Steffen Börm
Christian Mehl

Numerical Methods for Eigenvalue Problems

De Gruyter

Mathematics Subject Classification 2010: Primary: 15A18, 15A22, 15A23, 15A42, 65F15, 65F25; *Secondary:* 65N25, 15B57.

ISBN 978-3-11-025033-6
e-ISBN 978-3-11-025037-4

Library of Congress Cataloging-in-Publication Data

A CIP catalog record for this book has been applied for at the Library of Congress.

Bibliographic information published by the Deutsche Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the internet at http://dnb.d-nb.de.

© 2012 Walter de Gruyter GmbH & Co. KG, Berlin/Boston

Typesetting: Da-TeX Gerd Blumenstein, Leipzig, www.da-tex.de
Printing and binding: Hubert & Co. GmbH & Co. KG, Göttingen
∞ Printed on acid-free paper

Printed in Germany

www.degruyter.com

Preface

Eigenvalues and eigenvectors of matrices and linear operators play an important role when solving problems from structural mechanics, and electrodynamics, e.g., by describing the resonance frequencies of systems, when investigating the long-term behaviour of stochastic processes, e.g., by describing invariant probability measures, and as a tool for solving more general mathematical problems, e.g., by diagonalizing ordinary differential equations or systems from control theory.

This book presents a number of the most important numerical methods for finding eigenvalues and eigenvectors of matrices. It is based on lecture notes of a short course for third-year students in mathematics, but it should also be accessible to students of physics or engineering sciences.

We discuss the central ideas underlying the different algorithms and introduce the theoretical concepts required to analyze their behaviour. Our goal is to present an easily accessible introduction to the field, including rigorous proofs of all important results, but not a complete overview of the vast body of research.

For an in-depth coverage of the theory of eigenvalue problems, we can recommend the following monographs:

- J. H. Wilkinson, "The Algebraic Eigenvalue Problem" [52]
- B. N. Parlett, "The Symmetric Eigenvalue Problem" [33]
- G. H. Golub and C. F. Van Loan, "Matrix Computations" [18]
- G. W. Stewart, "Matrix Algorithms" [43, 44]
- D. S. Watkins, "The matrix eigenvalue problem" [49]

We owe a great debt of gratitude to their authors, since this book would not exist without their work.

The book is intended as the basis for a short course (one semester or trimester) for third- or fourth-year undergraduate students. We have organized the material mostly in short sections that should each fit one session of a course. Some chapters and sections are marked by an asterisk *. These contain additional results that we consider optional, e.g., rather technical proofs of general results or algorithms for special problems. With one exception, the results of these optional sections are not required for the remainder of the book. The one exception is the optional Section 2.7 on non-unitary transformations, which lays the groundwork for the optional Section 4.8 on the convergence of the power iteration for general matrices.

In order to keep the presentation self-contained, a number of important results are proven only for special cases, e.g., for self-adjoint matrices or a spectrum consisting only of simple eigenvalues. For the general case, we would like to refer the reader to the monographs mentioned above.

Deviating from the practice of collecting fundamental results in a separate chapter, we introduce some of these results when they are required. An example is the Bauer–Fike theorem given as Proposition 3.11 in Section 3.4 on error estimates for the Jacobi iteration instead of in a separate chapter on perturbation theory. While this approach is certainly not adequate for a reference work, we hope that it improves the accessibility of lecture notes like this book that are intended to be taught in sequence.

We would like to thank Daniel Kressner for his valuable contributions to this book.

Kiel, December 2011
Steffen Börm

Berlin, December 2011
Christian Mehl

Contents

Chapter 1

Introduction

Eigenvalue problems play an important role in a number of fields of numerical mathematics: in structural mechanics and electrodynamics, eigenvalues correspond to resonance frequencies of systems, i.e., to frequencies to which these systems respond particularly well (or badly, depending on the context). When studying stochastic processes, invariant probability measures correspond to eigenvectors for the eigenvalue 1, and finding them yields a description of the long-term behaviour of the corresponding process.

This book gives an introduction to the basic theory of eigenvalue problems and focuses on important algorithms for finding eigenvalues and eigenvectors.

1.1 Example: Structural mechanics

Before we consider abstract eigenvalue problems, we turn our attention to some applications that lead naturally to eigenvalue problems.

The first application is the investigation of resonance frequencies. As an example, we consider the oscillations of a string of unit length. We represent the string as a function

$$u : \mathbb{R} \times [0,1] \to \mathbb{R}, \quad (t,x) \mapsto u(t,x),$$

where $u(t,x)$ denotes the deflection of the string at time t and position x (cf. Figure 1.1).

The oscillations are then described by the *wave equation*

$$\frac{\partial^2 u}{\partial t^2}(t,x) = c\frac{\partial^2 u}{\partial x^2}(t,x) \quad \text{for all } t \in \mathbb{R},\ x \in (0,1),$$

where $c > 0$ is a parameter describing the string's properties (e.g., its thickness). We assume that the string is fixed at both ends, i.e., that

$$u(t,0) = u(t,1) = 0 \quad \text{holds for all } t \in \mathbb{R}.$$

Since the differential equation is linear, we can separate the variables: we write u in the form

$$u(t,x) = u_0(x)\cos(\omega t) \quad \text{for all } t \in \mathbb{R},\ x \in [0,1]$$

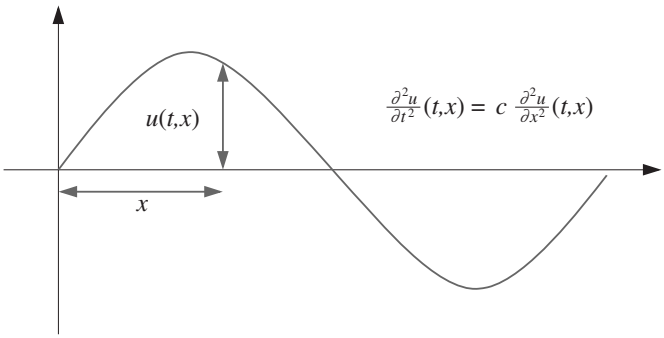

Figure 1.1. Mathematical model of a string.

with a *frequency parameter* $\omega \in \mathbb{R}_{\geq 0}$ and a function

$$u_0 : [0,1] \to \mathbb{R}, \quad x \mapsto u_0(x),$$

depending only on the location, but not on the time. The differential equation takes the form

$$\begin{aligned} -\omega^2 u_0(x)\cos(\omega t) &= \frac{\partial^2 u}{\partial t^2}(t,x) = c\frac{\partial^2 u}{\partial x^2}(t,x) \\ &= c u_0''(x)\cos(\omega t) \quad \text{for all } t \in \mathbb{R},\ x \in (0,1), \end{aligned}$$

and eliminating the time-dependent factor yields

$$-c u_0''(x) = \omega^2 u_0(x) \quad \text{for all } x \in (0,1). \tag{1.1}$$

We introduce $\lambda := \omega^2 \in \mathbb{R}_{\geq 0}$ and define the differential operator L by

$$L[u_0](x) := -c u_0''(x) \quad \text{for all } u_0 \in C^2(0,1),\ x \in (0,1)$$

in order to obtain

$$L[u_0] = \lambda u_0.$$

This is an eigenvalue problem in the infinite-dimensional space $C^2(0,1)$.

In order to be able to treat it by a numerical method, we have to *discretize* the problem. A simple approach is the *finite difference method*: Taylor expansion yields

$$\begin{aligned} u_0(x+h) &= u_0(x) + h u_0'(x) + \frac{h^2}{2}u_0''(x) + \frac{h^3}{6}u_0^{(3)}(x) + \frac{h^4}{24}u_0^{(4)}(\eta_+), \\ u_0(x-h) &= u_0(x) - h u_0'(x) + \frac{h^2}{2}u_0''(x) - \frac{h^3}{6}u_0^{(3)}(x) + \frac{h^4}{24}u_0^{(4)}(\eta_-) \end{aligned}$$

for $h \in \mathbb{R}_{>0}$ with $0 \leq x-h \leq x+h \leq 1$, where $\eta_+ \in [x, x+h]$ and $\eta_- \in [x-h, x]$. Adding both equations and using the intermediate value theorem yields

$$u_0(x-h) - 2u_0(x) + u_0(x+h) = h^2 u_0''(x) + \frac{h^4}{12} u_0^{(4)}(\eta)$$

with $\eta \in [x-h, x+h]$. Dividing by h^2 gives us an equation for the second derivative:

$$\frac{u_0(x-h) - 2u_0(x) + u_0(x+h)}{h^2} = u_0''(x) + \frac{h^2}{12} u_0^{(4)}(\eta).$$

We obtain a useful approximation by dropping the right-most term: fixing $n \in \mathbb{N}$ and setting

$$x_k := hk, \quad h := \frac{1}{n+1} \quad \text{for all } k \in \{0, \ldots, n+1\},$$

we find

$$\frac{u_0(x_{k-1}) - 2u_0(x_k) + u_0(x_{k+1})}{h^2} \approx u_0''(x_k) \quad \text{for all } k \in \{1, \ldots, n\},$$

and the term on the left-hand side requires only values of u_0 in the discrete points $x_0, \ldots, x_{n+1}$. We collect these values in a vector

$$e := \begin{pmatrix} u_0(x_1) \\ \vdots \\ u_0(x_n) \end{pmatrix}, \quad e_0 = e_{n+1} = 0,$$

and replace $u_0''(x)$ in (1.1) by the approximation to get

$$c\frac{2e_k - e_{k-1} - e_{k+1}}{h^2} \approx \lambda e_k \quad \text{for all } k \in \{1, \ldots, n\}.$$

This system can be written as the algebraic eigenvalue problem

$$\frac{c}{h^2} \begin{pmatrix} 2 & -1 & & \\ -1 & \ddots & \ddots & \\ & \ddots & \ddots & -1 \\ & & -1 & 2 \end{pmatrix} \begin{pmatrix} e_1 \\ \vdots \\ \vdots \\ e_n \end{pmatrix} \approx \lambda \begin{pmatrix} e_1 \\ \vdots \\ \vdots \\ e_n \end{pmatrix}, \tag{1.2}$$

and solving the system yields approximations $u_0(x_k) \approx e_k$ of the values of u_0 in the points $x_1, \ldots, x_n$.

In order to reach a high accuracy, we have to ensure that h is small, so we have to be able to handle large values of n. We are typically only interested in computing a small number of the smallest eigenvalues, and this problem can be solved efficiently by specialized algorithms (e.g., the inverse iteration discussed in Chapter 4).

1.2 Example: Stochastic processes

The next example is not motivated by physics, but by computer science: we are interested in determining the "most important" pages of the world wide web. Let $n \in \mathbb{N}$ be the number of web pages, and let $L \in \mathbb{R}^{n\times n}$ represent the hyperlinks between these pages in the following way:

$$\ell_{ij} = \begin{cases} 1 & \text{if page } j \text{ contains a link to page } i, \\ 0 & \text{otherwise} \end{cases} \qquad \text{for all } i, j \in \{1, \dots, n\}.$$

We follow the PageRank [31] approach: we consider a "random web user" that moves from page to page and compute the probability $p_j^{(m)} \in [0, 1]$ of him visiting a certain page j in his m-th step. We denote the number of links on page j by

$$\ell_j := \sum_{i=1}^{n} \ell_{ij} \qquad \text{for all } j \in \{1, \dots, n\}$$

and assume that the random user chooses each of the links with equal probability $1/\ell_j$. In order to make this approach feasible, we have to assume $\ell_j \neq 0$ for all $j \in \{1, \dots, n\}$, i.e., we have to assume that each page contains at least one link.

This means that the probability of switching from page j to page i is given by

$$s_{ij} := \frac{\ell_{ij}}{\ell_j} = \begin{cases} 1/\ell_j & \text{if } \ell_{ij} = 1, \\ 0 & \text{otherwise} \end{cases} \qquad \text{for all } i, j \in \{1, \dots, n\}.$$

The probability of visiting page i in step $m + 1$ is given by

$$p_i^{(m+1)} = \sum_{j=1}^{n} s_{ij}\, p_j^{(m)} \qquad \text{for all } i \in \{1, \dots, n\},\ m \in \mathbb{N}_0.$$

Since this equation corresponds to a matrix-vector multiplication by $S = (s_{ij})_{i,j=1}^{n}$, we can write it in the compact form

$$p^{(m+1)} = Sp^{(m)} \qquad \text{for all } m \in \mathbb{N}_0. \tag{1.3}$$

In order to ensure that the result does not depend on the arbitrarily chosen starting vector $p^{(0)}$, the PageRank algorithm uses the limit

$$p^* := \lim_{m\to\infty} p^{(m)} \tag{1.4}$$

to determine the "importance" of a web page: if p_j^* is large, the probability of a user visiting the j-th web page is high, therefore it is assumed to be important. Due to

$$p^* = \lim_{m\to\infty} p^{(m)} = \lim_{m\to\infty} p^{(m+1)} = \lim_{m\to\infty} Sp^{(m)} = S \lim_{m\to\infty} p^{(m)} = Sp^*,$$

the vector $p^* \in \mathbb{R}^n$ is an eigenvector of the matrix S for the eigenvalue one.

In this example, we not only reduce a problem related to stochastic processes (the "random walk" of the web user) to an algebraic eigenvalue problem, but we also find a simple algorithm for computing at least the eigenvector: due to (1.4), we can hope to approximate p^* by computing $p^{(m)}$ for a sufficiently large value of m. Due to (1.3), this requires only m matrix-vector multiplications, and since we can assume that each web page contains only a small number of links, these multiplications can be performed very efficiently.

In order to ensure convergence, the PageRank algorithm replaces the matrix S by the matrix $\widehat{S} = (\hat{s}_{ij})_{i,j=1}^n$ given by

$$\hat{s}_{ij} = (1-\alpha)s_{ij} + \alpha u_i \quad \text{for all } i, j \in \{1, \ldots, n\},$$

where $\alpha \in (0, 1]$ (a typical value is 0.15) is a parameter controlling how close $\widehat{S}$ is to the original matrix S, while $u \in \mathbb{R}^n$ is a vector (the "teleportation vector") describing the event that a user switches to a different page without following a link: when visiting page j, the user either follows one of the links with a total probability of $1 - \alpha$ or switches to another page with a total probability of α. In the latter case, u_i is the relative probability of switching to page i. For a suitable choice of u (e.g., a vector with strictly positive entries and $u_1 + \ldots + u_n = 1$), the *Perron–Frobenius theorem* [34] ensures convergence of the sequence $(p^{(m)})_{m=0}^{\infty}$ to p^*.

1.3 Example: Systems of linear differential equations

Another, more general example for eigenvalue problems are systems of linear differential equations that appear frequently in natural and engineering sciences. If $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$ and $A \in \mathbb{F}^{n\times n}$ then

$$y' = Ay \tag{1.5}$$

is a *(homogeneous) system of linear differential equations* and a solution is defined to be a continuously differentiable function $y : \mathbb{R} \to \mathbb{F}^n$ satisfying $y'(t) = Ay(t)$ for all $t \in \mathbb{R}$. For some vector $y_0 \in \mathbb{F}^n \setminus \{0\}$ the ansatz $y(t) = e^{\lambda t} y_0$ yields the identity

$$\lambda e^{\lambda t} y_0 = y'(t) = Ay(t) = Ae^{\lambda t} y_0$$

which, after division on both sides by $e^{\lambda t}$, reduces to the characteristic equation

$$\lambda y_0 = Ay_0.$$

Thus, if $\lambda \in \mathbb{F}$ is an eigenvalue of A and $y_0 \in \mathbb{F}^n \setminus \{0\}$ is an associated eigenvector, then $y(t) = e^{\lambda t} y_0$ is a solution of the corresponding system of differential equations. It is well known from the theory of differential equations that if A is diagonalizable and if $v_1, \ldots, v_n$ is a basis of $\mathbb{F}^n$ consisting of eigenvectors of A associated with

the eigenvalues $\lambda_1, \ldots, \lambda_n$, then any solution y of the system of differential equations (1.5) has the form

$$y(t) = \sum_{i=1}^{n} c_i e^{\lambda_i t} v_i$$

for some coefficients $c_1, \ldots, c_n \in \mathbb{F}$. In the non-diagonalizable case, the general solution can be constructed from the so called *Jordan normal form*.

Instead of systems of linear differential equations of first order as in the form (1.5), one may also consider systems of linear differential equations of higher order having the general form

$$\sum_{k=0}^{\ell} A_k y^{(k)} = A_\ell y^{(\ell)} + A_{\ell-1} y^{(\ell-1)} + \cdots + A_2 y'' + A_1 y' + A_0 y = 0,$$

where $A_0, \ldots, A_\ell \in \mathbb{F}^{n\times n}$. In this case the ansatz $y(t) = e^{\lambda t} y_0$ for some nonzero vector $y_0 \in \mathbb{F}^n$ yields the identity

$$\sum_{k=0}^{\ell} A_k \lambda^k e^{\lambda t} y_0 = 0,$$

or equivalently, after division by $e^{\lambda t}$,

$$\left(\sum_{k=0}^{\ell} \lambda^k A_k \right) y_0 = 0. \tag{1.6}$$

The problem of solving (1.6) is called a *polynomial eigenvalue problem*.

A particular example in applications can be found in the theory of mechanical vibration. The equations of motion for a viscously damped linear system with n degrees of freedom are given by

$$My''(t) + Cy'(t) + Ky(t) = 0,$$

where M, C, K are $n \times n$ matrices called mass matrix, damping matrix, and stiffness matrix, respectively. The corresponding *quadratic eigenvalue problem* has the form

$$(\lambda^2 M + \lambda C + K)x = 0.$$

A simple example for the case $n = 1$ is the spring-mass system with damping by friction, see Figure 1.2.

By Hooke's law, the equation of motion for this system without friction is

$$my''(t) + ky(t) = 0,$$

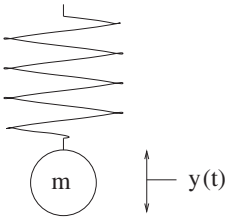

Figure 1.2. Mass-spring system.

where m is the mass attached to the spring and k is the spring constant. If friction is considered, it is usually modeled in such a way that the friction is assumed to be proportional to the velocity $y'(t)$ thus yielding the equation of motion

$$my''(t) + cy'(t) + ky(t) = 0,$$

for some constant c.

Chapter 2

Existence and properties of eigenvalues and eigenvectors

Summary

This chapter investigates existence and uniqueness of eigenvalues and eigenvectors for a given matrix. The key result is the *Schur decomposition* introduced in Theorem 2.46, a very useful tool for the investigation of eigenvalue problems. One of its most important consequences is the fact that a matrix can be diagonalized unitarily if and only if it is normal. The optional Section 2.7 presents a block-diagonalization result for general square matrices

Learning targets

- ✓ Recall the definition and some of the most important properties of eigenvalues, eigenvectors, similarity transformations and the characteristic polynomial corresponding to a matrix.
- ✓ Introduce a number of basic concepts of Hilbert space theory, e.g., the Cauchy–Schwarz inequality, self-adjoint, normal, isometric and unitary matrices.
- ✓ Prove the existence of the Schur decomposition.
- ✓ Use it to establish the existence of eigenvector bases for normal and self-adjoint matrices and of invariant subspaces in the general case.

2.1 Eigenvalues and eigenvectors

Let $n, m \in \mathbb{N}$, and let $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$ be the field of real or complex numbers.

We denote the space of matrices with n rows and m columns by $\mathbb{F}^{n\times m}$. The coefficients of a matrix $A \in \mathbb{F}^{n\times m}$ are given by

$$A = \begin{pmatrix} a_{11} & \dots & a_{1m} \\ \vdots & \ddots & \vdots \\ a_{n1} & \dots & a_{nm} \end{pmatrix}.$$

Zero coefficients in a matrix are frequently omitted in our notation, e.g., the n-dimensional identity matrix is usually represented by

$$I_n = \begin{pmatrix} 1 & & \\ & \ddots & \\ & & 1 \end{pmatrix}.$$

If the dimension is clear from the context, we denote the identity matrix by I.

The product of a matrix by a vector $x \in \mathbb{F}^m$ is given by

$$Ax = \begin{pmatrix} a_{11} & \dots & a_{1m} \\ \vdots & \ddots & \vdots \\ a_{n1} & \dots & a_{nm} \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_m \end{pmatrix} = \begin{pmatrix} a_{11}x_1 + \cdots + a_{1m}x_m \\ \vdots \\ a_{n1}x_1 + \cdots + a_{nm}x_m \end{pmatrix}$$

or, more precisely, by

$$(Ax)_i = \sum_{j=1}^{m} a_{ij}x_j \quad \text{for all } i \in \{1, \dots, n\}. \tag{2.1}$$

The mapping

$$\mathbb{F}^m \to \mathbb{F}^n, \quad x \mapsto Ax$$

is a linear operator mapping $\mathbb{F}^m$ to $\mathbb{F}^n$, and we use the corresponding notations:

Definition 2.1 (Null space and range). Let $A \in \mathbb{F}^{n\times m}$. The space

$$\mathcal{N}(A) := \{x \in \mathbb{F}^m \ : \ Ax = 0\}$$

is called the *null space* of A, and its dimension is called the *nullity* of A. The space

$$\mathcal{R}(A) := \{y \in \mathbb{F}^n \ : \ \text{there exists a vector } x \in \mathbb{F}^m \text{ with } Ax = y\}$$

is called the *range* of A, and its dimension is called the *rank* of A.

The matrix A is called *injective* if $\mathcal{N}(A) = \{0\}$ holds, and it is called *surjective* if $\mathcal{R}(A) = \mathbb{F}^n$ holds.

We recall the *rank-nullity theorem*: let $(y_i)_{i=1}^k$ be a basis of $\mathcal{R}(A)$. By definition, we can find $(\widehat{x}_i)_{i=1}^k$ in $\mathbb{F}^m$ such that

$$A\widehat{x}_i = y_i \quad \text{for all } i \in \{1, \dots, k\}.$$

Since the family $(y_i)_{i=1}^k$ is linearly independent, the same holds for $(\widehat{x}_i)_{i=1}^k$. This means that we can expand the family to a basis $(\widehat{x}_i)_{i=1}^m$ of $\mathbb{F}^m$. For each

$j \in \{k+1, \dots, m\}$, we obviously have $A\widehat{x}_j \in \mathcal{R}(A)$ and can therefore find $z_j \in \operatorname{span}\{\widehat{x}_1, \dots, \widehat{x}_k\}$ such that $A\widehat{x}_j = Az_j$ holds, i.e., $\widehat{x}_j - z_j \in \mathcal{N}(A)$. We define

$$x_i := \begin{cases} \widehat{x}_i & \text{if } i \leq k, \\ \widehat{x}_i - z_i & \text{otherwise} \end{cases} \qquad \text{for all } i \in \{1, \dots, m\}$$

and see that $(x_i)_{i=1}^m$ is a basis of $\mathbb{F}^m$ such that $\operatorname{span}\{x_{k+1}, \dots, x_m\} \subseteq \mathcal{N}(A)$ holds. This yields $\dim \mathcal{N}(A) = m - k$, i.e.,

$$\dim \mathcal{R}(A) + \dim \mathcal{N}(A) = m \qquad \text{for all } A \in \mathbb{F}^{n \times m}. \tag{2.2}$$

Definition 2.2 (Eigenvalue and Eigenvector). Let $A \in \mathbb{F}^{n \times n}$, and let $\lambda \in \mathbb{F}$. λ is called an *eigenvalue* of A, if there is a vector $x \in \mathbb{F}^n \setminus \{0\}$ such that

$$Ax = \lambda x \tag{2.3}$$

holds. Any such vector is called an *eigenvector* of A for the eigenvalue λ. A pair (λ, x) consisting of an eigenvalue and a corresponding eigenvector is called an *eigenpair*.

The set

$$\sigma(A) := \{\lambda \in \mathbb{F} \ : \ \lambda \text{ is an eigenvalue of } A\}$$

is called the *spectrum* of A.

Let $k \in \mathbb{N}$. The *product* AB of two matrices $A \in \mathbb{F}^{n \times k}$ and $B \in \mathbb{F}^{k \times m}$ is given by

$$(AB)_{ij} = \sum_{\ell=1}^{k} a_{i\ell} b_{\ell j} \qquad \text{for all } i \in \{1, \dots, n\},\ j \in \{1, \dots, m\}. \tag{2.4}$$

The definition ensures

$$ABx = A(Bx) \qquad \text{for all } x \in \mathbb{F}^m,$$

i.e., it is compatible with the matrix-vector multiplication.

Exercise 2.3 (Polynomials). Let $A \in \mathbb{F}^{n \times n}$. We define the *$\ell$-th power* of the matrix by

$$A^\ell := \begin{cases} I & \text{if } \ell = 0, \\ AA^{\ell-1} & \text{otherwise} \end{cases} \qquad \text{for all } \ell \in \mathbb{N}_0.$$

This definition allows us to apply polynomials to matrices: for each polynomial

$$p(t) = a_0 + a_1 t + \cdots + a_m t^m,$$

we define

$$p(A) := a_0 A^0 + a_1 A^1 + \cdots + a_m A^m.$$

Prove $p(\sigma(A)) \subseteq \sigma(p(A))$. Can you find A and p with $\sigma(p(A)) \neq p(\sigma(A))$? Hint: consider Exercise 2.6.

Exercise 2.4 (Projection). Let $P \in \mathbb{F}^{n\times n}$ be a *projection*, i.e., let it satisfy $P^2 = P$. Prove $\sigma(P) \subseteq \{0, 1\}$.

Is any matrix $A \in \mathbb{F}^{n\times n}$ satisfying $\sigma(A) \subseteq \{0, 1\}$ a projection?

Hint: consider $A \in \mathbb{F}^{2\times 2}$ with $a_{21} = 0$.

Exercise 2.5 (Nil-potent matrix). Let $N \in \mathbb{F}^{n\times n}$ be a *nil-potent* matrix, i.e., let there be a $k \in \mathbb{N}$ with $N^k = 0$. Prove $\sigma(N) \subseteq \{0\}$.

Exercise 2.6 (Empty spectrum). Let $\mathbb{F} = \mathbb{R}$. Consider the matrix

$$A := \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \in \mathbb{R}^{2\times 2}.$$

Prove $\sigma(A) = \emptyset$, i.e., show that A has no eigenvalues. Does the situation change if we let $\mathbb{F} = \mathbb{C}$?

If x is an eigenvector of a matrix $A \in \mathbb{F}^{n\times n}$, multiplying x by any non-zero number will again yield an eigenvector. Instead of dealing with non-unique eigenvectors, it is preferable to use an alternative characterization of eigenvalues:

Proposition 2.7 (Null space). *Let $A \in \mathbb{F}^{n\times n}$, and let $\lambda \in \mathbb{F}$. λ is an eigenvalue of A if and only if*

$$\mathcal{N}(\lambda I - A) \neq \{0\}$$

holds, i.e., if $\lambda I - A$ is not injective.

Proof. We first observe that for all $x \in \mathbb{F}^n$ and all $\lambda \in \mathbb{F}$, the following statements are equivalent:

$$\begin{aligned} \lambda x &= Ax, \\ \lambda x - Ax &= 0, \\ (\lambda I - A)x &= 0, \\ x &\in \mathcal{N}(\lambda I - A). \end{aligned}$$

If $\lambda \in \mathbb{F}$ is an eigenvalue, we can find a corresponding eigenvector $x \in \mathbb{F}^n \setminus \{0\}$, i.e., we have $Ax = \lambda x$, and therefore $x \in \mathcal{N}(\lambda I - A)$ and $\mathcal{N}(\lambda I - A) \supseteq \{0, x\} \neq \{0\}$.

If, on the other hand, $\mathcal{N}(\lambda I - A) \neq \{0\}$ holds, we can find $x \in \mathcal{N}(\lambda I - A) \setminus \{0\}$, and this vector x is an eigenvector. □

Since the null space of $\lambda I - A$ is uniquely determined by A and λ, working with it instead of individual eigenvectors offers significant advantages both for practical and theoretical investigations.

Definition 2.8 (Eigenspace). Let $A \in \mathbb{F}^{n\times n}$, and let $\lambda \in \sigma(A)$. Then

$$\mathcal{E}(A,\lambda) := \mathcal{N}(\lambda I - A)$$

is called the *eigenspace* of A for the eigenvalue λ.

Definition 2.9 (Geometric multiplicity). Let $A \in \mathbb{F}^{n\times n}$, and let $\lambda \in \sigma(A)$. The dimension of the eigenspace $\mathcal{E}(A,\lambda)$ is called the *geometric multiplicity* of λ and denoted by $\mu_g(A,\lambda)$.

Instead of looking for individual eigenvectors, we look for a basis of an eigenspace. This offers the advantage that we can change the basis during the course of our algorithms in order to preserve desirable properties like isometry or non-degeneracy.

Exercise 2.10 (Eigenspaces). Let $A \in \mathbb{F}^{n\times n}$, and let $\lambda, \mu \in \sigma(A)$ with $\lambda \neq \mu$. Prove

$$\mathcal{E}(A,\lambda) \cap \mathcal{E}(A,\mu) = \{0\}.$$

Exercise 2.11 (Geometric multiplicity). Let $A \in \mathbb{F}^{n\times n}$. Prove

$$\sum_{\lambda\in\sigma(A)} \mu_g(A,\lambda) \leq n.$$

2.2 Characteristic polynomials

By the rank-nullity theorem (2.2), λ is an eigenvalue of a matrix A if and only if $\lambda I - A$ is not invertible. Using this property, we can characterize the eigenvalues without explicitly constructing eigenvectors.

Let $n, m \in \mathbb{N}$, and let $j \in \{1,\dots,m\}$. The *j-th canonical unit vector* $\delta_j \in \mathbb{F}^m$ is given by

$$(\delta_j)_i := \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{otherwise} \end{cases} \qquad \text{for all } i \in \{1,\dots,m\}, \tag{2.5}$$

and for a matrix $A \in \mathbb{F}^{n\times m}$, we denote the *$j$-th column vector* by

$$a_j := A\delta_j = \begin{pmatrix} a_{1j} \\ \vdots \\ a_{nj} \end{pmatrix}.$$

The matrix A is injective if and only if its columns $a_1,\dots,a_m$ are linearly independent.

We can use the *determinant* to characterize tuples of linearly independent vectors. The determinant is a mapping

$$\det : (\mathbb{F}^n)^n \to \mathbb{F}$$

of n-tuples of n-dimensional vectors to scalar values that is *multilinear*, i.e., we have

$$\begin{aligned}&\det(x_1, \ldots, x_{j-1}, x_j + \alpha z, x_{j+1}, \ldots, x_n)\\ &\quad = \det(x_1, \ldots, x_{j-1}, x_j, x_{j+1}, \ldots, x_n)\\ &\qquad + \alpha \det(x_1, \ldots, x_{j-1}, z, x_{j+1}, \ldots, x_n)\\ &\quad \text{for all } x_1, \ldots, x_n, z \in \mathbb{F}^n,\ \alpha \in \mathbb{F},\ j \in \{1, \ldots, n\}.\end{aligned}$$

The determinant is also *alternating*, i.e., we have

$$\begin{aligned}&\det(x_1, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_{j-1}, x_j, x_{j+1}, \ldots, x_n)\\ &\quad = -\det(x_1, \ldots, x_{i-1}, x_j, x_{i+1}, \ldots, x_{j-1}, x_i, x_{j+1}, \ldots, x_n)\\ &\qquad \text{for all } x_1, \ldots, x_n \in \mathbb{F}^n,\ i, j \in \{1, \ldots, n\} \text{ with } i < j\end{aligned}$$

and satisfies $\det(\delta_1, \ldots, \delta_n) = 1$.

Let $x_1, \ldots, x_n \in \mathbb{F}^n$. If there are $i, j \in \{1, \ldots, n\}$ with $i < j$ and $x_i = x_j$, the fact that the determinant is alternating implies

$$\begin{aligned}&\det(x_1, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_{j-1}, x_j, x_{j+1}, \ldots, x_n)\\ &\quad = -\det(x_1, \ldots, x_{i-1}, x_j, x_{i+1}, \ldots, x_{j-1}, x_i, x_{j+1}, \ldots, x_n)\\ &\quad = -\det(x_1, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_{j-1}, x_j, x_{j+1}, \ldots, x_n),\end{aligned}$$

and therefore $\det(x_1, \ldots, x_n) = 0$. Since the determinant is also multilinear, we can add any multiple of any argument to any other argument without changing the value of the determinant. In particular, it is possible to prove that the determinant is equal to zero if and only if its arguments are linearly dependent.

This means that a matrix $A \in \mathbb{F}^{n \times n}$ is invertible if and only if $\det(a_1, \ldots, a_n) \neq 0$, i.e., if the determinant of its column vectors vanishes. We extend the definition of the determinant to quadratic matrices by setting

$$\det : \mathbb{F}^{n \times n} \to \mathbb{F}, \quad A \mapsto \det(a_1, \ldots, a_n),$$

and have $\det(A) \neq 0$ if and only if A is invertible. Combining this property with Proposition 2.7, we obtain a characterization of eigenvalues that does not require eigenvectors:

Definition 2.12 (Characteristic polynomial). Let $A \in \mathbb{F}^{n \times n}$.

$$p_A : \mathbb{F} \to \mathbb{F}, \quad t \mapsto \det(tI - A),$$

is a polynomial of degree n. We call it the *characteristic polynomial* of A.

Proposition 2.13 (Zeros of p_A). *Let $A \in \mathbb{F}^{n\times n}$. $\lambda \in \mathbb{F}$ is an eigenvalue of A if and only if $p_A(\lambda) = 0$ holds.*

Proof. Let $\lambda \in \mathbb{F}$. If λ is an eigenvalue of A, Proposition 2.7 implies that $\lambda I - A$ is not injective, therefore this matrix has to be non-invertible, and we have $p_A(\lambda) = \det(\lambda I - A) = 0$.

If, on the other hand, we have $0 = p_A(\lambda) = \det(\lambda I - A)$, the matrix $\lambda I - A$ is non-invertible. Since it is a quadratic matrix, the rank-nullity theorem (2.2) implies that it cannot be injective, and Proposition 2.7 yields that λ has to be an eigenvalue. □

This result allows us to characterize the spectrum of a matrix $A \in \mathbb{F}^{n\times n}$ as the set of zeros of its characteristic polynomial:

$$\sigma(A) = \{\lambda \in \mathbb{F} \ : \ p_A(\lambda) = 0\}.$$

Given a polynomial p and a λ with $p(\lambda) = 0$, we can find a polynomial q such that

$$p(t) = (\lambda - t)q(t) \quad \text{for all } t \in \mathbb{F}.$$

If λ is a zero of q, we can apply this construction repeatedly to find the maximal power $k \in \mathbb{N}$ such that $(\lambda - t)^k$ is a divisor of p: the *multiplicity* of λ is the number $k \in \mathbb{N}$ uniquely defined by the property that there is a polynomial q satisfying

$$\begin{aligned} p(t) &= (\lambda - t)^k q(t) \quad \text{for all } t \in \mathbb{F}, \\ q(\lambda) &\neq 0. \end{aligned}$$

Definition 2.14 (Algebraic multiplicity). Let $A \in \mathbb{F}^{n\times n}$ and $\lambda \in \sigma(A)$. By Proposition 2.13, λ is a zero of p_A. We call its multiplicity the *algebraic multiplicity* of the eigenvalue λ and denote it by $\mu_a(A, \lambda)$.

If $\mu_a(A, \lambda) = 1$, λ is called a *simple eigenvalue*.

Exercise 2.15 (Algebraic multiplicity). Let $n \in \mathbb{N}$ and $A \in \mathbb{C}^{n\times n}$. Prove

$$\sum_{\lambda\in\sigma(A)} \mu_a(A, \lambda) = n. \tag{2.6}$$

Hint: the fundamental theorem of algebra states that every non-constant complex-valued polynomial has at least one root.

Exercise 2.16 (Companion matrix). Let $n \in \mathbb{N}_{\geq 2}$ and $c_0, c_1, \ldots, c_{n-1} \in \mathbb{F}$. Let

$$C := \begin{pmatrix} 0 & 1 & & & \\ 0 & 0 & \ddots & & \\ \vdots & \vdots & \ddots & \ddots & \\ 0 & 0 & \ldots & 0 & 1 \\ -c_0 & -c_1 & \ldots & -c_{n-2} & -c_{n-1} \end{pmatrix}.$$