tilde{x}_2 \end{pmatrix} = \begin{pmatrix} 2 \cos u \\ \sin u \end{pmatrix},$$

for some value of the angle u . By letting u run through the interval $0 \leq u \leq 2\pi$ we can generate all points on the ellipse. This is the background for the following MATLAB code, which was used to plot the ellipse in the figure.

```
u = linspace(0,2*pi,201);
xt = [2*cos(u); sin(u)];
x = [0.6 -0.8; 0.8 0.6] * xt; % original coordinates
plot(x(1,:),x(2,:))
```





Figure 1.6. Rotated ellipse.

The matrix used in the relation between \mathbf{x} and $\tilde{\mathbf{x}}$ is a rotation matrix

$$\mathbf{A} = \begin{pmatrix} \cos v & -\sin v \\ \sin v & \cos v \end{pmatrix},$$

cf Example 1.19. The angle is defined by

$$\cos v = 0.6, \quad \sin v = \sqrt{1 - \cos^2 v} = 0.8.$$

This corresponds to $v \simeq 53$ degrees, which is the angle between (1) and \vec{b}_1 in Figure 1.6.

In conclusion, the points satisfying the given quadratic are on an ellipse, which is rotated around its centre. The rotation angle is $v \simeq 53$ degrees.

In Example 3.9 we show how we found out that the specific choice of basis vectors \mathbf{b}_1 and \mathbf{b}_2 would simplify the quadratic equation. Also we show how we can avoid the tiresome calculations in the reformulation from $73x_1^2 - 72x_1x_2 + 52x_2^2 = 100$ to $25\tilde{x}_1^2 + 100\tilde{x}_2^2 = 100$. ■

2. Linear Systems of Equations

A linear system of equations has the form

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= b_1 \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n &= b_2 \\ &\vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n &= b_m \end{aligned}, \quad (2.1)$$

where the coefficients $\{a_{ij}\}$ and right-hand sides $\{b_i\}$ are given numbers. We seek the solution x_1, x_2, \dots, x_n .

From the discussion of matrix–vector products in Section 1.3 it follows that we can write the system in the compact form

$$\mathbf{Ax} = \mathbf{b}, \quad (2.2)$$

where the *coefficient matrix* $\mathbf{A} \in \mathbb{R}^{m \times n}$ has the elements a_{ij} , and the right-hand side vector \mathbf{b} and the solution vector \mathbf{x} contain respectively $\{b_i\}$ and $\{x_j\}$.

Example 2.1. Figure 2.1 shows three points (t_i, y_i) in the plane: $(1, 0.8)$, $(2, 0.7)$ and $(3, 1.6)$.




Figure 2.1. Three given points on a parabola.

We use a model which says that the points are on a parabola, $y_i = p(t_i)$, where

$$p(t) = x_1 t^2 + x_2 t + x_3.$$

However, we do not know the coefficients x_1 , x_2 and x_3 .

Using the same arguments as in Example 1.11, we see that the following equations must be satisfied,

$$t_i^2 x_1 + t_i x_2 + x_3 = y_i, \quad i = 1, 2, 3.$$

This is three linear equations with the three unknowns x_1 , x_2 and x_3 . We can express the three equations in the form

$$\begin{aligned} t_1^2 x_1 + t_1 x_2 + x_3 &= y_1 \\ t_2^2 x_1 + t_2 x_2 + x_3 &= y_2 \\ t_3^2 x_1 + t_3 x_2 + x_3 &= y_3, \end{aligned}$$

or as $\mathbf{Ax} = \mathbf{y}$, where

$$\mathbf{A} = \begin{pmatrix} t_1^2 & t_1 & 1 \\ t_2^2 & t_2 & 1 \\ t_3^2 & t_3 & 1 \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix}.$$

By inserting the given values for (t_i, y_i) we get the system

$$\begin{pmatrix} 1 & 1 & 1 \\ 4 & 2 & 1 \\ 9 & 3 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 0.8 \\ 0.7 \\ 1.6 \end{pmatrix}.$$

■

Example 2.2. In many textbooks the “determinant method” is used to solve a system of two linear equations with two unknowns,

$$a_{11}x_1 + a_{12}x_2 = b_1,$$

$$a_{21}x_1 + a_{22}x_2 = b_2.$$

The *determinant* of the coefficient matrix \mathbf{A} is defined by

$$\det \mathbf{A} = \begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix} = a_{11}a_{22} - a_{21}a_{12}.$$

This is zero if and only if there is a scalar β such that

$$\begin{pmatrix} a_{11} \\ a_{21} \end{pmatrix} = \beta \begin{pmatrix} a_{12} \\ a_{22} \end{pmatrix},$$

ie, the two columns in \mathbf{A} are linearly dependent, cf page 9. Such a matrix is said to be *singular*. This case is discussed in Section 2.3.

If $\det \mathbf{A} \neq 0$, we say that \mathbf{A} is *nonsingular*, and the solution is

$$x_1 = \frac{a_{22}b_1 - a_{12}b_2}{\det \mathbf{A}}, \quad x_2 = \frac{a_{11}b_2 - a_{21}b_1}{\det \mathbf{A}}.$$

This can be verified by inserting these expressions in the above system.

As a specific example, consider the system $\mathbf{Ax} = \mathbf{b}$ with

$$\mathbf{A} = \begin{pmatrix} 2 & -4 \\ 2.5 & -2 \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} 4 \\ 6.5 \end{pmatrix}.$$

We get

$$\det \mathbf{A} = 2 \cdot (-2) - 2.5 \cdot (-4) = 6,$$

This is nonzero, so the matrix is nonsingular, and the solution is

$$x_1 = \frac{(-2) \cdot 4 - (-4) \cdot 6.5}{6} = 3, \quad x_2 = \frac{2 \cdot 6.5 - 2.5 \cdot 4}{6} = 0.5.$$

The determinant method may be generalized to larger systems, but as n grows it becomes very inefficient. ■

The solution of linear systems of equations occurs so often in scientific computation that it probably is the type of task that uses most computer effort. There are many efficient programs for the solution of the problem, and in surroundings like MATLAB or R there are simple commands that call such a program. In the next sections we give a brief introduction to the mathematical background of these programs. We first discuss some particularly simple cases, and then proceed to a generally applicable solution method.

2.1. Simple systems

The simplest case of a linear system of equations is when the coefficient matrix is *diagonal*, ie when the equations have the form

$$a_{ii}x_i = b_i, \quad i = 1, 2, \dots, n.$$

Assuming that all the a_{ii} are nonzero, the solution is

$$x_i = b_i/a_{ii}, \quad i = 1, 2, \dots, n.$$

Another simple case is when the system has the form

$$\mathbf{Q}\mathbf{x} = \mathbf{b}, \tag{2.3}$$

where \mathbf{Q} is an *orthogonal matrix*. This means that the columns in \mathbf{Q} are *orthonormal*:

$$\mathbf{Q}_{:,i}^T \mathbf{Q}_{:,j} = \begin{cases} 0 & \text{if } i \neq j, \\ 1 & \text{if } i = j. \end{cases},$$

cf (1.13). The system (2.3) is equivalent to

$$x_1 \mathbf{Q}_{:,1} + x_2 \mathbf{Q}_{:,2} + \dots + x_n \mathbf{Q}_{:,n} = \mathbf{b},$$

and as we did in Section 1.5, we see that

$$x_k = \mathbf{Q}_{:,k}^T \mathbf{b}, \quad k = 1, 2, \dots, n.$$

It is easy to verify that this means that the solution to the orthogonal system in Eq. (2.3) can be expressed as

$$\mathbf{x} = \mathbf{Q}^T \mathbf{b}, \tag{2.4}$$

ie the solution is found simply by multiplying the right-hand side \mathbf{b} by the transposed of the coefficient matrix \mathbf{Q} . We shall use this several times in the remainder of the note.

Example 2.3. In Example 1.19 we showed that a rotation matrix

$$\mathbf{A} = \begin{pmatrix} \cos v & -\sin v \\ \sin v & \cos v \end{pmatrix}.$$

has orthonormal columns. This means that the vectors in the two columns can be used as basis in \mathbb{R}^2 , and letting $\tilde{\mathbf{x}}$ denote the coordinates with respect to this basis, we see that

$$\mathbf{x} = \mathbf{A} \tilde{\mathbf{x}} \Leftrightarrow \tilde{\mathbf{x}} = \mathbf{A}^T \mathbf{x}. \quad \blacksquare$$

Finally, we consider *triangular systems*. This means that the coefficient matrix is triangular. We start by looking at $\mathbf{L}\mathbf{x} = \mathbf{b}$, where \mathbf{L} is a lower triangular matrix, so the system of equations has the form

$$\begin{aligned} \ell_{11}x_1 &= b_1 \\ \ell_{21}x_1 + \ell_{22}x_2 &= b_2 \\ &\vdots \\ \ell_{n1}x_1 + \ell_{n2}x_2 + \dots + \ell_{nn}x_n &= b_n \end{aligned}$$

The first equation has x_1 as the only unknown, and assuming that the coefficient $\ell_{11} \neq 0$, we get $x_1 = b_1/\ell_{11}$. We insert this value in the second equation and if $\ell_{22} \neq 0$, we get $x_2 = (b_2 - \ell_{21}x_1)/\ell_{22}$. The known values for x_1 and x_2 are inserted in the third equation, etc. We can summarize this in the form

$$\begin{array}{l} x_1 = b_1/\ell_{11} \\ \mathbf{for} \quad i = 2, 3, \dots, n \\ \quad x_i = (b_i - \ell_{i1}x_1 - \dots - \ell_{i,i-1}x_{i-1})/\ell_{ii} \\ \mathbf{end} \end{array} \quad (2.5)$$

This *algorithm* is called *forward substitution*. It can be used only if all $\ell_{ii} \neq 0$.

It is equally simple to solve the problem $\mathbf{U}\mathbf{x} = \mathbf{c}$, where \mathbf{U} is an upper triangular matrix,

$$\begin{array}{r} u_{11}x_1 + u_{12}x_2 + \dots + u_{1n}x_n = c_1 \\ u_{22}x_2 + \dots + u_{2n}x_n = c_2 \\ \vdots \\ u_{nn}x_n = c_n \end{array}$$

We solve the last equation first and move up. Assuming that all $u_{ii} \neq 0$, we get

$$\begin{array}{l} x_n = c_n/u_{nn} \\ \mathbf{for} \quad i = n-1, \dots, 2, 1 \\ \quad x_i = (c_i - u_{i,i+1}x_{i+1} - \dots - u_{i,n}x_n)/u_{ii} \\ \mathbf{end} \end{array} \quad (2.6)$$

This algorithm is known as *back substitution*.

Example 2.4. Consider the upper triangular system

$$\begin{pmatrix} 9 & 3 & 1 \\ & \frac{2}{3} & \frac{5}{9} \\ & & \frac{3}{9} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 1.6 \\ \frac{-0.1}{9} \\ \frac{5.7}{9} \end{pmatrix}.$$

By means of Algorithm (2.6) we get

$$\begin{aligned} x_3 &= \frac{5.7}{9} / \frac{3}{9} = 1.9, \\ x_2 &= \left(\frac{-0.1}{9} - \frac{5}{9} \cdot 1.9 \right) / \frac{2}{3} = -1.6, \\ x_1 &= (1.6 - 3 \cdot (-1.6) - 1 \cdot 1.9) / 9 = 0.5. \quad \blacksquare \end{aligned}$$

2.2. Gaussian Elimination

The idea is to transform the given general, square system of equations $\mathbf{A}\mathbf{x} = \mathbf{b}$ to an upper triangular system $\mathbf{U}\mathbf{x} = \mathbf{c}$, which can then be solved by back substitution. The transformation must be made so that the solution is not changed, and this is ensured if we use a sequence of one or both of the two elementary operations

- 1° Interchange two equations.
- 2° Modify an equation by subtracting a multiple of one of the other equations.

Example 2.5. Given the system from Example 2.2

$$\begin{aligned} 2x_1 - 4x_2 &= 4 \\ 2.5x_1 - 2x_2 &= 6.5 \end{aligned}$$

We modify the second equation by subtracting 1.25 times the first equation:

$$2.5x_1 - 2x_2 - 1.25(2x_1 - 4x_2) = 6.5 - 1.25 \cdot 4.$$

This reduces to $3x_2 = 1.5$, so the original system is equivalent to the upper triangular system

$$\begin{aligned} 2x_1 - 4x_2 &= 4 \\ 3x_2 &= 1.5 \end{aligned}$$

By means of back substitution we get the same solution as in Example 2.2: $x_2 = 0.5$, $x_1 = 3$. Note that the *elimination factor* $1.25 = 2.5/2$ was chosen in order to eliminate x_2 from the modified equation.

Figure 2.2 gives a geometric interpretation: The two equations correspond to the straight lines $L_1 : 2x_1 - 4x_2 = 4$ and $L_2 : 2.5x_1 - 2x_2 = 6.5$, and the solution is the coordinates of the point where they intersect. The modification of the system corresponds to replacing L_2 by $\hat{L}_2 : x_2 = 0.5$.




Figure 2.2. *Gaussian elimination.*

Any linear equation with two unknowns and at least one nonzero coefficient can be interpreted as the equation for a straight line, and the solution to a 2×2 system is the the point of intersection between the two lines. The figure suggests possible complications:

L_1 may be parallel to the (1)-axis. This happens if $a_{11} = 0$, and this problem might be cured simply by interchanging the two equations, corresponding to renumbering the two lines.

The two lines may be parallel, in which case the system has no solution.

The two line may be coinciding, in which case there is infinitely many solutions: all points on the line satisfy both equations.

We will return to these problems in Section 2.3. ■

Let us now consider a general $n \times n$ system $\mathbf{Ax} = \mathbf{b}$. It is convenient for the presentation (and for hand calculation) to introduce the so-called *augmented matrix* \mathbf{T} , which is the $n \times (n+1)$ matrix

$$\mathbf{T} = \left(\begin{array}{cccc|c} a_{11} & a_{12} & \dots & a_{1n} & b_1 \\ a_{21} & a_{22} & \dots & a_{2n} & b_2 \\ \vdots & \vdots & & \vdots & \vdots \\ a_{n1} & a_{n2} & \dots & a_{nn} & b_n \end{array} \right). \quad (2.7)$$

The vertical line is put there only to separate the coefficients and the elements of the right-hand side. Note that the i^{th} row $\mathbf{T}_{i,:}$ holds all information about the i^{th} equation. During computation the elements in \mathbf{T} change, but in order to avoid a very complicated notation, we let a_{ij} denote the current value of the coefficient of x_j in the i^{th} equation, and b_i is the current value of the i^{th} right-hand side. Accordingly we use “ \leftarrow ”, for instance $a \leftarrow a + b$, to indicate that the old value of a is replaced by $a+b$.

The first step in Gaussian elimination is to keep x_1 in one equation and eliminate it from all the others. This is done by first finding p_1 , the number of the equation with the largest contribution from x_1 ,

$$|a_{p_1,1}| = \max_{i=1,2,\dots,n} |a_{i1}| \quad (2.8)$$

If $a_{p_1,1} = 0$, then x_1 does not appear in any of the equations, and we proceed to the next step. Otherwise, if $p_1 \neq 1$, then interchange the 1st and the p_1 th equation,

$$\mathbf{T}_{1,:} \leftrightarrow \mathbf{T}_{p_1,:},$$

and modify equations

$$\left. \begin{array}{l} \ell_{i1} = a_{i1}/a_{11} \\ \mathbf{T}_{i,:} \leftarrow \mathbf{T}_{i,:} - \ell_{i1}\mathbf{T}_{1,:} \end{array} \right\} \quad i = 2, 3, \dots, n.$$

After this step the original system has been changed to

$$\mathbf{T} = \left(\begin{array}{cccc|c} a_{11} & a_{12} & \dots & a_{1n} & b_1 \\ 0 & a_{22} & \dots & a_{2n} & b_2 \\ \vdots & \vdots & & \vdots & \vdots \\ 0 & a_{n2} & \dots & a_{nn} & b_n \end{array} \right),$$

where all the a_{ij} and b_i may have values different from the original values. Note that the *active part* of the system: rows $2, \dots, n$, columns $2, \dots, n, n+1$, has the same structure as the original system, but it is smaller.

The second step is similar to the first: Keep x_2 in one of the active equations and eliminate it from all the other active equations. This continues. For $k = 1, 2, \dots, n-1$ the active part of the system consists of rows k, \dots, n and columns $k, \dots, n+1$. We keep x_k in one of the active equations (if possible) and delete it

from the others. The entire elimination algorithm can be expressed by

```

for  $k = 1, 2, \dots, n-1$ 
  find  $p_k : |a_{p_k,k}| = \max_{i=k+1,\dots,n} |a_{i,k}|$ 
  if  $a_{p_k,k} \neq 0$  then
    if  $p_k \neq k$  then  $\mathbf{T}_{k,:} \leftrightarrow \mathbf{T}_{p_k,:}$  end
    for  $i = k+1, \dots, n$ 
       $\ell_{ik} = a_{ik}/a_{kk}$ 
       $\mathbf{T}_{i,:} \leftarrow \mathbf{T}_{i,:} - \ell_{ik}\mathbf{T}_{k,:}$ 
    end
  end
end

```

(2.9)

After executing this algorithm the augmented matrix has the structure

$$\mathbf{T} = (\mathbf{U} \mid \mathbf{c}) ,$$

where \mathbf{U} is upper triangular and \mathbf{c} contains the final values of the elements in the right-hand side vector.

Example 2.6. Consider the parabola problem from Example 2.1,

$$\begin{pmatrix} 1 & 1 & 1 \\ 4 & 2 & 1 \\ 9 & 3 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 0.8 \\ 0.7 \\ 1.6 \end{pmatrix} .$$

We use the augmented matrix notation, and indicate the operations that lead from one version of \mathbf{T} to the next:

$$\begin{pmatrix} 1 & 1 & 1 & \mid & 0.8 \\ 4 & 2 & 1 & \mid & 0.7 \\ 9 & 3 & 1 & \mid & 1.6 \end{pmatrix} \xrightarrow{p_1=3} \begin{pmatrix} 9 & 3 & 1 & \mid & 1.6 \\ 4 & 2 & 1 & \mid & 0.7 \\ 1 & 1 & 1 & \mid & 0.8 \end{pmatrix} \xrightarrow{\substack{\ell_{21} = \frac{4}{9} \\ \ell_{31} = \frac{1}{9}}} \begin{pmatrix} 9 & 3 & 1 & \mid & 1.6 \\ 0 & \frac{2}{3} & \frac{5}{9} & \mid & -\frac{0.1}{9} \\ 0 & \frac{2}{3} & \frac{8}{9} & \mid & \frac{5.6}{9} \end{pmatrix} \xrightarrow{\substack{p_2=2 \\ \ell_{3,2}=1}} \begin{pmatrix} 9 & 3 & 1 & \mid & 1.6 \\ 0 & \frac{2}{3} & \frac{5}{9} & \mid & -\frac{0.1}{9} \\ 0 & 0 & \frac{3}{9} & \mid & \frac{5.7}{9} \end{pmatrix} .$$

Thus, the solution \mathbf{x} satisfies the upper triangular system

$$\begin{pmatrix} 9 & 3 & 1 \\ & \frac{2}{3} & \frac{5}{9} \\ & & \frac{3}{9} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 1.6 \\ -\frac{0.1}{9} \\ \frac{5.7}{9} \end{pmatrix} .$$

We solved this system in Example 2.4. The solution is

$$\mathbf{x} = (0.5 \quad -1.6 \quad 1.9)^T .$$

Note that this agrees with the coefficients in the parabola given in Figure 1.3, so the three given points identified the correct parabola. ■

The term *Gaussian elimination* is used to cover the entire solution method, involving both the elimination algorithm and back substitution in the transformed, upper triangular system.

As presented, the algorithm involves *pivoting*. This means that each diagonal element is determined as the largest element in the leading column of the active part of the system. In hand calculation this part is needed only, if the current diagonal

element is zero. When the calculations are made on a computer, the use of it is very important. The reason is that a computer has a limited accuracy, and if an elimination factor ℓ_{ik} is very large, then there is a risk that the result of

$$\mathbf{T}_{i,:} \leftarrow \mathbf{T}_{i,:} - \ell_{ik}\mathbf{T}_{k,:}$$

is so dominated by the contribution $\ell_{ik}\mathbf{T}_{k,:}$, that the “old” information in $\mathbf{T}_{i,:}$ is lost. With pivoting all the elimination factors satisfy $|\ell_{ik}| \leq 1$, and this reduces the risk of loss of information.

Example 2.7. Consider the matrix \mathbf{A} from Example 1.15 with a right-side $\mathbf{b} = \mathbf{Az}$:

$$\begin{array}{l} \left(\begin{array}{ccc|c} 1 & 2 & 3 & 8 \\ 4 & 5 & 6 & 20 \\ 7 & 8 & 9 & 32 \end{array} \right) \quad p_1 = 3 \quad \sim \quad \left(\begin{array}{ccc|c} 7 & 8 & 9 & 32 \\ 4 & 5 & 6 & 20 \\ 1 & 2 & 3 & 8 \end{array} \right) \quad \begin{array}{l} \sim \\ \ell_{21} = \frac{4}{7} \\ \ell_{31} = \frac{1}{7} \end{array} \\ \left(\begin{array}{ccc|c} 7 & 8 & 9 & 32 \\ 0 & \frac{3}{7} & \frac{6}{7} & \frac{12}{7} \\ 0 & \frac{6}{7} & \frac{12}{7} & \frac{24}{7} \end{array} \right) \quad p_2 = 3 \quad \sim \quad \left(\begin{array}{ccc|c} 7 & 8 & 9 & 32 \\ 0 & \frac{6}{7} & \frac{12}{7} & \frac{24}{7} \\ 0 & 0 & 0 & 0 \end{array} \right) \quad \begin{array}{l} \sim \\ \ell_{3,2} = \frac{1}{2} \end{array} \end{array} .$$

We see that $u_{33} = 0$, so there is a problem that we need to discuss. ■

2.3. Complete Solution

The elimination algorithm from the previous section transforms the given system $\mathbf{Ax} = \mathbf{b}$ to an upper triangular system $\mathbf{Ux} = \mathbf{c}$,

$$\begin{array}{rcl} u_{11}x_1 + u_{12}x_2 + \cdots + u_{1n}x_n & = & c_1 \\ u_{22}x_2 + \cdots + u_{2n}x_n & = & c_2 \\ & \vdots & \\ u_{nn}x_n & = & c_n \end{array} , \quad (2.10)$$

without changing the solution \mathbf{x} .

If all the $u_{ii} \neq 0$, then the solution is found simply by back substitution. Note that in this case the solution \mathbf{x} is unique.

The case where one or more $u_{ii} = 0$ needs further investigation. We say that the matrix \mathbf{A} is *singular* (and that it is *nonsingular* if all the $u_{ii} \neq 0$). We shall only give a detailed discussion of the simplest case, where $u_{nn} = 0$ is the only zero element on the diagonal of \mathbf{U} . Then the last equation in (2.10) has the form

$$0 \cdot x_n = c_n .$$

If $c_n \neq 0$, then this equation, and therefore the original system $\mathbf{Ax} = \mathbf{b}$, has no solution. If $c_n = 0$, we say that the system is *consistent*. Then the last equation is satisfied by $x_n = \alpha$, where α is any real number. In order to find the other elements in \mathbf{x} we see from Eq. (2.10) that they must satisfy the reduced system

$$\begin{array}{rcl} u_{11}x_1 + u_{12}x_2 + \cdots + u_{1,n-1}x_{n-1} & = & c_1 - u_{1,n}\alpha \\ u_{22}x_2 + \cdots + u_{2,n-1}x_{n-1} & = & c_2 - u_{2,n}\alpha \\ & \vdots & \\ u_{n-1,n-1}x_{n-1} & = & c_{n-1} - u_{n-1,n}\alpha \end{array} . \quad (2.11)$$

All the diagonal coefficients in this system are nonzero, so we can use back substitution and get

$$x_i = z_i + \alpha v_i, \quad i = n-1, \dots, 2, 1,$$

where

$$z_{n-1} = c_{n-1}/u_{n-1,n-1}, \quad v_{n-1} = -u_{i,n-1}/u_{n-1,n-1},$$

and for $i = n-2, \dots, 2, 1$:

$$\begin{aligned} z_i &= (c_i - u_{i,i+1}z_{i+1} - \dots - u_{i,n-1}z_{n-1})/u_{ii}, \\ v_i &= (-u_{i,n} - u_{i,i+1}v_{i+1} - \dots - u_{i,n-1}v_{n-1})/u_{ii}. \end{aligned}$$

So, in this case the *complete solution* to $\mathbf{Ax} = \mathbf{b}$ is

$$\mathbf{x} = \mathbf{x}_b + \alpha \mathbf{w}, \quad \alpha \in \mathbb{R}, \quad (2.12)$$

where

$$\mathbf{x}_b = \begin{pmatrix} z_1 \\ z_2 \\ \vdots \\ z_{n-1} \\ 0 \end{pmatrix}, \quad \mathbf{w} = \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_{n-1} \\ 1 \end{pmatrix}.$$

Example 2.8. In Example 2.7 we saw that

$$\left(\begin{array}{ccc|c} 1 & 2 & 3 & 8 \\ 4 & 5 & 6 & 20 \\ 7 & 8 & 9 & 32 \end{array} \right) \sim \left(\begin{array}{ccc|c} 7 & 8 & 9 & 32 \\ 0 & \frac{6}{7} & \frac{12}{7} & \frac{24}{7} \\ 0 & 0 & 0 & 0 \end{array} \right).$$

Using the method described above we get the complete solution

$$\mathbf{x} = \mathbf{x}_b + \alpha \mathbf{w} = \begin{pmatrix} 0 \\ 4 \\ 0 \end{pmatrix} + \alpha \begin{pmatrix} 1 \\ -2 \\ 1 \end{pmatrix}, \quad \alpha \in \mathbb{R}.$$

We recognize the vector \mathbf{w} from Example 1.15. The vector \mathbf{z} from the same example is obtained for $\alpha = 1$. ■

The solution $\mathbf{x} = \mathbf{x}_b + \alpha \mathbf{w}$ was constructed so that

$$\mathbf{A}(\mathbf{x}_b + \alpha \mathbf{w}) = \mathbf{b}$$

for any choice of α . By use of the linearity condition of a linear mapping, Eq. (1.7), we see that $\mathbf{A}(\mathbf{x}_b + \alpha \mathbf{w}) = \mathbf{b}$ is equivalent to the condition

$$\mathbf{A} \mathbf{x}_b + \alpha \mathbf{A} \mathbf{w} = \mathbf{b}$$

for any choice of α . This can be satisfied only, if

$$\mathbf{A} \mathbf{x}_b = \mathbf{b} \quad \text{and} \quad \mathbf{A} \mathbf{w} = \mathbf{0}. \quad (2.13)$$

We say that \mathbf{w} is in the *nullspace* (or *kernel*) of \mathbf{A} . The system

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

is said to be a *homogeneous system*. The zero vector $\mathbf{x} = \mathbf{0}$ satisfies this system for any \mathbf{A} , and if \mathbf{A} is nonsingular, then $\mathbf{0}$ is the only solution. If \mathbf{A} is singular, however, then it has a proper nullspace that contains nonzero vectors, for instance \mathbf{w} .

By rewriting the condition $\mathbf{A}\mathbf{w} = \mathbf{0}$ we see that

$$w_1\mathbf{A}_{:,1} + \cdots + w_{n-1}\mathbf{A}_{:,n-1} + \mathbf{A}_{:,n} = \mathbf{0}.$$

This shows (cf page 9) that the columns in a singular matrix are linearly dependent. We can write

$$\mathbf{A}_{:,n} = -w_1\mathbf{A}_{:,1} - \cdots - w_{n-1}\mathbf{A}_{:,n-1},$$

and this implies that the image of any $\mathbf{x} \in \mathbb{R}^n$ can be expressed as

$$\begin{aligned} \mathbf{A}\mathbf{x} &= x_1\mathbf{A}_{:,1} + \cdots + x_{n-1}\mathbf{A}_{:,n-1} + x_n\mathbf{A}_{:,n} \\ &= \hat{x}_1\mathbf{A}_{:,1} + \cdots + \hat{x}_{n-1}\mathbf{A}_{:,n-1}, \end{aligned} \tag{2.14}$$

where $\hat{x}_j = x_j - w_j x_n$. In other words, the image is a linear combination of the first $n-1$ columns in \mathbf{A} . This is a subspace of \mathbb{R}^n , called the *column range* of \mathbf{A} .

The condition for $\mathbf{A}\mathbf{x} = \mathbf{b}$ having a solution is that \mathbf{b} is “reachable”, ie that \mathbf{b} is in the column range of \mathbf{A} . If the $n \times n$ matrix is nonsingular, then the nullspace of \mathbf{A} has $\mathbf{0}$ as the only element, and the column range of \mathbf{A} is the entire \mathbb{R}^n .

Example 2.9. We shall briefly discuss a problem, $\mathbf{A}\mathbf{x} = \mathbf{b} \sim \mathbf{U}\mathbf{x} = \mathbf{c}$, where \mathbf{U} has more than one $u_{ii} = 0$ or the zero diagonal element is not in the last row.

In all of these cases the matrix \mathbf{A} is singular, and the solution – if it exists – has at least one free parameter. Intuitively, the number of free parameters is equal to the number of zero diagonal elements in \mathbf{U} , but this guess may be wrong! To illustrate that, consider the two examples

$$\begin{aligned} \left(\begin{array}{ccccc|c} 4 & 2 & 10 & 0 & 4 & 50 \\ 2 & 4 & 8 & 6 & 14 & 64 \\ 1 & 0 & 2 & 3 & 3 & 10 \\ 3 & 2 & 8 & 9 & 13 & 52 \\ 0 & 3 & 3 & 3 & 9 & 36 \end{array} \right) &\sim \left(\begin{array}{ccccc|c} 4 & 2 & 10 & 0 & 4 & 50 \\ 0 & 3 & 3 & 6 & 12 & 39 \\ 0 & 0 & 0 & 4 & 4 & 4 \\ 0 & 0 & 0 & 8 & 8 & 8 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right), \\ \left(\begin{array}{ccccc|c} 4 & 6 & 8 & 8 & 4 & 4 \\ 2 & 13 & 10 & 0 & 6 & 36 \\ 1 & -1 & 0.5 & 11 & 6 & -9.5 \\ 3 & 7 & 7.5 & 9 & 6 & 5.5 \\ 0 & -5 & -3 & 4 & -1 & -20 \end{array} \right) &\sim \left(\begin{array}{ccccc|c} 4 & 6 & 8 & 8 & 4 & 4 \\ 0 & 10 & 6 & -4 & 4 & 34 \\ 0 & 0 & 0 & 8 & 6 & -2 \\ 0 & 0 & 0 & 4 & 2 & -6 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right), \end{aligned}$$

In both cases \mathbf{U} has two zero diagonals, and for each \mathbf{A} the right-hand side was generated as $\mathbf{b} = \mathbf{A}\check{\mathbf{x}}$, where

$$\check{\mathbf{x}} = (1 \ -2 \ 3 \ -4 \ 5)^T,$$

so both systems have a solution. For the first system the complete solution is

$$\mathbf{x} = \begin{pmatrix} 7 \\ 11 \\ 0 \\ 1 \\ 0 \end{pmatrix} + \begin{pmatrix} -2 & 0 \\ -1 & -2 \\ 1 & 0 \\ 0 & -1 \\ 0 & 1 \end{pmatrix} \boldsymbol{\alpha}, \quad \boldsymbol{\alpha} \in \mathbb{R}^2,$$

so the solution involves two free parameters. The vector $\check{\mathbf{x}}$ is obtained for $\alpha_1 = -13$, $\alpha_2 = 5$.

The complete solution to the second system is

$$\mathbf{x} = \begin{pmatrix} 4.3 \\ -0.2 \\ 0 \\ -4 \\ 5 \end{pmatrix} + \begin{pmatrix} -1.1 \\ -0.6 \\ 1 \\ 0 \\ 0 \end{pmatrix} \alpha, \quad \alpha \in \mathbb{R},$$

so for this problem there is only one free parameter. The vector $\check{\mathbf{x}}$ is obtained for $\alpha = 3$.

It is outside the scope of this lecture note to describe how we computed the complete solution. In Section 4.3 we discuss a method, which is better than Gaussian elimination for computing nullspace, column range and complete solution. ■

2.4. Matrix inverse

Let $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{B} \in \mathbb{R}^{n \times q}$, and consider the *matrix equation*

$$\mathbf{A}\mathbf{X} = \mathbf{B}.$$

If \mathbf{A} is nonsingular, then there is a unique solution $\mathbf{X} \in \mathbb{R}^{n \times q}$, whose columns satisfy the linear system of equations

$$\mathbf{A}\mathbf{X}_{:,j} = \mathbf{B}_{:,j}, \quad j = 1, 2, \dots, q.$$

The algorithms in Gaussian elimination are simple to adjust to cope with matrix equations: If we let the augmented matrix include all the columns of \mathbf{B} ,

$$\mathbf{T} = (\mathbf{A} \mid \mathbf{B}),$$

then the elimination algorithm (2.9) needs not be changed, and it leads to

$$\mathbf{T} \sim (\mathbf{U} \mid \mathbf{C}).$$

If there are no zeros on the diagonal of \mathbf{U} , then the solution can be found by the following modified version of the back substitution algorithm (2.6),

```

 $\mathbf{X}_{n,:} = \mathbf{C}_{:,j}/u_{nn}$ 
for  $i = n-1, \dots, 2, 1$ 
   $\mathbf{X}_{i,:} = (\mathbf{C}_{i,:} - u_{i,i+1}\mathbf{X}_{i+1,:} - \dots - u_{i,n}\mathbf{X}_{n,:})/u_{ii}$ 
end

```

Example 2.10. We would like to solve the matrix equation

$$\begin{pmatrix} 2 & 3 \\ 4 & 5 \end{pmatrix} \mathbf{X} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

With the notation from Example 2.6 we get

$$\left(\begin{array}{cc|cc} 2 & 3 & 1 & 0 \\ 4 & 5 & 0 & 1 \end{array} \right) \underset{p_1=2}{\sim} \left(\begin{array}{cc|cc} 4 & 5 & 0 & 1 \\ 2 & 3 & 1 & 0 \end{array} \right) \underset{\ell_{21}=\frac{1}{2}}{\sim} \left(\begin{array}{cc|cc} 4 & 5 & 0 & 1 \\ 0 & \frac{1}{2} & 1 & -\frac{1}{2} \end{array} \right),$$

and back substitution gives

$$\begin{aligned} \mathbf{X}_{2,:} &= (1 \quad -\frac{1}{2}) / \frac{1}{2} = (2 \quad -1), \\ \mathbf{X}_{1,:} &= \{(0 \quad 1) - 5(2 \quad -1)\} / 4 = (\frac{-5}{2} \quad \frac{3}{2}). \end{aligned}$$

This means that the solution is $\mathbf{X} = \frac{1}{2} \begin{pmatrix} -5 & 3 \\ 4 & -2 \end{pmatrix}$. ■

The matrix equation with $\mathbf{B} = \mathbf{I}$, the identity matrix, is of special interest. Its solution is called “ \mathbf{A} inverse” and is written \mathbf{A}^{-1} ,

$$\mathbf{A} \mathbf{A}^{-1} = \mathbf{I}. \tag{2.15}$$

The matrix inverse is defined only for nonsingular matrices.

Now, let \mathbf{A} be a nonsingular $n \times n$ matrix, \mathbf{b} an n -vector, and define $\mathbf{y} = \mathbf{A}^{-1}\mathbf{b}$. Then

$$\mathbf{A} \mathbf{y} = \mathbf{A} \mathbf{A}^{-1} \mathbf{b} = \mathbf{I} \mathbf{b} = \mathbf{b}.$$

This shows that $\mathbf{y} = \mathbf{x}$, the unique solution to $\mathbf{A} \mathbf{x} = \mathbf{b}$, ie, if \mathbf{A} is nonsingular, then

$$\mathbf{A} \mathbf{x} = \mathbf{b} \Leftrightarrow \mathbf{x} = \mathbf{A}^{-1} \mathbf{b}. \tag{2.16}$$

If we multiply first equation by \mathbf{A}^{-1} from the left, we get

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

and combining this with Eqs. (2.15) and (2.16) we see that the matrix inverse satisfies

$$\mathbf{A} \mathbf{A}^{-1} = \mathbf{A}^{-1} \mathbf{A} = \mathbf{I}. \tag{2.17}$$

This is one of the rare examples, where the interchange of the factors in a matrix-matrix product does not change the result.

Example 2.11. In Example 2.10 we actually computed a matrix inverse:

$$\mathbf{A} = \begin{pmatrix} 2 & 3 \\ 4 & 5 \end{pmatrix} \Leftrightarrow \mathbf{A}^{-1} = \frac{1}{2} \begin{pmatrix} -5 & 3 \\ 4 & -2 \end{pmatrix}.$$

Now consider the system $\mathbf{A} \mathbf{x} = \begin{pmatrix} 4 \\ 6 \end{pmatrix}$. By means of Eq. (2.16) we can find the solution:

$$\mathbf{x} = \frac{1}{2} \begin{pmatrix} -5 & 3 \\ 4 & -2 \end{pmatrix} \begin{pmatrix} 4 \\ 6 \end{pmatrix} = \begin{pmatrix} -1 \\ 2 \end{pmatrix}. \quad \blacksquare$$

There are matrices for which it is almost embarrassingly simple to compute the inverse. We shall mention diagonal and orthogonal matrices:

$$\mathbf{D} = \begin{pmatrix} d_{11} & & & \\ & d_{22} & & \\ & & \ddots & \\ & & & d_{nn} \end{pmatrix} \Leftrightarrow \mathbf{D}^{-1} = \begin{pmatrix} d_{11}^{-1} & & & \\ & d_{22}^{-1} & & \\ & & \ddots & \\ & & & d_{nn}^{-1} \end{pmatrix}.$$

Example 2.12.

$$D = \begin{pmatrix} 1 & & \\ & \frac{1}{2} & \\ & & 3 \end{pmatrix} \Leftrightarrow D^{-1} = \begin{pmatrix} 1 & & \\ & 2 & \\ & & \frac{1}{3} \end{pmatrix} .$$

An orthogonal matrix Q satisfies

$$Q^T Q = I , \tag{2.18}$$

of the discussion at the start of Section 2.1. According to Eq. (2.17) this implies that

$$Q^{-1} = Q^T ,$$

and that also Q^T is orthogonal

$$Q Q^T = I .$$

It is wasteful to use Gaussian elimination to solve a system $Qx = y$ with an orthogonal matrix. It is both faster and slightly more accurate to exploit that

$$Qx = b \Leftrightarrow x = Q^T b . \tag{2.19}$$

Example 2.13. We actually exploited Eq. (2.19) in connection with Eq. (1.14). The general equation for basis shift is $x = B\alpha$, where the columns in B are the coordinates of the “new” basis vectors. If they are orthonormal, then B is orthogonal, and the “new” coordinates are $\alpha = B^T x$. ■

For a general square, nonsingular matrix A it must be said that Eq. (2.16) is a convenient notation when discussing a mathematical model, but “ $x = A^{-1}b$ ” should be read as shorthand for “find x as the solution to $Ax = b$ ”. The matrix inverse itself is only needed in special cases, and the solution via Gaussian elimination is both faster and usually also less affected by rounding errors than the solution via the matrix inverse.

Exercise 2.14. We can use this comment on the matrix inverse to derive the following result: Let A and B be square and nonsingular. Show that

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

Hint: Consider the matrix equation $ABX = I$ and introduce the auxiliary matrix $C = BX$. ■

Example 2.15. It sometimes happens that one has to solve a series of problems

$$Ax_k = b_k , \quad k = 1, 2, \dots, K ,$$

with the same matrix but different right-hand sides. It is, of course, wasteful to start from scratch with each problem, since the operations on the matrix elements are the same.

If all b_k are known, then the simplest method is to solve the problem as a matrix equation $AX = B$, where the $n \times K$ matrix B has b_k as its k^{th} column, $B_{:,k} = b_k$.

In some cases, however, b_k is a function of x_1, x_2, \dots, x_{k-1} , so it is not known from the start, and we cannot use the matrix equation approach. One might argue that

the use of \mathbf{A}^{-1} would be appropriate here, since the inverse only has to be computed once, and then each \mathbf{x}_k is computed by the matrix-vector product $\mathbf{A}^{-1}\mathbf{b}_k$. In the next section we describe a better method for solving this problem. ■

2.5. LU factorization

Example 2.16. In Example 2.7 we saw that

$$\mathbf{A} = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \begin{array}{l} \sim \\ p_1 = 3 \\ \ell_{21} = \frac{4}{7} \\ \ell_{31} = \frac{1}{7} \end{array} \begin{pmatrix} 7 & 8 & 9 \\ 0 & \frac{3}{7} & \frac{6}{7} \\ 0 & \frac{6}{7} & \frac{12}{7} \end{pmatrix} \\ \begin{array}{l} \sim \\ p_2 = 3 \\ \ell_{32} = \frac{1}{2} \end{array} \begin{pmatrix} 7 & 8 & 9 \\ 0 & \frac{6}{7} & \frac{12}{7} \\ 0 & 0 & 0 \end{pmatrix} = \mathbf{U} .$$

We introduce a lower triangular matrix \mathbf{L} with ones on the diagonal and the elimination factors stored column-wise in the strictly lower triangle. The storage is affected by row interchanges: if $p_k \neq k$, we use the row interchange $\mathbf{T}_{k,:} \leftrightarrow \mathbf{T}_{p_k,:}$ in Algorithm (2.9), and the same interchange should be applied to the “preceding” elimination factors,

$$(\mathbf{L})_{kj} \leftrightarrow (\mathbf{L})_{p_k j}, \quad j = 1, \dots, k-1 .$$

With the above matrix we have $p_2 = 3$, so ℓ_{21} and ℓ_{31} end up in $(\mathbf{L})_{31}$ and $(\mathbf{L})_{21}$, respectively, while $(\mathbf{L})_{32} = \ell_{32}$:

$$\mathbf{L} = \begin{pmatrix} 1 & 0 & 0 \\ \frac{1}{7} & 1 & 0 \\ \frac{4}{7} & \frac{1}{2} & 1 \end{pmatrix} .$$

It is easily verified that

$$\mathbf{L}\mathbf{U} = \begin{pmatrix} 7 & 8 & 9 \\ 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} .$$

This product is equal to the matrix obtained after the row interchange $\mathbf{A}_{1,:} \leftrightarrow \mathbf{A}_{3,:}$, followed by $\mathbf{A}_{2,:} \leftrightarrow \mathbf{A}_{3,:}$. ■

The result in this example is no coincidence. We used the elimination algorithm Eq. (2.9), and it can be shown that the effect on the coefficient matrix can be expressed by

$$\mathbf{P}\mathbf{A} = \mathbf{L}\mathbf{U} . \tag{2.20}$$

This is the so-called *LU factorization* of \mathbf{A} . The matrix \mathbf{P} is a *permutation matrix*, characterized by each row and each column having one element equal to 1, and all the other elements are zero. \mathbf{L} is a lower triangular matrix with ones on the diagonal and the elimination factors in the strictly lower triangle, stored as outlined in the example. Finally, \mathbf{U} is the upper triangular matrix resulting from the elimination algorithm.

Example 2.17. The permutation matrix \mathbf{P} can be obtained from the identity matrix \mathbf{I} by applying the row interchanges defined by the elimination algorithm. For the

problem in Example 2.16 we get

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \underset{p_1=3}{\sim} \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \underset{p_2=3}{\sim} \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} = \mathbf{P}.$$

Note that the columns in \mathbf{P} are orthogonal. This is a general property of a permutation matrix, and implies that $\mathbf{P}^{-1} = \mathbf{P}^T$.

Let j_i denote the position of the 1 in the i^{th} row of \mathbf{P} . Then the result of $\mathbf{P}\mathbf{A}$ is that the rows of \mathbf{A} come in the sequence j_1, j_2, \dots, j_n . In particular, for the problem of Example 2.16 the above \mathbf{P} gives $j_1 = 3, j_2 = 1, j_3 = 2$, and

$$\mathbf{P}\mathbf{A} = \begin{pmatrix} 7 & 8 & 9 \\ 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix}.$$

This is equal to the result of the product $\mathbf{L}\mathbf{U}$ in Example 2.16. ■

Now, assume that we know the LU factorization of \mathbf{A} and want to solve the system $\mathbf{A}\mathbf{x} = \mathbf{b}$. This is obviously equivalent to $\mathbf{P}\mathbf{A}\mathbf{x} = \mathbf{P}\mathbf{b}$, or

$$\mathbf{L}\mathbf{U}\mathbf{x} = \mathbf{P}\mathbf{b}.$$

We introduce $\mathbf{c} = \mathbf{U}\mathbf{x}$, and it follows that the original problem can be solved in the two steps

$$\begin{aligned} 1^\circ & \quad \text{solve } \mathbf{L}\mathbf{c} = \mathbf{P}\mathbf{b}, \\ 2^\circ & \quad \text{solve } \mathbf{U}\mathbf{x} = \mathbf{c}. \end{aligned} \tag{2.21}$$

The two systems involve a lower and an upper triangular matrix, so their solution is found via forward and back substitution, respectively. The ones on the diagonal ensure that \mathbf{L} is nonsingular, so \mathbf{c} is unique. A singular system is reflected in zero diagonal elements in \mathbf{U} .

Example 2.18. The nonsingular matrix in Example 2.6 has the LU factorization

$$\mathbf{P} = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \quad \mathbf{L} = \begin{pmatrix} 1 & & \\ \frac{4}{9} & 1 & \\ \frac{1}{9} & 1 & 1 \end{pmatrix}, \quad \mathbf{U} = \begin{pmatrix} 9 & 3 & 1 \\ & \frac{6}{9} & \frac{2}{9} \\ & & \frac{3}{9} \end{pmatrix}.$$

The vector $\mathbf{P}\mathbf{b}$ is obtained by interchanging the first and third element in \mathbf{b} . We combine this with the forward substitution algorithm, Eq. (2.5), and exploit that all the diagonal elements are one, so there is no division involved:

$$\begin{aligned} c_1 &= 1.6, \\ c_2 &= 0.7 - \frac{4}{9} = \frac{-0.1}{9}, \\ c_3 &= 0.8 - \frac{1}{9} \cdot 1.6 - 1 \cdot \frac{-0.1}{9} = \frac{5.7}{9}. \end{aligned}$$

This \mathbf{c} is identical to the modified right-hand side in Example 2.6, which is no wonder: The calculations involved in step 1 $^\circ$ are the same as were used to get from $(\mathbf{A} \mid \mathbf{b})$ to $(\mathbf{U} \mid \mathbf{c})$, except that they are done in a different order. Step 2 $^\circ$ was performed already in Example 2.4. ■

Example 2.19. Algorithm 2.21 can be interpreted as an application of the relation treated in Exercise 2.14:

$$\mathbf{L}\mathbf{U}\mathbf{x} = \mathbf{P}\mathbf{b} \quad \Leftrightarrow \quad \mathbf{x} = (\mathbf{L}\mathbf{U})^{-1}\mathbf{P}\mathbf{b} = \mathbf{U}^{-1}\mathbf{L}^{-1}\mathbf{P}\mathbf{b} = \mathbf{U}^{-1}\mathbf{c}. \quad \blacksquare$$

Example 2.20. MATLAB has a built-in function `lu` for computing the LU factorization.

Let the MATLAB variable `A` contain the matrix \mathbf{A} , then the command

$$[\mathbf{L}, \mathbf{U}, \mathbf{P}] = \text{lu}(\mathbf{A})$$

returns \mathbf{L} , \mathbf{U} and \mathbf{P} in `L`, `U` and `P`, respectively. The command

$$[\mathbf{L}, \mathbf{U}] = \text{lu}(\mathbf{A})$$

returns $\mathbf{P}^T \mathbf{L}$ in `L` and \mathbf{U} in `U`. Thus, `L` contains a row permuted lower triangular matrix. We return to this aspect in Example 2.28 at the end of the the next section. ■

Example 2.21. It is relevant to ask the question: If $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $\mathbf{b} \in \mathbb{R}^n$, how much work is needed to solve the system $\mathbf{Ax} = \mathbf{b}$? The elimination algorithm Eq. (2.9) involves pivot searches, row interchanges and a number of “flops”. A *flop* (short for floating point operation) is a simple arithmetic operation (addition, subtraction, multiplication or division) between two real numbers. Traditionally, only the flops are counted, and it can be shown that the transformation from \mathbf{A} to \mathbf{U} involves

$$\left(\frac{2}{3}n^3 - \frac{1}{2}n^2 - \frac{1}{6}n\right) \text{ flops} .$$

For large n the first term dominates, and we say that the “cost” of the transformation from \mathbf{A} to \mathbf{U} is $\frac{2}{3}n^3$ flops. Note that this is also the cost of computing the LU factorization. The cost of solving each of the triangular systems $\mathbf{Lc} = \mathbf{Pb}$ and $\mathbf{Ux} = \mathbf{c}$ is n^2 flops. For comparison, the matrix-vector multiplication \mathbf{Az} costs $2n^2$ flops, ie the same as the solution of both triangular systems.

For a matrix equation $\mathbf{AX} = \mathbf{B}$ we need $\frac{2}{3}n^3$ flops to transform \mathbf{A} to \mathbf{U} and $2n^2$ flops per column in \mathbf{B} . If there are q columns, the solution cost is $\left(\frac{2}{3}n^3 + 2q \cdot n^2\right)$ flops. If $q = n$, then the cost is $\frac{8}{3}n^3$ flops. In the special case where $\mathbf{B} = \mathbf{I}$ it is possible to exploit the zeros in the off-diagonal positions, and the cost of computing \mathbf{A}^{-1} reduces to $2n^3$ flops.

For the problem in Example 2.15,

$$\mathbf{Ax}_k = \mathbf{b}_k, \quad k = 1, 2, \dots, K,$$

it follows that the solution via the matrix inverse costs $(2n^3 + 2Kn^2)$ flops. Instead we should compute the LU factorization once and then use Eq. (2.21) for each \mathbf{b}_k . With this method the cost reduces to $\left(\frac{2}{3}n^3 + 2Kn^2\right)$ flops, so we save $\frac{4}{3}n^3$ flops. ■

2.6. SPD systems

In an SPD system the $n \times n$ coefficient matrix \mathbf{A} is *symmetric* and *positive definite* (the matrix is an SPD matrix). This means that $\mathbf{A}^T = \mathbf{A}$ and

$$\mathbf{x}^T \mathbf{Ax} > 0 \quad \text{for all nonzero } \mathbf{x} \in \mathbb{R}^n . \quad (2.22)$$

Note that the left hand side is the scalar product between the vectors \mathbf{x} and \mathbf{Ax} , so it is a real number, and this is positive. Note also, that that this implies that $\mathbf{Ax} \neq \mathbf{0}$ for all nonzero \mathbf{x} , so the columns in \mathbf{A} are linearly independent and \mathbf{A} is nonsingular.

If \mathbf{A} is symmetric and

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

then \mathbf{A} is an SPSD (symmetric, positive semidefinite) matrix.

Example 2.22. In Chapter 4 we introduce the *normal equations* matrix

$$\mathbf{A} = \mathbf{F}^T \mathbf{F} ,$$

where \mathbf{F} is an $m \times n$ matrix with $m \geq n$. We immediately see that \mathbf{A} is an $n \times n$ matrix, and using Eq. (1.6) we get

$$\mathbf{A}^T = \mathbf{F}^T (\mathbf{F}^T)^T = \mathbf{F}^T \mathbf{F} = \mathbf{A} ,$$

so \mathbf{A} is symmetric. Next, we let $\mathbf{y} = \mathbf{F} \mathbf{x}$ and get

$$\mathbf{x}^T \mathbf{A} \mathbf{x} = \mathbf{x}^T \mathbf{F}^T \mathbf{F} \mathbf{x} = \mathbf{y}^T \mathbf{y} = y_1^2 + \cdots + y_m^2 \geq 0 .$$

This shows that \mathbf{A} is an SPSD matrix.

If the columns in \mathbf{F} are linearly independent, then $\mathbf{y} = \mathbf{F} \mathbf{x} \neq \mathbf{0}$ for all nonzero \mathbf{x} , and it follows that

$$\mathbf{x}^T \mathbf{A} \mathbf{x} = y_1^2 + \cdots + y_m^2 > 0 .$$

So in this case \mathbf{A} is an SPD matrix. ■

It can be shown that pivoting does not improve the accuracy of Gaussian elimination applied to an SPD system. If we omit this, there are no row interchanges, and formula (2.20) simplifies to

$$\mathbf{A} = \mathbf{L} \mathbf{U} .$$

Further, one can show that the symmetry is preserved in the active part of the matrix during the transformation, and this has the effect that $\mathbf{U} = \mathbf{D} \mathbf{L}^T$, where \mathbf{D} is the diagonal matrix given by the diagonal elements in \mathbf{U} . This means that Gaussian elimination applied to an SPD system is equivalent to the so-called *LDL factorization*

$$\mathbf{A} = \mathbf{L} \mathbf{D} \mathbf{L}^T . \tag{2.23}$$

Example 2.23. With the notation from for instance Example 2.6 we get

$$\begin{pmatrix} 4 & 12 & -8 & 16 \\ 12 & 45 & -33 & 66 \\ -8 & -33 & 41 & 30 \\ 16 & 66 & 30 & 502 \end{pmatrix} \begin{matrix} \sim \\ \ell_{21} = 3 \\ \ell_{31} = -2 \\ \ell_{41} = 4 \end{matrix} \begin{pmatrix} 4 & 12 & -8 & 16 \\ 0 & 9 & -9 & 18 \\ 0 & -9 & 25 & 62 \\ 0 & 18 & 62 & 438 \end{pmatrix} \begin{matrix} \sim \\ \ell_{32} = -1 \\ \ell_{42} = 2 \end{matrix} \\ \begin{pmatrix} 4 & 12 & -8 & 16 \\ 0 & 9 & -9 & 18 \\ 0 & 0 & 16 & 80 \\ 0 & 0 & 80 & 402 \end{pmatrix} \begin{matrix} \sim \\ \ell_{43} = 5 \end{matrix} \begin{pmatrix} 4 & 12 & -8 & 16 \\ 0 & 9 & -9 & 18 \\ 0 & 0 & 16 & 80 \\ 0 & 0 & 0 & 2 \end{pmatrix} .$$

We see that the active parts of the matrix are symmetric. Collecting the elimination factors in \mathbf{L} we get

$$\mathbf{L} = \begin{pmatrix} 1 & & & \\ 3 & 1 & & \\ -2 & -1 & 1 & \\ 4 & 2 & 5 & 1 \end{pmatrix} ,$$

and \mathbf{U} is the resulting upper triangular matrix,

$$\mathbf{U} = \begin{pmatrix} 4 & 12 & -8 & 16 \\ & 9 & -9 & 18 \\ & & 16 & 80 \\ & & & 2 \end{pmatrix} = \begin{pmatrix} 4 & & & \\ & 9 & & \\ & & 16 & \\ & & & 2 \end{pmatrix} \begin{pmatrix} 1 & 3 & -2 & 4 \\ & 1 & -1 & 2 \\ & & 1 & 5 \\ & & & 1 \end{pmatrix}.$$

We see that $\mathbf{U} = \mathbf{DL}^T$. ■

The symmetry can be exploited: We only need to compute the modified elements on the diagonal and in either the strictly lower or the strictly upper triangle of the matrix. Then we also know the modified elements in the other triangle. This means that the cost of computing the LDL factorization is only half the cost of computing the LU factorization, $\frac{1}{3}n^3$ flops.

The LDL factorization can be computed for any symmetric matrix. However, if the matrix is not positive definite, we should use pivoting in order to avoid excessive effects of rounding errors. Fortunately, it is simple to check the condition Eq. (2.22) during the computation of the factorization:

$$\begin{aligned} \mathbf{x}^T \mathbf{A} \mathbf{x} &= \mathbf{x}^T \mathbf{L} \mathbf{D} \mathbf{L}^T \mathbf{x} \\ &= \mathbf{y}^T \mathbf{D} \mathbf{y} = d_{11}y_1^2 + d_{22}y_2^2 + \cdots + d_{nn}y_n^2. \end{aligned}$$

Here, we introduced $\mathbf{y} = \mathbf{L}^T \mathbf{x}$. The ones along the diagonal in \mathbf{L} ensure that this matrix is nonsingular, and therefore $\mathbf{y} \neq \mathbf{0}$ for any nonzero \mathbf{x} . The right-hand side must be positive for any nonzero \mathbf{y} , and this is clearly the case if and only if all the $d_{ii} > 0$. This means that if we meet a $d_{kk} \leq 0$ during the computation of the factorization, then we know that the matrix is not SPD, and can stop.

Example 2.24. We change the second diagonal element in the matrix from Example 2.23 to $a_{22} = 35$ and get

$$\begin{pmatrix} 4 & 12 & -8 & 16 \\ 12 & 35 & -33 & 66 \\ -8 & -33 & 41 & 30 \\ 16 & 66 & 30 & 502 \end{pmatrix} \sim \begin{pmatrix} 4 & 12 & -8 & 16 \\ 0 & -1 & \star & \star \\ \star & \star & \star & \star \\ \star & \star & \star & \star \end{pmatrix} \begin{matrix} \\ \ell_{21} = 3 \\ \ell_{31} = \star \\ \ell_{41} = \star \end{matrix}$$

The \star indicates values that we do not need to compute. $d_{22} = -1 \leq 0$ tells us that the matrix is not positive definite, and we stop the computation. A system $\mathbf{Ax} = \mathbf{b}$ with this matrix should be solved by Gaussian elimination with pivoting. ■

An SPD system is often solved via the *Cholesky factorization* of the matrix,

$$\mathbf{A} = \mathbf{C} \mathbf{C}^T, \tag{2.24}$$

where \mathbf{C} is a lower triangular matrix. We can derive this from the LDL factorization,

$$\mathbf{A} = \mathbf{L} \mathbf{D} \mathbf{L}^T = \left(\mathbf{L} \mathbf{D}^{1/2} \right) \left(\mathbf{D}^{1/2} \mathbf{L}^T \right) = \mathbf{C} \mathbf{C}^T.$$

Here, we have introduced the matrix

$$\mathbf{D}^{1/2} = \begin{pmatrix} \sqrt{d_{11}} & & & \\ & \sqrt{d_{22}} & & \\ & & \ddots & \\ & & & \sqrt{d_{nn}} \end{pmatrix}.$$

Since the d_{ii} are positive, the square roots exist.

Example 2.25. For the matrix from Example 2.23 the Cholesky factor is

$$\mathbf{C} = \begin{pmatrix} 1 & & & \\ 3 & 1 & & \\ -2 & -1 & 1 & \\ 4 & 2 & 5 & 1 \end{pmatrix} \begin{pmatrix} 2 & & & \\ & 3 & & \\ & & 4 & \\ & & & \sqrt{2} \end{pmatrix} = \begin{pmatrix} 2 & & & \\ 6 & 3 & & \\ -4 & -3 & 4 & \\ 8 & 6 & 20 & \sqrt{2} \end{pmatrix} .$$

■

The Cholesky factor may be computed directly, ie without the detour around the LDL factorization.

Once the Cholesky factorization is known, the SPD system

$$\mathbf{A} \mathbf{x} = \mathbf{b} \quad \Leftrightarrow \quad \mathbf{C} \mathbf{C}^T \mathbf{x} = \mathbf{b}$$

can be solved in the two steps

$$\begin{aligned} 1^\circ & \quad \text{solve } \mathbf{C} \mathbf{z} = \mathbf{b} , \\ 2^\circ & \quad \text{solve } \mathbf{C}^T \mathbf{x} = \mathbf{z} . \end{aligned} \tag{2.25}$$

The solution of these two systems is found via forward and back substitution, respectively.

Example 2.26. Consider the system

$$\begin{pmatrix} 4 & 12 & -8 & 16 \\ 12 & 45 & -33 & 66 \\ -8 & -33 & 41 & 30 \\ 16 & 66 & 30 & 502 \end{pmatrix} \mathbf{x} = \begin{pmatrix} 108 \\ 441 \\ -61 \\ 2034 \end{pmatrix} .$$

In Examples 2.23 and 2.25 we found that the symmetric coefficient matrix is positive definite and computed its Cholesky factor \mathbf{C} . The solution to the system is found in the two steps

$$\begin{pmatrix} 2 & & & \\ 6 & 3 & & \\ -4 & -3 & 4 & \\ 8 & 6 & 20 & \sqrt{2} \end{pmatrix} \mathbf{z} = \begin{pmatrix} 108 \\ 441 \\ -61 \\ 2034 \end{pmatrix} \quad \Rightarrow \quad \mathbf{z} = \begin{pmatrix} 54 \\ 39 \\ 684\sqrt{2} \end{pmatrix} ,$$

$$\begin{pmatrix} 2 & 6 & -4 & 8 \\ & 3 & -3 & 6 \\ & & 4 & 20 \\ & & & \sqrt{2} \end{pmatrix} \mathbf{x} = \begin{pmatrix} 54 \\ 39 \\ 684\sqrt{2} \end{pmatrix} \quad \Rightarrow \quad \mathbf{x} = \begin{pmatrix} -1 \\ 2 \\ -3 \\ 4 \end{pmatrix} .$$

■

Example 2.27. Traditionally the Cholesky factorization is defined as in Eq. (2.24), $\mathbf{A} = \mathbf{C} \mathbf{C}^T$, where \mathbf{C} is lower triangular. We might as well define it as

$$\mathbf{A} = \mathbf{R}^T \mathbf{R} ,$$

where $\mathbf{R} = \mathbf{C}^T$ is upper triangular. This is the definition used in MATLAB. The built-in function `chol` can be used to compute the Cholesky factor. Let \mathbf{A} contain a square matrix, then

$$\mathbf{R} = \text{chol}(\mathbf{A})$$

attempts to compute the Cholesky factor for the matrix. The matrix is assumed to be symmetric and only the elements in the upper triangle are considered, for instance

$$R = \text{chol}([4 \ 2; 0 \ 10]) \quad R = \text{chol}([4 \ 2; 2 \ 10])$$

both give $R = [2 \ 1; 0 \ 3]$. If the matrix is not positive definite, you get an error return. ■

Example 2.28. Let the MATLAB variables \mathbf{A} and \mathbf{b} contain respectively an $n \times n$ matrix \mathbf{A} and column vector \mathbf{b} with n elements. Then the MATLAB command

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

returns the solution to the system $\mathbf{Ax} = \mathbf{b}$. The method depends on possible special properties of \mathbf{A} :

- 1° If \mathbf{A} is lower triangular (or a permuted lower triangular matrix, cf Example 2.20), then the system is solved via (permuted) forward substitution. Similarly a (permuted) upper triangular system is solved by (permuted) back substitution.
- 2° If \mathbf{A} is symmetric, then a Cholesky factorization of it is attempted. If it succeeds, then it is used in the two step algorithm Eq. (2.25) to compute the solution. Otherwise the general method is employed.
- 3° General method: Compute the LU factorization of \mathbf{A} , and use Eq. (2.21) to solve the system. ■

3. Eigenvalues and Eigenvectors

Example 3.1. Let λ be a constant. The simple differential equation

$$\frac{dy}{dt} = \lambda y$$

has the solution

$$y(t) = c e^{\lambda t},$$

where c is an arbitrary constant. If $y(0) = \beta$ is known, then $c = \beta$ gives the desired solution.

A linear system of differential equations has the form

$$\frac{d\mathbf{y}}{dt} = \mathbf{A} \mathbf{y},$$

where each element in \mathbf{y} is a function of time and the elements in the square matrix \mathbf{A} are constant. It is tempting to postulate that the solution to this system is

$$\mathbf{y}(t) = e^{\mathbf{A}t} \mathbf{c},$$

where \mathbf{c} is an arbitrary vector. Is there an interpretation of the “*matrix exponential*” $e^{\mathbf{A}t}$ that makes this postulate true ?

Some other interesting questions are

- How did we find a good change of basis for the ellipse problem in Example 1.20 ?
 - Why do the most interesting links always come out close to the top in a *Google* search ?
 - For a given $n \times n$ mapping matrix \mathbf{A} , is there a nonzero vector \mathbf{v} whose image $\mathbf{A}\mathbf{v}$ is proportional to \mathbf{v} ?
-

Surprisingly, the key question is the last one. This is the so-called *eigenproblem*

$$\mathbf{A} \mathbf{v} = \lambda \mathbf{v}, \tag{3.1}$$

where the scalar λ is the *eigenvalue* and the nonzero \mathbf{v} is a corresponding *eigenvector*. We say that the pair (λ, \mathbf{v}) is an *eigensolution*.

3.1. Properties of eigensolutions

The definition $\mathbf{A}\mathbf{v} = \lambda\mathbf{v}$ is equivalent to

$$(\mathbf{A} - \lambda\mathbf{I})\mathbf{v} = \mathbf{0}. \tag{3.2}$$

We look for a nonzero solution to this homogeneous system of equations, and it follows from the discussion in Section 2.3 that the matrix $\mathbf{A} - \lambda\mathbf{I}$ must be singular. Especially, if $\lambda = 0$ is an eigenvalue for \mathbf{A} , then the matrix itself is singular.

Example 3.2. For a 2×2 matrix \mathbf{A} we can use the *determinant* defined in Example 2.2 to express that $\mathbf{A} - \lambda\mathbf{I}$ must be singular:

$$\begin{aligned} 0 = \det(\mathbf{A} - \lambda\mathbf{I}) &= \begin{vmatrix} a_{11} - \lambda & a_{12} \\ a_{21} & a_{22} - \lambda \end{vmatrix} \\ &= (a_{11} - \lambda)(a_{22} - \lambda) - a_{21}a_{12} \\ &= \lambda^2 - (a_{11} + a_{22})\lambda + a_{11}a_{22} - a_{21}a_{12} . \end{aligned}$$

This means that the eigenvalues are the roots of the so-called *characteristic polynomial*

$$P_A(\lambda) = \lambda^2 - (a_{11} + a_{22})\lambda + a_{11}a_{22} - a_{21}a_{12} .$$

We shall consider three specific examples,

$$\mathbf{B} = \begin{pmatrix} 73 & -36 \\ -36 & 52 \end{pmatrix}, \quad \mathbf{C} = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}, \quad \mathbf{H} = \begin{pmatrix} 73 & -36 \\ 36 & 52 \end{pmatrix} .$$

With the first matrix we get

$$P_B(\lambda) = \lambda^2 - 125\lambda + 2500 .$$

Using a well known technique for solving the quadratic equation $P_B(\lambda) = 0$, we first find the discriminant

$$D = (-125)^2 - 4 \cdot 1 \cdot 2500 = 5625 ,$$

and the roots are

$$\lambda_1 = \frac{125 + \sqrt{D}}{2} = 100, \quad \lambda_2 = \frac{125 - \sqrt{D}}{2} = 25 .$$

In order to find an eigenvector corresponding to λ_2 we solve the singular system $(\mathbf{B} - 25\mathbf{I})\mathbf{v}_2 = \mathbf{0}$. With the notation from Section 2.2 we get

$$\left(\begin{array}{cc|c} 48 & -36 & 0 \\ -36 & 27 & 0 \end{array} \right) \begin{array}{l} \sim \\ p_1 = 1 \\ \ell_{21} = \frac{-36}{48} \end{array} \left(\begin{array}{cc|c} 48 & -36 & 0 \\ 0 & 0 & 0 \end{array} \right) .$$

The reduced system has $u_{22} = 0$, confirming that $\mathbf{B} - 25\mathbf{I}$ is singular. Proceeding as in Section 2.3 we find the complete solution

$$\mathbf{v}_2 = \alpha \begin{pmatrix} 3 \\ 4 \end{pmatrix}, \quad \alpha \in \mathbb{R} .$$

Similarly we find that

$$\mathbf{v}_1 = \alpha \begin{pmatrix} -4 \\ 3 \end{pmatrix}, \quad \alpha \in \mathbb{R} ,$$

are the eigenvectors corresponding to $\lambda_1 = 100$.

Next, we find

$$P_C(\lambda) = (1 - \lambda)^2 .$$

The only root is $\lambda = 1$. We say that this eigenvalue has *algebraic multiplicity 2*. The eigenvectors are $(\alpha \ 0)^T$ for any $\alpha \in \mathbb{R}$.

Finally,

$$P_H(\lambda) = \lambda^2 - 125\lambda + 5092 .$$

The discriminant is

$$D = (-125)^2 - 4 \cdot 1 \cdot 5092 = -4743 .$$

This is negative, so \mathbf{H} has no eigenvalue $\lambda \in \mathbb{R}$. ■

Example 3.3. The condition that $\mathbf{A} - \lambda \mathbf{I}$ be singular has the immediate consequence that the diagonal elements in a triangular (especially a diagonal) matrix are eigenvalues of that matrix. For instance

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

has the eigenvalues $\lambda_1 = 1$, $\lambda_2 = 4$ and $\lambda_3 = 6$.

The $n \times n$ identity matrix \mathbf{I} gives $\mathbf{I}\mathbf{v} = \mathbf{v}$ for any $\mathbf{v} \in \mathbb{R}^n$, cf Example 1.10. This means that \mathbf{I} has the eigenvalue $\lambda = 1$ with arithmetic multiplicity n , and any $\mathbf{v} \in \mathbb{R}^n$ is an eigenvector. ■

These examples show that if \mathbf{A} has an eigensolution, then the eigenvector is not unique: From the linearity conditions in Eq. (1.7) and the above definition of the eigenproblem it follows that

$$\mathbf{A}(\alpha\mathbf{v}) = \alpha(\mathbf{A}\mathbf{v}) = \alpha(\lambda\mathbf{v}) = \lambda(\alpha\mathbf{v}) ,$$

so $\alpha\mathbf{v}$ is also an eigenvector for any choice of the scalar α . Therefore it may be useful to think of an eigenvector as a direction in \mathbb{R}^n , cf the geometric view in Example 1.1. When we talk about “the eigenvector” in the following, we assume that it is normalized so that $\|\mathbf{v}\| = 1$. This makes \mathbf{v} unique, except for sign.

We saw in Examples 3.2 and 3.3 that a matrix \mathbf{A} may have several different eigenvalues λ_j with corresponding eigenvectors \mathbf{v}_j ,

$$\mathbf{A}\mathbf{v}_j = \lambda_j\mathbf{v}_j , \quad j = 1, 2, \dots, p . \tag{3.3}$$

Further, it may happen (as for instance with \mathbf{I}) that there are linearly independent eigenvectors corresponding to the same eigenvalue. This case is also covered by the formulation (3.3). If $\lambda_i = \lambda_j$, then

$$\mathbf{A}(\alpha\mathbf{v}_i + \beta\mathbf{v}_j) = \alpha(\lambda_i\mathbf{v}_i) + \beta(\lambda_j\mathbf{v}_j) = \lambda_i(\alpha\mathbf{v}_i + \beta\mathbf{v}_j)$$

for any choice of scalars α and β . This shows that any linear combination of eigenvectors corresponding to a *multiple eigenvalue* is also an eigenvector.

Without proof we give the following facts about the eigenproblem for an $n \times n$ matrix \mathbf{A} . We shall use 5° in Section 3.2.3.

- 1° \mathbf{A} has at most n different eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_p$, $p \leq n$.
- 2° If $\lambda_1, \lambda_2, \dots, \lambda_p$ are mutually different, then the corresponding eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p$ are linearly independent.
- 3° If \mathbf{A} is symmetric and $\lambda_i \neq \lambda_j$ then \mathbf{v}_i and \mathbf{v}_j are orthogonal.

- 4° If \mathbf{A} is symmetric, then it is possible to choose an orthonormal set of n eigenvectors for \mathbf{A} .
- 5° If all the elements in \mathbf{A} are positive and the sum of the elements in each column is one, then $\lambda = 1$ is an eigenvalue for \mathbf{A} .

Example 3.4. Let

$$\mathbf{A} = \begin{pmatrix} 2 & -1 & -2 \\ -1 & 2 & -2 \\ -2 & -2 & -1 \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} 15 & -6 & -12 \\ 12 & -3 & -12 \\ 12 & -6 & -9 \end{pmatrix}.$$

Both matrices have the eigenvalues $\lambda_1 = -3$ and $\lambda_2 = \lambda_3 = 3$. For the symmetric matrix \mathbf{A} we may choose the orthonormal eigenvectors

$$\mathbf{v}_1 = \frac{1}{\sqrt{6}} \begin{pmatrix} 1 \\ 1 \\ 2 \end{pmatrix}, \quad \mathbf{v}_2 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}, \quad \mathbf{v}_3 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}.$$

The normalized eigenvectors for \mathbf{B} may be chosen as

$$\mathbf{v}_1 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \quad \mathbf{v}_2 = \frac{1}{\sqrt{5}} \begin{pmatrix} 1 \\ 2 \\ 0 \end{pmatrix}, \quad \mathbf{v}_3 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}. \quad \blacksquare$$

Now, assume that \mathbf{A} has n linearly independent eigenvectors, and make a matrix \mathbf{V} with \mathbf{v}_j as its j^{th} column,

$$\mathbf{V} = (\mathbf{v}_1 \quad \mathbf{v}_2 \quad \cdots \quad \mathbf{v}_n). \quad (3.4)$$

Then

$$\mathbf{A}\mathbf{V} = (\lambda_1\mathbf{v}_1 \quad \lambda_2\mathbf{v}_2 \quad \cdots \quad \lambda_n\mathbf{v}_n) = \mathbf{V}\mathbf{\Lambda}, \quad (3.5)$$

where $\mathbf{\Lambda}$ is the diagonal matrix with the eigenvalues,

$$\mathbf{\Lambda} = \begin{pmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_n \end{pmatrix}.$$

In Eq. (3.5) we used the column-wise formulation of a matrix–matrix product, Eq. (1.5), and the result from Exercise 1.13 about the product of a general and a diagonal matrix.

The assumption about the columns in \mathbf{V} being linearly independent, implies that \mathbf{V} is nonsingular. Therefore \mathbf{V}^{-1} exists, and if we multiply with this matrix from the right on both sides in Eq. (3.5), we get

$$\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}. \quad (3.6)$$

This *eigen factorization* of \mathbf{A} is essential for the applications that we give in the next section.

If \mathbf{A} is symmetric, then the assumption about n linearly independent eigenvectors is satisfied, and if we choose them to be orthonormal, cf fact 4° above, then \mathbf{V} is an orthogonal matrix, and Eq. (3.6) simplifies to

$$\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T. \quad (3.7)$$

Example 3.5. We return to the question from Section 2.6: Let \mathbf{A} be symmetric; is it also *positive definite* ?

To answer that question we may use the factorization in Eq. (3.7).

$$\begin{aligned}\mathbf{x}^T \mathbf{A} \mathbf{x} &= \mathbf{x}^T \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T \mathbf{x} \\ &= \mathbf{y}^T \mathbf{\Lambda} \mathbf{y} = \lambda_1 y_1^2 + \lambda_2 y_2^2 + \cdots + \lambda_n y_n^2.\end{aligned}$$

Here, we introduced $\mathbf{y} = \mathbf{V}^T \mathbf{x}$. Since \mathbf{V} is nonsingular, there exists a nonzero \mathbf{y} corresponding to any nonzero \mathbf{x} . The right-hand side must be positive for any nonzero \mathbf{y} , and this is clearly the case if and only if all the $\lambda_i > 0$.

Therefore, if all the eigenvalues for a symmetric matrix are strictly positive, then it is an SPD matrix. It should be said, that it is computationally cheaper to check via the LDL^T or Cholesky factorization, as described in Section 2.6. ■

Example 3.6. MATLAB has a built-in function `eig` for computing eigensolutions. Let \mathbf{A} be a square matrix, then the command

$$\mathbf{la} = \text{eig}(\mathbf{A})$$

returns a vector `la` with the eigenvalues of \mathbf{A} . The command

$$[\mathbf{V}, \mathbf{La}] = \text{eig}(\mathbf{A})$$

returns the arrays `V` and `La`, representing respectively \mathbf{V} and $\mathbf{\Lambda}$ in Eq. (3.6), or Eq. (3.7) if \mathbf{A} is symmetric.

The algorithm used is based on the following relation: If \mathbf{S} is a nonsingular matrix, then

$$\mathbf{A} \mathbf{v} = \lambda \mathbf{v} \quad \Leftrightarrow \quad \mathbf{S} \mathbf{A} \mathbf{S}^{-1} \mathbf{S} \mathbf{v} = \lambda \mathbf{S} \mathbf{v}.$$

The transformation from \mathbf{A} to $\tilde{\mathbf{A}} = \mathbf{S} \mathbf{A} \mathbf{S}^{-1}$ is called a *similarity transformation*. We see that it preserves the eigenvalues, but changes the eigenvectors. If we know the eigensolutions and if the eigenvectors are linearly independent, then it follows from Eq. (3.6) that if we take $\mathbf{S} = \mathbf{V}^{-1}$, then $\tilde{\mathbf{A}} = \mathbf{\Lambda}$.

We do not know the eigensolutions, but it is possible to compute a sequence of similarity transformations in such a way that the sequence of $\tilde{\mathbf{A}}$ -matrices converge to an upper triangular matrix \mathbf{U} . According to Example 3.3 the diagonal elements in \mathbf{U} are the eigenvalues, and it is fairly easy to compute the eigenvectors of \mathbf{U} and transform them back to eigenvectors of \mathbf{A} .

The algorithm used in MATLAB uses orthogonal matrices for the similarity transformations, ie $\tilde{\mathbf{A}} = \mathbf{Q} \mathbf{A} \mathbf{Q}^T$. In case \mathbf{A} is symmetric, the symmetry is preserved, and the sequence converges to a diagonal matrix. ■

3.2. Applications of eigensolutions

In this section we shall answer the first three questions raised in Example 3.1.

3.2.1. Differential equations

First, we look at a model that involves n functions of t , $y_1(t)$, $y_2(t)$, \dots , $y_n(t)$. We

do not know closed form expressions for the functions, but know their values at $t = 0$,

$$y_i(0) = \beta_i, \quad i = 1, 2, \dots, n.$$

Further, we know that for $t > 0$ the functions satisfy a linear system of coupled differential equations,

$$\begin{aligned} y_1' &= a_{11}y_1 + a_{12}y_2 + \cdots + a_{1n}y_n \\ y_2' &= a_{21}y_1 + a_{22}y_2 + \cdots + a_{2n}y_n \\ &= \vdots \\ y_n' &= a_{n1}y_1 + a_{n2}y_2 + \cdots + a_{nn}y_n \end{aligned}.$$

We introduce the vector function $\mathbf{y}(t)$ and the vector $\boldsymbol{\beta}$ with the i^{th} elements $y_i(t)$ and β_i , respectively. Further, we let the $n \times n$ matrix \mathbf{A} hold the coefficients a_{ij} , and see that the problem can be written in the form

$$\mathbf{y}'(t) = \mathbf{A} \mathbf{y}(t) \quad \text{for } t > 0, \quad \mathbf{y}(0) = \boldsymbol{\beta}. \quad (3.8)$$

We want to solve this *initial value problem* for a *linear system of differential equations*.

The solution is quite easy to find with the help of the eigensolutions of \mathbf{A} , provided that there are n linearly independent eigenvectors, $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$. In that case we can use them as basis in \mathbb{R}^n , ie we can write

$$\mathbf{y}(t) = \tilde{y}_1(t)\mathbf{v}_1 + \tilde{y}_2(t)\mathbf{v}_2 + \cdots + \tilde{y}_n(t)\mathbf{v}_n = \mathbf{V} \tilde{\mathbf{y}}(t), \quad (3.9)$$

where \mathbf{V} is the matrix with the eigenvectors as columns, cf Eq. (3.4), and the vector $\tilde{\mathbf{y}}$ has \tilde{y}_i as the i^{th} element. By differentiation we get

$$\mathbf{y}'(t) = \tilde{y}_1'(t)\mathbf{v}_1 + \tilde{y}_2'(t)\mathbf{v}_2 + \cdots + \tilde{y}_n'(t)\mathbf{v}_n = \mathbf{V} \tilde{\mathbf{y}}'(t).$$

We insert these relations in the differential equation and use the factorization Eq. (3.6) for \mathbf{A} ,

$$\mathbf{V} \tilde{\mathbf{y}}'(t) = \mathbf{A} \mathbf{V} \tilde{\mathbf{y}}(t) = (\mathbf{V} \boldsymbol{\Lambda} \mathbf{V}^{-1}) \mathbf{V} \tilde{\mathbf{y}}(t).$$

So the differential equation is equivalent to

$$\tilde{\mathbf{y}}'(t) = \boldsymbol{\Lambda} \tilde{\mathbf{y}}(t).$$

The starting values are given by $\mathbf{V} \tilde{\mathbf{y}}(0) = \boldsymbol{\beta}$, or

$$\tilde{\mathbf{y}}(0) = \tilde{\boldsymbol{\beta}} \quad \text{with} \quad \tilde{\boldsymbol{\beta}} = \mathbf{V}^{-1} \boldsymbol{\beta}.$$

We see that this change of variables transforms the original system of coupled differential equations to a system of scalar differential equations,

$$\begin{aligned} &\text{for } j = 1, 2, \dots, n \\ \tilde{y}_j'(t) &= \lambda_j \tilde{y}_j(t) \quad \text{for } t > 0, \quad \tilde{y}_j(0) = \tilde{\beta}_j. \end{aligned}$$

The solution is

$$\tilde{y}_j(t) = \tilde{\beta}_j e^{\lambda_j t}, \quad j = 1, 2, \dots, n.$$

In vector form we can express this as

$$\tilde{\mathbf{y}}(t) = [\tilde{\boldsymbol{\beta}}]\boldsymbol{\varphi}(t), \quad (3.10)$$

where

$$\boldsymbol{\varphi}(t) = \begin{pmatrix} e^{\lambda_1 t} \\ e^{\lambda_2 t} \\ \vdots \\ e^{\lambda_n t} \end{pmatrix}$$

and $[\mathbf{x}]$ denotes the diagonal matrix with diagonal elements given by the vector \mathbf{x} ,

$$[\mathbf{x}] = \begin{pmatrix} x_1 & & & \\ & x_2 & & \\ & & \ddots & \\ & & & x_n \end{pmatrix}. \quad (3.11)$$

Finally, by combining Eqs. (3.9) – (3.11) we see that the components of the solution can be expressed in the form

$$y_i(t) = \alpha_{i1}e^{\lambda_1 t} + \alpha_{i2}e^{\lambda_2 t} + \cdots + \alpha_{in}e^{\lambda_n t}, \quad \text{with } \alpha_{ij} = v_{ij}\tilde{\beta}_j. \quad (3.12)$$

Note that every component in both $\tilde{\mathbf{y}}(t)$ and $\mathbf{y}(t)$ is a linear combination of the exponential functions $e^{\lambda_j t}$, where the λ_j are the eigenvalues of \mathbf{A} . The coefficients in the linear combinations depend on the initial values and on the eigenvectors of \mathbf{A} .

Example 3.7. The drug concentration problem in Example 1.2 can be modelled by a *compartment model*. This has the form of Eq. (3.8), with n being the number of compartments.

As an example consider the initial value problem

$$\mathbf{y}'(t) = \begin{pmatrix} -0.18 & 0.15 \\ 0.024 & -0.27 \end{pmatrix} \mathbf{y}(t) \quad \text{for } t > 0, \quad \mathbf{y}(0) = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

The eigensolutions of the matrix are

$$\lambda_1 = -0.15, \quad \mathbf{v}_1 = \begin{pmatrix} 0.981 \\ 0.196 \end{pmatrix}, \quad \lambda_2 = -0.30, \quad \mathbf{v}_2 = \begin{pmatrix} -0.781 \\ 0.625 \end{pmatrix}.$$

We used the MATLAB function `eig` from Example 3.6 to compute the eigensolutions, and display the results with at most three decimals. Going through the steps of computing $\tilde{\boldsymbol{\beta}}$ and the coefficients α_{ij} , we find that the two components of $\mathbf{y}(t)$ are

$$y_1(t) = e^{-0.15t} - e^{-0.30t}, \quad y_2(t) = 0.2e^{-0.15t} + 0.8e^{-0.30t}.$$




Figure 3.1. Two components of $\mathbf{y}(t)$.

The points in figure 1.2 have coordinates $(t_r, y_1(t_r))$ for 17 different values of t_r . ■

Example 3.8. In Example 3.1 we asked the question: If we want to write the solution to $\mathbf{y}' = \mathbf{A}\mathbf{y}$ as $\mathbf{y}(t) = e^{\mathbf{A}t}\mathbf{c}$, how shall we interpret $e^{\mathbf{A}t}$?

The answer is not easy to see from Eq. (3.12). However, we can replace Eq. (3.10) by the equivalent formula

$$\tilde{\mathbf{y}}(t) = [\varphi(t)]\tilde{\boldsymbol{\beta}} = [\varphi(t)]\mathbf{V}^{-1}\boldsymbol{\beta} .$$

Then $\mathbf{y}(t) = \mathbf{V}\tilde{\mathbf{y}}(t)$ leads to $\mathbf{y}(t) = \mathbf{V}[\varphi(t)]\mathbf{V}^{-1}\boldsymbol{\beta}$. This implies that a valid interpretation is

$$e^{\mathbf{A}t} = \mathbf{V}[\varphi(t)]\mathbf{V}^{-1} ,$$

where $[\varphi(t)]$ is the diagonal matrix with $([\varphi(t)])_{ii} = e^{\lambda_i t}$ and \mathbf{V} is the matrix with $\mathbf{V}_{:,i} = \mathbf{v}_i$. ■

3.2.2. Simplification of quadratics

This application also relies on using eigenvectors as basis. In Example 1.20 we asked a question of the form: which points (x_1, x_2) satisfy the equation $q(\mathbf{x}) = 0$, where $q(\mathbf{x})$ is the quadratic

$$q(\mathbf{x}) = c_1x_1^2 + c_2x_2^2 + c_3x_1x_2 + c_4x_1 + c_5x_2 + c_6 .$$

It simplifies the problem if we can get rid of the mixed term $c_3x_1x_2$ involving the product of different elements in \mathbf{x} . To achieve that we introduce the matrix \mathbf{A} and vector \mathbf{b} ,

$$\mathbf{A} = \begin{pmatrix} c_1 & \frac{1}{2}c_3 \\ \frac{1}{2}c_3 & c_2 \end{pmatrix} , \quad \mathbf{b} = \begin{pmatrix} c_4 \\ c_5 \end{pmatrix} . \quad (3.13)$$

It follows that

$$\begin{aligned} \mathbf{x}^T \mathbf{A} \mathbf{x} &= \mathbf{x}^T \begin{pmatrix} c_1x_1 + \frac{1}{2}c_3x_2 \\ \frac{1}{2}c_3x_1 + c_2x_2 \end{pmatrix} = c_1x_1^2 + c_2x_2^2 + c_3x_1x_2 , \\ \mathbf{b}^T \mathbf{x} &= c_4x_1 + c_5x_2 , \end{aligned}$$

so that

$$\begin{aligned} q(\mathbf{x}) &= \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x} + c_6 \\ &= \mathbf{x}^T \mathbf{V} \boldsymbol{\Lambda} \mathbf{V}^T \mathbf{x} + \mathbf{b}^T \mathbf{x} + c_6 \\ &= \tilde{\mathbf{x}}^T \boldsymbol{\Lambda} \tilde{\mathbf{x}} + \tilde{\mathbf{b}}^T \tilde{\mathbf{x}} + c_6 \\ &= \lambda_1 \tilde{x}_1^2 + \lambda_2 \tilde{x}_2^2 + \tilde{b}_1 \tilde{x}_1 + \tilde{b}_2 \tilde{x}_2 + c_6 . \end{aligned} \quad (3.14)$$

Here we used the factorization Eq. (3.7) for the symmetric matrix \mathbf{A} , and introduced

$$\tilde{\mathbf{x}} = \mathbf{V}^T \mathbf{x} \Leftrightarrow \mathbf{x} = \mathbf{V} \tilde{\mathbf{x}} , \quad \tilde{\mathbf{b}} = \mathbf{V}^T \mathbf{b} \Leftrightarrow \mathbf{b} = \mathbf{V} \tilde{\mathbf{b}} ,$$

so the elements in $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{b}}$ are the coordinates of \mathbf{x} and \mathbf{b} with respect to the basis given by the eigenvectors of \mathbf{A} . The change of variables makes it easier to discuss the solution to the equation $q(\mathbf{x}) = 0$. We shall illustrate that with two examples.

Example 3.9. The quadratic in Example 1.20 has the coefficients

$$\mathbf{c} = (73 \ 52 \ -72 \ 0 \ 0 \ -100)^T .$$

The corresponding matrix is

$$\mathbf{A} = \begin{pmatrix} 73 & -36 \\ -36 & 52 \end{pmatrix} ,$$

which has the eigensolutions

$$\lambda_1 = 25 , \quad \mathbf{v}_1 = \begin{pmatrix} 0.6 \\ 0.8 \end{pmatrix} , \quad \lambda_2 = 100 , \quad \mathbf{v}_2 = \begin{pmatrix} -0.8 \\ 0.6 \end{pmatrix} .$$

Therefore, Eq. (3.14) leads to

$$q(\mathbf{x}) = 25\tilde{x}_1^2 + 100\tilde{x}_2^2 - 100 = 0 .$$

This agrees with the equation that we derived in Example 1.20, using some rather tedious calculations. ■

Example 3.10. For the quadratic

$$q(\mathbf{x}) = 18x_1^2 + 32x_2^2 + 48x_1x_2 - 40x_1 - 95x_2 + 37.5$$

we use Eq. (3.13) to get

$$\mathbf{A} = \begin{pmatrix} 18 & 24 \\ 24 & 32 \end{pmatrix} , \quad \mathbf{b} = \begin{pmatrix} -40 \\ -95 \end{pmatrix} .$$

The eigensolutions of \mathbf{A} are (presented in the form defined by Eq. (3.5))

$$\mathbf{V} = \begin{pmatrix} 0.6 & -0.8 \\ 0.8 & 0.6 \end{pmatrix} , \quad \mathbf{\Lambda} = \begin{pmatrix} 50 & \\ & 0 \end{pmatrix} .$$

So $\lambda_1 = 50$, $\lambda_2 = 0$, and the eigenvectors are the same as in the previous example. Further,

$$\tilde{\mathbf{b}} = \mathbf{V}^T \mathbf{b} = \begin{pmatrix} -100 \\ -25 \end{pmatrix} ,$$

and Eq. (3.14) takes the form

$$q(\mathbf{x}) = 50\tilde{x}_1^2 - 100\tilde{x}_1 - 25\tilde{x}_2 + 37.5 .$$

Therefore, the equation $q(\mathbf{x}) = 0$ is equivalent to

$$\tilde{x}_2 = 2\tilde{x}_1^2 - 4\tilde{x}_1 + 1.5 = 2(\tilde{x}_1 - 1)^2 - 0.5 .$$

This is the equation for the parabola shown in Figure 3.2 below.




Figure 3.2. Rotated parabola. ■

3.2.3. Google's PageRank

Finally, we give a simplified description of the PageRank principle used in Google. Figure 3.3 shows a very small web with only $n = 6$ pages and 10 *links*, indicated by arrows. For instance there are two links from page 4: to pages 5 and 6.




Figure 3.3. Web with $n = 6$ pages.

Let page i have the PageRank p_i . This number should reflect the importance of that page in such a way that it is better for the page to be pointed at from a page with a large PageRank than from a less important page. In order to quantify the PageRank we set up a *model* that simulates a web-surfer.

Suppose that the surfer is at page j , which has $q = q_j$ *out-links*. The model says that the surfer may either follow one of these links or choose a new page at random. In both cases it is assumed that all possible candidates for being the next page have the same probability to be chosen. In the follow-a-link case this means that each of the q target pages “inherits” the PageRank contribution $\frac{1}{q}p_j$. This can be represented by an $n \times n$ matrix \mathbf{H} , where all elements are zero, except for $h_{ij} = \frac{1}{q}$ for the pages pointed at from the j^{th} page, which has q out-links. For instance

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

is the \mathbf{H} -matrix for the web in Figure 3.3.

The web may contain “*dangling pages*”, as for instance page 2 in Figure 3.3, from which there are no out-links. From such a page the surfer chooses any of the n pages at random. This corresponds to a modified matrix \mathbf{S} , which is equal to \mathbf{H} except for the columns where all elements are zero. These columns are replaced by $\frac{1}{n}\mathbf{e}$, where \mathbf{e} is the vector of all ones. The \mathbf{S} -matrix corresponding to the above \mathbf{H} is

$$\mathbf{S} = \begin{pmatrix} 0 & \frac{1}{6} & \frac{1}{3} & 0 & 0 & 0 \\ \frac{1}{2} & \frac{1}{6} & \frac{1}{3} & 0 & 0 & 0 \\ \frac{1}{2} & \frac{1}{6} & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{6} & \frac{1}{3} & 0 & 0 & \frac{1}{2} \\ 0 & \frac{1}{6} & 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & \frac{1}{6} & 0 & \frac{1}{2} & 1 & 0 \end{pmatrix}.$$

Finally, the surfer may occasionally want to pick a new page at random. If this was done from every page, the corresponding \mathbf{S} -matrix would be $\frac{1}{n}\mathbf{E}$, where \mathbf{E} is the matrix of all ones. However, let us assume this random choice has probability $1-\alpha$. Then the matrix to use is

$$\mathbf{G} = \alpha\mathbf{S} + \frac{1-\alpha}{n}\mathbf{E}. \quad (3.15)$$

This is known as the *Google matrix*.

Returning to the PageRank, $g_{ij}p_j$ is the contribution from page j to the PageRank of page i , and we demand that

$$\begin{aligned} &\text{for } i = 1, 2, \dots, n : \\ &g_{i1}p_1 + g_{i2}p_2 + 0 \cdots + g_{in}p_n = p_i, \end{aligned}$$

or, in matrix-vector notation:

$$\mathbf{G}\mathbf{p} = \mathbf{p}. \quad (3.16)$$

We recognize this as an eigenproblem with $\lambda = 1$ and \mathbf{p} , the vector of PageRanks, is a corresponding eigenvector. This may seem a bit magical, but the matrix \mathbf{G} satisfies the conditions of 5° on page 38, so $\lambda = 1$ is indeed one of its eigenvalues. The corresponding eigenvector depends on the choice of α , as shown below for the 6×6 matrix \mathbf{G} corresponding to the web in Figure 3.3.

α	0.5	0.6	0.7	0.8	0.9
\mathbf{p}	0.116	0.102	0.085	0.064	0.037
	0.145	0.133	0.115	0.090	0.054
	0.124	0.111	0.093	0.071	0.042
	0.176	0.181	0.187	0.195	0.206
	0.199	0.213	0.230	0.254	0.286
	0.239	0.262	0.290	0.326	0.375

The eigenvector is normalized such that the sum of the PageRanks is one. We see that all these choices of α agree that the most important pages are 6, 5 and 4 in that order. It is believed that $\alpha = 0.85$ is used by Google. In Section 4.3 we describe how the relevant pages are found.

Example 3.11. The number n of pages registered by Google is not known, except that $n > 10^{10}$, so the matrices are huge, and it is impossible, just to store \mathbf{G} . However, \mathbf{H} is very *sparse*, ie most of the elements are zero, and this can be exploited. The size of the matrix prevents the use of the method mentioned in Example 3.6. Instead a very simple algorithm is employed. This is the so-called *power method*, which has the form

```

Choose  $\mathbf{p}^{[0]}$  with  $p_1^{[0]} + \cdots + p_n^{[0]} = 1$    (for instance  $\mathbf{p}^{[0]} = \frac{1}{n}\mathbf{e}$  )
for  $k = 1, 2, \dots, K$  :
     $\mathbf{p}^{[k]} = \mathbf{G}\mathbf{p}^{[k-1]}$ 
end
```

It can be shown that not only does \mathbf{G} have the eigenvalue $\lambda = 1$, but also the other eigenvalues satisfy $|\lambda_j| < 1$, and that these properties imply that the sequence $\mathbf{p}^{[1]}, \mathbf{p}^{[2]}, \dots$ converges to \mathbf{p} . For our 6×6 example with $\alpha = 0.85$ we get

$\mathbf{p}^{[0]}$	$\mathbf{p}^{[1]}$	$\mathbf{p}^{[2]}$	$\mathbf{p}^{[3]}$	$\mathbf{p}^{[4]}$	$\mathbf{p}^{[5]}$
0.167	0.096	0.082	0.068	0.062	0.057
0.167	0.167	0.123	0.103	0.090	0.083
0.167	0.119	0.089	0.077	0.068	0.064
0.167	0.167	0.193	0.187	0.198	0.196
0.167	0.190	0.230	0.244	0.255	0.261
0.167	0.261	0.281	0.321	0.327	0.339

Already $\mathbf{p}^{[2]}$ shows that most important pages are 6, 5 and 4 in that order.

In the computation of $\mathbf{G}\mathbf{p}^{[k-1]}$ we can exploit the sparsity of \mathbf{H} when computing of the contribution $\alpha\mathbf{H}\mathbf{p}^{[k-1]}$. The contribution from a dangling page (no j) is to add $\alpha/n \cdot p_j^{[k-1]}$ to every component of $\mathbf{p}^{[k]}$. Finally, the condition on the sum of the elements in $\mathbf{p}^{[0]}$ is preserved:

$$\mathbf{e}^T \mathbf{p}^{[k-1]} = p_1^{[k-1]} + \dots + p_n^{[k-1]} = 1, \quad k = 1, 2, \dots$$

This is useful when computing the contribution involving $\mathbf{E}\mathbf{p}^{[k-1]}$: Exploiting the outer product defined on page 3 we get

$$\mathbf{E}\mathbf{p}^{[k-1]} = \mathbf{e}\mathbf{e}^T \mathbf{p}^{[k-1]} = \mathbf{e}.$$

Therefore, the contribution $\frac{1-\alpha}{n} \mathbf{E}\mathbf{p}^{[k-1]}$ consists simply in adding $\frac{1-\alpha}{n}$ to each element in $\alpha\mathbf{S}\mathbf{p}^{[k-1]}$. ■

4. Linear Least Squares Problems

Example 4.1. We have applied different loads p_i to an elastic spring, and measured the resulting lengths ℓ_i .

i	p_i	ℓ_i
1	0.8	7.97
2	1.6	10.2
3	2.4	14.2
4	3.2	16.0
5	4.0	21.2




Figure 4.1. Measured length (ℓ) as function of load (p).

The figure shows the points (p_i, ℓ_i) . They seem to be close to a straight line with the equation

$$\ell = x_1 + x_2 p ,$$

for some values of the *parameters* x_1 and x_2 . In physics this mathematical model is known as *Hooke's law* of elasticity. The first parameter, x_1 , is the undisturbed length of the spring, and $1/x_2$ is the so-called *spring constant*. ■

The problem of determining (a good approximation to) the parameters in Hooke's law is an example of a *linear data fitting problem*: Given *measurements* (t_i, y_i) , $i = 1, 2, \dots, m$, where y_i is assumed to be a perturbed value from some “background” function Y ,

$$y_i = Y(t_i) + \varepsilon_i .$$

The ε_i are measurement errors. Further, we are given a *fitting model*

$$M(\mathbf{x}, t) = x_1 f_1(t) + x_2 f_2(t) + \dots + x_n f_n(t) , \quad (4.1)$$

where the f_j are given functions of t . We wish to choose the parameters \mathbf{x} such that $M(\mathbf{x}, t)$ is close to $Y(t)$ in a certain range of t -values, which includes the given t_i . In other words, we wish to find \mathbf{x} such that the error function

$$e(\mathbf{x}, t) = Y(t) - M(\mathbf{x}, t)$$

is close to zero for all t in the desired range. This task, however, is impossible: We only have discrete information about Y , and furthermore, this information, the y_i , is inflicted with errors. The best we can do is to find \mathbf{x} such that the *residuals*

$$r_i(\mathbf{x}) = y_i - M(\mathbf{x}, t_i) \quad (4.2)$$

are close to zero for all $i = 1, 2, \dots, m$. Traditionally, this goal is interpreted as: Find $\hat{\mathbf{x}}$, the set of parameters that minimizes the function

$$\varphi(\mathbf{x}) = r_1(\mathbf{x})^2 + r_2(\mathbf{x})^2 + \dots + r_m(\mathbf{x})^2 . \quad (4.3)$$

The minimizer $\hat{\mathbf{x}}$ is the so-called *least squares solution*. In the next sections we shall describe how it is computed.

4.1. Overdetermined systems

Combining (4.2) and (4.1) we see that the i^{th} residual is

$$r_i(\mathbf{x}) = y_i - (x_1 f_1(t_i) + x_2 f_2(t_i) + \dots + x_n f_n(t_i)) .$$

This means that the *residual vector* $\mathbf{r} = \mathbf{r}(\mathbf{x})$ is

$$\mathbf{r}(\mathbf{x}) = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{pmatrix} - \begin{pmatrix} f_1(t_1)x_1 + f_2(t_1)x_2 + \dots + f_n(t_1)x_n \\ f_1(t_2)x_1 + f_2(t_2)x_2 + \dots + f_n(t_2)x_n \\ \vdots \\ f_1(t_m)x_1 + f_2(t_m)x_2 + \dots + f_n(t_m)x_n \end{pmatrix} .$$

Further, by comparison with Eq. (1.3) we see that

$$\mathbf{r}(\mathbf{x}) = \mathbf{y} - \mathbf{F}\mathbf{x} , \quad (4.4)$$

where \mathbf{F} is the $m \times n$ matrix

$$\mathbf{F} = \begin{pmatrix} f_1(t_1) & f_2(t_1) & \dots & f_n(t_1) \\ f_1(t_2) & f_2(t_2) & \dots & f_n(t_2) \\ \vdots & \vdots & \ddots & \vdots \\ f_1(t_m) & f_2(t_m) & \dots & f_n(t_m) \end{pmatrix} .$$

Example 4.2. For the problem with the elastic spring in Example 4.1 the free variable is denoted p instead of t and the observations are called ℓ_i instead of y_i . Taking this into account, the fitting model has the form Eq. (4.1) with $f_1(t) = 1$, $f_2(t) = t$ and we see that

$$\mathbf{y} = \begin{pmatrix} 7.97 \\ 10.2 \\ 14.2 \\ 16.0 \\ 21.2 \end{pmatrix} , \quad \mathbf{F} = \begin{pmatrix} 1 & 0.8 \\ 1 & 1.6 \\ 1 & 2.4 \\ 1 & 3.2 \\ 1 & 4.0 \end{pmatrix} .$$

■

Ideally, we would like to find \mathbf{x} so that $\mathbf{r}(\mathbf{x}) = \mathbf{0}$, ie so that

$$\mathbf{F}\mathbf{x} = \mathbf{y} . \quad (4.5)$$

In data fitting problems it is always the case that $m > n$, ie the system involves more equations than unknowns. We say that Eq. (4.5) is an *overdetermined system of equations*.

Example 4.3. Let $m = 3$ and $n = 2$. Figure 4.2 is a geometric illustration in the 3-dimensional space.




Figure 4.2. Geometric illustration of an overdetermined system.

As discussed on page 24, a vector $\mathbf{z} = \mathbf{F}\mathbf{x}$ is in the column range of \mathbf{F} . This is the shaded plane in the figure, spanned by the columns of \mathbf{F} . The vector $\vec{\mathbf{y}}$ in the figure is outside this plane, and therefore the system (4.5) has no solution in this case. We shall use the somewhat sloppy notation, that $\hat{\mathbf{x}}$ is the least squares solution to

$$\mathbf{F}\mathbf{x} \simeq \mathbf{y} .$$

Combining Eqs. (4.3) and (1.2) we see that

$$\varphi(\mathbf{x}) = \|\mathbf{r}(\mathbf{x})\|^2 .$$

The least squares solution $\hat{\mathbf{x}}$ minimizes this, and it follows, that an equivalent definition is that $\hat{\mathbf{x}}$ minimizes $\|\mathbf{y} - \mathbf{F}\hat{\mathbf{x}}\|$. In words: find the vector $\mathbf{F}\hat{\mathbf{x}}$ in the column range of \mathbf{F} , which is closest to the right-hand side vector \mathbf{y} . The corresponding residual is

$$\hat{\mathbf{r}} = \mathbf{y} - \mathbf{F}\hat{\mathbf{x}} . \quad (4.6)$$

Example 4.4. Proceeding with the $m = 3, n = 2$ case from Example 4.3, we recall from Example 1.9 that $\|\mathbf{r}\| = |\vec{\mathbf{r}}|$, the Euclidean length of the residual vector. Intuitively this is smallest when \mathbf{r} is orthogonal to the column range, cf Figure 4.3.




Figure 4.3. Geometric illustration of least squares solution.

A vector in the column range has the form $\mathbf{z} = \mathbf{F}\mathbf{x}$, where \mathbf{x} is an arbitrary n -vector. Generalizing from the example, the vectors \mathbf{z} and $\hat{\mathbf{r}}$ should be orthogonal, so their scalar product should be zero:

$$0 = (\mathbf{F}\mathbf{x})^T \hat{\mathbf{r}} = \mathbf{x}^T \mathbf{F}^T \hat{\mathbf{r}} .$$

This is satisfied for all \mathbf{x} if

$$\mathbf{0} = \mathbf{F}^T \hat{\mathbf{r}} = \mathbf{F}^T (\mathbf{y} - \mathbf{F}\hat{\mathbf{x}}) .$$

By rearranging the terms we see the least squares solution $\hat{\mathbf{x}}$ satisfies

$$(\mathbf{F}^T \mathbf{F}) \hat{\mathbf{x}} = \mathbf{F}^T \mathbf{y} . \quad (4.7)$$

This system of linear equations is known as the *normal equations*.

The matrix

$$\mathbf{A} = \mathbf{F}^T \mathbf{F}$$

is $n \times n$, and in Example 2.22 we saw that it is symmetric and positive definite, provided that the columns in \mathbf{F} are linearly independent. We shall assume that this is the case. Then we can show that $\hat{\mathbf{x}}$ is indeed the minimizer of $\|\mathbf{r}(\mathbf{x})\|$: If we change \mathbf{x} from $\hat{\mathbf{x}}$ to $\hat{\mathbf{x}} + \mathbf{u}$, then the residual changes to

$$\mathbf{r} = \mathbf{y} - \mathbf{F}(\hat{\mathbf{x}} + \mathbf{u}) = \hat{\mathbf{r}} - \mathbf{F}\mathbf{u} ,$$

and, exploiting Eq. (1.2),

$$\begin{aligned} \|\mathbf{r}\|^2 &= \mathbf{r}^T \mathbf{r} = (\hat{\mathbf{r}} - \mathbf{F}\mathbf{u})^T (\hat{\mathbf{r}} - \mathbf{F}\mathbf{u}) \\ &= \hat{\mathbf{r}}^T \hat{\mathbf{r}} - \hat{\mathbf{r}}^T \mathbf{F}\mathbf{u} - \mathbf{u}^T \mathbf{F}^T \hat{\mathbf{r}} + \mathbf{u}^T \mathbf{F}^T \mathbf{F}\mathbf{u} \\ &= \hat{\mathbf{r}}^T \hat{\mathbf{r}} + \mathbf{u}^T \mathbf{A}\mathbf{u} > \hat{\mathbf{r}}^T \hat{\mathbf{r}} \quad \text{if } \mathbf{u} \neq \mathbf{0} . \end{aligned}$$

Summarizing: If the columns in \mathbf{F} are linearly independent, then the least squares solution $\hat{\mathbf{x}}$ can be found by solving the normal equations (4.7). Since $\mathbf{A} = \mathbf{F}^T \mathbf{F}$ is an **spd** matrix, the least squares solution is unique, and it may be found via the Cholesky factorization of \mathbf{A} , cf Section 2.6.

Example 4.5. For the problem with the elastic spring in Example 4.1 we saw in Example 4.1 that

$$\mathbf{y} = \begin{pmatrix} 7.97 \\ 10.2 \\ 14.2 \\ 16.0 \\ 21.2 \end{pmatrix} , \quad \mathbf{F} = \begin{pmatrix} 1 & 0.8 \\ 1 & 1.6 \\ 1 & 2.4 \\ 1 & 3.2 \\ 1 & 4.0 \end{pmatrix} .$$


The columns in \mathbf{F} are linearly independent. The normal equations and their solution are

$$\begin{pmatrix} 5 & 12 \\ 12 & 35.2 \end{pmatrix} \begin{pmatrix} \hat{x}_1 \\ \hat{x}_2 \end{pmatrix} = \begin{pmatrix} 69.57 \\ 192.776 \end{pmatrix} , \quad \hat{\mathbf{x}} = \begin{pmatrix} 4.2360 \\ 4.0325 \end{pmatrix} .$$

This shows, cf Example 4.1, that the undisturbed length of the spring and the spring constant are estimated to be 4.2360 and $1/4.0325 \simeq 0.2480$, respectively.

The data and the fit $M(\hat{\mathbf{x}}, p) = \hat{x}_1 + \hat{x}_2 p$ is shown in Figure 4.4.

Figure 4.4. Spring data and least squares fit.



Example 4.6. We consider the problem from Example 1.2: We have measured the drug concentration y_i in the kidneys at times t_i after the drug was taken. The points shown in Figure 1.2 are $(t_i, Y(t_i))$, where Y is the background function. In practice the measurements have errors, and Figure 4.5 gives an example of results from a “practical” experiment.

The phenomenon may be modelled by a compartment model with two compartments, and according to Eq. (3.12) the background function has the form

$$Y(t) = c_1 e^{\lambda_1 t} + c_2 e^{\lambda_2 t}$$


for some values of $c_1, c_2, \lambda_1, \lambda_2$. This should be exploited in the choice of fitting model. If λ_1 and λ_2 are known, then

$$M(\mathbf{x}, t) = x_1 e^{\lambda_1 t} + x_2 e^{\lambda_2 t}$$

is appropriate. With the data in Figure 4.5 and $\lambda_1 = -0.15, \lambda_2 = -0.30$, cf Example 3.7, we get $\hat{x}_1 = 1.006, \hat{x}_2 = -1.011$

Figure 4.5. Data and fit.

$$M(\hat{\mathbf{x}}, t) = 1.006 e^{-0.15t} - 1.011 e^{-0.30t} .$$



If we do not know λ_1 and λ_2 , then it is appropriate to use a fitting model with four parameters

$$M_{[4]}(\mathbf{x}, t) = x_1 e^{x_3 t} + x_2 e^{x_4 t} .$$

In this model x_3 and x_4 appear nonlinearly, and we cannot use the approach leading to the linear system known as the normal equations. (The matrix \mathbf{F} would depend on the unknowns x_3 and x_4). There is software, for instance in MATLAB and R, for finding the least squares solution also in this case, but it is outside the scope of this lecture note to discuss it. We only mention that for the data in Figure 4.5 and the above 4-parameter model the least squares solution is

$$\hat{\mathbf{x}}_{[4]} = (1.4291 \quad -1.4196 \quad -0.16157 \quad -0.25906)^T .$$

The corresponding sum of the squared residuals is $\varphi(\hat{\mathbf{x}}_4) = 0.001458$. This is indeed

smaller than the result from the 2-parameter model, $\varphi(\hat{\mathbf{x}}) = 0.001512$. If we plotted the graph of $M_{[4]}(\hat{\mathbf{x}}_{[4]}, t)$, it would be almost indistinguishable from the graph of the least squares 2-parameter model $M(\hat{\mathbf{x}}, t)$. ■

4.2. QR factorization

Let \mathbf{F} be an $m \times n$ matrix with $m \geq n$. A *QR factorization* of \mathbf{F} is a relation of the form

$$\mathbf{F} = \mathbf{Q} \begin{pmatrix} \mathbf{R} \\ \mathbf{0} \end{pmatrix}, \quad (4.8)$$

where \mathbf{Q} is an $m \times m$ orthogonal matrix and the $n \times n$ matrix \mathbf{R} is upper triangular. The $\mathbf{0}$ indicates an $(m-n) \times n$ matrix with all elements equal to zero. According to the rules of matrix–matrix products (see Eq. (1.5) and Eq. (1.9)) these zeros imply that the last $m-n$ columns in \mathbf{Q} do not affect the product in Eq. (4.8). Let $\check{\mathbf{Q}}$ denote this part of \mathbf{Q} and let $\hat{\mathbf{Q}}$ be the $m \times n$ matrix consisting of the first n columns in \mathbf{Q} . It follows that Eq. (4.8) is equivalent to the so-called “*economy-size*” (or “*skinny*”) QR factorization,

$$\mathbf{F} = \check{\mathbf{Q}} \mathbf{R}, \quad (4.9)$$

This is illustrated in Figure 4.6




Figure 4.6. QR factorization.

It is outside the scope of this lecture note to show that any \mathbf{F} has a QR factorization and to discuss methods for computing it. Instead we shall demonstrate how it can be used to solve the problem

$$\text{minimize } \|\mathbf{y} - \mathbf{F}\mathbf{x}\|.$$

First, we let \mathbf{Q} be an arbitrary, orthogonal matrix, and use Eq. (1.2): $\|\mathbf{r}\|^2 = \mathbf{r}^T \mathbf{r}$, Eq. (1.6): $(\mathbf{Q}\mathbf{r})^T = \mathbf{r}^T \mathbf{Q}^T$, Eq. (2.18): $\mathbf{Q}^T \mathbf{Q} = \mathbf{I}$ and Example 1.10: $\mathbf{I}\mathbf{r} = \mathbf{r}$:

$$\|\mathbf{Q}\mathbf{r}\|^2 = (\mathbf{Q}\mathbf{r})^T (\mathbf{Q}\mathbf{r}) = \mathbf{r}^T \mathbf{Q}^T \mathbf{Q} \mathbf{r} = \mathbf{r}^T \mathbf{I} \mathbf{r} = \mathbf{r}^T \mathbf{r} = \|\mathbf{r}\|^2.$$

This shows that $\mathbf{Q}\mathbf{r}$ has the same norm as \mathbf{r} , so a minimizer of $\|\mathbf{Q}\mathbf{r}(\mathbf{x})\|$ is also a minimizer for $\|\mathbf{r}(\mathbf{x})\|$. This holds for any orthogonal matrix, and if we use \mathbf{Q}^T with

\mathbf{Q} from the QR factorization, we get

$$\mathbf{Q}^T \mathbf{r}(\mathbf{x}) = \mathbf{Q}^T(\mathbf{y} - \mathbf{F}\mathbf{x}) = \mathbf{Q}^T \mathbf{y} - \begin{pmatrix} \mathbf{R} \\ \mathbf{0} \end{pmatrix} \mathbf{x} = \begin{pmatrix} \hat{\mathbf{Q}}^T \mathbf{y} - \mathbf{R}\mathbf{x} \\ \check{\mathbf{Q}}^T \mathbf{y} \end{pmatrix}.$$

Finally, the definition of the norm, Eq. (1.2), implies that

$$\left\| \begin{pmatrix} \hat{\mathbf{Q}}^T \mathbf{y} - \mathbf{R}\mathbf{x} \\ \check{\mathbf{Q}}^T \mathbf{y} \end{pmatrix} \right\|^2 = \|\hat{\mathbf{Q}}^T \mathbf{y} - \mathbf{R}\mathbf{x}\|^2 + \|\check{\mathbf{Q}}^T \mathbf{y}\|^2.$$

This is clearly minimized when \mathbf{x} is chosen such that $\hat{\mathbf{Q}}^T \mathbf{y} - \mathbf{R}\mathbf{x} = \mathbf{0}$.

Summarizing, we have derive a method for computing the least squares solution $\hat{\mathbf{x}}$. This method consists in computing the economy-size factorization $\mathbf{F} = \hat{\mathbf{Q}}\mathbf{R}$ followed by solving the quadratic system

$$\mathbf{R}\hat{\mathbf{x}} = \hat{\mathbf{Q}}^T \mathbf{y}. \quad (4.10)$$

Since \mathbf{R} is upper triangular, we can use back substitution to solve this system.

Example 4.7. In Example 2.28 we described what happens in MATLAB when we issue the command `x = A \ b`, with a square \mathbf{A} . If \mathbf{F} and \mathbf{y} are the MATLAB representations of $\mathbf{F} \in \mathbb{R}^{m \times n}$ and $\mathbf{y} \in \mathbb{R}^m$, then

$$\mathbf{x} = \mathbf{F} \setminus \mathbf{y}$$

returns \mathbf{x} , the MATLAB representation for $\hat{\mathbf{x}}$, the least squares solution to $\mathbf{F}\mathbf{x} \simeq \mathbf{y}$. The solution is found via the QR factorization, because this method is less affected by rounding errors than the solution via the normal equations.

MATLAB has a built-in function `qr` for computing the QR factorization. The command

$$[\mathbf{Q}, \mathbf{R}] = \text{qr}(\mathbf{F}, 0)$$

returns the economy-size QR factorization. ■

4.3. Singular value decomposition

Let \mathbf{F} be an $m \times n$ matrix with $m \geq n$. The *SVD: singular value decomposition* of \mathbf{F} is the factorization

$$\mathbf{F} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \quad (4.11)$$

where the matrices $\mathbf{U} \in \mathbb{R}^{m \times n}$ and $\mathbf{V} \in \mathbb{R}^{n \times n}$ have orthonormal columns and $\mathbf{\Sigma}$ is the diagonal matrix

$$\mathbf{\Sigma} = \begin{pmatrix} \sigma_1 & & & \\ & \sigma_2 & & \\ & & \ddots & \\ & & & \sigma_n \end{pmatrix},$$

The σ_j are the *singular values*. They are nonnegative and ordered so that

$$\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_p > 0, \quad \sigma_{p+1} = \cdots = \sigma_n = 0. \quad (4.12)$$

If the columns in \mathbf{F} are linearly independent, then all the σ_j are strictly positive, ie $p = n$.

Let $\mathbf{u}_j = \mathbf{U}_{:,j}$, the j^{th} column in \mathbf{U} . This is known as the j^{th} *left singular vector*. Similarly, $\mathbf{v}_j = \mathbf{V}_{:,j}$, $j = 1, \dots, n$ are the *right singular vectors*.

SVD is a powerful tool for analyzing linear mappings: First, we see that the $\{\mathbf{v}_j\}$ form an orthonormal basis for \mathbb{R}^n . This means that any n -vector \mathbf{x} can be expressed as

$$\mathbf{x} = \mathbf{V} \tilde{\mathbf{x}}, \quad (4.13)$$

where $\tilde{\mathbf{x}}$ is the vector of coordinates with respect to the basis $\{\mathbf{v}_j\}$, cf Section 1.5. The matrix \mathbf{V} is orthogonal, and therefore Eq. (4.13) is equivalent to

$$\tilde{\mathbf{x}} = \mathbf{V}^T \mathbf{x}. \quad (4.14)$$

Next, by means of Equations (4.11) and (4.14) we get

$$\begin{aligned} \mathbf{F} \mathbf{x} &= \mathbf{U} \Sigma \mathbf{V}^T \mathbf{x} \\ &= \mathbf{U} \Sigma \tilde{\mathbf{x}} \\ &= \sigma_1 \tilde{x}_1 \mathbf{u}_1 + \sigma_2 \tilde{x}_2 \mathbf{u}_2 + \dots + \sigma_p \tilde{x}_p \mathbf{u}_p. \end{aligned} \quad (4.15)$$

In the last reformulation we exploited that $(\Sigma \tilde{\mathbf{x}})_j = \sigma_j \tilde{x}_j$, Eq. (1.9) and that $\sigma_j = 0$ for $j > p$. Eq. (4.15) shows that vectors $\mathbf{u}_1, \dots, \mathbf{u}_p$ are in the *column range* of \mathbf{F} , cf page 24. They actually form an orthonormal basis for this subspace of \mathbb{R}^m . Further, if $p < n$, then the vectors $\mathbf{v}_{p+1}, \dots, \mathbf{v}_n$ form an orthonormal basis for the *nullspace* of \mathbf{F} .

Now we look at the *inverse problem*: Given \mathbf{F} and \mathbf{y} , find \mathbf{x} such that the norm of the difference $\mathbf{r} = \mathbf{y} - \mathbf{F} \mathbf{x}$ is as small as possible. From the discussion in Section 4.1 we know that

$$\mathbf{F} \hat{\mathbf{x}} = \hat{\mathbf{y}}, \quad \hat{\mathbf{r}} = \mathbf{y} - \mathbf{F} \hat{\mathbf{x}},$$

where $\hat{\mathbf{y}}$ is in the range of \mathbf{F} and $\hat{\mathbf{r}}$ is orthogonal to this subspace of \mathbb{R}^m . This condition implies that

$$\begin{aligned} 0 &= \mathbf{u}_j^T (\mathbf{y} - \mathbf{F} \hat{\mathbf{x}}) \\ &= \mathbf{u}_j^T (\mathbf{y} - \mathbf{U} \Sigma \mathbf{V}^T \hat{\mathbf{x}}) \\ &= \mathbf{u}_j^T (\mathbf{y} - \mathbf{U} \Sigma \tilde{\mathbf{x}}) \\ &= \mathbf{u}_j^T \mathbf{y} - \sigma_j \tilde{x}_j, \quad j = 1, 2, \dots, p. \end{aligned}$$

The relation $\mathbf{u}_j^T \mathbf{U} \Sigma \tilde{\mathbf{x}} = \sigma_j \tilde{x}_j$ follows from the orthonormality of the columns in \mathbf{U} and the result of Exercise 1.13. Since $\mathbf{F} \mathbf{v}_j = \mathbf{0}$ for $j > p$, it follows that $\hat{\mathbf{x}}$ can have arbitrary coordinates \tilde{x}_j , $j = p+1, \dots, n$. Summarizing, we have

$$\hat{\mathbf{x}} = \mathbf{V} \tilde{\mathbf{x}} \quad \text{with} \quad \tilde{x}_j = \begin{cases} \mathbf{u}_j^T \mathbf{y} / \sigma_j & j = 1, 2, \dots, p, \\ \alpha_j & j = p+1, \dots, n, \end{cases} \quad (4.16)$$

where the α_j are arbitrary.

Example 4.8. In Example 2.7 we looked at the problem $\mathbf{A} \mathbf{x} = \mathbf{y}$ with

$$\mathbf{A} = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} 8 \\ 20 \\ 32 \end{pmatrix}.$$

The SVD of the matrix is $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ with (rounded to three decimals)

$$\mathbf{U} = \begin{pmatrix} 0.215 & 0.887 & 0.408 \\ 0.521 & 0.250 & -0.816 \\ 0.826 & -0.388 & 0.408 \end{pmatrix}, \quad \mathbf{\Sigma} = \begin{pmatrix} 16.848 & & \\ & 1.068 & \\ & & 0 \end{pmatrix},$$

$$\mathbf{V} = \begin{pmatrix} 0.480 & -0.777 & -0.408 \\ 0.572 & -0.076 & 0.816 \\ 0.665 & 0.625 & -0.408 \end{pmatrix}.$$

We see that $p=2$ and that the nullspace of \mathbf{A} is spanned by the last column in \mathbf{V} . This is equivalent to the nullspace basis $(1 \ -2 \ 1)^T$ given in Example 1.15.

The coordinates of \mathbf{y} with respect to the basis given by the columns in \mathbf{U} are

$$\tilde{\mathbf{y}} = \mathbf{U}^T \mathbf{y} = \begin{pmatrix} 38.573 \\ -0.323 \\ 0 \end{pmatrix}.$$

$\tilde{y}_3 = 0$ shows that \mathbf{y} has no component outside the range of \mathbf{A} , ie the system is consistent. ■

Example 4.9. There is a close relation between the SVD for a matrix \mathbf{F} and the eigensolutions of the corresponding normal equations matrix $\mathbf{A} = \mathbf{F}^T \mathbf{F}$. If we insert Eq. (4.11) and exploit that $\mathbf{U}^T \mathbf{U} = \mathbf{I}$, we get

$$\mathbf{A} = \mathbf{V} \mathbf{\Sigma}^T \mathbf{U}^T \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T = \mathbf{V} \mathbf{\Sigma}^2 \mathbf{V}^T,$$

where

$$\mathbf{\Sigma}^2 = \begin{pmatrix} \sigma_1^2 & & & \\ & \sigma_2^2 & & \\ & & \ddots & \\ & & & \sigma_n^2 \end{pmatrix}.$$

Comparing this with Eq. (3.7) we see that the right singular vectors are eigenvectors of \mathbf{A} , and that $\sigma_j = \sqrt{\lambda_j}$. We know that the eigenvalues are nonnegative, and assume that they are numbered in decreasing order. The left singular vectors may afterwards be computed by means of Eq. (4.15):

$$\mathbf{u}_j = \frac{1}{\sigma_j} \mathbf{F} \mathbf{v}_j, \quad j = 1, 2, \dots, p. \quad \blacksquare$$

Example 4.10. MATLAB has a built-in function `svd` for computing SVD. The command

$$\mathbf{s} = \text{svd}(\mathbf{F})$$

returns a vector \mathbf{s} with the singular values of \mathbf{F} . The command

$$[\mathbf{U}, \mathbf{S}, \mathbf{V}] = \text{svd}(\mathbf{F}, 0)$$

returns the matrices \mathbf{U} , $\mathbf{\Sigma}$ and \mathbf{V} . Instead of the method outlined in the previous example `svd` uses a special purpose algorithm that avoids the possible loss of information that may result from forming the normal equations matrix. ■

We shall mention another application of SVD, which is closely related to a statistical method known as *PCA: principal component analysis*. It is based on an alternative formulation of the decomposition. Using the rules from Section 1.3 we see that $\mathbf{F} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ is equivalent to

$$\mathbf{F} = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T + \sigma_2 \mathbf{u}_2 \mathbf{v}_2^T + \cdots + \sigma_n \mathbf{u}_n \mathbf{v}_n^T. \quad (4.17)$$

In words: Any matrix can be expressed as a weighted sum of outer products. The weights are the singular values, and they reflect the importance of each contribution. It sometimes happens that most of the information in \mathbf{F} is accounted for by the approximation

$$\mathbf{F}_k = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T + \sigma_2 \mathbf{u}_2 \mathbf{v}_2^T + \cdots + \sigma_k \mathbf{u}_k \mathbf{v}_k^T,$$

where k is maybe much smaller than n . We can write this approximation in the form

$$\mathbf{F}_k = \mathbf{U}_k \mathbf{\Sigma}_k \mathbf{V}_k^T, \quad (4.18)$$

where \mathbf{U}_k and \mathbf{V}_k consist of the first k columns in \mathbf{U} and \mathbf{V} , respectively, and $\mathbf{\Sigma}_k$ is the upper left $k \times k$ corner of $\mathbf{\Sigma}$.

Example 4.11. PCA is widely used in, for instance, medicine to identify certain combinations of parameters that can be used to explain most of the variation in the observed data.

As a specific example, however, we shall give a simplified description of a search machine like Google. We set up an $m \times n$ matrix, where m is the number of pages and n is the number of selected *keywords*. The element f_{ij} is a measure of the importance of the j^{th} keyword in the i^{th} page. As mentioned in Example 3.11, $m > 10^{10}$, and n is maybe 50 000, so \mathbf{F} is a huge matrix (but it is very sparse).

Now, given a query, we set up a vector \mathbf{q} , according to the keywords, and compute its coordinates with respect to the columns in \mathbf{V}_k ,

$$\tilde{\mathbf{q}} = \mathbf{V}_k^T \mathbf{q}.$$

Let J denote the set of indices for which $\tilde{q}_j \geq \tau_1 \|\tilde{\mathbf{q}}\|$ where τ_1 is a certain threshold, for instance $\tau_1 = 0.05$. If J has only one element, then $\mathbf{y} = \mathbf{U}_{:,j}^{[k]}$ is the corresponding vector in \mathbb{R}^m . If J has more than one element, then

$$\mathbf{y} = \mathbf{U}_J^{[k]} \mathbf{\Sigma}_J^{[k]} \tilde{\mathbf{q}}_J,$$

where index J signifies that we only include the contributions corresponding to the indices from J .

The *hits* are those pages for which $y_i \geq \tau_2 \|\mathbf{y}\|$, where for instance $\tau_2 = 0.01$. The ranking of page i is some combination of the value y_i and the PageRank p_i , cf Section 3.2.3.

If k in Eq. (4.18) is too small, we may both overlook interesting pages and get hits that we really do not want. For larger k these misclassifications are reduced, but the computational cost increases. See for instance the book by Eldén for a discussion of this. ■

Literature

There are many books on linear algebra. For those readers that want to know more about the subject we can recommend

G. Strang, *Linear Algebra and its Application*, 4th edition, Cengage, 2006

J. Eising, *Lineær algebra*, Matematisk Institut, DTU, 1993. (In Danish)

The following book focuses on numerical aspects applications

L.N. Trefethen and D. Bau, *Numerical linear algebra*, SIAM 1997,

while

L. Eldén, *Matrix Methods In Data Mining and Pattern Recognition*, SIAM, 2007

treats both numerical linear algebra and applications. This includes the mathematics involved in Google.

Further, we would like to recommend

E.S. Allman and J.A. Rhodes, *Mathematical Models in Biology*, Cambridge University Press, 2004

M.W. Berry and M. Browne, *Understanding Search Engines*, SIAM, 1999

A.N. Langville and C.D. Meyer, *Google's PageRank and Beyond*, Princeton University Press, 2006

Notation

Vectors are normally written as lower case bold letters, and matrices are upper case, bold letters.

\mathbf{A}	Typical name for a matrix.
\mathbf{A}^T	The transpose of matrix \mathbf{A} .
\mathbf{A}^{-1}	Matrix \mathbf{A} inverse.
$a_{ij}, (\mathbf{A})_{ij}$	The element in position (i, j) in the matrix \mathbf{A} .
$\mathbf{A}_{i,:}$	The row vector with the elements in the i^{th} row of \mathbf{A} .
$\mathbf{A}_{:,j}$	The column vector with the elements in the j^{th} column of \mathbf{A} .
m, n	Positive integers.
\mathbb{R}	The set of real numbers.
\mathbb{R}^n	The set of vectors with n real elements.
$\mathbb{R}^{m \times n}$	The set of matrices with m rows and n columns, where all elements are real.
\vec{v}	Typical name for a geometric vector.
\mathbf{x}	Typical name for a vector. Unless otherwise stated, \mathbf{x} is a column vector (cf page 2).
x_i	i^{th} element in the vector \mathbf{x} . i is the <i>index</i> of the element.
α, β	Typical names for a scalar.
λ	Eigenvalue.
$\mathbf{\Lambda}$	Diagonal matrix with $(\mathbf{\Lambda})_{jj} = \lambda_j$, the j^{th} eigenvalue.

In MATLAB examples we use typewriter font to indicate the MATLAB representation of a variable. For instance `A` is the MATLAB variable that represents matrix \mathbf{A} .

List of Danish Words

I mange tilfælde er det danske ord for et begreb identisk med det engelske på nær en mindre ændring i stavemåde, fx *vector* \sim *vektor*, *orthogonal* \sim *ortogonal*. I andre tilfælde fremkommer den danske betegnelse ved direkte oversættelse af den engelske. Nedenfor giver vi en liste over de danske betegnelser i de tilfælde, hvor oversættelsen ikke er umiddelbar.

<i>English</i>	<i>Dansk</i>
augmented matrix	totalmatrix
back substitution	tilbage-løsning
column	søjle
column range	billedrum
complete solution	fuldstændig løsning
eigenvalue	egenværdi
eigenvector	egenvektor
forward substitution	fremad-løsning
Gaussian elimination	Gauss-elimination
identity matrix	enhedsmatrix
kernel	kærne
mapping	afbildning
matrix	matrix. Én matrix, matricen, flere matricer. En <i>matrice</i> er en støbeform
nonsingular	regulær (om matrix)
quadratic	2. grads polynomium
rank	rang
row	række
singular	singulær (om matrix)
square	kvadratisk eller kvadratformet (om matrix)
transpose	transponeret
triangular matrix	trekantmatrix

Index

- active part, 20
- affine mapping, 10
- algebraic multiplicity, 36
- arithmetic operation, 30
- augmented matrix, 20

- back substitution, 18, 29, 34, 53, 59
- background function, 47, 51
- backslash, 34, 53
- basis, 10, 35, 40, 54

- change of basis, 12, 35
- characteristic polynomial, 36
- `chol`, 33
- Cholesky factorization, 32, 50
- coefficient matrix, 15
- column range, 24, 49, 59
 - vector, 2f
- compartment model, 41, 51
- complete solution, 22ff, 36, 59
- consistent system, 22, 55
- coordinates, 1, 10, 27, 43, 55
- cost, 30
- CT scanning, 8

- dangling page, 44
- determinant, 16, 36
- diagonal, 2, 37, 39
- diagonal matrix, 2, 6, 17, 26, 31, 37, 41
- differential equation, 40, 44
- discriminant, 36
- dot product, 3, 5
- drug concentration, 1, 41, 51

- economy-size, 52f
- `eig`, 39
- eigen factorization, 38
- eigenproblem, 35
- eigensolution, 35
- elimination factor, 19
- ellipse, 13, 35, 42

- factorization, Cholesky, 32, 50
 - eigen, 38
 - LDL, 31
 - LU, 28
 - QR, 52
 - SVD, 53
- fitting model, 47, 51
- flop, 30
- forward substitution, 18, 29, 34, 59

- Gaussian elimination, 18
- geometric illustration, 49
 - vector, 1, 5, 10
- Google, 35, 44, 56
 - matrix, 45

- hits, 56
- homogeneous system, 24, 36
- Hooke's law, 47

- identity matrix, 2, 6, 10, 26, 37
- image, 8
- initial value problem, 40
- inner product, 3
- inverse problem, 11, 54

- kernel, 23
- keywords, 56

- LDL factorization, 31
- least squares solution, 48f
- linear combination, 9, 37
 - data fitting, 47
 - independency, 9f, 30, 37, 50, b 54
 - mapping, 8f, 35, 54
- linearity conditions, 8, 23, 37
- links, 44
- lower triangular matrix, 2
- `lu`, 30
- LU factorization, 28

- mapping matrix, 8, 35
- MATLAB, 1, 3f, 12, 16, 30, 33f, 39, 41, 51, 53, 55
- matrix, 2
 - addition, 4
 - column, 2, 9
 - equation, 25
 - exponential, 35,42
 - row, 2
- matrix, diagonal, 2, 6, 17, 26, 31, 37, 41
 - identity, 3
 - inverse, 26
 - nonsingular, 16, 22f, 38, 59
 - orthogonal, 17, 27, 52
 - singular, 16, 22f, 29, 36, 59
 - sparse, 45, 56
 - SPD, 31, 39, 50
 - square, 2, 18, 35, 39
 - symmetric, 30, 37
 - triangular, 2, 17f, 28, 37, 52
- matrix-matrix product, 7, 27
 - -vector product, 5, 8, 30
- model, 2, 8, 15, 26, 39, 41, 44, 47, 51
- MR scanning, 8
- multiple eigenvalue, 37

- nonsingular matrix, 16, 22f, 38, 59
- norm, 4, 11, 38
- norm conditions, 5
- normal equations, 31, 50f
- nullspace, 23, 54

- orthogonal, 4, 29, 37, 50
 - matrix, 17, 27, 52
- orthonormal, 11, 17, 27, 38, 53
- out-links, 44
- outer product, 4, 56
- overdetermined, 48

- PageRank, 44, 56
- parabola, 6, 15, 21, 43
- PCA, 56
- permutation matrix, 33
- pivoting, 25

- positive definite, 30, 39
- positive semidefinite, 30
- principal component, 56

- qr, 53
- QR factorization, 52f
- quadratic equation, 36

- residual, 47
- rotation matrix, 12, 17
- rounding errors, 27, 32, 53
- row vector, 2

- scalar product, 3, 6
- search machine, 56
- similarity transformation, 39
- singular value decomposition, 53
- singular vector, 54
- skinny, 52
- sparse matrix, 45, 56
- SPD matrix, 31, 39, 50
- spring, 47, 50
- square matrix, 2, 18, 35, 39
- surfer, 44
- SVD, 53
- svd, 55
- symmetric matrix, 30, 37

- threshold, 56
- transpose, 3, 58
- triangle inequality, 5
- triangular system, 17

- usual basis, 10

- vector, 1
 - addition, 3
 - norm, 4
 - transpose, 3
- vector-vector product, 3f

- web, 44

- zero vector, 2