Applied Linear Algebra

Carl de Boor

draft 15jan03

TABLE OF CONTENTS

overview

Here is a quick run-down on these notes, with various terms to be learned in **boldface**.

Much of scientific work involves relationships called **map**s:

$$f : X \to Y : x \mapsto f(x)$$

For example,

- time $\mapsto$ the population of the US;
- temperature $\mapsto$ pressure in a bottle;
- location (longitude, latitude, altitude) $\mapsto$ (barometric pressure, humidity, temperature);
- mother's age $\mapsto$ frequency of newborn with Down syndrom
- available resources (capital, raw materials, labor pool, etc) $\mapsto$ output of the US economy
- etc.

All this is part of our hope to understand effects in terms of causes.

Once we feel we understand such a relationship, we are eager to put it to use in order to find out how to cause certain effects. Mathematically, we are trying to solve the equation:

$$f(?) = y$$

for given $f : X \to Y$ and given $y \in Y$.

In this generality and vagueness, nothing much can be said other than to urge familiarity with basic map terms, such as, **domain**, **target** and **range** of a map, the map properties **1-1** (equivalent to **uniqueness** of solutions), **onto** (equivalent to **existence** of a solution for any y), **invertible** (equivalent to having exactly one solution for any $y \in Y$, the best-possible situation), and the notions of **left inverse**, **right inverse** and **inverse** related to the earlier notions by the concept of **map composition**.

Often, though, the map f is a **smooth** map, from some subset X of **real n-dimensional coordinate space** $\mathbb{R}^n$ to $\mathbb{R}^m$, say. With the list $x = (x_1, \ldots, x_n)$ our notation for $x \in \mathbb{R}^n$, this means that, first of all,

$$f(x) = (f_1(x), f_2(x), \ldots, f_m(x)) \ \in \mathbb{R}^m$$

with each f_j a **scalar**-valued function, and, secondly, at any point $p \in X$, we can expand each f_j into a Taylor series:

$$f_j(p+h) = f_j(p) + Df_j(p)^{\mathrm{t}}h + o(h), \quad j = 1, \ldots, m,$$

with

$$Df_j(p) = (D_1 f_j(p), \ldots, D_n f_j(p)) \ \in \mathbb{R}^n$$

the **gradient** of f_j at p, and $x^{\mathrm{t}}y$ the **scalar product** of the **n-vector**s x and y, and the $o(h)$ denoting 'higher-order' terms that we eventually are going to ignore in best scientific fashion.

This implies that

$$f(p+h) = f(p) + Df(p)h + o(h),$$

with

$$Df(p) = \begin{bmatrix} D_1 f_1(p) & \cdots & D_n f_1(p) \\ \vdots & \cdots & \vdots \\ D_1 f_m(p) & \cdots & D_n f_m(p) \end{bmatrix}$$

the **Jacobian** matrix of f at p.

With this, a standard approach to finding a solution to the equation

$$f(?) = y$$

is **Newton's method**: If x is our **current guess** at the solution, we are looking for a **correction** h so that

$$y = f(x+h) = f(x) + Df(x)h + o(h);$$

we ignore the 'higher-order' terms that hide behind the expression $o(h)$, and so get a *linear equation* for h:

$$y - f(x) = Df(x)?,$$

which we solve for h, add this correction to our current x to get a new guess

$$x \leftarrow x + h = x + Df(x)^{-1}(y - f(x))$$

and repeat. Under suitable circumstances, the process converges, to a solution.

The *key idea* here is the reduction, from solving a general equation $f(?) = y$ to solving a sequence of **linear** equations, $Df(x)? = z$. This works since, in principle, we can always solve a linear system.

Most equations $f(?) = y$ that can be solved are actually solved by this process or a variant thereof, hence the importance of knowing how to solve *linear* equations.

For this reason, our first task will be to introduce **linear map**s and **linear space**s, especially **linear spaces of functions**., i.e., vector spaces in which the basic **vector operations**, namely **vector addition** and **multiplication by a scalar**, are defined **pointwise**. These provide the proper means for expressing the concept of **linearity**. Then we recall **elimination** as the method for solving a **homogeneous** linear system

$$A? = 0$$

with $A \in \mathbb{R}^{m\times n}$. Specifically, we recall that elimination classifies the unknowns as **bound** and **free**, and this leads to **row echelon forms**, in particular the **rrref** or **really reduced row echelon form**, from which we can obtain a complete description of the solution set of $A? = 0$, i.e., for $\operatorname{null} A$, the **nullspace** of A, as well as an efficient description of $\operatorname{ran} A$, the **range** of A. Thus equipped, we deal with the general linear system $A? = b$ via the homogeneous linear system $[A, b]? = 0$.

Both $\operatorname{null} A$ and $\operatorname{ran} A$ are typical examples of **linear subspace**s, and these efficient descriptions for them are in terms of a **basis**, i.e., in terms of an invertible linear map V from some **coordinate space** $\mathbb{F}^n$ to the linear subspace in question. This identifies bases as particular **column map**s, i.e., linear maps from some coordinate space, i.e., maps of the form

$$\mathbb{F}^n \to X : a \mapsto a_1v_1 + \cdots a_nv_n =: [v_1, \ldots, v_n]a$$

for some sequence $v_1, \ldots, v_n$ in the linear space X in question.

We'll spend some time recalling various details about bases, how to construct them, how to use them, and will also mention their generalization, **direct sum**s and their associated **linear projector**s or **idempotent**s. We stress the notion of **dimension** (= number of columns or elements in a basis), in particular the **Dimension Formula**

$$\dim \operatorname{dom} A = \dim \operatorname{ran} A + \dim \operatorname{null} A,$$

valid for any linear map A, which summarizes much of what is important about dimension.

We'll also worry about how to determine the **coordinates** of a given $x \in X$ with respect to a given basis V for X, i.e., how to solve the equation

$$V? = x.$$

This will lead us to **row map**s, i.e., linear maps from some linear space to coordinate space, i.e., maps of the form

$$X \to \mathbb{F}^n : x \mapsto (\lambda_1 x, \cdots, \lambda_n x) =: [\lambda_1, \ldots, \lambda_n]^{\mathrm{t}} x$$

for some sequence $\lambda_1, \ldots, \lambda_n$ of **linear functional**s on the linear space X in question. It will also lead us to **interpolation** aka **change of basis**, and will make us single out **inner product space**s as spaces with a ready supply of suitable row maps, and thence to **least-squares**, to particularly good bases, namely **o.n.** (:= **orthonormal**) bases (which are the **isometries** for the standard **norm**, the **Euclidean norm** $\|x\|_2 = \sqrt{x^t x}$ associated with the standard **inner product** and which can be constructed from an arbitrary basis by **Gram-Schmidt**).

We'll find that bases also show up naturally when we try to **factor** a given linear map $A \in L(X, Y)$ in the most efficient way, as a product

$$A = V\Lambda^t$$

with $\Lambda^t \in L(X, \mathbb{F}^r)$ and $V \in L(\mathbb{F}^r, Y)$ and r as small as possible. It will be one of my tasks to convince you that you have actually carried out such factorizations, in fact had to do this in order to do certain standard operations, like differentiating or integrating polynomials and other functions. Such factorizations are intimately connected with the **rank** of A (since the smallest possible r is the rank of A) and lead, for a matrix A, to the **SVD**, or **Singular Value Decomposition**,

$$A = V\Sigma W^c$$

with V, W o.n. and Σ diagonal, a factorization that is, in a certain sense, a best way of describing the action of the linear map A. Other common factorizations for matrices are the **PLU factorization** with P a **permutation matrix**, L **unit lower triangular**, and U **upper triangular** (generated during elimination); and the (more stable) **QR factorization**, with Q **unitary** (i.e., an o.n. basis) and R **upper**, or, **right triangular**, obtained by elimination with the aid of specific **elementary matrices** called **Householder reflections**.

For *square* matrices, one hopes to (but does not always) get factorizations of the form $A = V\Sigma V^{-1}$ with Σ diagonal (the simplest example of a matrix without such a factorization is the **nilpotent** matrix $\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$), but often must be (and is) content to get the **Schur form**, which is available for any square matrix and is of the form $A = VUV^c$ with V an o.n. basis and U upper triangular. In either case, A is then said to be **similar** to Σ and U, respectively. These latter factorizations, or **similarities**, are essential for an understanding of the **power sequence**

$$A^0 = \text{ id}, A^1 = A, A^2 = AA, A^3 = AAA, ...$$

of the square matrix A and, more generally, for an understanding of the **matrix polynomial** $p(A)$, since, e.g.,

$$A = V\,\text{diag}(\mu_1, \ldots, \mu_n)V^{-1} \quad \Longrightarrow \quad p(A) = V\,\text{diag}(p(\mu_1), \ldots, p(\mu_n))V^{-1},$$

for any polynomial p and even for some well-behaved functions p like the **exponential** $p : t \mapsto \exp(t)$. In particular, then

$$A^k = V\,\text{diag}(\mu_1^k, \ldots, \mu_n^k)V^{-1}, \quad k = 0, 1, 2, \ldots,$$

therefore we can describe the behavior of the *matrix* sequence $(A^k : k = 0, 1, 2, \ldots)$ entirely in terms of the *scalar* sequences $(\mu_j^k : k = 0, 1, 2, \ldots)$. Specifically, we can characterize **power-boundedness**, **convergence**, and **convergence to** 0.

There are many reasons for wanting to understand the power sequence of a matrix; here is one. Often, elimination is not the most efficient way to solve a linear system. Rather, the linear system

$$A? = b$$

itself is solved by iteration, by splitting $A =: M - N$ with M 'easily' invertible, and looking at the **equivalent equation**

$$M? = N? + b$$

which leads to the **iteration**

$$x \leftarrow M^{-1}(Nx + b) =: Bx + c.$$

Convergence of this process depends crucially on the behavior of the power sequence for B (and does not at all depend on the particular **vector norm** or **map norm** used).

The factorization

$$A = V \operatorname{diag}(\mu_1, \ldots, \mu_n) V^{-1}$$

is equivalent to having $AV = V \operatorname{diag}(\mu_1, \ldots, \mu_n)$, i.e.,

$$[Av_1, \ldots, Av_n] = [\mu_1 v_1, \ldots, \mu_n v_n]$$

for some invertible $V = [v_1, \ldots, v_n] : \mathbb{F}^n \to \operatorname{dom} A$, i.e., to having a basis V consisting of **eigenvector**s for A, with the μ_j the corresponding **eigenvalue**s. For this reason, we'll study the **eigenstructure** of A and the **spectrum** of A, as well as **similarity**, i.e., the **equivalence relation**

$$A \sim C \ := \ \exists V \ \ A = VCV^{-1}.$$

In this study, we make use of polynomials, particular the **annihilating polynomial**s (which are the nontrivial polynomials p for which $p(A) = 0$) and their cousins, the nontrivial polynomials p for which $p(A)x = 0$ for some $x \neq 0$, and the unique **monic** annihilating polynomial of minimal degree, called the **minimal polynomial** for A, as well as the **Krylov sequence** $x, Ax, A^2x, \ldots$.

We'll discuss the most important classification of eigenvalues, into **defective** and **non-defective** eigenvalues, and give a complete description of the asymptotic behavior of the power sequence $A^0, A^1, A^2, \ldots$ in terms of the eigenstructure of A, even when A is not **diagonalizable**, i.e., is not similar to a diagonal matrix (which is equivalent to some eigenvalue of A being defective).

We'll also discuss standard means for locating the spectrum of a matrix, such as **Gershgorin's circles** and the **characteristic polynomial** of a matrix, and give the Perron-Frobenius theory concerning the dominant eigenvalue of a positive matrix.

From the Schur form (*vide supra*), we derive the basic facts about the eigenstructure of **hermitian** and of **normal** matrices. We give the **Jordan form** only because of its mathematical elegance since, in contrast to the Schur form, it cannot be constructed reliably numerically.

As a taste of the many different applications of Linear Algebra, we discuss briefly: the solution of a system of constant-coefficient ODEs, Markov processes, subdivision in CAGD, Linear Programming, the Discrete Fourier Transform, approximation by broken lines, and the use of flats in analysis and CAGD.

Further, we also consider briefly **minimization** of a real-valued map

$$f : K \to \mathbb{R}$$

with $K \subset \mathbb{R}^n$. Returning to our Taylor expansion

$$f(p+h) = f(p) + Df(p)^{\mathrm{t}} h + o(h),$$

we notice that, usually, p cannot be a minimum point for f unless it is a **critical point**, i.e., unless the gradient, $Df(p)$, is the zero vector. However, even with $Df(p) = 0$, we only know that f is 'flat' at p. In particular, a critical point could also be a (local) maximum point, or a saddle point, etc. To distinguish between the various possibilities, we must look at the **second-order** terms, i.e., we must write and know, more explicitly, that

$$f(p+h) = f(p) + Df(p)^{\mathrm{t}} h + h^{\mathrm{t}} D^2 f(p) h/2 + o(h^{\mathrm{t}} h),$$

with

$$H := D^2 f := \begin{bmatrix} D_1 D_1 f & \cdots & D_1 D_n f \\ \vdots & \ldots & \vdots \\ D_n D_1 f & \cdots & D_n D_n f \end{bmatrix}$$

the **Hessian** for f, hence

$$h \mapsto h^{\mathrm{t}} D^2 f(p) h = \sum_{i,j} D_i D_j f(p) h_i h_j$$

the associated **quadratic form**.

We will learn to distinguish between maxima, minima, and saddle points by the signs of the eigenvalues of the Hessian, mention **Sylvester's Law of Inertia**, and show how to estimate the effect of **perturbations** on H on the spectrum of H, using ideas connected with the **Rayleigh quotient**.

At this point, you will realize that these notes are strongly influenced by the use of Linear Algebra in Analysis, with important applications, e.g., in Graph Theory, ???, or ???, being ignored (partly through ignorance).

Finally, although **determinant**s have little to contribute to Linear Algebra at the level of this book, we'll give a complete introduction to this very important Linear Algebra tool, and then discuss the **Schur complement**, **Sylvester's determinant identity**, and the **Cauchy-Binet formula**.

Throughout, we'll rely on needed material from prerequisite courses as collected in an appendix called **Background**.

this page purposely put in here

1. Sets, assignments, lists, and maps

The basic objects of Mathematics are sets and maps. Linear Algebra is perhaps the first course where this fact becomes evident and where it can be illustrated in a relative straightforward context. Since a complete understanding of the course material requires a thorough appreciation of the basic facts about maps, we begin with these and their simpler cousins, lists and assignments, after a brief review of standard language and notation concerning sets.

Sets

Sets of interest in these notes include

- the **natural numbers** : $\mathbb{N} := \{1, 2, \ldots\}$;
- the **integers** : $\mathbb{Z} := \{\ldots, -1, 0, 1, \ldots\} = (-\mathbb{N}) \cup \{0\} \cup \mathbb{N}$;
- the nonnegative integers : $\mathbb{Z}_+ := \{p \in \mathbb{Z} : p \geq 0\}$;
- the **rational numbers** : $\mathbb{Z} \div \mathbb{N} := \{p/q : p \in \mathbb{Z}, q \in \mathbb{N}\}$;
- the **real numbers** and the nonnegative reals : $\mathbb{R}$, $\qquad \mathbb{R}_+ := \{x \in \mathbb{R} : x \geq 0\}$;
- the **complex numbers** : $\mathbb{C} := \mathbb{R} + \mathrm{i}\mathbb{R} = \{x + \mathrm{i}y : x, y \in \mathbb{R}\}$, $\qquad \mathrm{i} := \sqrt{-1}$.

As these examples show, a set is often specified in the form $\{x : P(x)\}$ which is read 'the set of all x that have the property $P(x)$'. Note the use of the colon, ':', (rather than a vertical bar, '|') to separate, the initial, provisional, description of the typical element of the set, from the conditions imposed on it for membership in the set. In these notes, braces, '{', '}', are used solely in the description of sets.

Standard notation concerning sets includes:

- $\#S$ denotes the **cardinality** of the set S, i.e., the count of its elements.
- $x \in S$ and $S \ni x$ both mean that x is an element of S.
- $S \subset T$, $T \supset S$ both mean that S is a **subset** of T, i.e., all the elements of S are also elements of T; if we want to convey that S is a **proper subset** of T, meaning that $S \subset T$ but $S \neq T$, we write $S \subsetneq T$.
- $\{\}$ denotes the **empty set**, the set with no elements.
- $S \cap T := \{x : x \in S \text{ and } x \in T\}$ is the **intersection** of S and T.
- $S \cup T := \{x : x \in S \text{ or } x \in T\}$ is the **union** of S and T.
- $S \backslash T := \{x : x \in S \text{ but not } x \in T\}$ is the **difference** of S from T and is often read 'S take away T'. In these notes, this difference is *never* written $S - T$, as the latter is reserved for the set $\{s - t : s \in S, t \in T\}$ formable when both S and T are subsets of the same vector space.

1.1 What is the standard name for the elements of $\mathbb{R}\backslash(\mathbb{Z} \div \mathbb{N})$?

1.2 What is the standard name for the elements of $\mathrm{i}\mathbb{R}$?

1.3 Work out each of the following sets. (a) $(\{-1, 0, 1\} \cap \mathbb{N}) \cup \{-2\}$; (b) $(\{-1, 0, 1\} \cup \{-2\}) \cap \mathbb{N}$; (c) $\mathbb{Z}\backslash(2\mathbb{Z})$; (d) $\{z^2 : z \in \mathrm{i}\mathbb{R}\}$.

1.4 Determine $\#((\mathbb{R}_+\backslash\{x \in \mathbb{R} : x^2 > 16\}) \cap \mathbb{N})$.

Assignments

Definition: An **assignment** or, more precisely, an **assignment on** I or I**-assignment**

$$f = (f_i)_{i \in I} = (f_i : i \in I)$$

associates with each element i in its **domain** (or, **index set**)

$$\operatorname{dom} f := I$$

some **term** or **item** or **entry** or **value** f_i. In symbols:

$$f : \operatorname{dom} f : i \mapsto f_i.$$

The set

$$\operatorname{ran} f := \{f_i : i \in \operatorname{dom} f\}$$

of all items appearing in the assignment f is called the **range** of the assignment.

If also g is an assignment, then $f = g$ exactly when $f_i = g_i$ for all $i \in \operatorname{dom} f = \operatorname{dom} g$.

Very confusingly, many mathematicians call an assignment an *indexed set*, even though it is most certainly not a set. The term **family** is also used; however it, too, smacks too much of a set or collection.

We call the assignment f **1-1** if $f_i = f_j \Longrightarrow i = j$.

The simplest assignment is the **empty assignment**, $()$, i.e., the unique assignment whose domain is the empty set. Note that the empty assignment is 1-1 (why??).

An assignment with domain the set

$$\underline{n} := \{1, 2, \ldots, n\}$$

of the first n natural numbers is called a **list**, or, more explicitly, an n**-list**.

To specify an n-list f, it is sufficient to list its terms or values:

$$f = (f_1, f_2, \ldots, f_n).$$

For example, the **cartesian product**

$$\times_{i=1}^n X_i := X_1 \times X_2 \times \cdots \times X_n := \{(x_1, x_2, \ldots, x_n) : x_i \in X_i, i = 1{:}n\}$$

of the set sequence $X_1, \ldots, X_n$ is, by definition, the collection of all n-lists with the ith item or **coordinate** taken from X_i, all i.

In these notes, we deal with n*-vectors*, i.e., n-lists of *numbers*, such as the 3-lists $(1, 3.14, -14)$ or $(3, 3, 3)$. (Note that the *list* $(3, 3, 3)$ is quite different from the *set* $\{3, 3, 3\}$. The list $(3, 3, 3)$ has three terms, while the set $\{3, 3, 3\}$ has exactly one element.)

Definition: An n**-vector** is a list of n scalars (numbers). The collection of all **real** (**complex**) n-vectors is denoted by $\mathbb{R}^n$ ($\mathbb{C}^n$).

In MATLAB, there are (at least) two ways to specify an n-vector, namely as a one-row matrix (colloquially known as a **row vector**), or as a one-column matrix (colloquially known as a **column vector**). For example, we can record the 3-vector $x = (1.3, 3.14, -15)$ as the one-row matrix

```
x_as_row = [1.3,3.14,-15];
```

or as the one-column matrix

```
x_as_col = [1.3;3.14;-15];
```

One can also write a one-column matrix as a column, without the need for the semicolons, e.g.,

```
x_as_col = [1.3
     3.14
   -15];
```

□

Back to general assignments. If $\operatorname{dom} f$ is finite, say $\#\operatorname{dom} f = n$, then we could always describe f by listing the n pairs (i, f_i), $i \in \operatorname{dom} f$, in some fashion. However, that may not always be the most helpful thing to do. Here is a famous example.

During the Cholera outbreak in 1854 in London, Dr. John Snow recorded the deaths by address, thus setting up an assignment whose domain consisted of all the houses in London. But he did not simply make a list of all the addresses and then record the deaths in that list. Rather, he took a map of London and marked the number of deaths at each address right on the map (not bothering to record the value 0 of no death). He found that the deaths clustered around one particular public water pump, jumped to a conclusion (remember that this was well before Pasteur's discoveries), had the handle of that pump removed and had the satisfaction of seeing the epidemic fade.

Thus, one way to think of an assignment is to visualize its domain in some convenient fashion, and, 'at' each element of the domain, its assigned item or value.

This is routinely done for matrices, another basic object in these notes.

1.5 In some courses, students are assigned to specific seats in the class room. (a) If you were the instructor in such a class, how would you record this seating assignment? (b) What are the range and domain of this assignment?

1.6 A **relation** between the sets X and Y is any subset of $X \times Y$. Each such relation relates or associates with some elements of X one or more elements of Y. For each of the following relations, determine whether or not it provides an assignment on the set $X := \underline{3} =: Y$. (i) $R = X \times Y$; (ii) $R = \{(x, x) : x \in X\}$; (iii) $R = \{(1, 2), (2, 2)\}$; (iv) $R = \{(1, 2), (2, 1)\}$; (v) $R = \{(1, 2), (3, 1), (2, 1)\}$; (vi) $R = \{(1, 2), (2, 2), (3, 1), (2, 1)\}$.

Matrices

Definition: A **matrix**, or, more precisely, an $m \times n$**-matrix**, is any assignment with domain the cartesian product

$$\underline{m} \times \underline{n} \;=\; \{(i, j) : i \in \underline{m}, j \in \underline{n}\}$$

of $\underline{m}$ with $\underline{n}$, for some nonnegative m and n.

The collection of all **real**, resp. **complex** $m \times n$-matrices is denoted by $\mathbb{R}^{m \times n}$, resp. $\mathbb{C}^{m \times n}$.

In other words, a matrix has a rectangular domain. Correspondingly, it is customary to display such an $m \times n$-matrix A as a rectangle of items:

$$A \;=\; \begin{bmatrix} A_{1,1} & A_{1,2} & \cdots & A_{1,n} \\ A_{2,1} & A_{2,2} & \cdots & A_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ A_{m,1} & A_{m,2} & \cdots & A_{m,n} \end{bmatrix}$$

rather than as a list of pairs. This means that we must think of its domain rotated clockwise 90° when compared to the ordinary $(x, y)-$plane, i.e., the domain of many other bivariate assignments (or maps).

This way of displaying a matrix has led to the following language.

Let A be an $m \times n$-matrix. The item

$$A_{i,j}$$

corresponding to the index (i, j) is also called the **(i, j)-entry** of A. The list $A_{i:} := (A_{i,j} : j \in \underline{n})$ is called the **ith row of** A, the list $A_{:j} := (A_{i,j} : i \in \underline{m})$ is called the **jth column of** A, and the list $(A_{ii} = A_{i,i} : 1 \le i \le \min\{m, n\})$ is called the **(main) diagonal of** A.

A matrix with nonzero entries only on or above (below) the diagonal is called **upper triangular** (**lower triangular**). A **diagonal matrix** is one that is both upper and lower triangular.

By definition, A^{t} denotes the **transpose** of the matrix A, i.e., the $n \times m$-matrix whose (i, j)-entry is A_{ji}, all i, j. Because of its importance in the later parts of these notes, we usually use the **conjugate transpose** $A^{\mathrm{c}} := \overline{A}^{\mathrm{t}}$ whose (i, j)-entry is the scalar $\overline{A_{ji}}$, with $\overline{\alpha}$ the complex conjugate of the scalar α.

When $m = n$, A is called a **square matrix** of **order** n.

The notation $A_{i:}$ for the ith row and $A_{:j}$ for the jth column of the matrix A is taken from `MATLAB`, where, however, `A(i,:)` is a one-row matrix and `A(:,j)` is a one-column matrix (rather than just a vector). The (main) diagonal of a matrix `A` is obtained in `MATLAB` by the command `diag(A)`, which returns, in a one-column matrix, the list of the diagonal elements. The upper (lower) triangular part of a matrix `A` is provided by the command `triu(A)` (`tril(A)`). The conjugate transpose of a matrix `A` is obtained by `A'`. This is the same as the transpose if `A` is real. To get the mere *transpose* $\mathtt{A}^{\mathrm{t}}$ in the contrary case, you must use the notation `A.'` which is strange since there is nothing *pointwise* about this operation.

The above-mentioned need to look at displays of matrices sideways is further compounded when we use `MATLAB` to plot a matrix. Here, for example, is the 'picture' of the 8×16-matrix $A :=$ `eye(8,16)` as generated by the command `mesh(eye(8,16))`. This matrix has all its diagonal entries equal to 1 and all other entries equal to 0. But note that a careless interpretation of this figure would lead one to see a matrix with 16 rows and only 8 columns, due to the fact that `MATLAB`'s `mesh(A)` command interprets `A(i,j)` as the value of a bivariate function at the point `(j,i)`.

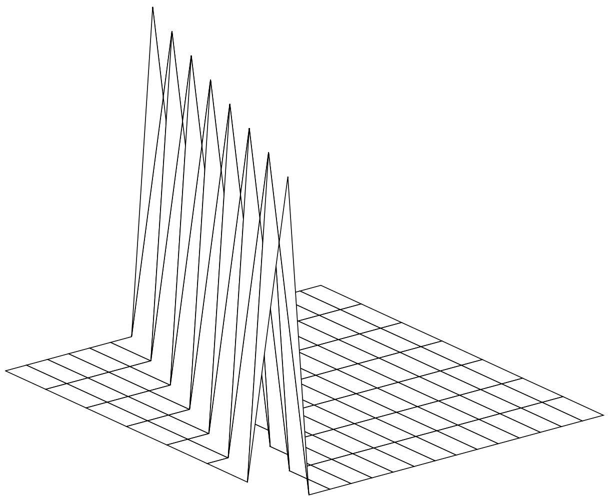

The rectangular identity matrix `eye(8,16)` as plotted in `MATLAB`

□

While lists can be concatenated in just one way, by letting one follow the other, matrices can be 'concatenated' by laying them next to each other and/or one underneath the other. The only requirement is that the result be again a matrix. If, for example,

$$A := [1 \quad 2], \qquad B := \begin{bmatrix} 3 \\ 6 \\ 9 \end{bmatrix} \qquad C := \begin{bmatrix} 4 & 5 \\ 7 & 8 \end{bmatrix},$$

then there are four different ways to 'concatenate' these three matrices, namely

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

In MATLAB, one would write the three matrices

```
A = [1 2];   B = [3;6;9];   C = [4 5; 7 8];
```

and would describe the four possible 'concatenations'

```
[[A;C],B];   [[C;A],B];   [B,[A;C]];   [B,[C;A]];
```

We saw earlier that even vectors are described in MATLAB by matrices since MATLAB only knows matrices.

□

1.7 For the matrix A given by `[[0 0 0 0];eye(2,4)]`, determine the following items: (a) the main diagonal; (b) the second column; (c) the third row; (d) $A_{3,2}$; (e) A^{t}; (f) A^{c}; (g) is A lower or upper triangular?

Lists of lists

Matrices are often used to record or represent a list $f = (f_1, f_2, \ldots, f_n)$ in which all the items f_j are themselves lists. This can always be done if all the items f_j in that list have the same length, i.e., for some m and all j, $\#f_j = m$. Further, it can be done in two ways, by columns or by rows.

Offhand, it seems more natural to think of a matrix as a list of its rows, particularly since we are used to writing things from left to right. Nevertheless, in these notes, it will always be done by columns, i.e., the sequence $(f_1, f_2, \ldots, f_n)$ of m-vectors will be associated with the $m \times n$-matrix A whose jth column is f_j, all j. We write this fact in this way:

$$A = [f_1, f_2, \ldots, f_n]; \qquad \text{i.e., } A_{:j} = f_j,\ j = 1{:}n.$$

This makes it acceptable to denote by

$$\#A$$

the number of columns of the matrix A. If I need to refer to the number of rows of A, I will simply count the number of columns of its transpose, A^{t}, or its conjugate transpose, A^{c}, i.e., write

$$\#A^{\mathrm{t}} \quad \text{or} \quad \#A^{\mathrm{c}},$$

rather than introduce yet another notation.

Here is a picturesque example of a list of lists, concerning the plotting of a **polyhedron**, specifically the regular **octahedron**. Its vertex set consists of the three unit vectors and their negatives, i.e.:

```
vs = [1 0 0; -1 0 0; 0 1 0; 0 -1 0; 0 0 1; 0 0 -1]';
```


Each face is a triangle, and we specify it here by giving the index, in the vertex array `vs`, of each of its three vertices:

```
ff = [2 4 5; 2 5 3; 4 1 5; 2 6 4]';
bf = [2 3 6; 6 1 4; 6 3 1; 5 1 3]';
```

The faces have been organized into front faces and back faces, in anticipation of the plotting about to be done, in which we want to plot the front faces strongly, but only lightly indicate the back faces. Be sure to look for specific faces in the figure below, in which the six vertices are numbered as in `vs`. E.g., the first front face, specified by the first column of `ff`, involves the vertices numbered `2, 4, 5`; it is the face we are viewing head-on.

First, we set the frame:

```
axis([-1 1 -1 1 -1 1])
hold on, axis off
```

Then we plot the back-faces first (using `r=[1 2 3 1]` to make sure that we plot closed triangles):

```
r = [1 2 3 1];
for j=1:4
   plot3(vs(1,bf(r,j)),vs(2,bf(r,j)),vs(3,bf(r,j)),':')
end
```

Then, finally, we plot the front faces and finish the picture:

```
for j=1:4
   plot3(vs(1,ff(r,j)),vs(2,ff(r,j)),vs(3,ff(r,j)),'linew',1.5);
end
hold off
```

Here is the resulting figure (obtained by the command `print -deps2 figoctah.eps` which generates a `postscript` file). I have labeled all the vertices by their index in the vertex list `vs`.

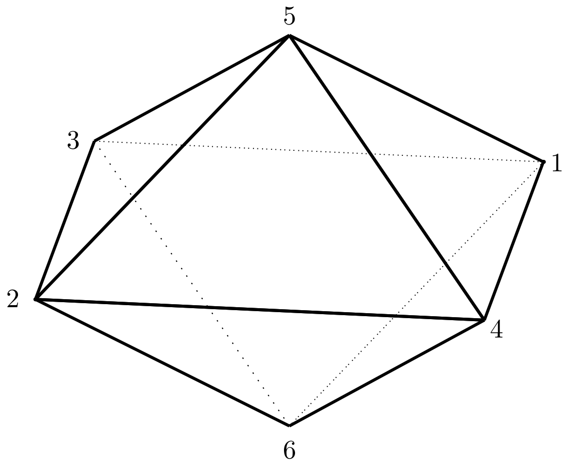

The regular octahedron.

1.8 The regular octahedron is one of five regular solids. Write a `MATLAB` function `drawrs(n)` that will, for input $n \in (1, 2, 3, 4, 5)$, draw the regular (tetrahedron, cube, octahedron, dodecahedron, icosahedron).

□

Maps

Definition: A **map**

$$f : X \to Y : x \mapsto f(x)$$

associates with each element x of its **domain** $\operatorname{dom} f := X$ a unique element $y = f(x)$, called the **value of f at x**, from its **target** $\operatorname{tar} f := Y$. If g is also a map, then $f = g$ means that $\operatorname{dom} f = \operatorname{dom} g$, $\operatorname{tar} f = \operatorname{tar} g$, and $f(x) = g(x)$ for all $x \in \operatorname{dom} f$.

The collection

$$\operatorname{ran} f := \{f(x) : x \in X\}$$

of all values taken by f is called the **range of f**. More generally, for any subset Z of X,

$$fZ := f(Z) := \{f(z) : z \in Z\}$$

is called the **image of Z under f**. In these terms,

$$\operatorname{ran} f = f(\operatorname{dom} f).$$

Also, for any $U \subset Y$, the set

$$f^{-1}U := \{x \in X : f(x) \in U\}$$

is called the **pre-image of U under f**. The collection of all maps from X to Y is denoted by

$$Y^X \quad \text{or} \quad (X \to Y).$$

Names other than **map** are in use, such as **mapping**, **operator**, **morphism**, **transformation** etc., all longer than ‘map’. A scalar-valued map is often called a **function**. Somewhat confusingly, many mathematicians use the term ‘range’ for what we have called here ‘target’; the same mathematicians use the term **image** for what we have called here ‘range’.

Every map $f : X \to Y$ gives rise to an assignment on X, namely the assignment $(f(x) : x \in X)$. On the other hand, an assignment f on X gives rise to *many* maps, one for each Y that contains $\operatorname{ran} f$, by the prescription $X \to Y : x \mapsto f_x$. We call this the **map into Y given by the assignment f**.

If X is empty, then Y^X consists of exactly one element, namely the map given by the empty assignment, and this even holds if Y is empty.

However, if Y is empty and X is not, then there can be no map from X to Y, since any such map would have to associate with each $x \in X$ some $y \in Y$, yet there are no $y \in Y$ to associate with.

“Wait a minute!”, you now say, “How did we manage when X was empty?” Well, if X is empty, then there is no $x \in X$, hence the question of what element of Y to associate with never comes up. Isn’t Mathematics slick?

1.9 Which of the following lists of pairs describes a map from {o,u,i,a} to {t,h,s}? A: ((u,s), (i,s), (a,t), (o,h), (i,s)); B: ((i,t), (a,s), (o,h), (i,s), (u,s)); C: ((a,s), (i,t), (u,h), (a,s), (i,t)).

1.10 For each of the following `MATLAB` maps, determine their range, as maps on real 2-by-3 matrices: (a) `A` $\mapsto$ `max(A)`; (b) `A` $\mapsto$ `A(:,2)`; (c) `A` $\mapsto$ `diag(A)`; (d) `A` $\mapsto$ `size(A)`; (e) `A` $\mapsto$ `length(A)`; (f) `A` $\mapsto$ `cos(A)`; (g) `A` $\mapsto$ `ones(A)`; (h) `A` $\mapsto$ `sum(A)`.

1.11 The **characteristic function** χ_S of the subset S of the set T is, by definition, the function on T that is 1 on S and 0 otherwise:

$$\chi_S : T \to \{0,1\} : t \mapsto \begin{cases} 1, & \text{if } t \in S; \\ 0, & \text{otherwise.} \end{cases}$$

Let R and S be subsets of T. Prove that (a) $\chi_{R\cup S} = \max(\chi_R, \chi_S)$; (b) $\chi_{R\cap S} = \min(\chi_R, \chi_S) = \chi_R \chi_S$; (c) $\chi_{R\backslash S} = \chi_R(1-\chi_S)$. (d) $R \subset S$ iff $\chi_R \le \chi_S$.

1.12 Let $f : T \to U$, and consider the map from subsets of U to subsets of T given by the rule

$$R \mapsto f^{-1}R := \{t \in T : f(t) \in R\}.$$

Prove that this map commutes with the set operations of union, intersection and 'take away', i.e., for any subsets R and S of U, (a) $f^{-1}(R \cup S) = (f^{-1}R) \cup (f^{-1}S)$; (b) $f^{-1}(R \cap S) = (f^{-1}R) \cap (f^{-1}S)$; (c) $f^{-1}(R \backslash S) = (f^{-1}R) \backslash (f^{-1}S)$.

1-1 and onto

In effect, a map is an assignment together with a target, with the target necessarily containing the range of the assignment. A major reason for introducing the concept of *map* (as distinct from the notion of *assignment*) is in order to raise the following basic question:

Given the map $f : X \to Y$ and $y \in Y$, find $x \in X$ for which $f(x) = y$, i.e., solve the equation

$$f(?) = y. \tag{1.1}$$

Existence occurs if this equation has a solution for every $y \in Y$, i.e., if $\operatorname{ran} f = \operatorname{tar} f$. **Uniqueness** occurs if there is at most one solution for every $y \in Y$, i.e., if $f(x) = f(z)$ implies that $x = z$, i.e., the assignment $(f(x) : x \in X)$ is 1-1.

Here are the corresponding map properties:

Definition: The map $f : X \to Y$ is **onto** in case $\operatorname{ran} f = Y$.

Definition: The map $f : X \to Y$ is **1-1** in case $f(x) = f(y) \Longrightarrow x = y$.

Not surprisingly, these two map properties will play a major role throughout these notes. (At last count, '1-1' appears over 300 times in these notes, and 'onto' over 200 times.) − There are other names in use for these properties: An onto map is also called **surjective** or **epimorph(ic)**, while a 1-1 map is also called **injective** or **monomorph(ic)**.

Perhaps the simplest useful examples of maps are those derived from lists, i.e., maps from some $\underline{n}$ into some set Y. Here is the basic observation concerning such maps being 1-1 or onto.

(1.2) If $g : \underline{n} \to Y$ is 1-1 and $f : \underline{m} \to Y$ is onto, then $n \le m$, with equality if and only if g is also onto and f is also 1-1.

Proof: The sequence $(f(1), \dots, f(m))$ contains every element of Y, but may also contain duplicates of some. Throw out all duplicates to arrive at the sequence $(h(1), \dots, h(q))$ which still contains all elements of Y but each one only once. In effect, we have 'thinned' f to a map $h : \underline{q} \to Y$ that is still onto but also 1-1. In particular, $q \le m$, with equality if and only if there were no duplicates, i.e., f is also 1-1.

Now remove from $(h(1), \dots, h(q))$ every entry of the sequence $(g(1), \dots, g(n))$. Since h is onto and 1-1, each of the n distinct entries $g(j)$ does appear in h's sequence exactly once, hence the remaining sequence $(k(1), \dots, k(r))$ has length $r = q - n$. Thus, $n \le q$, with equality, i.e., with $r = 0$, if and only if g is onto. In any case, the concatenation $(g(1), \dots, g(n), k(1), \dots, k(r))$ provides an 'extension' of the 1-1 map g to a map to Y that is still 1-1 but also onto.

Put the two arguments together to get that $n \le q \le m$, with equality if and only if f is also 1-1 and g is also onto. □

Note the particular conclusion that if both $g : \underline{n} \to Y$ and $f : \underline{m} \to Y$ are 1-1 and onto, then necessarily $n = m$. This number is called the **cardinality** of Y and is denoted

$$\#Y.$$

Hence, if we know that $\#Y = n$, i.e., that there is some invertible map from $\underline{n}$ to Y, then we know that any map $f : \underline{n} \to Y$ is onto if and only if it is 1-1. This is the

(1.3) Pigeonhole principle: If $f : \underline{n} \to Y$ with $\#Y = n$, then f is 1-1 if and only if f is onto.

Any map from $\underline{n}$ to $\underline{n}$ that is 1-1 and onto is called a **permutation of order** n since its list is a reordering of the first n integers. Thus $(3, 2, 1)$ or $(3, 1, 2)$ are permutations of order 3 while the map into $\underline{3}$ given by the 3-vector $(3, 3, 1)$ is not a permutation, as it is neither 1-1 nor onto.

By the Pigeonhole principle, in order to check whether an n-vector represents a permutation, we only have to check whether its range is $\underline{n}$ (which would mean that it is onto, as a map into $\underline{n}$), or we only have to check whether all its values are different and in $\underline{n}$ (which would mean that it is a 1-1 map into its domain, $\underline{n}$).

The finiteness of $\underline{n}$ is essential here. For example, consider the **right shift**

$$r : \mathbb{N} \to \mathbb{N} : n \mapsto n + 1. \tag{1.4}$$

This maps different numbers to different numbers, i.e., is 1-1, but fails to be onto since the number 1 is not in its range. On the other hand, the **left shift**

$$l : \mathbb{N} \to \mathbb{N} : n \mapsto \max\{n - 1, 1\} \tag{1.5}$$

is onto, but fails to be 1-1 since it maps both 1 and 2 to 1.

In light of this example, it is all the more impressive that such a pigeonhole principle continues to hold for certain special maps $f : X \to Y$ with both X and Y infinite. Specifically, according to (4.16)Corollary, if X and Y are *vector spaces* of the same finite *dimension* and $f : X \to Y$ is a *linear* map, then f is 1-1 if and only f is onto. This result is one of the high points of basic linear algebra. A more down-to-earth formulation of it, as in (3.16)Theorem, is the following: *A linear system with as many equations as unknowns has a solution for every right-hand side if and only if it has only the trivial solution when the right-hand side is 0.*

1.13 Prove: any $g : \underline{n} \to Y$ with $n > \#Y$ cannot be 1-1.

1.14 Prove: any $f : \underline{m} \to Y$ with $m < \#Y$ cannot be onto.

1.15 Let $g : \underline{n} \to Y$ be 1-1, and $f : \underline{m} \to Y$ be onto. Prove that

(i) for some $k \ge n$, g can be 'extended' to a map $h : \underline{k} \to Y$ that is 1-1 and onto;

(ii) for some $k \le m$, f can be 'thinned' to a map $h : \underline{k} \to Y$ that is onto and 1-1.

1.16 Prove: *If T is finite and $S \subset T$, then S is finite, too.* (Hint: consider the set N of all $n \in \mathbb{N} \cup \{0\}$ for which there is a 1-1 map $g : \underline{n} \to S$.)

1.17 Prove that $S \subset T$ and $\#T < \infty$ implies that $\#S \le \#T$, with equality if and only if $S = T$.

Some examples

The next simplest maps after those given by lists are probably those that come to you in the form of a *list of pairs*. For example, at the end of the semester, I am forced to make up a grade map. The authorities send me the domain of that map, namely the students in this class, in the form of a list, and ask me to assign, to each student, a grade, thus making up a list of pairs of the form

$$\text{name} \quad | \quad \text{grade}$$

Here at UW, the target of the grade map is the set

$$\{\text{A, AB, B, BC, C, D, F, I}\},$$

but there is no requirement to make this map onto. In fact, I could not meet that requirement if there were fewer than 8 students in the class. Neither is it required to make the grade map 1-1. In fact, it is not possible to make the grade map 1-1 if the class has more than 8 students in it. But if the class has exactly 8 students in it, then a grade map that is onto is automatically also 1-1, and a grade map that is 1-1 is automatically also onto.

There are many maps in your life that are given as a list of pairs, such as the list of dorm-room assignments or the price list in the cafeteria. The dorm-room assignment list usually has the set of students wanting a dorm room as its domain and the set of available dorm rooms as its target, is typically not 1-1, but the authorities would like it to be onto. The price list at the cafeteria has all the items for sale as its domain, and the set $\mathbb{N}/100 := \{m/100 : m \in \mathbb{N}\}$ of all positive reals with at most two digits after the decimal point as its target. There is little sense in wondering whether this map is 1-1 or onto.

1.18 Describe an interesting map (not already discussed in class) that you have made use of in the last month or so (or, if nothing comes to mind, a map that someone like you might have used recently). Be sure to include domain and target of your map in your description and state whether or not it is 1-1, onto.

Maps and their graphs

One successful mental image of a 'map' is to imagine both domain and target as sets of some possibly indistinct shape, with curved arrows indicating with which particular element in the target the map f associates a particular element in the domain.

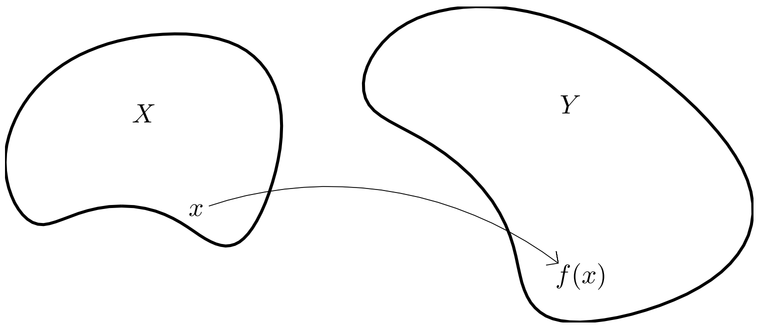

One way to visualize the map $f : X \to Y : x \mapsto f(x)$.

Another successful mental (and more successful mathematical) image of a map $f : X \to Y$ is in terms of its **graph**, i.e., in terms of the set of pairs

$$\{(x, f(x)) : x \in X\}.$$

In fact, the mathematically most satisfying definition of 'map from X to Y' is: *a subset of $X \times Y$ that, for each $x \in X$, contains exactly one pair (x, y)*. In this view, *a map is its graph.*

Here, for example, is the (graph of the) grade map G for a graduate course I taught recently. I abbreviated the students' names, to protect the innocent.

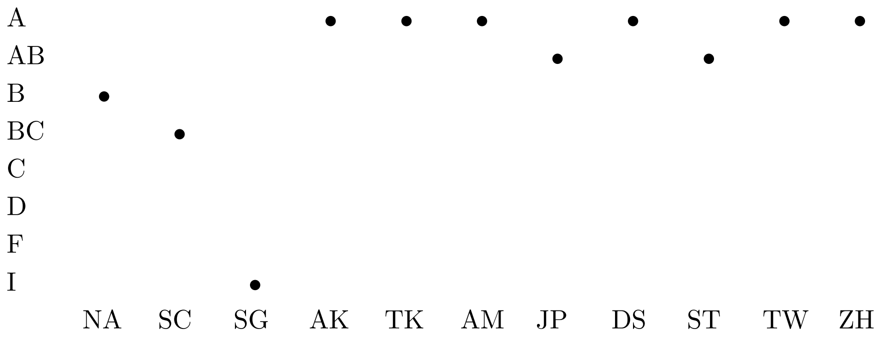

You may be more familiar with the graphs of real functions, such as the 'squaring' map

$$()^2 : [0 \mathinner{..} 2] \to [0 \mathinner{..} 4] : x \mapsto x^2,$$

whose graph is shown in the next figure.

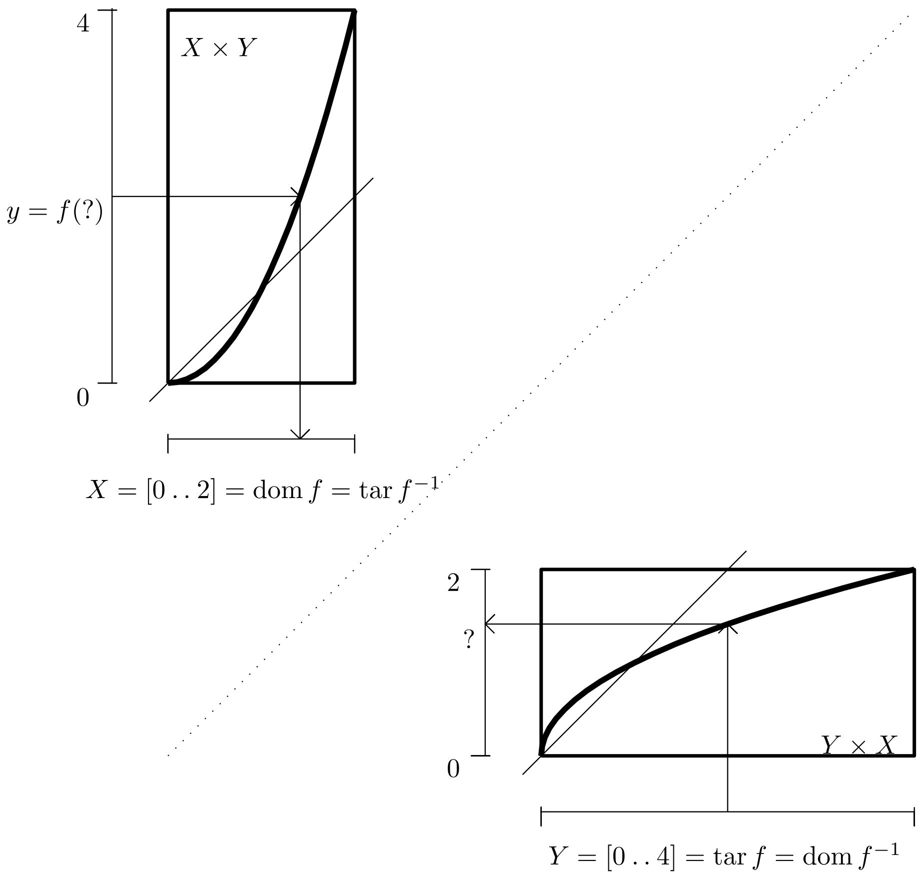

The graph of the squaring map $f := ()^2 : [0 \mathinner{..} 2] \to [0 \mathinner{..} 4] : x \mapsto x^2$ and of its inverse $f^{-1} = \sqrt{\ } : [0 \mathinner{..} 4] \to [0 \mathinner{..} 2] : x \mapsto \sqrt{x}$.

1.19 For each of the following subsets R of the cartesian product $X \times Y$ with $X = [0\,..\,2]$ and $Y = [0\,..\,4]$, determine whether it is the graph of a map from X to Y and, if it is, whether that map is 1-1 and/or onto or neither.

(a) $R = \{(x, y) : y = (x - 1/2)^2\}$; (b) $R = \{(x, y) : x \geq 1, y = (2x - 2)^2\}$; (c) $R = \{(x, y) : y = (2x - 2)^2\}$; (d) $R = \{(x, y) : x = y\}$.

1.20 Same as previous problem, but with X and Y interchanged and, correspondingly, R replaced by $R^{-1} := \{(y, x) \in Y \times X : (x, y) \in R\}$. Also, discuss any connections you see between the answers in these two problems.

Invertibility

The graph of a map f helps us solve the standard 'computational' problem involving maps, namely the problem of finding an $x \in X$ that solves the equation

$$f(?) = y$$

for given $f : X \to Y$ and $y \in Y$. The solution set is the pre-image of $\{y\}$ under f, i.e., the set

$$f^{-1}\{y\} = \{x \in X : f(x) = y\}.$$

For example, when looking at the graph of the above grade map G, we see that $G^{-1}\{\text{AB}\} = \{\text{JP, ST}\}$, while $G^{-1}\{\text{D}\} = \{\}$ (the empty set). In the first case, we have two solutions, in the second case, we have none.

In effect, when looking for solutions to the equation $f(?) = y$, we are looking at the graph of f with the roles of domain and target interchanged: We are trying to associate with each $y \in Y$ some $x \in X$ in such a way that $f(x) = y$. If f is onto, then there is *at least* one solution for every $y \in Y$, and conversely (**existence**). If f is 1-1, then there is *at most* one solution for any $y \in Y$, and conversely (**uniqueness**). Ideally, there is, for each $y \in Y$, exactly one $x \in X$ for which $f(x) = y$.

Definition: The map $f : X \to Y$ is **invertible** := for every $y \in Y$ there exists exactly one $x \in X$ for which $f(x) = y$.

Let $f : X \to Y$.

f is invertible if and only if f is 1-1 and onto.

f is invertible if and only if the **inverse of its graph**, i.e., the set

$$\{(f(x), x) : x \in X\} \subset Y \times X,$$

is the graph of a map from Y to X. This latter map is called the **inverse of** f and is denoted by f^{-1}.

Any 1-1 assignment f, taken as a map into its range, is invertible, since it is both 1-1 and onto. The above grade map G fails on both counts to be invertible, it is neither 1-1 nor onto. The squaring map $()^2 : [0\,..\,2] \to [0\,..\,4] : x \mapsto x^2$, on the other hand, is invertible since it is both 1-1 and onto. The earlier figure shows the graph of its inverse, obtained from the graph of the squaring map by reversing the roles of domain and target. In effect, we obtain the inverse of the graph of f by looking at the graph of f sideways and can often tell at a glance whether or not it is the graph of a map, i.e., whether f is 1-1 and onto.

A map may be 'half' invertible, i.e., it may be either 1-1 or onto, without being both. For example, the right shift (1.4) is 1-1, but not onto, while the left shift (1.5) is onto, but not 1-1. Only if domain and target happen to have the same *finite* number of elements, then being 1-1 is guaranteed to be the same as being onto, by the pigeonhole principle (see Problem 1.34).

(1.6) If $f : X \to Y$, with $\#X = \#Y < \infty$, then f 1-1 *or* onto implies f 1-1 *and* onto, i.e., invertible.

In particular, for any *finite* X, any map $f : X \to X$ that is 1-1 *or* onto is automatically invertible.

The notion of f being 'half' invertible is made precise by the notions of left and right inverse. Their definition requires the **identity map**, often written

$$\mathrm{id}$$

if its domain (which is also its target) is clear from the context. The full definition is:

$$\mathrm{id}_X : X \to X : x \mapsto x.$$

In other words, the identity map is a particularly boring map, it leaves everything unchanged.

We also need **map composition**:

Definition: The **composition** $f \circ g$ of two maps $f : X \to Y$ and $g : U \to W \subset X$ is the map

$$f \circ g : U \to Y : u \mapsto f(g(u)).$$

We write fg instead of $f \circ g$ whenever there is no danger of confusion.
Map composition is **associative**, i.e., whenever fg and gh are defined, then

$$(fg)h = f \circ (gh).$$

There is a corresponding definition for the composition $x \circ y$ of two assignments, x and y, under the assumption that $\operatorname{ran} y \subset \operatorname{dom} x$. Thus,

$$x_y := x \circ y = (x_{y_i} : i \in \operatorname{dom} y)$$

is an assignment whose domain is $\operatorname{dom} y$ and whose range is contained in $\operatorname{ran} x$.

As a simple *example*, if x is an n-vector and y is an m-vector with $\operatorname{ran} y \subset \underline{n} = \{1, \ldots, n\}$, then

$$z := x_y := x \circ y = (x_{y_1}, \ldots, x_{y_m}).$$

In MATLAB, if `x` describes the n-vector x and `y` describes the m-vector y with entries in $\underline{n} = \{1, \ldots, n\}$, then `z=x(y)` describes the m-vector $z = x_y = x \circ y$.

In the same way, if $\mathtt{A} \in \mathbb{F}^{m\times n}$, and `b` is a k-list with entries from $\underline{m} = \{1, \ldots, m\}$, and `c` is an l-list with entries from $\underline{n} = \{1, \ldots, n\}$, then $\mathtt{A}(\mathtt{b}, \mathtt{c})$ is a $k \times l$-matrix, namely the matrix $\mathtt{D} := \mathtt{A}(\mathtt{b}, \mathtt{c}) \in \mathbb{F}^{k\times l}$ with

$$\mathtt{D}(i, j) = \mathtt{A}(\mathtt{b}(i), \mathtt{c}(j)), \quad i \in \underline{k}, j \in \underline{l}.$$

In effect, the matrix $\mathtt{D} = \mathtt{A}(\mathtt{b}, \mathtt{c})$ is obtained from `A` by choosing rows $\mathtt{b}(1), \mathtt{b}(2), \ldots, \mathtt{b}(k)$ and columns $\mathtt{c}(1), \mathtt{c}(2), \ldots, \mathtt{c}(l)$ of `A`, in that order.

If *all* rows, in their natural order, are to be chosen, then use `A(:,c)`. Again if *all* columns, in their natural order, are to be chosen, then use `A(b,:)`.

In particular, `A(1,:)` is the matrix having the first row of `A` as its sole row, and `A(:,end)` is the matrix having the last column of `A` as its sole column. The matrix `A(1:2:end,:)` is made up from all the odd rows of `A`. `A(end:-1:1,:)` is the matrix obtained from `A` by reversing the order of the rows (as could also be obtained by the command `fliplr(A)`). `A(:,2:2:end)` is obtained by removing from `A` all odd-numbered columns. If `x` is a one-row matrix, then `x(ones(1,m),:)` and `x(ones(m,1),:)` both give the matrix having all its `m` rows equal to the single row in `x` (as would the expression `repmat(x,m,1)`).

MATLAB permits the expression `A(b,c)` to appear on the *left* of the equality sign: If `A(b,c)` and `D` are matrices of the same size, then the statement

```
A(b,c) = D;
```

changes, for each $(\mathtt{i},\mathtt{j}) \in \mathrm{dom}\mathtt{D}$, the entry `A(b(i),c(j))` of `A` to the value of `D(i,j)`. What if, e.g., `b` is not 1-1? MATLAB does the replacement for each entry of `b`, from the first to the last. Hence, the last time is the one that sticks. For example, if `a=1:4`, then the statement `a([2,2,2])=[1,2,3]` changes `a` to `[1,3,3,4]`. On the other hand, if `A` appears on both sides of such an assignment, then the one on the right is taken to be as it is at the outset of that assignment. For example,

```
A([i,j],:) = A([j,i],:);
```

is a nice way to interchange the ith row of `A` with its jth.

□

As a first use of map composition, here is the standard test for a map being onto or 1-1.

> If fg is onto, then f is onto; if fg is 1-1, then g is 1-1.

Proof: Since $\mathrm{ran}(fg) \subset \mathrm{ran}\, f \subset \mathrm{tar}\, f = \mathrm{tar}\, fg$, fg onto implies f onto. Also, if $g(y) = g(z)$, then $(fg)(y) = (fg)(z)$, hence fg 1-1 implies $y = z$, i.e., g is 1-1. □

For example, the composition lr of the left shift (1.5) with the right shift (1.4) is the identity, hence l is onto and r is 1-1 (as observed earlier).

Remark. The only practical way to check whether a given g is 1-1 is to come up with an f so that fg is 'obviously' 1-1, e.g., invertible. The only practical way to check whether a given f is onto is to come up with a g so that fg is 'obviously' onto, e.g., invertible.

> **Definition:** If $f \in Y^X$ and $g \in X^Y$ and $fg = \mathrm{id}$, then f (being to the left of g) is a **left inverse** of g, and g is a **right inverse** of f. In particular, any left inverse is onto and any right inverse is 1-1.

To help you remember which of f and g is onto and which is 1-1 in case $fg = \mathrm{id}$, keep in mind that being onto provides conclusions about elements of the target of the map while being 1-1 provides conclusions about elements in the domain of the map.

Now we consider the converse statements.

> If $f : X \to Y$ is 1-1, then f has a left inverse.

Proof: If f is 1-1 and $x \in X$ is some element, then

$$g : Y \to X : y \mapsto \begin{cases} f^{-1}\{y\} & \text{if } y \in \mathrm{ran}\, f; \\ x & \text{otherwise,} \end{cases}$$

is well-defined since each $y \in \mathrm{ran}\, f$ is the image of exactly one element of X. With g so defined, $gf = \mathrm{id}$ follows. □

The corresponding statement: *If $f : X \to Y$ is onto, then f has a right inverse* would have the following 'proof': Since f is onto, we can define $g : Y \to X : y \mapsto$ some point in $f^{-1}\{y\}$. Regardless of how we pick that point $g(y) \in f^{-1}\{y\}$, the resulting map is a right inverse for f. – Some object to this argument since it requires us to pick, for each y, a particular element from that set $f^{-1}\{y\}$. The **belief** that this can *always* be done is known as "The Axiom of Choice".

If f is an invertible map, then f^{-1} is both a right inverse and a left inverse for f. Conversely, if g is a right inverse for f and h is a left inverse for f, then f is invertible and $h = f^{-1} = g$.
Consequently, if f is invertible, then: (i) f^{-1} is also invertible, and $(f^{-1})^{-1} = f$; and, (ii) if also g is an invertible map, with $\operatorname{tar} g = \operatorname{dom} f$, then fg is invertible, and $(fg)^{-1} = g^{-1}f^{-1}$ (note the order reversal).

Proof: Let $f : X \to Y$ be invertible. Since, for every $y \in Y$, $f^{-1}(y)$ solves the equation $f(?) = y$, we have $ff^{-1} = \mathrm{id}_Y$, while, for any $x \in X$, x is a solution of the equation $f(?) = f(x)$, hence necessarily $x = f^{-1}(f(x))$, thus also $f^{-1}f = \mathrm{id}_X$.

As to the converse, if f and has both a left and a right inverse, then it must be both 1-1 and onto, hence invertible. Further, if $hf = \mathrm{id}_X$ and $fg = \mathrm{id}_Y$, then (using the associativity of map composition),

$$h = h\,\mathrm{id}_Y = h \circ (fg) = (hf)g = \mathrm{id}_X g = g,$$

showing that $h = g$, hence $h = f^{-1} = g$.

As to the consequences, the identities $ff^{-1} = \mathrm{id}_Y$ and $f^{-1}f = \mathrm{id}_X$ explicitly identify f as a right and left inverse for f^{-1}, hence f must be the inverse of f^{-1}. Also, by map associativity, $(fg)g^{-1}f^{-1} = f\,\mathrm{id}_X f^{-1} = ff^{-1} = \mathrm{id}_Y$, etc. □

While $fg = \mathrm{id}$ implies $gf = \mathrm{id}$ in general only in case $\#\operatorname{dom} f = \#\operatorname{tar} f < \infty$, it does imply that gf is as much of an identity map as it can be: Indeed, if $fg = \mathrm{id}$, then $(gf)g = g \circ (fg) = g\,\mathrm{id} = g$, showing that $(gf)x = x$ for every $x \in \operatorname{ran} g$. There is no such hope for $x \notin \operatorname{ran} g$, since such x cannot possibly be in $\operatorname{ran} gf = g(\operatorname{ran} f) \subset \operatorname{ran} g$. However, since $gf(x) = x$ for all $x \in \operatorname{ran} g$, we conclude that $\operatorname{ran} gf = \operatorname{ran} g$. This makes gf the identity on its range, $\operatorname{ran} g$. In particular, $(gf) \circ (gf) = gf$, i.e., gf is **idempotent** or, a **projector**.

(1.7) Proposition: If $f : X \to Y$ and $fg = \mathrm{id}_Y$, then gf is a projector, i.e., the identity on its range, and that range equals $\operatorname{ran} g$.

For example, the composition lr of the left shift (1.5) with the right shift (1.4) is the identity, hence rl must be the identity on $\operatorname{ran} r = \{2, 3, \ldots\}$ and, indeed, it is.

If the n-vector `c` in `MATLAB` describes a permutation, i.e., if the map $c : \underline{n} \to \underline{n} : j \mapsto$ `c(`j`)` is 1-1 or onto, hence invertible, then the n-vector `cinv` giving its inverse can be obtained with the commands

```
cinv = c; cinv(c) = 1:length(c);
```

The first command makes sure that `cinv` starts out as a vector of the same size as `c`. With that, the second command changes `cinv` into one for which `cinv(c) = [1,2,...,length(c)]`. In other words, `cinv` describes a *left* inverse for (the map given by) `c`, hence the inverse (by the pigeonhole principle).

A second, more expensive, way to construct `cinv` is with the help of the command `sort`, as follows:

```
[d, cinv] = sort(c);
```

For (try `help sort`), whether or not `c` describes a permutation, this command produces, in the n-vector `d`, the list of the items in `c` in nondecreasing order, and provides, in `cinv`, the recipe for this re-ordering:

`d(i)=c(cinv(i))`, `i`$=1{:}n$.

In particular, if `c` describes a permutation, then, necessarily, `d = [1,2,3,...]`, therefore `c(cinv) = [1,2,...,length(c)]`, showing that `cinv` describes a *right* inverse for (the map given by) `c`, hence the inverse (by the pigeonhole principle).

Both of these methods extend, to the construction of a left, respectively a right, inverse, in case the map given by `c` has only a left, respectively a right, inverse.

□

1.21 Let $f : \underline{2} \to \underline{3}$ be given by the list $(2, 3)$, and let $g : \underline{3} \to \underline{2}$ be the map given by the list $(2, 1, 2)$.

(a) Describe fg and gf (e.g., by giving their lists).

(b) Verify that fg is a projector, i.e., is the identity on its range.

1.22 For each of the following maps, state whether or not it is 1-1, onto, invertible. Also, describe a right inverse or a left inverse or an inverse for it or else state why such right inverse or left inverse or inverse does not exist.

The maps are specified in various ways, e.g., by giving their list and their target or by giving both domain and target and a rule for constructing their values.

(a) a is the map to $\{1, 2, 3\}$ given by the list $(1, 2, 3)$.

(b) b is the map to $\{1, 2, 3, 4\}$ given by the list $(1, 2, 3)$.

(c) c is the map to $\{1, 2\}$ given by the list $(1, 2, 1)$.

(d) $d : \mathbb{R}^2 \to \mathbb{R} : x \mapsto 2x_1 - 3x_2$.

(e) $f : \mathbb{R}^2 \to \mathbb{R}^2 : x \mapsto (-x_2, x_1)$.

(f) $g : \mathbb{R}^2 \to \mathbb{R}^2 : x \mapsto (x_1 + 2, x_2 - 3)$.

(g) $h : \mathbb{R} \to \mathbb{R}^2 : y \mapsto (y/2, 0)$.

1.23 Verify that, in the preceding problem, $dh = \mathrm{id}$, and explain geometrically why one would call hd a projector.

1.24 Prove: *If $fg = fh$ for $g, h : S \to T$ and with $f : T \to U$ 1-1, then $g = h$.*

1.25 Prove: *If $fh = gh$ for $f, g : T \to U$ and with $h : S \to T$ onto, then $f = g$.*

1.26 Use the preceding two problems to prove the following converse of (1.7)Proposition: *If $f : X \to Y$ and gf is a projector, then f is onto and g is 1-1 iff $fg = \mathrm{id}_Y$.*

1.27 If both f and g are maps from $\underline{n}$ to $\underline{n}$, then so are both fg and gf. In particular, for any $f \in \underline{n}^{\underline{n}}$, its power sequence

$$f^0 := \mathrm{id}_{\underline{n}}, f^1 := f, f^2 := f \circ f, f^3 := f \circ f^2, \ldots$$

is well defined. Further, since $\underline{n}^{\underline{n}}$ is finite, the sequence $f^0, f^1, f^2, \ldots$ of powers must eventually repeat itself. In other words, there must be a first r such that $f^r = f^j$ for some $j < r$. Let's call the difference $d := r - j$ between these two exponents the **cycle length** of f.

(a) Find the cycle length for the map given by the sequence $(2, 3, 4, 1, 1)$. (Feel free to use MATLAB.)

(b) Also determine the cycle lengths for the following maps:

A:=(2,3,4,5,1);	B:=(2,3,1,5,4);	C:=(1,2,3,4,5);
D:=(2,5,2,2,1);	E:=(2,5,2,5,2);	F:=(2,5,2,2,5).

(c) Given all these examples (and any others you care to try), what is your *guess* as to the special nature of the map f^d in case the cycle length of f is d and f is invertible?

1.28 Finish appropriately the following MATLAB function

```
function b = ii(a)
% If  ran(a) = N := {1,2,...,length(a)} , hence  a  describes
% the invertible map
%            f:N --> N : j -> a(j)
% then  b  describes the inverse of  f , i.e., the map  g:N --> N  for which
%  fg = id_N  and  gf = id_N .
% Otherwise, the message
%  The input doesn't describe an invertible map
% is printed and an empty  b  is returned.
```


1.29 Let $f_i : X \to X$ for $i = 1{:}n$, hence $g := f_1 \cdots f_n$ is also a map from X to X. Prove that *g is invertible if, but not only if, each f_i is invertible, and, in that case, $g^{-1} = f_n^{-1} \cdots f_1^{-1}$.* (Note the order reversal!)

1.30 If $f : S \to T$ is invertible, then f has exactly one left inverse. Is the converse true?

1.31 Let g be a left inverse for $f : S \to T$, and assume that $\#S > 1$. Prove that g is the unique left inverse for f iff g is 1-1. (Is the assumption that $\#S > 1$ really needed?)

1.32 Let g be a right inverse for f. Prove that g is the unique right inverse for f iff g is onto.

1.33 If $f : S \to T$ is invertible, then f has exactly one right inverse. Is the converse true?

1.34

(i) Prove: *If $g : Z \to X$ is invertible, then, for any $f : X \to Y$, f is 1-1 (onto) if and only if the map fg is 1-1 (onto).*

(ii) Derive (1.6) from (1.3).

The inversion of maps

The notions of 1-1 and onto, and the corresponding notions of right and left inverse, are basic to the discussion of the standard 'computational' problem already mentioned earlier: for $f : X \to Y$ and $y \in Y$, solve

$$f(?) = y. \tag{1.1}$$

When we try to solve (1.1), we are really trying to find, for each $y \in Y$, some $x \in X$ for which $f(x) = y$, i.e., we are trying to come up with a right inverse for f. *Existence* of a solution for every right side is the same as having f onto, and is ensured by the existence of a right inverse for f. Existence of a left inverse for f ensures *uniqueness*: If $hf = \mathrm{id}$, then $f(x) = f(y)$ implies that $x = h(f(x)) = h(f(y)) = y$. Thus existence of a left inverse implies that f is *1-1*. But existence of a left inverse does *not*, in general, provide a solution.

When f has its domain in $\mathbb{R}^n$ and and its target in $\mathbb{R}^m$, then we can think of solving (1.1) *numerically.* Under the best of circumstances, this still means that we must proceed by *approximation.* The solution is found as the limit of a sequence of solutions to *linear* equations, i.e., equations of the form $A? = b$, with A a *linear* map. This is so because linear (algebraic) equations are the only kind of equations we can actually solve exactly (ignoring roundoff). This is one reason why Linear Algebra is so important. It provides the mathematical structures, namely vector spaces and linear maps, needed to deal efficiently with linear equations and, thereby, with other equations.

1.35 T/F

(a) 0 is a natural number.

(b) $\#\{3, 3, 3\} = 1$.

(c) $\#(3, 3, 3) = 3$.

(d) $(\{3, 1, 3, 2, 4\} \cap \{3, 5, 4\}) \cup \{3, 3\} = \{4, 3, 3, 3, 3\}$.

(e) If A, B are finite sets, then $\#(A \cup B) = \#A + \#B - \#(A \cap B)$.

(f) $\#\{\} = 1$.

(g) $\{3, 3, 1, 6\}\backslash\{3, 1\} = \{3, 6\}$.

(h) If $f : X \to X$ for some finite X, then f is 1-1 if and only if f is onto.

(i) The map $f : \underline{3} \to \underline{3}$ given by the list $(3, 1, 2)$ is invertible, and its inverse is given by the list $(2, 3, 1)$.

(j) The map $f : \underline{3} \to \underline{2}$ given by the list $(1, 2, 1)$ has a right inverse.

(k) If $U \subset \operatorname{tar} f$, then f maps $f^{-1}U$ onto U.

(l) The map f is invertible if and only if f^{-1} is the graph of a map.

(m) If $f, g \in X^X$ and $h := fg$ is invertible, then both f and g are invertible.

(n) The matrix $\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$ is diagonal.

(o) The matrix $\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ is upper triangular.

2. Vector spaces and linear maps

Vector spaces, especially spaces of functions

Linear algebra is concerned with vector spaces. These are sets on which two operations, *vector addition* and *multiplication by a scalar*, are defined in such a way that they satisfy various laws. Here they are, for the record:

(2.1) Definition: To say that X is a **linear space** (of **vector**s), or a **vector space**, over the commutative field $\mathbb{F}$ (of **scalar**s) means that there are two maps, (i) $X \times X \to X : (x, y) \mapsto x + y$ called **(vector) addition**; and (ii) $\mathbb{F} \times X \to X : (\alpha, x) \mapsto \alpha x =: x\alpha$ called **scalar multiplication**, which satisfy the following rules.

(a) X is a **commutative group with respect to addition**; i.e., addition
- (a.1) is **associative**: $x + (y + z) = (x + y) + z$;
- (a.2) is **commutative**: $x + y = y + x$;
- (a.3) has **neutral** element: $\exists 0\ \forall x\ x + 0 = x$;
- (a.4) has **inverse**: $\forall x\ \exists y\ x + y = 0$.

(s) scalar multiplication is
- (s.1) **associative**: $\alpha(\beta x) = (\alpha\beta)x$;
- (s.2) **field-addition distributive**: $(\alpha + \beta)x = \alpha x + \beta x$;
- (s.3) **vector-addition distributive**: $\alpha(x + y) = \alpha x + \alpha y$;
- (s.4) **unitary**: $1x = x$.

It is standard to denote the element $y \in X$ for which $x + y = 0$ by $-x$ since such y is uniquely determined by the requirement that $x + y = 0$. I will denote the neutral element in X by the same symbol, 0, used for the zero scalar. For reasons to become clear, I often write $x\alpha$ for αx.

While the **scalar**s can come from some abstract field, we will only be interested in the real scalars $\mathbb{R}$ and the complex scalars $\mathbb{C}$. Also, from a practical point of view, the most important linear spaces consist of **function**s, i.e., of scalar-valued maps all on some common domain. This means that the typical linear space we will deal with is (a subset of) the collection of all maps $\mathbb{F}^T$ from some fixed domain T into the **scalar field** $\mathbb{F}$ (either $\mathbb{F} = \mathbb{R}$ or $\mathbb{F} = \mathbb{C}$), with **pointwise** addition and multiplication by scalars. Here is the definition:

(2.2) Definition of pointwise vector operations:

(a) The **sum** $f + g$ of $f, g \in \mathbb{F}^T$ is the function

$$f + g : T \to \mathbb{F} : t \mapsto f(t) + g(t).$$

(s) The **product** αf of the **scalar** $\alpha \in \mathbb{F}$ with the function $f \in \mathbb{F}^T$ is the function

$$\alpha f : T \to \mathbb{F} : t \mapsto \alpha f(t).$$

With respect to these operations, $\mathbb{F}^T$ is a linear space (over $\mathbb{F}$). In particular, the function

$$0 : T \to \mathbb{F} : t \mapsto 0$$

is the neutral element, or zero vector, and, for $f \in \mathbb{F}^T$,

$$-f : T \to \mathbb{F} : t \mapsto -f(t)$$

is the additive inverse for f.

Note that it is not possible to add two functions unless they have the same domain!

Standard examples include:

(i) $T = \underline{n}$, in which case we get n-**dimensional coordinate space** $\mathbb{F}^n$ whose elements (vectors) we call n-vectors.

(ii) $T = \underline{m} \times \underline{n}$, in which case we get the space $\mathbb{F}^{m \times n}$, whose elements we call m-by-n matrices.

(iii) $T = \mathbb{R}$, $\mathbb{F} = \mathbb{R}$, in which case we get the space of all real-valued functions on the real line.

(iv) $T = \mathbb{R}^n$, $\mathbb{F} = \mathbb{R}$, in which case we get the space of all real-valued functions of n real variables.

The most common way to get a vector space is as a *linear subspace*:

Definition: A nonempty subset Y of a vector space X is a **linear subspace** (of X) in case it is closed under addition and multiplication by a scalar. This means that the two sets

$$Y + Y := \{y + z : y, z \in Y\} \text{ and } \mathbb{F}Y := \{\alpha y : \alpha \in \mathbb{F}, y \in Y\}$$

are in Y.

Standard examples include:

(i) The **trivial space** $\{0\}$, consisting of the zero vector alone; it's a great space for testing one's understanding.

(ii) $\Pi_k :=$ the set of all **polynomials of degree** $\le k$ as a subset of $\mathbb{F}^{\mathbb{F}}$.

(iii) The set $C([a\,..\,b])$ of all **continuous functions** on the interval $[a\,..\,b]$.

(iv) The set of all real symmetric matrices of order n as a subset of $\mathbb{R}^{n \times n}$.

(v) The set of all real-valued functions on $\mathbb{R}$ that vanish on some fixed set S.

(vi) The set $\mathrm{BL}_\xi \subset C([\xi_1\,..\,\xi_{\ell+1}])$ of all **broken lines** with (interior) breaks at $\xi_2 < \cdots < \xi_\ell$.

It is a good exercise to check that, according to the abstract definition of a vector space, any linear subspace of a vector space is again a vector space. Conversely, if a subset of a vector space is *not* closed under vector addition or under multiplication by scalars, then it cannot be a vector space (with respect to the given operations) since it violates the basic assumption that the sum of any two elements and the product of any scalar with any element is again an element of the space. (To be sure, the empty subset $\{\}$ of a linear space is vacuously closed under the two vector operations but fails to be a linear subspace since it fails to be nonempty.)

Proposition: A subset Y of a vector space X is a vector space (with respect to the same addition and multiplication by scalars) if and only if Y is a linear subspace (of X), i.e., Y is nonempty and is closed under addition and multiplication by scalars.

Corollary: The sum, $Y + Z := \{y + z : y \in Y, z \in Z\}$, and the intersection, $Y \cap Z$, of two linear subspaces, Y and Z, of a vector space is a linear subspace.

We saw that pointwise addition and multiplication by a scalar makes the collection $\mathbb{F}^T$ of all maps from some set T to the scalars a vector space. The same argument shows that the collection X^T of all maps from some set T into a *vector space* X (over the scalar field $\mathbb{F}$) is a vector space under pointwise addition and multiplication by scalars. This means, explicitly, that we define the sum $f + g$ of $f, g \in X^T$ by

$$f + g : T \to X : t \mapsto f(t) + g(t)$$

and define the product αf of $f \in X^T$ with the scalar $\alpha \in \mathbb{F}$ by

$$\alpha f : T \to X : t \mapsto \alpha f(t).$$

Thus, we can generate from one vector space X many different vector spaces, namely all the linear subspaces of the vector space X^T, with T an arbitrary set.

2.1 For each of the following sets of real-valued assignments or maps, determine whether or not they form a vector space (with respect to pointwise addition and multiplication by scalars), and give a reason for your answer. (a) $\{x \in \mathbb{R}^3 : x_1 = 4\}$; (b) $\{x \in \mathbb{R}^3 : x_1 = x_2\}$; (c) $\{x \in \mathbb{R}^3 : 0 \le x_j, j = 1, 2, 3\}$; (d) $\{(0, 0, 0)\}$; (e) $\{x \in \mathbb{R}^3 : x \notin \mathbb{R}^3\}$; (f) $C([0\,..\,2])$; (g) The collection of all 3×3 matrices with all diagonal entries equal to zero; (h) $\{(x, 0) : x \in \mathbb{R}\} \cup \{(0, y) : y \in \mathbb{R}\}$.

2.2 Prove that, *for every x in the vector space X, $(-1)x = -x$, and $0x = 0$.*

2.3 Prove that *the intersection of any collection of linear subspaces of a vector space is a linear subspace.*

2.4 Prove: *The union of two linear subspaces is a linear subspace if and only if one of them contains the other.*

2.5 Prove: *The finite union of linear subspaces is a linear subspace if and only if one of them contains all the others.* (Hint: reduce to the situation that no subspace is contained in the union of the other subspaces and, assuming this leaves you with at least two subspaces, take from each a point that is in none of the others and consider the straight line through these two points.)

2.6 Provide a proof of the above Proposition.

Linear maps

Definition: Let X, Y be vector spaces (over the same scalar field $\mathbb{F}$). The map $f : X \to Y$ is called **linear** if it is

(a) **additive**, i.e.,

$$\forall\{x, z \in X\}\ f(x + z) = f(x) + f(z);$$

and

(s) **homogeneous**, i.e.,

$$\forall\{x \in X, \alpha \in \mathbb{F}\}\ f(\alpha x) = \alpha f(x).$$

We denote the collection of all linear maps from X to Y by

$$L(X, Y).$$

Many books call a linear map a **linear transformation** or a **linear operator**. It is customary to denote linear maps by capital letters. Further, if A is a linear map and $x \in \operatorname{dom} A$, then it is customary to write Ax instead of $A(x)$.

Examples: If X is a linear subspace of $\mathbb{F}^T$, then, for every $t \in T$, the map

$$\delta_t : X \to \mathbb{F} : f \mapsto f(t)$$

of evaluation at t is linear since the vector operations are pointwise.

The map $D : C^{(1)}(\mathbb{R}) \to C(\mathbb{R}) : g \mapsto Dg$ that associates with each continuously differentiable function g its first derivative Dg is a linear map.

The map $C([a\,..\,b]) \to \mathbb{R} : g \mapsto \int_a^b g(t)\,\mathrm{d}t$ is linear.

Let $\mathbf{c} := \{a : \mathbb{N} \to \mathbb{F} : \lim_{n\to\infty} a_n \text{ exists}\}$, i.e., $\mathbf{c}$ is the vector space of all convergent sequences. Then the map $\mathbf{c} \to \mathbb{F} : a \mapsto \lim_{n\to\infty} a_n$ is linear.

These examples show that the basic operations in Calculus are linear. This is the reason why so many people outside Algebra, such as Analysts and Applied Mathematicians, are so interested in Linear Algebra.

The simplest linear map on a vector space X to a vector space Y is the so-called **trivial map**. It is the linear map that maps every element of X to 0; it is, itself, denoted by

$$0.$$

It is surprising how often this map serves as a suitable illustration or counterexample.

Example: If $a \in \mathbb{R}^n$, then

$$a^{\mathrm{t}} : \mathbb{R}^n \to \mathbb{R} : x \mapsto a^{\mathrm{t}}x := a_1x_1 + a_2x_2 + \cdots + a_nx_n \tag{2.3}$$

is a linear map of great practical importance. Indeed, any (real) linear algebraic equation in n unknowns has the form

$$a^{\mathrm{t}}? = y$$

for some **coefficient vector** $a \in \mathbb{R}^n$ and some **right side** $y \in \mathbb{R}$. Such an equation has solutions for arbitrary y if and only if $a \neq 0$. You have already learned that the general solution can always be written as the sum of a particular solution and an arbitrary solution of the corresponding **homogeneous** equation

$$a^{\mathrm{t}}? = 0.$$

In particular, the map a^{t} cannot be 1-1 unless $n = 1$.

Assume that $a \neq 0$. For $n = 2$, it is instructive and easy to visualize the solution set as a straight line, parallel to the straight line

$$\operatorname{null} a^{\mathrm{t}} := \{x \in \mathbb{R}^2 : a^{\mathrm{t}}x = 0\}$$

through the origin formed by all the solutions to the corresponding homogeneous problem, and perpendicular to the coefficient vector a. Note that the 'nullspace' $\operatorname{null} a^{\mathrm{t}}$ splits $\mathbb{R}^2$ into the two **half-spaces**

$$\{x \in \mathbb{R}^2 : a^{\mathrm{t}}x > 0\} \qquad \{x \in \mathbb{R}^2 : a^{\mathrm{t}}x < 0\},$$

one of which contains a. Here is such a figure, for the particular equation

$$2x_1 + 3x_2 = 6.$$

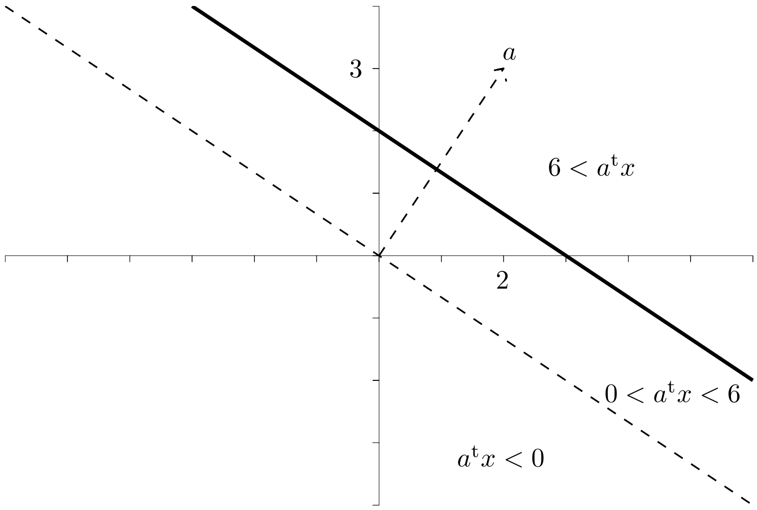

(2.4) Figure. One way to visualize all the parts of the equation $a^t x = 6$ with $a = (2, 3)$.

□

By adding or composing two linear maps (if appropriate) or by multiplying a linear map by a scalar, we obtain further linear maps. Here are the details.

The (pointwise) sum $A+B$ of $A, B \in L(X,Y)$ and the product αA of $\alpha \in \mathbb{F}$ with $A \in L(X,Y)$ are again in $L(X,Y)$, hence $L(X,Y)$ is closed under (pointwise) addition and multiplication by a scalar, therefore a linear subspace of the vector space Y^X of all maps from X into the vector space Y.

$L(X,Y)$ is a vector space under pointwise addition and multiplication by a scalar.

Linearity is preserved not only under (pointwise) addition and multiplication by a scalar, but also under map *composition*.

The composition of two linear maps is again linear (if it is defined).

Indeed, if $A \in L(X,Y)$ and $B \in L(Y,Z)$, then BA maps X to Z and, for any $x, y \in X$,

$$(BA)(x+y) = B(A(x+y)) = B(Ax + Ay) = B(Ax) + B(Ay) = (BA)(x) + (BA)(y).$$

Also, for any $x \in X$ and any scalar α,

$$(BA)(\alpha x) = B(A(\alpha x)) = B(\alpha Ax) = \alpha B(Ax) = \alpha (BA)(x).$$

2.7 For each of the following maps, determine whether or not it is linear (give a reason for your answer).

(a) $\Pi_{<k} \to \mathbb{Z}_+ : p \mapsto \#\{x : p(x) = 0\}$ (i.e., the map that associates with each polynomial of degree $< k$ the number of its zeros).

(b) $C([a\,..\,b]) \to \mathbb{R} : f \mapsto \max_{a \le x \le b} f(x)$

(c) $\mathbb{F}^{3\times 4} \to \mathbb{F} : A \mapsto A_{2,2}$

(d) $L(X,Y) \to Y : A \mapsto Ax$, with x a fixed element of X (and, of course, X and Y vector spaces).

(e) $\mathbb{R}^{m\times n} \to \mathbb{R}^{n\times m} : A \mapsto A^c$ (with A^c the (conjugate) transpose of the matrix A)

(f) $\mathbb{R} \to \mathbb{R}^2 : x \mapsto (x, \sin(x))$

2.8 *The linear image of a vector space is a vector space*: Let $f : X \to T$ be a map on some vector space X into some set T on which addition and multiplication by scalars is defined in such a way that

$$f(\alpha x + \beta y) = \alpha f(x) + \beta f(y), \qquad \alpha, \beta \in \mathbb{F},\ x, y \in X. \tag{2.5}$$

Prove that $\operatorname{ran} f$ is a vector space (with respect to the addition and multiplication as restricted to $\operatorname{ran} f$). (See Problem 4.27 for an important application.)

Linear maps from $\mathbb{F}^n$

As a ready source of many examples, we now give a complete description of $L(\mathbb{F}^n, X)$.

For any sequence $v_1, v_2, \ldots, v_n$ in the vector space X, the map

$$f : \mathbb{F}^n \to X : a \mapsto v_1 a_1 + v_2 a_2 + \cdots + v_n a_n$$

is linear.

Proof: The proof is a boring but necessary verification.

(a) additivity:

$$\begin{aligned}
f(a+b) &= v_1\,(a+b)_1 + v_2\,(a+b)_2 + \cdots + v_n\,(a+b)_n && \text{(definition of } f\text{)}\\
&= v_1\,(a_1+b_1) + v_2\,(a_2+b_2) + \cdots + v_n\,(a_n+b_n) && \text{(addition of } n\text{-vectors)}\\
&= v_1a_1 + v_1b_1 \ +\ v_2a_2 + v_2b_2 \ +\ \cdots \ +\ v_na_n + v_nb_n && \text{(multipl. by scalar distributes)}\\
&= v_1a_1 + v_2a_2 + \cdots + v_na_n \quad + \quad v_1b_1 + v_2b_2 + \cdots + v_nb_n && \text{(vector addition commutes)}\\
&= f(a) \ +\ f(b) && \text{(definition of } f\text{)}
\end{aligned}$$

(s) homogeneity:

$$\begin{aligned}
f(\alpha a) &= v_1\,(\alpha a)_1 + v_2\,(\alpha a)_2 + \cdots + v_n\,(\alpha a)_n && \text{(definition of } f\text{)}\\
&= v_1\alpha a_1 + v_2\alpha a_2 + \cdots + v_n\alpha a_n && \text{(multipl. of scalar with } n\text{-vectors)}\\
&= \alpha(v_1a_1 + v_2a_2 + \cdots + v_na_n) && \text{(multipl. by scalar distributes)}\\
&= \alpha f(a) && \text{(definition of } f\text{)}
\end{aligned}$$

□

Definition: The weighted sum

$$v_1a_1 + v_2a_2 + \cdots + v_na_n$$

is called the **linear combination of the vectors** $v_1, v_2, \ldots, v_n$ **with weights** $a_1, a_2, \ldots, a_n$. I will use the suggestive abbreviation

$$[v_1, v_2, \ldots, v_n]a := v_1a_1 + v_2a_2 + \cdots + v_na_n,$$

hence use

$$[v_1, v_2, \ldots, v_n]$$

for the map $V : \mathbb{F}^n \to X : a \mapsto v_1a_1 + v_2a_2 + \cdots + v_na_n$. I call such a map a **column map**, and call v_j its j**th column**. Further, I denote its number of columns by

$$\#V.$$

The most important special case of this occurs when also X is a coordinate space, $X = \mathbb{F}^m$ say. In this case, each v_j is an m-vector, and

$$v_1 a_1 + v_2 a_2 + \cdots + v_n a_n = Va,$$

with V the $m \times n$-matrix with columns $v_1, v_2, \ldots, v_n$. This explains why I chose to write the weights in the linear combination $v_1 a_1 + v_2 a_2 + \cdots + v_n a_n$ to the right of the vectors v_j rather than to the left. For, it suggests that working with the map $[v_1, v_2, \ldots, v_n]$ is rather like working with a matrix with columns $v_1, v_2, \ldots, v_n$.

> Note that `MATLAB` uses the notation $[v_1, v_2, \ldots, v_n]$ for the matrix with columns $v_1, v_2, \ldots, v_n$, as do some textbooks. This stresses the fact that it is customary to think of the *matrix* $C \in \mathbb{F}^{m\times n}$ with columns $c_1, c_2, \ldots, c_n$ as the *linear map* $[c_1, c_2, \ldots, c_n] : \mathbb{F}^n \to \mathbb{F}^m : x \mapsto c_1 x_1 + c_2 x_2 + \cdots + c_n x_n$.

□

Agreement: For any sequence $v_1, v_2, \ldots, v_n$ of m-vectors,

$$[v_1, v_2, \ldots, v_n]$$

denotes both the $m \times n$-matrix V with columns $v_1, v_2, \ldots, v_n$ and the linear map

$$V : \mathbb{F}^n \to \mathbb{F}^m : a \mapsto [v_1, v_2, \ldots, v_n]a = v_1 a_1 + v_2 a_2 + \cdots + v_n a_n.$$

Thus,

$$\mathbb{F}^{m\times n} = L(\mathbb{F}^n, \mathbb{F}^m).$$

Thus, a matrix $V \in \mathbb{F}^{m\times n}$ is associated with two rather different maps: (i) since it is an assignment with domain $\underline{m} \times \underline{n}$ and values in $\mathbb{F}$, we could think of it as a map on $\underline{m} \times \underline{n}$ to $\mathbb{F}$; (ii) since it is the n-list of its columns, we can think of it as the linear map from $\mathbb{F}^n$ to $\mathbb{F}^m$ that carries the n-vector a to the m-vector $Va = v_1 a_1 + v_2 a_2 + \cdots + v_n a_n$. From now on, we will stick to the second interpretation when we talk about the domain, the range, or the target, of a matrix. Thus, for $V \in \mathbb{F}^{m\times n}$, $\operatorname{dom} V = \mathbb{F}^n$ and $\operatorname{tar} V = \mathbb{F}^m$, and $\operatorname{ran} V \subset \mathbb{F}^m$. – If we want the first interpretation, we call $V \in \mathbb{F}^{m\times n}$ a (two-dimensional) **array**.

Next, we prove that there is nothing special about the linear maps of the form $[v_1, v_2, \ldots, v_n]$ from $\mathbb{F}^n$ into the vector space X, i.e., *every* linear map from $\mathbb{F}^n$ to X is necessarily of that form. The identity map

$$\mathrm{id}_n : \mathbb{F}^n \to \mathbb{F}^n : a \to a$$

is of this form, i.e.,

$$\mathrm{id}_n = [e_1, e_2, \ldots, e_n]$$

with e_j the j**th unit vector**, i.e.,

$$e_j := (\underbrace{0, \ldots, 0}_{j-1 \text{ zeros}}, 1, 0, \ldots, 0)$$

the vector (with the appropriate number of entries) all of whose entries are 0, except for the jth, which is 1. Written out in painful detail, this says that

$$a = e_1 a_1 + e_2 a_2 + \cdots + e_n a_n, \qquad \forall a \in \mathbb{F}^n.$$

Further,

(2.6) Proposition: If $V = [v_1, v_2, \ldots, v_n] : \mathbb{F}^n \to X$ and $f \in L(X, Y)$, then $fV = [f(v_1), \ldots, f(v_n)]$.

Proof: If $\operatorname{dom} f = X$ and f is linear, then fV is linear and, for any $a \in \mathbb{F}^n$,

$$(fV)a = f(Va) = f(v_1a_1 + v_2a_2 + \cdots + v_na_n) = f(v_1)a_1 + f(v_2)a_2 + \cdots + f(v_n)a_n = [f(v_1), \ldots, f(v_n)]a.$$

□

Consequently, for any $f \in L(\mathbb{F}^n, X)$,

$$f = f\ \mathrm{id}_n = f\,[e_1, e_2, \ldots, e_n] = [f(e_1), \ldots, f(e_n)].$$

This proves:

(2.7) Proposition: The map f from $\mathbb{F}^n$ to the vector space X is linear if and only if

$$f = [f(e_1), f(e_2), \ldots, f(e_n)].$$

In other words,

$$L(\mathbb{F}^n, X) = \{[v_1, v_2, \ldots, v_n] : v_1, v_2, \ldots, v_n \in X\} \qquad (\simeq X^n).$$

As a simple example, recall from (2.3) the map $a^{\mathrm{t}} : \mathbb{R}^n \to \mathbb{R} : x \mapsto a_1x_1 + a_2x_2 + \cdots + a_nx_n = [a_1, \ldots, a_n]x$, and, in this case, $a^{\mathrm{t}}e_j = a_j$, all j. This confirms that a^{t} is linear and shows that

(2.8) $$a^{\mathrm{t}} = [a_1, \ldots, a_n] = [a]^{\mathrm{t}}.$$

Notation: I follow `MATLAB` notation. E.g., $[V, W]$ denotes the column map in which first all the columns of V are used and then all the columns of W. Also, if V and W are column maps, then I write

$$V \subset W$$

to mean that V is obtained by omitting (zero or more) columns from W; i.e., $V = W(:, c)$ for some subsequence c of $1{:}\#W$.

Finally, if W is a column map and M is a set, then I'll write

$$W \ \subset\ M$$

to mean that the columns of W are elements of M. For example:

(2.9) Proposition: If Z is a linear subspace of Y and $W \in L(\mathbb{F}^m, Y)$, then $W \subset Z \implies \operatorname{ran} W \subset Z$.

The important (2.6)Proposition is the reason we define the **product of matrices** the way we do, namely as

$$(AB)(i,j) := \sum_k A(i,k)B(k,j), \quad \forall i,j.$$

For, if $A \in \mathbb{F}^{m\times n} = L(\mathbb{F}^n, \mathbb{F}^m)$ and $B = [b_1, b_2, \ldots, b_r] \in \mathbb{F}^{n\times r} = L(\mathbb{F}^r, \mathbb{F}^n)$, then $AB \in L(\mathbb{F}^r, \mathbb{F}^m) = \mathbb{F}^{m\times r}$, and

$$AB = A[b_1, b_2, \ldots, b_r] = [Ab_1, \ldots, Ab_r].$$

Notice that the product AB of two maps A and B makes sense if and only if $\operatorname{dom} A \supset \operatorname{tar} B$. For matrices A and B, this means that the number of columns of A must equal the number of rows of B; we couldn't apply A to the columns of B otherwise.

In particular, *the 1-column matrix* $[Ax]$ *is the product of the matrix* A *with the 1-column matrix* $[x]$, i.e.,

$$A[x] = [Ax], \quad \forall A \in \mathbb{F}^{m\times n}, x \in \mathbb{F}^n.$$

For this reason, most books on elementary linear algebra and most users of linear algebra *identify* the n-vector x with the $n \times 1$-matrix $[x]$, hence write simply x for what I have denoted here by $[x]$. I will feel free from now on to use the same identification. However, I will not be doctrinaire about it. In particular, I will continue to specify a particular n-vector x by writing down its entries in a list, like $x = (x_1, x_2, \ldots)$, since that uses much less space than does the writing of

$$[x] = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \end{bmatrix}.$$

It is consistent with the standard identification of the n-vector x with the $n \times 1$-matrix $[x]$ to mean by x^{t} the $1 \times n$-matrix $[x]^{\mathrm{t}}$. Further, with y also an n-vector, one identifies the $(1,1)$-*matrix* $[x]^{\mathrm{t}}[y] = x^{\mathrm{t}}y$ with the *scalar*

$$\sum_j x_j y_j = y^{\mathrm{t}} x.$$

On the other hand,

$$yx^{\mathrm{t}} = [y][x]^{\mathrm{t}} \;=\; (y_i x_j : (i,j) \in \underline{n} \times \underline{n})$$

is an $n \times n$-*matrix* (and identified with a scalar only if $n = 1$).

However, I will *not* use the terms 'column vector' or 'row vector', as they don't make sense to me. Also, whenever I want to stress the fact that x or x^{t} is meant to be a matrix, I will write $[x]$ and $[x]^{\mathrm{t}}$, respectively.

For example, what about the expression $xy^{\mathrm{t}}z$ in case x, y, and z are vectors? It makes sense only if y and z are vectors of the same length, say $y, z \in \mathbb{F}^n$. In that case, it is $[x][y]^{\mathrm{t}}[z]$, and this we can compute in two ways: we can apply the matrix xy^{t} to the vector z, or we can multiply the vector x with the scalar $y^{\mathrm{t}}z$. Either way, we obtain the vector $x(y^{\mathrm{t}}z) = (y^{\mathrm{t}}z)x$, i.e., the $(y^{\mathrm{t}}z)$-multiple of x. However, while the product $x(y^{\mathrm{t}}z)$ of x with $(y^{\mathrm{t}}z)$ makes sense both as a matrix product and as multiplication of the vector x by the scalar $y^{\mathrm{t}}z$, the product $(y^{\mathrm{t}}z)x$ *only* makes sense as a product of the scalar $y^{\mathrm{t}}z$ with the vector x.

(2.10) Example: Here is an example, of help later. Consider the socalled **elementary row operation**

$$E_{i,k}(\alpha)$$

on n-vectors, in which one adds α times the kth entry to the ith entry. Is this a linear map? What is a formula for it?

We note that the kth entry of any n-vector x can be computed as ${e_k}^{\mathrm{t}}x$, while adding β to the ith entry of x is accomplished by adding βe_i to x. Hence, adding α times the kth entry of x to its ith entry replaces x by $x + e_i(\alpha {e_k}^{\mathrm{t}}x) = x + \alpha e_i {e_k}^{\mathrm{t}}x$. This gives the handy formula

$$E_{i,k}(\alpha) = \mathrm{id}_n + \alpha e_i {e_k}^{\mathrm{t}}. \tag{2.11}$$

Now, to check that $E_{i,j}(\alpha)$ is linear, we observe that it is the sum of two maps, and the first one, id_n, is certainly linear, while the second is the composition of the three maps,

$${e_k}^{\mathrm{t}} : \mathbb{F}^n \to \mathbb{F} \simeq \mathbb{F}^1 : z \mapsto {e_k}^{\mathrm{t}}z, \quad [e_i] : \mathbb{F}^1 \to \mathbb{F}^n : \beta \to e_i\beta, \quad \alpha : \mathbb{F}^n \to \mathbb{F}^n : z \mapsto \alpha z,$$

and each of these is linear (the last one because we assume $\mathbb{F}$ to be a *commutative* field).

Matrices of the form

$$E_{y,z}(\alpha) := \mathrm{id} + \alpha y z^{\mathrm{t}} \tag{2.12}$$

are called **elementary**. They are very useful since, if invertible, their inverse has the same simple form; see (2.19)Proposition below. □

2.9 Use the fact that the jth column of the matrix A is the image of e_j under the linear map A to construct the matrices that carry out the given action.

(i) The matrix A of order 2 that rotates the plane clockwise 90 degrees;

(ii) The matrix B that reflects $\mathbb{R}^n$ across the hyperplane $\{x \in \mathbb{R}^n : x_n = 0\}$;

(iii) The matrix C that keeps the hyperplane $\{x \in \mathbb{R}^n : x_n = 0\}$ pointwise fixed, and maps e_n to $-e_n$;

(iv) The matrix D of order 2 that keeps the y-axis fixed and maps $(1,1)$ to $(2,1)$.

2.10 Use the fact that the jth column of the matrix $A \in \mathbb{F}^{m\times n}$ is the image of e_j under A to derive the four matrices A^2, AB, BA, and B^2 for each of the given pair A and B: (i) $A = [e_1, 0]$, $B = [0, e_1]$; (ii) $A = [e_2, e_1]$, $B = [e_2, -e_1]$; (iii) $A = [e_2, e_3, e_1]$, $B = A^2$.

2.11 For each of the following pairs of matrices A, B, determine their products AB and BA if possible, or else state why it cannot be done.

(a) $A = \begin{bmatrix} 1 & -1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$, B the matrix `eye(2)`; (b) $A = \begin{bmatrix} 2 & 1 & 4 \\ 0 & 1 & 2 \end{bmatrix}$, $B = A^t$; (c) $A = \begin{bmatrix} 2 & 1 & 4 \\ 0 & 1 & 2 \\ 0 & 0 & -1 \end{bmatrix}$, $B = \begin{bmatrix} -1 & -1 & 2 \\ 0 & 2 & -1 \\ 0 & 0 & 3 \end{bmatrix}$;
(d) $A = \begin{bmatrix} 2+\mathrm{i} & 4-\mathrm{i} \\ 3-\mathrm{i} & 3+\mathrm{i} \end{bmatrix}$, $B = \begin{bmatrix} 2-\mathrm{i} & 3+\mathrm{i} & 3\mathrm{i} \\ 3-\mathrm{i} & 4+\mathrm{i} & 2 \end{bmatrix}$.

2.12 For any $A, B \in L(X)$, the products AB and BA are also linear maps on X, as are $A^2 := AA$ and $B^2 := BB$. Give an example of $A, B \in L(X)$ for which $(A+B)^2$ does not equal $A^2 + 2AB + B^2$. (Hint: keep it as simple as possible, by choosing X to be $\mathbb{R}^2$, hence both A and B are 2-by-2 matrices.)

2.13 Give an example of matrices A and B for which both $AB = 0$ and $BA = 0$, while neither A nor B is a zero matrix.

2.14 Prove: *If A and B are matrices with r rows, and C and D are matrices with c columns, and AC and BD are defined, then the product of the two partitioned matrices $[A, B]$ and $[C; D]$ is defined and equals $AC + BD$.*

2.15 Prove that both $\mathbb{C} \to \mathbb{R} : z \mapsto \operatorname{Re} z$ and $\mathbb{C} \to \mathbb{R} : z \mapsto \operatorname{Im} z$ are linear maps when we consider $\mathbb{C}$ as a vector space over the real scalar field.

The linear equation $A? = y$, and $\operatorname{ran} A$ and $\operatorname{null} A$

We are ready to recognize and use the fact that the general system

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= y_1 \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n &= y_2 \\ \cdots &= \cdot \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n &= y_m \end{aligned} \tag{2.13}$$

of m linear equations in the n unknowns $x_1, \ldots, x_n$ is equivalent to the vector equation

$$Ax = y,$$

provided

$$x := (x_1, \ldots, x_n), \quad y := (y_1, \ldots, y_m), \quad A := \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix}.$$

Here, **equivalence** means that the entries $x_1, \ldots, x_n$ of the n-vector x solve the system of linear equations (2.13) if and only if x solves the vector equation $A? = y$. This equivalence is not only a notational convenience. Switching from (2.13) to $A? = y$ is the conceptual shift that started Linear Algebra. It shifts the focus, from the scalars $x_1, \ldots, x_n$, to the vector x formed by them, and to the map A given by the coefficients in (2.13), its range and nullspace (about to be defined), and this makes for simplicity, clarity, and generality.

To stress the generality, we now give a preliminary discussion of the equation

$$A? = y$$

in case A is a *linear* map, from the vector space X to the vector space Y say, with y some element of Y.

Existence of a solution for every $y \in Y$ is equivalent to having A be *onto*, i.e., to having $\operatorname{ran} A = Y$. Now, the range of A is the linear image of a vector space, hence itself a vector space. Indeed, if $v_1, v_2, \ldots, v_m$

are elements of $\operatorname{ran} A$, then there must be a sequence $w_1, w_2, \ldots, w_m$ in X with $Aw_j = v_j$, all j. Since X is a vector space, it contains Wa for arbitrary $a \in \mathbb{F}^m$, therefore the corresponding linear combination $Va = [Aw_1, Aw_2, \ldots, Aw_m]a = (AW)a = A(Wa)$ must be in $\operatorname{ran} A$. In other words, if $V \subset \operatorname{ran} A$, then $\operatorname{ran} V \subset \operatorname{ran} A$.

Hence, if we wonder whether A is onto, and we happen to know an *onto* column map $[v_1, v_2, \ldots, v_m] = V \in L(\mathbb{F}^m, Y)$, then we only have to check that the *finitely many* columns, $v_1, v_2, \ldots, v_m$, of V are in $\operatorname{ran} A$. For, if some are not in $\operatorname{ran} A$, then, surely, A is not onto. However, if they all are in $\operatorname{ran} A$, then $Y = \operatorname{ran} V \subset \operatorname{ran} A \subset \operatorname{tar} A = Y$, hence $\operatorname{ran} A = Y$ and A is onto.

(2.14)Proposition: The range of a linear map $A \in L(X, Y)$ is a linear subspace, i.e., is nonempty and closed under vector addition and multiplication by a scalar.

If Y is the range of the column map V, then A is onto if and only if the finitely many columns of V are in $\operatorname{ran} A$.

Uniqueness of a solution for every $y \in Y$ is equivalent to having A be *1-1*, i.e., to have $Ax = Az$ imply that $x = z$. For a *linear* map $A : X \to Y$, we have $Ax = Az$ if and only if $A(x - z) = 0$. In other words, if $y = Ax$, then

$$A^{-1}\{y\} = x + \{z \in X : Az = 0\}. \tag{2.15}$$

In particular, A is 1-1 if and only if $\{z \in X : Az = 0\} = \{0\}$. In other words, to check whether a *linear* map is 1-1, we only have to check whether it is 1-1 'at' one particular point, e.g., 'at' 0. For this reason, the set $A^{-1}\{0\} = \{z \in X : Az = 0\}$ of all elements of X mapped by A to 0 is singled out.

Definition: The set

$$\operatorname{null} A := \{z \in \operatorname{dom} A : Az = 0\}$$

is called the **nullspace** or **kernel** of the linear map A.
The linear map is 1-1 if and only if its nullspace is **trivial**, i.e., contains only the zero vector.
The nullspace of a linear map is a linear subspace.

Almost all linear subspaces you'll meet will be of the form $\operatorname{ran} A$ or $\operatorname{null} A$ for some linear map A. These two ways of specifying a linear subspace are very different in character.

If we are told that our linear subspace Z of X is of the form $\operatorname{null} A$, for a certain linear map A on X, then we know, offhand, exactly one element of Z for sure, namely the element 0 which lies in every linear subspace. On the other hand, it is easy to *test* whether a given $x \in X$ lies in $Z = \operatorname{null} A$: simply compute Ax and check whether it is the zero vector.

If we are told that our linear subspace Z of X is of the form $\operatorname{ran} A$ for some linear map from some U into X, then we can 'write down' explicitly every element of $\operatorname{ran} A$: they are all of the form Au for some $u \in \operatorname{dom} A$. On the other hand, it is much harder to *test* whether a given $x \in X$ lies in $Z = \operatorname{ran} A$: Now we have to check whether the equation $A? = x$ has a solution (in U).

As a simple example, the vector space Π_k of all polynomials of degree $\le k$ is usually specified as the range of the column map

$$[()^0, ()^1, \ldots, ()^k] : \mathbb{R}^{k+1} \to \mathbb{R}^{\mathbb{R}},$$

with

$$()^j : \mathbb{R} \to \mathbb{R} : t \mapsto t^j$$

a convenient (though non-standard!) notation for the **monomial of degree** j, i.e., as the collection of all real-valued functions that are of the form

$$t \mapsto a_0 + a_1 t + \cdots + a_k t^k$$

for some coefficient-vector a. On the other hand, Π_k can also be defined as $\operatorname{null} D^{k+1}$, i.e., as the collection of all real-valued functions that are $k+1$-times continuously differentiable and have their $(k+1)$st derivative identically zero.

(2.16) Remark: The nullspace $\operatorname{null} A$ of the linear map $A : X \to Y$ consists exactly of the solutions to the *homogeneous* equation

$$A? = 0.$$

The linear equation $A? = y$ is readily associated with a *homogeneous* linear equation, namely the equation

$$[A, y]? = 0,$$

with

$$[A, y] : X \times \mathbb{F} : (z, \alpha) \mapsto Az + y\alpha.$$

If $Ax = y$, then $(x, -1)$ is a nontrivial element of $\operatorname{null}[A, y]$. Conversely, if $(z, \alpha) \in \operatorname{null}[A, y]$ *and* $\alpha \neq 0$, then $z/(-\alpha)$ is a solution to $A? = y$. Hence, for the construction of solutions to linear equations, it is sufficient to know how to solve *homogeneous* linear equations, i.e., how to construct the nullspace of a linear map.

2.16 For each of the following systems of linear equations, determine A and y of the equivalent vector equation $A? = y$.
(a) $\begin{array}{rcrcr} 2x_1 & - & 3x_2 & = & 4 \\ 4x_1 & + & 2x_2 & = & -6 \end{array}$; (b) $\begin{array}{rcrcr} 2u_1 & - & 3u_2 & = & 4 \\ 4u_1 & + & 2u_2 & = & -6 \end{array}$; (c) $\begin{array}{rcrcrcr} & & & & -4c & = & 16 \\ 2a & + & 3b & & & = & 9 \end{array}$.

2.17 For each of the following A and y, write out a system of linear equations equivalent to the vector equations $A? = y$.
(a) $A = \begin{bmatrix} 2 & 3 \\ 6 & 4 \\ \mathrm{e} & -2 \end{bmatrix}$, $y = (9, -\sqrt{3}, 1)$; (b) $A = \begin{bmatrix} 1 & 2 & 3 & 4 \\ 4 & 3 & 2 & 1 \end{bmatrix}$, $y = (10, 10)$. (c) $A = [] \in \mathbb{R}^{0\times 3}$, $y = () \in \mathbb{R}^0$.

Inverses

We have agreed to think of the *matrix* $A \in \mathbb{F}^{m\times n}$ as the *column map* $[A(:,1), \ldots, A(:,n)]$, i.e., as the linear map $\mathbb{F}^n \to \mathbb{F}^m : a \mapsto Aa := \sum_j A(:,j)a_j$. For this reason, it is also customary to refer to the range $\operatorname{ran} A$ of a matrix A as the **column space** of that matrix, while the range $\operatorname{ran} A^{\mathrm{t}}$ of its transpose is known as its **row space**. Further, we have found that, in these terms, the *matrix product* AB is also the *composition* $A \circ B$, i.e.,

$$(A \circ B)a = A(B(a)) = (AB)a = \sum_j (AB)(:,j)a_j.$$

In these terms, the identity map id_n on $\mathbb{F}^n$ corresponds to the **identity matrix** $[e_1, e_2, \ldots, e_n]$, hence the name for the latter.

(2.17) Proposition: The inverse of a linear map is again a linear map.

Proof: Let $A \in L(X, Y)$ be invertible and $y, z \in Y$. By additivity of A, $A(A^{-1}y + A^{-1}z) = A(A^{-1}y) + A(A^{-1}z) = y + z$. Hence, applying A^{-1} to both sides, we get $A^{-1}y + A^{-1}z = A^{-1}(y + z)$, thus A^{-1} is additive. Also, from $A(\alpha A^{-1}y) = \alpha A(A^{-1}y) = \alpha y$, we conclude that $\alpha A^{-1}y = A^{-1}(\alpha y)$, hence A^{-1} is homogeneous. □

Thus, if $A \in \mathbb{F}^{n\times n}$ is invertible (as a linear map from $\mathbb{F}^n$ to $\mathbb{F}^n$), then also its inverse is a linear map (from $\mathbb{F}^n$ to $\mathbb{F}^n$), hence a square matrix of order n. We call it the **inverse** matrix for A, and denote it by A^{-1}. Being the inverse for A, it is both a right and a left inverse for A, i.e., it satisfies

$$A^{-1}A = \mathrm{id}_n = AA^{-1}.$$

More generally, we would call $A \in \mathbb{F}^{m\times n}$ invertible if there were $B \in \mathbb{F}^{n\times m}$ so that

$$AB = \mathrm{id}_m \quad \text{and} \quad BA = \mathrm{id}_n.$$

However, we will soon prove (cf. (3.17)) that this can only happen when $m = n$.

We will also soon prove (cf. (3.16)Theorem below) the *pigeonhole principle for square matrices*, i.e., that a linear map from $\mathbb{F}^n$ to $\mathbb{F}^n$ is 1-1 if and only if it is onto. In other words, if $A, B \in \mathbb{F}^{n\times n}$ and, e.g., $AB = \mathrm{id}_n$, hence A is onto, then A must also be 1-1, hence invertible, and therefore its right inverse must be its inverse, therefore we must also have $BA = \mathrm{id}_n$. In short:

(2.18) Amazing Fact: If $A, B \in \mathbb{F}^{n\times n}$ and $AB = \mathrm{id}_n$, then also $BA = \mathrm{id}_n$.

To me, this continues to be one of the most remarkable results in basic Linear Algebra. Its proof uses nothing more than the identification of matrices with linear maps (between coordinate spaces) and the numerical process called *elimination*, for solving a homogeneous linear system $A? = 0$, i.e., for constructing $\operatorname{null} A$.

In preparation, and as an exercise in invertible matrices, we verify the following useful fact about elementary matrices.

(2.19) Proposition: For $x, y \in \mathbb{F}^n$ and $\alpha \in \mathbb{F}$, the elementary matrix

$$E_{y,z}(\alpha) = \mathrm{id}_n + \alpha y z^{\mathrm{t}}$$

is invertible if and only if $1 + \alpha z^{\mathrm{t}} y \neq 0$, and, in that case

$$E_{y,z}(\alpha)^{-1} = E_{y,z}\left(\frac{-\alpha}{1 + \alpha z^{\mathrm{t}} y}\right). \qquad (2.20)$$

Proof: We compute $E_{y,z}(\alpha)E_{y,z}(\beta)$ for arbitrary α and β. Since

$$\alpha y z^{\mathrm{t}}\ \beta y z^{\mathrm{t}} = \alpha\beta\ (z^{\mathrm{t}} y)\ y z^{\mathrm{t}},$$

we conclude that

$$E_{y,z}(\alpha)E_{y,z}(\beta) = (\mathrm{id}_n + \alpha y z^{\mathrm{t}})(\mathrm{id}_n + \beta y z^{\mathrm{t}}) = \mathrm{id}_n + (\alpha + \beta + \alpha\beta(z^{\mathrm{t}} y)) y z^{\mathrm{t}}.$$

In particular, since the factor $(\alpha + \beta + \alpha\beta(z^{\mathrm{t}} y))$ is symmetric in α and β, we conclude that

$$E_{y,z}(\alpha)E_{y,z}(\beta) = E_{y,z}(\beta)E_{y,z}(\alpha).$$

Further, if $1 + \alpha z^{\mathrm{t}} y \neq 0$, then the choice

$$\beta = \frac{-\alpha}{1 + \alpha z^{\mathrm{t}} y}$$

will give $\alpha + \beta + \alpha\beta(z^{\mathrm{t}} y) = 0$, hence $E_{y,z}(\beta)E_{y,z}(\alpha) = E_{y,z}(\alpha)E_{y,z}(\beta) = \mathrm{id}_n$. This proves that $E_{y,z}(\alpha)$ is invertible, with its inverse given by (2.20).

Conversely, assume that $1 + \alpha z^{\mathrm{t}} y = 0$. Then $y \neq 0$, yet

$$E_{y,z}(\alpha) y = y + \alpha(z^{\mathrm{t}} y) y = 0,$$

showing that $E_{y,z}(\alpha)$ is not 1-1 in this case, hence not invertible. □

2.18 Prove: *If two matrices commute (i.e., $AB = BA$), then they are square matrices, of the same order.*

2.19 Give a noninvertible 2-by-2 matrix without any zero entries.

2.20 Prove that the matrix $A := \begin{bmatrix} 1 & 2 \\ 4 & -1 \end{bmatrix}$ satisfies the equation $A^2 = 9\,\mathrm{id}_2$. Use this to show that A is invertible, and to write down the matrix A^{-1}.

2.21 Prove: *The matrix* $A := \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ *is invertible if and only if* $ad \neq bc$, *in which case* $\begin{bmatrix} d & -b \\ -c & a \end{bmatrix}/(ad-bc)$ *is its inverse.*

2.22 Consider the map $f : \mathbb{C} \to \mathbb{R}^{2\times 2} : z = a + \mathrm{i}b \mapsto \begin{bmatrix} a & -b \\ b & a \end{bmatrix}$. Show that f is a 1-1 linear map when we think of $\mathbb{C}$ as a vector space over the real scalar field.

2.23 Let $A, B \in L(X)$. Show that $(AB)^2 = A^2B^2$ can hold without necessarily having $AB = BA$. Show also that $(AB)^2 = A^2B^2$ implies that $AB = BA$ in case both A and B are invertible.

2.24 Give an example of matrices A and B, for which both AB and BA are defined and for which $AB = \mathrm{id}$, but neither A nor B is invertible.

2.25 Prove: *If A and C are invertible matrices, and B has as many rows as does A and as many columns as does C, then also $[A, B; 0, C]$ is invertible and*

$$[A, B; 0, C]^{-1} = \begin{bmatrix} A & B \\ 0 & C \end{bmatrix}^{-1} = \begin{bmatrix} A^{-1} & -A^{-1}BC^{-1} \\ 0 & C^{-1} \end{bmatrix}.$$

2.26 A square matrix A is called **diagonally dominant** if $|A_{ii}| > \sum_{j\neq i} |A_{ij}|$ for all i. Prove: *a diagonally dominant matrix is invertible.* (Hint: Prove the contrapositive: if $0 \neq x \in \mathrm{null}\, A$, then, for some i, $|A_{ii}| \leq \sum_{j\neq i} |A_{ij}|$.)

2.27 Use (2.19)Proposition to prove the **Sherman-Morrison Formula**: *If $A \in \mathbb{F}^{n\times n}$ is invertible and $y, z \in \mathbb{F}^n$ are such that $\alpha := 1 + z^{\mathrm{t}}A^{-1}y \neq 0$, then $A + yz^{\mathrm{t}}$ is invertible, and*

$$(A + yz^{\mathrm{t}})^{-1} = A^{-1} - \alpha^{-1}A^{-1}yz^{\mathrm{t}}A^{-1}.$$

(Hint: $A + yz^{\mathrm{t}} = A(\,\mathrm{id} + (A^{-1}y)z^{\mathrm{t}})$.)

2.28 Prove the **Woodbury** generalization of the Sherman-Morrison Formula: *if A and $\mathrm{id} + D^{\mathrm{t}}AC$ are invertible, then so is $A + CD^{\mathrm{t}}$, and*

$$(A + CD^{\mathrm{t}})^{-1} = A^{-1} - A^{-1}C(\,\mathrm{id} + D^{\mathrm{t}}A^{-1}C)^{-1}D^{\mathrm{t}}A^{-1}.$$

2.29 T/F

(a) If $A, B \in L(X, Y)$ are both invertible, then so is $A + B$.

(b) If $AB = 0$ for $A, B \in \mathbb{F}^{n\times n}$, then $B = 0$.

(c) If A and B are matrices with $AB = \mathrm{id}_m$ and $BA = \mathrm{id}_n$, then $B = A^{-1}$.

(d) If $A = \begin{bmatrix} B & C \\ 0 & 0 \end{bmatrix}$ with both A and B square matrices and 0 standing for zero matrices of the appropriate size, then $A^n = \begin{bmatrix} B^n & B^{n-1}C \\ 0 & 0 \end{bmatrix}$ for all n.

(e) If $A \in \mathbb{R}^{m\times n}$ and $A^{\mathrm{t}}A = 0$, then $A = 0$.

(f) If the matrix product AB is defined, then $(AB)^{\mathrm{t}} = A^{\mathrm{t}}B^{\mathrm{t}}$.

(g) If A is an invertible matrix, then so is A^{t}, and $(A^{\mathrm{t}})^{-1} = (A^{-1})^{\mathrm{t}}$.

(h) $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}$ is an elementary matrix.

(i) If Y is a subset of some vector space X, x, y, z are particular elements of X, and x and $2y - 3x$ are in Y, but $3y - 2x$ and y are not, then Y cannot be a linear subspace.

3. Elimination, or: The determination of $\operatorname{null} A$ and $\operatorname{ran} A$

Elimination and Backsubstitution

Elimination has as its goal an efficient description of the solution set for the *homogeneous* linear system $A? = 0$, i.e., of the nullspace of the matrix A. It is based on the following observation:

(3.1) Lemma: If B is obtained from A by subtracting some multiple of some row of A from some *other* row of A, then $\operatorname{null} B = \operatorname{null} A$.

Proof: Assume, more specifically, that B is obtained from A by subtracting α times row k from row i, for some $k \neq i$. Then, by (2.10)Example,

$$B = E_{i,k}(-\alpha)\, A,$$

with $E_{i,k}(-\alpha) = \mathrm{id}_m - \alpha e_i {e_k}^{\mathrm{t}}$. Consequently, $\operatorname{null} B \supset \operatorname{null} A$, and this holds even if $i = k$.

However, since $i \neq k$, we have ${e_k}^{\mathrm{t}} e_i = 0$, hence, for any α, $1 + \alpha({e_k}^{\mathrm{t}} e_i) = 1 \neq 0$. Therefore, by (2.19), also

$$E_{i,k}(\alpha)\, B = A,$$

hence also $\operatorname{null} B \subset \operatorname{null} A$. □

One solves the homogeneous linear system $A? = 0$ by **elimination**. This is an *inductive* process, and it results in a classification of the unknowns as *free* or *bound*. A **bound** unknown has associated with it a **pivot row** or **pivot equation** which determines this unknown uniquely once all later unknowns are determined. Any unknown without a pivot equation is a **free** unknown; its value can be chosen arbitrarily. We call the jth *column* of A bound (free) if the jth unknown is bound (free). The classification proceeds inductively, from the first to the last unknown or column, i.e., for $k = 1, 2, \ldots$, with the kth step as follows.

At the beginning of the kth **elimination step**, we have in hand a matrix B, called the **working-array**, which is **equivalent** to our initial matrix A in that $\operatorname{null} B = \operatorname{null} A$. Further, we have already classified the first $k-1$ unknowns as either bound or free, with each bound unknown associated with a particular row of B, its *pivot row*, and this row having a nonzero entry at the position of its associated bound unknown and zero entries for all previous unknowns. All other rows of B are nonpivot rows; they do not involve the unknowns already classified, i.e., they have nonzero entries only for unknowns not yet classified. (Note that, with the choice $B := A$, this description also fits the situation at the beginning of the first step.) We now classify the kth unknown or column and, correspondingly, change B, as follows:

bound case: We call the kth unknown or column **bound** (some would say **basic**) in case we can find some nonpivot row $B(h,:)$ for which $B(h,k) \neq 0$. We pick one such row and call it the **pivot row** for the kth unknown. Further, we use it to eliminate the kth unknown from all the remaining nonpivot rows $B(i,:)$ by the calculation

$$B(i,:) \leftarrow B(i,:) - \frac{B(i,k)}{B(h,k)} B(h,:).$$

free case: In the contrary case, we call the kth unknown or column **free** (some would say **nonbasic**). No action is required in this case, since none of the nonpivot rows involves the kth unknown.

By (3.1)Lemma, the changes (if any) made in B will not change $\operatorname{null} B$. This finishes the kth elimination step.

For future reference, here is a formal description of the entire algorithm. This description relies on a sequence p to keep track of which row, if any, is used as pivot row for each of the unknowns. If row h is the pivot row for the kth unknown, then $p(k) = h$ after the kth elimination step. Since p is initialized to have all its entries equal to 0, this means that, at any time, the rows k not yet used as pivot rows are exactly those for which $p(k) = 0$.

(3.2) Elimination Algorithm:
input: $A \in \mathbb{F}^{m\times n}$.
$B \leftarrow A$, $p \leftarrow (0,\ldots,0) \in \mathbb{R}^n$.
for $k = 1{:}n$, **do**:
  for some $h \notin \operatorname{ran} p$ with $B(h,k) \neq 0$, **do**:
    $p(k) \leftarrow h$
    for all $i \notin \operatorname{ran} p$, **do**:

$$B(i,:) \leftarrow B(i,:) - \frac{B(i,k)}{B(h,k)}\, B(h,:)$$

    enddo
  enddo
enddo
output: B, p, and, possibly, `free` $\leftarrow$ `find(`p`==0)`, `bound` $\leftarrow$ `find(`p`>0)`.

Note that nothing is done at the kth step if there is no $h \notin \operatorname{ran} p$ with $B(h,k) \neq 0$, i.e., if $B(h,k) = 0$ for all $h \notin \operatorname{ran} p$. In particular, $p(k)$ will remain 0 in that case.

A numerical example: We start with

$$A := \begin{bmatrix} 0 & 2 & 0 & 2 & 5 & 4 & 0 & 6 \\ 0 & 1 & 0 & 1 & 2 & 2 & 0 & 3 \\ 0 & 2 & 0 & 2 & 5 & 4 & -1 & 7 \\ 0 & 1 & 0 & 1 & 3 & 2 & -1 & 4 \end{bmatrix}, \qquad p = (0,0,0,0,0,0,0,0).$$

The first unknown is free. We take the second row as pivot row for the second unknown and eliminate it from the remaining rows, to get

$$B = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & \underline{\mathbf{1}} & 0 & 1 & 2 & 2 & 0 & 3 \\ 0 & 0 & 0 & 0 & 1 & 0 & -1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & -1 & 1 \end{bmatrix}, \qquad p = (0,2,0,0,0,0,0,0).$$

Thus the third unknown is free as is the fourth, but the fifth is not, since there are nonzero entries in the fifth column of some nonpivotal row, e.g., the first row. We choose the first row as pivot row for the fifth unknown and use it to eliminate this unknown from the remaining nonpivot rows, i.e., from rows 3 and 4. This gives

$$B = \begin{bmatrix} 0 & 0 & 0 & 0 & \underline{\mathbf{1}} & 0 & 0 & 0 \\ 0 & \mathbf{1} & 0 & 1 & 2 & 2 & 0 & 3 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 \end{bmatrix}, \qquad p = (0,2,0,0,1,0,0,0).$$

The sixth unknown is free, but there are nonzero entries in the seventh column of the remaining nonpivot rows, so the seventh unknown is bound, with, e.g., the fourth row as its pivot row. We use that row to eliminate the seventh unknown from the remaining nonpivot row. This gives

$$B = \begin{bmatrix} 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 \\ 0 & \mathbf{1} & 0 & 1 & 2 & 2 & 0 & 3 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \underline{\mathbf{-1}} & 1 \end{bmatrix}, \qquad p = (0,2,0,0,1,0,4,0).$$

With that, there are no nontrivial nonpivot rows left. In particular, the eighth unknown is free, hence we have already in hand the final array.

Hence, altogether `bound` $= (2,5,7)$ ($=$ `find(p>0)`) and `free` $= (1,3,4,6,8)$ ($=$ `find(p==0)`). □

After the n steps of this elimination process (which started with $B = A$), we have in hand a matrix B with $\operatorname{null} B = \operatorname{null} A$ and with each unknown classified as bound or free. The two increasing sequences, `bound` and `free`, containing the indices of the bound and free unknowns respectively, will be much used in the sequel. Each bound unknown has associated with it a particular row of B, its pivot row. All nonpivot rows of B (if any) are entirely zero.

Neat minds would reorder the rows of B, listing first the pivot rows in order, followed by the nonpivot rows and, in this way, obtain a **row echelon form** for A. In any case, in determining $x \in \operatorname{null} B$, we only have to pay attention to the pivot rows. This means that we can determine a particular element x of $\operatorname{null} B = \operatorname{null} A$ by *backsubstitution*, i.e., from its last entry to its first as follows:

For $k = n{:}{-1}{:}1$, if the kth unknown is bound, i.e., $k \in$ `bound`, determine x_k from its pivot equation (since that equation only involves $x_k, \cdots, x_n$); else, pick x_k arbitrarily (as then the kth unknown is free, i.e., $k \in$ `free`).

Here is a more formal description, for future reference.

(3.3) Backsubstitution Algorithm:
input: $B \in \mathbb{F}^{m \times n}$ and p (both as output from (3.2)), $z \in \mathbb{F}^n$.
$x \leftarrow z$
for $k = n{:}{-1}{:}1$, **do**:
  if $p(k) \neq 0$, **then** $x_k \leftarrow -\left(\sum_{j>k} B(p(k), j) x_j\right) / B(p(k), k)$ **endif**
enddo
output: x, which is the unique solution of $A? = 0$ satisfying $x_i = z_i$ for all i with $p(i) = 0$.

Notice that the value of every free unknown is arbitrary and that, once these are chosen somehow, then the bound unknowns are uniquely determined by the requirement that we are seeking an element of $\operatorname{null} B = \operatorname{null} A$. In other words, the general element of $\operatorname{null} B$ has exactly as many degrees of freedom as there are free unknowns. Since there are #`free` unknowns, $\operatorname{null} B$ is said to be 'of dimension #`free`'.

In particular, for any k, the kth entry, x_k, of an $x \in \operatorname{null} B$ can be nonzero only in one of two ways: (a) the kth unknown is free, i.e., $k \in$ `free`; (b) the kth unknown is bound, but $x_j \neq 0$ for some $j > k$. It follows that x_k can be the rightmost nonzero entry of such an x only if the kth unknown is free. Conversely, if the kth unknown is free, and x is the element of $\operatorname{null} B = \operatorname{null} A$ computed by setting $x_k = 1$ and setting all other free entries equal to 0, then x_k is necessarily the rightmost nonzero entry of x (since all free entries to the right of it were chosen to be zero, thus preventing any bound entry to the right of it from being nonzero).

This proves

(3.4) Observation: There exists $x \in \operatorname{null} A$ with rightmost nonzero entry x_k if and only if the kth unknown is free.

This simple observation gives a *characterization* of the sequence `free` entirely in terms of the nullspace of the matrix A we started with. This implies that *the classification into free and bound unknowns or columns is independent of all the details of the elimination*. More than that, since, for any 1-1 matrix M with m columns, $\operatorname{null}(MA) = \operatorname{null} A$, it implies that, for any such matrix MA, we get exactly the same sequences `free` and `bound` as we would get for A. This is the major reason for the uniqueness of a more disciplined echelon form, the 'really reduced row echelon form', to be discussed in the next section.

Since $A(:, k) \in \operatorname{ran} A(:, [1{:}k{-}1])$ if and only if there is some $x \in \operatorname{null} A$ whose rightmost nonzero entry is its kth, we have the following reformulation of (3.4)Observation and consequences.

(3.5) Corollary:

(i) The kth column of A is free if and only if it is a weighted sum of the columns strictly to the left of it, i.e., $A(:,k) \in \operatorname{ran} A(:,1{:}k-1)$.

(ii) $A(:,1{:}k)$ is 1-1 if and only if all its columns are bound.

(iii) $\operatorname{null} A$ is nontrivial if and only if there are free columns.

Perhaps the most widely used consequence of (iii) here is the following. If there are more unknowns than equations, then there are not enough equations to go around, i.e., some unknowns must be free, therefore there are nontrivial solutions to our homogeneous equation $A? = 0$. We remember this fundamental result of elimination in the following form:

(3.6) Theorem: Any matrix with more columns than rows has a nontrivial nullspace.

3.1 Determine the bound and free columns for each of the following matrices A.

(a) $0 \in \mathbb{R}^{m\times n}$; (b) $[e_1,\ldots,e_n] \in \mathbb{R}^{n\times n}$; (c) $[e_1, 0, e_2, 0] \in \mathbb{R}^{6\times 4}$; (d) $\begin{bmatrix} 2 & 2 & 5 & 6 \\ 1 & 1 & -2 & 2 \end{bmatrix}$; (e) $\begin{bmatrix} 0 & 2 & 1 & 4 \\ 0 & 0 & 2 & 6 \\ 1 & 0 & -3 & 2 \end{bmatrix}$; (f) $[x][y]^{\mathrm{t}}$, with $x = (1,2,3,4) = y$.

3.2 (3.5)Corollary assures you that $y \in \operatorname{ran} A$ if and only if the last column of $[A, y]$ is free. Use this fact to determine, for each of the following y and A, whether or not $y \in \operatorname{ran} A$.

(a) $y = (\pi, 1-\pi)$, $A = \begin{bmatrix} 1 & -2 \\ -1 & 2 \end{bmatrix}$; (b) $y = e_2$, $A = \begin{bmatrix} 1 & 2 & -1 \\ 2 & 3 & -4 \\ 3 & 4 & -8 \end{bmatrix}$; (c) $y = e_2$, $A = \begin{bmatrix} 1 & 2 & -1 \\ 2 & 3 & -4 \\ 3 & 4 & -7 \end{bmatrix}$.

3.3 Prove (3.1)Lemma directly, i.e., without using (2.19)Proposition. (Hint: Prove that $\operatorname{null} B \supset \operatorname{null} A$. Then prove that also A is obtainable from B by the same kind of step, hence also $\operatorname{null} A \supset \operatorname{null} B$.)

3.4 Prove: *If M and A are matrices for which MA is defined and, furthermore, M is 1-1, then $MA? = 0$ has exactly the same free and bound unknowns as does $A? = 0$.*

3.5 Assuming the matrix A has exactly α bound columns and the matrix B has exactly β bound columns and both have the same number of rows, how many bound columns does the matrix $[A, B]$ have (a) at least? (b) at most? (c) How, if at all, would your answers to (a), (b) change if I told you that A has m rows?

The really reduced row echelon form and other reduced forms

The construction of the *really reduced row echelon form* takes elimination four steps further, none of which changes the nullspace:

(i) When the hth pivot row is found, and it is not the hth row, then it is exchanged with the current hth row to make it the hth row. (This keeps things neat; all the rows not yet used as pivot rows lie below all the rows already picked as pivot rows.)

(ii) Each pivot row is divided by its **pivot element**, i.e., by its left-most nonzero entry. (This helps with the elimination of the corresponding unknown from other rows: if $B(h,k)$ is the pivot element in question (i.e., $\mathtt{bound}(h) = k$, i.e., x_k is the hth bound unknown), then, after this normalization, one merely subtracts $B(i,k)$ times $B(h,:)$ from $B(i,:)$ to eliminate the kth unknown from row i.)

(iii) One eliminates each bound unknown from *all* rows (other than its pivot row), i.e., also from pivot rows belonging to earlier bound unknowns, and not just from the rows not yet used as pivot rows. For real efficiency, though, this additional step should be carried out after elimination is completed; it starts with the elimination of the *last* bound unknown, proceeds to the second-last bound unknown, etc., and ends with the *second* bound unknown (the first bound unknown was eliminated from all other rows already).

The resulting matrix B is called the **reduced row echelon form for** A, and this is written:

$$B = \operatorname{rref}(A).$$

However, it turns out to be very neat to add the following final step:

(iv) Remove all rows that are entirely zero, thus getting the matrix

$$R := B(1{:}\#\texttt{bound}, :) =: \operatorname{rrref}(A)$$

called the *really reduced row echelon form* of A.

Here is a formal description (in which we talk about *the* rrref for A even though we prove its *uniqueness* only later, in (3.12)):

(3.7) Definition: We say that R is the **really reduced row echelon form for** $A \in \mathbb{F}^{m\times n}$ and write $R = \operatorname{rrref}(A)$, in case $R \in \mathbb{F}^{r\times n}$ for some r and there is a strictly increasing r-sequence `bound` (provided by the `MATLAB` function `rref` along with $\operatorname{rref}(A)$) so that the following is true:

1. R is a **row echelon form for** A: This means that (i) $\operatorname{null} R = \operatorname{null} A$; and (ii) for each $k = \texttt{bound}(i)$, $R(i,:)$ is the pivot row for the kth unknown, i.e., $R(i,:)$ is the unique row in R for which $R(i,k)$ is the first (or, leftmost) nonzero entry.
2. R is **really reduced** or normalized, in the sense that $R(:,\texttt{bound})$ is the identity matrix, i.e., for each i, the pivot element $R(i,\texttt{bound}(i))$ equals 1 and is the only nonzero entry in its column, and R has only these $r = \#\texttt{bound}$ rows.

A numerical example, continued: For the earlier numerical example, the rref and the rrref would look like this:

$$\operatorname{rref}(A) = \begin{bmatrix} 0 & \mathbf{1} & 0 & 1 & 0 & 2 & 0 & 3 \\ 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \underline{\mathbf{1}} & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \qquad \operatorname{rrref}(A) = \begin{bmatrix} 0 & \mathbf{1} & 0 & 1 & 0 & 2 & 0 & 3 \\ 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \underline{\mathbf{1}} & -1 \end{bmatrix}.$$

Recall (or observe directly) that, for this example, $\texttt{bound} = (2,5,7)$ and $\texttt{free} = (1,3,4,6,8)$. □

Finally, for most purposes, it is sufficient to have a **b-form** for A, i.e., a matrix R that satisfies the following two conditions:

(3.8)(i) $\operatorname{null} R = \operatorname{null} A$;

(3.8)(ii) $R(:,\texttt{b}) = \operatorname{id}$ for some sequence `b`.

Certainly, in these terms, the $\operatorname{rrref}(A)$ is a `bound`-form for A, but a matrix A may have a `b`-form for many different `b` and, as we shall see, only the two conditions (3.8)(i-ii) really matter. Moreover, we have, in effect, a `b`-form for A in hand well before we get to $\operatorname{rrref}(A)$. For, there is no need to reorder the rows of the working array; we merely eliminate each bound unknown from all rows but its pivot row, being sure to divide each pivot row by its pivot element, drop any non-pivot rows, and then have in hand the `b`-form for A, with `b` the inverse of `p(p>0)`.

3.6 For each of the matrices A in H.P. 3.1 , determine its rrref.

A complete description for $\operatorname{null} A$ obtained from a `b`-form

If R is a `b`-form for A, then it is easy to determine all solutions of the homogeneous linear system $A? = 0$, i.e., all the elements of $\operatorname{null} A$.

In recognition of the special case $R = \operatorname{rrref}(A)$, I'll use `f` for a sequence **complementary to** `b` in the sense that it contains all the indices in $\underline{n}$ but not in `b`.

In `MATLAB`, one would obtain `f` from n and `b` by the commands `f = 1:n; f(b) = [];` □

We now obtain from any $\texttt{b}$-form R for A a 1-1 matrix C with the property that $\operatorname{null} A = \operatorname{ran} C$, thus getting a description both as a range and as a nullspace. Since such a C is 1-1 onto $\operatorname{null} A$, this implies that every $x \in \operatorname{null} A$ can be written *in exactly one way* in the form $x = Ca$. We will soon agree to call such a C a 'basis' for the vector space $\operatorname{null} A$.

In the discussion, we use the following notation introduced earlier: If x is an n-vector and p is a list of length r with range in $\underline{n}$, then x_p is the r-vector

$$x_p = (x_{p(i)} : i = 1{:}r).$$

With this, by property (3.8)(i),

$$x \in \operatorname{null} A \quad \Longleftrightarrow \quad 0 = Rx = \sum_j R(:,j)x_j = R(:,\texttt{b})x_{\texttt{b}} \;+\; R(:,\texttt{f})x_{\texttt{f}}\ .$$

Since $R(:,\texttt{b}) = \mathrm{id}$ by property (3.8)(ii), we conclude that

$$x \in \operatorname{null} A \quad \Longleftrightarrow \quad x_{\texttt{b}} = -R(:,\texttt{f})x_{\texttt{f}}.$$

We can write this even more succinctly in matrix form as follows:

$$\operatorname{null} A = \operatorname{ran} C,$$

with C the $(n \times \#\texttt{f})$-matrix whose '$\texttt{f}$-rows' form an identity matrix, and whose '$\texttt{b}$-rows' are formed by the '$\texttt{f}$-columns' of $-R$:

(3.9) $$C(\texttt{f},:) = \mathrm{id}, \qquad C(\texttt{b},:) = -R(:,\texttt{f}).$$

E.g., for the earlier numerical example and with $R = \mathrm{rrref}(A)$,

$$C = \begin{bmatrix} \mathbf{1} & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & -2 & -3 \\ 0 & \mathbf{1} & 0 & 0 & 0 \\ 0 & 0 & \mathbf{1} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \mathbf{1} & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & \mathbf{1} \end{bmatrix} = \begin{bmatrix} \mathbf{1} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & \mathbf{1} & 0 & 0 & 0 \\ 0 & 0 & \mathbf{1} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \mathbf{1} & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbf{1} \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ -0 & -0 & -1 & -2 & -3 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ -0 & -0 & -0 & -0 & -0 \\ 0 & 0 & 0 & 0 & 0 \\ -0 & -0 & -0 & -0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Note that C is 1-1, since $x := Ca = 0$ implies that $0 = x_{\texttt{f}} = C(\texttt{f},:)a = a$. Therefore, C is (or, the columns of C form) a 'basis' for $\operatorname{null} A$, in the sense that C is a 1-1 onto column map to $\operatorname{null} A$.

Finally, when $R = \mathrm{rrref}(A)$, then the resulting C is 'upper triangular' in the sense that then

(3.10) $$i > \texttt{free}(j) \quad \Longrightarrow \quad C(i,j) = 0.$$

3.7 Determine a 'basis' for the nullspace of $A := \begin{bmatrix} 1 & 1 \\ 2 & 2 \end{bmatrix}$ and use it to describe the solution set of the system $A? = (1,2)$. Draw a picture indicating both the solution set and $\operatorname{null} A$.

3.8 For each of the matrices A in H.P. 3.1 , give a 'basis' for $\operatorname{null} A$.

The factorization $A = A(:,\mathtt{bound})\operatorname{rrref}(A)$

Continuing with our $\mathtt{b}$-form R for A, we claim that

$$A(:,\mathtt{b})R = A.$$

For the proof, we compare $A(:,\mathtt{b})R =: M$ and A column by column. First, $M(:,\mathtt{b}) = A(:,\mathtt{b})R(:,\mathtt{b}) = A(:,\mathtt{b})$, by property (3.8)(ii). As to $M(:,\mathtt{f}) = A(:,\mathtt{b})R(:,\mathtt{f})$, we observe that, for any c (of length $\#\mathtt{f}$), the vector x with

$$x_{\mathtt{b}} := R(:,\mathtt{f})c, \qquad x_{\mathtt{f}} := -c,$$

is in $\operatorname{null} R = \operatorname{null} A$, hence

$$0 = Ax = A(:,\mathtt{b})x_{\mathtt{b}} + A(:,\mathtt{f})x_{\mathtt{f}} = A(:,\mathtt{b})R(:,\mathtt{f})c + A(:,\mathtt{f})(-c).$$

In other words,

$$M(:\mathtt{f})c = A(:,\mathtt{b})R(:,\mathtt{f})c = A(:,\mathtt{f})c, \quad \forall c \in \mathbb{F}^{\#\mathtt{f}},$$

showing that also $M(:,\mathtt{f}) = A(:,\mathtt{f})$. This proves our claim that $A(:,\mathtt{b})R = A$, hence, in particular,

(3.11) $$A = A(:,\mathtt{bound})\operatorname{rrref}(A).$$

3.9 Prove: *If M is such that $MA = \operatorname{rrref}(A) =: R$, and $\mathtt{bound}$ is the increasing sequence of indices of bound columns of A, then M is a left inverse for $A(:,\mathtt{bound})$.*

A 'basis' for $\operatorname{ran} A$

Here is a first consequence of the factorization $A = A(:,\mathtt{b})R$ (with R satisfying (3.8)(i–ii)): The factorization implies that $\operatorname{ran} A \subset \operatorname{ran} A(:,\mathtt{b})$, while certainly $\operatorname{ran} A(:,\mathtt{b}) \subset \operatorname{ran} A$. Hence

$$\operatorname{ran} A = \operatorname{ran} A(:,\mathtt{b}).$$

Also, $A(:,\mathtt{b})$ is 1-1: For, if $A(:,\mathtt{b})a = 0$, then the n-vector x with $x_{\mathtt{b}} = a$ and with $x_{\mathtt{f}} = 0$ is in $\operatorname{null} A = \operatorname{null} R$, hence $a = x_{\mathtt{b}} = -R(:,\mathtt{f})x_{\mathtt{f}} = -R(:,\mathtt{f})0 = 0$. Consequently, $A(:,\mathtt{b})$ is (or, the columns of $A(:,\mathtt{b})$ form) a 'basis' for $\operatorname{ran} A$.

3.10 For each of the matrices A in H.P. 3.1 , give a 'basis' for $\operatorname{ran} A$.

3.11 Let A be the $n \times n$ matrix $[0, e_1, \ldots, e_{n-1}]$ (with e_j denoting the jth unit vector, of the appropriate length). (a) What is its rref? (b) In the equation $A? = 0$, which unknowns are bound, which are free? (c) Give a 'basis' for $\operatorname{null} A$ and a 'basis' for $\operatorname{ran} A$.

3.12 Let M be the 6×3-matrix $[e_3, e_2, e_1]$. (a) What is its rref? (b) Use (a) to prove that M is 1-1. (c) Construct a left inverse for M. (d) (off the wall:) Give a matrix P for which $\operatorname{null} P = \operatorname{ran} M$.

3.13 Let $N := M^{\mathrm{t}}$, with M the matrix in the previous problem. (a) What is its rref? (b) Use (a) to prove that N is onto. (c) Construct a right inverse for N.

3.14 Use the rref to prove that $\operatorname{ran} U = \operatorname{ran} V$, with

$$U := \begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \\ -1 & 1 & 3 \end{bmatrix}, \qquad V := \begin{bmatrix} 1 & 2 \\ 2 & 4 \\ -4 & -5 \end{bmatrix}.$$

(Hints: Proving two sets to be equal usually involves showing that each is a subset of the other. In this case, applying elimination to $[V, U]$ as well as to $[U, V]$ should provide all the information you need.)

Uniqueness of the rrref(A)

If R is a $\mathtt{b}$-form for A, then, as we just proved, $A = A(:,\mathtt{b})R$ and $A(:,\mathtt{b})$ is 1-1. Hence, if also S is a $\mathtt{b}$-form for A, then we have $A(:,\mathtt{b})R = A = A(:,\mathtt{b})S$ and, since $A(:,\mathtt{b})$ is 1-1, this implies that $R = S$. In other words, the matrix R is uniquely determined by the condition that $A(:,\mathtt{b})R = A$. In particular, rrref(A) is uniquely determined, since we already observed that, by (3.4), the sequence $\mathtt{bound}$ only depends on null A.

Further, since rref(A) differs from rrref(A) only by those additional $m - \#\mathtt{bound}$ zero rows, it follows that each A also has a *unique* rref.

This finishes the proof of the following summarizing theorem.

(3.12) Theorem: For given $A \in \mathbb{F}^{m\times n}$, there is exactly one matrix R having the properties 1. and 2. (listed in (3.7)) of a rrref for A. Further, with $\mathtt{bound}$ and $\mathtt{free}$ the indices of bound and free unknowns, $A(:,\mathtt{bound})$ is 1-1 onto ran A, and $C \in \mathbb{F}^{n\times\#\mathtt{free}}$, given by $C(\mathtt{free},:) = \mathrm{id}$, $C(\mathtt{bound},:) = -R(:,\mathtt{free})$, is 1-1 onto null A, and C is 'upper triangular' in the sense that $C(i,j) = 0$ for $i > \mathtt{free}(j)$.

The rrref(A) **and the solving of** $A? = y$

(3.5)Corollary(i) is exactly what we need when considering the linear system

$$A? = y \tag{3.13}$$

for given $A \in \mathbb{F}^{m\times n}$ and given $y \in \mathbb{F}^m$. For, here we are hoping to write y as a linear combination of the columns of A, and (3.5) tells us that this is possible exactly when the last unknown in the *homogeneous* system

$$[A, y]? = 0 \tag{3.14}$$

is free. Further, the factorization (3.11), applied to the **augmented** matrix $[A, y]$, tells us how to write y as a linear combination of the columns of A in case that can be done. For, with $R = \mathrm{rrref}([A, y])$, it tells us that

$$y = [A, y](:,\mathtt{bound})R(:, n+1),$$

and this gives us y in terms of the columns of A precisely when $n + 1 \notin \mathtt{bound}$, i.e., when the $(n + 1)$st unknown is free.

(3.15) Proposition: For $A \in \mathbb{F}^{m\times n}$ and $y \in \mathbb{F}^m$, the equation

$$A? = y$$

has a solution if and only if the last column of $[A, y]$ is free, in which case the last column of rrref$([A, y])$ provides the unique solution to

$$A(:,\mathtt{bound})? = y.$$

More generally, if $R = \mathrm{rrref}([A, B])$ for some arbitrary matrix $B \in \mathbb{F}^{m\times s}$ and all the unknowns corresponding to columns of B are free, then, by (3.11), applied to $[A, B]$ rather than A, we have

$$B = A(:,\mathtt{bound})R(:, n + (1{:}s)).$$

3.15 Prove that rrref$(\mathrm{id}_n) = \mathrm{id}_n$.

A numerical example, continued: Recall our earlier example in which we used elimination to convert a given matrix to its rrref, as follows:

$$\begin{bmatrix} 0 & 2 & 0 & 2 & 5 & 4 & 0 & 6 \\ 0 & 1 & 0 & 1 & 2 & 2 & 0 & 3 \\ 0 & 2 & 0 & 2 & 5 & 4 & -1 & 7 \\ 0 & 1 & 0 & 1 & 3 & 2 & -1 & 4 \end{bmatrix} \rightarrow \begin{bmatrix} 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 \\ 0 & \mathbf{1} & 0 & 1 & 2 & 2 & 0 & 3 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{-1} & 1 \end{bmatrix} \rightarrow \begin{bmatrix} 0 & \mathbf{1} & 0 & 1 & 0 & 2 & 0 & 3 \\ 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & -1 \end{bmatrix},$$

hence $\texttt{bound} = (2, 5, 7)$, $\texttt{free} = (1, 3, 4, 6, 8)$. Now, the elimination algorithm is entirely unaware of how we got the initial matrix. In particular, we are free to interpret in various ways the array on the left as being of the form $[A, B]$. As soon as we specify the number of columns, in A or B, we know A and B exactly.

First, choose B to be a one-column matrix. Then, since the last unknown is free, we conclude that

$$(6, 3, 7, 4) = A(:, \texttt{bound})R(:, 8) = \begin{bmatrix} 2 & 5 & 0 \\ 1 & 2 & 0 \\ 2 & 5 & -1 \\ 1 & 3 & -1 \end{bmatrix} (3, 0, -1).$$

If we choose B to be a three-column matrix instead, then the linear system $A? = B$ is unsolvable since now one of the columns of B (the second one) corresponds to a bound unknown. What about the other two columns of this B? The first one corresponds to a free unknown, hence is a weighted sum of the columns to the left of it, hence is in ran A. But the last one fails to be in ran A since its unknown is free only because of the presence of the seventh column, and this seventh column is *not* a weighted sum of the columns to the left of it, hence neither is the eighth column. Indeed, the corresponding column of R has its last entry nonzero, showing that $A(:, \texttt{bound}(3))$ is needed to write the last column of A as a weighted sum of columns to the left of it. □

3.16 Use elimination to show that $\begin{bmatrix} 2 & -1 & 0 \\ 1 & 2 & 1 \\ 0 & 2 & -1 \end{bmatrix}$ is 1-1 and onto.

3.17 Use elimination to settle the following assertions, concerning the linear system $A? = y$, with the (square) matrix A and the right side y given by

$$[A, y] := \begin{bmatrix} 1 & -2 & 3 & 1 \\ 2 & k & 6 & 6 \\ -1 & 3 & k-3 & 0 \end{bmatrix}.$$

(a) If $k = 0$, then the system has an infinite number of solutions. (b) For another specific value of k, which you must find, the system has no solutions. (c) For all other values of k, the system has a unique solution.

(To be sure, there probably is some preliminary work to do, after which it is straightforward to answer all three questions.)

3.18 Here are three questions that can be settled **without doing any arithmetic**. Please do so.

(i) Can both of the following equalities be right?

$$\begin{bmatrix} -5 & 2 \\ 3 & -1 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 3 & 5 \end{bmatrix} = \mathrm{id}_2 = \begin{bmatrix} 1 & 2 \\ 3 & 5 \end{bmatrix} \begin{bmatrix} -4 & 2 \\ 3 & 5 \end{bmatrix}$$

(ii) How does one find the coordinates of $e_1 \in \mathbb{R}^2$ with respect to the vector sequence $(1, 3), (2, 5)$ (i.e., numbers α, β for which $e_1 = (1, 3)\alpha + (2, 5)\beta$), given that

$$AV := \begin{bmatrix} -5 & 2 \\ 3 & -1 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 3 & 5 \end{bmatrix} = \mathrm{id}_2\,?$$

(iii) How does one conclude at a glance that the following equation must be wrong?

$$\begin{bmatrix} -5 & 2 \\ 3 & -1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 & 1 \\ 3 & 5 & 0 \end{bmatrix} = \mathrm{id}_3\,?$$

The pigeonhole principle for square matrices

We are ready for a discussion of our basic problem, namely solving $A? = y$, in case $A \in \mathbb{F}^{m\times n}$, hence $y \in \mathbb{F}^m$. When is A 1-1, onto, invertible? We answer all these questions by applying elimination to the augmented matrix $[A, y]$.

If A is 1-1, then, by (3.5)Corollary, all its columns must be *bound*. In particular, there must be enough rows to bind them, i.e., $m \ge n$. Further, if $m = n$, then, by the time we reach the last column of $[A, y]$, there is no row left to bind it. Therefore, the last column must be free regardless of the choice of y, hence, by (3.5)Corollary, $y \in \operatorname{ran} A$ for every $y \in \mathbb{F}^m = \operatorname{tar} A$, i.e., A is onto.

If A is onto, then, for $i = 1{:}m$, there is $b_i \in \mathbb{F}^n$ so that $Ab_i = e_i \in \mathbb{F}^m$. Hence, with $B := [b_1, \ldots, b_m] \in \mathbb{F}^{n\times m}$, we have $AB = A[b_1, \ldots, b_m] = [Ab_1, \ldots, Ab_m] = [e_1, \ldots, e_m] = \mathrm{id}_m$. It follows that B is 1-1, hence B has at least as many rows as columns, i.e., $n \ge m$, and A is a left inverse for B. Further, if $n = m$, then, by the previous paragraph, B is also onto, hence invertible, hence any left inverse must be its inverse. In particular $A = B^{-1}$ and therefore, in particular, A is 1-1.

Note that the argument just given provides the proof of the 'Amazing Fact' (2.18), since it concludes from $AB = \mathrm{id}$ (with A, B square) that A must be the inverse of B, and this implies, in particular, that also $BA = \mathrm{id}$.

But we have proved much more, namely the following basic Theorem.

(3.16) Theorem (pigeonhole principle for square matrices): A square matrix is 1-1 if and only if it is onto.

In other words, when dealing with a *square* matrix, 1-1 *or* onto is already enough to have 1-1 *and* onto, i.e., to have invertibility.

We also now know that only square matrices are invertible.

(3.17) Proposition: An invertible matrix is necessarily square. More precisely, if $A \in \mathbb{F}^{m\times n}$, then (i) A 1-1 implies that $m \ge n$; and (ii) A onto implies that $m \le n$.

(3.18) Example: Constructing the inverse by elimination If $A \in \mathbb{F}^{n\times n}$ is invertible, then the first n columns of $[A, \mathrm{id}_n]$ are necessarily bound and the remaining n columns are necessarily free. Therefore, if $R := \operatorname{rrref}([A, \mathrm{id}_n])$, then $R = [\,\mathrm{id}_n, ?]$ and, with (3.11), necessarily $[A, \mathrm{id}_n] = AR = [A\,\mathrm{id}_n, A?]$, hence $? = A^{-1}$, i.e., $R = [\,\mathrm{id}_n, A^{-1}]$.

practical note: Although `MATLAB` provides the function `inv(A)` to generate the inverse of `A`, there is usually no reason to compute the inverse of a matrix, nor would you solve the linear system $A? = y$ in practice by computing $\operatorname{rref}([A, y])$ or by computing `inv(A)*y`. Rather, in `MATLAB` you would compute the solution of `A`$? =$`y` as `A\y`. For this, `MATLAB` also uses elimination, but in a more sophisticated form, to keep rounding error effects as small as possible. In effect, the choice of pivot rows is more elaborate than we discussed above.

□

3.19 For each of the following matrices A, use elimination (to the extent necessary) to (a) determine whether it is invertible and, if it is, to (b) construct the inverse.

(a) $\begin{bmatrix}1&2&3\\2&3&4\end{bmatrix}$; (b) $\begin{bmatrix}1&2\\2&3\\3&4\end{bmatrix}$; (c) $\begin{bmatrix}1&2&3\\2&3&4\\3&4&5\end{bmatrix}$; (d) $\begin{bmatrix}1&2&3\\2&3&4\\3&4&4\end{bmatrix}$; (e) $\begin{bmatrix}1&1&1\\1&2&4\\1&3&8\end{bmatrix}$; (f) $[e_1-e_3, e_2, e_3+e_4, e_4] \in \mathbb{R}^{4\times 4}$.

3.20 Prove that A *is invertible iff* $\operatorname{rrref}(A) = \mathrm{id}_n$.

3.21 One way to solve *Laplace's equation*, $\Delta f := D_1^f + D_2^2 f = y$ on some domain G in $\mathbb{R}^2$ with $f = g$ on the boundary, ∂G of G, numerically is to choose a regular grid $T = \{(ih, jh) : i \in I, j \in J\}$ of points, with I and J chosen so that (ih, jh) is either strictly inside G or else is next to one such, and then to try to compute $u \in \mathbb{R}^T$ so that $u(t) = (u(t+(h,0)) + u(t-(h,0)) + u(t+(0,h)) + u(t-(0,h)))/4$ for all t strictly inside G, while, for the other points in T, $u(t)$ is determined from the given boundary values g in a linear manner.

Prove that the resulting linear system $Au = b$ for the 'vector' $u = (u(t) : t \in T)$ has exactly one solution. (Hint: if $u(t) = \max u(T)$ for some t inside G, then, $u(t)$ being the average of its four neighbors, those neighbors must have the same value.)

3.22 Let $L \in \mathbb{R}^{n\times n}$ be the lower triangular matrix with all diagonal entries equal to 1 and all the strictly lower triangular entries equal to -1, and let $n > 1$. Prove that $(L^{-1})_{n1} = 2^{n-2}$.

(3.19) Example: Triangular matrices There is essentially only one class of square matrices whose invertibility can be settled by inspection, namely the class of triangular matrices.

Assume that the square matrix A is upper triangular, meaning that $i > j \implies A(i,j) = 0$. If all its diagonal elements are nonzero, then each of its unknowns has a pivot row, hence is bound and, consequently, A is 1-1, hence, by (3.16)Theorem, it is invertible. Conversely, if some of its diagonal elements are zero, then there must be a first zero diagonal entry, say $A(i,i) = 0 \ne A(k,k)$ for $k < i$. Then, for $k < i$, row k is a pivot row for x_k, hence, when it comes time to decide whether x_i is free or bound, all rows not yet used as pivot rows do not involve x_i explicitly, and so x_i is free. Consequently, null A is nontrivial and A fails to be 1-1.

Exactly the same argument can be made in case A is lower triangular, meaning that $i < j \implies A(i,j) = 0$, provided you are now willing to carry out the elimination process from right to left, i.e., in the order x_n, x_{n-1}, etc., and, correspondingly, recognize a row as pivot row for x_k in case x_k is the last unknown that appears explicitly (i.e., with a nonzero coefficient) in that row.

(3.20) Proposition: A square triangular matrix is invertible if and only if all its diagonal entries are nonzero.

□

(3.21) Example: Interpolation If $V \in L(\mathbb{F}^n, X)$ and $Q \in L(X, \mathbb{F}^n)$, then QV is a linear map from $\mathbb{F}^n$ to $\mathbb{F}^n$, i.e., a square matrix, of order n. If QV is 1-1 or onto, then (3.16)Theorem tells us that QV is invertible. In particular, V is 1-1 and Q is onto, and so, for every $y \in \mathbb{F}^n$, there exists exactly one $p \in \operatorname{ran} V$ for which $Qp = y$. This is the essence of *interpolation*.

For example, take $X = \mathbb{R}^{\mathbb{R}}$, $V = [()^0, ()^1, \ldots, ()^{k-1}]$, hence ran V equals $\Pi_{<k}$, the collection of all polynomials of degree $< k$. Further, take $Q : X \to \mathbb{R}^k : f \mapsto (f(\tau_1), \ldots, f(\tau_k))$ for some fixed sequence $\tau_1 < \cdots < \tau_k$ of points. Then the equation

$$QV? = Qf$$

asks for the (power) coefficients of a polynomial of degree $< k$ that agrees with the function f at the k distinct points τ_j.

We investigate whether QV is 1-1 or onto, hence invertible. For this, consider the matrix QW, with the columns of $W := [w_1, \ldots, w_k]$ the so-called **Newton polynomials**

$$w_j : t \mapsto \prod_{h<j} (t - \tau_h).$$

Observe that $(QW)(i,j) = (Qw_j)(\tau_i) = \prod_{h<j}(\tau_i - \tau_h) = 0$ if and only if $i < j$. Therefore, QW is square and lower triangular with nonzero diagonal entries, hence invertible by (3.20)Proposition, while w_j is a polynomial of exact degree $j-1 < k$, hence $w_j = Vc_j$ for some k-vector c_j. It follows that the invertible matrix QW equals

$$QW = [Qw_1, \ldots, Qw_k] = [QVc_1, \ldots, QVc_k] = (QV)[c_1, \ldots, c_k].$$

In particular, QV is onto, hence invertible, hence also V is 1-1, therefore invertible as a linear map from $\mathbb{R}^k$ to its range, $\Pi_{<k}$. We have proved:

(3.22) Proposition: For every $f : \mathbb{R} \to \mathbb{R}$ and every k distinct points $\tau_1, \ldots, \tau_k$ in $\mathbb{R}$, there is exactly one choice of coefficient vector a for which the polynomial $[()^0, \ldots, ()^{k-1}]a$ of degree $< k$ agrees with f at these τ_j.

In particular, (i) the column map $[()^0, \ldots, ()^{k-1}] : \mathbb{R}^k \to \Pi_{<k}$ is invertible, and (ii) any polynomial of degree $< k$ with more than $k-1$ distinct zeros must be 0. (Do not confuse this simple result with the **Fundamental Theorem of Algebra** which claims that every nonconstant polynomial with complex cofficients has a zero.)

□

3.23 (a) Construct the unique element of $\operatorname{ran}[()^0, ()^2, ()^4]$ that agrees with $()^1$ at the three points 0, 1, 2.

(b) Could (a) have been carried out if the pointset had been -1, 0, 1 (instead of 0, 1, 2)?

3.24 Let $\tau_1 \neq \tau_2$. Prove that, for an arbitrary $a \in \mathbb{R}^4$, there exists exactly one cubic polynomial p for which

$$(p(\tau_1), Dp(\tau_1), p(\tau_2), Dp(\tau_2)) = a.$$

(Hint: Try $W := [()^0, (\cdot - \tau_1), (\cdot - \tau_1)^2, (\cdot - \tau_1)^2(\cdot - \tau_2)]$.)

3.25 T/F

(a) $\begin{bmatrix} 1 & 0 & 1 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ is in row echelon form.

(b) If all unknowns in the linear system $A? = 0$ are free, then $A = 0$;

(c) If all unknowns in the linear system $A? = 0$ are bound, then A is invertible.

(d) If some unknowns in the linear system $A? = 0$ are free, then A cannot be invertible.

(e) The inverse of an upper triangular matrix is lower triangular.

(f) A linear system of n equations in $n+1$ unknowns always has solutions.

(g) Any square matrix in row echelon form is upper triangular.

(h) If A and B are square matrices of the same order, then $AB? = 0$ has the same number of bound unknowns as does $BA? = 0$.

(i) If A and B are square matrices of the same order, and AB is invertible, then also BA is invertible.

(j) If $\operatorname{null} A = \operatorname{null} B$, then $A? = 0$ and $B? = 0$ have the same free and bound unknowns.

4. The dimension of a vector space

Bases

The only vector spaces in which we can carry out calculations are the coordinate spaces $\mathbb{F}^n$. To calculate with other vector spaces, we have to relate them first to some coordinate space. This is true even when X is a proper subspace of $\mathbb{F}^n$, e.g., the nullspace of some matrix.

For example, we do not really compute with polynomials, we usually compute with the coefficients of the polynomial. Precisely (see (3.22)Proposition), one sets up the invertible linear map

$$\mathbb{F}^n \to \Pi_{<n} : a \mapsto a_1 + a_2 t + a_3 t^2 + \cdots + a_n t^{n-1}$$

where I have, temporarily, followed the (ancient and sometimes confusing) custom of describing the *monomials* by the list of symbols $(\ , t, t^2, t^3, \ldots)$ rather than by the nonstandard symbols $()^j,\ j = 0, 1, 2, 3, \ldots$ introduced earlier. One adds polynomials by adding their coefficients, or evaluates polynomials from their coefficients, etc. You may be so used to that, that you haven't even noticed until now that you do not work with the polynomials themselves, but only with their coefficients.

It is therefore a practically important goal to provide ways of **representing** the elements of a given vector space X by n-vectors. We do this by using linear maps from some $\mathbb{F}^n$ that have X as their range, i.e., we look for sequences $v_1, v_2, \ldots, v_n$ in X for which the linear map $[v_1, v_2, \ldots, v_n] : \mathbb{F}^n \to X$ is onto. If there is such a map for some n, then we call X **finitely generated**.

Among such onto maps $V \in L(\mathbb{F}^n, X)$, those that are also 1-1, hence invertible, are surely the most desirable ones since, for such V, there is, for any $x \in X$, exactly one $a \in \mathbb{F}^n$ with $x = Va$. Any *invertible* column map to X is, by definition, a **basis** for X.

Since $\mathrm{id}_n \in L(\mathbb{F}^n)$ is trivially invertible, it is a basis for $\mathbb{F}^n$. It is called the **natural basis for** $\mathbb{F}^n$.

The bound part, $A(:, \texttt{bound})$, of $A \in \mathbb{F}^{m\times n}$ is a basis for $\operatorname{ran} A$. You also know (from pages 36ff) how to construct a basis for the nullspace of any $A \in \mathbb{F}^{m\times n}$ from its $\operatorname{rrref}(A)$.

Here is a small difficulty with this (and any other) definition of dimension: What is the dimension of the **trivial space**, i.e., the vector space that consists of the zero vector alone? It is a perfectly well-behaved vector space (though a bit limited, – except as a challenge to textbook authors when it comes to discussing its basis).

We deal with it here by considering $V \in L(\mathbb{F}^n, X)$ even when $n = 0$. Since $\mathbb{F}^n$ consists of lists of n items (each item an element from $\mathbb{F}$), the peculiar space $\mathbb{F}^0$ must consist of lists of *no* items, i.e., of *empty* lists. There is only one empty list (of scalars), hence $\mathbb{F}^0$ has just one element, the empty list, $(\,)$, and this element is necessarily the neutral element (or, zero vector) for this space. Correspondingly, there is exactly one *linear* map from $\mathbb{F}^0$ into X, namely the map $\mathbb{F}^0 \to X : () = 0 \mapsto 0$. Since this is a linear map from $\mathbb{F}^0$, we call it the column map into X with *no* columns, and denote it by $[\,]$. Thus,

(4.1)
$$[\,] : \mathbb{F}^0 \to X : () = 0 \mapsto 0.$$

Note that $[\,]$ *is 1-1*. Note also that the range of $[\,]$ consists of the trivial subspace, $\{0\}$. In particular, the column map $[\,]$ is *onto* $\{0\}$, hence is invertible, as map from $\mathbb{F}^0$ to $\{0\}$. It follows that $[\,]$ is a basis for $\{0\}$. Isn't Mathematics wonderful! - As it turns out, the column map $[\,]$ will also be very helpful below.

Here are some standard terms related to bases of a vector space:

Definition: The range of $V := [v_1, v_2, \ldots, v_n]$ is called the **span of the sequence** $v_1, v_2, \ldots, v_n$:

$$\operatorname{span}(v_1, v_2, \ldots, v_n) := \operatorname{ran} V.$$

$x \in X$ is said to be **linearly dependent on** $v_1, v_2, \ldots, v_n$ in case $x \in \operatorname{ran} V$, i.e., in case x is a **linear combination of the** v_j. Otherwise x is said to be **linearly independent of** $v_1, v_2, \ldots, v_n$.

$v_1, v_2, \ldots, v_n$ is said to be **linearly independent** in case V is 1-1, i.e., in case $Va = 0$ implies $a = 0$ (i.e., the only way to write the zero vector as a linear combination of the v_j is to choose all the weights equal to 0).

$v_1, v_2, \ldots, v_n$ is said to be **spanning for** X in case V is onto, i.e., in case $\operatorname{span}(v_1, v_2, \ldots, v_n) = X$.

$v_1, v_2, \ldots, v_n$ is said to be a **basis for** X in case V is invertible, i.e., 1-1 and onto.

If V is invertible, then $V^{-1}x$ is an n-vector, called the **coordinate vector for** x **with respect to the basis** $v_1, v_2, \ldots, v_n$.

You may wonder why there are all these terms in use for the *sequence* $v_1, v_2, \ldots, v_n$, particularly when the corresponding terms for the *map* V are so much shorter and to the point. I don't know the answer. However, bear in mind that the terms commonly used are those for sequences. An even greater puzzle is the fact that many textbooks present bases as *sets* rather than *sequences*. At least, that is what they say. But, not surprisingly, whenever there is some action involving a basis, the basis is written $\{v_1, \ldots, v_n\}$, i.e., as a sequence in everything but in name. It is for you to ask such authors whether $\{3, 3\}$ is a basis for $\mathbb{R}^1 = \mathbb{R}$. They will say that it is not even though it is since, after all, $3 = 3$, hence $\{3, 3\} = \{3\}$.

A major use of the basis concept is the following which generalizes the way we earlier constructed arbitrary linear maps from $\mathbb{F}^n$.

(4.2) Proposition: Let $V = [v_1, v_2, \ldots, v_n]$ be a basis for the vector space X, and let Y be an arbitrary vector space. Any map $f : \{v_1, \ldots, v_n\} \to Y$ has exactly one extension to a linear map A from X to Y. In other words, we can choose the values of a linear map on the columns of a basis arbitrarily and, once chosen, this pins down the linear map everywhere.

Proof: The map $A := [f(v_1), \ldots, f(v_n)]V^{-1}$ is linear, from X to Y, and carries v_j to $f(v_j)$ since $V^{-1}v_j = e_j$, all j. This shows existence. Further, if also $B \in L(X, Y)$ with $Bv_j = f(v_j)$, all j, then $BV = [f(v_1), \ldots, f(v_n)] = AV$, therefore $B = A$ (since V is invertible). □

4.1 Describe what the $n \times n$-matrix $A = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \cdot & \cdot & \cdot & \cdots & \cdot & \cdot \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & 0 & \cdots & 0 & 0 \end{bmatrix}$ does to all the vectors e_j, i.e., give a simple formula for Ae_j. Deduce from your formula that $\operatorname{ran} A^n = \{0\}$, hence that $A^n = 0$.

4.2 Prove: $A \in L(X)$ commutes with every $B \in L(X)$ if and only if $A = \alpha\, \mathrm{id}_X$, i.e., a scalar multiple of the identity.

4.3 Let $X \times Y$ be the product space of the vector spaces X and Y. The map $f : X \times Y \to \mathbb{F}$ is **bilinear** if it is linear in each slot, i.e., if $f(\cdot, y) \in L(X, \mathbb{F})$ for all $y \in Y$, and $f(x, \cdot) \in L(Y, \mathbb{F})$ for every $x \in X$.

(i) Prove that, for every $A \in \mathbb{F}^{m \times n}$, the map $f_A : \mathbb{F}^m \times \mathbb{F}^n : (x, y) \mapsto y^t Ax$ is bilinear.

(ii) Prove that, for every bilinear $f : \mathbb{F}^m \times \mathbb{F}^n \to \mathbb{F}$, there exists exactly one $A \in \mathbb{F}^{m \times n}$ with $f_A = f$.

(iii) Prove that the map $A \mapsto f_A$ is an invertible linear map on $\mathbb{F}^{m \times n}$ to the vector space $BL(\mathbb{F}^m, \mathbb{F}^n)$ of all bilinear maps on $\mathbb{F}^m \times \mathbb{F}^n$ under pointwise vector operations.

4.4 `MATLAB`'s command `yy = interp1(x,y,xx,'spline')` returns the value(s) at `xx` of a certain function f that matches the data given by `x`, `y`, in the sense that f(`x(i)`) = `y(i)` for `i=1:n`, with n the length of both `x` and `y` (and assuming that the entries of `x` are pairwise distinct). (If you wanted to look at f on the interval $[a \,..\, b]$, you might choose `xx = linspace(a,b,N+1);` with `N` some suitably large number, and then `plot(xx,yy)`.)

(a) Generate some numerical evidence for the claim that (up to roundoff) the map `y`$\mapsto f$ provided by this command is linear.

(b) Assuming that the map is linear, deduce from the above description of the map that it must be 1-1, hence a basis for its range.

(c) Still assuming that the map `y` $\mapsto f$ provided by that command is indeed linear, hence a column map, provide a plot of each of its columns, as functions on the interval $[0 \,..\, 3]$, for the specific choice `0:3` for `x`.

(d) (quite open-ended) Determine as much as you can about the elements of the range of this column map.

(e) Is the map still linear if you replace `'spline'` by `'cubic'`?

Construction of a basis

Next, we consider the construction of a basis. This can be done either by *extending a 1-1 column map* V to a basis, or by *thinning an onto column map* W to a basis. For this, remember that, for two column maps V and W into some vector space X, we agreed to mean by $V \subset W$ that V can be obtained from W by thinning, i.e., by omitting zero or more columns from W, and W can be obtained from V by extending, i.e., by inserting zero or more columns.

In the discussion to follow, it is convenient to classify the columns of a column map as *bound* or *free*, using (3.5)Corollary as a guide. Specifically, we call a column **free** if it is a weighted sum of the columns to its left; otherwise, we call it **bound**.

For example, if $V \subset W$, then any free column of V is also free as a column of W, while a bound column of V may possibly be free as a column of W unless $W = [V, U]$.

(4.3) Lemma: The kth column of the column map V is free if and only if $\operatorname{null} V$ contains a vector whose last nonzero entry is its kth.

Proof: The kth column of $V = [v_1, \ldots, v_n] \in L(\mathbb{F}^n, X)$ is free iff $v_k \in \operatorname{ran}[v_1, \ldots, v_{k-1}]$. In particular, the first column is free iff it is 0 (recall that $\operatorname{ran}[\,] = \{0\}$).

If the kth column is free, then $v_k = [v_1, \ldots, v_{k-1}]a$ for some $a \in \mathbb{F}^{k-1}$, hence $(a, -1, 0, \ldots, 0) \in \mathbb{F}^n$ is a vector in $\operatorname{null} V$ whose last nonzero entry is its kth. Conversely if $x \in \operatorname{null} V$ with $x_k \neq 0 = x_{k+1} = \cdots = x_n$, then $[v_1, \ldots, v_{k-1}]x_{1:k-1} + v_k x_k = 0$, therefore, as $x_k \neq 0$, $v_k = [v_1, \ldots, v_{k-1}](x_{1:k-1}/(-x_k))$ showing that the kth column is free. □

(4.4) Corollary: A column map is 1-1 if and only if all of its columns are bound.

We are ready for the following algorithm which extracts from any column map W a basis for its range.

(4.5) Basis Selection Algorithm:
input: the column map W
$V \leftarrow [\,]$;
for $w \in W$ **do:**
  if $w \notin \operatorname{ran} V$**, then** $V \leftarrow [V, w]$**; endif**
enddo
output: the column map V

Proposition: The output of the Basis Selection Algorithm is a basis for the range of its input.

Proof: The resulting V has the same range as W (since the only columns of W not explicitly columns of V are those that are already in the range of V). In addition, by construction, every column of V is bound, hence V is 1-1 by (4.4)Corollary, therefore a basis for its range. □

(4.6) Proposition: Any onto column map can be thinned to a basis.

Now note that the Basis Selection Algorithm will put any bound column of W into the resulting basis, V. In particular, if $W = [U, Z]$ with U 1-1, then, as already remarked just prior to (4.3)Lemma, all columns of U will be bound also as columns of W, hence will end up in the resulting basis. This proves

(4.7) Proposition: Any 1-1 column map into a finitely generated vector space can be extended to a basis for that space.

If V is a 1-1 column map into X then, by (4.4)Corollary, all its columns are bound. Hence if V is **maximally 1-1** into X, meaning that $[V, w]$ fails to be 1-1 for every $w \in X$, then that additional column must be free, i.e., $w \in \operatorname{ran} V$ for all $w \in X$, showing that then V is also onto, hence a basis. This proves

(4.8) Corollary: Any maximally 1-1 column map into a vector space is a basis for that space.

If W is a column map onto X, then, by (4.6), it can always be thinned to a basis. Hence, if W is **minimally onto**, meaning that no $V \subset W$ (other than W) is onto, then W itself must be that basis.

(4.9) Corollary: Any minimally onto column map into a vector space is a basis for that space.

4.5 How would you carry out the (4.5) Basis Selection Algorithm for the special case that W is a matrix? (Hint: (3.2)).

4.6 Try out your answer to the previous problem on the specific matrix $W = \begin{bmatrix} 0 & 2 & 0 & 2 & 5 & 4 & 0 & 6 \\ 0 & 1 & 0 & 1 & 2 & 2 & 0 & 3 \\ 0 & 2 & 0 & 2 & 5 & 4 & -1 & 7 \end{bmatrix}$.

Dimension

(4.10) Lemma: Any two bases for a vector space have the same number of columns.

This number of columns in any basis for X is denoted

$$\dim X$$

and is called the **dimension of** X.

Proof: Let $V \in L(\mathbb{F}^n, X)$ and $W \in L(\mathbb{F}^m, X)$ be bases for X. Then, $W^{-1}V$ is an invertible linear map from $\mathbb{F}^n$ to $\mathbb{F}^m$, hence an invertible matrix and therefore, by (3.17)Proposition(i), necessarily a square matrix, i.e., $n = m$. □

See H.P. 4.10 for the classical proof of this lemma.

Notice that we have actually proved the stronger statement

(4.11) Lemma: If V and W are column maps into X, and V is 1-1 and W is onto, then $\#V \le \#W$.

Again, also this stronger result is an immediate consequence of something proved in the previous chapter: Since W is onto, each column v_j of V can be written as $v_j = Wc_j$ for some vector c_j. Hence $V = WC$ for some matrix C and, since V is 1-1, so must C be. By (3.17)Proposition(i) or its antecedent, (3.6)Theorem, this implies that C cannot have more columns than rows, i.e., $\#V = \#C \le \dim \operatorname{tar} C = \dim \operatorname{dom} W = \#W$.

Since id_n is a basis for $\mathbb{F}^n$ and has n columns, we conclude that the *n-dimensional coordinate space has, indeed, dimension n*. In effect, $\mathbb{F}^n$ is the prototypical vector space of dimension n. Any n-dimensional vector space X is connected to $\mathbb{F}^n$ by invertible linear maps, the bases for X.

Note that the trivial vector space, $\{0\}$, has dimension 0 since its (unique) basis has no columns.

(4.12) Example: The dimension of $\Pi_k(\mathbb{R}^d)$. The space $\Pi_k(\mathbb{R}^d)$ of **d-variate polynomials of degree** $\le k$ is, by definition, the range of the column map $V := [()^\alpha : |\alpha| \le k]$, with

$$()^\alpha : \mathbb{R}^d \to \mathbb{R} : t \mapsto t^\alpha := t_1^{\alpha_1} \cdots t_d^{\alpha_d}$$

a nonstandard notation for the α-power function, with $\alpha \in \mathbb{Z}_+^d$, i.e., α any d-vector with nonnegative integer entries, and with $|\alpha| := \sum_j \alpha_j$. For $d = 1$, it is the space of univariate polynomials of degree $\le k$, and we showed its dimension to be $k+1$ by showing that, in that case, V is 1-1.

When $d = 1$, then V can be seen to be 1-1 also by considering the 'data map'

$$Q : \Pi_k \to \mathbb{R}^{k+1} : p \mapsto (p(0), Dp(0), D^2p(0)/2, \ldots, D^kp(0)/k!),$$

for which we have $QV = \mathrm{id}$, hence V is 1-1.

An analogous argument, involving the 'data map'

$$p \mapsto (D^\alpha p(0)/\alpha! : \alpha \in \mathbb{Z}_+^d, |\alpha| \le k),$$

with $\alpha! := \alpha_1! \cdots \alpha_d!$, shows that

$$\dim \Pi_k(\mathbb{R}^d) = \#\{\alpha \in \mathbb{Z}_+^d : |\alpha| \le k\},$$

and the latter number can be shown to equal $\binom{k+d}{d}$. □

4.7 Prove that the space $\Pi_2(\mathbb{R}^2)$ of bivariate polynomials of total degree ≤ 2 has dimension 6.

4.8 Prove that a vector space of dimension n has subspaces of dimension j for each $j = 0{:}n$.

4.9 Prove (by induction on n) **Steinitz Exchange**: *If $V \in L(\mathbb{F}^n, X)$ is 1-1 and $W \in L(\mathbb{F}^m, X)$ is onto, then, for some $U \subset W$ with $\#U = \#W - \#V$, also $[V, U]$ is onto.*

4.10 Use the previous homework to prove (4.11)Lemma.

Some uses of the dimension concept

Here is a major use of the dimension concept as it relates to *vector spaces*.

(4.13) Proposition: If X, Y are vector spaces with $X \subset Y$ and $\dim Y < \infty$, then $\dim X \le \dim Y$, with equality iff $X = Y$.

Proof: Since there is *some* 1-1 column map into X (e.g., the unique linear map from $\mathbb{F}^0$ into X), while $\dim Y$ is an upper bound on the number of columns in any 1-1 column map into $X \subset Y$ (by (4.7)Proposition), there exists a maximally 1-1 column map V into X. By (4.8)Corollary, any such V is necessarily a basis for X, hence X is finitely generated. By (4.7)Proposition, we can extend V to a basis $[V, W]$ for Y. Hence, $\dim X \le \dim Y$ with equality iff $W = [\,]$, i.e., iff $X = Y$. □

Note the following important (nontrivial) part of (4.13)Proposition:

(4.14) Corollary: Any linear subspace of a finite-dimensional vector space is finite-dimensional.

The dimension concept is usually applied to *linear maps* by way of the following formula.

(4.15) Dimension Formula: For any linear map A with finite-dimensional domain,

$$\dim \operatorname{dom} A = \dim \operatorname{ran} A + \dim \operatorname{null} A.$$

Proof: Since $\operatorname{dom} A$ is finite-dimensional, so is $\operatorname{null} A$ (by (4.14)Corollary), hence $\operatorname{null} A$ has a basis, $V \in L(\mathbb{F}^n, \operatorname{null} A)$ say. By (4.7)Proposition, we can extend this to a basis $[V, U]$ for $\operatorname{dom} A$. Let $r := \#U$. Then, $[V, U]$ is invertible and $\dim \operatorname{dom} A - \dim \operatorname{null} A = (n + r) - n = r$.

It remains to prove that $\dim \operatorname{ran} A = r$. For this, we prove that $AU : \mathbb{F}^r \to \operatorname{ran} A$ is invertible.

Since $A[V, U] = [AV, AU]$ maps onto $\operatorname{ran} A$ and $AV = 0$, already AU must map onto $\operatorname{ran} A$, i.e., AU is onto.

Moreover, AU is 1-1: For, if $AUa = 0$, then $Ua \in \operatorname{null} A$, hence, since V maps onto $\operatorname{null} A$, there is some b so that $Ua = Vb$. This implies that $[V, U](b, -a) = 0$ and, since $[V, U]$ is 1-1, this shows that, in particular, $a = 0$. □

4.11 Prove: *If the product AB of the two linear maps A and B is defined, then* $\dim \operatorname{ran}(AB) \le \min\{\dim \operatorname{ran} A, \dim \operatorname{ran} B\}$.

4.12 Prove: *If the product AB of the two linear maps A and B is defined, then* $\dim \operatorname{ran}(AB) = \dim \operatorname{ran} B - \dim(\operatorname{null} A \cap \operatorname{ran} B)$.

4.13 Give an example, of two square matrices A and B, that show that $\dim \operatorname{ran}(AB)$ need not equal $\dim \operatorname{ran}(BA)$ when both AB and BA are defined.

(4.16) Corollary: Let $A \in L(X, Y)$.
(i) If $\dim X < \dim Y$, then A cannot be onto.
(ii) If $\dim X > \dim Y$, then A cannot be 1-1.
(iii) If $\dim X = \dim Y < \infty$, then A is onto if and only if A is 1-1. (This implies (2.18)!)

Proof: (i) $\dim \operatorname{ran} A \le \dim \operatorname{dom} A = \dim X < \dim Y = \dim \operatorname{tar} A$, hence $\operatorname{ran} A \ne \operatorname{tar} A$.

(ii) $\dim \operatorname{null} A = \dim \operatorname{dom} A - \dim \operatorname{ran} A = \dim X - \dim \operatorname{ran} A \ge \dim X - \dim Y > 0$, hence $\operatorname{null} A \ne \{0\}$.

(iii) If $\dim X = \dim Y$, then $\dim \operatorname{tar} A = \dim \operatorname{dom} A = \dim \operatorname{ran} A + \dim \operatorname{null} A$, hence A is onto (i.e., $\operatorname{tar} A = \operatorname{ran} A$) if and only if $\dim \operatorname{null} A = 0$, i.e., A is 1-1. □

(4.17) Lemma: Let X, Y be vector spaces, and assume that X is finite-dimensional. Then $\dim X = \dim Y$ if and only if there exists an invertible $A \in L(X, Y)$.

Proof: Let $n := \dim X$. Since $n < \infty$, there exists an invertible $V \in L(\mathbb{F}^n, X)$ (, a basis for X). If now $A \in L(X, Y)$ is invertible, then AV is an invertible linear map from $\mathbb{F}^n$ to Y, hence $\dim Y = n = \dim X$. Conversely, if $\dim Y = \dim X$, then there exists an invertible $W \in L(\mathbb{F}^n, Y)$; but then WV^{-1} is an invertible linear map from X to Y. □

For the next general result concerning the dimension concept, recall that both the sum

$$Y + Z := \{y + z : y \in Y, z \in Z\}$$

and the intersection $Y \cap Z$ of two linear subspaces is again a linear subspace.

(4.18) Proposition: If Y and Z are linear subspaces of the finite-dimensional vector space X, then

$$\dim(Y + Z) = \dim Y + \dim Z - \dim(Y \cap Z). \tag{4.19}$$

Proof 1: $Y \cap Z$ is a linear subspace of X, hence is finite-dimensional (by (4.14)Corollary), hence $Y \cap Z$ has a basis, V say. Extend it, as we may (by (4.7)Proposition), to a basis $[U, V]$ of Y and to a basis $[V, W]$ of Z, and consider the column map $[U, V, W]$.

We claim that $[U, V, W]$ is 1-1. Indeed, if $[U, V, W](a, b, c) = 0$, then $[U, V](a, b) = -Wc$, with the left side in Y and the right side in Z, hence both are in $Y \cap Z = \operatorname{ran} V$. Therefore, $-Wc = Vd$ for some d, hence $[V, W](d, c) = 0$, and as $[V, W]$ is 1-1, it follows, in particular, that $c = 0$. This leaves $[U, V](a, b) = 0$ and, since $[U, V]$ is 1-1 by construction, now also $(a, b) = 0$.

We conclude that $[U, V, W]$ is a basis for its range, and that range is $\operatorname{ran}[U, V, W] = \operatorname{ran}[U, V, V, W] = \operatorname{ran}[U, V] + \operatorname{ran}[V, W] = Y + Z$. Therefore, $\dim(Y + Z) = \#U + \#V + \#W = \#[U, V] + \#[V, W] - \#V = \dim Y + \dim Z - \dim(Y \cap Z)$. □

Proof 2: The following alternative proof shows (4.19) to be a special case of the (4.15)Dimension Formula, and provides a way to construct a basis for $Y \cap Z$ from bases for Y and Z.

Consider the column map $A := [U, W]$ with U a basis for Y and W a basis for Z. Since $\dim \operatorname{dom} A = \#U + \#W = \dim Y + \dim Z$ and $\operatorname{ran} A = Y + Z$, the formula (4.19) follows from the (4.15)Dimension Formula, once we show that $\dim \operatorname{null} A = \dim Y \cap Z$. For this, let $x \in Y \cap Z$. Then $x = Ua = Wb$ for some a and b, therefore $A(a, -b) = [U, W](a, -b) = Ua - Wb = x - x = 0$, hence $(a, -b) \in \operatorname{null} A$. Hence, $(a, -b) = Cc$ for some c and with $C =: [C_U; C_W]$ a basis for $\operatorname{null} A$. In particular, $a = C_U c$, showing that the column map UC_U has all of $Y \cap Z$ in its range. On the other hand, $0 = AC = UC_U + WC_W$, hence $UC_U = -WC_W$ and, in particular, UC_U maps into $Y \cap Z$, hence onto $Y \cap Z$. Finally, UC_U is 1-1: for, if $UC_U a = 0$, then $C_U a = 0$ since U is 1-1, but then also $WC_W a = -UC_U a = 0$, hence also $C_W a = 0$, therefore $Ca = 0$ and so $a = 0$ since C is 1-1 by assumption. Altogether, this shows that UC_U is a basis for $Y \cap Z$, hence $\dim Y \cap Z = \#UC_U = \#C = \dim \operatorname{null} A$. □

Here are three of several corollaries of this basic proposition to be used in the sequel.

(4.20) Corollary: If $[V, W]$ is 1-1, then $\operatorname{ran} V \cap \operatorname{ran} W$ is trivial.

(4.21) Corollary: If $\dim Y + \dim Z > \dim X$ for some linear subspaces Y and Z of the finite-dimensional vector space X, then $Y \cap Z$ is a *nontrivial* linear subspace, i.e., $Y \cap Z$ contains nonzero elements.

(4.22) Corollary: If Y and Z are linear subspaces of the finite-dimensional vector space X, and $Y \cap Z = \{0\}$, then

$$\dim Y + \dim Z \le \dim X,$$

with equality if and only if $X = Y + Z$, in which case $\dim Z = \dim X - \dim Y =: \operatorname{codim} Y$ is called the **codimension** of Y (in X).

4.14 Prove: *if AB is defined, then* $\dim \operatorname{ran}(AB) \le \min\{\dim \operatorname{ran} A, \dim \operatorname{ran} B\}$. (Hint: $\operatorname{ran}(AB) = A(\operatorname{ran} B)$.)

4.15 Make use of the dimension concept to shorten the solution of H.P. 3.14 .

4.16 For each of the following linear maps, determine its range and its nullspace. Make as much use of the Dimension Formula as possible. (You may, if need be, use the fact that, by (3.22)Proposition, $V_k := [()^0, ()^1, \ldots, ()^k]$ is a basis for Π_k.) (a) $D : \Pi_k \to \Pi_{k-1} : p \mapsto Dp$, with Dp the first derivative of p. (b) $I : \Pi_{k-1} \to \Pi_k : p \mapsto \int_0^{\cdot} p(s)ds$, i.e., Ip is the primitive or antiderivative of p that vanishes at 0, i.e., $(Ip)(t) = \int_0^t p(s)ds$. (c) $A : \Pi_k \to \Pi_k : p \mapsto Dp + p$.

4.17 Prove that $V := [()^0, ()^1, ()^2 - 1, 4()^3 - 3()^1, 8()^4 - 8()^2 + 1]$ is a basis for Π_4.

4.18 Prove: *For any finite-dimensional linear subspace Y of the domain of a linear map A*, $\dim A(Y) \le \dim Y$.

4.19 Prove: *If V and W are 1-1 column maps into the vector space X, then* $\operatorname{ran} V$ *and* $\operatorname{ran} W$ *have a nontrivial intersection if and only if $[V, W]$ is not 1-1.*

4.20 Use the preceding homework and elimination to determine for each of the matrices given whether $\operatorname{ran} A$ and $\operatorname{null} A$ have nontrivial intersection: (a) $A := \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix}$; (b) $A := \begin{bmatrix} -2 & -1 \\ 4 & 2 \end{bmatrix}$.

4.21 Call $(Y_0, \ldots, Y_r)$ a **proper chain** in the vector space X if each Y_j is a subspace and $Y_0 \subsetneq Y_1 \subsetneq \cdots \subsetneq Y_r$. Prove that, for any such proper chain, $r \le \dim X$, with equality if and only if $\dim Y_j = j$, $j = 0{:}\dim X$.

4.22 Let d be any scalar-valued map, defined on the collection of all linear subspaces of a finite-dimensional vector space X, that satisfies the following two conditions: (i) $Y \cap Z = \{0\} \implies d(Y + Z) = d(Y) + d(Z)$; (ii) $\dim Y = 1 \implies d(Y) = 1$.

Prove that $d(Y) = \dim Y$ for every linear subspace Y of X.

4.23 Prove: for any $A \in L(X, Y)$ and any linear subspace Z of X, $\dim A(Z) = \dim Z - \dim(Z \cap (\operatorname{null} A))$.

4.24 The **defect** of a linear map is the dimension of its nullspace: $\operatorname{defect}(A) := \dim \operatorname{null} A$. (a) Prove that $\operatorname{defect}(B) \le \operatorname{defect}(AB) \le \operatorname{defect}(A) + \operatorname{defect}(B)$. (b) Prove: If $\dim \operatorname{dom} B = \dim \operatorname{dom} A$, then also $\operatorname{defect}(A) \le \operatorname{defect}(AB)$. (c) Give an example of linear maps A and B for which AB is defined and for which $\operatorname{defect}(A) > \operatorname{defect}(AB)$.

4.25 Let $A \in L(X, Y)$, $B \in L(X, Z)$, with Y finite-dimensional. There exists $C \in L(Y, Z)$ with $A = CB$ if and only if $\operatorname{null} B \subset \operatorname{null} A$.

4.26 Prove: Assuming that the product ABC of three linear maps is defined, $\dim \operatorname{ran}(AB) + \dim \operatorname{ran}(BC) \le \dim \operatorname{ran} B + \dim \operatorname{ran}(ABC)$.

4.27 Factor space: Let Y be a linear subspace of the vector space X and consider the collection

$$X/Y := \{x + Y : x \in X\}$$

of subsets of X, with

$$x + Y := \{x\} + Y = \{x + y : y \in Y\}.$$

(i) Prove that the map

$$f : X \to X/Y : x \mapsto x + Y$$

is linear with respect to the addition

$$M + N := \{m + n : m \in M, n \in N\}$$

and the multiplication by a scalar

$$\alpha M := \begin{cases} \{\alpha m : m \in M\}, & \text{if } \alpha \ne 0; \\ Y, & \text{if } \alpha = 0, \end{cases}$$

and has Y as its nullspace.

(ii) Prove that, with these vector operations, X/Y is a linear space. (X/Y is called a **factor space**.)

(iii) Prove that $\dim X/Y = \operatorname{codim} Y$.

The dimension of $\mathbb{F}^T$

Recall from (2.2) that $\mathbb{F}^T$ is the set of all scalar-valued maps on the set T, with the set T, offhand, arbitrary.

The best known instance is n-dimensional coordinate space

$$\mathbb{F}^n := \mathbb{F}^{\underline{n}},$$

with $T = \underline{n} := \{1, 2, \ldots, n\}$. The vector space $\mathbb{F}^{m\times n}$ of all $(m \times n)$-matrices is another instance; here $T = \underline{m} \times \underline{n} := \{(i, j) : i = 1{:}m; j = 1{:}n\}$.

(4.23) Proposition: If T is a finite set, then $\dim \mathbb{F}^T = \#T$.

Proof: Since T is finite, $\#T =: n$ say, we can order its elements, i.e., there is an invertible map $s : \underline{n} \to T$ (in fact, there are $n! = 1 \cdot 2 \cdots n$ such). This induces the map

$$V : \mathbb{F}^n \to \mathbb{F}^T : f \mapsto f \circ s^{-1}$$

which is linear (since, in both spaces, the vector operations are pointwise), and is invertible since it has

$$\mathbb{F}^T \to \mathbb{F}^n : g \mapsto g \circ s$$

as its inverse. Hence, V is a basis for $\mathbb{F}^T$ (the **natural basis**). □

Note how we managed this without even exhibiting the columns of V. To be sure, the jth column V is the function $v_j : T \to \mathbb{F} : s_k \mapsto \delta_{kj}$ that maps s_j to 1 and maps any other $t \in T$ to 0.

Corollary: $\dim \mathbb{F}^{m\times n} = mn$.

Proof: In this case, $\mathbb{F}^{m\times n} = \mathbb{F}^T$ with $T = \underline{m} \times \underline{n} := \{(i, j) : i = 1{:}m; j = 1{:}n\}$, hence $\#T = mn$. □

(4.24) Corollary: $\dim L(X, Y) = \dim X \cdot \dim Y$.

Proof: Assuming that $n := \dim X$ and $m := \dim Y$ are finite, we can represent every $A \in L(X, Y)$ as a matrix $\widehat{A} := W^{-1}AV \in \mathbb{F}^{m\times n}$, with V a basis for X and W a basis for Y. This sets up a map

$$R : L(X, Y) \to \mathbb{F}^{m\times n} : A \mapsto \widehat{A} = W^{-1}AV,$$

and this map is linear and invertible (indeed, its inverse is the map $\mathbb{F}^{m\times n} \to L(X, Y) : B \mapsto WBV^{-1}$). Consequently, by (4.17)Lemma, $L(X, Y)$ and $\mathbb{F}^{m\times n}$ have the same dimension. □

Corollary: If $\#T \not< \infty$, then $\mathbb{F}^T$ is not finite-dimensional.

Proof: For every finite $S \subset T$, $\mathbb{F}^T$ contains the linear subspace

$$\{f \in \mathbb{F}^T : f(t) = 0, \text{all } t \notin S\}$$

of dimension equal to $\dim \mathbb{F}^S = \#S$. If $\#T \not< \infty$, then T contains finite subsets S of arbitrarily large size, hence $\mathbb{F}^T$ contains linear subspaces of arbitrarily large dimension, hence cannot itself be finite-dimensional, by (4.13)Proposition. □

4.28 Prove: *The dimension of the vector space of all upper triangular matrices of order n is $(n+1)n/2$.*

Direct sums

A very useful coarsening of the basis concept concerns the sum of subspaces.

Let $Y_1, \ldots, Y_r$ be linear subspaces of the vector space X, let V_j be a column map onto Y_j, all j, and consider the column map

$$V := [V_1, \ldots, V_r].$$

To be sure, we could have also started with some arbitrary column map V into X, arbitrarily grouped its columns to obtain $V = [V_1, \ldots, V_r]$, and then defined $Y_j := \operatorname{ran} V_j$, all j.

Either way, any $a \in \operatorname{dom} V$ is of the form $(a_1, \ldots, a_r)$ with $a_j \in \operatorname{dom} V_j$, all j. Hence

$$\operatorname{ran} V = \{V_1 a_1 + \cdots + V_r a_r : a_j \in \operatorname{dom} V_j, j = 1{:}r\} = \{y_1 + \cdots + y_r : y_j \in Y_j, j = 1{:}r\} =: Y_1 + \cdots + Y_r,$$

the *sum* of the subspaces $Y_1, \ldots, Y_r$.

Think of this sum, as you may, as the range of the map

$$A : Y_1 \times \cdots \times Y_r \to X : (y_1, \ldots, y_r) \mapsto y_1 + \cdots + y_r. \tag{4.25}$$

Having this map A onto says that every $x \in X$ can be written in the form $y_1 + \cdots + y_r$ with $y_j \in Y_j$, all j. In other words, X is the sum of the Y_j. In symbols,

$$X = Y_1 + \cdots + Y_r.$$

Having A also 1-1 says that there is *exactly one way* to write each $x \in X$ as such a sum. In this case, we write

$$X = Y_1 \dotplus \cdots \dotplus Y_r,$$

and say that X is the **direct sum** of the subspaces Y_j. Note the dot atop the plus sign, to indicate the special nature of this sum. Some books would use instead the encircled plus sign, $\oplus$, but we reserve that sign for an even more special direct sum in which the summands Y_j are 'orthogonal' to each other; see the chapter on inner product spaces.

(4.26) Proposition: Let V_j be a basis for the linear subspace Y_j of the vector space X, $j = 1{:}r$, and set $V := [V_1, \ldots, V_r]$. Then, the following are equivalent.

(i) $X = Y_1 \dotplus \cdots \dotplus Y_r$.

(ii) V is a basis for X.

(iii) $X = Y_1 + \cdots + Y_r$ and $\dim X \ge \dim Y_1 + \cdots + \dim Y_r$.

(iv) For each j, $Y_j \cap Y_{\backslash j} = \{0\}$, with $Y_{\backslash j} := Y_1 + \cdots + Y_{j-1} + Y_{j+1} + \cdots + Y_r$, and $\dim X \le \dim Y_1 + \cdots + \dim Y_r$.

Proof: Since $\operatorname{dom} V = \operatorname{dom} V_1 \times \cdots \times \operatorname{dom} V_r$, and V_j is a basis for Y_j, all j, the linear map

$$C : \operatorname{dom} V \to Y_1 \times \cdots \times Y_r : a = (a_1, \ldots, a_r) \mapsto (V_1 a_1, \ldots, V_r a_r)$$

is invertible and $V = AC$, with A as given in (4.25). Hence, V is invertible if and only if A is invertible. This proves that (i) and (ii) are equivalent.

Also, (ii) implies (iii). As to (iii) implying (ii), the first assumption of (iii) says that V is onto X, and the second assumption says that $\dim \operatorname{dom} V = \#V \le \dim X$, hence V is minimally onto and therefore a basis for X.

As to (ii) implying (iv), the first claim of (iv) is a special case of (4.20)Corollary, and the second claim is immediate.

Finally, as to (iv) implying (ii), assume that $0 = Va = \sum_j V_j a_j$. Then, for any j, $y := V_j a_j = -\sum_{i \ne j} V_i a_i \in Y_j \cap Y_{\backslash j}$, hence $y = 0$ by the first assumption and, since V_j is a basis for Y_j, hence 1-1, this implies that $a_j = 0$. In other words, V is 1-1, while, by the second assumption, $\#V = \sum_j \dim Y_j \ge \dim X$, hence V is maximally 1-1, therefore a basis for X. □

(4.27) Corollary: If V is a basis for X, then, for any grouping $V =: [V_1, \ldots, V_r]$ of the columns of V, X is the direct sum of the linear subspaces $\operatorname{ran} V_j$, $j = 1{:}r$.

One particular grouping is, of course, $V_j = [v_j]$, all j, in which case each $Y_j := \operatorname{ran} V_j$ is a one-dimensional linear subspace, i.e., a straight line through the origin, and we see $X = \operatorname{ran} V$ as the direct sum of these straight lines, each of which we are accustomed to think of as a **coordinate axis**.

This is illustrated in (4.28)Figure for the special case $\operatorname{ran} V = \mathbb{R}^2$, hence V has just two columns. We see each $x \in \mathbb{R}^2$ written as the sum $x = y_1 + y_2$, with $y_j = a_j v_j \in Y_j = \operatorname{ran}[v_j]$ the Y_j-**component** of x (and, of course, $a = (a_1, a_2)$ the coordinate vector of x with respect to the basis V).

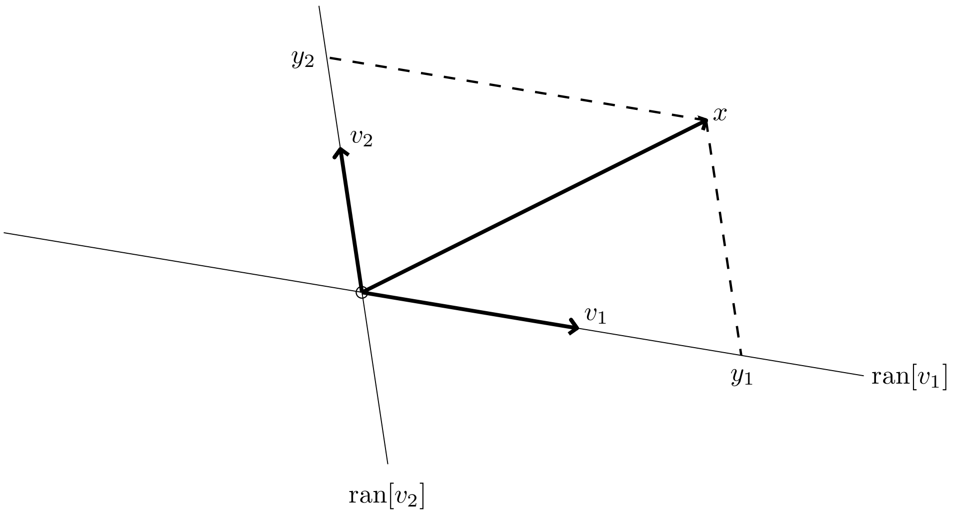

(4.28) Figure. A basis provides a coordinate system.

The direct sum construct is set up in just the same way, except that the Y_j may be planes or even higher-dimensional subspaces rather than just straight lines.

4.29 When X is the direct sum of Y and Z, then Z is said to be a **complement** of Y. With Y and Z linear subspaces of the finite-dimensional vector space X, prove the following assertions concerning complements.

(i) Y has a complement.
(ii) If both Z and Z_1 complement Y, then $\dim Z = \dim Z_1 = \operatorname{codim} Y$. In particular, $\operatorname{codim} Y = \dim X - \dim Y$.
(iii) $\operatorname{codim}(Y + Z) = \operatorname{codim} Y + \operatorname{codim} Z - \operatorname{codim}(Y \cap Z)$.
(iv) If Y has only one complement, then $Y = \{0\}$ or $Y = X$.
(v) If $\operatorname{codim} Y > \dim Z$, then $Y + Z \neq X$.
(vi) If $\dim Y > \operatorname{codim} Z$, then $Y \cap Z \neq \{0\}$.

4.30 Let $(d_1, \ldots, d_r)$ be a sequence of natural numbers, and let X be an n-dimensional vector space. There exists a direct sum decomposition

$$X = Y_1 \dot{+} \cdots \dot{+} Y_r$$

with $\dim Y_j = d_j$, all j, if and only if $\sum_j d_j = n$.

4.31 Let d be any scalar-valued map, defined on the collection of all linear subspaces of a finite-dimensional vector space X, that satisfies the following two conditions: (i) $Y \cap Z = \{0\} \implies d(Y + Z) = d(Y) + d(Z)$; (ii) $\dim Y = 1 \implies d(Y) = 1$.

Prove that $d(Y) = \dim(Y)$ for every linear subspace of X.

4.32 Prove that the cartesian product $Y_1 \times \cdots \times Y_r$ of vector spaces, all over the same scalar field $\mathbb{F}$, becomes a vector space under *pointwise* or **slotwise** addition and multiplication by a scalar.

This vector space is called the **product space** with **factor**s $Y_1, \ldots, Y_r$.

Elimination in vector spaces

In the discussion of the (4.5)Basis Selection Algorithm, we left unanswered the unspoken question of just how one would tell which columns of $W \in L(\mathbb{F}^m, X)$ are bound, hence end up in the resulting 1-1 map V.

The answer is immediate in case $X \subset \mathbb{F}^r$ for some r, for then W is just an $r \times m$-matrix, and elimination does the trick since it is designed to determine the bound columns of a matrix. It works just as well when X is, more generally, a subset of $\mathbb{F}^T$ for some set T, as long as T is finite, since we can then apply elimination to the 'matrix'

$$W = (w_j(t) : (t, j) \in T \times \underline{m}) \tag{4.29}$$

whose rows are indexed by the (finitely many) elements of T.

Elimination even works when T is not finite, since looking for a pivot row in the matrix (4.29) with *infinitely* many rows is only a *practical* difficulty. If τ_i is the row 'index' of the pivot row for the ith bound column of W, $i = 1{:}r$, then we know that W has the same nullspace as the (finite-rowed) matrix $(w_j(\tau_i) : i = 1{:}r, j = 1{:}m)$. This proves, for arbitrary T, the following important

(4.30) Proposition: For any $W \in L(\mathbb{F}^m, \mathbb{F}^T)$, there exists a sequence $(\tau_1, \ldots, \tau_r)$ in T, with r equal to the number of bound columns in W, so that $\operatorname{null} W$ is equal to the nullspace of the matrix $(w_j(\tau_i) : i = 1{:}r, j = 1{:}m)$.

In particular, W is 1-1 if and only if the matrix $(w_j(\tau_i) : i, j = 1{:}m)$ is invertible for some sequence $(\tau_1, \ldots, \tau_m)$ in T.

If T is not finite, then we may not be able to determine in finite time whether or not a given column is bound since we may have to look at infinitely many rows not yet used as pivot rows. The only efficient way around this is to have W given to us in the form

$$W = UA,$$

with U some 1-1 column map, hence A a matrix. Under these circumstances, the kth column of W is free if and only if the kth column of A is free, and the latter we can determine by elimination applied to A.

Indeed, if U is 1-1, then both W and A have the same nullspace, hence, by (4.3)Lemma, the kth column of W is bound if and only if the kth column of A is bound.

As an example, consider $W = [w_1, w_2, w_3, w_4]$, with $w_j : \mathbb{R} \to \mathbb{R} : t \mapsto \sin(t - j)$, $j = 1, 2, 3, 4$. Hence, by the addition formula,

$$W = UA, \quad \text{with } U := [\sin, \cos], \quad A := \begin{bmatrix} \cos(-1) & \cos(-2) & \cos(-3) & \cos(-4) \\ \sin(-1) & \sin(-2) & \sin(-3) & \sin(-4) \end{bmatrix},$$

and we see at once that U is 1-1 (e.g. from the fact that $QU = \mathrm{id}_2$, with $Q : f \mapsto (f(\pi/2), f(0))$). We also see at once that the first two columns of A are bound (e.g., since $\cos(1)\cos(2) < 0$ while $\sin(1)\sin(2) > 0$), hence the remaining columns of A must be free (since there are no rows left to bind them). Consequently, the first two columns of W are bound, while the last two columns are free.

Note that, necessarily, U is a basis for $\operatorname{ran} W$ since $W = UA$ implies that $\operatorname{ran} W \subset \operatorname{ran} U$, hence having two columns of W bound implies that $2 \le \dim \operatorname{ran} W \le \dim \operatorname{ran} U \le \#U = 2$, and so U is 1-1 onto $\operatorname{ran} W$.

In general, it may be hard to find such a handy factorization $W = UA$ for given $W \in L(\mathbb{F}^m, X)$. In that case, we may have to *discretize* our problem by finding somehow some $Q \in L(X, \mathbb{F}^n)$ that is 1-1 on $\operatorname{ran} W$. With such a 'data map' Q in hand, we know that $\operatorname{null} W$ equals the nullspace of the *matrix* QW. In particular, the kth column of W is bound if and only if the kth column of the *matrix* QW is bound, and elimination applied to QW will ferret out all those columns.

The need for suitable 'data maps' here in the general case is one of many reasons why we now turn to the study of this second way of connecting our vector space X to some coordinate space, namely via linear maps from X to $\mathbb{F}^n$.

4.33 For each of the following column maps $V = [v_1, \ldots, v_r]$ into the vector space Π_4 of all real polynomials of degree ≤ 4, determine whether or not it is 1-1 and/or onto.

(a) $[()^3 - ()^1 + 1, ()^2 + 2()^1 + 1, ()^1 - 1]$; (b) $[()^4 - ()^1, ()^3 + 2, ()^2 + ()^1 - 1, ()^1 + 1]$; (c) $[1 + ()^4, ()^4 + ()^3, ()^3 + ()^2, ()^2 + ()^1, ()^1 + 1]$.

4.34 For each of the specific column maps $V = [f_j : j = 0{:}r]$ given below (with f_j certain real-valued functions on the real line), determine which columns are bound and which are free. Use this information to determine (i) a basis for $\operatorname{ran} V$; and (ii) the smallest n so that $f_n \in \operatorname{ran}[f_0, f_1, \ldots, f_{n-1}]$.

(a) $r = 6$, and $f_j : t \mapsto (t - j)^2$, all j.

(b) $r = 4$ and $f_j : t \mapsto \sin(t - j)$, all j.

(c) $r = 4$ and $f_j : t \mapsto \exp(t - j)$, all j. (If you know enough about the exponential function, then you need not carry out any calculation on this problem.)

4.35 Assume that $\tau_1 < \cdots < \tau_{2k+1}$. Prove that $W = [w_0, \ldots, w_k]$ with $w_j : t \mapsto (t - \tau_{j+1}) \cdots (t - \tau_{j+k})$ is a basis for Π_k. (Hint: Consider QW with $Q : p \mapsto (p(\tau_{k+1+i}) : i = 0{:}k)$.)

4.36 Assume that $(\tau_1, \ldots, \tau_{2k+1})$ is nondecreasing. Prove that $W = [w_0, \ldots, w_k]$ with $w_j : t \mapsto (t - \tau_{j+1}) \cdots (t - \tau_{j+k})$ is a basis for Π_k if and only if $\tau_k < \tau_{k+1}$.

4.37 T/F

(a) If one of the columns of a column map is 0, then the map cannot be 1-1.

(b) If the column map V into $\mathbb{R}^n$ is 1-1, then V has at most n columns.

(c) If the column map V into $\mathbb{R}^n$ is onto, then V has at most n columns.

(d) If a column map fails to be 1-1, then it has a zero column.

(e) If a vector space has only one basis, then it must be the trivial space.

(f) If a column of a matrix A is free, then it cannot be part of a basis for $\operatorname{ran} A$.

5. The inverse of a basis, and interpolation

Data maps

There are two ways to connect a given vector space X with the coordinate space $\mathbb{F}^n$ in a linear way, namely by a linear map from $\mathbb{F}^n$ to X, and by a linear map to $\mathbb{F}^n$ from X. By now, you are thoroughly familiar with the first kind, the column maps. It is time to learn something about the other kind.

A very important example of such a map is the inverse of a basis $V : \mathbb{F}^n \to X$ for the vector space X, also known as the **coordinate map** for that basis because it provides, for each $x \in X$, its **coordinates with respect to the basis**, i.e., the n-vector $a := V^{-1}x$ for which $x = Va$. In effect, every *invertible* linear map from X to $\mathbb{F}^n$ is a coordinate map, namely the coordinate map for its inverse. However, (nearly) every linear map from X to $\mathbb{F}^n$, invertible or not, is of interest, as a means of extracting numerical information from the elements of X. For, we can, offhand, only compute with numbers, hence can 'compute' with elements of an abstract vector space only in terms of numerical data about them.

Any linear map from the vector space X to $\mathbb{F}^n$ is necessarily of the form

$$f : X \to \mathbb{F}^n : x \mapsto (f_i(x) : i = 1{:}n),$$

with each $f_i = e_i{}^{\mathrm{t}} \circ f$ a **linear functional** on X, i.e., a scalar-valued linear map on X.

5.1 For each of the following maps, determine whether or not it is a linear functional. (a) $\Pi_k \to \mathbb{R} : p \mapsto \deg p$; (b) $\mathbb{R}^3 \to \mathbb{R} : x \mapsto 3x_1 - 2x_3$; (c) $C([a\,..\,b]) \to \mathbb{R} : f \mapsto \max_{a\le t\le b} f(t)$; (d) $C([a\,..\,b]) \to \mathbb{R} : f \mapsto \int_a^b f(s)w(s)\,\mathrm{d}s$, with $w \in C([a\,..\,b])$; (e) $C^{(2)}(\mathbb{R}) \to \mathbb{R} : f \mapsto a(t)D^2 f(t) + b(t)Df(t) + c(t)f(t)$, for some functions a, b, c defined on $[a\,..\,b]$ and some $t \in [a\,..\,b]$. (f) $C^{(2)}(\mathbb{R}) \to C(\mathbb{R}) : f \mapsto aD^2 f + bDf + cf$, for some $a, b, c \in C(\mathbb{R})$.

Here are some standard examples of linear functionals. Assume that X is a space of functions, hence X is a linear subspace of $\mathbb{F}^T$ for some set T. Then, for each $t \in T$,

$$\delta_t : X \to \mathbb{F} : x \mapsto x(t)$$

is a linear functional on X, the linear functional of evaluation at t. For any n-sequence $s = (s_1, \ldots, s_n)$ in T,

$$X \to \mathbb{F}^n : f \mapsto (f(s_1), \ldots, f(s_n))$$

is a standard linear map from X to $\mathbb{F}^n$.

If, more concretely, X is a linear subspace of $C^{(n-1)}[a\,..\,b]$ and $s \in [a\,..\,b]$, then

$$X \to \mathbb{F}^n : f \mapsto (f(s), Df(s), \ldots, D^{n-1}f(s))$$

is another standard linear map from such X to $\mathbb{F}^n$.

Finally, if $X = \mathbb{F}^m$, then any linear map from X to $\mathbb{F}^n$ is necessarily a matrix. But it is convenient to write this matrix in the form A^{t} for some $A \in \mathbb{F}^{n\times m}$, as such A^{t} acts on X via the rule

$$X \mapsto \mathbb{F}^n : x \mapsto A^{\mathrm{t}}x = (A(:,j)^{\mathrm{t}}x : j = 1{:}n).$$

Because of this last example, we will call all linear maps from a vector space to a coordinate space **row maps**, and use the notation

$$\Lambda^{\mathrm{t}} : X \to \mathbb{F}^n : x \mapsto (\lambda_i x : i = 1{:}n) =: [\lambda_1, \ldots, \lambda_n]^{\mathrm{t}}x, \tag{5.1}$$

calling the linear functional λ_i the ith **row** of this map. We will also call such maps **data maps** since they extract numerical information from the elements of X. There is no hope of doing any practical work with the vector space X unless we have a ready supply of such data maps on X. For, by and large, we can only compute with numbers.

(5.2)Proposition: If $\Lambda^{\mathrm{t}} = [\lambda_1, \lambda_2, \ldots, \lambda_n]^{\mathrm{t}} : X \to \mathbb{F}^n$ and $B \in L(U, X)$, then $\Lambda^{\mathrm{t}}B = [\lambda_1 B, \ldots, \lambda_n B]^{\mathrm{t}}$.

A formula for the coordinate map

Let $V \in L(\mathbb{F}^n, X)$ be a basis for the vector space X. How do we find the coordinates

$$a = V^{-1}x \tag{5.3}$$

for given $x \in X$?

Offhand, we solve the (linear) equation $V? = x$ for a. Since V is a basis, we know that this equation has exactly one solution. But that is not the same thing as having a concrete formula for a in terms of x.

If $X = \mathbb{F}^n$, then V^{-1} is a matrix; in this case, (5.3) is an explicit formula. However, even if $X \subset \mathbb{F}^n$ but $X \neq \mathbb{F}^n$, then (5.3) is merely a formal expression.

(5.4) Example: If V is a basis for some linear subspace X of $\mathbb{F}^n$, then we can obtain a formula for V^{-1} via elimination as follows.

Act as if V were invertible, i.e., apply elimination to $[V, \mathrm{id}_n]$. Let $r := \#V$. Since V is 1-1, the first r columns in $[V, \mathrm{id}_n]$ are bound, hence we are able to produce, via elimination, an equivalent matrix R for which $R(\mathtt{q}, 1{:}r) = \mathrm{id}_r$, for some r-sequence $\mathtt{q}$. Since we obtain R from $[V, \mathrm{id}_n]$ by (invertible) row operations, we know that $R = M[V, \mathrm{id}_n] = [MV, M]$ for some invertible matrix M. In particular,

$$\mathrm{id}_r = R(\mathtt{q}, 1{:}r) = (MV)(\mathtt{q}, :) = M(\mathtt{q}, :)V,$$

showing $M(\mathtt{q}, :) = R(\mathtt{q}, r + (1{:}n))$ to be a left inverse for V, hence equal to V^{-1} when restricted to $\operatorname{ran} V$.

Suppose, in particular, that we carry elimination all the way through, to obtain $R = \operatorname{rref}([V, \mathrm{id}_n])$. Then, $\mathtt{q} = 1{:}r$ and, with $r + \mathtt{b}$ and $r + \mathtt{f}$ the bound and free columns of $[V, \mathrm{id}_n]$ other than the columns of V, we necessarily have $M(\mathtt{q}, \mathtt{b}) = 0$, hence, for this choice of M, we get

$$V^{-1}x = M(\mathtt{q}, :)x = M(\mathtt{q}, \mathtt{f})x(\mathtt{f}), \qquad x \in X := \operatorname{ran} V.$$

In effect, we have replaced here the equation $V? = x$ by the *equivalent* equation

$$V(\mathtt{f}, :)? = x(\mathtt{f})$$

whose coefficient matrix is invertible. In particular, $\#\mathtt{f} = \#V$; see H.P. 5.3 . □

5.2 For each of the following bases V of the linear subspace $\operatorname{ran} V$ of $\mathbb{F}^n$, give a matrix U for which Ux gives the coordinates of $x \in \operatorname{ran} V$ with respect to the basis V. How would you check your answer?

(a) $V = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$; (b) $V = [e_2, e_1, e_3] \in \mathbb{R}^{3\times 3}$; (c) $V = \begin{bmatrix} 1 & 2 \\ 2 & 4 \\ 0 & 6 \end{bmatrix}$; (d) $V = \begin{bmatrix} 1 & 0 \\ 0 & 0 \\ -1 & 1 \\ 2 & -2 \end{bmatrix}$.

5.3 Prove the claim at the end of (5.4)Example.

The reduction in (5.4)Example, of the abstract linear equation $V? = x$ to a uniquely solvable square linear system, also works in the general setting.

To obtain a concrete expression, we **discretize** the abstract equation $V? = x$ by considering instead the *numerical* equation

$$\Lambda^{\mathrm{t}} V? = \Lambda^{\mathrm{t}} x$$

for some suitable data map $\Lambda^{\mathrm{t}} \in L(Y, \mathbb{F}^n)$ defined on some vector space $Y \supset X$. Here, 'suitable' means that the *matrix* $\Lambda^{\mathrm{t}} V$ is invertible, for then the unique solution of this equation must be the sought-for coordinate vector for $x \in X$ with respect to the basis V, i.e.,

$$a = V^{-1}x = (\Lambda^{\mathrm{t}} V)^{-1}\Lambda^{\mathrm{t}} x.$$

In (5.4)Example, we simply chose the linear map $y \mapsto y(\mathtt{f})$ as our Λ^{t}, i.e., $\Lambda^{\mathrm{t}} = \mathrm{id}_n(\mathtt{f}, :) = [e_j : j \in \mathtt{f}]^{\mathrm{t}}$, with $\mathtt{f}$ chosen, in effect, to ensure that $\Lambda^{\mathrm{t}} V = V(\mathtt{f}, :)$ is invertible. We indeed obtained there V^{-1} as

$$x \mapsto U(:, \mathtt{f})x(\mathtt{f}) = V(\mathtt{f}, :)^{-1}x(\mathtt{f}) = (\Lambda^{\mathrm{t}} V)^{-1}\Lambda^{\mathrm{t}} x.$$

How would one find a 'suitable' data map in general? That depends on the particular circumstances. For example, if $V \in L(\mathbb{F}^n, Y)$ and $\Lambda^t \in L(Y, \mathbb{F}^n)$, and we somehow know that Λ^t maps $X := \operatorname{ran} V = V(\mathbb{F}^n)$ *onto* $\mathbb{F}^n$, then we know that $\Lambda^t V$ maps $\mathbb{F}^n$ onto $\mathbb{F}^n$, hence, being a square matrix, $\Lambda^t V$ must be invertible. Conversely, if $\Lambda^t V$ is invertible, then V must be 1-1, hence provides a basis for its range, and Λ^t must map $\operatorname{ran} V$ onto $\mathbb{F}^n$.

(5.5) Proposition: If the linear map $V : \mathbb{F}^n \to X \subset Y$ is onto, and $\Lambda^t \in L(Y, \mathbb{F}^n)$ is such that their (square) **Gramian matrix**, $\Lambda^t V$, is 1-1 or onto, hence invertible, then V is a basis for X, and its inverse is

$$V^{-1} : X \to \mathbb{F}^n : x \mapsto (\Lambda^t V)^{-1} \Lambda^t x.$$

Change of basis

To be sure, under the assumptions of (5.5)Proposition, we also know that Λ^t maps X onto $\mathbb{F}^n$, hence, since both X and $\mathbb{F}^n$ are of the same finite dimension, the restriction of Λ^t to X must be invertible as a linear map to $\mathbb{F}^n$. Consequently, there must be an invertible $W \in L(\mathbb{F}^n, X)$, i.e., a basis W for X, with $\Lambda^t W = \mathrm{id}_n$.

Hence, the right side in our numerical equation $\Lambda^t V? = \Lambda^t x$ is the coordinate vector for $x \in X$ with respect to this basis W for X. In other words, our great scheme for computing the coordinates of $x \in X$ with respect to the basis V for X requires us to know the coordinates of x with respect to some basis for X. In other words, the entire calculation is just a *change of basis*, with $\Lambda^t V = W^{-1} V$ the socalled **transition matrix** that carries the V-coordinates of x to the W-coordinates of x.

However, this in no way diminishes its importance. For, we really have no choice in this matter. We cannot compute without having numbers to start with. Also, we often have ready access to the coordinate vector $\Lambda^t x$ without having in hand the corresponding basis W. At the same time, we may much prefer to know the coordinates of x with respect to a different basis.

For example, we know from (3.22)Proposition that, with $(\tau_1, \ldots, \tau_k)$ any sequence of pairwise distinct real numbers, the linear map $\Lambda^t : p \mapsto (p(\tau_1), \ldots, p(\tau_k))$ is 1-1 on the k-dimensional space $\Pi_{<k}$, hence provides the coordinates of $p \in \Pi_{<k}$ with respect to a certain basis W of $\Pi_{<k}$, namely the socalled **Lagrange basis** whose columns can be verified to be the so-called **Lagrange fundamental polynomials**

$$\ell_j : t \mapsto \prod_{h \neq j} \frac{t - \tau_h}{\tau_j - \tau_h}, \qquad j = 1{:}k. \tag{5.6}$$

However, you can imagine that it is a challenge to differentiate or integrate a polynomial written in terms of this basis. Much better for that to have the coordinates of the polynomial with respect to the power basis $V = [()^0, \ldots, ()^{k-1}]$.

5.4 What are the coordinates of $p \in \Pi_k$ with respect to the Lagrange basis for $\Pi_{<k}$ for the points $\tau_1, \ldots, \tau_k$?

5.5 Find the value at 0 of the quadratic polynomial p, for which $p(-1) = p(1) = 3$ and $Dp(1) = 6$.

5.6 Find a formula for $p(0)$ in terms of $p(-1)$, $p(1)$ and $Dp(1)$, assuming that p is a quadratic polynomial.

5.7 Find the coordinates for the polynomial $q(t) = 3 - 4t + 2t^2$ with respect to the basis $W := [()^0, ()^0 + ()^1, ()^0 + ()^1 + ()^2]$ of the space of quadratic polynomials. (Hint: you are given the coordinates for q wrto $V := [()^0, ()^1, ()^2] = W(W^{-1}V)$ and can easily determine $(W^{-1}V)^{-1} = V^{-1}W$.)

5.8 Find the coordinates for the polynomial $q(t) = 3 - 4t + 2t^2$ with respect to the basis $W := [()^0, ()^0 + ()^1, ()^0 + ()^1 + ()^2]$ of the space of quadratic polynomials. (Hint: you are given the coordinates for q wrto $V := [()^0, ()^1, ()^2] = W(W^{-1}V)$ and can easily determine $W^{-1}V)^{-1} = V^{-1}W$.)

5.9 Let $v_1, \ldots, v_n$ be a sequence of $(n-1)$-times continuously differentiable functions, all defined on the interval $[a \,.\,.\, b]$. For $x \in [a \,.\,.\, b]$, the matrix

$$W(v_1, \ldots, v_n; x) := (D^{i-1} v_j(x) : i, j = 1{:}n)$$

is called the **Wronski matrix at** x for the sequence $(v_j : j = 1{:}n)$.

Prove that $V := [v_1, \ldots, v_n]$ is 1-1 in case, for some $x \in [a \,.\,.\, b]$, $W(v_1, \ldots, v_n; x)$ is invertible. (Hint: Consider the Gram matrix $\Lambda^t V$ with $\Lambda^t f := (f(x), f'(x), \ldots, D^{n-1} f(x))$.)

Interpolation and linear projectors

As (3.21)Example already intimates, our formula in (5.5) for the inverse of a basis $V \in L(\mathbb{F}^n, X)$ can be much more than that. It is useful for *interpolation* in the following way. Assuming that $\Lambda^{\mathrm{t}}V$ is invertible, it follows that, for any $y \in Y$, $x = V(\Lambda^{\mathrm{t}}V)^{-1}\Lambda^{\mathrm{t}}y$ is the unique element in X that **agrees with** y **at** Λ^{t} in the sense that

$$\Lambda^{\mathrm{t}}x = \Lambda^{\mathrm{t}}y.$$

To recall the specifics of (3.21)Example, if $X = \Pi_{<k}$ and $\Lambda^{\mathrm{t}} : g \mapsto (g(\tau_i) : i = 1{:}k)$, with $\tau_1 < \cdots < \tau_k$, then, by (3.22)Proposition, for arbitrary $g : \mathbb{R} \to \mathbb{R}$, there is exactly one polynomial p of degree $< k$ for which $p(\tau_i) = g(\tau_i)$, $i = 1{:}k$.

One can readily imagine other examples.

Example: In **Hermite interpolation**, one specifies not only values but also derivatives. For example, in **two-point** Hermite interpolation from $\Pi_{<k}$, one picks two points, $t \neq u$, and two nonnegative integers r and s with $r + 1\ + s + 1 = k$, and defines

$$\Lambda^{\mathrm{t}} : g \mapsto (g(t), Dg(t), \ldots, D^r g(t), g(u), Dg(u), \ldots, D^s g(u)).$$

Now the requirement that $\Lambda^{\mathrm{t}}p = \Lambda^{\mathrm{t}}g$ amounts to looking for $p \in \Pi_{<k}$ that agrees with g in the sense that p and g have the same derivative values of order $0, 1, \ldots, r$ at t and the same derivative values of order $0, 1, \ldots, s$ at u. □

Example: Recall from Calculus the bivariate **Taylor series**

$$g(s,t) = g(0) + D_s g(0)\, s + D_t g(0)\, t + \big({D_s}^2 g(0)s^2 + D_s D_t g(0)st + D_t D_s g(0)ts + {D_t}^2 g(0)t^2\big)/2 + h.o.t.$$

In particular, for any smooth function g, the quadratic polynomial

$$p : (s,t) \mapsto g(0) + D_s g(0)\, s + D_t g(0)\, t + \big({D_s}^2 g(0)s^2 + 2D_s D_t g(0)st + {D_t}^2 g(0)t^2\big)/2$$

is the unique quadratic polynomial that matches the information about g given by the data map

$$\Lambda^{\mathrm{t}} : g \mapsto (g(0), D_s g(0), D_t g(0), D_s^2 g(0), D_s D_t g(0), D_t^2 g(0)).$$

□

Example: When dealing with **Fourier series**, one uses the data map

$$\Lambda^{\mathrm{t}} : g \mapsto (\int_0^{2\pi} g(t)\,\mathrm{cis}(jt)\,\mathrm{d}t : j = 0{:}N),$$

with cis standing for 'sine and cosine'. One looks for a **trigonometric polynomial**

$$p = [\mathrm{cis}(j\cdot) : j = 0{:}N]a$$

that satisfies $\Lambda^{\mathrm{t}}p = \Lambda^{\mathrm{t}}g$, and finds it in the **truncated Fourier series** for g. □

Directly from (5.5)Proposition, we obtain (under the assumptions of that proposition) the following pretty formula

$$x = Py := V(\Lambda^{\mathrm{t}}V)^{-1}\Lambda^{\mathrm{t}}y \tag{5.7}$$

for the interpolant $x \in X$ to given $y \in Y$ with respect to the data map Λ^{t}. The linear map $P := V(\Lambda^{\mathrm{t}}V)^{-1}\Lambda^{\mathrm{t}}$ so defined on Y is very special:

(5.8) Proposition: Let the linear map $V : \mathbb{F}^n \to Y$ be onto $X \subset Y$, and let $\Lambda^{\mathrm{t}} \in L(Y, \mathbb{F}^n)$ be such that their Gramian matrix, $\Lambda^{\mathrm{t}}V$, is invertible. Then $P := V(\Lambda^{\mathrm{t}}V)^{-1}\Lambda^{\mathrm{t}}$ is a linear map on Y with the following properties:

(i) P is the identity on $X = \operatorname{ran} V$.

(ii) $\operatorname{ran} P = \operatorname{ran} V = X$.

(iii) P is a **projector** or **idempotent**, i.e., $PP = P$, hence $P(\mathrm{id} - P) = 0$.

(iv) $\operatorname{null} P = \operatorname{null} \Lambda^{\mathrm{t}} = \operatorname{ran}(\mathrm{id} - P)$.

(v) Y is the direct sum of $\operatorname{ran} P$ and $\operatorname{null} P$, i.e., $Y = \operatorname{ran} P \dotplus \operatorname{null} P$.

Proof: (i) $PV = V(\Lambda^{\mathrm{t}}V)^{-1}\Lambda^{\mathrm{t}}\,V = V\,\mathrm{id} = V$, hence $P(Va) = Va$ for all $a \in \mathbb{F}^n$.

(ii) Since $P = VA$ for some A, we have that $\operatorname{ran} P \subset \operatorname{ran} V$, while $PV = V$ implies that $\operatorname{ran} P \supset \operatorname{ran} V$.

(iii) By (i) and (ii), P is the identity on its range, hence, in particular, $PP = P$, or, equivalently, $P(\mathrm{id} - P) = 0$.

(iv) The fact that $P = A\Lambda^{\mathrm{t}}$ for some A implies that $\operatorname{null} P \supset \operatorname{null} \Lambda^{\mathrm{t}}$, while also

$$\Lambda^{\mathrm{t}}P = \Lambda^{\mathrm{t}}\,V(\Lambda^{\mathrm{t}}V)^{-1}\Lambda^{\mathrm{t}} = \mathrm{id}_n\Lambda^{\mathrm{t}} = \Lambda^{\mathrm{t}},$$

hence also $\operatorname{null} P \subset \operatorname{null} \Lambda$. As to $\operatorname{null} P = \operatorname{ran}(\mathrm{id} - P)$, note that $x \in \operatorname{null} P$ implies that $x = x - Px = (\mathrm{id} - P)x \in \operatorname{ran}(\mathrm{id} - P)$, while, conversely, $\operatorname{null} P \supset \operatorname{ran}(\mathrm{id} - P)$ since, by (iii), $P(\mathrm{id} - P) = 0$.

(v) For any $y \in Y$, $y = Py + (\mathrm{id} - P)y \in \operatorname{ran} P + \operatorname{null} P$, by (iv), hence $Y = \operatorname{ran} P + \operatorname{null} P$. If also $y = x + z$ for some $x \in \operatorname{ran} P$ and some $z \in \operatorname{null} P$, then, by (i) and (iv), $Py = P(x + z) = Px + Pz = x$, therefore also $z = y - x = y - Py = (\mathrm{id} - P)y$, showing such a decomposition to be unique. □

5.10 Let $P \in L(X)$. (i) Prove that P is a projector if and only if $R := \mathrm{id} - 2P$ is **involutory** or **self-inverse** (meaning that $RR = \mathrm{id}$). (ii) For the linear projector P of (5.9)Example, work out the corresponding map R, and add to (5.10)Figure the point Ry.

5.11 Consider the linear map Q given on $X = \{f : \mathbb{R} \to \mathbb{R}\}$ by $Qf(t) = (f(t) + f(-t))/2$. Prove that Q is a linear projector. Also, give a succinct description of its range and its nullspace. (Hint: consider the map $F : X \to X$ defined by $(Ff)(t) = -f(t)$.)

(5.9) Example: We specialize the general situation of (5.8)Proposition to the case $V : \mathbb{R}^1 \to X \subset \mathbb{R}^2$, so we can draw a figure like (5.10)Figure.

Take $Y = \mathbb{R}^2$, and let $v \in \mathbb{R}^2 \neq 0$, hence $X := \operatorname{ran} V$ with $V := [v]$ is 1-dimensional. The general linear map $\Lambda^{\mathrm{t}} : \mathbb{R}^2 \to \mathbb{R}^1$ is of the form $[w]^{\mathrm{t}}$ for some $w \in \mathbb{R}^2$, and the requirement that $\Lambda^{\mathrm{t}}V$ be invertible reduces to the requirement that $[w]^{\mathrm{t}}[v] = w^{\mathrm{t}}v \neq 0$.

With $V = [v]$ and $\Lambda^{\mathrm{t}} = [w]^{\mathrm{t}}$ so chosen, the linear projector P becomes

$$P := \frac{vw^{\mathrm{t}}}{w^{\mathrm{t}}v} : y \mapsto v\frac{w^{\mathrm{t}}y}{w^{\mathrm{t}}v}.$$

We readily verify directly that

$$PP = \frac{vw^{\mathrm{t}}}{w^{\mathrm{t}}v}\frac{vw^{\mathrm{t}}}{w^{\mathrm{t}}v} = \frac{v\,w^{\mathrm{t}}v\,w^{\mathrm{t}}}{(w^{\mathrm{t}}v)\,(w^{\mathrm{t}}v)} = \frac{vw^{\mathrm{t}}}{w^{\mathrm{t}}v} = P,$$

i.e., that P is a linear projector. Its range equals $\mathrm{ran}[v]$, i.e., the straight line through the origin in the direction of v. Its nullspace equals $\mathrm{null}[w]^{\mathrm{t}}$ and this is necessarily also 1-dimensional, by (4.15)Dimension Formula, hence is the straight line through the origin perpendicular to w. The two lines have only the origin in common since $y \in \mathrm{ran}\, P \cap \mathrm{null}\, P$ implies that $y = v\alpha$ for some scalar α, therefore $0 = w^{\mathrm{t}}y = w^{\mathrm{t}}v\alpha$ and this implies that $\alpha = 0$ since $w^{\mathrm{t}}v \neq 0$ by assumption.

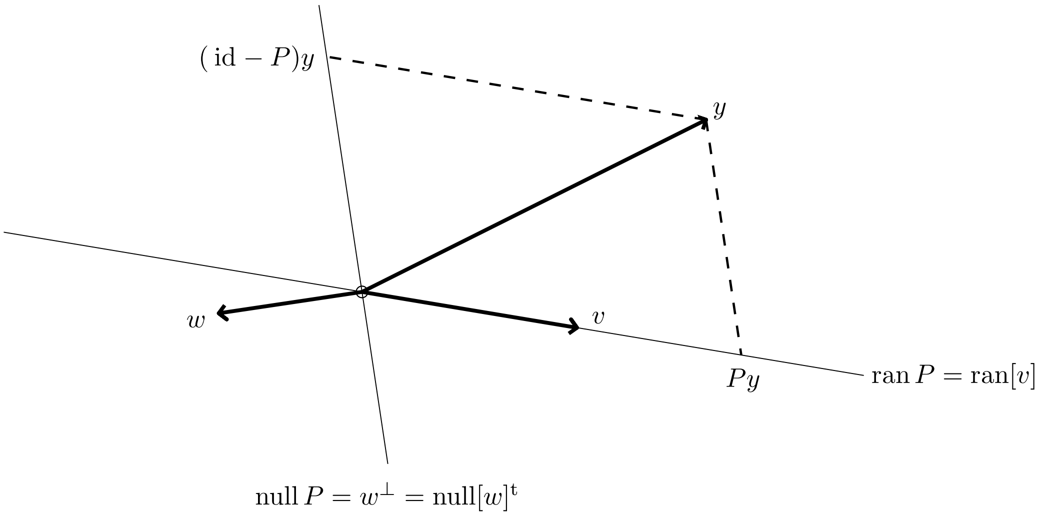

(5.10) Figure. The direct sum decomposition provided by a certain linear projector. Compare this to (4.28)Figure.

We can locate the two summands in the split

$$y = Py + (\mathrm{id} - P)y$$

graphically (see (5.10)Figure): To find Py, draw the line through y parallel to $\mathrm{null}\, P$; its unique intersection with $\mathrm{ran}\, P = \mathrm{ran}[v]$ is Py. The process of locating $(\mathrm{id} - P)y$ is the same, with the roles of $\mathrm{ran}\, P$ and $\mathrm{null}\, P$ reversed: Now draw the line through y parallel to $\mathrm{ran}\, P$; its unique intersection with $\mathrm{null}\, P$ is the element $(\mathrm{id} - P)y$.

This shows graphically that, for each y, Py is the unique element of $\mathrm{ran}\, P$ for which $w^{\mathrm{t}}Py = w^{\mathrm{t}}y$, i.e., the unique point in the intersection of $\mathrm{ran}\, P$ and $y + \mathrm{null}[w]^{\mathrm{t}}$. □

It is useful to note that, for any linear projector P, also $(\mathrm{id} - P)$ is a linear projector (since $(\mathrm{id} - P)(\mathrm{id} - P) = \mathrm{id} - P - P + PP = \mathrm{id} - P$), and that any direct sum decomposition $Y = X \dot{+} Z$ of a finite-dimensional Y necessarily has $X = \mathrm{ran}\, P$ and $Z = \mathrm{null}\, P$ for some linear projector P. The following is a more general such claim, of use later.

(5.11) Proposition: Let $X_1, \ldots, X_r$ be linear subspaces of the finite-dimensional vector space Y. Then the following are equivalent.

(i) Y is the direct sum of the X_j, i.e., $Y = X_1 \dotplus \cdots \dotplus X_r$.

(ii) There exist $P_j \in L(Y)$ with $\operatorname{ran} P_j = X_j$ so that

$$\mathrm{id}_Y = P_1 + \cdots + P_r \tag{5.12}$$

and

$$P_j P_k = \begin{cases} P_j = P_k & \text{if } j = k; \\ 0 & \text{otherwise.} \end{cases} \tag{5.13}$$

In particular, each P_j is a linear projector.

Also, the conditions in (ii) uniquely determine the P_j.

Proof: Let V_j be a basis for X_j, all j. By (4.26)Proposition, (i) is equivalent to having $V := [V_1, \ldots, V_r]$ be a basis for Y.

'(i) $\Longrightarrow$ (ii)': By assumption, V is a basis for Y. Let $V^{-1} =: \Lambda^{\mathrm t} =: [\Lambda_1, \ldots, \Lambda_r]^{\mathrm t}$ be its inverse, grouped correspondingly. Then

$$\mathrm{id}_{\dim Y} = \Lambda^{\mathrm t} V = [\Lambda_1, \ldots, \Lambda_r]^{\mathrm t}[V_1, \ldots, V_r] = ({\Lambda_i}^{\mathrm t} V_j : i, j = 1{:}r),$$

i.e.,

$${\Lambda_i}^{\mathrm t} V_j = \begin{cases} \mathrm{id} & \text{if } i = j; \\ 0 & \text{otherwise.} \end{cases}$$

Hence, the linear maps

$$P_j := V_j {\Lambda_j}^{\mathrm t}, \quad j = 1{:}r,$$

satisfy (5.13), and $\operatorname{ran} P_j = X_j$, for all j. But also

$$\mathrm{id}_Y = V\Lambda^{\mathrm t} = [V_1, \ldots, V_r][\Lambda_1, \ldots, \Lambda_r]^{\mathrm t} = \sum_j V_j {\Lambda_j}^{\mathrm t},$$

showing (5.12).

'(ii) $\Longrightarrow$ (i)': By assumption, $\operatorname{ran} P_j = \operatorname{ran} V_j$, all j. Therefore, by assumption (5.13),

$$P_j V_i = \begin{cases} V_j & \text{if } j = i; \\ 0 & \text{otherwise.} \end{cases} \tag{5.14}$$

Therefore, $0 = Va = \sum_i V_i a_i$ implies, for any particular j, that $0 = P_j 0 = P_j V a = \sum_i P_j V_i a_i = P_j V_j a_j = V_j a_j$, hence $a_j = 0$ (since V_j is 1-1). It follows that V is 1-1. On the other hand, the assumption (5.12) implies that V is onto. Hence, V is a basis for Y.

Finally, to prove the uniqueness of the P_j satisfying (ii), notice that (5.14) pins down P_j on all the columns of V. Since (ii) implies that V is a basis for Y, this therefore determines P_j uniquely (by (4.2)Proposition). □

Returning to the issue of interpolation, this gives the following

(5.15)Corollary: If $V \in L(\mathbb{F}^n, Y)$ is 1-1, and $\Lambda^t \in L(Y, \mathbb{F}^n)$ is such that $\operatorname{ran} V \cap \operatorname{null} \Lambda^t = \{0\}$, then $P := V(\Lambda^t V)^{-1}\Lambda^t$ is well-defined; it is the *unique* linear projector P with

$$\operatorname{ran} P = \operatorname{ran} V, \qquad \operatorname{null} P = \operatorname{null} \Lambda^t. \tag{5.16}$$

In particular, then Λ^t is onto, and

$$Y = \operatorname{ran} V \dotplus \operatorname{null} \Lambda^t. \tag{5.17}$$

For an arbitrary abstract vector space, it may be very hard to come up with suitable concrete data maps. For that reason, we now consider a particular kind of vector space for which it is very easy to provide suitable data maps, namely the inner product spaces.

6. Inner product spaces

Definition and examples

Inner product spaces are vector spaces with an additional operation, the *inner product*. Here is the definition.

(6.1) Definition: An **inner product space** is a vector space Y (over the field $\mathbb{F} = \mathbb{R}$ or $\mathbb{C}$) and an **inner product**, meaning a map

$$\langle , \rangle : Y \times Y \to \mathbb{F} : (x, y) \mapsto \langle x, y \rangle$$

that is

(a) **positive definite**, i.e., $\|x\|^2 := \langle x, x\rangle \ge 0$, with equality iff $x = 0$;

(b) **linear in its first argument**, i.e., $\langle \cdot, y\rangle \in L(Y, \mathbb{F})$;

(c) **hermitian**, or **skew-symmetric**, i.e., $\langle y, x\rangle = \overline{\langle x, y\rangle}$.

You already know an inner product space, namely n-dimensional Euclidean space, i.e., the space of n-vectors with the inner product

$$\langle x, y\rangle := \overline{y}^{\mathrm{t}} x = \sum_j x_j \overline{y_j} =: y^{\mathrm{c}} x,$$

though you may know it under the name **scalar product** or **dot product**. In particular, (b) and (c) are evident in this case. As to (a), observe that, for any complex number $z = u + \mathrm{i}v$,

$$\overline{z} z = (u - \mathrm{i}v)(u + \mathrm{i}v) = u^2 + v^2 = |z|^2 \ge 0,$$

with equality if and only if $u = 0 = v$, i.e., $z = 0$. Hence, for any $x \in \mathbb{F}^n$,

$$\langle x, x\rangle = \overline{x}^{\mathrm{t}} x = |x_1|^2 + \cdots + |x_n|^2 \ge 0,$$

with equality iff all the x_j are zero, i.e., $x = 0$.

Of course, if the scalar field is $\mathbb{R}$, we can forget about taking complex conjugates since then $\overline{x} = x$. But if $\mathbb{F} = \mathbb{C}$, then it is essential that we define $\langle x, y\rangle$ as $y^{\mathrm{c}} x = \overline{y}^{\mathrm{t}} x$ rather than as $y^{\mathrm{t}} x$ since we would not get positive definiteness otherwise. Indeed, if z is a complex number, then there is no reason to think that z^2 is nonnegative, and the following calculation

$$(1, \mathrm{i})^{\mathrm{t}} (1, \mathrm{i}) = 1^2 + (\mathrm{i})^2 = 1 - 1 = 0$$

shows that, for a complex x, $x^{\mathrm{t}} x$ can be zero without x being zero.

So, why not simply stick with $\mathbb{F} = \mathbb{R}$? Work on eigenvalues requires consideration of *complex* scalars (since it relies on zeros of polynomials, and a polynomial may have complex zeros even if all its coefficients are real). For this reason, we have taken the trouble all along to take into account the possibility that $\mathbb{F}$ might be $\mathbb{C}$. It is a minor nuisance at this point, but will save time later.

Another example of an inner product space of great practical interest is the space $Y = \overset{\circ}{C}$ of all continuous 2π-periodic functions, with the inner product

$$\langle f, g\rangle := \int_0^{2\pi} f(t)\overline{g(t)}\, \mathrm{d}t.$$

Of course, we can also think of the space $C([a\,..\,b])$ as an inner product space, with respect to the inner product

$$\langle f, g\rangle := \int_a^b f(t)\overline{g(t)}\,\mathrm{d}t.$$

Often, it is even useful to consider on $C([a\,..\,b])$ the more general inner product

$$\langle f, g\rangle := \int_a^b f(t)\overline{g(t)}w(t)\,\mathrm{d}t$$

with w some positive function on $[a\,..\,b]$, and there are analogous inner product spaces consisting of functions of several variables.

In order to stress the fact that a general inner product space Y behaves just like $\mathbb{F}^n$ with the standard inner product, I will use the notation

$$y^{\mathrm{c}} : Y \to \mathbb{F} : x \mapsto \langle x, y\rangle, \qquad \forall y \in Y,$$

for the linear functional provided, according to (6.1)(b), by the inner product, hence will feel free to write $y^{\mathrm{c}}x$ rather than $\langle x, y\rangle$ for the inner product of x with y. Correspondingly, you can read the rest of this chapter as if we were just talking about the familiar space of n-vectors with the dot product, yet be certain that, when the time comes, you will have in hand very useful facts about an arbitrary inner product space, for example the space $\mathring{C}$.

The conjugate transpose

Here is the promised ready supply of data maps available for an inner product space.

Any column map $W = [w_1, \ldots, w_n] \in L(\mathbb{F}^n, Y)$ into an inner product space Y provides the corresponding data map

$$W^{\mathrm{c}} : Y \mapsto \mathbb{F}^n : x \mapsto ({w_j}^{\mathrm{c}}x : j = 1{:}n),$$

called its **conjugate transpose** or **Hermitian**.

The terminology comes from the special case $Y = \mathbb{F}^m$. In that case, $W \in \mathbb{F}^{m\times n}$, and then W^{c} is, indeed, just the conjugate transpose of the matrix W since then $w_j = W(:, j)$, hence

$${w_j}^{\mathrm{c}}x = W(:, j)^{\mathrm{c}}x = \sum_k \overline{W(k, j)}x_k = \sum_k (W^{\mathrm{c}})(j, k)x_k = (W^{\mathrm{c}}x)_j.$$

Further, if $W \in L(\mathbb{F}^n, Y)$ and $A \in \mathbb{F}^{n\times m}$, then, with $WA = [u_j := \sum_k w_k A(k, j) : j = 1{:}n]$, one verifies that

$$((WA)^{\mathrm{c}}x)_j = {u_j}^{\mathrm{c}}x = \sum_k \overline{A(k, j)}{w_k}^{\mathrm{c}}x = \sum_k A^{\mathrm{c}}(j, k){w_k}^{\mathrm{c}}x = (A^{\mathrm{c}}(W^{\mathrm{c}}x))_j.$$

This proves

(6.2): If $W \in L(\mathbb{F}^n, Y)$ and $A \in \mathbb{F}^{n\times m}$, then $WA \in L(\mathbb{F}^m, Y)$ and $(WA)^{\mathrm{c}} = A^{\mathrm{c}}W^{\mathrm{c}}$.

This observation shows that *the above definition of the conjugate transpose of a column map is a special case of the abstract definition of the conjugate transpose of $A \in L(X, Y)$ as the unique map $A^{\mathrm{c}} : Y \to X$ (necessarily linear) for which*

$$\langle x, A^{\mathrm{c}}y\rangle = \langle Ax, y\rangle, \qquad \forall (x, y) \in X \times Y. \tag{6.3}$$

Indeed, if also $\langle x, z\rangle = \langle Ax, y\rangle$ for all $x \in X$, then $\langle x, z - A^c y\rangle = 0$ for all $x \in X$, including $x = z - A^c y$, hence, by the definiteness of the inner product, $z - A^c y = 0$, showing that $A^c y$ is uniquely determined by (6.3). With that, it is easy to see that A^c is a linear map, and that the conjugate transpose of an n-column map into Y is, indeed, the conjugate transpose in the sense of (6.3) (with $X = \mathbb{F}^n$), and that

$$(BA)^c = A^c B^c \tag{6.4}$$

in case BA makes sense, hence, in particular,

$$A^{-c} := (A^{-1})^c = (A^c)^{-1}. \tag{6.5}$$

The only fly in the ointment is the fact that, for some $A \in L(X, Y)$, there may not be any map $A^c : Y \to X$ satisfying (6.3) unless X is 'complete', a condition that is beyond the scope of these notes. However, if both X and Y are finite-dimensional inner-product spaces, then, with V and W bases for X and Y, respectively, we can write any $A \in L(X, Y)$ as $A = W\widehat{A}V^{-1}$ (using the *matrix* $\widehat{A} := W^{-1}AV$), hence, with (6.4), have available the formula

$$A^c = (W\widehat{A}V^{-1})^c = V^{-c}\widehat{A}^c W^c$$

for the conjugate transpose of A, – another nice illustration of the power of the basis concept.

With that, we are ready for the essential fact about the conjugate transpose needed now.

(6.6) Lemma: If the range of the 1-1 column map V is contained in the range of some column map W, then $W^c V$ is 1-1, i.e., W^c is 1-1 on $\operatorname{ran} V$.

Proof: Assume that $W^c V a = 0$ and let $b := Va$. Then $b \in \operatorname{ran} V \subset \operatorname{ran} W$, hence we must have $b = Wc$ for some vector c. Therefore, using (6.2),

$$0 = c^c 0 = c^c\, W^c V a = (Wc)^c V a = b^c b.$$

By the definiteness of the inner product, this implies that $b = 0$, i.e., $Va = 0$, therefore that $a = 0$, since V is assumed to be 1-1. □

By taking now, in particular, $W = V$, it follows that, for any basis V of the linear subspace X of the inner product space Y, the linear map $(V^c V)^{-1} V^c$ is well-defined, hence provides a formula for V^{-1}.

In `MATLAB`, the conjugate transpose of a matrix `A` is obtained as `A'`, hence the corresponding formula is `inv(V'*V)*V'`. It is, in effect, used there to carry out the operation `V\` for a matrix `V` that is merely 1-1.

□

Orthogonal projectors and closest points

We conclude that, with V a basis for the linear subspace X of the inner product space Y, the linear projector

$$P_V := V(V^c V)^{-1} V^c$$

is well-defined. Moreover, by (5.8), $\operatorname{null} P_V = \operatorname{null} V^c = \{y \in Y : V^c y = 0\}$. Since $x \in \operatorname{ran} P_V = \operatorname{ran} V$ is necessarily of the form $x = Va$, it follows that, for any $x \in \operatorname{ran} P_V$ and any $y \in \operatorname{null} P_V$,

$$x^c y = (Va)^c y = a^c (V^c y) = 0.$$

In other words, $\operatorname{ran} P_V$ *and* $\operatorname{null} P_V = \operatorname{ran}(\operatorname{id} - P_V)$ *are perpendicular or orthogonal to each other*, in the sense of the following definition.

Definition: We say that the elements u, v of the inner product space Y are **orthogonal** or **perpendicular** to each other, and write this

$$u \perp v,$$

in case $\langle u, v\rangle = 0$.

More generally, for any $F, G \subset Y$, we write $F \perp G$ to mean that, $\forall f \in F, g \in G,\ f \perp g$.

The **orthogonal complement**

$$F^\perp := \{y \in Y : y \perp F\}$$

of F is the largest set G perpendicular to F.

Note that $u \perp v$ iff $v \perp u$ since $\langle v, u\rangle = \overline{\langle u, v\rangle}$.

Because of the orthogonality

$$\operatorname{null} P_V = \operatorname{ran}(\operatorname{id} - P_V) \ \perp\ \operatorname{ran} P_V$$

just proved, P_V is called the **orthogonal** projector onto $\operatorname{ran} V$. Correspondingly, we write

$$Y = \operatorname{ran} P_V \oplus \operatorname{null} P_V \tag{6.7}$$

to stress the fact that, in this case, the summands in this direct sum are orthogonal to each other. Since they sum to Y, it follows (see H.P. 6.10 below) that each is the orthogonal complement of the other.

This orthogonality, as we show in a moment, has the wonderful consequence that, for any $y \in Y$, $P_V y$ is the unique element of $\operatorname{ran} P_V = \operatorname{ran} V$ that is closest to y in the sense of the **(Euclidean) norm**

$$\|\cdot\| : Y \to \mathbb{R} : y \mapsto \sqrt{y^{\mathrm{c}} y}. \tag{6.8}$$

Thus, whether or not $y \in Y$ is in $\operatorname{ran} V$, the coordinate vector $a = (V^{\mathrm{c}} V)^{-1} V^{\mathrm{c}} y$ supplied by our formula gives the coordinates of the point in $\operatorname{ran} V$ closest to y. If $y \in \operatorname{ran} V$, then this is, of course, y itself.

(6.9) Example: We continue with (5.9)Example. In that example, the choice $\Lambda^{\mathrm{t}} = V^{\mathrm{c}}$ amounts to choosing $w = v$. Now P becomes $P = vv^{\mathrm{c}}/v^{\mathrm{c}}v$, and, correspondingly,

$$Py = v\frac{v^{\mathrm{c}} y}{v^{\mathrm{c}} v},$$

which we recognize as the standard formula for the orthogonal projection of the vector y onto the line spanned by the vector v.

Correspondingly, (5.10)Figure changes to the following.

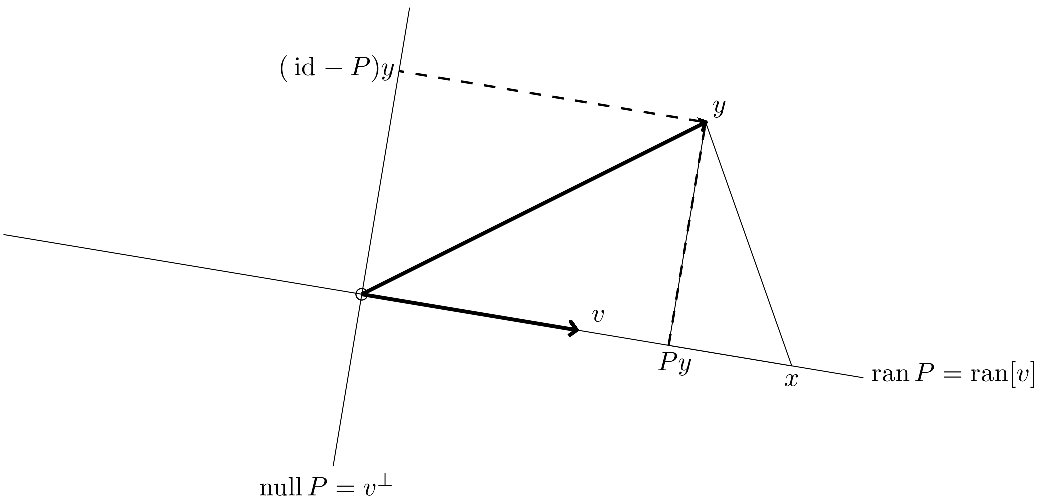

(6.10) Figure. If $y - Py$ is perpendicular to $\mathrm{ran}\,P$, then Py is the closest point to y from $\mathrm{ran}\,P$ since then, for any $x \in \mathrm{ran}\,P$, $\|y-x\|^2 = \|y-Py\|^2 + \|x-Py\|^2$.

□

The *proof* that, for any $y \in Y$, $P_V y$ is the unique element of $\mathrm{ran}\,V$ closest to y in the sense of the norm (6.8) is based on nothing more than the following little calculation.

(6.11) $$\|u+v\|^2 = (u+v)^{\mathrm{c}}(u+v) = \|u\|^2 + v^{\mathrm{c}}u + u^{\mathrm{c}}v + \|v\|^2.$$

Since $v^{\mathrm{c}}u = \overline{u^{\mathrm{c}}v}$, this proves

(6.12) Pythagoras: $u \perp v \quad \Longrightarrow \quad \|u+v\|^2 = \|u\|^2 + \|v\|^2.$

Since, for any $x \in X$, $y-x = (y-P_V y) + (P_V y - x)$, while $(y-P_V y) \in \mathrm{null}\,P_V \perp \mathrm{ran}\,P_V = X \ni (P_V y - x)$ we conclude that

(6.13) $$\|y-x\|^2 = \|y-P_V y\|^2 + \|P_V y - x\|^2.$$

Here, the first term on the right is *independent of* x. This shows that $\|y-x\|$ is uniquely minimized over $x \in X$ by the choice $x = P_V y$, as we claimed.

Here is the formal statement.

(6.14) Theorem: For any basis V for the linear subspace X of the inner product space Y, the linear map
$$P_V = V(V^{\mathrm{c}}V)^{-1}V^{\mathrm{c}}$$
equals P_X, the **orthogonal projector onto** X, in the sense that, for all $y \in Y$, $P_V y \in X$ and $y - P_V y \perp X$.

Therefore, Y is the **orthogonal direct sum**

$$Y = \operatorname{ran} V \oplus \operatorname{null} V^{c} = \operatorname{ran} P_V \oplus \operatorname{null} P_V = X \oplus \operatorname{ran}(\operatorname{id} - P_V),$$

and

$$\forall\{y \in Y,\ x \in X\}\quad \|y - x\| \ge \|y - P_V y\|,$$

with equality if and only if $x = P_V y$.

Incidentally, by choosing $x = 0$ in (6.13), – legitimate since $\operatorname{ran} V$ is a linear subspace, – we find the following very useful fact.

(6.15) Proposition: For any 1-1 column map V into Y and any $y \in Y$,

$$\|y\| \ge \|P_V y\|,$$

with equality if and only if $y = P_V y$, i.e., if and only if $y \in \operatorname{ran} V$.

This says that P_V *strictly reduces norms*, except for those elements that it doesn't change at all.

6.1 Construct the orthogonal projection of the vector $(1, 1, 1)$ onto the line $L = \operatorname{ran}[1; -1; 1]$.

6.2 Construct the orthogonal projection of the vector $x := (1, 1, 1)$ onto the straight line $y + \operatorname{ran}[v]$, with $y = (2, 0, 1)$ and $v = (1, -1, 1)$. (Hint: you want to minimize $\|x - (y + \alpha v)\|$ over all $\alpha \in \mathbb{R}$.)

6.3 Compute the distance between the two straight lines $y + \operatorname{ran}[v]$ and $z + \operatorname{ran}[w]$, with $y = (2, 0, 1)$, $v = (1, 1, 1)$, $z = (-1, 1, -1)$ and $w = (0, 1, 1)$. (Hint: you want to minimize $\|y + \alpha v - (z + \beta w)\|$ over α, β.)

6.4 With $v_1 = (1, 2, 2)$, $v_2 = (-2, 2, -1)$, (a) construct the matrix that provides the orthogonal projection onto the subspace $\operatorname{ran}[v_1, v_2]$ of $\mathbb{R}^3$; (b) compute the orthogonal projection of the vector $y = (1, 1, 1)$ onto $\operatorname{ran}[v_1, v_2]$.

6.5 Taking for granted that the space $Y := C([-1\,..\,1])$ of real-valued continuous functions on the interval $[-1\,..\,1]$ is an inner product space with respect to the inner product

$$\langle f, g\rangle := \int_{-1}^{1} f(t)g(t)\,\mathrm{d}t,$$

do the following: (a) Construct (a formula for) the orthogonal projector onto $X := \Pi_1$, using the power basis, $V = [()^0, ()^1]$ for X. (b) Use your formula to compute the orthogonal projection of $()^2$ onto Π_1.

6.6 (a) Prove: *If $\mathbb{F} = \mathbb{R}$, then $u \perp v$ if and only if $\|u + v\|^2 = \|u\|^2 + \|v\|^2$.* (b) What goes wrong with your argument when $\mathbb{F} = \mathbb{C}$?

6.7 For each of the following maps $f : \mathbb{F}^n \times \mathbb{F}^n \to \mathbb{F}$, determine whether or not it is an inner product.

(a) $\mathbb{F} = \mathbb{R}$, $n = 3$, and $f(x, y) = x_1y_1 + x_3y_3$; (b) $\mathbb{F} = \mathbb{R}$, $n = 3$, and $f(x, y) = x_1y_1 - x_2y_2 + x_3y_3$; (c) $\mathbb{F} = \mathbb{R}$, $n = 2$, and $f(x, y) = x_1^2 + y_1^2 + x_2y_2$; (d) $\mathbb{F} = \mathbb{C}$, $n = 3$, and $f(x, y) = x_1y_1 + x_2y_2 + x_3y_3$; (e) $\mathbb{F} = \mathbb{R}$, $n = 3$, and $f(x, y) = x_1y_2 + x_2y_3 + x_3y_1$;

6.8 Prove that, for any invertible $A \in \mathbb{F}^{n\times n}$, $\langle\cdot,\cdot\rangle : \mathbb{F}^n \times \mathbb{F}^n \to \mathbb{F} : (x, y) \mapsto (Ay)^{c}Ax = y^{c}(A^{c}A)x$ is an inner product on $\mathbb{F}^n$.

6.9 Prove that, for any subset F of the inner product space Y, the orthogonal complement $F^{\perp}$ is a linear subspace. (Hint: $F^{\perp} = \cap_{f\in F} \operatorname{null} f^{c}$.)

6.10 Prove that, whenever $Y = X \oplus Z$, then $X^{\perp} = Z$ and $Z^{\perp} = X$.

6.11 Prove that, *for any linear subspace X of a finite-dimensional inner product space Y, $(\operatorname{id} - P_X) = P_{X^{\perp}}$.*

6.12 Prove that, *for any finite-dimensional linear subspace X of an inner product space Y, $(X^{\perp})^{\perp} = X$.*

6.13 An **isometry** or **rigid motion** in an inner product space X is any map $f : X \to X$ that preserves distances, i.e., for which $\|f(x) - f(y)\| = \|x - y\|$ for all $x, y \in X$. Prove that any rigid motion of $\mathbb{R}^2$ that maps the origin to itself is necessarily a linear map. (Hint: you might prove first that, for any $x \ne y$ and any $\alpha \in \mathbb{R}$, the point $(1 - \alpha)x + \alpha y$ is the unique point in the plane whose distance from x is $|\alpha|\|y - x\|$ and from y is $|1 - \alpha|\|y - x\|$.)

Least-squares

Note that $P_V y = Va$, with the coefficient vector a the unique solution to the linear equation

$$V^c Va = V^c y.$$

This equation is also referred to as the **normal equation** since it requires that $V^c(y - Va) = 0$, i.e., that the residual, $y - Va$, be perpendicular or *normal* to every column of V, hence to all of $\operatorname{ran} V$ (see (6.10)Figure). In effect, given that the equation $V? = y$ doesn't have a solution for $y \in Y \backslash X$, our particular $Va = P_V y$ gives us the closest thing to a solution.

In particular, if $y \in Y = \mathbb{R}^n$ and $V \in \mathbb{R}^{n\times r}$ is 1-1, then $P_V y$ minimizes $\|y - Va\|$ over all $a \in \mathbb{R}^r$. For that reason, the coefficient vector $a := V^{-1}P_V y$ is called the **least-squares solution** to the (usually inconsistent or overdetermined) linear system $V? = y$.

In MATLAB, the vector $P_V y$ is computed as `V*(V\y)`, in line with the fact mentioned earlier that the action of the matrix $(V^c V)^{-1}V^c$ is provided by the operator `V\`, i.e., (up to roundoff and for any vector `y`) the three vectors

```
a1 = V\y, a2 = inv(V'*V)*V'*y, a3 = (V'*V)\(V'*y)
```

are all the same. However, the first way is preferable since it avoids actually forming the matrix `V'*V` (or its inverse) and, therefore, is less prone to roundoff effects.

□

A practically very important special case of this occurs when $X = \operatorname{ran} V$ consists of functions on some domain T and, for some finite subset S of T,

$$Q_S : X \to \mathbb{R}^S : f \mapsto (f(s) : s \in S)$$

is 1-1. Then

$$(6.16)\qquad \langle f, g\rangle_S := \sum_{s\in S} f(s)g(s) =: (Q_S f)^t(Q_S g)$$

is an inner product on X since it is evidently linear in the first argument and also hermitian and nonnegative, and is definite since $\langle f, f\rangle_S = 0$ implies $Q_S f = 0$, hence $f = 0$ since Q_S is 1-1. Then, for arbitrary $g \in \mathbb{R}^S$, we can compute

$$V^c g := (\langle g, v\rangle_S : v \in V) = (Q_S V)^t g,$$

hence can construct

$$P_{V,S}g := V(V^c V)^{-1}V^c g$$

as the unique element Va of $\operatorname{ran} V$ closest to g in the sense that the sum of squares $\sum_{s\in S} |g(s) - (Va)(s)|^2$ is as small as possible. For this reason, $P_{V,S}g$ is also called the **discrete least-squares approximation** from $\operatorname{ran} V$ to g, or, more explicitly, to the data $((s, g(s)) : s \in S)$. If $\#V = \#S$, then $P_{V,S}g$ is the unique interpolant to these data from $\operatorname{ran} V$.

In any calculation of such a discrete least-squares approximation, we would, of course, have to list the elements of S in some fashion, say as the entries s_j of the sequence $s := (s_1, \ldots, s_n)$. Then we can think of Q_S as the data map into $\mathbb{R}^n$ given by $f \mapsto (f(s_j) : j = 1{:}n)$. Correspondingly, $Q_S V$ becomes an $n \times r$-matrix, and this matrix is 1-1, by the assumption that Q_S is 1-1 on $X = \operatorname{ran} V$. Further, the coefficient vector $a := (V^c V)^{-1}V^c g$ for $P_{V,S}g$ is the least-squares solution to the linear equation

$$Q_S V? = g$$

which seeks a coefficient vector a so that Va interpolates to the data $((s_j, g(s_j)) : j = 1{:}n)$. Such an interpolant exists if and only if the matrix $Q_S V$ is invertible. Otherwise, one has to be content with a least-squares solution, i.e., a discrete least-squares approximation to these data, from $\operatorname{ran} V$.

6.14 Compute the discrete least squares approximation by straight lines (i.e., from Π_1) to the data (j, j^2), $j = 1{:}10$ using (a) the basis $[()^0, ()^1]$; (b) the basis $[()^0, ()^1 - 5.5()^0]$. (c) Why might one prefer (b) to (a)?

Orthonormal column maps

The formula

$$P_V = V(V^{c}V)^{-1}V^{c}$$

for the orthogonal projector onto the range of the 1-1 column map V becomes particularly simple in case

$$V^{c}V = \mathrm{id}; \tag{6.17}$$

it then reduces to

$$P_V = VV^{c}.$$

We call V **orthonormal** (or, **o.n.**, for short) in this case since, written out entry by entry, (6.17) reads

$$\langle v_j, v_k\rangle = \left\{\begin{matrix} 1 & \text{if } j = k; \\ 0 & \text{otherwise,}\end{matrix}\right\} =: \delta_{jk}.$$

In other words, each column of V is *normalized*, meaning that it has norm 1, and different columns are orthogonal to each other. Such bases are special in that they provide their own inverse, i.e.,

$$x = V(V^{c}x), \quad \forall x \in \operatorname{ran} V.$$

The term 'orthonormal' can be confusing, given that earlier we mentioned the normal equation, $V^{c}V? = V^{c}y$, socalled because it expresses the condition that the residual, $y - Va$, be orthogonal or 'normal' to the columns of V. In fact, *norma* is the Latin name for a mason's tool for checking that a wall is at right angles to the ground. In the same way, the *normal* to a surface at a point is a vector at right angles to the surface at that point. Nevertheless, to **normalize** the vector y does not mean to change it into a vector that is perpendicular to some subspace or set. Rather, it means to divide it by its norm, thereby obtaining the vector $y/\|y\|$ that points in the same direction as y but has norm 1. To be sure, this can only be done for $y \neq 0$ and then $\|\, y/\|y\|\,\| = 1$ because the Euclidean norm is **absolutely homogeneous**, meaning that

$$\|\alpha y\| = |\alpha|\|y\|, \quad \forall (\alpha, y) \in \mathbb{F} \times Y. \tag{6.18}$$

We now show that every finite-dimensional linear subspace of an inner-product space Y has o.n. bases.

(6.19) Proposition: For every 1-1 $V \in L(\mathbb{F}^n, Y)$, there exists an o.n. $Q \in L(\mathbb{F}^n, Y)$ so that, for all j, $\operatorname{ran}[q_1, q_2, \ldots, q_j] = \operatorname{ran}[v_1, v_2, \ldots, v_j]$, hence $V = QR$ with R (invertible and) upper triangular, a **QR factorization** for V.

Proof: For $j = 1{:}n$, define $u_j := v_j - P_{V_{<j}} v_j$, with $V_{<j} := V_{j-1} := [v_1, \ldots, v_{j-1}]$. By (6.14)Theorem, $u_j \perp \operatorname{ran} V_{<j}$, all j, hence $u_j \perp u_k$ for $j \neq k$. Also, each u_j is nonzero (since $u_j = V_j(a, 1)$ for some $a \in \mathbb{F}^{j-1}$, and V_j is 1-1), hence $q_j := u_j/\|u_j\|$ is well-defined and, still, $q_j \perp q_k$ for $j \neq k$.

It follows that $Q := [q_1, \ldots, q_n]$ is o.n., hence, in particular, 1-1. Finally, since $q_j = u_j/\|u_j\| \in \operatorname{ran} V_j$, it follows that, for each j, the 1-1 map $[q_j, \ldots, q_j]$ has its range in the j-dimensional space $\operatorname{ran} V_j$, hence must be a basis for it. □

Since $Q_{<j} = [q_1, \ldots, q_{j-1}]$ is an o.n. basis for $\operatorname{ran} V_{<j}$, it is of help in constructing q_j since it gives

$$u_j = v_j - P_{V_{<j}} v_j, \quad \text{with} \quad P_{V_{<j}} v_j = P_{Q_{<j}} v_j = \sum_{k<j} q_k \langle v_j, q_k \rangle = \sum_{k<j} u_k \frac{\langle v_j, u_k \rangle}{\langle u_k, u_k \rangle}. \tag{6.20}$$

For this reason, it is customary to construct the u_j or the q_j's one by one, from the first to the last, using (6.20). This process is called **Gram-Schmidt orthogonalization**. To be sure, as (6.20) shows, there is no real need (other than neatness) to compute the q_j from the u_j and, by skipping the calculation of q_j, one avoids taking square-roots.

Since any 1-1 column map into a finite-dimensional vector space can be extended to a basis for that vector space, we have also proved the following.

(6.21) Corollary: Every o.n. column map Q into a finite-dimensional inner product space can be extended to an o.n. basis for that space.

Given any 1-1 matrix `V`, the `MATLAB` command `[q,r] = qr(V,0)` provides an o.n. basis, `q`, for ran `V`, along with the upper triangular matrix `r` for which `q*r` equals `V`. The (simpler) statement `[Q,R]=qr(V)` provides a *unitary*, i.e., a *square* o.n., matrix `Q` and an upper triangular matrix `R` so that `Q*R` equals `V`. If `V` is itself square, then `q` equals `Q`. In the contrary case, `Q` equals `[q,U]` for some o.n. basis `U` of the orthogonal complement of ran`V`. Finally, the simplest possible statement, `p = qr(V)`, gives the most complicated result, namely a matrix `p` of the same size as `V` that contains `r` in its upper triangular part and complete information about the various Householder matrices used in its strictly lower triangular part.

While, for each $j = 1{:}\#$`V`, ran`V`$(:, [1{:}j])$ = ran`Q`$(:, [1{:}j])$, the construction of `q` or `Q` does not involve the Gram-Schmidt algorithm, as that algorithm is not reliable numerically when applied to an arbitrary 1-1 matrix `V`. Rather, the matrix `V` is factored column by column with the aid of certain elementary matrices, the so-called **Householder reflection**s $\mathrm{id} - 2ww^c/w^c w$.

□

As already observed, it is customary to call a square o.n. matrix **unitary**. It is also customary to call a *real* unitary matrix **orthogonal**. However, the columns of such an 'orthogonal matrix' are not just orthogonal to each other, they are also normalized. Thus it would be better to call such a matrix 'orthonormal', freeing the term 'orthogonal matrix' to denote one whose columns are merely orthogonal to each other. But such naming conventions are hard to change. I will simply not use the term 'orthogonal matrix', but use 'real unitary matrix' instead.

An o.n. column map Q has many special properties, all of which derive from the defining property, $Q^c Q = \mathrm{id}$, by the observation that therefore, for any $a, b \in \mathbb{F}^n$,

$$\langle Qa, Qb \rangle \;=\; \langle Q^c Q a, b \rangle = \langle a, b \rangle. \tag{6.22}$$

This says that Q is **inner-product preserving**. In particular, any o.n. $Q \in L(\mathbb{F}^n, X)$ is an **isometry** in the sense that

$$\forall a \in \mathbb{F}^n \quad \|Qa\| = \|a\|. \tag{6.23}$$

More than that, any o.n. $Q \in L(\mathbb{F}^n, X)$ is **angle-preserving** since a standard definition of the **angle** φ between two real nonzero n-vectors x and y is the following implicit one:

$$\cos(\varphi) := \frac{\langle x, y \rangle}{\|x\| \|y\|}.$$

To be sure, this definition makes sense only if we can be sure that the righthand side lies in the interval $[-1 \,..\, 1]$. But this is a consequence of the

Cauchy-Bunyakovski-Schwarz or CBS Inequality: For any u, v in the inner product space Y,

$$|\langle u, v\rangle| = |v^c u| \le \|u\|\|v\|, \tag{6.24}$$

with equality if and only if $[u, v]$ is not 1-1.

Be sure to remember not only the *in*equality, but also exactly when it is an *equal*ity.

Proof: If $v = 0$, then there is equality in (6.24) and $[u, v]$ is not 1-1. Otherwise, $v \ne 0$ and, in that case, by (6.15)Proposition, the orthogonal projection $P_{[v]}u = v(v^c u)/\|v\|^2$ onto $\operatorname{ran}[v]$ of an arbitrary $u \in Y$ has norm smaller than $\|u\|$ unless $u = P_{[v]}u$. In other words, $|v^c u|/\|v\| = \|v(v^c u)/\|v\|^2\| \le \|u\|$, showing that (6.24) holds in this case, with equality if and only if $u \in \operatorname{ran}[v]$. □

6.15 Prove (6.18).

6.16 Prove that $V = \begin{bmatrix} 1 & 1 & 1 \\ -1 & 1 & -1 \\ 0 & 1 & 2 \end{bmatrix}$ is a basis for $\mathbb{R}^3$ and compute the coordinates of $x := (1, 1, 1)$ with respect to V.

6.17 Verify that $V = \begin{bmatrix} 1 & -1 & 1 \\ 1 & -1 & -1 \\ 1 & 2 & 0 \\ 0 & 0 & 2 \end{bmatrix}$ is an orthogonal basis for its range, and extend it to an orthogonal basis for $\mathbb{R}^4$.

6.18 (a) Use the calculations in H.P. 6.14 to construct an orthogonal basis for Π_2 from the power basis $V = [()^0, ()^1, ()^2]$ with respect to the (discrete) inner product in H.P. 6.14 .

(b) Use (a) to compute the discrete least-squares approximation from Π_2 to the data (j, j^3), $j = 1{:}10$.

6.19 Use the result of H.P. 6.5 to construct an o.n. basis for Π_2 wrto the inner product $\langle f, g\rangle := \int_{-1}^{1} f(t)g(t)\,\mathrm{d}t$.

6.20 What is the angle between $(1, 2, 2)$ and $(3, -1, -2)$?

6.21 Consider the **Vandermonde** matrix

$$A := [\delta_{z_0}, \ldots, \delta_{z_k}]^c[()^0, \ldots, ()^k] = (z_i^j : i, j = 0{:}k)$$

for some sequence $z_0, \ldots, z_k$ of complex numbers.

Prove that A is a scalar multiple of a unitary matrix if and only if, after some reordering and for some real α,

$$\{z_0, \ldots, z_k\} = \{\exp(2\pi \mathrm{i}(\alpha + i/(k+1))) : i = 0{:}k\}.$$

$\operatorname{ran} A$ and $\operatorname{null} A^c$ form an orthogonal direct sum for $\operatorname{tar} A$

The two basic linear subspaces associated with $A \in L(X, Y)$ are its range, $\operatorname{ran} A$, and its kernel or nullspace, $\operatorname{null} A$. However, when X and Y are inner product spaces, it is also very useful to consider the range of A and the nullspace of the (conjugate) transpose A^c of A together. For, then, by the definiteness of the inner product, $A^c y = 0$ iff $\langle x, A^c y\rangle = 0$ for all $x \in X$, while, by (6.3), $\langle x, A^c y\rangle = \langle Ax, y\rangle$, hence

$$\operatorname{null} A^c = \{y \in Y : y \perp \operatorname{ran} A\}.$$

Recalling the notation

$$M^\perp := \{y \in Y : y \perp M\}$$

for the it orthogonal complement of the subset M of Y, we get the following.

(6.25) Proposition: For any $A \in L(X, Y)$, $(\operatorname{ran} A)^\perp = \operatorname{null} A^c$.

(6.26) Corollary: For any $A \in L(X,Y)$, Y is the *orthogonal* direct sum $Y = \operatorname{ran} A \oplus \operatorname{null} A^c$. Hence

$$\dim \operatorname{tar} A = \dim \operatorname{ran} A + \dim \operatorname{null} A^c.$$

Proof: Let V be any basis for $\operatorname{ran} A$. By (6.14)Theorem,

$$Y = \operatorname{ran} V \oplus \operatorname{null} V^c,$$

while, by choice of V, $\operatorname{ran} V = \operatorname{ran} A$, and so, by (6.25), $\operatorname{null} V^c = (\operatorname{ran} V)^\perp = (\operatorname{ran} A)^\perp = \operatorname{null} A^c$. □

In particular, A is onto if and only if A^c is 1-1. Further, since $(A^c)^c = A$, we also have the following complementary statement.

(6.27) Corollary: For any $A \in L(X,Y)$, X is the *orthogonal* direct sum $X = \operatorname{ran} A^c \oplus \operatorname{null} A$. Hence,

$$\dim \operatorname{dom} A = \dim \operatorname{ran} A^c + \dim \operatorname{null} A.$$

In particular, A^c is onto if and only if A is 1-1. Also, on comparing (6.27) with the Dimension Formula, we see that $\dim \operatorname{ran} A = \dim \operatorname{ran} A^c$.

The fact (see (6.26)Corollary) that $\operatorname{tar} A = \operatorname{ran} A \oplus \operatorname{null} A^c$ is often used as a characterization of the elements $y \in \operatorname{tar} A$ for which the equation $A? = y$ has a solution. For, it says that $y \in \operatorname{ran} A$ if and only if $y \perp \operatorname{null} A^c$. Of course, since $\operatorname{null} A^c$ consists exactly of those vectors that are orthogonal to all the columns of A, this is just a special case of the fact (see H.P. 6.12) that the orthogonal complement of the orthogonal complement of a linear subspace is that linear subspace itself.

The inner product space $\mathbb{F}^{m\times n}$ and the trace of a matrix

At the outset of these notes, we introduced the space $\mathbb{F}^{m\times n}$ as a special case of the space $\mathbb{F}^T$ of all scalar-valued functions on some set T, namely with

$$T = \underline{m} \times \underline{n}.$$

This set being finite, there is a natural inner product on $\mathbb{F}^{m\times n}$, namely

$$\langle A, B\rangle := \sum_{i,j} \overline{B(i,j)} A(i,j).$$

This inner product can also be written in the form

$$\langle A, B\rangle = \sum_{i,j} B^c(j,i)A(i,j) = \sum_j (B^c A)(j,j) = \operatorname{trace}(B^c A).$$

Here, the **trace** of a square matrix C is, by definition, the sum of its diagonal entries,

$$\operatorname{trace} C := \sum_j C(j,j).$$

The norm in this inner product space is called the **Frobenius norm**,

(6.28) $$\|A\|_F := \sqrt{\operatorname{trace} A^c A} = \sum_{i,j} |A(i,j)|^2.$$

The Frobenius norm is **compatible** with the Euclidean norm $\| \, \|$ on $\mathbb{F}^n$ and $\mathbb{F}^m$ in the sense that

$$\|Ax\| \le \|A\|_F \|x\|, \quad x \in \mathbb{F}^n. \tag{6.29}$$

Not surprisingly, the map $\mathbb{F}^{m\times n} \to \mathbb{F}^{n\times m} : A \mapsto A^{\mathrm{t}}$ is unitary, i.e., inner-product preserving:

$$\langle A^{\mathrm{t}}, B^{\mathrm{t}} \rangle = \sum_{i,j} \overline{B^{\mathrm{t}}(i,j)} A^{\mathrm{t}}(i,j) = \sum_{i,j} \overline{B(j,i)} A(j,i) = \langle A, B \rangle. \tag{6.30}$$

In particular,

$$\operatorname{trace}(B^{\mathrm{c}} A) = \operatorname{trace}(A B^{\mathrm{c}}), \quad A, B \in \mathbb{F}^{m\times n}. \tag{6.31}$$

6.22 T/F

(a) $(x,y) \mapsto y^{\mathrm{c}} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} x$ is an inner product on $\mathbb{R}^2$.

(b) $\|x+y\|^2 \le \|x\|^2 + \|y\|^2$;

7. Norms, map norms, and the condition of a basis

Assume that V is a basis for the nontrivial linear subspace X of the inner product space Y. The coordinate vector a for $x \in X$ is the unique solution of the equation

$$V? = x.$$

We may not be able to compute the solution exactly. Even if we know the entries of the solution exactly, as common fractions say, we may not be able to use them exactly if we use some floating-point arithmetic, as is common. It is for this reason that one is interested in gauging the effect of an erroneous coordinate vector $\widehat{a}$ on the accuracy of $V\widehat{a}$ as a representation for $x = Va$.

How to judge the error by the residual

Since, presumably, we do not know a, we cannot compute the **error**

$$\varepsilon := a - \widehat{a};$$

we can only compute the **residual**

$$r := x - V\widehat{a}.$$

Nevertheless, can we judge the error by the residual? Does a 'small' **relative residual**

$$\|r\|/\|x\|$$

imply a 'small' **relative error**

$$\|\varepsilon\|/\|a\| \ ?$$

By definition, the **condition** (or, **condition number**) $\kappa(V)$ of the basis V is the greatest factor by which the relative error, $\|\varepsilon\|/\|a\|$, can exceed the relative residual, $\|r\|/\|x\| = \|V\varepsilon\|/\|Va\|$; i.e.,

$$\kappa(V) := \sup_{a,\varepsilon} \frac{\|\varepsilon\|/\|a\|}{\|V\varepsilon\|/\|Va\|}. \tag{7.1}$$

However, by interchanging here the roles of a and ε and then taking reciprocals, this also says that

$$1/\kappa(V) = \inf_{\varepsilon,a} \frac{\|\varepsilon\|/\|a\|}{\|V\varepsilon\|/\|Va\|}.$$

Hence, altogether,

$$\frac{1}{\kappa(V)} \frac{\|r\|}{\|x\|} \ \le \ \frac{\|\varepsilon\|}{\|a\|} \ \le \ \kappa(V) \frac{\|r\|}{\|x\|}. \tag{7.2}$$

In other words, *the larger the condition number, the less information about the size of the relative error is provided by the size of the relative residual.*

For a better feel for the condition number, note that we can also write the formula (7.1) for $\kappa(V)$ in the following fashion:

$$\kappa(V) = \sup_{\varepsilon} \frac{\|\varepsilon\|}{\|V\varepsilon\|} \sup_{a} \frac{\|Va\|}{\|a\|}.$$

Also,

$$\|Va\|/\|a\| = \|V(a/\|a\|)\|,$$

with $a/\|a\|$ *normalized*, i.e., of norm 1. Hence, altogether,

$$\kappa(V) = \frac{\sup\{\|Va\| : \|a\| = 1\}}{\inf\{\|Va\| : \|a\| = 1\}}. \tag{7.3}$$

This says that we can visualize the condition number $\kappa(V)$ in the following way; see (7.5)Figure. Consider the image

$$\{Va : \|a\| = 1\} \tag{7.4}$$

under V of the **unit sphere**

$$\{a \in \mathbb{F}^n : \|a\| = 1\}$$

in $\mathbb{F}^n$. It will be some kind of ellipsoid, symmetric with respect to the origin. In particular, there will be a point $a_{\max}$ with $\|a_{\max}\| = 1$ for which $Va_{\max}$ will be as far from the origin as possible. There will also be a point $a_{\min}$ with $\|a_{\min}\| = 1$ for which $Va_{\min}$ will be as close to the origin as possible. In other words,

$$\kappa(V) = \|Va_{\max}\|/\|Va_{\min}\|,$$

saying that the condition number gives the ratio of the largest to the smallest diameter of the ellipsoid (7.4). The larger the condition number, the skinnier is the ellipsoid.

In particular, if $a = a_{\max}$ while $\varepsilon = a_{\min}$, then the relative error is 1 while the relative residual is $\|Va_{\min}\|/\|Va_{\max}\|$, and this is tiny to the extent that the ellipsoid is 'skinny'.

On the other hand, if $a = a_{\min}$ while $\varepsilon = a_{\max}$, then the relative error is still 1, but now the relative residual is $\|Va_{\max}\|/\|Va_{\min}\|$, and this is large to the extent that the ellipsoid is 'skinny'.

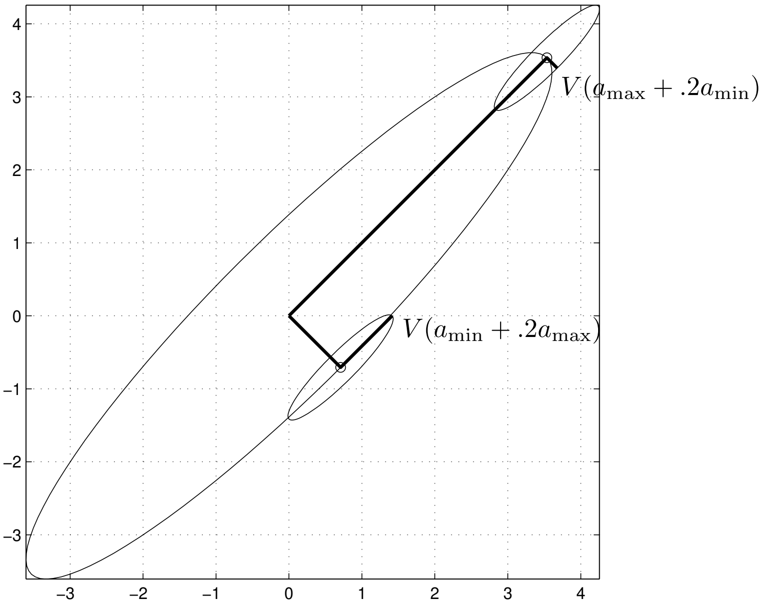

(7.5) Figure. Extreme effects of a 20% relative error on the relative residual, for $V = \begin{bmatrix} 3 & 2 \\ 2 & 3 \end{bmatrix}$.

The worst-conditioned column maps V are those that fail to be 1-1 since, for them, $Va_{\min} = 0$, hence $\kappa(V) = \infty$.

On the other extreme, it follows directly from (7.3) that $\kappa(V) \geq 1$, and this lower bound is reached by any o.n. basis V since any o.n. basis is an isometry, by (6.23), i.e., $\|Va\| = \|a\|$ for all $a \in \mathbb{F}^n$. Thus o.n. bases are best-conditioned, and rightfully prized for that. It was for this reason that we took the trouble to prove that every finite-dimensional linear subspace of an inner product space has o.n. bases, and even discussed just how to construct such bases.

The map norm

As we now explain, the numbers $\|Va_{\max}\| = \max\{\|Va\| : \|a\| = 1\}$ and $1/\|Va_{\min}\| = 1/\min\{\|Va\| : \|a\| = 1\}$ both are examples of a map norm according to the following

(7.6) Definition: The **map norm**, $\|A\|$, of $A \in L(X, Y)$ is the smallest nonnegative number c for which

$$\|Ax\| \le c\|x\|, \quad \forall x \in X.$$

If X is trivial, then $\|A\| = 0$ for the sole $A \in L(X, Y)$. Otherwise

$$\|A\| = \sup_{x \ne 0} \|Ax\|/\|x\| = \sup\{\|Ax\| : \|x\| = 1\}. \tag{7.7}$$

Here, the last equality follows from the absolute homogeneity of the norm and the homogeneity of A which combine to permit the conclusions that

$$\|Ax\|/\|x\| = \|A(x/\|x\|)\| \quad \text{and} \quad \|(x/\|x\|)\| = 1.$$

In these notes, we are only interested in *finite-dimensional* X and, for such X,

$$\|A\| = \max_{x \ne 0} \|Ax\|/\|x\| = \max\{\|Ax\| : \|x\| = 1\}. \tag{7.8}$$

The reason for this is beyond the scope of these notes, but is now stated for the record: If X is finite-dimensional, then

$$F : x \mapsto \|Ax\|$$

is continuous and the *unit sphere*

$$\{x \in X : \|x\| = 1\}$$

is compact, hence F achieves its maximum value on that sphere. (For the same reason, F also achieves its minimum value on the unit sphere, and this justifies the existence of $a_{\max}$ and $a_{\min}$ in the preceding section.)

We conclude that *determination* of the map norm is a two-part process, as formalized in the following.

(7.9) Calculation of $\|A\|$: The number c equals the norm $\|A\|$ if and only if

(i) for all x, $\|Ax\| \le c\|x\|$; and

(ii) for some $x \ne 0$, $\|Ax\| \ge c\|x\|$.

The first says that $\|A\| \le c$, while second says that $\|A\| \ge c$, hence, together they say that $\|A\| = c$.

(7.10) Example: We compute $\|A\|$ in case $A \in \mathbb{F}^{m \times n}$ is of the simple form

$$A = [v][w]^c = vw^c$$

for some $v \in \mathbb{F}^m$ and some $w \in \mathbb{F}^n$. Since

$$Ax = (vw^c)x = v(w^c x),$$

we have

$$\|(vw^c)x\| = \|v\||w^c x| \le \|v\|\|w\|\|x\|,$$

the equality by the absolute homogeneity of the norm, and the inequality by (6.24)Cauchy's Inequality. This shows that $\|vw^c\| \le \|v\|\|w\|$. On the other hand, for the specific choice $x = w$, we get $(vw^c)w = v(w^c w) = v\|w\|^2$, hence $\|(vw^c)w\| = \|v\|\|w\|\|w\|$. Assuming that $w \ne 0$, this shows that $\|vw^c\| \ge \|v\|\|w\|$. However, this inequality is trivially true in case $w = 0$ since then $vw^c = 0$. So, altogether, we have that

$$\|vw^c\| = \|v\|\|w\|.$$

Note that we have, incidentally, proved that, for any $v \in \mathbb{F}^n$,

$$\|[v]\| = \|v\| = \|[v]^c\|. \tag{7.11}$$

□

As another example, note that, if also $B \in L(Y, Z)$ for some inner product space Z, then BA is defined and

$$\|(BA)x\| = \|B(Ax)\| \le \|B\|\,\|Ax\| \le \|B\|\,\|A\|\|x\|.$$

Therefore,

$$\|BA\| \le \|B\|\,\|A\|. \tag{7.12}$$

We are ready to discuss the condition (7.3) of a basis V in terms of map norms.

Directly from (7.8), $\max\{\|Va\| : \|a\| = 1\} = \|V\|$.

(7.13) Proposition: If $A \in L(X, Y)$ is invertible and $X \ne \{0\}$ is finite-dimensional, then

$$\|A^{-1}\| = 1/\min\{\|Ax\| : \|x\| = 1\}.$$

Proof: Since A is invertible, $y \in Y$ is nonzero if and only if $y = Ax$ for some nonzero $x \in X$. Hence,

$$\|A^{-1}\| = \max_{y \ne 0} \frac{\|A^{-1}y\|}{\|y\|} = \max_{x \ne 0} \frac{\|A^{-1}Ax\|}{\|Ax\|} = 1/\min_{x \ne 0} \frac{\|Ax\|}{\|x\|},$$

and this equals $1/\min\{\|Ax\| : \|x\| = 1\}$ by the absolute homogeneity of the norm and the homogeneity of A. □

In particular, $1/\|A^{-1}\|$ is the largest number c for which

$$c\|x\| \le \|Ax\|, \quad \forall x \in X.$$

We conclude that

$$\kappa(V) = \|V\|\|V^{-1}\|. \tag{7.14}$$

7.1 Complement (7.13)Proposition by discussing the situation when $X = \{0\}$.

7.2 Prove that $\kappa(V) \ge 1$ for any basis V with at least one column.

7.3 Determine $\kappa([\,])$.

Vector norms and their associated map norms

`MATLAB` provides the map norm of the matrix `A` by the statement `norm(A)` (or by the statement `norm(A,2)`, indicating that there are other map norms available).

The `norm` command gives the Euclidean norm when its argument is a 'vector'. Specifically, `norm(v)` and `norm(v,2)` both give $\|v\| = \sqrt{v^{c}v}$. However, since in (present-day) `MATLAB`, everything is a matrix, there is room here for confusion since experimentation shows that `MATLAB` defines a 'vector' to be any 1-column matrix and any 1-row matrix. Fortunately, there is no problem with this, since, by (7.11), the norm of the *vector* v equals the norm of the *matrices* v and v^{c}.

□

The best explicit expression available for $\|A\|$ for an arbitrary $A \in \mathbb{F}^{m\times n}$ is the following:

$$\|A\| = \sqrt{\rho(A^{\mathrm{c}}A)} = \sigma_1(A). \tag{7.15}$$

This formula cannot be evaluated in finitely many steps since the number $\rho(A^{\mathrm{c}}A)$ is, by definition, the 'spectral radius' of $A^{\mathrm{c}}A$, i.e., the smallest possible radius of a disk centered at the origin that contains all the eigenvalues of $A^{\mathrm{c}}A$. The 2-norm of A also equals $\sigma_1(A)$ which is, by definition, the largest 'singular value' of A. In general, one can only compute approximations to this number.

For this reason (and others), other vector norms are in common use, among them the **max-norm**

$$\|x\|_\infty := \max_j |x_j|, \quad \forall x \in \mathbb{F}^n,$$

for which the associated map norm is easily computable. It is

$$\|A\|_\infty := \max_{x\neq 0} \|Ax\|_\infty/\|x\|_\infty = \max_i \sum_j |A(i,j)| = \max_i \|A(i,:)\|_1, \tag{7.16}$$

with

$$\|v\|_1 := \sum_j |v_j| \tag{7.17}$$

yet another vector norm, the socalled **1-norm**. The map norm associated with the 1-norm is also easily computable. It is

$$\|A\|_1 := \max_{x\neq 0} \|Ax\|_1/\|x\|_1 = \max_j \sum_i |A(i,j)| = \max_j \|A(:,j)\|_1 = \|A^{\mathrm{t}}\|_\infty = \|A^{\mathrm{c}}\|_\infty. \tag{7.18}$$

In this connection, the Euclidean norm is also known as the **2-norm**, since

$$\|x\| = \sqrt{x^{\mathrm{c}}x} = \sqrt{\sum_j |x_j|^2} =: \|x\|_2.$$

Therefore, when it is important, one writes the corresponding map-norm with a subscript 2, too. For example, compare (7.18) with

$$\|A\| = \|A\|_2 = \|A^{\mathrm{c}}\|_2 = \|A^{\mathrm{t}}\|_2. \tag{7.19}$$

For the proof of these identities, recall from (6.24) that

$$\|x\|_2 = \max_{y\neq 0} |\langle x, y\rangle|/\|y\|_2, \qquad x \in \mathbb{F}^n. \tag{7.20}$$

Hence,

$$\|A\|_2 = \max_{x\neq 0} \frac{\|Ax\|_2}{\|x\|_2} = \max_{x\neq 0}\max_{y\neq 0} \frac{|\langle Ax, y\rangle|}{\|x\|_2\|y\|_2} = \max_{y\neq 0}\max_{x\neq 0} \frac{|\langle x, A^{\mathrm{c}}y\rangle|}{\|x\|_2\|y\|_2} = \max_{y\neq 0} \frac{\|A^{\mathrm{c}}y\|_2}{\|y\|_2} = \|A^{\mathrm{c}}\|_2. \tag{7.21}$$

The equality $\|A^{\mathrm{t}}\| = \|A^{\mathrm{c}}\|$ holds in any of the map-norms discussed since they all depend only on the absolute values of the entries of the matrix A.

> The `MATLAB` statement `norm(A,inf)` provides the norm $\|\mathtt{A}\|_\infty$ in case `A` is a 'matrix', i.e., not a 'vector'. If `A` happens to equal $[v]$ or $[v]^{\mathrm{t}}$ for some vector v, then `norm(A,inf)` returns the max-norm of that vector, i.e., the number $\|v\|_\infty$. By (7.16), this is ok if `A` $= [v]$, but gives, in general, the wrong result if `A` $= v^{\mathrm{t}}$. This is an additional reason for sticking with the rule of using only $(n,1)$-matrices for representing n-vectors in `MATLAB`.
>
> The 1-norm, $\|\mathtt{A}\|_1$, is supplied by the statement `norm(A,1)`.

□

All three (vector-)norms mentioned so far are, indeed, norms in the sense of the following definition.

(7.22) Definition: The map $\| \ \| : X \to \mathbb{R} : x \mapsto \|x\|$ is a **vector norm**, provided it is

(i) **positive definite**, i.e., $\forall\{x \in X\}\ \|x\| \ge 0$ with equality if and only if $x = 0$;

(ii) **absolutely homogeneous**, i.e., $\forall\{\alpha \in \mathbb{F}, x \in X\}\ \|\alpha x\| = |\alpha|\|x\|$;

(iii) **subadditive**, i.e., $\forall\{x, y \in X\}\ \|x + y\| \le \|x\| + \|y\|$.

This last inequality is called the **triangle inequality**, and the vector space X supplied with a vector norm is called a **normed vector space**.

The absolute value is a vector norm for the vector space $\mathbb{F} = \mathbb{F}^1$. From this, it is immediate that both the max-norm and the 1-norm are vector norms for $\mathbb{F}^n$. As to the norm $x \mapsto \sqrt{x^c x}$ on an inner product space and, in particular, the Euclidean or 2-norm on $\mathbb{F}^n$, only the triangle inequality might still be in doubt, but it is an immediate consequence of (6.24)Cauchy's Inequality, which gives that

$$\langle x, y\rangle + \langle y, x\rangle = 2\,\mathrm{Re}\langle x, y\rangle \le 2|\langle x, y\rangle| \le 2\|x\|\|y\|,$$

and therefore:

$$\|x + y\|^2 = \|x\|^2 + \langle x, y\rangle + \langle y, x\rangle + \|y\|^2 \le (\|x\| + \|y\|)^2.$$

Also, for X finite-dimensional, and both X and Y normed vector spaces, with norms $\| \ \|_X$ and $\| \ \|_Y$ respectively, the vector space $L(X, Y)$ is a normed vector space with respect to the corresponding map norm

$$\|A\| := \|A\|_{X,Y} := \max_{x \ne 0} \frac{\|Ax\|_Y}{\|x\|_X}. \tag{7.23}$$

All statements about the map norm $\|A\|$ made in the preceding section hold for any of the map norms $\|A\|_{X,Y}$ since their proofs there use only the fact that $x \mapsto \sqrt{x^c x}$ is a norm according to (7.22)Definition. In particular, we will feel free to consider

$$\kappa(A)_p := \|A\|_p \|A^{-1}\|_p, \quad p = 1, 2, \infty, \quad A \in \mathbb{F}^n.$$

Why all these different norms? Each norm associates with a vector just one number, and, as with bases, any particular situation may best be handled by a particular norm.

For example, in considering the condition of the power basis $V := [()^{j-1} : j = 1{:}k]$ for $\Pi_{<k}$, we might be more interested in measuring the size of the residual $p - V\widehat{a}$ in terms of the max-norm

$$\|f\|_{[c..d]} := \max\{|f(t)| : c \le t \le d\}$$

over the interval $[c \mathinner{..} d]$ of interest, rather than in the averaging way supplied by the corresponding 2-norm

$$\left(\int_a^b |f(t)|^2 \,\mathrm{d}t\right)^{1/2}.$$

In any case, any two norms on a finite-dimensional vector space are equivalent in the following sense.

(7.24) Proposition: For any two norms, $\| \ \|'$ and $\| \ \|''$, on a finite-dimensional vector space X, there exists a positive constant c so that

$$\|x\|'' \le c\|x\|', \quad \forall x \in X.$$

This is just the statement that the map norm

$$\|\operatorname{id}_X\| := \max_{x\neq 0} \|x\|''/\|x\|'$$

is finite.

For example, for any $x \in \mathbb{F}^n$,

$$\|x\|_1 \le \sqrt{n}\|x\|_2, \quad \text{and} \quad \|x\|_2 \le \sqrt{n}\|x\|_\infty, \quad \text{while} \quad \|x\|_1 \ge \|x\|_2 \ge \|x\|_\infty. \tag{7.25}$$

Finally, given that it is very easy to compute the max-norm $\|A\|_\infty$ of $A \in \mathbb{F}^{m\times n}$ and much harder to compute the 2-norm $\|A\| = \|A\|_2$, why does one bother at all with the 2-norm? One very important reason is the availability of a large variety of *isometries*, i.e., matrices A with

$$\|Ax\| = \|x\|, \quad \forall x.$$

Each of these provides an o.n. basis for its range, and, by (6.19)Proposition, each finite-dimensional linear subspace of an inner product space has o.n. bases.

In contrast, the only $A \in \mathbb{F}^{n\times n}$ that are isometries in the max-norm, i.e., satisfy

$$\|Ax\|_\infty = \|x\|_\infty, \quad \forall x \in \mathbb{F}^n,$$

are of the form

$$\operatorname{diag}(\varepsilon_1, \ldots, \varepsilon_n)P,$$

with P a permutation matrix and each ε_j a scalar of absolute value 1.

For this reason, we continue to rely on the 2-norm. In fact, any norm without a subscript or other adornment is meant to be the 2-norm (or, more generally, the norm in the relevant inner product space).

7.4 Prove that, for any $\alpha \in \mathbb{F}$, the linear map $M_\alpha : X \to X : x \mapsto \alpha x$ on the normed vector space $X \ne \{0\}$ has map norm $|\alpha|$.

7.5 Prove that, for any diagonal matrix $D \in \mathbb{F}^{m\times n}$ and for $p = 1, 2, \infty$, $\|D\|_p = \max_j |D(j,j)|$.

8. Factorization and rank

The need for factoring linear maps

In order to compute with a linear map $A \in L(X, Y)$, we have to factor it through a coordinate space. This means that we have to write it as

$$A = V\Lambda^{\mathrm{t}}, \qquad \text{with } \ V \in L(\mathbb{F}^r, Y), \text{ hence } \Lambda^{\mathrm{t}} \in L(X, \mathbb{F}^r) \ .$$

The following picture might be helpful:

$$\begin{tikzcd} X \arrow[rr, "A"] \arrow[dr, "\Lambda^{\mathrm{t}}"'] & & Y \\ & \mathbb{F}^r \arrow[ur, "V"'] & \end{tikzcd}$$

For example, recall how you apply the linear map D of differentiation to a polynomial $p \in \Pi_k$: First you get the polynomial coefficients of that polynomial, and then you write down Dp in terms of those coefficients.

To test my claim, carry out the following thought experiment: You know that there is exactly one polynomial p of degree $\le k$ that matches given ordinates at given $k+1$ distinct abscissae, i.e., that satisfies

$$p(\tau_i) = y_i, \qquad i = 0{:}k$$

for given data $(\tau_i, y_i), i = 0{:}k$. Now, try, e.g., to compute the first derivative of the polynomial p of degree ≤ 3 that satisfies $p(j) = (-1)^j$, $j = 1, 2, 3, 4$. Can you do it without factoring the linear map $D : \Pi_3 \to \Pi_3$ through some coordinate space?

As another *example*, recall how we dealt with **coordinate maps**, i.e., the inverse of a basis. We saw that, even though a basis $V : \mathbb{F}^n \to \mathbb{F}^m$ for some linear subspace X of $\mathbb{F}^m$ is a concrete matrix, its inverse, V^{-1} is, offhand, just a formal expression. For actual work, we made use of any *matrix* $\Lambda^{\mathrm{t}} : \mathbb{F}^m \to \mathbb{F}^n$ that is 1-1 on X, thereby obtaining the *factorization*

$$V^{-1} = (\Lambda^{\mathrm{t}} V)^{-1} \Lambda^{\mathrm{t}}$$

in which $\Lambda^{\mathrm{t}} V$ is a square matrix, hence $(\Lambda^{\mathrm{t}} V)^{-1}$ is also a matrix.

The smaller one can make $\#V$ in a factorization $A = V\Lambda^{\mathrm{t}}$ of $A \in L(X, Y)$, the cheaper is the calculation of A.

Definition: The smallest r for which $A \in L(X, Y)$ can be factored as $A = V\Lambda^{\mathrm{t}}$ with $V \in L(\mathbb{F}^r, Y)$ (hence $\Lambda^{\mathrm{t}} \in L(X, \mathbb{F}^r)$) is called the **rank** of A. This is written

$$r = \operatorname{rank} A.$$

Any factorization $A = V\Lambda^{\mathrm{t}}$ with $\#V = \operatorname{rank} A$ is called **minimal**.

As an *example*,

$$A := \begin{bmatrix} 1 & 2 & 3 & 4 & 5 \\ 1 & 2 & 3 & 4 & 5 \\ 1 & 2 & 3 & 4 & 5 \\ 1 & 2 & 3 & 4 & 5 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} [\,1 \quad 2 \quad 3 \quad 4 \quad 5\,],$$

hence this A has rank 1 (since we can write it as $A = V\Lambda^{\mathrm{t}}$ with $\operatorname{dom} V = \mathbb{F}^1$, but we couldn't do it with $\operatorname{dom} V = \mathbb{F}^0$). To calculate Ax, we merely need to calculate the number $\alpha := (1,2,3,4,5)^{\mathrm{t}}x$, and then obtain Ax as the particular scalar multiple $y\alpha$ of the vector $y := (1,1,1,1)$. That is much cheaper than computing the matrix product of the 4×5-matrix A with the 1-column matrix $[x]$.

As the example illustrates, any matrix

$$A := [v][w]^{\mathrm{t}} = vw^{\mathrm{t}}$$

with $v \in \mathbb{F}^m$ and $w \in \mathbb{F}^n$ has rank 1 unless it is trivial, i.e., unless either v or w is the zero vector. This explains why an *elementary* matrix is also called a **rank-one perturbation of the identity**.

The only linear map of rank 0 is the zero map. If A is not the zero map, then its range contains some nonzero vector, hence so must the range of any V for which $A = V\Lambda^{\mathrm{t}}$ with $\operatorname{dom} V = \mathbb{F}^r$, therefore such r must be > 0.

As another *example*, for any vector space X,

$$\dim X = \operatorname{rank} \operatorname{id}_X.$$

Indeed, if $n = \dim X$, then, for any basis $V \in L(\mathbb{F}^n, X)$ for X, $\operatorname{id}_X = VV^{-1}$, therefore $\operatorname{rank} \operatorname{id}_X \le n$, while, for any factorization $\operatorname{id}_X = V\Lambda^{\mathrm{t}}$ for some $V \in L(\mathbb{F}^r, X)$, V must necessarily be onto, hence $\dim X \le r$, by (4.6)Proposition, and therefore $\dim X \le \operatorname{rank} \operatorname{id}_X$. In fact, it is possible to make the rank concept the primary one and *define* $\dim X$ as the rank of id_X.

When A is an $m \times n$-matrix, then, trivially, $A = A \operatorname{id}_n = \operatorname{id}_m A$, hence $\operatorname{rank} A \le \min\{m,n\}$.

At times, particularly when A is a matrix, it is convenient to write the factorization $A = V\Lambda^{\mathrm{t}}$ more explicitly as

$$(8.1) \qquad A =: [v_1, v_2, \ldots, v_r][\lambda_1, \lambda_2, \ldots, \lambda_r]^{\mathrm{t}} = \sum_{j=1}^{r} [v_j]\lambda_j.$$

Since each of the maps

$$v_j\lambda_j := [v_j]\lambda_j = [v_j] \circ \lambda_j : x \mapsto (\lambda_j x)v_j$$

has rank ≤ 1, this shows that *the rank of A gives the smallest number of terms necessary to write A as a sum of rank-one maps.*

(8.2) Proposition: $A = V\Lambda^{\mathrm{t}}$ is minimal if and only if V is a basis for $\operatorname{ran} A$. In particular,

$$\operatorname{rank} A = \dim \operatorname{ran} A.$$

Proof: Let $A = V\Lambda^{\mathrm{t}}$. Then $\operatorname{ran} A \subset \operatorname{ran} V$, hence

$$\dim \operatorname{ran} A \le \dim \operatorname{ran} V \le \#V,$$

with equality in the first $\le$ iff $\operatorname{ran} A = \operatorname{ran} V$ (by (4.13)Proposition), and in the second $\le$ iff V is 1-1. Thus, $\dim \operatorname{ran} A \le \#V$, with equality iff V is a basis for $\operatorname{ran} A$. □

One can prove in a similar way that $A = V\Lambda^{\mathrm{t}}$ is minimal if and only if Λ^{t} is onto and $\operatorname{null} A = \operatorname{null} \Lambda^{\mathrm{t}}$.

(8.3) Corollary: The factorization $A = A(:, \texttt{bound})\operatorname{rrref}(A)$ provided by elimination (see (3.11)) is minimal.

(8.4) Corollary: If $A = V\Lambda^{\mathrm{t}}$ is minimal and A is invertible, then also V and Λ^{t} are invertible.

Proof: By (8.2)Proposition, $V \in L(\mathbb{F}^r, Y)$ is a basis for $\operatorname{ran} A$, while $\operatorname{ran} A = Y$ since A is invertible. Hence, V is invertible. Therefore, also $\Lambda^{\mathrm{t}} = V^{-1}A$ is invertible. □

But note that the matrix $[\,1\,] = [\,1 \quad 0\,] \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ is invertible, even though neither of its two factors is.

8.1 Determine a minimal factorization for the matrix

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

8.2 With A the matrix of the previous problem, give a basis for $\operatorname{ran} A$ and a basis for $\operatorname{ran} A^{\mathrm{t}}$.

8.3 Give an example of a pair of matrices, A and B, of order 4, each of rank 2, yet $\operatorname{ran} A \cap \operatorname{ran} B = \{0\}$.

8.4 Prove: *For any two linear maps A and B for which AB is defined,* $\operatorname{rank}(AB) \le \min\{\operatorname{rank} A, \operatorname{rank} B\}$. (Hint: If $A = V_A {\Lambda_A}^{\mathrm{t}}$ and $B = V_B {\Lambda_B}^{\mathrm{t}}$, then $AB = V_A({\Lambda_A}^{\mathrm{t}} V_B {\Lambda_B}^{\mathrm{t}}) = (V_A {\Lambda_A}^{\mathrm{t}} V_B){\Lambda_B}^{\mathrm{t}}$. Totally different hint: Use the Dimension Formula together with the fact that $\operatorname{rank} C = \dim \operatorname{ran} C$.)

8.5 Prove: *If $A = V\Lambda^{\mathrm{t}}$ is a minimal factorization and A is a projector (i.e., $A^2 = A$), then* $\Lambda^{\mathrm{t}} V = \mathrm{id}$. (Hint: H.P. 1.26 .)

The trace of a linear map

Each $A \in L(X)$ can be factored in possibly many different ways as

$$A = V\Lambda^{\mathrm{t}} = [v_1, \ldots, v_n][\lambda_1, \ldots, \lambda_n]^{\mathrm{t}}$$

for some n (necessarily $\ge \operatorname{rank} A$). It may therefore be surprising that, nevertheless, the number

$$\sum_j \lambda_j v_j$$

only depends on A. For the proof of this claim, we notice that

$$\sum_j \lambda_j v_j = \operatorname{trace}(\Lambda^{\mathrm{t}} V).$$

Now, let W be a basis for X, with dual basis $\mathrm{M} := W^{-1}$. Then

$$\hat{A} := \mathrm{M}^{\mathrm{t}} A W = \mathrm{M}^{\mathrm{t}} V \Lambda^{\mathrm{t}} W,$$

while

$$\Lambda^{\mathrm{t}} W \mathrm{M}^{\mathrm{t}} V = \Lambda^{\mathrm{t}} V.$$

Hence, by (6.31),

$$\operatorname{trace}(\hat{A}) = \operatorname{trace}(\mathrm{M}^{\mathrm{t}} V \Lambda^{\mathrm{t}} W) = \operatorname{trace}(\Lambda^{\mathrm{t}} W \mathrm{M}^{\mathrm{t}} V) = \operatorname{trace}(\Lambda^{\mathrm{t}} V).$$

By holding our factorization $A = V\Lambda^{\mathrm{t}}$ fixed, this implies that $\operatorname{trace}(\hat{A})$ does not depend on the particular basis W for X we happen to use here, hence only depends on the linear map A. With that, holding now this linear map A fixed, we see that also $\operatorname{trace}(\Lambda^{\mathrm{t}} V)$ does not depend on the particular factorization $A = V\Lambda^{\mathrm{t}}$ we picked, but only depends on A. This number is called the trace of A, written

$$\operatorname{trace}(A).$$

The problems provide the basic properties of the trace of a linear map.

8.6 $\operatorname{trace}(\mathrm{id}_X) = \dim X$.

8.7 If $P \in L(X)$ is a projector (i.e., $P^2 = P$), then $\operatorname{trace}(P) = \dim \operatorname{ran} P$.

8.8 $A \mapsto \operatorname{trace}(A)$ is the unique scalar-valued linear map on $L(X)$ for which $\operatorname{trace}([x]\lambda) = \lambda x$ for all $x \in X$ and $\lambda \in X'$.

8.9 If $A \in L(X,Y)$ and $B \in L(Y,X)$, then (both AB and BA are defined and) $\operatorname{trace}(AB) = \operatorname{trace}(BA)$.

8.10 Prove that, for column maps V, W into X, and row maps Λ^{t}, M^{t} from X, $V\Lambda^{\mathrm{t}} = W\mathrm{M}^{\mathrm{t}}$ implies that $\operatorname{trace}(\Lambda^{\mathrm{t}} V) = \operatorname{trace}(\mathrm{M}^{\mathrm{t}} W)$ even if X is not finite-dimensional.

The rank of a matrix and of its (conjugate) transpose

In this section, let A' denote either the transpose or the conjugate transpose of the matrix A. Then, either way, $A = VW'$ iff $A' = WV'$. This trivial observation implies all kinds of things about the relationship between a matrix and its (conjugate) transpose.

As a starter, it says that $A = VW'$ is minimal if and only if $A' = WV'$ is minimal. Therefore:

Proposition: $\operatorname{rank} A = \operatorname{rank} A^{\mathrm{c}} = \operatorname{rank} A^{\mathrm{t}}$.

(8.5) Corollary: If A is a matrix, then $\dim \operatorname{ran} A = \dim \operatorname{ran} A^{\mathrm{c}} = \dim \operatorname{ran} A^{\mathrm{t}}$.

(8.6) Corollary: For any matrix A, A' is 1-1 (onto) if and only if A is onto (1-1).

Proof: If $A \in \mathbb{F}^{m\times n}$, then A is onto iff $\operatorname{rank} A = m$ iff $\operatorname{rank} A' = m$ iff the natural factorization $A' = A' \operatorname{id}_m$ is minimal, i.e., iff A' is 1-1.

The other equivalence follows from this since $(A')' = A$. □

For a different proof of these results, see the comments that follow (6.26)Corollary and (6.27)Corollary.

Elimination as factorization

The description (3.2) of elimination does not rely on any particular ordering of the rows of the given $(m \times n)$-matrix A. At any stage, it only distinguishes between pivot rows and those rows not yet used as pivot rows. We may therefore imagine that we initially place the rows of A into the workarray B in exactly the order in which they are going to be used as pivot rows, followed, in any order whatsoever, by those rows (if any) that are never going to be used as pivot rows.

In terms of the n-vector `p` provided by the (3.2)Elimination Algorithm, this means that we start with $B = A(\mathtt{q},:)$, with `q` obtained from `p` by

```
q = p(find(p>0)); 1:m; ans(q) = []; q = [q, ans];
```

Indeed, to recall, `p(j)` is positive if and only if the `j`th unknown is bound, in which case row `p(j)` is the pivot row for that unknown. Thus the assignment `q = p(find(p>0))` initializes `q` so that `A(q,:)` contains the pivot rows in order of their use. With that, `1:m; ans(q) = [];` leaves, in `ans`, the indices of all rows not used as pivot rows.

Note that `q` is a permutation of order m. Hence $B = QA$, with Q the corresponding permutation matrix, meaning the matrix obtained from the identity matrix by the very same reordering, $Q =$`eye(m)(q,:)`.

We prefer to write this as $A = PB$, with P the inverse of Q, hence obtainable from `q` by

```
P = eye(m); P(q,:) = P;
```

□

With that done, we have, at the beginning of the algorithm,

$$B = P^{-1}A$$

for some permutation matrix P, and all the work in the algorithm consists of repeatedly subtracting some multiple α of some row h of B from some *later* row, i.e., some row i with $i > h$. In terms of matrices, this means the repeated replacement

$$B \leftarrow E_{i,h}(-\alpha)B$$

with $i > h$. Since, by (2.19), $E_{i,h}(-\alpha)^{-1} = E_{i,h}(\alpha)$, this implies that

$$A = PLU,$$

with L the product of all those elementary matrices $E_{i,h}(\alpha)$ (in the appropriate order), and U the final state of the workarray B. Specifically, U is in row-echelon form (as defined in (3.7)); in particular, U is upper triangular.

Each $E_{i,h}(\alpha)$ is **unit lower triangular**, i.e., of the form $\mathrm{id} + N$ with N **strictly lower triangular**, i.e.,

$$N(r,s) \neq 0 \quad \Longrightarrow \quad r > s.$$

For, because of the initial ordering of the rows in B, only $E_{i,h}(\alpha)$ with $i > h$ appear. This implies that L, as the product of unit lower triangular matrices, is itself unit lower triangular.

If we apply the elimination algorithm to the matrix $[A, C]$, with $A \in \mathbb{F}^{m\times m}$ invertible, then the first m columns are bound, hence the remaining columns are free. In particular, both P and L in the resulting factorization depend only on A and not at all on C.

In particular, in solving $A? = y$, there is no need to subject all of $[A, y]$ to the elimination algorithm. If elimination just applied to A gives the factorization

$$A = PLU \tag{8.7}$$

for an invertible A, then we can find the unique solution x to the equation $A? = y$ by the two-step process:

$$c \leftarrow L^{-1}P^{-1}y$$
$$x \leftarrow U^{-1}c$$

and these two steps are easily carried out. The first step amounts to subjecting the rows of the matrix $[y]$ to all the row operations (including reordering) used during elimination applied to A. The second step is handled by the Backsubstitution Algorithm (3.3), with input $B = [U, c]$, $p = (1, 2, \ldots, m, 0)$, and $z = (0, \ldots, 0, -1)$.

Once it is understood that the purpose of elimination for solving $A? = y$ is the factorization of A into a product of "easily" invertible factors, then it is possible to seek factorizations that might serve the same goal in a better way. The best-known alternative is the QR factorization, in which one obtains

$$A = QR,$$

with R upper triangular and Q o.n., i.e., $Q^{\mathrm{c}}Q = \mathrm{id}$. Such a factorization is obtained by doing elimination a column at a time, usually with the aid of **Householder matrices**. These are elementary matrices of the form

$$H_w := E_{w,w}(-2/w^{\mathrm{c}}w) = \mathrm{id} - \frac{2}{w^{\mathrm{c}}w}ww^{\mathrm{c}},$$

and are easily seen to be **self-inverse** or **involutory** (i.e., $H_wH_w = \mathrm{id}$), **hermitian** (i.e., ${H_w}^{\mathrm{c}} = H_w$), hence **unitary** (i.e., ${H_w}^{\mathrm{c}}H_w = \mathrm{id} = H_w{H_w}^{\mathrm{c}}$).

While the computational cost of constructing the QR factorization is roughly double that needed for the PLU factorization, the QR factorization has the advantage of being more impervious to the effects of rounding errors. Precisely, the relative rounding error effects in both a PLU factorization $A = PLU$ and in a QR factorization $A = QR$ can be shown to be proportional to the condition numbers of the factors. Since Q is o.n., $\kappa(Q) = 1$ and $\kappa(R) = \kappa(A)$, while, for a PLU factorization $A = PLU$, only the permutation matrix, P, is o.n., and $\kappa(L)$ and $\kappa(U)$ can be quite large.

8.11 Prove: *If $L_1D_1U_1 = A = L_2D_2U_2$, with L_i unit lower triangular, D_i invertible diagonal, and U_i unit upper triangular matrices, then $L_1 = L_2$, $D_1 = D_2$, and $U_1 = U_2$.*

SVD

Let $A = VW^c$ be a minimal factorization for the $m \times n$-matrix A of rank r. Then $A^c = WV^c$ is a minimal factorization for A^c. By (8.2), this implies that V is a basis for $\operatorname{ran} A$ and W is a basis for $\operatorname{ran} A^c$.

Can we choose both these bases to be o.n.?

Well, if both V and W are o.n., then, for any x, $\|Ax\| = \|VW^c x\| = \|W^c x\|$, while, for $x \in \operatorname{ran} A^c$, $x = WW^c x$, hence $\|x\| = \|W^c x\|$. Therefore, altogether, in such a case, A is an isometry on $\operatorname{ran} A^c$, a very special situation.

Nevertheless and, perhaps, surprisingly, there is an o.n. basis W for $\operatorname{ran} A^c$ for which the columns of AW are *orthogonal*, i.e., $AW = V\Sigma$ with V o.n. and Σ diagonal, hence $A = V\Sigma W^c$ with also V o.n.

(8.8) Theorem: For every $A \in \mathbb{F}^{m\times n}$, there exist o.n. bases V and W for $\operatorname{ran} A$ and $\operatorname{ran} A^c$, respectively, and a diagonal matrix Σ with positive diagonal entries, so that

(8.9) $$A = V\Sigma W^c.$$

Proof: For efficiency, the proof given here uses results, concerning the eigenstructure of hermitian positive definite matrices, that are established only later in these notes. This may help to motivate the study to come of the eigenstructure of matrices.

For motivation of the proof, assume for the moment that $A = V\Sigma W^c$ is a factorization of the kind we claim to exist. Then, with $\Sigma =: \operatorname{diag}(\sigma_1, \dots, \sigma_r)$, it follows that

$$A^c A = W\Sigma^c V^c\, V\Sigma W^c = W\Sigma^c\Sigma W^c,$$

hence

(8.10) $$A^c A W = W\mathrm{T}, \quad \text{with} \quad \mathrm{T} := \operatorname{diag}(\tau_1, \tau_2, \dots, \tau_r)$$

and W o.n., and the $\tau_j = \overline{\sigma_j}\sigma_j = |\sigma_j|^2$ all positive.

Just such an o.n. $W \in \mathbb{F}^{n\times r}$ and positive scalars τ_j do exist by (12.2)Corollary and (15.2)Proposition, since the matrix $A^c A$ is **hermitian** (i.e., $(A^c A)^c = A^c A$) and **positive semidefinite** (i.e., $\langle A^c A x, x\rangle \ge 0$ for all x) and has rank r.

With W and the τ_j so chosen, it follows that W is an o.n. basis for $\operatorname{ran} A^c$, since (8.10) implies that $\operatorname{ran} W \subset \operatorname{ran} A^c$, and W is a 1-1 column map of order $r = \dim \operatorname{ran} A^c$. Further, $U := AW$ satisfies $U^c U = W^c A^c A W = W^c W \mathrm{T} = \mathrm{T}$, hence

$$V := AW\Sigma^{-1}, \quad \text{with} \quad \Sigma := \mathrm{T}^{1/2} := \operatorname{diag}(\sqrt{\tau_j} : j = 1{:}r),$$

is o.n., and so $V\Sigma W^c = A$, because $WW^c = P := P_{\operatorname{ran} A^c}$, hence $\operatorname{ran}(\operatorname{id} - P) = \operatorname{null} P = \operatorname{ran} A^{c\perp} = \operatorname{null} A$, and so $AWW^c = AP = A(P + (\operatorname{id} - P)) = A$. □

It is customary to order the numbers

$$\sigma_j := \sqrt{\tau_j}, \quad j = 1{:}r.$$

Specifically, one assumes that

$$\sigma_1 \ge \sigma_2 \ge \cdots \ge \sigma_r.$$

These numbers σ_j are called the (nonzero) **singular value**s of A, and with this ordering, the factorization

$$A = \sum_{j=1}^{\operatorname{rank} A} v_j \sigma_j {w_j}^c$$

is called a (**reduced**) **singular value decomposition** or **svd** for A.

Offhand, a svd is *not* unique. E.g., *any* o.n. basis V for $\mathbb{F}^n$ provides the svd $V\,\mathrm{id}_n V^c$ for id_n.

Some prefer to have a factorization $A = \tilde{V}\tilde{\Sigma}\tilde{W}^c$ in which both $\tilde{V}$ and $\tilde{W}$ are o.n. bases for all of $\mathbb{F}^m$ and $\mathbb{F}^n$, respectively (rather than just for $\operatorname{ran} A$ and $\operatorname{ran} A^c$, respectively). This can always be achieved by extending V and W from (8.9) in any manner whatsoever to o.n. bases $\tilde{V} := [V, V_1]$ and $\tilde{W} := [W, W_1]$ and, correspondingly, extending Σ to

$$\tilde{\Sigma} := \operatorname{diag}(\Sigma, 0) = \begin{bmatrix} \Sigma & 0 \\ 0 & 0 \end{bmatrix} \in \mathbb{F}^{m\times n}$$

by the adjunction of blocks of 0 of appropriate size. With this, we have

$$A = \tilde{V}\tilde{\Sigma}\tilde{W}^c = \sum_{j=1}^{\min\{m,n\}} v_j \sigma_j {w_j}^c, \tag{8.11}$$

and the diagonal entries

$$\sigma_1 \ge \sigma_2 \ge \cdots \ge \sigma_r > 0 = \sigma_{r+1} = \cdots = \sigma_{\min\{m,n\}}$$

of $\tilde{\Sigma}$ are altogether referred to as the **singular value**s of A. Note that this sequence is still ordered. We will refer to (8.11) as a **Singular Value Decomposition** or **SVD**.

The `MATLAB` command `svd(A)` returns the SVD rather than the svd of `A` when issued in the form `[V,S,W] = svd(A)`. Specifically, `A = V*S*W'`, with `V` and `W` both unitary, of order m and n, respectively, if `A` is an $m \times n$-matrix. By itself, `svd(A)` returns, in a one-column matrix, the (ordered) sequence of singular values of `A`.

□

The Pseudo-inverse

Here is a first of many uses to which the svd has been put. It concerns the solution of the equation

$$A? = y$$

in case A is not invertible (for whatever reason). In a previous chapter (see page 71), we looked in this case for a solution of the 'projected' problem

$$A? = P_{\operatorname{ran} A}\, y =: \widehat{y} \tag{8.12}$$

for the simple reason that any solution x of this equation makes the **residual** $\|Ax - y\|_2$ as small as it can be made by any x. For this reason, any solution of (8.12) is called a **least-squares solution** for $A? = y$.

If now A is 1-1, then (8.12) has exactly one solution. The question is what to do in the contrary case. One proposal is to get the **best least-squares solution**, i.e., the solution of minimal norm. The svd for A makes it easy to find this particular solution.

If $A = V\Sigma W^c$ is a svd for A, then V is an o.n. basis for $\operatorname{ran} A$, hence

$$\widehat{b} = P_{\operatorname{ran} A} b = VV^c b.$$

Therefore, (8.12) is equivalent to the equation

$$V\Sigma W^c ? = VV^c b.$$

Since V is o.n., hence 1-1, and Σ is invertible, this equation is, in turn, equivalent to

$$W^c ? = \Sigma^{-1} V^c b,$$

hence to

$$WW^c? = W\Sigma^{-1}V^c b. \tag{8.13}$$

Since W is also o.n., $WW^c = P_W$ is an o.n. projector, hence, by (6.15)Proposition, strictly reduces norms unless it is applied to something in its range. Since the right-hand side of (8.13) is in $\operatorname{ran} W$, it follows that the solution of smallest norm of (8.13), i.e., the best least-squares solution of $A? = y$, is that right-hand side, i.e., the vector

$$\widehat{x} := A^+ y,$$

with the matrix

$$A^+ := W\Sigma^{-1}V^c$$

the **Moore-Penrose pseudo-inverse** of A.

Note that

$$A^+A = W\Sigma^{-1}V^c V\Sigma W^c = WW^c,$$

hence A^+ is a left inverse for A in case W is square, i.e., in case $\operatorname{rank} A = \#A$. Similarly,

$$AA^+ = V\Sigma W^c W\Sigma^{-1}V^c = VV^c,$$

hence A^+ is a right inverse for A in case V is square, i.e., in case $\operatorname{rank} A = \#A^c$. In any case,

$$A^+A = P_{\operatorname{ran} A^c}, \qquad AA^+ = P_{\operatorname{ran} A},$$

therefore, in particular,

$$AA^+A = A.$$

2-norm and 2-condition of a matrix

Recall from (6.23) that o.n. matrices are 2-norm-preserving, i.e.,

$$\|x\|_2 = \|Ux\|_2, \qquad \forall x \in \mathbb{F}^n, \text{o.n. } U \in \mathbb{F}^{m\times n}.$$

This implies that

$$\|TB\|_2 = \|B\|_2 = \|BU^c\|_2, \qquad \forall \text{ o.n. } T \in \mathbb{F}^{r\times m},\ B \in \mathbb{F}^{m\times n},\ \text{o.n. } U \in \mathbb{F}^{r\times n}.$$

Indeed,

$$\|TB\|_2 = \max_{x\neq 0} \frac{\|TBx\|_2}{\|x\|_2} = \max_{x\neq 0} \frac{\|Bx\|_2}{\|x\|_2} = \|B\|_2.$$

By (7.21), this implies that also

$$\|BU^c\|_2 = \|UB^c\|_2 = \|B^c\|_2 = \|B\|_2.$$

It follows that, with $A = V\Sigma W^c \in \mathbb{F}^{m\times n}$ a svd for A,

$$\|A\|_2 = \|\Sigma\|_2 = \sigma_1, \tag{8.14}$$

the last equality because of the fact that $\Sigma = \operatorname{diag}(\sigma_1, \ldots, \sigma_r)$ with $\sigma_1 \geq \sigma_2 \geq \cdots \geq 0$.

Assume that, in addition, A is invertible, therefore $r = \operatorname{rank} A = n = m$, making also V and W square, hence A^+ is both a left and a right inverse for A, therefore necessarily $A^{-1} = A^+ = V\Sigma^{-1}W^c$. It follows that $\|A^{-1}\|_2 = 1/\sigma_n$. Hence, the 2-condition of $A \in \mathbb{F}^{n\times n}$ is

$$\kappa_2(A) = \|A\|_2 \|A^{-1}\|_2 = \sigma_1/\sigma_n,$$

and this is how this condition number is frequently *defined*.

The effective rank of a noisy matrix

The problem to be addressed here is the following. If we construct a matrix in the computer, we have to deal with the fact that the entries of the constructed matrix are not quite exact; rounding errors during the calculations may have added some noise. This is even true for a matrix merely entered into the computer, in case some of its entries cannot be represented exactly by the floating point arithmetic used (as is the case, e.g., for the number .1 or the number 1/3 in any of the standard binary-based floatingpoint arithmetics).

This makes it impossible to use, e.g., the rref algorithm to determine the rank of the underlying matrix. However, if one has some notion of the size of the noise involved, then one can use the svd to determine a sharp *lower* bound on the rank of the underlying matrix, because of the following.

(8.15) Proposition: If $A = V\Sigma W^{\mathrm{c}}$ is a svd for A and $\mathrm{rank}(A) > k$, then $\min\{\|A-B\|_2 : \mathrm{rank}(B) \le k\} = \sigma_{k+1} = \|A - A_k\|_2$, with

$$A_k := \sum_{j=1}^{k} v_j \sigma_j {w_j}^{\mathrm{c}}.$$

Proof: If $B \in \mathbb{F}^{m\times n}$ with $\mathrm{rank}(B) \le k$, then $\dim \mathrm{null}(B) > n - (k+1) = \dim \mathbb{F}^n - \dim \mathrm{ran}\, W_{k+1}$, with

$$W_{k+1} := [w_1, w_2, \ldots, w_{k+1}].$$

Therefore, by (4.21)Corollary, the intersection $\mathrm{null}(B) \cap \mathrm{ran}\, W_{k+1}$ contains a vector z of norm 1. Then $Bz = 0$, and $W^{\mathrm{c}} z = {W_{k+1}}^{\mathrm{c}} z$, and $\|{W_{k+1}}^{\mathrm{c}} z\|_2 = \|z\|_2 = 1$. Therefore, $Az = V\Sigma W^{\mathrm{c}} z = V_{k+1}\Sigma_{k+1}{W_{k+1}}^{\mathrm{c}} z$, hence

$$\|A - B\|_2 \ge \|Az - Bz\|_2 = \|Az\|_2 = \|\Sigma_{k+1}{W_{k+1}}^{\mathrm{c}} z\|_2 \ge \sigma_{k+1}\|{W_{k+1}}^{\mathrm{c}} z\|_2 = \sigma_{k+1}.$$

On the other hand, for the specific choice $B = A_k$, we get $\|A - A_k\|_2 = \sigma_{k+1}$ by (8.14), since $A - A_k = \sum_{j>k} v_j\sigma_j {w_j}^{\mathrm{c}}$ is a svd for it, hence its largest singular value is σ_{k+1}. □

In particular, if we have in hand a svd

$$A + E = V\,\mathrm{diag}(\hat\sigma_1, \hat\sigma_2, \ldots, \hat\sigma_{\hat r})W^{\mathrm{c}}$$

for the *perturbed* matrix $A + E$, and know (or believe) that $\|E\|_2 \le \varepsilon$, then the best we can say about the rank of A is that it must be at least

$$r_\varepsilon := \max\{j : \hat\sigma_j > \varepsilon\}.$$

For example, the matrix

$$A = \begin{bmatrix} 2/3 & 1 & 1/3 \\ 4/3 & 2 & 2/3 \\ 1 & 1 & 1 \end{bmatrix}$$

is readily transformed by elimination into the matrix

$$B = \begin{bmatrix} 0 & 1/3 & -1/3 \\ 0 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix},$$

hence has rank 2. However, on entering A into a computer correct to four decimal places after the decimal point, we get (more or less) the matrix

$$A_c = \begin{bmatrix} .6667 & 1 & .3333 \\ 1.3333 & 2 & .6667 \\ 1 & 1 & 1 \end{bmatrix},$$

and for it, MATLAB correctly returns id_3 as its `rref`. However, the singular values of A_c, as returned by `svd`, are

```
(3.2340..., 0.5645..., 0.000054...)
```

indicating that there is a rank-2 matrix B with $\|A_c - B\|_2 < .000055$. Since entries of A_c are only accurate to within 0.00005, the safe conclusion is that A has rank ≥ 2; it happens to have rank 2 in this particular example.

The polar decomposition

The svd can also be very helpful in establishing results of a more theoretical flavor, as the following discussion is intended to illustrate.

This discussion concerns a useful extension to square matrices of the polar form (see Backgrounder)

$$z = |z| \exp(\mathrm{i}\varphi)$$

of a complex number z, i.e., a factorization of z into a nonnegative number $|z| = \sqrt{z\overline{z}}$ (its modulus or absolute value) and a number whose absolute value is equal to 1, a socalled **unimodular** number.

There is, for any $A \in \mathbb{C}^{n\times n}$, a corresponding decomposition

$$A = \sqrt{AA^{\mathrm{c}}}E, \tag{8.16}$$

called a **polar decomposition**, with $\sqrt{AA^{\mathrm{c}}}$ 'nonnegative' in the sense that it is hermitian and positive semidefinite, and E 'unimodular' in the sense that it is unitary, hence norm-preserving, i.e., an isometry.

The polar decomposition is almost immediate, given that we already have a SVD $A = \tilde{V}\tilde{\Sigma}\tilde{W}^{\mathrm{c}}$ for A (see (8.11)) in hand. Indeed, from that,

$$A = \tilde{V}\tilde{\Sigma}\tilde{V}^{\mathrm{c}}\,\tilde{V}\tilde{W}^{\mathrm{c}},$$

with $P := \tilde{V}\tilde{\Sigma}\tilde{V}^{\mathrm{c}}$ evidently hermitian, and also positive semidefinite since

$$\langle Px, x\rangle = x^{\mathrm{c}}\tilde{V}\tilde{\Sigma}\tilde{V}^{\mathrm{c}}x = \sum_j \tilde{\sigma}_j |(\tilde{V}^{\mathrm{c}}x)_j|^2$$

is nonnegative for all x, given that $\tilde{\sigma}_j \ge 0$ for all j; and

$$P^2 = \tilde{V}\tilde{\Sigma}\tilde{V}^{\mathrm{c}}\tilde{V}\tilde{\Sigma}\tilde{V}^{\mathrm{c}} = \tilde{V}\tilde{\Sigma}\tilde{\Sigma}^{\mathrm{c}}\tilde{V}^{\mathrm{c}} = \tilde{V}\tilde{\Sigma}\tilde{W}^{\mathrm{c}}\tilde{W}\tilde{\Sigma}^{\mathrm{c}}\tilde{V}^{\mathrm{c}} = AA^{\mathrm{c}};$$

and, finally, $E := \tilde{V}\tilde{W}^{\mathrm{c}}$ unitary as the product of unitary maps.

Equivalence and similarity

The SVD provides a particularly useful example of *equivalence*. The linear maps A and $\widehat{A}$ are called **equivalent** if there are *invertible* linear maps V and W so that

$$A \;=\; V\widehat{A}W^{-1}.$$

Since both V and W are invertible, such equivalent linear maps share all essential properties, such as their rank, being 1-1, or onto, or invertible.

Equivalence is particularly useful when the domains of V and W are coordinate spaces, i.e., when V and W are *bases*, and, correspondingly, $\widehat{A}$ is a matrix, as in the following diagram:

$$\begin{array}{ccc} X & \xrightarrow{A} & Y \\ W\,\uparrow & & \uparrow\,V \\ \mathbb{F}^n & \xrightarrow[\widehat{A}]{} & \mathbb{F}^m \end{array}$$

In this situation, $\widehat{A} = V^{-1}AW$ is called a **matrix representation for** A.

For example, we noted earlier that the matrix

$$\widehat{D}_k := \begin{bmatrix} 0 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 2 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & k \end{bmatrix}$$

is the standard matrix representation used in Calculus for the linear map $D : \Pi_k \to \Pi_{k-1}$ of differentiation of polynomials of degree $\le k$.

In practice, one looks, for given $A \in L(X, Y)$, for matrix representations $\widehat{A}$ that are as simple as possible. If that means a matrix with as many zero entries as possible and, moreover, all the nonzero entries the same, say equal to 1, then a simplest such matrix representation is of the form

$$\widehat{A} = \operatorname{diag}(\,\mathrm{id}_{\operatorname{rank} A}, 0) = \begin{bmatrix} \mathrm{id}_{\operatorname{rank} A} & 0 \\ 0 & 0 \end{bmatrix},$$

with 0 indicating zero matrices of the appropriate size to make $\widehat{A}$ of size $\dim \operatorname{tar} A \times \dim \operatorname{dom} A$.

The situation becomes much more interesting and challenging when $\operatorname{dom} A = \operatorname{tar} A$ and, correspondingly, we insist that also $V = W$. Linear maps A and $\widehat{A}$ for which there exists an invertible linear map V with

$$A = V\widehat{A}V^{-1}$$

are called *similar*. Such similarity will drive much of the rest of these notes.

8.12 For the given linear maps $A, B, C : \mathbb{F}^{2\times 3}$, find their matrix representation with respect to the basis $V = [e_1+e_2, e_2+e_3, e_3+e_1]$ for $\mathbb{F}^3$ and $W := \begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix}$ for $\mathbb{F}^2$: (a) $Ax = (5x_1 + 2x_2 + 7x_3, x_1 + x_2 - x_3)$; (b) $Bx = (x_1 + x_2 + x_3, x_2 - x_3)$; (c) $Cx = (-x_1 - x_2 - x_3, x_3)$.

8.13 What is the matrix representation of the linear map $\mathbb{C} \to \mathbb{C} : x \mapsto zx$ with respect to the basis $[1, \mathrm{i}]$ for $\mathbb{C}$ (as a vector space with $\mathbb{F} = \mathbb{R}$) and with $z =: a + ib$ a given complex number?

8.14 T/F

(a) If A, B, M are matrices such that $\operatorname{rank} AM = \operatorname{rank} B$, then M is invertible.

(b) If M is invertible and $AM = B$, then $\operatorname{rank} A = \operatorname{rank} B$.

(c) If M is invertible and $MA = B$, then $\operatorname{rank} A = \operatorname{rank} B$.

9. Duality

This short chapter can be skipped without loss of continuity. Much of it can serve as a review of what has been covered so far. It owes much to the intriguing book [citGL].

Complementary mathematical concepts

Duality concerns mathematical concepts that come in pairs, that *complement* one another. Examples of interest in these notes include:

- $\subset$ and $\supset$;
- A *subset* S of T and its **complement**, $\backslash S := T \backslash S$;
- $\cap$ and $\cup$;
- $\forall$ and $\exists$;
- *1-1* and *onto*;
- *right* and *left* inverse;
- *bound* and *free*;
- *nullspace* and *range* of a linear map;
- an *invertible map* and its *inverse*;
- *column map* and *row map*;
- *synthesis* and *analysis*;
- a *basis* and its *inverse*;
- *columns* and *rows* of a matrix;
- a *matrix* and its (conjugate) *transpose*;
- a linear *subspace* and one of its *complement*s;
- dim and codim;
- the *vector space* X and its **dual**, $X' := L(X, \mathbb{F})$;
- the linear map $A \in L(X, Y)$ and its **dual**, $A' : Y' \to X' : \lambda \mapsto \lambda A$;
- a *norm* on the vector space X and the *dual norm* on X'.

Each such pair expresses a kind of *symmetry*. Such symmetry provides, with each result, also its ‘dual’, i.e., the result obtained by replacing one or more concepts appropriately by its complement. This leads to efficiency, both in the proving and in the remembering of results.

A classical example is that of *points* and *lines* in a geometry, and results concerning lines through points. E.g., *through every two distinct points there goes exactly one line*; its ‘dual’ statement is: *any two distinct lines have exactly one point in common.*

Another classical example is **DeMorgan’s Law**, according to which any statement concerning the union, intersection and containment of subsets is true if and only if its ‘dual’ statement is true, i.e., the statement obtained by replacing each set by its complement and replacing $(\subset, \supset, \cap, \cup)$ by $(\supset, \subset, \cup, \cap)$, respectively. For example, the two ‘distributive’ laws

$$(R \cap S) \cup T = (R \cup T) \cap (S \cup T), \quad (R \cup S) \cap T = (R \cap T) \cup (S \cap T)$$

are ‘dual’ to each other. Again, having verified that *the intersection of a collection of sets is the largest set contained in all of them*, we have, by ‘duality’, also verified that *the union of a collection of sets is the smallest set containing all of them.*

Here are some specific examples concerning the material covered in these notes so far.

Let V, W be column maps. *If $V \subset W$ and W is 1-1, then so is V.* Its ‘dual’: *if $V \supset W$ and W is onto, then so is V.* This makes maximally 1-1 maps and a minimally onto maps particularly interesting as, by now, you know very well: *A column map is maximally 1-1 if and only if it is minimally onto if and only if it is a basis.*

Let $A \in \mathbb{F}^{m\times n}$. Then, *$A$ is 1-1(onto) if and only if A^{t} is onto(1-1).* In terms of the rows and columns of the matrix A and in more traditional terms, this says that the columns form a linearly independent (spanning) sequence if and only if the rows form a spanning (linearly independent) sequence. This is a special case of the result that $\operatorname{null} A = (\operatorname{ran} A^{\mathrm{t}})^{\perp}$, hence that $\dim \operatorname{null} A = \operatorname{codim} \operatorname{ran} A^{\mathrm{t}}$. By going from A to A^{t}, and from a subspace to its orthogonal complement, we obtain from these the 'dual' result that $\operatorname{ran} A = (\operatorname{null} A^{\mathrm{t}})^{\perp}$, hence that $\dim \operatorname{ran} A = \operatorname{codim} \operatorname{null} A^{\mathrm{t}}$.

Recall from (3.11) the factorization $A = A(:, \texttt{bound})\operatorname{rrref}(A)$. It supplies the corresponding factorization $A^{\mathrm{t}} = A^{\mathrm{t}}(:, \texttt{rbound})\operatorname{rrref}(A^{\mathrm{t}})$ with `rbound` the index sequence of bound columns of A^{t}, i.e. of bound *rows* of A. By combining these two factorizations, we get the more symmetric factorization

$$A = (\operatorname{rrref}(A^{\mathrm{t}}))^{\mathrm{t}} A(\texttt{rbound}, \texttt{bound})\operatorname{rrref}(A),$$

which is called the **car**-factorization in [citGS].

9.1 Prove that, for any $A \in L(X, Y)$, $\operatorname{codim} \operatorname{null} A = \dim \operatorname{ran} A$.

9.2 In the list of pairs of complementary concepts, given at the beginning of this chapter, many of the pairs have been ordered so as to have the first term in each pair naturally correspond to the first term in any related pair.

For example, a right (left) inverse is necessarily 1-1 (onto).

Discover as many such correspondences as you can.

The dual of a vector space

The **dual of the vector space** X is, by definition, the vector space

$$X' := L(X, \mathbb{F})$$

of all linear maps into the underlying scalar field. Each such map is called a **linear functional** on X. (The term 'functional' is used to indicate a map, on a vector space, whose target is the underlying scalar field. Some books use the term 'form' instead.)

We have made much use of linear functionals, namely as the rows $\lambda_1, \ldots, \lambda_n$ of specific row maps (or data maps)

$$\Lambda^{\mathrm{t}} = [\lambda_1, \ldots, \lambda_n]^{\mathrm{t}} \in L(X, \mathbb{F}^n)$$

from the vector space X to n-dimensional coordinate space.

Example: If $X = \mathbb{F}^n$, then

$$X' = L(\mathbb{F}^n, \mathbb{F}) = \mathbb{F}^{1\times n} \sim \mathbb{F}^n,$$

and it has become standard to identify $(\mathbb{F}^n)'$ with $\mathbb{F}^n$ via

$$\mathbb{F}^n \to (\mathbb{F}^n)' : a \mapsto a^{\mathrm{t}}.$$

While this identification is often quite convenient, be aware that, strictly speaking, $\mathbb{F}^n$ and its dual are quite different objects. □

Here is a quick discussion of X' for an arbitrary finite-dimensional vector space, X. X being finite-dimensional, it has a basis, $V \in L(\mathbb{F}^n, X)$ say. Let

$$V^{-1} =: \Lambda^{\mathrm{t}} =: [\lambda_1, \ldots, \lambda_n]^{\mathrm{t}}$$

be its inverse. Each of its rows λ_i is a linear functional on X, hence

$$\Lambda := [\lambda_1, \ldots, \lambda_n]$$

is a column map into X'.

Λ is 1-1: Indeed, if $\Lambda a = 0$, then $\sum_i a(i)\lambda_i$ is the zero functional, hence, in particular, $\sum_i a(i)\lambda_i v_j = 0$ for all columns v_j of V. This implies that $0 = (\sum_i a(i)\lambda_i v_j : j = 1{:}n) = a^{\rm t}(\Lambda^{\rm t} V) = a^{\rm t}\,\mathrm{id}_n = a^{\rm t}$, hence $a = 0$.

It follows that $\dim \operatorname{ran} \Lambda = \dim \operatorname{dom} \Lambda = n$, hence we will know that Λ is also onto as soon as we know that the dimension of its target is $\le n$, i.e.,

$$\dim X' \le n.$$

For the proof of this inequality, observe that, for each $\lambda \in X'$, the composition λV is a linear map from $\mathbb{F}^n$ to $\mathbb{F}$, hence a 1-by-n matrix. Moreover, the resulting map

$$X' \to \mathbb{F}^{1\times n} \sim \mathbb{F}^n : \lambda \to \lambda V$$

is linear. It is also 1-1 since $\lambda V = 0$ implies that $\lambda = 0$ since V is invertible. Hence, indeed, $\dim X' \le n$.

(9.1) Proposition: For each basis V of the n-dimensional vector space X, the rows of its inverse, $V^{-1} =: \Lambda^{\rm t} =: [\lambda_1, \ldots, \lambda_n]^{\rm t}$, provide the columns for the basis $\Lambda = [\lambda_1, \ldots, \lambda_n]$ for X'. In particular, $\dim X' = \dim X$.

The two bases, Λ and V, are said to be **dual** or **bi-orthonormal** to signify that

$$\lambda_i v_j = \delta_{ij}, \qquad i, j = 1{:}n.$$

Here is the 'dual' claim.

(9.2)Proposition: Let X be an n-dimensional linear subspace of the vector space Y. Then, for each $\Lambda^{\rm t} \in L(Y, \mathbb{F}^n)$ that is 1-1 on X, there exists exactly one basis, V, for X that is dual or bi-orthonormal to Λ.

For every $\lambda \in Y'$, there exists exactly one $a \in \mathbb{F}^n$ so that

$$\lambda = \Lambda a \quad \text{on } X. \tag{9.3}$$

In particular, each $\lambda \in X'$ has a unique such **representation** Λa in $\operatorname{ran} \Lambda$.

Proof: Since $\dim X = \dim \operatorname{tar} \Lambda^{\rm t}$ and the restriction of $\Lambda^{\rm t} =: [\lambda_1, \ldots, \lambda_n]^{\rm t}$ to X is 1-1, it must be invertible, i.e., there exists exactly one basis V for X with $\Lambda^{\rm t} V = \mathrm{id}_n$, hence with Λ and V dual to each other.

In particular, $\Lambda := [\lambda_1, \ldots, \lambda_n]$ is a basis for its range. Let now $\lambda \in Y'$ and consider the equation

$$\Lambda ? = \lambda \qquad \text{on } X.$$

Since V is a basis for X, this equation is equivalent to the equation $(\Lambda ?)V = \lambda V$. Since

$$(\Lambda a)V = (\sum_i a(i)\lambda_i v_j : j = 1{:}n) = a^{\rm t}(\Lambda^{\rm t} V),$$

this equation, in turn, is equivalent to

$$?^{\rm t}\Lambda^{\rm t} V = \lambda V,$$

and, since $\Lambda^{\rm t} V = \mathrm{id}_n$, this has the unique solution $? = \lambda V = (\lambda v_j : j = 1{:}n)$. □

If X is not finite-dimensional, it may be harder to provide a complete description of its dual. In fact, in that case, one calls X' the **algebraic dual** and, for even some very common vector spaces, like $C([a \,..\, b])$, there is no constructive description for its algebraic dual. If X is a normed vector space, one focuses attention instead on its **topological dual**. The topological dual consists of all a *continuous* linear functionals on X, and this goes beyond the level of these notes. Suffice it to say that, for any finite-dimensional normed vector space, the algebraic dual coincides with the topological dual.

The very definition of $0 \in L(X, \mathbb{F})$ ensures that $\lambda \in X'$ is 0 if and only if $\lambda x = 0$ for all $x \in X$. What about its dual statement: *$x \in X$ is 0 if and only if $\lambda x = 0$ for all $\lambda \in X'$?* For an arbitrary vector space, this turns out to require the Axiom of Choice. However, if X is a linear subspace of $\mathbb{F}^T$ for some set T, then, in particular,

$$\delta_t : X \to \mathbb{F} : x \mapsto x(t)$$

is a linear functional on X, hence the vanishing at x of all linear functionals in X' implies that, in particular, $x(t) = 0$ for all $t \in T$, hence $x = 0$.

(9.4) Fact: For any x in the vector space X, $x = 0$ if and only if $\lambda x = 0$ for all $\lambda \in X'$.

Proof: If X is finite-dimensional, then, by (9.1), the condition $\lambda x = 0$ for all $\lambda \in X'$ is equivalent, for any particular basis V for X with dual basis Λ for X', to having $b^{\mathrm{t}}\Lambda^{\mathrm{t}}Va = 0$ for all $b \in \mathbb{F}^n$ and for $x =: Va$. Since $\Lambda^{\mathrm{t}}V = \mathrm{id}_n$, it follows that $a = \Lambda^{\mathrm{t}}Va$ must be zero, hence $x = 0$. □

Finally, one often needs the following

(9.5) Fact: Every linear functional on some linear subspace of a vector space can be extended to a linear functional on the whole vector space.

Proof: If X is a linear subspace of the finite-dimensional vector space Y, then there is a basis $[V, W]$ for Y with V a basis for X. If now $\lambda \in X'$, then there is a unique $\mu \in Y'$ with $\mu[V, W] = [\lambda V, 0]$, and it extends λ to all of Y.

If Y is not finite-dimensional, then it is, once again a job for the Axiom of Choice to aid in the proof. □

The dual of an inner product space

We introduced inner-product spaces as spaces with a ready supply of linear functionals. Specifically, the very definition of an inner product $\langle , \rangle$ on the vector space Y requires that, for each $y \in Y$, $y^{\mathrm{c}} := \langle \cdot, y \rangle$ be a linear functional on Y. This sets up a map

$$^{\mathrm{c}} : Y \to Y' : y \mapsto y^{\mathrm{c}}$$

from the inner product space to its dual. This map is additive. It is also homogeneous in case $\mathbb{F} = \mathbb{R}$. If $\mathbb{F} = \mathbb{C}$, then the map is **skew-homogeneous**, meaning that

$$(\alpha y)^{\mathrm{c}} = \overline{\alpha} y^{\mathrm{c}}, \qquad \alpha \in \mathbb{F},\ y \in Y.$$

Either way, this map is 1-1 if and only if its nullspace is trivial. But, since $y^{\mathrm{c}} = 0$ implies, in particular, that $y^{\mathrm{c}}y = 0$, the positive definiteness required of the inner product guarantees that then $y = 0$, hence the map $y \mapsto y^{\mathrm{c}}$ is 1-1.

If now $n := \dim Y < \infty$, then, by (9.1)Proposition, $\dim Y' = \dim Y = n$, hence, by the Dimension Formula, $y \mapsto y^c$ must also be onto. This proves

(9.6) Proposition: If Y is a finite-dimensional inner product space, then every $\lambda \in Y'$ can be written in exactly one way as $\lambda = y^c$ for some $y \in Y$.

We say in this case that y^c **represent**s λ.

If Y is not finite-dimensional, then the conclusion of this proposition still holds, provided we consider only the topological dual of Y and provided Y is 'complete', the very concept we declared beyond the scope of these notes when, earlier, we discussed the Hermitian (aka conjugate transpose) of a linear map between two inner product spaces.

The dual of a linear map

Any $A \in L(X, Y)$ induces in a natural way the linear map

$$A' : Y' \to X' : \lambda \mapsto \lambda A.$$

This map is called the **dual** to A.

If also $B \in L(Y, Z)$, then $BA \in L(X, Z)$ and, for every $\lambda \in Z'$, $\lambda(BA) = (\lambda B)A = A'(B'(\lambda))$, hence

$$(BA)' = A'B', \quad A \in L(X, Y),\ B \in L(Y, Z). \tag{9.7}$$

If both X and Y are coordinate spaces, hence A is a matrix, then, with the identification of a coordinate space with its dual, the dual of A coincides with its transpose i.e.,

$$A' = A^t, \qquad A \in \mathbb{F}^{m \times n} = L(\mathbb{F}^n, \mathbb{F}^m).$$

If $Y = \mathbb{F}^m$, hence A is a row map, $A = \Lambda^t = [\lambda_1, \ldots, \lambda_m]^t$ say, then, with the identification of $(\mathbb{F}^m)'$ with $\mathbb{F}^m$, $(\Lambda^t)'$ becomes the column map

$$(\Lambda^t)' = [\lambda_1, \ldots, \lambda_m] = \Lambda.$$

In this way, we now recognize a row map on X as the **pre-dual** of a column map into X'.

If $X = \mathbb{F}^n$, hence A is a column map, $A = V = [v_1, \ldots, v_n]$ say, then, with the identification of $(\mathbb{F}^n)'$ with $\mathbb{F}^n$, V' becomes a row map on Y', namely the row map that associates $\lambda \in Y'$ with the n-vector $(\lambda v_j : j = 1{:}n)$. Its rows are the linear functionals

$$Y' \to \mathbb{F} : \lambda \mapsto \lambda v_j$$

on Y' 'induced' by the columns of V. Each of these rows is therefore a linear functional on Y', i.e., an element of $(Y')'$, the **bidual** of Y. Also *if, in addition, V is 1-1, the V' is onto*. Indeed, in that case, V is a basis for its range, hence has an inverse, Λ^t say. Now, for arbitrary $b^t \in (\mathbb{F}^n)'$, $b^t = b^t(\Lambda^t V) = (\Lambda b)V$, with Λb a linear functional on $\operatorname{ran} V$. By (9.5)Fact, there is some $\lambda \in Y'$ that agrees with Λb on $\operatorname{ran} V$. In particular, $\lambda V = (\Lambda b)V = b^t$, showing that V' is, indeed onto.

Finally, if X and Y are arbitrary vector spaces but A is of finite rank, then, for any basis V for $\operatorname{ran} A$ with dual basis M, we have

$$A = V \mathrm{M}^t A =: V \Lambda^t,$$

and, by (8.2)Proposition, this is a minimal factorization for A. It follows that

$$A' = \Lambda V',$$

and, since V is 1-1, hence V' is onto, and also Λ is 1-1, we conclude that Λ is a basis for $\operatorname{ran} A'$, hence $\Lambda V'$ is a minimal factorization for $\operatorname{ran} A'$.

In particular, $\operatorname{rank} A' = \operatorname{rank} A$. Also, if A is onto, then A' 1-1.

10. The powers of a linear map and its spectrum

If $\operatorname{tar} A = \operatorname{dom} A$, then we can form the powers

$$A^k := \underbrace{AA\cdots A}_{k \text{ factors}}$$

of A. Here are some examples that show the importance of understanding the powers of a linear map.

Examples

Fixed-point iteration: A standard method for solving a large linear system $A? = y$ (with $A \in \mathbb{F}^{n\times n}$) is to split the matrix A suitably as

$$A = M - N$$

with M 'easily invertible', and to generate the sequence $x_0, x_1, x_2, \ldots$ of approximate solutions by the **iteration**

(10.1) $$x_k := M^{-1}(Nx_{k-1} + y), \quad k = 1, 2, \ldots.$$

Assuming this iteration to converge, with $x := \lim_{k\to\infty} x_k$ its limit, it follows that

(10.2) $$x = M^{-1}(Nx + y),$$

hence that $Mx = Nx + y$, therefore finally that $Ax = (M - N)x = y$, i.e., the limit solves our original problem $A? = y$.

Let $\varepsilon_k := x - x_k$ be the **error** in our kth approximate solution. Then on subtracting the iteration equation (10.1) from the exact equation (10.2), we find that

$$\varepsilon_k = x - x_k = M^{-1}(Nx + b\ - (Nx_{k-1} + y)) = M^{-1}N\varepsilon_{k-1}.$$

Therefore, by induction,

$$\varepsilon_k = B^k\varepsilon_0, \quad \text{with } B := M^{-1}N$$

the **iteration map**. Since we presumably don't know the solution x, we have no way of choosing the **initial guess** x_0 in any special way. For convergence, we must therefore demand that

$$\lim_{k\to\infty} B^k z = 0 \quad \text{for all } z \in \mathbb{F}^n.$$

It turns out that this will happen if and only if all eigenvalues of B are less than 1 in absolute value.

random walk: Consider a random walk on a graph G. The specifics of such a random walk are given by a **stochastic** matrix M of order n, with n the number of vertices in the graph. This means that all the entries of M are nonnegative, and all the entries in each row add up to 1, i.e.,

$$M \ge 0, \qquad Me = e,$$

with e the vector with all entries equal to 1,

$$e := (1, 1, 1, \ldots, 1).$$

The entries of M are interpreted as probabilities: $M_{i,j}$ gives the probability that, on finding ourselves at vertex i, we would proceed to vertex j. Thus, the probability that, after two steps, we would have gone from vertex i to vertex j is the sum of the probabilities that we would have gone from i to some k in the first step and thence to j in the second step, i.e., the number

$$\sum_k M_{i,k} M_{k,j} = M^2_{i,j}.$$

More generally, the probability that we have gone after m steps from vertex i to vertex j is the number $M^m_{i,j}$, i.e., the (i,j)-entry of the mth power of the matrix M.

A study of the powers of such a stochastic matrix reveals that, for large m, all the rows of M^m look more and more alike. Precisely, for each row i,

$$\lim_{m\to\infty} M^m_{i:} = x_\infty$$

for a certain (i-independent) vector x_∞ with nonnegative entries that sum to one; this is part of the so-called Perron-Frobenius Theory. In terms of the random walk, this means that, for large m, the probability that we will be at vertex j after m steps is more or less independent of the vertex we started off from. One can find this limiting probability distribution x_∞ as a properly scaled eigenvector of the transpose M^{t} of M belonging to the eigenvalue 1.

As the simple example $M = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ shows, the last paragraph isn't quite correct. Look for the discussion of the Perron-Frobenius theorem later in these notes (see pages 135ff).

polynomials in a map: Once we know the powers A^k of A, we can also construct polynomials in A, in the following way. If p is the polynomial

$$p : t \mapsto c_0 + c_1 t + c_2 t^2 + \cdots + c_k t^k,$$

then we define the linear map $p(A)$ to be what we get when we substitute A for t:

$$p(A) := c_0 \,\mathrm{id} + c_1 A + c_2 A^2 + \cdots + c_k A^k.$$

We can even consider power series. The most important example is the **matrix exponential**:

$$\exp(A) := \mathrm{id} + A + A^2/2 + A^3/6 + \cdots + A^k/k! + \cdots. \tag{10.3}$$

The matrix exponential is used in solving the first-order system

$$Dy(t) = Ay(t) \text{ for } t > 0, \quad y(0) = b \tag{10.4}$$

of constant-coefficient ordinary differential equations. Here A is a square matrix, of order n say, and $y(t)$ is an n-vector that depends on t. Further,

$$Dy(t) := \lim_{h\to 0} (y(t+h) - y(t))/h$$

is the first derivative at t of the vector-valued function y. One verifies that the particular function

$$y(t) := \exp(tA)b, \quad t \ge 0,$$

solves the differential equation (10.4). Practical application does require efficient ways for evaluating the power series

$$\exp((tA)) := \mathrm{id} + tA + (tA)^2/2 + (tA)^3/6 + \cdots + (tA)^k/k! + \cdots,$$

hence for computing the powers of tA.

Eigenvalues and eigenvectors

The calculation of $A^k x$ is simplest if A maps x to a scalar multiple of itself, i.e., if

$$Ax = \mu x = x\mu$$

for some scalar μ. For, in that case, $A^2 x = A(Ax) = A(x\mu) = Ax\mu = x\,\mu^2$ and, more generally,

$$(10.5) \qquad Ax = x\mu \quad \Longrightarrow \quad A^k x = x\,\mu^k, \quad k = 0, 1, 2, \ldots.$$

If $x = 0$, this will be so for any scalar μ. If $x \neq 0$, then this will be true for at most one scalar μ. That scalar is called an *eigenvalue for* A with associated *eigenvector* x.

(10.6) Definition: Let $A \in L(X)$. Any scalar μ for which there is a *nontrivial* vector $x \in X$ so that $Ax = x\mu$ is called an **eigenvalue** of A, with (μ, x) the corresponding **eigenpair**. The collection of all eigenvalues of A is called the **spectrum** of A and is denoted $\mathrm{spec}(A)$.
Thus

$$\mathrm{spec}(A) = \{\mu \in \mathbb{F} : \ A - \mu\,\mathrm{id} \ \text{ is not invertible}\}.$$

All the elements of $\mathrm{null}(A - \mu\,\mathrm{id})\backslash 0$ are called the **eigenvector**s of A *associated with* μ. The number

$$\rho(A) := \max|\,\mathrm{spec}(A)| = \max\{|\mu| : \mu \in \mathrm{spec}(A)\}$$

is called the **spectral radius of** A.

Since $\mu \in \mathrm{spec}(A)$ exactly when $(A - \mu\,\mathrm{id})$ is not invertible, this puts a premium on knowing whether or not a given linear map is invertible. We pointed out in Chapter 3 that the only matrices for which we could tell this at a glance are the triangular matrices. To recall, by (3.20)Proposition, a triangular matrix is invertible if and only if none of its diagonal entries is zero. Since $(A - \mu\,\mathrm{id})$ is triangular for any μ in case A is triangular, this gives the important

(10.7) Proposition: For any triangular matrix of order n, $\mathrm{spec}(A) = \{A_{jj} : n = 1{:}n\}$.

In the best of circumstances, there is an entire basis $V = [v_1, v_2, \ldots, v_n]$ for $X = \mathrm{dom}\,A$ consisting of eigenvectors for A. In this case, it is very easy to compute $A^k x$ for any $x \in X$. For, in this situation, $Av_j = v_j\mu_j$, $j = 1{:}n$, hence

$$AV = [Av_1, \ldots, Av_n] = [v_1\mu_1, \ldots, v_n\mu_n] = V\mathrm{M},$$

with M the *diagonal* matrix

$$\mathrm{M} := \mathrm{diag}(\mu_1, \mu_2, \ldots, \mu_n).$$

Therefore, for any k,

$$A^k V = V\mathrm{M}^k = V\,\mathrm{diag}(\mu_1^k, \ldots, \mu_n^k).$$

Also, since V is a basis for X, any $x \in X$ can be written (uniquely) as $x = Va$ for some n-vector a and thus

$$A^k x = A^k Va = V\mathrm{M}^k a = v_1\mu_1^k a_1 + v_2\mu_2^k a_2 + \cdots + v_n\mu_n^k a_n$$

for any k. For example, for such a matrix and for any t,

$$\exp(tA) = V\exp(t\mathrm{M})V^{-1} = V\,\mathrm{diag}(\ldots, \exp(t\mu_j), \ldots)V^{-1}.$$

To be sure, if A is not 1-1, then at least one of the μ_j must be zero, but this doesn't change the fact that M is a diagonal matrix.

(10.8) Example: The matrix $A := \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$ maps the 2-vector $x := (1,1)$ to $3x$ and the 2-vector $y := (1,-1)$ to itself. Hence, $A[x,y] = [3x,y] = [x,y]\operatorname{diag}(3,1)$ or

$$A = V\operatorname{diag}(3,1)V^{-1}, \quad \text{with} \;\; V := [x,y] = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}.$$

Elimination gives

$$[V, \mathrm{id}] = \begin{bmatrix} 1 & 1 & 1 & 0 \\ 1 & -1 & 0 & 1 \end{bmatrix} \to \begin{bmatrix} \mathbf{1} & 1 & 1 & 0 \\ 0 & -2 & -1 & 1 \end{bmatrix} \to \begin{bmatrix} 1 & 0 & 1/2 & 1/2 \\ 0 & \mathbf{-2} & -1 & 1 \end{bmatrix} \to \begin{bmatrix} 1 & 0 & 1/2 & 1/2 \\ 0 & 1 & 1/2 & -1/2 \end{bmatrix},$$

hence

$$V^{-1} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}/2.$$

It follows that, for any k,

$$A^k = V\operatorname{diag}(3^k,1)V^{-1} = \begin{bmatrix} 3^k & 1 \\ 3^k & -1 \end{bmatrix}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}/2 = \begin{bmatrix} 3^k+1 & 3^k-1 \\ 3^k-1 & 3^k+1 \end{bmatrix}/2.$$

In particular,

$$A^{-1} = \begin{bmatrix} 1/3+1 & 1/3-1 \\ 1/3-1 & 1/3+1 \end{bmatrix}/2 = \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix}/3.$$

Also,

$$\exp(tA) = V\operatorname{diag}(\mathrm{e}^{3t},\mathrm{e}^{t})V^{-1} = \begin{bmatrix} \mathrm{e}^{3t}+\mathrm{e}^{t} & \mathrm{e}^{3t}-\mathrm{e}^{t} \\ \mathrm{e}^{3t}-\mathrm{e}^{t} & \mathrm{e}^{3t}+\mathrm{e}^{t} \end{bmatrix}.$$

□

10.1 Let $A = \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix}$. (i) Find a basis V and a diagonal matrix M so that $A = V\mathrm{M}V^{-1}$. (ii) Determine the matrix $\exp(A)$.

10.2 Let $A = \begin{bmatrix} 4 & 1 & -1 \\ 2 & 5 & -2 \\ 1 & 1 & 2 \end{bmatrix}$.

Use elimination to determine *all* eigenvectors for this A belonging to the eigenvalue 3, and all eigenvectors belonging to the eigenvalue 5. (It is sufficient to give a basis for $\operatorname{null}(A - 3\,\mathrm{id})$ and for $\operatorname{null}(A - 5\,\mathrm{id})$.)

10.3 If A is a triangular matrix, then one of its eigenvectors can be determined without any calculation. Which one?

10.4

(a) Prove that the matrix $A = \begin{bmatrix} 4 & 1 & -1 \\ 2 & 5 & -2 \\ 1 & 1 & 2 \end{bmatrix}$ maps the vector space $Y := \operatorname{ran} V$ with $V := \begin{bmatrix} 0 & 2 \\ 3 & 1 \\ 1 & 1 \end{bmatrix}$ into itself, hence the **restriction** of A to Y, i.e.,

$$A_{|Y} := B : Y \to Y : y \mapsto Ay$$

is a well-defined linear map. (You will have to verify that $\operatorname{ran} AV \subseteq \operatorname{ran} V$; looking at $\operatorname{rref}([V\ AV])$ should help.)

(b) Determine the matrix representation of B with respect to the basis V for $\operatorname{dom} B = Y$, i.e., compute the matrix $V^{-1}BV$. (Hint: (5.4)Example tells you how to read off this matrix from the calculations in (a).)

(c) Determine the spectrum of the linear map $B = A_{|Y}$ defined in (a). (Your answer in (b) could be helpful here since similar maps have the same spectrum.)

10.5 Prove that 0 is the only eigenvalue of the matrix $A = \begin{bmatrix} 0 & 1 & 2 \\ 0 & 0 & 3 \\ 0 & 0 & 0 \end{bmatrix}$ and that, up to scalar multiples, e_1 is the only eigenvector for A.

10.6 Let $\mu \in \operatorname{spec}(A)$ (hence $Ax = \mu x$ for some $x \neq 0$). Prove:

(i) For any scalar α, $\alpha\mu \in \operatorname{spec}(\alpha A)$.

(ii) For any scalar α, $\mu + \alpha \in \operatorname{spec}(A + \alpha\,\mathrm{id})$.

(iii) For any natural number k, $\mu^k \in \operatorname{spec}(A^k)$.

(iv) If A is invertible, then $\mu \neq 0$ and $\mu^{-1} \in \operatorname{spec}(A^{-1})$.

(v) If A is a matrix, then $\mu \in \operatorname{spec}(A^{\mathrm{t}})$ and $\overline{\mu} \in \operatorname{spec}(A^{\mathrm{c}})$.

Diagona(liza)bility

Definition: A linear map $A \in L(X)$ is called **diagona(liza)ble** if it has an **eigenbasis**, i.e., if there is a basis for its domain X consisting entirely of eigenvectors for A.

(10.9) Lemma: If V_μ is a basis for $\mathrm{null}(A - \mu\,\mathrm{id})$, then $[V_\mu : \mu \in \mathrm{spec}(A)]$ is 1-1.

Proof: Note that, for any $\mu \in \mathrm{spec}(A)$ and any ν,

$$(A - \nu\,\mathrm{id})V_\mu = (\mu - \nu)V_\mu,$$

and, in particular, $(A - \mu\,\mathrm{id})V_\mu = 0$. Hence, if $\sum_\mu V_\mu a_\mu = 0$, then, for each $\mu \in \mathrm{spec}(A)$, after applying to both sides of this equation the product of all $(A - \nu\,\mathrm{id})$ with $\nu \in \mathrm{spec}(A)\backslash\mu$, we are left with the equation $(\prod_{\nu\neq\mu}(\mu - \nu))V_\mu a_\mu = 0$, and this implies that $a_\mu = 0$ since V_μ is 1-1 by assumption. In short, $[V_\mu : \mu \in \mathrm{spec}(A)]a = 0$ implies $a = 0$. □

(10.10) Corollary: $\#\mathrm{spec}(A) \le \dim\mathrm{dom}\, A$, with equality only if A is diagonalizable.

(10.11) Proposition: A linear map $A \in L(X)$ is diagonalizable if and only if

(10.12) $$\dim X = \sum_{\mu\in\mathrm{spec}(A)} \dim\mathrm{null}(A - \mu\,\mathrm{id}).$$

Proof: By (10.9)Lemma, (10.12) implies that $\mathrm{dom}\, A$ has a basis consisting of eigenvectors for A.

Conversely, if V is a basis for $X = \mathrm{dom}\, A$ consisting entirely of eigenvectors for A, then $A = V\mathrm{M}V^{-1}$ for some diagonal matrix $\mathrm{M} =: \mathrm{diag}(\mu_1, \dots, \mu_n)$, hence, for any scalar μ, $(A - \mu\,\mathrm{id}) = V(\mathrm{M} - \mu\,\mathrm{id})V^{-1}$. In particular, $\mathrm{null}(A - \mu\,\mathrm{id}) = \mathrm{ran}[v_j : \mu = \mu_j]$, hence $\sum_{\mu\in\mathrm{spec}(A)} \dim\mathrm{null}(A - \mu\,\mathrm{id}) = \sum_{\mu\in\mathrm{spec}(A)} \#\{j : \mu_j = \mu\} = n = \#V = \dim X$. □

(10.11)Proposition readily identifies a circumstance under which A is *not* diagonable, namely when $\mathrm{null}(A-\mu\,\mathrm{id})\cap\mathrm{ran}(A-\mu\,\mathrm{id}) \neq \{0\}$ for some μ. For, with V_ν a basis for $\mathrm{null}(A-\nu\,\mathrm{id})$ for any $\nu \in \mathrm{spec}(A)$, we compute $AV_\nu = \nu V_\nu$, hence $(A-\mu\,\mathrm{id})V_\nu = (\nu-\mu)V_\nu$ and therefore, for any $\nu \neq \mu$, $V_\nu = (A-\mu\,\mathrm{id})V_\nu/(\nu-\mu) \subset \mathrm{ran}(A - \mu\,\mathrm{id})$. This places all the columns of the 1-1 map $V_{\backslash\mu} := [V_\nu : \nu \neq \mu]$ in $\mathrm{ran}(A - \mu\,\mathrm{id})$ while, by (10.9)Lemma, $\mathrm{ran}\, V_\mu \cap \mathrm{ran}\, V_{\backslash\mu}$ is trivial. Hence, if $\mathrm{ran}\, V_\mu = \mathrm{null}(A - \mu\,\mathrm{id})$ has nontrivial intersection with $\mathrm{ran}(A - \mu\,\mathrm{id})$, then $\mathrm{ran}\, V_{\backslash\mu}$ cannot be all of $\mathrm{ran}(A - \mu\,\mathrm{id})$, and therefore

$$\sum_{\nu\neq\mu} \dim\mathrm{null}(A - \nu\,\mathrm{id}) = \#V_{\backslash\mu} \quad < \quad \dim\mathrm{ran}(A - \mu\,\mathrm{id}) = \dim X - \dim\mathrm{null}(A - \mu\,\mathrm{id}),$$

hence, by (10.11)Proposition, such A is not diagonable.

This has motivated the following

Definition: The scalar μ is a **defective eigenvalue** of A if

$$\operatorname{null}(A - \mu\,\mathrm{id}) \cap \operatorname{ran}(A - \mu\,\mathrm{id}) \neq \{0\}.$$

Any such μ certainly is an eigenvalue (since, in particular, $\operatorname{null}(A-\mu\,\mathrm{id}) \neq \{0\}$), but I don't care for such *negative labeling*; if it were up to me, I would call such μ an **interesting eigenvalue**, since the existence of such eigenvalues makes for a richer theory. Note that, by (4.18)Proposition, μ is a defective eigenvalue for A iff, for some, hence for every, bases V and W for $\operatorname{ran}(A - \mu\,\mathrm{id})$ and $\operatorname{null}(A - \mu\,\mathrm{id})$ respectively, $[V, W]$ is not 1-1.

(10.13) Corollary: If A has a defective eigenvalue, then A is not diagonable.

10.7 Prove: if $A \in L(X)$ is diagonalizable and $\#\mathrm{spec}(A) = 1$, then $A = \mu\,\mathrm{id}_X$ for some $\mu \in \mathbb{F}$.

10.8 What is a simplest matrix A with $\mathrm{spec}(A) = \{1, 2, 3\}$?

10.9 For each of the following matrices $A \in \mathbb{F}^{2\times 2}$, determine whether or not 0 is a defective eigenvalue (give a reason for your answer). For a mechanical approach, see H.P. 4.20 . (a) $A = 0$. (b) $A = \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix}$. (c) $A = \begin{bmatrix} -2 & -1 \\ 4 & 2 \end{bmatrix}$. (d) $A = \mathrm{id}_2$.

10.10 Prove that, *for every linear map A on the finite-dimensional vector space X, if A is diagonable, then so is $p(A)$ for every polynomial p.*

10.11 Prove that *any linear projector P on a finite-dimensional vector space X is diagonalizable.* (Hint: Show that, for any basis U for $\operatorname{ran} P$ and any basis W for $\operatorname{null} P$, $V := [U, W]$ is a basis for X, and that all the columns of V are eigenvectors for P. All of this should follow from the fact that $P^2 = P$.)

10.12 Prove that any linear involutory map R on a finite-dimensional vector space X is diagonalizable. (Hint: H.P. 5.10 .)

Are all square matrices diagonable?

By (10.13)Corollary, this will be so only if all square matrices have only nondefective eigenvalues.

(10.14) Example: The simplest example of a matrix with a defective eigenvalue is provided by the matrix

$$N := \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} = [0, e_1].$$

By (10.7)Proposition, $\mathrm{spec}(N) = \{0\}$. Yet $\operatorname{null} N = \operatorname{ran}[e_1] = \operatorname{ran} N$, hence the only eigenvalue of N is defective, and N fails to be diagonable, by (10.13)Corollary.

Of course, for this simple matrix, one can see directly that it cannot be diagonable, since, if it were, then some basis V for $\mathbb{R}^2$ would consist entirely of eigenvectors for the sole eigenvalue, 0, for N, hence, for this basis, $NV = 0$, therefore $N = 0$, contrary to fact. □

We will see shortly that, on a finite-dimensional vector space over the complex scalars, almost all linear maps are diagonable, and all linear maps are almost diagonable.

Does every square matrix have an eigenvalue?

Since an eigenvalue for A is any *scalar* μ for which $\text{null}(A - \mu\,\text{id})$ is not trivial, the answer necessarily depends on what we mean by a scalar.

If we only allow *real* scalars, i.e., if $\mathbb{F} = \mathbb{R}$, then not every matrix has eigenvalues. The simplest example is a rotation of the plane, e.g., the matrix

$$A := \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} = [e_2, -e_1].$$

This linear map rotates every $x \in \mathbb{R}^2$ 90 degrees counter-clockwise, hence the only vector x mapped by it to a scalar multiple of itself is the zero vector. In other words, this linear map has no eigenvectors, hence no eigenvalues.

The situation is different when we also allow *complex* scalars, i.e., when $\mathbb{F} = \mathbb{C}$, and this is the reason why we considered complex scalars all along in these notes. Now every (square) matrix has eigenvalues, as follows from the following simple argument.

(10.15) Theorem: Any linear map A on some nontrivial finite-dimensional vector space X over the *complex* scalar field $\mathbb{F} = \mathbb{C}$ has eigenvalues.

Proof: Let $n := \dim X$, pick any $x \in X\backslash 0$ and consider the column map

$$V := [x, Ax, A^2x, \ldots, A^nx].$$

Since $\#V > \dim \operatorname{tar} V$, V cannot be 1-1. This implies that some column of V is free. Let A^dx be the first free column, i.e., the first column that is in the range of the columns preceding it. Then $\text{null}\, V$ contains exactly one vector of the form

$$a = (a_0, a_1, \ldots, a_{d-1}, 1, 0, \cdots, 0),$$

and this is the vector we choose. Then, writing the equation $Va = 0$ out in full, we get

$$a_0x + a_1Ax + \cdots + a_{d-1}A^{d-1}x + A^dx \;=\; 0. \qquad (10.16)$$

Now here comes the trick: Consider the *polynomial*

$$p : t \mapsto a_0 + a_1t + \cdots + a_{d-1}t^{d-1} + t^d. \qquad (10.17)$$

Then, substituting for t our map A, we get the linear map

$$p(A) := a_0\,\text{id} + a_1A + \cdots + a_{d-1}A^{d-1} + A^d.$$

With this, (10.16) can be written, very concisely,

$$p(A)x = 0.$$

This is not just notational convenience. Since $a_d = 1$, p isn't the zero polynomial, and since $x \neq 0$, d must be greater than 0, i.e., p cannot be just a constant polynomial. Thus, by the *Fundamental Theorem of Algebra*, p has zeros. More precisely,

$$p(t) = (t - z_1)(t - z_2)\cdots(t - z_d)$$

for certain (possibly *complex*) scalars $z_1, \ldots, z_d$. This implies (see (10.19)Lemma below) that

$$p(A) = (A - z_1\,\text{id})(A - z_2\,\text{id})\cdots(A - z_d\,\text{id}).$$

Now, $p(A)$ is not 1-1 since it maps the nonzero vector x to zero. Therefore, *not all the maps* $(A - z_j\,\text{id})$, $j = 1{:}d$, *can be 1-1.* In other words, for some j, $(A - z_j\,\text{id})$ fails to be 1-1, i.e., has a nontrivial nullspace, and that makes z_j an eigenvalue for A. □

(10.18) Example: Let's try this out on our earlier example, the rotation matrix

$$A := [e_2, -e_1].$$

Choosing $x = e_1$, we have

$$[x, Ax, A^2x] = [e_1, e_2, -e_1],$$

hence the first free column is $A^2x = -e_1$, and, by inspection,

$$x + A^2x = 0.$$

Thus the polynomial of interest is

$$p : t \mapsto 1 + t^2 = (t - \mathrm{i})(t + \mathrm{i}),$$

with

$$\mathrm{i} := \sqrt{-1}$$

the *imaginary unit* (see the Backgrounder on complex numbers). In fact, we conclude that, with $y := (A + \mathrm{i}\,\mathrm{id})x$, $(A - \mathrm{i}\,\mathrm{id})y = p(A)x = 0$, while $y = Ae_1 + \mathrm{i}e_1 = e_2 + \mathrm{i}e_1 \neq 0$, showing that $(\mathrm{i}, e_2 + \mathrm{i}e_1)$ is an eigenpair for this A.

Polynomials in a linear map

The proof of (10.15)Theorem uses in an essential way the following fact.

(10.19) Lemma: If r is the product of the polynomials p and q, i.e., $r(t) = p(t)q(t)$ for all t, then, for any linear map $A \in L(X)$,

$$r(A) = p(A)q(A) = q(A)p(A).$$

Proof: If you wanted to check that $r(t) = p(t)q(t)$ for the polynomials r, p, q, then you would multiply p and q term by term, collect like terms and then compare coefficients with those of r. For example, if $p(t) = t^2 + t + 1$ and $q(t) = t - 1$, then

$$p(t)q(t) = (t^2 + t + 1)(t - 1) = t^2(t - 1) + t(t - 1) + (t - 1) = t^3 - t^2 + t^2 - t + t - 1 = t^3 - 1,$$

i.e., the product of these two polynomials is the polynomial r given by $r(t) = t^3 - 1$. The only facts you use are: (i) free reordering of terms (commutativity of addition), and (ii) things like $tt = t^2$, i.e., the fact that

$$t^i t^j = t^{i+j}.$$

Both of these facts hold if we replace t by A. □

Here is a further use of this lemma. We now prove that the polynomial p constructed in the proof of (10.15) has the property that every one of its roots is an eigenvalue for A. This is due to the fact that we constructed it in the form (10.17) with d the *smallest* integer for which $A^d x \in \mathrm{ran}[x, Ax, \ldots, A^{d-1}x]$. Thus, with μ any zero of p, we can write

(10.20) $$p(t) = (t - \mu)q(t)$$

for some polynomial q necessarily of the form

$$q(t) = b_0 + b_1 t + \cdots + b_{d-2}t^{d-2} + t^{d-1}.$$

The crucial point here is that q is of degree $< d$. This implies that $q(A)x \neq 0$ since, otherwise, $(b_0, b_1, \ldots, 1)$ would be a nontrivial vector in $\mathrm{null}[x, Ax, \ldots, A^{d-1}x]$ and this would contradict the choice of d as the index for which $A^d x$ is the *first* free column in $[x, Ax, A^2, \ldots]$. Since

$$0 = p(A)x = (A - \mu\,\mathrm{id})q(A)x,$$

it follows that μ is an eigenvalue for A with associated eigenvector $q(A)x$.

This is exactly how we got an eigenvector for the eigenvalue i in (10.18)Example.

(10.21) Example: As another example, consider again the matrix $A = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$ from (10.8)Example. We choose $x = e_1$ and consider

$$[x, Ax, \ldots, A^n x] = [e_1, Ae_1, A(Ae_1)] = \begin{bmatrix} 1 & 2 & 5 \\ 0 & 1 & 4 \end{bmatrix}.$$

Since $[e_1, Ae_1, A^2 e_1]$ is in row echelon form, we conclude that the first two columns are bound. Elimination gives the rref

$$\begin{bmatrix} 1 & 0 & -3 \\ 0 & 1 & 4 \end{bmatrix},$$

hence $(3, -4, 1) \in \operatorname{null}[e_1, Ae_1, A^2 e_1]$. Correspondingly, $p(A)e_1 = 0$, with

$$p(t) = 3 - 4t + t^2 = (t-3)(t-1).$$

Consequently, $\mu = 3$ is an eigenvalue for A, with corresponding eigenvector

$$(A - \mathrm{id})e_1 = (1, 1);$$

also, $\mu = 1$ is an eigenvalue for A, with corresponding eigenvector

$$(A - 3\,\mathrm{id})e_1 = (-1, 1).$$

Note that the resulting basis $\begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$ for $\mathbb{F}^2$ consisting of eigenvectors for A differs in some detail from the one we found in (10.8)Example. After all, if v is an eigenvector, then so is αv for any scalar α. □

Here is some standard language concerning the items in our discussion so far. One calls any nontrivial polynomial r for which $r(A)x = 0$ an **annihilating polynomial for** A **at** x. We may assume without loss of generality that this polynomial is **monic**, i.e., its highest nonzero coefficient is 1, since we can always achieve this by dividing the polynomial by its highest nonzero coefficient without changing the fact that it is an annihilating polynomial for A at x. If such a polynomial is of exact degree k, say, then it has the form

$$r(t) = b_0 + b_1 t + \cdots + b_{k-1} t^{k-1} + t^k.$$

Since $r(A)x = 0$, we conclude that

$$b_0 x + b_1 Ax + \cdots + b_{k-1} A^{k-1} x + A^k x = 0.$$

In particular, $A^k x$ is in $\operatorname{ran}[x, Ax, \ldots, A^{k-1}x]$, i.e., the column $A^k x$ of $[x, Ax, A^2 x, \ldots]$ is free. This implies that $k \ge d$, with d the degree of the polynomial p constructed in the proof of (10.15)Theorem. For, there we chose d as the *smallest* index for which $A^d x$ is a free column of $[x, Ax, A^2, \ldots]$. In particular, all prior columns must be bound. This makes p the unique monic polynomial of smallest degree for which $p(A)x = 0$.

Here, for the record, is a formal account of what we have proved.

(10.22) Proposition: For every $A \in L(X)$ with $\dim X < \infty$ and every $x \in X \backslash 0$, there is a unique monic polynomial p of smallest degree for which $p(A)x = 0$. This polynomial is called the **minimal polynomial for** A **at** x and is denoted

$$p_{A,x}.$$

It can be constructed in the form

$$p_{A,x}(t) = a_0 + a_1 t + \cdots + a_{d-1}t^{d-1} + t^d,$$

with d the smallest index for which $A^d x$ is a free column of $[x, Ax, A^2x, \ldots]$. Moreover, $(a_0, a_1, \ldots, 1)$ is the unique vector in $\mathrm{null}[x, Ax, \ldots, A^d x]$ with its last entry equal to 1.

Assuming that X is a vector space over $\mathbb{F} = \mathbb{C}$, every zero μ of $p_{A,x}$ is an eigenvalue of A, with associated eigenvector $q(A)x$, where $p_{A,x}(t) =: (t - \mu)q(t)$. (See the Backgrounder on Horner's method for the standard way to compute q from $p_{A,x}$ and μ.)

For *example*, consider the *permutation matrix* $P = [e_2, e_3, e_1]$ and take $x = e_1$. Then

$$[x, Px, P^2x, P^3x] = [e_1, e_2, e_3, e_1].$$

Hence, P^3x is the first free column here. The element in the nullspace corresponding to it is the vector $(-1, 0, 0, 1)$. Hence, the minimal polynomial for P at $x = e_1$ is of degree 3; it is the polynomial $p(t) = t^3 - 1$. It has the zero $\mu = 1$, which therefore is an eigenvalue of P. A corresponding eigenvector is obtainable in the form $q(P)e_1$ with $q(t) := p(t)/(t-1) = t^2 + t + 1$, hence the eigenvector is $e_3 + e_2 + e_1$.

10.13 Use Elimination as in (10.21) to determine all the eigenvalues and, for each eigenvalue, a corresponding eigenvector, for each of the following matrices: (i) $\begin{bmatrix} 7 & -4 \\ 5 & -2 \end{bmatrix}$; (b) $[0, e_1, e_2] \in \mathbb{R}^{3\times 3}$ (try $x = e_3$); (iii) $\begin{bmatrix} -1 & 1 & -3 \\ 20 & 5 & 10 \\ 2 & -2 & 6 \end{bmatrix}$.

10.14

(a) Prove: If p is any nontrivial polynomial and A is any square matrix for which $p(A) = 0$, then $\mathrm{spec}(A) \subseteq \{\mu \in \mathbb{C} : p(\mu) = 0\}$. (Hint: prove first that, for any eigenvector x for A with eigenvalue μ and any polynomial p, $p(A)x = p(\mu)x$.)

(b) What can you conclude about $\mathrm{spec}(A)$ in case you know that A is *idempotent*, i.e., a linear projector, i.e., $A^2 = A$?

(c) What can you conclude about $\mathrm{spec}(A)$ in case you know that A is **nilpotent**, i.e., $A^q = 0$ for some integer q?

(d) What can you conclude about $\mathrm{spec}(A)$ in case you know that A is **involutory**, i.e., $A^{-1} = A$?

(e) What is the spectrum of the linear map $D : \Pi_k \to \Pi_k$ of differentiation, as a map on polynomials of degree $\le k$?

10.15 The **companion matrix** for the monic polynomial $p : t \mapsto a_1 + a_2 t + \cdots + a_n t^{n-1} + t^n$ is, by definition, the matrix $A_p := [e_2, \ldots, e_n, -a] \in \mathbb{F}^{n\times n}$. (a) Prove that p is the minimal polynomial for A at e_1. (b) Use (a) and MATLAB's `eig` command to find all the zeros of the polynomial $p : t \mapsto 1 + t + t^2 + \cdots + t^9$. Check your answer.

10.16 Use the minimal polynomial at e_1 to determine the spectrum of the following matrices: (i) $[e_2, 0]$; (ii) $[e_2, e_3, e_1]$; (iii) $[e_2, e_2]$; (iv) $[e_2, e_1, 2e_3]$.

10.17 Prove: (i) for any $A, B \in L(X)$, $\mathrm{null}\, A \cap \mathrm{null}\, B \subset \mathrm{null}(A + B)$. (ii) for any $A, B \in L(X)$ with $AB = BA$, $\mathrm{null}\, A + \mathrm{null}\, B \subset \mathrm{null}(AB)$. (iii) If g is the greatest common divisor of the nontrivial polynomials $p_1, \ldots, p_r$ and m is their smallest common multiple, then, for any $A \in L(X)$, $\mathrm{null}\, g(A) = \cap_j p_j(A)$ and $\mathrm{null}\, m(A) = \sum_j \mathrm{null}\, p_j(A)$.

10.18 Let A be a matrix of order n, let $b \in \mathbb{F}^n \backslash 0$, and let P be the orthogonal projector of $\mathbb{F}^n$ onto the space $Y := \mathrm{ran}[b, Ab, \ldots, A^{r-1}b]$, the **Krylov subspace of order** r **for** A **generated by** b. Assume that Y is r-dimensional, and let $PA^r b =: \sum_{j<r} a_j A^j b$. (i) Prove that $K := [x, PAx, (PA)^2x, \ldots, (PA)^r x] = [x, Ax, \ldots, A^{r-1}x, PA^r x]$. (ii) Prove that $q(t) := t^r - \sum_{j<r} a_j t^j$ is the minimal polynomial at x for the linear map $PA : Y \to Y : y \mapsto PAy$. (iii) Conclude that q is the unique monic polynomial of degree r for which $\|q(A)x\|_2$ is as small as possible.

It is enough to understand the eigenstructure of matrices

So far, we know how to find *some* eigenvalues and corresponding eigenvectors for a given $A \in L(X)$, making use of minimal polynomials found by elimination. But can we be sure to find all the eigenvalues that way? By (10.10)Corollary, we know that we have found them all if we have found $n := \dim X$ of them. But if we find fewer than that, then we can't be sure.

The standard approach to finding the entire spectrum of A is by searching for linear maps B that have the same spectrum as A but carry that spectrum more openly, like triangular matrices (see (10.7)Proposition). This search makes essential use of the notion of similarity.

Definition: We say that $A \in L(X)$ and $B \in L(Y)$ are **similar to each other** and write

$$A \sim B$$

in case there is an invertible $V \in L(Y, X)$ so that

$$A = VBV^{-1}.$$

In particular, *a linear map is diagonable if and only if it is similar to a diagonal matrix*.

In trying to decide whether or not a given linear map A is diagonable, it is sufficient to decide this question for any convenient linear map B similar to A. For, if such a B is diagonable, i.e., similar to a diagonal matrix, then A is similar to that very same diagonal matrix. This follows from the fact that similarity is an equivalence relation:

(10.23) Proposition: Similarity is an **equivalence relation**. Specifically,

(i) $A \sim A$ (**reflexive**);

(ii) $A \sim B$ implies $B \sim A$ (**symmetric**);

(iii) $A \sim B$ and $B \sim C$ implies $A \sim C$ (**transitive**).

Proof: Certainly, $A \sim A$, since $A = \mathrm{id} A \,\mathrm{id}$. Also, if $A = VBV^{-1}$ for some invertible V, then also $W := V^{-1}$ is invertible, and $B = V^{-1}AV = WAW^{-1}$. Finally, if $A = VBV^{-1}$ and $B = WCW^{-1}$, then $U := VW$ is also invertible, and $A = V(WCW^{-1})V^{-1} = UCU^{-1}$. □

Now, any linear map $A \in L(X)$ on a *finite-dimensional* vector space X is similar (in many ways if X is not trivial) to a *matrix*. Indeed, for any basis V for X, $\widehat{A} := V^{-1}AV$ is a matrix similar to A. The map $\widehat{A}$ so defined is indeed a matrix since both its domain and its target is a coordinate space (the same one, in fact; hence $\widehat{A}$ is a *square* matrix). We conclude that, in looking for ways to decide whether or not a linear map is diagonable, it is sufficient to know how to do this for square *matrices*.

Every complex (square) matrix is similar to an upper triangular matrix

While having in hand a diagonal matrix similar to a given $A \in L(X)$ is very nice indeed, for most purposes it is sufficient to have in hand an *upper triangular* matrix similar to A. There are several reasons for this.

One reason is that, as soon as we have an upper triangular matrix similar to A, then we can easily manufacture from this a matrix similar to A and with off-diagonal elements as small as we please (except that, in general, we can't make them all zero).

A more fundamental reason is that, once we have an upper triangular matrix similar to A, then we know the entire spectrum of A since, by (10.7)Proposition, the spectrum of a triangular matrix is the set of its diagonal entries. Here are the various facts.

(10.24) Proposition: If A and $\widehat{A}$ are similar, then $\mathrm{spec}(A) = \mathrm{spec}(\widehat{A})$.

Proof: If $\widehat{A} = V^{-1}AV$ for some invertible V, then, for any scalar μ, $\widehat{A} - \mu\,\mathrm{id} = V^{-1}(A - \mu\,\mathrm{id})V$. In particular, $\widehat{A} - \mu\,\mathrm{id}$ is not invertible (i.e., $\mu \in \mathrm{spec}(\widehat{A})$) if and only if $A - \mu\,\mathrm{id}$ is not invertible (i.e., $\mu \in \mathrm{spec}(A)$). □

(10.25) Corollary: If $A \in L(X)$ is similar to a triangular matrix $\widehat{A}$, then μ is an eigenvalue for A if and only if $\mu = \widehat{A}_{j,j}$ for some j. In a formula,

$$\mathrm{spec}(A) = \{\widehat{A}_{j,j} : \text{ all } j\}.$$

More precisely, if $\widehat{A} = V^{-1}AV$ is upper triangular and j is the smallest index for which $\mu = \widehat{A}_{j,j}$, then there is an eigenvector for A belonging to μ available in the form $w = Va$, with a the element in the standard basis for $\mathrm{null}(\widehat{A} - \mu\,\mathrm{id})$ associated with the (free) jth column, i.e., $a \in \mathrm{null}(\widehat{A} - \mu\,\mathrm{id})$, $a_j = 1$, and all other entries corresponding to free columns are 0; cf. (3.9).

The now-standard algorithm for computing the eigenvalues of a given matrix A is the **QR method**. It generates a sequence $B_1, B_2, B_3, \ldots$ of matrices all similar to A that converges to an upper triangular matrix. To the extent that the lower triangular entries of B_k are small (compared to $\|B_k\|$, say), the diagonal entries of B_k are close to eigenvalues of B_k, hence of A. The actual version of the QR method used in `MATLAB` is quite sophisticated, as much care has gone into making the algorithm reliable in the presence of round-off as well as fast.

The `MATLAB` command `eig(`A`)` gives you the list of eigenvalues of A. The more elaborate command `[V,M]=eig(`A`)` gives you, in `V`, a list of corresponding 'eigenvectors', in the sense that, approximately, A`V`$(:,j) =$ `V`$(:,j)$`M`(j,j), all j. □

(10.26) Theorem: Every complex (square) matrix is similar to an upper triangular matrix.

Proof: The proof is by induction on the order, n, of the given matrix A.

If $n = 1$, then A is a 1×1-matrix, hence trivially upper triangular. Assume that we have proved the theorem for all matrices of order $n-1$ and let A be of order n. Since the scalar field is $\mathbb{C}$, we know that A has an eigenvector, u_1, say, with corresponding eigenvalue, μ_1 say. Extend u_1 to a basis $U = [u_1, u_2, \ldots, u_n]$ for $\mathbb{C}^n$. Then

$$AU = [Au_1, \cdots, Au_n] = [u_1\mu_1, Au_2, \cdots, Au_n].$$

We want to compute $U^{-1}AU$. For this, observe that $U^{-1}u_1 = U^{-1}Ue_1 = e_1$. Therefore,

$$U^{-1}AU = [e_1\mu_1, U^{-1}Au_2, \cdots, U^{-1}Au_n].$$

Writing this out in detail, we have

$$U^{-1}AU = \widehat{A} := \begin{bmatrix} \mu_1 & \times & \cdots & \times \\ 0 & \times & \cdots & \times \\ \vdots & \vdots & \cdots & \vdots \\ 0 & \times & \cdots & \times \end{bmatrix} =: \begin{bmatrix} \mu_1 & C \\ 0 & A_1 \end{bmatrix}.$$

Here, C is some $1 \times (n-1)$ matrix of no further interest, A_1 is a matrix of order $n-1$, hence, by induction hypothesis, there is some invertible W so that $\widehat{A}_1 := W^{-1}A_1W$ is upper triangular. We compute

$$\mathrm{diag}(1, W^{-1})\widehat{A}\ \mathrm{diag}(1, W) = \begin{bmatrix} 1 & 0 \\ 0 & W^{-1} \end{bmatrix} \begin{bmatrix} \mu_1 & C \\ 0 & A_1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & W \end{bmatrix} = \begin{bmatrix} \mu_1 & CW \\ 0 & W^{-1}A_1W \end{bmatrix}.$$

The computation uses the fact that multiplication from the left (right) by a block-diagonal matrix multiplies the corresponding rows (columns) from the left (right) by the corresponding diagonal blocks. Since $\operatorname{diag}(1, W^{-1}) \operatorname{diag}(1, W) = \operatorname{diag}(1, \mathrm{id}_{n-1}) = \mathrm{id}_n$, this shows that $\widehat{A}$ is similar to an upper triangular matrix. Since A is similar to $\widehat{A}$, this finishes the proof. □

Various refinements in this proof are possible (as we will show later, in the discussion of the Schur form), to give more precise information about possible upper triangular matrices similar to a given A. For the present, though, this is sufficient for our needs since it allows us to prove the following:

(10.27) Corollary: Every complex (square) matrix is similar to an 'almost diagonal' matrix. Precisely, for every complex matrix A and every $\varepsilon > 0$, there exists an upper triangular matrix B_ε similar to A whose off-diagonal entries are all $< \varepsilon$ in absolute value.

Proof: By (10.26)Theorem, we know that any such A is similar to an upper triangular matrix. Since similarity is transitive (see (10.23)Proposition), it is therefore sufficient to prove this Corollary in case A is upper triangular, of order n, say.

The proof in this case is a simple trick: Consider the matrix

$$B := W^{-1}AW,$$

with

$$W := \operatorname{diag}(\delta^1, \delta^2, \ldots, \delta^n),$$

and the *scalar* δ to be set in a moment. W is indeed invertible as long as $\delta \neq 0$, since then

$$W^{-1} = \operatorname{diag}(\delta^{-1}, \delta^{-2}, \ldots, \delta^{-n}).$$

Now, multiplying a matrix by a diagonal matrix from the *left* (*right*) multiplies the *rows* (*columns*) of that matrix by the diagonal entries of the diagonal matrix. Therefore,

$$B_{i,j} = (W^{-1}AW)_{i,j} = A_{i,j}\delta^{j-i}, \quad \text{all } i, j.$$

In particular, B is again upper triangular, and its diagonal entries are those of A. However, all its possibly nontrivial off-diagonal entries lie above the diagonal, i.e., are entries $B_{i,j}$ with $j > i$, hence are the corresponding entries of A multiplied with some *positive* power of the scalar δ. Thus, if

$$c := \max_{i<j} |A_{i,j}|$$

and we choose $\delta := \min\{\varepsilon/c, 1\}$, then, we can be certain that

$$|B_{i,j}| \le \varepsilon, \qquad \text{all } i \neq j,$$

regardless of how small we choose that positive ε. □

10.19 T/F

(a) The only diagonalizable matrix A having just one factorization $A = VMV^{-1}$ with M diagonal is the empty matrix.

(b) If A is the linear map of multiplication by a scalar, then any basis for its domain is an eigenbasis for A.

(c) A triangular matrix of order n is diagonalizable if and only if it has n different diagonal entries.

(d) Any (square) triangular matrix is diagonalizable.

(e) Any matrix of order 1 is diagonalizable.

(f) A matrix of order n has n eigenvalues.

(g) Similar linear maps have the same spectrum.

(h) The linear map of differentiation on Π_k is nilpotent.

(i) The identity map is idempotent.

(j) If the matrix A has 3 eigenvalues, then it must have at least 3 columns.

(k) If $\operatorname{null}(A - \mu\,\mathrm{id})$ is not trivial, then every one of its elements is an eigenvector for A belonging to the eigenvalue μ.

11. Convergence of the power sequence

Convergence of sequences in a normed vector space

Our discussion of the power sequence $A^0, A^1, A^2, \cdots$ of a linear map naturally involves the *convergence* of such a sequence.

Convergence of a vector sequence or a map sequence is most conveniently described with the aid of a norm, as introduced earlier, starting at page 80.

Suppose $z_1, z_2, z_3, \cdots$ is an infinite sequence of n-vectors. In order to avoid confusion, I refer to the jth entry of the kth term z_k in such a vector sequence by $z_k(j)$. We say that this sequence **converges to the n-vector** z_∞ and write

$$z_\infty = \lim_{k\to\infty} z_k,$$

in case

$$\lim_{k\to\infty} \|z_\infty - z_k\| = 0.$$

It is not hard to see that

$$z_\infty = \lim_{k\to\infty} z_k \quad\iff\quad \forall\{i\}\ z_\infty(i) = \lim_{k\to\infty} z_k(i).$$

Note that $z_\infty = \lim_{k\to\infty} z_k$ if and only if, for every $\varepsilon > 0$, there is some k_0 so that, for all $k > k_0$, $\|z_\infty - z_k\| < \varepsilon$. This says that, for any given $\varepsilon > 0$ however small, all the terms in the sequence from a certain point on lie in the "ball"

$$B_\varepsilon(z_\infty) := \{y \in \mathbb{F}^n : \|y - z_\infty\| < \varepsilon\}$$

whose center is z_∞ and whose radius is ε.

(11.1) Lemma: A convergent sequence is necessarily bounded. More explicitly, if the sequence (x_k) of n-vectors converges, then $\sup_k \|x_k\| < \infty$, i.e., there is some c so that, for all k, $\|x_k\| \le c$.

The proof is a verbatim repeat of the proof of this assertion for scalar sequences, as given in the Backgrounder on scalar sequences.

Analogously, we say that the sequence $A_1, A_2, A_3, \cdots$ of matrices **converges** to the matrix B and write

$$\lim_{k\to\infty} A_k = B,$$

in case

$$\lim_{k\to\infty} \|B - A_k\|_\infty = 0.$$

As in the case of vector sequences, a convergent sequence of matrices is necessarily bounded.

Here, for convenience, we have used the map norm associated with the max-norm since we have the simple and explicit formula (7.16) for it. Yet we know from (7.24)Proposition that any two norms on any finite-dimensional normed vector space are equivalent. In particular, if $\|\ \|'$ is any norm on $L(\mathbb{F}^n) = \mathbb{F}^{n\times n}$, then there is a positive constant c so that

$$\|A\|_\infty / c \le \|A\|' \le c\|A\|_\infty, \quad \forall A \in \mathbb{F}^{n\times n}.$$

This implies that $\lim_{k\to\infty} \|B - A_k\|_\infty = 0$ if and only if

$$\lim_{k\to\infty} \|B - A_k\|' = 0,$$

showing that our definition of what it means for A_k to converge to B is independent of the particular matrix norm we use. We might even have chosen the matrix norm

$$\|A\|' := \max_{i,j} |A(i,j)| = \max_{x\neq 0} \frac{\|Ax\|_\infty}{\|x\|_1},$$

and so explicitly confirmed that convergence of matrices is entry-wise, i.e., $\lim_{k\to\infty} A_k = B$ if and only if

$$\lim_{k\to\infty} A_k(i,j) = B(i,j), \quad \forall i,j.$$

Note that, in this chapter, I am using `MATLAB`'s way of writing matrix entries, writing $A_k(i,j)$ instead of $(A_k)_{i,j}$ for the (i,j)-entry of A_k, in order to keep the number of subscripts down.

11.1 For each of the following matrices A, work out A^k for arbitrary $k \in \mathbb{N}$ and, from that, determine directly whether or not the power sequence $A^0, A^1, A^2, \ldots$ converges; if it does, also determine that limit. (i) $A := \alpha\,\mathrm{id}_X$; (ii) $A := \begin{bmatrix} 1/2 & 2^{10} \\ 0 & 1/2 \end{bmatrix}$; (iii) $A := [-e_1, e_2]$; (iv) $A = \begin{bmatrix} a & b \\ 0 & c \end{bmatrix}$.

Three interesting properties of the power sequence of a linear map

We have already most of the tools in hand needed to analyze the following three interesting properties that the **power sequence of** A, i.e., the sequence

$$A^0, A^1, A^2, \cdots \tag{11.2}$$

may have.

Let $A \in L(X)$ with $\dim X < \infty$. Then, for any basis V of X,

$$\widehat{A} := V^{-1}AV$$

is a matrix similar to A, and, for any k,

$$A^k = V\widehat{A}^k V^{-1}.$$

Thus, if we understand the sequence (11.2) for any square *matrix* A, then we understand (11.2) for any $A \in L(X)$ with $\dim X < \infty$.

For this reason, we state here the three interesting properties only for a *matrix* A.

We call the matrix A **power-bounded** in case its power sequence is bounded, i.e., $\sup_k \|A^k\|_\infty < \infty$, i.e., there is a constant c so that, for all k, $\|A^k\|_\infty \le c$.

We call the matrix A **convergent** in case its power sequence converges, i.e., in case, for some matrix B, $B = \lim_{k\to\infty} A^k$.

We call the matrix A **convergent to** 0 in case

$$\lim_{k\to\infty} A^k = 0.$$

See the Backgrounder on the convergence of scalar sequences and, in particular, on the scalar sequence $(\zeta^0, \zeta^1, \zeta^2, \cdots)$.

The first property is fundamental in the study of evolutionary (i.e., time-dependent) processes, such as weather or fluid flow. In the *simplest* approximation, the state of the system (be it the weather or waves on the ocean or whatever) at time t is described by some vector $y(t)$, and the state $y(t+\Delta t)$ at some slightly later time $t + \Delta t$ is computed as

$$y(t+\Delta t) = Ay(t),$$

with A some time-independent matrix. Such a process is called **stable** if $\|y(t)\|$ remains bounded for all time regardless of the initial state, $y(0)$, of the system. Since $y(k\Delta t) = A^k y(0)$, the requirement of stability is equivalent to the power boundedness of A.

The third property is fundamental in the study of iterative processes, as discussed earlier.

The second property is in between the other two. In other words, we have listed the three properties here in the order of increasing strength: if A is convergent to 0, then it is, in particular, convergent. Again, if A is convergent, then it is, in particular, power-bounded.

Suppose now that x is an eigenvector for A, with corresponding eigenvalue μ. Then $Ax = \mu x$, hence $A^k x = \mu^k x$ for $k = 1, 2, 3, \ldots$. Suppose A is powerbounded. Then, in particular, for some c, we should have $c\|x\|_\infty \ge \|A^k\|_\infty \|x\|_\infty \ge \|A^k x\|_\infty = \|\mu^k x\|_\infty = |\mu|^k \|x\|_\infty$. Since $\|x\|_\infty \ne 0$, this implies that the scalar sequence $(|\mu|^k : k = 1, 2, 3, \ldots)$ must be bounded, hence $|\mu| \le 1$. Since we took an arbitrary eigenvector, we conclude that

$$\text{(11.3)} \qquad A \text{ powerbounded} \quad \Longrightarrow \quad \rho(A) \le 1.$$

Actually, more is true. Suppose that μ is a *defective* eigenvalue for A, which, to recall, means that

$$\operatorname{null}(A - \mu\,\mathrm{id}) \cap \operatorname{ran}(A - \mu\,\mathrm{id}) \ne \{0\}.$$

In other words, there exists an eigenvector for A belonging to μ of the form $x = (A - \mu\,\mathrm{id})y$. This implies that

$$Ay = x + \mu y.$$

Therefore

$$A^2 y = Ax + \mu Ay = \mu x + \mu(x + \mu y) = 2\mu x + \mu^2 y.$$

Therefore

$$A^3 y = 2\mu Ax + \mu^2 Ay = 2\mu^2 x + \mu^2(x + \mu y) = 3\mu^2 x + \mu^3 y.$$

By now, the pattern is clear:

$$A^k y = k\mu^{k-1} x + \mu^k y.$$

This also makes clear the difficulty: If $|\mu| = 1$, then

$$\|A^k\|_\infty \|y\|_\infty \ge \|A^k y\|_\infty \ge k\|x\|_\infty - \|y\|_\infty.$$

This shows that A cannot be powerbounded.

We conclude:

(11.4) Proposition: If the matrix A is powerbounded, then, for all $\mu \in \operatorname{spec}(A)$, $|\mu| \le 1$, with equality only if μ is a nondefective eigenvalue for A.

Now we consider the case that A is convergent (hence, in particular, powerbounded). If A is convergent, then, for any eigenvector x with associated eigenvalue μ, the sequence $(\mu^k x : k = 0, 1, 2, \ldots)$ must converge. Since x stays fixed, this implies that the scalar sequence $(\mu^k : k = 0, 1, 2, \ldots)$ must converge. This, to recall, implies that $|\mu| \le 1$ with equality only if $\mu = 1$.

Finally, if A is convergent to 0, then, for any eigenvector x with associated eigenvalue μ, the sequence $(\mu^k x)$ must converge to 0. Since x stays fixed (and is nonzero), this implies that the scalar sequence (μ^k) must converge to 0. This, to recall, implies that $|\mu| < 1$.

Remarkably, these simple necessary conditions just derived, for powerboundedness, convergence, and convergence to 0, are also sufficient; see (11.10)Theorem.

For the proof, we need one more piece of information, namely a better understanding of the distinction between defective and nondefective eigenvalues.

11.2 For each of the following four matrices, determine whether or not it is (a) powerbounded, (b) convergent, (c) convergent to zero. (i) id_n; (ii) $[1, 1; 0, 1]$; (iii) $[8/9, 10^{10}; 0, 8/9]$; (iv) $-\mathrm{id}_n$.

Splitting off the nondefective eigenvalues

Recall that the scalar μ is called a *defective* eigenvalue for $A \in L(X)$ in case

$$\operatorname{null}(A - \mu\,\mathrm{id}) \cap \operatorname{ran}(A - \mu\,\mathrm{id}) \neq \{0\}.$$

(11.5) Proposition: If M is a set of nondefective eigenvalues of $A \in L(X)$, for some finite-dimensional vector space X, then X has a basis $U = [V, W]$, with V consisting entirely of eigenvectors of A belonging to these nondefective eigenvalues, and W any basis for the subspace $Z := \operatorname{ran} p(A)$, with $p(t) := \prod_{\mu \in M}(t - \mu)$.

Further, Z is **A-invariant**, meaning that $A(Z) \subset Z$, hence $A|_Z : Z \to Z : z \mapsto Az$ is a well-defined map on Z, and $\operatorname{spec}(A|_Z) = \operatorname{spec}(A)\backslash M$.

Proof: Since $Ap(A) = p(A)A$, we have $AZ = A(\operatorname{ran} p(A)) = \operatorname{ran} Ap(A) = p(A) \operatorname{ran} A \subset \operatorname{ran} p(A) = Z$, showing Z to be A-invariant. This implies that $A|_Z : Z \to Z : z \mapsto Az$ is a well-defined linear map on Z.

We claim that X is the direct sum of $\operatorname{null} p(A)$ and $\operatorname{ran} p(A)$, i.e.,

(11.6) $$X = \operatorname{null} p(A) \dotplus \operatorname{ran} p(A).$$

Since, by (4.15)Dimension Formula, $\dim X = \dim \operatorname{null} p(A) + \dim \operatorname{ran} p(A)$, it is, by (4.26)Proposition, sufficient to prove that

(11.7) $$\operatorname{null} p(A) \cap \operatorname{ran} p(A) = \{0\}.$$

For its proof, let

$$p_\mu : t \mapsto p(t)/(t - \mu), \qquad \mu \in M,$$

and recall from (5.6) that

$$(p_\mu / p_\mu(\mu) : \mu \in M)$$

is a Lagrange basis for the polynomials of degree $< \#M$. In particular,

$$1 = \sum_{\mu \in M} p_\mu / p_\mu(\mu).$$

Hence, with (10.19)Lemma, $\mathrm{id} = \sum_{\mu \in M} p_\mu(A)/p_\mu(\mu)$ and so, for any $x \in X$,

$$x = \sum_{\mu \in M} x_\mu,$$

with

$$x_\mu := p_\mu(A)x / p_\mu(\mu)$$

in $\operatorname{null}(A - \mu\,\mathrm{id})$ in case $x \in \operatorname{null} p(A)$ (since $(A - \mu\,\mathrm{id})x_\mu = p(A)x/p_\mu(\mu)$), but also in $\operatorname{ran}(A - \mu\,\mathrm{id})$ in case also $x \in \operatorname{ran} p(A) \subset \operatorname{ran}(A - \mu\,\mathrm{id})$, hence then $x_\mu = 0$ since we assumed that each $\mu \in M$ is not defective. This shows (11.7), hence (11.6).

More than that, we just saw that $x \in \operatorname{null} p(A)$ implies that $x = \sum_\mu x_\mu$ with $x_\mu \in \operatorname{null}(A - \mu\,\mathrm{id})$, all $\mu \in M$, hence, $\operatorname{null} p(A) \subset \operatorname{ran} V$, with

$$V := [V_\mu : \mu \in M]$$

and V_μ a basis for $\operatorname{null}(A - \mu\,\mathrm{id})$, all μ. On the other hand, each column of V is in $\operatorname{null} p(A)$, hence also $\operatorname{ran} V \subset \operatorname{null} p(A)$, therefore V is onto $\operatorname{null} p(A)$ and, since it is 1-1 by (10.9)Lemma, it is a basis for $\operatorname{null} p(A)$. Therefore, by (11.6), $U := [V, W]$ is a basis for X for any basis W for $Z = \operatorname{ran} p(A)$.

Finally, let $\nu \in \operatorname{spec}(A)$. If ν were in both M and $\operatorname{spec}(A|_Z)$, then $Ax = \nu x$ for some $x \in Z\backslash 0$, yet also $p(A)x = 0$, hence $0 \neq x \in \operatorname{null} p(A) \cap \operatorname{ran} p(A)$, contradicting (11.7). Thus, if $\nu \in M$, then $\nu \notin \operatorname{spec}(A|_Z)$. If, on the other hand, $\nu \notin M$, then, with x any eigenvector for ν, we have $p(A)x = \alpha x$ with

$$\alpha := \prod_{\mu \in M} (\nu - \mu) \ \neq 0,$$

and so, $x = \alpha^{-1} p(A)x \in \operatorname{ran} p(A) = Z$, hence $\nu \in \operatorname{spec}(A|_Z)$. This proves that $\operatorname{spec}(A|_Z) = \operatorname{spec}(A)\backslash M$. □

It follows that the matrix representation for A with respect to this basis $U = [V, W]$ has the simple form

$$U^{-1}AU = \begin{bmatrix} \mathrm{M} & 0 \\ 0 & \widehat{B} \end{bmatrix} := \operatorname{diag}(\mu_1, \ldots, \mu_r, \widehat{B}),$$

with $\mu_1, \ldots, \mu_r$ a sequence taken from M, and $\widehat{B}$ some square matrix, namely $\widehat{B} = W^{-1}AW$.

(11.8) Theorem: Let $A \in L(X)$, with X a finite-dimensional vector space.
(i) If A is diagonable, then all its eigenvalues are nondefective.
(ii) If $\mathbb{F} = \mathbb{C}$ and all of A's eigenvalues are nondefective, then A is diagonable.

Proof: (i) This is just a restatement of (10.13)Corollary.

(ii) If none of the eigenvalues of A is defective, then we can choose $M = \operatorname{spec}(A)$ in (11.5)Proposition, leaving $A|_Z$ as a linear map with an empty spectrum. Hence, if also $\mathbb{F} = \mathbb{C}$, then we know from (10.15)Theorem that $\operatorname{ran} W = \operatorname{dom} A|_Z$ must be trivial, hence V is a basis for X. □

Here is a simple *example*. Let $A = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$. Then A maps $x := (1, 1)$ to $(3, 3) = 3x$. Hence, $\mu := 3 \in \operatorname{spec}(A)$. We compute

$$\operatorname{ran}(A - \mu\,\mathrm{id}) = \operatorname{ran} \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix} = \operatorname{ran} \begin{bmatrix} -1 \\ 1 \end{bmatrix},$$

since the first column of $(A - \mu\,\mathrm{id})$ is bound and the second is free. This also implies that $\operatorname{null}(A - \mu\,\mathrm{id})$ is one-dimensional, with $V := \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ a basis for it.

It follows, by inspection, that $\operatorname{null}(A - \mu\,\mathrm{id}) \cap \operatorname{ran}(A - \mu\,\mathrm{id}) = \{0\}$ since the only vector of the form $(1, 1)\alpha$ *and* of the form $(-1, 1)\beta$ is the zero vector. Equivalently, the matrix $U := \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$ is 1-1, hence a basis for $\mathbb{R}^2$. Consequently, 3 is a nondefective eigenvalue for A.

Now, what about $A|_Z$, with $Z = \operatorname{ran}(A - \mu\,\mathrm{id})$? In this case, things are very simple since Z is one-dimensional. Since $A(Z) \subset Z$, A must map any $z \in Z$ to a scalar multiple of itself! In particular, since $z = (-1, 1) \in \operatorname{ran}(A - \mu\,\mathrm{id})$, A must map this z into a scalar multiple of itself, and this is readily confirmed by the calculation that A maps z to $-(2, 1) + (1, 2) = z$, i.e., to itself. This shows that z is an eigenvector for A belonging to the eigenvalue 1.

Altogether therefore,

$$AU = [Ax, Az] = [3x, z] = U \operatorname{diag}(3, 1),$$

showing that A is actually diagonable.

This simple example runs rather differently when we change A to $A := \begin{bmatrix} 2 & 1 \\ 0 & 2 \end{bmatrix}$. Since A is upper triangular, its sole eigenvalue is $\mu = 2$. But $(A - \mu\,\mathrm{id}) = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$, and we saw earlier that its range and nullspace have the nontrivial vector e_1 in common. Hence, 2 is a defective eigenvalue for this matrix A.

(11.9) Example: Let $A := [x][y]^{\mathrm{t}}$ with $x, y \in \mathbb{R}^n \backslash 0$. Then $\operatorname{rank} A = 1$, hence $\operatorname{ran} A = \operatorname{ran}[x]$ is one-dimensional, therefore x is an eigenvector for A. Since $Az = x\ (y^{\mathrm{t}}z)$, we have, in particular,

$$Ax = x\ (y^{\mathrm{t}}x),$$

hence x is an eigenvector for A belonging to the eigenvalue $\mu := y^{\mathrm{t}}x$.

Since A is of rank 1, $\dim\operatorname{null} A = n-1$. Let V be a basis for $\operatorname{null} A$, i.e., $V \in L(\mathbb{R}^{n-1}, \operatorname{null} A)$ invertible. Then $U := [V, x]$ is 1-1 (hence a basis for $\mathbb{R}^n$) if and only if $x \notin \operatorname{ran} V$, i.e., if and only if $x \notin \operatorname{null} A$.

case $x \notin \operatorname{ran} V$: Then $U = [V, x]$ is a basis for $\mathbb{R}^n$. Consider the representation $\widehat{A} = U^{-1}AU$ for A with respect to this basis: With $V =: [v_1, v_2, \ldots, v_{n-1}]$, we have $Au_j = Av_j = 0$ for $j = 1{:}n{-}1$, therefore

$$\widehat{A}e_j = 0, \qquad j = 1{:}n{-}1.$$

Further, we have $Ax = x\,(y^{\mathrm{t}}x)$, therefore

$$\widehat{A}e_n = U^{-1}AUe_n = U^{-1}Ax = (y^{\mathrm{t}}x)e_n,$$

(recall that, for any $z \in \mathbb{R}^n$, $U^{-1}z$ provides the coordinates of z with respect to the basis U, i.e., $U(U^{-1}z) = z$). Hence, altogether,

$$\widehat{A} = [0, \ldots, 0, (y^{\mathrm{t}}x)e_n].$$

In particular, A is diagonable, with eigenvalues 0 and $y^{\mathrm{t}}x$.

case $x \in \operatorname{ran} V$: Then $U = [V, x]$ is not a basis for $\mathbb{R}^n$. Worse than that, A is now not diagonable. This is due to the fact that, in this case, the eigenvalue 0 for A is *defective*: For, $x \neq 0$ while $Ax = 0$, hence

$$\{0\} \neq \operatorname{ran}(A - 0\,\mathrm{id}) = \operatorname{ran} A = \operatorname{ran}[x] \subset \operatorname{null} A = \operatorname{null}(A - 0\,\mathrm{id}).$$

Therefore $\operatorname{null}(A - 0\,\mathrm{id}) \cap \operatorname{ran}(A - 0\,\mathrm{id}) \neq \{0\}$. □

It is hard to tell just by looking at a matrix whether or not it is diagonable. There is one exception: If A is hermitian, i.e., equal to its conjugate transpose, then it is not only diagonable, but has an orthonormal basis of eigenvectors, as is shown in the next chapter.

11.3 Prove: *If* $A = \begin{bmatrix} B & C \\ 0 & D \end{bmatrix}$, *with B and D square matrices, then* $\operatorname{spec}(A) = \operatorname{spec}(B) \cup \operatorname{spec}(D)$. (Hint: Prove first that such a matrix A is invertible if and only if both B and D are invertible.)

11.4 Use H.P. 11.3 to determine the spectrum of the matrix $A := \begin{bmatrix} 1 & 2 & 3 & 4 \\ 2 & 4 & 5 & 6 \\ 0 & 0 & 1 & 2 \\ 0 & 0 & 2 & 1 \end{bmatrix}$.

11.5 (a) Use H.P. 11.3 to determine the spectrum of the matrix $A := \begin{bmatrix} 1 & 2 & a \\ 2 & 1 & b \\ 0 & 0 & 3 \end{bmatrix}$. (b) For which choices of a and b is A not diagonable?

Three interesting properties of the power sequence of a linear map: The sequel

(11.10) Theorem: Let $A \in \mathbb{C}^{n\times n}$. Then:

(i) A is powerbounded iff, for all $\mu \in \operatorname{spec}(A)$, $|\mu| \le 1$, with $|\mu| = 1$ only if μ is not defective.

(ii) A is convergent iff, for all $\mu \in \operatorname{spec}(A)$, $|\mu| \le 1$, with $|\mu| = 1$ only if μ is not defective and $\mu = 1$.

(iii) A is convergent to 0 iff $\rho(A) < 1$.

Proof: We only have to prove the implications '$\Longleftarrow$', since we proved all the implications '$\Longrightarrow$' in an earlier section (see pages 114ff).

We begin with (iii). Since A is a matrix over the complex scalars, we know from (10.27)Corollary that, for any $\varepsilon > 0$, we can find an upper triangular matrix B_ε similar to A and with all off-diagonal entries less than ε in absolute value. This means, in particular, that $A = VB_\varepsilon V^{-1}$ for some (invertible) matrix V, hence, for any k, $A^k = V(B_\varepsilon)^k V^{-1}$, therefore,

$$\|A^k\|_\infty \le \|V\|_\infty \|B_\varepsilon\|_\infty^k \|V^{-1}\|_\infty.$$

We compute

$$\|B_\varepsilon\|_\infty = \max_i \sum_j |B_\varepsilon(i,j)| \le \max_i |B_\varepsilon(i,i)| + (n-1)\varepsilon,$$

since each of those sums involves $n-1$ off-diagonal entries and each such entry is less than ε in absolute value. Further, B_ε is upper triangular and similar to A, hence

$$\max_i |B_\varepsilon(i,i)| = \max\{|\mu| : \mu \in \operatorname{spec}(A)\} = \rho(A).$$

By assumption, $\rho(A) < 1$. This makes it possible to choose ε positive yet so small that $\rho(A) + (n-1)\varepsilon < 1$. With this choice, $\|B_\varepsilon\|_\infty < 1$, hence $\lim_{k\to\infty} \|B_\varepsilon\|_\infty^k = 0$. Therefore, since $\|V\|_\infty$ and $\|V^{-1}\|_\infty$ stay fixed throughout, also $\|A^k\|_\infty \to 0$ as $k \to \infty$. In other words, A is convergent to 0.

With this, we are ready also to handle (i) and (ii). Both assume that all eigenvalues of A of modulus 1 are nondefective. By (11.5)Proposition, this implies the existence of a basis $U = [V, W]$ for $\mathbb{C}^n$ so that V consists of eigenvectors of A belonging to eigenvalues of modulus 1, while $Z := \operatorname{ran} W$ is A-invariant and $A|_Z$ has only eigenvalues of modulus < 1. In particular, $AV = V\mathrm{M}$ for some diagonal matrix M with all diagonal entries of modulus 1, and $AW = WB$ for some matrix B with $\operatorname{spec}(B) = \operatorname{spec}(A|_Z)$, hence $\rho(B) < 1$. Consequently, for any k,

$$A^k U = A^k[V, W] = [A^k V, A^k W] = [V\mathrm{M}^k, WB^k] = U \operatorname{diag}(\mathrm{M}^k, B^k).$$

In other words,

$$A^k = U \operatorname{diag}(\mathrm{M}^k, B^k) U^{-1}.$$

Therefore, $\|A^k\|_\infty \le \|U\|_\infty \max\{\|\mathrm{M}\|_\infty^k, \|B^k\|_\infty\} \|U^{-1}\|_\infty$, and this last expression converges since $\|\mathrm{M}\|_\infty = 1$ while $\|B^k\|_\infty \to 0$, by (iii). Since any convergent sequence is bounded, this implies that also the sequence $(\|A^k\|_\infty)$ must be bounded, hence we have finished the proof of (i).

Assume now, in addition, as in (ii) that all eigenvalues of A of modulus 1 are actually equal to 1. Then $\mathrm{M} = \mathrm{id}$, and so, $\lim_{k\to\infty} A^k = C := U \operatorname{diag}(\mathrm{M}, 0) U^{-1}$ since $A^k - C = U \operatorname{diag}(0, B^k) U^{-1}$, hence

$$\|A^k - C\|_\infty \le \|U\|_\infty \|B^k\|_\infty \|U^{-1}\|_\infty \le \mathrm{const}\|B^k\|_\infty \to 0$$

as $k \to \infty$. □

(11.11) Example: Here is a concrete example, chosen for its simplicity.

Let $A = \begin{bmatrix} 1 & 1 \\ 0 & \alpha \end{bmatrix}$. Then $\operatorname{spec}(A) = \{1, \alpha\}$. In particular, A is diagonable if $\alpha \ne 1$ (by (10.10)Corollary) since then A has two eigenvalues. On the other hand, if $\alpha = 1$, then A is not diagonable since it then looks like $\mathrm{id}_2 + N$, with $N := [0, e_1]$ the simplest example of a non-diagonable matrix. Also, in the latter case, the sole eigenvalue, 1, is certainly defective since e_1 is both in $\operatorname{null}(A - \mathrm{id})$ and in $\operatorname{ran}(A - \mathrm{id})$.

Also,

$$A^k = \begin{bmatrix} 1 & 1 + \alpha + \cdots + \alpha^{k-1} \\ 0 & \alpha^k \end{bmatrix} = \begin{cases} \begin{bmatrix} 1 & \frac{1-\alpha^k}{1-\alpha} \\ 0 & \alpha^k \end{bmatrix} & \text{if } \alpha \ne 1; \\ \begin{bmatrix} 1 & k \\ 0 & 1 \end{bmatrix} & \text{otherwise.} \end{cases}$$

We see that A is powerbounded whenever $|\alpha| \le 1$ *except* when $\alpha = 1$, i.e., except when there is a defective absolutely largest eigenvalue.

Further, A is convergent iff $|\alpha| < 1$, i.e., if, in addition, the sole eigenvalue of size 1 is equal to 1 and is nondefective. □

The power method

The simple background for the success of the **power method** is the following corollary to (11.10)Theorem (ii).

(11.12) Proposition: If A has just one eigenvalue μ of absolute value $\rho(A)$ and μ is nondefective, then, for almost any x and almost any y, the sequence

$$A^k x/(y^c A^k x), \qquad k = 1, 2, \ldots$$

converges to an eigenvector of A belonging to that absolutely maximal eigenvalue μ. In particular, for almost any vector y, the ratio

$$y^c A^{k+1} x / y^c A^k x$$

converges to μ.

Proof: By assumption, there is (by (11.5)Proposition) a basis $U := [V, W]$, with V a basis for the space $\mathrm{null}(A - \mu\,\mathrm{id})$ of all eigenvectors of A belonging to that absolutely largest eigenvalue μ of A, and $B := A|_{\mathrm{ran}\,W}$ having all its eigenvalues $< |\mu|$ in absolute value. This implies that $\rho(B/\mu) < 1$. Therefore, for any $x =: [V, W](a, b)$,

$$A^k x = \mu^k V a + B^k W b = \mu^k \left(V a + (B/\mu)^k W b \right)$$

and $(B/\mu)^k W b \to 0$ as $k \to \infty$. Consequently, for any y,

$$\frac{y^c A^{k+1} x}{y^c A^k x} = \frac{\mu^{k+1}(y^c V a + y^c (B/\mu)^{k+1} W b)}{\mu^k (y^c V a + y^c (B/\mu)^k W b)} = \mu \frac{y^c V a + y^c (B/\mu)^{k+1} W b}{y^c V a + y^c (B/\mu)^k W b} \to \mu$$

provided $y^c V a \neq 0$. □

Note that the speed with which $y^c A^{k+1} x / y^c A^k x$ converges to μ depends on the speed with which $(B/\mu)^k W b \to 0$ as $k \to \infty$, hence, ultimately, on $\rho(B/\mu)$.

In the **scaled power method**, one would, instead, consider the sequence

$$x_{k+1} := A(x_k/\|x_k\|), \quad k = 0, 1, \ldots,$$

or, more simply, the sequence

$$x_{k+1} := A(x_k/y^t x_k), \quad k = 0, 1, \ldots.$$

The power method is at the heart of good algorithms for the calculation of eigenvalues. In particular, the standard algorithm, i.e., the QR method, can be interpreted as a (very sophisticated) variant of the power method.

11.6 Using `MATLAB` if really necessary, try out the Power method on the following matrices A, starting at the specified vector x, and discuss success or failure. (Note: You can always use `eig(A)` to find out what the absolutely largest eigenvalue of A is (as well as some eigenvector for it), hence can tell whether or not the power method is working for you. If it isn't, identify the source of failure.) (a) $A = \begin{bmatrix} 0 & .2 & .2 & .3 \\ .2 & 0 & .2 & .3 \\ .5 & .4 & 0 & .4 \\ .3 & .4 & .6 & 0 \end{bmatrix}$, $x = (1, 1, 1, 1)$; (b) $A = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$, $x = (1, -1)$; (c) $A = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$, $x = e_1$; (d) $A = \begin{bmatrix} 4 & 1 & -1 \\ 2 & 5 & -2 \\ 1 & 1 & 2 \end{bmatrix}$, $x = (1, -2, -1)$.

11.7 T/F

(a) If the matrix A of order n has n eigenvalues, then none of its eigenvalues is defective.

(b) If, for some sequence $(x_n : n \in \mathbb{N})$ of m-vectors, $\lim_{n\to\infty} \|x_n\|_2 = 0$, then $\lim_{n\to\infty} \|x_n\| = 0$ for any norm $\|\cdot\|$ on $\mathbb{F}^m$.

(c) If all the eigenvalues of A are < 1, then $\lim_{k\to\infty} A^k \to 0$.

(d) If all the eigenvalues of A are ≤ 1 in absolute value, then A is power-bounded.

(e) If $p(A)x = 0$ for some polynomial p, $A \in L(X)$ and $x \in X\backslash\{0\}$, then every eigenvalue of A is a zero of p.

12. Canonical forms

Canonical forms exhibit essential aspects of a linear map. Of the three discussed in this chapter, only the Schur form has practical significance. But the mathematics leading up to the other two is too beautiful to be left out.

The only result from this chapter used later in these notes is the spectral theorem for hermitian matrices; see (12.2/) Corollary.

The Schur form

The discussion of the powers A^k of A used crucially the fact that any square matrix is similar to an upper triangular matrix. The argument we gave there for this fact is due to I. Schur, who used a refinement of it to show that the basis V for which $V^{-1}AV$ is upper triangular can even be chosen to be *unitary* or *orthonormal*, i.e., so that

$$V^{c}V = \text{ id}.$$

(12.1) Schur's theorem: Every $A \in L(\mathbb{C}^n)$ is **unitarily similar** to an upper triangular matrix, i.e., there exists a unitary basis U for $\mathbb{C}^n$ so that $\widehat{A} := U^{-1}AU = U^{c}AU$ is upper triangular.

Proof: Simply repeat the proof of (10.26)Theorem, with the following modifications: Normalize the eigenvector u_1, i.e., make it have (Euclidean) length 1, then extend it to an o.n. basis for $\mathbb{C}^n$ (as can always be done by applying Gram-Schmidt to an arbitrary basis $[u_1, \ldots]$ for $\mathbb{C}^n$). Also, observe that unitary similarity is also an equivalence relation since the product of unitary matrices is again unitary. Finally, if W is unitary, then so is $\text{diag}(1, W)$. □

Here is one of the many consequences of Schur's theorem. It concerns **hermitian** matrices, i.e., matrices A for which $A^{c} = A$. By Schur's theorem, such a matrix, like any other matrix, is unitarily similar to an upper triangular matrix, i.e., for some unitary matrix U , $\widehat{A} := U^{c}AU$ is upper triangular. On the other hand, for any matrix A and any unitary matrix U,

$$(U^{c}AU)^{c} = U^{c}(A^{c})U.$$

In other words: if $\widehat{A}$ is the matrix representation for A with respect to a *unitary* basis, then $\widehat{A}^{c}$ is the matrix representation for A^{c} with respect to the very same basis. For our hermitian matrix A with its upper triangular matrix representation $\widehat{A} = U^{c}AU$ with respect to the unitary basis U, this means that also $\widehat{A}^{c} = \widehat{A}$, i.e., that the *upper* triangular matrix $\widehat{A}$ is also *lower triangular* and that its diagonal entries are all real. This proves the hard part of the following remarkable

(12.2) Corollary: A matrix $A \in \mathbb{C}^n$ is hermitian if and only it is unitarily similar to a real diagonal matrix.

Proof: We still have to prove that if $\widehat{A} := U^{c}AU$ is real and diagonal for some unitary U, then A is necessarily hermitian. But that follows at once from the fact that then $\widehat{A}^{c} = \widehat{A}$, therefore $A^{c} = (U\widehat{A}U^{c})^{c} = U\widehat{A}^{c}U^{c} = U\widehat{A}U^{c} = A$. □

A slightly more involved argument makes it possible to characterize all those matrices that are unitarily similar to a diagonal matrix (real or not). Such a matrix has enough eigenvectors to make up an entire orthonormal basis from them. Here are the details.

Start with the observation that diagonal matrices commute with one another. Thus, if $\widehat{A} := U^{c}AU$ is diagonal, then

$$A^{c}A = (U\widehat{A}^{c}U^{c})(U\widehat{A}U^{c}) = U\widehat{A}^{c}\widehat{A}U^{c} = U\widehat{A}\widehat{A}^{c}U^{c} = (U\widehat{A}U^{c})(U\widehat{A}^{c}U^{c}) = AA^{c},$$

hence having $A^cA = AA^c$ is a necessary condition for A to be unitarily similar to a diagonal matrix. Remarkably, this condition is sufficient as well. Note that this condition can be directly tested by computing the two products and comparing them. It constitutes the only criterion for the diagona(liza)bility of a matrix available that can be tested for by finitely many calculations. Not surprisingly, matrices with this property are very convenient and have, correspondingly, been given a very positive label. They are called **normal**. (Another label might have been **boring**.)

One way to prove that normal matrices are unitarily similar to a diagonal matrix is by way of a refinement of Schur's theorem: It is possible to find a unitary basis that simultaneously upper-triangularizes two matrices A and B provided A and B **commute**, i.e., provided $AB = BA$. This is due to the fact that commuting matrices have some eigenvector in common.

Assuming this refinement of Schur's theorem (cf. (12.5)Theorem below), one would obtain, for a given normal matrix A, a unitary basis U so that both U^cAU and U^cA^cU are upper triangular. Since one of these is the conjugate transpose of the other, they must both be diagonal. This finishes the proof of

(12.3) Theorem: A matrix $A \in \mathbb{C}^n$ is unitarily similar to a diagonal matrix if and only if $AA^c = A^cA$.

Now for the proof of the refined Schur's theorem. Since the proof of Schur's theorem rests on eigenvectors, it is not surprising that a proof of its refinement rests on the following

(12.4) Lemma: If $A, B \in \mathbb{C}^n$ commute, then there exists a vector that is eigenvector for both of them.

Proof: Let x be an eigenvector for A, $Ax = x\mu$ say, and let $p = p_{B,x}$ be the minimal annihilating polynomial for B at x. Since $x \neq 0$, p has zeros. Let ν be one such and set $p =: (\cdot - \nu)q$. Since $\mathbb{F} = \mathbb{C}$, we know that $v := q(B)x$ is an eigenvector for B (for the eigenvalue ν). But then, since $AB = BA$, we also have $Aq(B) = q(B)A$, therefore

$$Av = Aq(B)x = q(B)Ax = q(B)x\mu = v\mu,$$

showing that our eigenvector v for B is also an eigenvector for A. □

(12.5) Schur's refined theorem: For every $A, B \in L(\mathbb{C}^n)$ that commute, there exists a unitary basis U for $\mathbb{C}^n$ so that both U^cAU and U^cBU are upper triangular.

Proof: This is a further refinement of the proof of (10.26)Theorem. The essential step in that proof was to come up with some eigenvector for A which was then extended to a basis, well, to an o.n. basis U for the proof of Schur's Theorem. Therefore, to have U simultaneously upper-triangularize both A and B, all that's needed is (i) to observe that, by (12.4)Lemma, we may choose u_1 to be a (normalized) eigenvector of A *and* B since, by assumption, $AB = BA$; and (ii) verify that the submatrices A_1 and B_1 obtained in the first step again commute (making it possible to apply the induction hypothesis to them). Here is the verification of this latter fact:

Assuming the eigenvalue of B corresponding to the eigenvector u_1 to be ν, we have

$$U^cAU = \begin{bmatrix} \mu & C \\ 0 & A_1 \end{bmatrix} \qquad U^cBU = \begin{bmatrix} \nu & D \\ 0 & B_1 \end{bmatrix}.$$

Therefore

$$\begin{bmatrix} \mu\nu & \mu D + CB_1 \\ 0 & A_1B_1 \end{bmatrix} = \begin{bmatrix} \mu & C \\ 0 & A_1 \end{bmatrix} \begin{bmatrix} \nu & D \\ 0 & B_1 \end{bmatrix}$$
$$= U^c AUU^c BU = U^c ABU = U^c BAU = U^c BU\, U^c AU = \begin{bmatrix} \nu\mu & \nu C + DA_1 \\ 0 & B_1A_1 \end{bmatrix},$$

hence also A_1 and B_1 commute. □

The primary decomposition

The following analysis goes back to Frobenius and could be viewed as a first step toward a finest A-invariant direct sum decomposition, aka the Jordan form, though the Jordan form is deduced in the next section without any reference to this section. We give the analysis here in the more general situation when the scalar field $\mathbb{F}$ may not be algebraically closed.

The 'primary decomposition' refers to the following facts (taken for granted here). The ring Π of (univariate) polynomials over the field $\mathbb{F}$ is a **unique factorization domain**. This means that each monic polynomial can be written in exactly one way (up to order of the factors) as a product of **irreducible** polynomials, i.e., monic polynomials that have no proper factors. Here, p is called a **proper factor of** q if (i) $0 < \deg p < \deg q$, and (ii) $q = hp$ for some polynomial h.

If $\mathbb{F} = \mathbb{C}$ (or any other algebraically closed field), then each such irreducible polynomial is a monic linear polynomial, i.e., of the form $(\cdot - \mu)$ for some scalar μ. Otherwise, irreducible polynomials may well be of higher than first degree. In particular, if $\mathbb{F} = \mathbb{R}$, then an irreducible polynomial may be of second degree, like the polynomial $()^2 + 1$, but no irreducible polynomial would be of higher degree than that.

The irreducible polynomials are the 'primes' in the 'ring' Π, hence the above-mentioned unique factorization is one into powers of 'primes', or a **prime factorization**.

To obtain the 'primary decomposition' of the linear space X with respect to the linear map $A \in L(X)$, it is convenient to start with the set

$$\mathcal{N}_A := \{p \in \Pi : \operatorname{null} p(A) \neq \{0\}\}$$

of all polynomials p for which $p(A)$ fails to be invertible. This set is not trivial, meaning that it contains more than just the zero polynomial, if, as we continue to assume, $\dim X < \infty$, since then

$$p_{A,x} \in \mathcal{N}_A, \quad \forall x \in X, \tag{12.6}$$

with $p_{A,x}$ the *minimal polynomial for A at x*, which, to recall, is the monic polynomial p of smallest degree for which $p(A)x = 0$.

Call an element of $\mathcal{N}_A$ **minimal** if it is monic and none of its proper factors is in $\mathcal{N}_A$, and let

$$\mathcal{Q}_A$$

be the collection of all minimal elements of $\mathcal{N}_A$.

The set $\mathcal{Q}_A$ is not empty since $\mathcal{N}_A$ is not empty and is closed under multiplication by a scalar, hence contains a monic polynomial of smallest degree. Any $q \in \mathcal{Q}_A$ is necessarily **irreducible**, since, otherwise, it would be the product of certain polynomials p with $p(A)$ 1-1, hence also $q(A)$ would be 1-1.

For every $q \in \mathcal{Q}_A$ and every $x \in \operatorname{null} q(A)\backslash 0$, necessarily $p_{A,x} = q$, by the minimality of $p_{A,x}$. This implies that

$$p, q \in \mathcal{Q}_A \ \text{ and } \ \operatorname{null} p(A) \cap \operatorname{null} q(A) \neq \{0\} \implies p = q. \tag{12.7}$$

(12.8) Lemma: Let p be a product of elements of $\mathcal{Q}_A$,

$$p =: \prod_{q\in\mathcal{Q}'_A} q(A)^{d_q}$$

say, with $d_q \in \mathbb{N}$ and $\mathcal{Q}'_A$ a finite subset of $\mathcal{Q}_A$. Then,

$$X_p := \operatorname{null} p(A) = \dot{+}_{q\in\mathcal{Q}'_A} \operatorname{null} q(A)^{d_q}, \tag{12.9}$$

i.e., $X_p = \operatorname{null} p(A)$ is the direct sum of the spaces $Y_q := \operatorname{null} q(A)^{d_q}$. In other words (by (4.26)Proposition), with V_q a basis for Y_q,

$$V_p := [V_q : q \in \mathcal{Q}'_A]$$

is a basis for X_p.

Proof: There is nothing to prove if $\mathcal{Q}'_A$ has just one element. So, assume that $\#\mathcal{Q}'_A > 1$, and consider the set

$$\mathcal{I} := \sum_{q\in\mathcal{Q}'_A} (p/q^{d_q})\Pi := \{\sum_{q\in\mathcal{Q}'_A} (p/q^{d_q})p_q : p_q \in \Pi\}$$

of all polynomials writable as a weighted sum of the polynomials

$$p/q^{d_q} = \prod_{g\in\mathcal{Q}'_A\backslash q} g^{d_g}$$

for $q \in \mathcal{Q}'_A$, with polynomial (rather than just scalar) weights. This set is a polynomial **ideal**, meaning that it is closed under addition, as well as under multiplication by polynomials. More than that, let q^* be the monic polynomial of smallest degree in $\mathcal{I}$. By Euclid's algorithm, for every $q \in \mathcal{I}$, there exist polynomials g and h with $q = hq^* + g$, hence $g = q - hq^* \in \mathcal{I}$, yet $\deg g < \deg q^*$, hence, by the minimality of q^*, $g = 0$. In other words, the monic polynomial q^* is a factor of every $q \in \mathcal{I}$, in particular of every p/q^{d_q} with $q \in \mathcal{Q}'_A$. But these polynomials have no proper factor in common. Therefore, q^* is necessarily the monic polynomial of degree 0, i.e., $q^* = ()^0$.

It follows that

$$()^0 = \sum_{q\in\mathcal{Q}'_A} (p/q^{d_q})h_q$$

for certain polynomials h_q. This implies that, for the corresponding linear maps

$$P_q : X_p \to X_p : y \mapsto (p/q^{d_q})(A)h_q(A)y, \quad q \in \mathcal{Q}'_A,$$

we have

$$\operatorname{id}_{X_p} = \sum_q P_q. \tag{12.10}$$

Also, for $q \ne g$, $P_qP_g = s(A)p(A) = 0$ for some $s \in \Pi$. Therefore also

$$P_q = P_q \operatorname{id}_{X_p} = P_q(\sum_g P_g) = \sum_g P_qP_g = P_qP_q.$$

This shows that each P_q is a linear projector, and, by (5.11), that X_p is the direct sum of the ranges of the P_q. It remains to show that

$$\operatorname{ran} P_q = Y_q = \operatorname{null} q(A)^{d_q}. \tag{12.11}$$

It is immediate that $\operatorname{ran} P_q \subset Y_q \subset X_p$. With that, $Y_q \subset \operatorname{null} P_g$ for all $g \in \mathcal{Q}'_A\backslash q$, and this implies (12.11), by (12.10). □

Now let $p = p_A$ be the **minimal (annihilating) polynomial for** A, meaning the monic polynomial p of smallest degree for which $p(A) = 0$.

To be sure, there is such a polynomial since X is finite-dimensional, hence so is $L(X)$ (by (4.24)Corollary), therefore $[A^r : r = 0{:}\dim L(X)]$ must fail to be 1-1, i.e., there must be some a for which

$$p(A) := \sum_{j \le \dim L(X)} a_j A^j = 0,$$

yet $a_j \ne 0$ for some $j > 0$, hence the set of all annihilating polynomials of positive degree is not empty, therefore must have an element of minimal degree, and it will remain annihilating and of that degree if we divide it by its leading coefficient.

By the minimality of p_A, every proper factor of p_A is necessarily in $\mathcal{N}_A$. Hence p_A is of the form

$$p_A = \prod_{q \in \mathcal{Q}'_A} q^{d_q}$$

for some $\mathcal{Q}'_A \subset \mathcal{Q}_A$. (In fact, it is immediate from (12.8)Lemma that necessarily $\mathcal{Q}'_A = \mathcal{Q}_A$, but we don't need that here.) This gives, with (12.8)Lemma, the primary decomposition for X wrto A:

$$X = \dot{+}_q \operatorname{null} q(A)^{d_q}. \tag{12.12}$$

Necessarily,

$$\operatorname{null} q(A)^{d_q} = \cup_r \operatorname{null} q(A)^r,$$

with d_q the smallest natural number for which this equality holds. Indeed, from (12.12), every $x \in X$ is uniquely writable as $x = \sum_g x_g$ with $x_g \in \operatorname{null} g(A)^{d_g}$, all $g \in \mathcal{Q}'_A$, and, since each $\operatorname{null} g(A)^{d_g}$ is A-invariant, we therefore have $q(A)^r x = \sum_g q(A)^r x_g = 0$ if and only if $q(A)^r x_g = 0$ for all $g \in \mathcal{Q}'_A$. However, as we saw before, for each $g \in \mathcal{Q}_A \backslash q$, $q(A)$ is 1-1 on $\operatorname{null} g(A)^{d_g}$, hence $q(A)^r x_g = 0$ if and only if $x_g = 0$. Therefore, altogether, $\operatorname{null} q(A)^{d_q} \supset \operatorname{null} q(A)^r$ for any r. This proves that

$$\operatorname{null} q(A)^{d_q} \supset \cup_r \operatorname{null} q(A)^r,$$

while the converse inclusion is trivial. If now $\operatorname{null} q(A)^r = \operatorname{null} q(A)^{d_q}$ for some $r < d_q$, then already $p := p_A/q^{d_q - r}$ would annihilate X, contradicting p_A's minimality.

If $\mathbb{F} = \mathbb{C}$, then each q is of the form $\cdot - \mu_q$ for some scalar μ_q and, correspondingly,

$$X = \dot{+}_q \operatorname{null}(A - \mu_q \operatorname{id})^{d_q}.$$

In particular, $A - \mu_q \operatorname{id}$ is nilpotent on

$$Y_q := \operatorname{null}(A - \mu_q \operatorname{id})^{d_q},$$

with degree of nilpotency equal to d_q. Since

$$A = \mu_q \operatorname{id} + (A - \mu_q \operatorname{id}),$$

it follows that

$$\exp(tA) = \exp(t\mu_q \operatorname{id}) \exp(t(A - \mu_q \operatorname{id})) = \exp(t\mu_q) \sum_{r<d_q} t^r (A - \mu_q \operatorname{id})^r / r! \quad \text{on } Y_q, \tag{12.13}$$

thus providing a very helpful detailed description of the solution $y : t \mapsto \exp(tA)c$ to the first-order ODE $y'(t) = Ay(t)$, $y(0) = c$, introduced in (10.4).

12.1 A subset F of the vector space $X := C^{(1)}(\mathbb{R})$ of continuously differentiable functions is called D-**invariant** if the derivative Df of any $f \in F$ is again in F.

Prove: Any finite-dimensional D-invariant linear subspace Y of $C^{(1)}(\mathbb{R})$ is necessarily the nullspace of a constant-coefficient ordinary differential operator.

The Jordan form

The Jordan form is the result of the search for the 'simplest' matrix representation for $A \in \mathbb{F}^{n\times n}$. It starts off from the following observation.

Suppose $X := \mathbb{F}^n$ is the direct sum

$$X = Y_1 \dotplus Y_2 \dotplus \cdots \dotplus Y_r \tag{12.14}$$

of r linear subspaces, each of which is A-invariant. Then $\operatorname{spec}(A) = \cup_j \operatorname{spec}(A|_{Y_j})$. More than that, with V_j a basis for Y_j, we have $AV_j \subset \operatorname{ran} V_j$, all j. This implies that the coordinate vector of any column of AV_j with respect to the basis $V := [V_1, \ldots, V_r]$ for X has nonzero entries only corresponding to columns of V_j, and these possibly nonzero entries can be found as the corresponding column in the matrix $V_j^{-1}AV_j$. Consequently, the matrix representation $\widehat{A} = V^{-1}AV$ for A with respect to the basis V is block-diagonal, i.e., of the form

$$\widehat{A} = \operatorname{diag}(V_j^{-1}AV_j : j = 1{:}r) = \begin{bmatrix} V_1^{-1}AV_1 & & \\ & \ddots & \\ & & V_r^{-1}AV_r \end{bmatrix}.$$

The smaller we can make the A-invariant summands Y_j, the simpler and more helpful is our overall description $\widehat{A}$ of the linear map A. Of course, the smallest possible A-invariant subspace of X is the trivial subspace, but it would not contribute any columns to V, hence we will assume from now on that our A-invariant direct sum decomposition (12.14) is **proper**, meaning that none of its summands Y_j is trivial.

With that, each Y_j has dimension ≥ 1, hence is as small as possible if it is 1-dimensional, $Y_j = \operatorname{ran}[v_j]$ say, for some nonzero v_j. In this case, A-invariance says that Av_j must be a scalar multiple of v_j, hence v_j is an eigenvector for A, and the sole entry of the matrix $[v_j]^{-1}A[v_j]$ is the corresponding eigenvalue for A.

Thus, at best, each Y_j is 1-dimensional, hence V consists entirely of eigenvectors for A, i.e., A is diagonable. Since we know that not every matrix is diagonable, we know that this best situation cannot always be attained. But we can try to make each Y_j as small as possible, in the following way.

(12.15) Jordan Algorithm:
input: $X = \mathbb{F}^n$, $A \in L(X)$.
$\mathcal{Y} \leftarrow \{X\}$
while $\exists Z_1 \dotplus Z_2 \in \mathcal{Y}$ with both Z_j nontrivial and A-invariant, **do**:
  replace $Z_1 \dotplus Z_2$ **in** $\mathcal{Y}$ **by** Z_1 **and** Z_2**.**
endwhile
output: the proper A-invariant direct sum decomposition $X = \dotplus_{Y\in\mathcal{Y}} Y$.

At all times, the elements of $\mathcal{Y}$ form a proper direct sum decomposition for X. Hence

$$\#\mathcal{Y} \le \sum_{Y\in\mathcal{Y}} \dim Y = \dim X = n.$$

Since each pass through the **while**-loop increases $\#\mathcal{Y}$ by 1, the algorithm must terminate after at most $n-1$ steps.

Now consider any particular Y in the collection $\mathcal{Y}$ output by the algorithm. It is a nontrivial A-invariant subspace. Hence, with the assumption that $\mathbb{F} = \mathbb{C}$, we know that $Y \to Y : y \mapsto Ay$ is a linear map with some eigenvalue, μ say. This implies that the linear map

$$N : Y \to Y : y \mapsto (A - \mu\,\mathrm{id})y$$

is well-defined and has a nontrivial nullspace.

Claim 1: For some $y \in Y$ and some $q \in \mathbb{N}$, $N^{q-1}y \ne 0 = N^q y$.

Proof: Indeed, since $\operatorname{null} N \ne \{0\}$, this holds, e.g., for $q = 1$ and $y \in \operatorname{null} N\backslash 0$. □

Claim 2: For any y and q as in Claim 1, there is $\lambda \in Y'$ with $\lambda N^{q-1}y \neq 0$ and, for any such λ, $Y = \operatorname{null}\Lambda^{\mathrm{t}} \oplus \operatorname{ran} V$, with $\Lambda := [\lambda N^{i-1} : i = 1{:}q]$ and $V := [N^{q-j}y : j = 1{:}q]$.

Proof: The Gramian matrix $\Lambda^{\mathrm{t}}V = (\lambda N^{i-1}N^{q-j}y : i,j = 1{:}q)$ is square and upper triangular, with all diagonal entries equal to $\lambda N^{q-1}y \neq 0$, hence $\Lambda^{\mathrm{t}}V$ is invertible. This implies, by (5.8), that Y is the direct sum of $\operatorname{null}\Lambda^{\mathrm{t}}$ and $\operatorname{ran} V$. □

Claim 3: There is a largest q satisfying Claim 1, and for that q, $\operatorname{null} N^q \cap \operatorname{ran} N^q = \{0\}$, hence $N^q = 0$.

Proof: The V of Claim 2 is 1-1, hence $q = \#V \leq \dim Y$, therefore there is a largest q satisfying Claim 1. For that q, $\operatorname{null} N^q \cap \operatorname{ran} N^q$ is trivial: indeed, if $x \in \operatorname{null}^q \cap \operatorname{ran} N^q$, then $x = N^q u$ for some $u \in Y$, and also $N^{2q}u = N^q x = 0$, but if $N^q u \neq 0$, then $N^{r-1}u \neq 0 = N^r u$ for some $r > q$, which would contradict the maximality of q. Hence $x = N^q u = 0$. With this, $Y = \operatorname{null} N^q \dotplus \operatorname{ran} N^q$ since, by the (4.15)Dimension formula, their dimensions sum to $\dim Y$. Both subspaces are also evidently N-invariant, hence A-invariant, therefore must provide a trivial A-invariant direct sum of Y, i.e., one of the summands must be trivial. Since $y \in \operatorname{null} N^q \backslash 0$, therefore $Y = \operatorname{null} N^q$, hence $N^q = 0$. □

Claim 4: For the largest q, V of Claim 2 is a basis for Y, hence $q = \dim Y$ and the matrix representation for $A|_Y$ with respect to the basis $V_Y := [N^{q-j}y : j = 1{:}q]$ for Y has the simple form

$$V_Y^{-1}(A|_Y)V_Y = \begin{bmatrix} \mu & 1 & 0 & \cdots & 0 & 0 \\ 0 & \mu & 1 & \cdots & 0 & 0 \\ 0 & 0 & \mu & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \cdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & \mu & 1 \\ 0 & 0 & 0 & \cdots & 0 & \mu \end{bmatrix} =: J(\mu, q). \tag{12.16}$$

Proof: Since $NV = V[0, e_1, \ldots, e_{q-1}]$, $\operatorname{ran} V$ is N-invariant for any q, while $\operatorname{null}\Lambda^{\mathrm{t}}$ is N-invariant when q is maximal since then $N^q = 0$ by Claim 3. Hence the direct sum decomposition of Claim 2 is A-invariant and so, one of the summands must be trivial. Since $y \in \operatorname{ran} V\backslash 0$, it follows that $Y = \operatorname{ran} V$ and, since V is 1-1 by Claim 2, it is a basis for Y, and $V^{-1}(A - \mu\,\mathrm{id})|_Y V = V^{-1}NV = [0, e_1, \ldots, e_{q-1}]$. □

It follows that the matrix representation for A with respect to the basis

$$V := [V_Y : Y \in \mathcal{Y}]$$

for $X = \mathbb{F}^n$ is block-diagonal, with each diagonal block a **Jordan block**, $J(\mu, q)$, i.e., of the form (12.16) for some scalar μ and some natural number q. Any such matrix representation for A is called a **Jordan (canonical) form** for A.

There is no reason to believe that such a Jordan form is unique. After all, it depends on the particular order we choose for the elements of $\mathcal{Y}$ when we make up the basis $V = [V_Y : Y \in \mathcal{Y}]$. More than that, there is, in general, nothing unique about $\mathcal{Y}$. For example, if $A = 0$ or, more generally $A = \alpha\,\mathrm{id}$, then any direct sum decomposition for X is A-invariant, hence V can be any basis for X whatsoever for this particular A.

Nevertheless, the Jordan form is canonical in the following sense.

(12.17) Proposition: Let $\widehat{A} =: \operatorname{diag}(J(\mu_Y, \dim Y) : Y \in \mathcal{Y})$ be a Jordan canonical form for $A \in \mathbb{F}^{n\times n}$. Then

(i) $\operatorname{spec}(A) = \{\widehat{A}(j,j) : j = 1{:}n\} = \cup_{Y\in\mathcal{Y}} \operatorname{spec}(A|_Y)$.

(ii) For each $\mu \in \operatorname{spec}(A)$ and each q,

$$n_q := \dim \operatorname{null}(A - \mu\,\mathrm{id})^q = \sum_{\mu_Y=\mu} \min(q, \dim Y), \tag{12.18}$$

hence $\Delta n_q := n_{q+1} - n_q$ equals the number of blocks for μ of order $> q$, giving the decomposition-independent number $-\Delta^2 n_{q-1}$ for the number of Jordan blocks of order q for μ.

In particular, the Jordan form is unique up to an ordering of its blocks.

While the Jordan form is mathematically quite striking, it is of no practical relevance since it does not depend continuously on the entries of A, hence cannot be determined reliably by numerical calculations.

13. Localization of eigenvalues

In this short chapter, we discuss briefly the standard techniques for 'localizing' the spectrum of a given linear map A. Such techniques specify regions in the complex plane that must contain parts or all of the spectrum of A. To give a simple example, we proved (in (12.2)Corollary) that all the eigenvalues of a hermitian matrix must be real, i.e., that $\operatorname{spec}(A) \subset \mathbb{R}$ in case $A^c = A$.

Since $\mu \in \operatorname{spec}(A)$ iff $(A - \mu\,\mathrm{id})$ is not invertible, it is not surprising that many localization theorems derive from a test for invertibility.

Gershgorin's circles

Let μ be an eigenvalue for A with corresponding eigenvector x. Without loss of generality, we may assume that $\|x\| = 1$ in whatever vector norm on $X = \operatorname{dom} A$ we are interested in at the moment. Then

$$|\mu| = |\mu|\|x\| = \|\mu x\| = \|Ax\| \le \|A\|\|x\| = \|A\|,$$

with $\|A\|$ the corresponding map norm for A. This proves that the spectrum of A must lie in the closed disk $B^-_{\|A\|}$ of radius $\|A\|$ centered at the origin. In other words,

$$\rho(A) \le \|A\| \qquad (13.1)$$

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

For example, no eigenvalue of $A = \begin{bmatrix} 1 & 2 \\ -2 & -1 \end{bmatrix}$ can be bigger than 3 in absolute value since $\|A\|_\infty = 3$.

A more refined containment set is obtained by the following more refined analysis.

If $E \in \mathbb{F}^{n\times n}$ has map-norm < 1, then $A := \mathrm{id}_n - E$ is surely 1-1 since then

$$\|Ax\| = \|x - Ex\| \ge \|x\| - \|Ex\| \ge \|x\| - \|E\|\|x\| = \|x\|(1 - \|E\|)$$

with the factor $(1 - \|E\|)$ *positive*, hence $Ax = 0$ implies that $\|x\| = 0$.

Now consider a **diagonally dominant** A, i.e., a matrix A with the property that

$$\forall i \quad |A(i,i)| > \sum_{j\ne i} |A(i,j)|. \qquad (13.2)$$

For example, of the three matrices

$$\begin{bmatrix} 2 & -1 \\ 2 & 3 \end{bmatrix}, \quad \begin{bmatrix} -2 & -1 \\ 3 & 3 \end{bmatrix}, \quad \begin{bmatrix} -2 & -1 \\ 4 & 3 \end{bmatrix}, \qquad (13.3)$$

only the first is diagonally dominant. Setting

$$D := \operatorname{diag} A = \operatorname{diag}(\ldots, A(i,i), \ldots),$$

we notice that (i) D is invertible (since all its diagonal entries are nonzero); and (ii) the matrix E defined by $D^{-1}A =: \mathrm{id} - E$ satisfies

$$E(i,j) = \begin{cases} -A(i,j)/A(i,i) & \text{if } i \ne j; \\ 0 & \text{otherwise,} \end{cases}$$

hence has norm

$$\|E\|_\infty = \max_i \sum_{j\ne i} |A(i,j)/A(i,i)| < 1,$$

by the assumed diagonal dominance of A. This implies that the matrix $\mathrm{id} - E = D^{-1}A$ is invertible, therefore also A is invertible. This proves

(13.4) Proposition: Any diagonally dominant matrix is invertible.

In particular, the first of the three matrices in (13.3) we now know to be invertible. As it turns out, the other two are also invertible; thus, diagonal dominance is only sufficient but not necessary for invertibility. Equivalently, a noninvertible matrix cannot be diagonally dominant.

In particular, for $(A - \mu\,\mathrm{id})$ to be *not* invertible, it must fail to be diagonally dominant, i.e.,

$$\exists i \quad |A(i,i) - \mu| \le \sum_{j \ne i} |A(i,j)|. \tag{13.5}$$

This gives the famous

(13.6) **Gershgorin Circle Theorem**: The spectrum of $A \in \mathbb{C}^{n\times n}$ is contained in the union of the disks

$$B_{r_i}(A(i,i)) := \{z \in \mathbb{C} : |A(i,i) - z| \le r_i := \sum_{j\ne i} |A(i,j)|\}, \quad i = 1{:}n.$$

For the three matrices in (13.3), this says that

$$\mathrm{spec}(\begin{bmatrix} 2 & -1 \\ 2 & 3 \end{bmatrix}) \subset B_1(2) \cup B_2(3) \quad \mathrm{spec}(\begin{bmatrix} -2 & -1 \\ 3 & 3 \end{bmatrix}) \subset B_1(-2) \cup B_3(3) \quad \mathrm{spec}(\begin{bmatrix} -2 & -1 \\ 4 & 3 \end{bmatrix}) \subset B_1(-2) \cup B_4(3).$$

More than that, according to a *refinement of the Gershgorin Circle Theorem*, the second matrix must have one eigenvalue in the closed disk $B_1^-(-2)$ and another one in the closed disk $B_3^-(3)$, since these two disks have an empty intersection. By the same refinement, if the third matrix disk only one eigenvalue, then it would necessarily have to be the point -1, i.e., the sole point common to the two disks $B_1^-(-2)$ and $B_4^-(3)$.

13.1 Does each of the two Gershgorin disks of the matrix $A := \begin{bmatrix} 5 & -1 \\ 6 & 0 \end{bmatrix}$ contain an eigenvalue of A?

The trace of a linear map

Recall that the *trace* of a square matrix A is given by

$$\mathrm{trace}(A) = \sum_j A(j,j).$$

Further, as already observed in (6.31), if the product of the two matrices B and C is square, then

$$\mathrm{trace}(BC) = \sum_j \sum_k B(j,k)C(k,j) = \sum_{jk} B(j,k)C(k,j) = \mathrm{trace}(CB). \tag{13.7}$$

Hence, if $A = V\widehat{A}V^{-1}$, then

$$\mathrm{trace}(A) = \mathrm{trace}(V(\widehat{A}V^{-1})) = \mathrm{trace}(\widehat{A}V^{-1}V) = \mathrm{trace}\,\widehat{A}.$$

This proves

(13.8) Proposition: Any two similar matrices have the same trace.

This permits the definition of the **trace** of an arbitrary linear map A on an arbitrary finite-dimensional vector space X as the trace of the matrices similar to it. In particular, $\operatorname{trace}(A)$ equals the sum of the diagonal entries of any Schur form for A, i.e., $\operatorname{trace}(A)$ is the sum of the eigenvalues of A, however with some of these eigenvalues possibly repeated.

For example, $\operatorname{trace}(\operatorname{id}_n) = n$, while $\operatorname{spec}(\operatorname{id}_n) = \{1\}$.

Offhand, such eigenvalue *multiplicity* seems to depend on the particular Schur form (or any other triangular matrix representation) for A. But, since all of these matrices have the same trace, you will not be surprised to learn that all these triangular matrix representations for A have each eigenvalue appear on its diagonal with exactly the same multiplicity. This multiplicity of $\mu \in \operatorname{spec}(A)$ is denoted

$$\#_a\mu$$

and is called the **algebraic multiplicity** of the eigenvalue, and is readily identified as the dimension of $\cup_r \operatorname{null}(A - \mu \operatorname{id})^r$. Further, the polynomial

$$\chi_A := \prod_{\mu \in \operatorname{spec}(A)} (\cdot - \mu)^{\#_a\mu}$$

is the much-studied **characteristic polynomial** for A.

It would not take much work to validate all these claims directly. But I prefer to obtain them along more traditional lines, namely via determinants.

Determinants

The **determinant** is, by definition, the unique multilinear alternating form

$$\det : [a_1, \ldots, a_n] \to \mathbb{F}$$

for which

$$\det(\operatorname{id}_n) = 1. \tag{13.9}$$

Here, **multilinear** means that det is *linear* in each of its n arguments, i.e.,

$$\det[\ldots, a + \alpha b, \ldots] = \det[\ldots, a, \ldots] + \alpha \det[\ldots, b, \ldots]. \tag{13.10}$$

(Here and below, the various ellipses ... indicate the other arguments, the ones that are kept fixed.) Further, **alternating** means that the interchange of two arguments reverses the sign, i.e.,

$$\det[\ldots, a, \ldots, b, \ldots] = -\det[\ldots, b, \ldots, a, \ldots].$$

In particular, $\det A = 0$ in case two columns of A are the same, i.e.,

$$\det[\ldots, b, \ldots, b, \ldots] = 0.$$

Combining this last with (13.10), we find that

$$\det[\ldots, a + \alpha b, \ldots, b, \ldots] = \det[\ldots, a, \ldots, b, \ldots],$$

i.e., *addition of a scalar multiple of one argument to a different argument does not change the determinant.*

In particular, *if* $A = [a_1, a_2, \ldots, a_n]$ *is not invertible, then* $\det A = 0$ since then there must be some column a_j of A writable as a linear combination of other columns, i.e.,

$$\det A = \det[\ldots, a_j, \ldots] = \det[\ldots, 0, \ldots] = 0,$$

the last equality by the multilinearity of the determinant.

Conversely, *if* A *is invertible, then* $\det A \neq 0$, and this follows from the fundamental determinantal identity

$$\det(AB) = \det(A)\det(B) \tag{13.11}$$

since, for an invertible A,

$$1 = \det(\mathrm{id}_n) = \det(AA^{-1}) = \det(A)\det(A^{-1}),$$

the first equality by (13.9).

(13.12) Theorem: For all $A \in \mathbb{C}^{n\times n}$, $\mathrm{spec}(A) = \{\mu \in \mathbb{C} : \det(A - \mu\,\mathrm{id}) = 0\}$.

Of course, this theorem is quite useless unless we have in hand an explicit formula for the determinant. Here is the standard formula:

$$\det[a_1, a_2, \ldots, a_n] = \sum_{\mathbf{i}\in \mathrm{S}_n} (-)^{\mathbf{i}} \prod_j a_j(\mathbf{i}(j)) \tag{13.13}$$

in which the sum is over all permutations of order n, i.e., all 1-1 (hence invertible) maps $\mathbf{i} : \{1, \ldots, n\} \to \{1, \ldots, n\}$, and the number $(-)^{\mathbf{i}}$ is 1 or -1 depending on the **parity** of the number of interchanges it takes to bring the sequence $\mathbf{i}$ back into increasing order.

For $n = 1$, we get the trivial fact that, for any scalar a, $\mathrm{spec}([a]) = \{a\}$.

For $n = 2$, (13.12) implies that

$$\mathrm{spec}(\begin{bmatrix} a & b \\ c & d \end{bmatrix}) = \{\mu \in \mathbb{C} : (a-\mu)(d-\mu) = bc\}.$$

For $n = 3$, we get

$$\mathrm{spec}(\begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix}) = \{\mu \in \mathbb{C} : p(\mu) = 0\},$$

with

$$p(\mu) := (a-\mu)(e-\mu)(i-\mu) + bfg + chd - c(e-\mu)g - (a-\mu)fh - bd(i-\mu).$$

For $n = 4$, (13.13) already involves 24 summands, and, for general n, we have $n! = 1 \cdot 2 \cdots n$ summands. Thus, even with this formula in hand, the theorem is mostly only of theoretical interest since already for modest n, the number of summands involved becomes too large for any practical computation.

In fact, the determinant $\det A$ of a given matrix A is usually computed with the aid of some factorization of A, relying on the fundamental identity (13.11) and on the following

(13.14) Lemma: The determinant of any triangular matrix is just the product of its diagonal entries.

Proof: This observation follows at once from (13.13) since any permutation $\mathbf{i}$ other than the identity $(1, 2, \ldots, n)$ must have $\mathbf{i}(k) < k$ for some k, hence the corresponding product $\prod_j a_j(\mathbf{i}(j))$ in (13.13) will be zero for any lower triangular matrix. Since any such $\mathbf{i}$ must also have $\mathbf{i}(h) > h$ for some h, the corresponding product will also vanish for any upper triangular matrix. Thus, in either case, only the product $\prod_j a_j(j)$ is possibly nonzero. □

So, with $A = PLU$ as constructed by Gauss-elimination, with L unit lower triangular and U upper triangular, and P a permutation matrix, we have

$$\det A = (-)^P \prod_j U(j,j),$$

with the number $(-)^P$ equal to 1 or -1 depending on the parity of the permutation carried out by P, i.e., whether the number of row interchanges made during Gauss elimination is even or odd.

Formula (13.13) is often taken as the definition of $\det A$. It is a simple consequence of the fundamental identity (13.11), and the latter follows readily from the multilinearity and alternation property of the determinant. For these and other details, see the chapter 'More on determinants'.

Annihilating polynomials

The nontrivial polynomial p is called **annihilating for** $A \in L(X)$ if $p(A) = 0$.

For example, A is *nilpotent* exactly when, for some k, the monomial $()^k$ annihilates A, i.e., $A^k = 0$. As another example, A is a linear projector (or, idempotent) exactly when the polynomial $p : t \mapsto t(t-1)$ annihilates A, i.e., $A^2 = A$.

Annihilating polynomials are of interest because of the following version of the **Spectral Mapping Theorem**:

(13.15) Theorem: For any polynomial p and any linear map $A \in L(X)$ with $\mathbb{F} = \mathbb{C}$,

$$\operatorname{spec}(p(A)) = p(\operatorname{spec}(A)) := \{p(\mu) : \mu \in \operatorname{spec}(A)\}.$$

Proof: If $\mu \in \operatorname{spec}(A)$, then, for some nonzero x, $Ax = \mu x$, therefore also $p(A)x = p(\mu)x$, hence $p(\mu) \in \operatorname{spec}(p(A))$. In other words, $p(\operatorname{spec}(A)) \subset \operatorname{spec}(p(A))$.

Conversely, if $\nu \in \operatorname{spec}(p(A))$, then $p(A) - \nu\,\mathrm{id}$ fails to be 1-1. However, assuming without loss of generality that p is a monic polynomial of degree r, we have $p(t) - \nu = (t-\mu_1)\cdots(t-\mu_r)$ for some scalars $\mu_1, \ldots, \mu_r$, therefore

$$p(A) - \nu\,\mathrm{id} = (A - \mu_1\,\mathrm{id})\cdots(A - \mu_r\,\mathrm{id}),$$

and, since the left-hand side is not 1-1, at least one of the factors on the right must fail to be 1-1. This says that some $\mu_j \in \operatorname{spec}(A)$, while $p(\mu_j) - \nu = 0$. In other words, $\operatorname{spec}(p(A)) \subset p(\operatorname{spec}(A))$. □

In particular, if p annihilates A, then $p(A) = 0$, hence $\operatorname{spec}(p(A)) = \{0\}$, therefore $\operatorname{spec}(A) \subset \{\mu \in \mathbb{C} : p(\mu) = 0\}$.

For example, 0 is the only eigenvalue of a nilpotent linear map. The only possible eigenvalues of an idempotent map are the scalars 0 and 1.

The best-known annihilating polynomial for a given $A \in \mathbb{F}^{n\times n}$ is its *characteristic polynomial*, i.e., the polynomial

$$\chi_A : t \mapsto \det(t\,\mathrm{id}_n - A).$$

To be sure, by (10.26), we can write any such A as the product $A = V\widehat{A}V^{-1}$ with $\widehat{A}$ upper triangular. Correspondingly,

$$\chi_A(t) = \det V \det(t\,\mathrm{id}_n - \widehat{A})(\det V)^{-1} = \det(t\,\mathrm{id}_n - \widehat{A}) = \chi_{\widehat{A}}(t) = \prod_j (t - \widehat{A}(j,j)),$$

the last equation by (13.14)Lemma. Consequently, $\chi_A(A) = V\chi_A(\widehat{A})V^{-1}$, with

$$\chi_A(\widehat{A}) = (\widehat{A} - \mu_1\,\mathrm{id})\cdots(\widehat{A} - \mu_n\,\mathrm{id}), \qquad \mu_j := \widehat{A}(j,j), \quad j = 1{:}n,$$

and this, we claim, is necessarily the zero map, for the following reason: The factor $(\widehat{A} - \mu_j \,\mathrm{id})$ is upper triangular, with the jth diagonal entry equal to zero. This implies that, for each i, $(\widehat{A} - \mu_j \,\mathrm{id})$ maps

$$T_i := \mathrm{ran}[e_1, \ldots, e_i]$$

into itself, but maps T_j into T_{j-1}. Therefore

$$\begin{aligned}\mathrm{ran}\,\chi_A(\widehat{A}) = \chi_A(\widehat{A})T_n &= (\widehat{A} - \mu_1 \,\mathrm{id}) \cdots (\widehat{A} - \mu_n \,\mathrm{id})T_n \\ &\subset (\widehat{A} - \mu_1 \,\mathrm{id}) \cdots (\widehat{A} - \mu_{n-1} \,\mathrm{id})T_{n-1} \\ &\subset (\widehat{A} - \mu_1 \,\mathrm{id}) \cdots (\widehat{A} - \mu_{n-2} \,\mathrm{id})T_{n-2} \\ &\cdots \\ &\subset (\widehat{A} - \mu_1 \,\mathrm{id})T_1 \subset T_0 = \{0\},\end{aligned}$$

or, $\chi_A(\widehat{A}) = 0$, therefore also $\chi_A(A) = 0$. This is known as the **Cayley-Hamilton Theorem**.

Note that the collection $I_A := \{p \in \Pi : p(A) = 0\}$ of all polynomials that annihilate a given linear map A is an **ideal**, meaning that it is a linear subspace that is also closed under multiplication by polynomials: if $p \in I_A$ and $\in \Pi$, then their product $qp : t \mapsto q(t)p(t)$ is also in I_A. Since I_A is not empty, it contains a monic polynomial of minimal degree. This polynomial is called the **minimal polynomial for** A and is denoted by p_A. Using the Euclidean algorithm (see Backgrounder), it is easy to see that p_A must be a factor of every $p \in I_A$; in technical terms, I_A is a **principal** ideal, more precisely the principal ideal generated by p_A.

In exactly the same way, the collection $I_{A,x} := \{p \in \Pi : p(A)x = 0\}$ is seen to be a principal ideal, with $p_{A,x}$ the unique monic polynomial of smallest degree in it. Since $I_A \subset I_{A,x}$, it follows that $p_{A,x}$ must be a factor for any $p \in I_A$ and, in particular, for χ_A.

13.2 (a) Prove: *If the minimal annihilating polynomial $p = p_{A,x}$ of the linear map $A \in L(X)$ at some $x \in X\backslash 0$ has degree equal to* $\dim X$, *then $p_{A,x}(A) = 0$.* (b) Prove that the spectrum of the companion matrix (see H.P. 10.15) of the monic polynomial p equals the zero set of p.

13.3 make one about the coefs of char.pol. being symmetric functions of evs, and one about the ith coeff. being the sum of the $n - i$th principal minors. all of these, including the trace, are invariant under similarity.

The multiplicities of an eigenvalue

Since χ_A is of exact degree n in case $A \in \mathbb{C}^n$, χ_A has exactly n zeros *counting multiplicities*. This means that

$$\chi_A(t) = (t - \mu_1) \cdots (t - \mu_n)$$

for a certain n-sequence μ. Further,

$$\mathrm{spec}(A) = \{\mu_j : j = 1{:}n\},$$

and this set may well contain only one number, as it does when $A = 0$ or $A = \mathrm{id}$. However, it is customary to associate with each eigenvalue, μ, its **algebraic multiplicity**, $\#_a\mu$, which, by definition, is its multiplicity as a zero of the characteristic polynomial, or, equivalently as we saw, as a diagonal entry of any triangular matrix similar to A. For example, the matrix id_n has only the eigenvalue 1, but it has algebraic multiplicity n. In this way, each $A \in \mathbb{C}^{n\times n}$ has n eigenvalues *counting algebraic multiplicity.*

I have been saying '*algebraic* multiplicity' rather than just 'multiplicity', since there is a second way of counting eigenvalues, and that is by geometric multiplicity. The **geometric multiplicity**, $\#_g\mu$, of the eigenvalue μ for A is, by definition, the dimension of the space of corresponding eigenvectors, i.e.,

$$\#_g\mu := \dim \mathrm{null}(A - \mu \,\mathrm{id}).$$

For the sole eigenvalue, 1, of id_n, algebraic and geometric multiplicity coincide. In contrast, the sole eigenvalue, 0, of $\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$ has algebraic multiplicity 2 but its geometric multiplicity is only 1.

An eigenvalue is (algebraically or geometrically) **simple** if it has (algebraic or geometric) multiplicity 1.

13.16 Proposition: For any eigenvalue, the algebraic multiplicity is no smaller than the geometric multiplicity, with equality if and only if the eigenvalue is not defective.

13.4

(i) Prove that *the multiplicity with which an eigenvalue μ of A appears as a diagonal entry of a triangular matrix T similar to A is the same for all such triangular matrices.* (Hint: Prove that it equals the multiplicity of μ as a zero of the characteristic polynomial $\chi_A : t \mapsto \det(t\,\mathrm{id}_n - A)$ of A; feel free to use what we proved about determinants, like: $\det(AB) = \det(A)\det(B)$, and $\det(T) = \prod_j T(j,j)$.)

(ii) Determine the algebraic and geometric multiplicities for all the eigenvalues of the following matrix. (Read off the eigenvalues; use elimination to determine geometric multiplicities.)

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

Perron-Frobenius

We call the matrix A **positive** (**nonnegative**) and write $A > 0$ ($A \ge 0$) in case all its entries are positive (nonnegative). A positive (nonnegative) matrix A of order n maps the **positive orthant**

$$\mathbb{R}^n_+ := \{y \in \mathbb{R}^n : y \ge 0\}$$

into its interior (into itself). Thus the (scaled) power method, started with a nonnegative vector, would converge to a nonnegative vector if it converges. This suggests that the absolutely largest eigenvalue for a nonnegative matrix is nonnegative, with a corresponding nonnegative eigenvector. The Perron-Frobenius theorem makes this intuition precise.

Since A maps $\mathbb{R}^n_+$ into itself, it makes sense to consider, for given $y \in \mathbb{R}^n_+\backslash 0$, scalars α for which $Ay \ge \alpha y$ (in the sense that $(Ay)_j \ge \alpha y_j$, all j), i.e., for which $Ay - \alpha y \ge 0$. The largest such scalar is the nonnegative number

$$r(y) := \min\{(Ay)_j/y_j : y_j > 0\}, \quad y \in \mathbb{R}^n_+\backslash 0.$$

The basic observation is that

$$(13.17) \qquad Ay - \alpha y > 0 \quad \Longrightarrow \quad r(y) > \alpha.$$

The function r so defined is dilation-invariant, i.e., $r(\alpha y) = r(y)$ for all $\alpha > 0$, hence r takes on all its values already on the set $S_+ := \{y \ge 0 : \|y\| = 1\}$. At this point, we need, once again, a result that goes beyond the scope of these notes, namely the fact that S_+ is compact, while r is continuous at any $y > 0$ and upper semicontinuous at any $y \ge 0$, hence r takes on its supremum over $\mathbb{R}^n_+\backslash 0$ at some point in S_+. I.e., there exists $x \in S_+$ for which

$$\mu := r(x) = \sup r(S_+) = \sup r(\mathbb{R}^n_+\backslash 0).$$

Assume now, in addition to $A \ge 0$, that also $p(A) > 0$ for some polynomial p.

Claim: $Ax = \mu x$.

Proof: Assume that $Ax \ne \mu x$. Since $\mu = r(x)$, we have $Ax - \mu x \ge 0$, therefore $A(p(A)x) - \mu p(A)x = p(A)(Ax - \mu x) > 0$, hence, by (13.17), $r(p(A)x) > \mu = \sup r(S_+)$, a contradiction. □

Claim: $x > 0$.

Proof: Since $0 \neq x \geq 0$ and $p(A) > 0$, we have $p(\mu)x = p(A)x > 0$, hence $x > 0$. □

Consequence: *x is the unique maximizer for r.*

Proof: If also $r(y) = \mu$ for some $y \in S_+$, then by the same argument $Ay = \mu y$, therefore $Az = \mu z$ for all $z = x + \alpha(y - x)$, and each of these z must be positive if it is nonnegative, and this is possible only if $y - x = 0$. □

Consequence: *For any eigenvalue ν of any matrix B with eigenvector y, if $|B| \leq A$, then $|\nu| \leq \mu$, with equality only if $|\, y/\|y\|\, | = x$ and $|B| = A$.* (More precisely, equality implies that $B = \exp(\mathrm{i}\varphi)DAD^{-1}$, with $D := \mathrm{diag}(\ldots, y_j/|y_j|, \ldots)$ and $\exp(\mathrm{i}\varphi) := \nu/|\nu|$.)

Proof: Observe that

$$|\nu||y| = |By| \leq |B|\,|y| \leq A|y|, \tag{13.18}$$

hence $|\nu| \leq r(|y|) \leq \mu$. If now there is equality, then, by the uniqueness of the minimizer x (and assuming without loss that $\|y\| = 1$), we must have $|y| = x$ and equality throughout (13.18), and this implies $|B| = A$. More precisely, $D := \mathrm{diag}(\ldots, y_j/|y_j|, \ldots)$ is then well defined and satisfies $y = D|y|$, hence $C|y| = \mu|y| = A|y|$, with $C := \exp(-\mathrm{i}\varphi)D^{-1}BD \leq A$ and $\nu =: \mu\exp(\mathrm{i}\varphi)$, therefore $C = A$. □

Consequences. By choosing $B = A$, we get that $\mu = \rho(A) := \max\{|\nu| : \nu \in \sigma(A)\}$, *and that μ has geometric multiplicity 1 (as an eigenvalue of A).*

We also get that $\rho(A)$ *is strictly monotone in the entries of* A, i.e., that $\rho(\widehat{A}) > \rho(A)$ in case $\widehat{A} \geq A \neq \widehat{A}$ (using the fact that $p(A) > 0$ and $\widehat{A} \geq A$ implies that also $q(\widehat{A}) > 0$ for some polynomial q; see below).

As a consequence, we find *computable* upper and lower bounds for the spectral radius of A:

Claim:

$$\forall\{y > 0\}\;\; r(y) \leq \rho(A) \leq R(y) := \max_j (Ay)_j/y_j,$$

with equality in one or the other if and only if there is equality throughout if and only if $y = \alpha x$ (for some positive α). In particular, $\rho(A)$ is the only eigenvalue of A with positive eigenvector.

Proof: Assume without loss that $\|y\| = 1$. We already know that for any such $y > 0$, $r(y) \leq \rho(A)$ with equality if and only if $y = x$. For the other inequality, observe that $R(y) = \|D^{-1}ADe\|_\infty$ with $D := \mathrm{diag}(\ldots, y_j, \ldots)$ and $e := (1, \ldots, 1)$. Since $D^{-1}AD \geq 0$, it takes on its max-norm at e, hence

$$R(y) = \|D^{-1}AD\|_\infty \geq \rho(D^{-1}AD) = \rho(A).$$

Now assume that $r(y) = R(y)$. Then $Ay = r(y)y$, hence $r(y) \leq r(x) = \rho(A) \leq R(y) = r(y)$, therefore equality must hold throughout and, in particular, $y = x$.

If, on the other hand, $r(y) < R(y)$, then we can find $\widehat{A} \neq A \leq \widehat{A}$ so that $\widehat{A}y = R(y)y$ (indeed, then $z := R(y)y - Ay$ is nonnegative but not 0, hence $\hat{A} := A + y_1^{-1}[z]e_1{}^{\mathrm{t}}$ does the job) therefore $r_{\widehat{A}}(y) = R(y) = R_{\widehat{A}}(y)$, hence $R(y) = \rho(\widehat{A}) > \rho(A)$. □

Claim: *μ has simple algebraic multiplicity.*

Proof: Since we already know that μ has simple geometric multiplicity, it suffices to show that μ is not a defective eigenvalue, i.e., that $\mathrm{null}(A - \mu\,\mathrm{id}) \cap \mathrm{ran}(A - \mu\,\mathrm{id}) = \{0\}$. So assume to the contrary that $Ay - \mu y$ is an eigenvector of A belonging to μ. Then, by the simple geometric multiplicity of μ, we may assume without loss that $Ay - \mu y = x$, or $Ay = \mu y + x$, therefore, for all k, $A^k y = \mu^k y + k\mu^{k-1}x$, hence, finally,

$$(A/\mu)^k y = y + k(x/\mu).$$

Hence, for large enough k, $z := (A/\mu)^k y$ has all its entries positive, and $Az = Ay + kx = \mu y + (k+1)x = \mu(z + x/\mu) > \mu z$, therefore $r(z) > \mu$, a contradiction. □

The collection of these claims/consequences constitutes the **Perron-Frobenius Theorem**. Oskar Perron proved all this under the assumption that $A > 0$ (i.e., $p(t) = t$). Frobenius extended it to all $A \geq 0$ that are **irreducible**. While this term has some algebraic and geometric meaning (see below), its most convenient definition for the present purpose is that $p(A) > 0$ for some polynomial p. In the contrary case, A is called **reducible**, and not(iv) below best motivates such a definition. Here are some equivalent statements:

Claim: *Let $A \geq 0$. Then the following are equivalent:*

(i) *$p(A) > 0$ for some polynomial p.*

(ii) *For all (i, j), there exists $k = k(i, j)$ so that $A^k(i, j) > 0$.*

(iii) *No proper A-invariant subspace is spanned by unit-vectors.*

(iv) *For no permutation matrix P is*

$$PAP^{-1} = \begin{bmatrix} B & C \\ 0 & D \end{bmatrix} \tag{13.19}$$

with B, D square matrices of positive order.

(v) *The directed graph for A is strongly connected.*

Proof: (ii)$\Longrightarrow$(i) since then $p(A) := \sum_{i,j} A^{k(i,j)} > 0$.

If (ii) does not hold, then there exists (i, j) so that $A^k(i, j) = 0$ for all k. But then also $p(A)(i, j) = 0$ for all polynomials p; in other words, (i)$\Longrightarrow$(ii). Further, it says that the set $J := J(j) := \{r : \exists\{k\}\ A^k(r, j) \neq 0\}$ is a proper subset of $\{1, \ldots, n\}$ (since it doesn't contain i), but neither is it empty (since it contains j, as $A^0(j, j) \neq 0$). Since $A^{k+\ell}(r, j) = \sum_m A^k(r, m)A^\ell(m, j)$, it follows that $J(m) \subset J(j)$ for all $m \in J(j)$. This implies, in particular, that $A(r, m) = 0$ for all $r \notin J(j), m \in J(j)$, hence that $\operatorname{span}(e_m)_{m\in J(j)}$ is a proper A-invariant subspace, thus implying not(iii). It also implies not(iv), since it shows that the columns $A(:, m), m \in J(j)$, have zero entries in rows $r, r \notin J(j)$, i.e., that (13.19) holds for the permutation $P = [(e_m)_{m\in J(j)}, (e_r)_{r\notin J(j)}]$, with both B and D of order $< n$.

Conversely, if e.g., (iii) does not hold, and $\operatorname{span}(e_m)_{m\in J(j)}$ is that proper A-invariant subspace, then it is also invariant under any $p(A)$, hence also $p(A)(r, m) = 0$ for every $r \notin J(j)$, $m \in J(j)$, i.e., (i) does not hold.

The final characterization is explicitly that given by Frobenius, – except that he did not formulate it in terms of graphs; that was done much later, by Rosenblatt (1957) and Varga (1962). Frobenius (???) observed that, since

$$A^k(i, j) = \sum_{j_1} \cdots \sum_{j_{k-1}} A(i, j_1) \cdots A(j_{k-1}, j),$$

therefore, for $i \neq j$, $A^k(i, j) \neq 0$ if and only if there exists some sequence $i =: i_0, i_1, \ldots, i_{k-1}, i_k := j$ so that $A(i_r, i_{r+1}) \neq 0$ for all r. Now, the **directed graph** of A is the graph with n vertices in which the directed edge (i, j) is present iff $A(i, j) \neq 0$. Such a graph is called **strongly connected** in case it contains, for each $i \neq j$, a path connecting vertex i with vertex j, and this, as we just observed, is equivalent to having $A^k(i, j) \neq 0$ for some $k > 0$. In short, (ii) and (v) are equivalent. □

There are various refinements of this last claim available. For example, in testing whether the directed graph of A is strongly connected, we only need to check paths involving distinct vertices, and such paths involve at most n vertices. Hence, in condition (ii), we need to check only for $k < n$. But, with that restriction, (ii) is equivalent to having $\mathrm{id}_n + A + \cdots + A^{n-1} > 0$ and, given that $A \geq 0$, this, in turn, is equivalent to having $(\mathrm{id}_n + A)^{n-1} > 0$, i.e., to having (i) hold for quite specific polynomials.

13.5 T/F

() If the sum $A + B$ of two matrices is defined, then $\det(A + B) = \det(A) + \det(B)$.

14. Some applications

3-space

In the vector space $X = \mathbb{R}^3$, the standard inner product is also called the **dot product**, because of the customary notation

$$y^{\mathrm{t}}x = \langle x, y\rangle =: x \cdot y, \qquad x, y \in \mathbb{R}^3.$$

In this most familiar vector space, another vector 'product' is of great use, the so-called **cross product** $x \times y$. It is most efficiently defined implicitly, i.e., by

$$(x \times y) \cdot z := \det[x, y, z], \qquad \forall x, y, z \in \mathbb{R}^3. \tag{14.1}$$

From (13.13) (see also page 164), we work out that

$$\det[x, y, z] = (x_2y_3 - x_3y_2)z_1 + (x_3y_1 - x_1y_3)z_2 + (x_1y_2 - x_2y_1)z_3,$$

hence

$$x \times y = (x_2y_3 - x_3y_2, x_3y_1 - x_1y_3, x_1y_2 - x_2y_1).$$

Given what you already know about determinants, the definition (14.1), though implicit, makes all the basic facts about the cross product immediate:

(i) *The cross product $x \times y$ is linear in its two arguments, x and y.*

(ii) *The cross product $x \times y$ is* **alternating**, *meaning that $y \times x = -(x \times y)$.*

(iii) Perhaps most importantly, *$x \times y$ is a vector perpendicular to both x and y.*

(iv) *$x \times y = 0$ if and only if $[x, y]$ is not 1-1.*

Indeed, if $[x, y]$ is 1-1, then we can always extend it to a basis $[x, y, z]$ for $\mathbb{R}^3$, and then $(x \times y)^{\mathrm{t}}z$ is not zero, hence then $x \times y \neq 0$. If $[x, y]$ fails to be 1-1, then, for any z, $[x, y, z]$ fails to be 1-1, hence then, necessarily, $x \times y = 0$.

So, assuming that $[x, y]$ is 1-1, we can compute the *unit* vector

$$u := (x \times y)/\|x \times y\|,$$

and so conclude that

$$\|x \times y\|_2^2 = \det[x, y, x \times y] = \|x \times y\| \det[x, y, u].$$

In other words,

(v) *the Euclidean length of $x \times y$ gives the (unsigned) area of the parallelepiped spanned by x and y.*

This also holds when $[x, y]$ fails to be 1-1 since then that area is zero.

When $[x, y]$ is 1-1, then there are exactly two unit vectors (or, directions) perpendicular to the plane $\operatorname{ran}[x, y]$ spanned by x and y, namely $u := (x \times y)/\|x \times y\|$ and $(y \times x)/\|y \times x\| = -u$, with u the choice that makes $\det(x, y, u)$ positive. If you imagine the thumb of your *right* hand to be x, and the pointer of that hand to be y, then the middle finger, bent to be perpendicular to both thumb and pointer, would be pointing in the direction of $x \times y$. For that reason, any basis $[x, y, z]$ for $\mathbb{R}^3$ with $\det[x, y, z] > 0$ is said to be **right-handed**.

14.1 Relate the standard choice $(x_2, -x_1)$ for a vector perpendicular to the 2-vector x to the above construction.

14.2 Give a formula for an n-vector $x_1 \times \cdots \times x_{n-1}$ that is perpendicular to the $n-1$ n-vectors $x_1, \ldots, x_{n-1}$ and whose Euclidean length equals the (unsigned) volume of the parallelepiped spanned by the vectors $x_1, \ldots, x_{n-1}$.

Rotation in 3-space

A particularly useful transformation of 3-space is counter-clockwise rotation by some angle θ around some given axis-vector a. Let $R = R_{\theta,a}$ be this rotation. We are looking for a computationally efficient way to represent this map.

This rotation leaves its **axis**, i.e., ran$[a]$, pointwise fixed, and rotates any vector in the plane $H := a^{\perp}$ counterclockwise θ radians. In other words, with

$$p = q + r, \quad q := P_{[a]}p, \quad r := p - q,$$

we have

$$Rp = q + Rr,$$

by the linearity of the rotation. To compute Rr, let s be the vector in H obtained by rotating r counter-clockwise $\pi/2$ radians. Then

$$Rr = \cos(\theta)r + \sin(\theta)s,$$

and that's it.

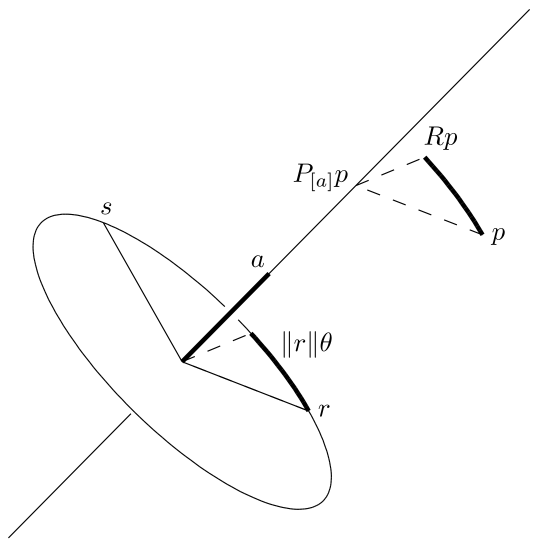

(14.2) Figure. Rotation of the point p counterclockwise θ radians around the axis spanned by the vector a. The orthogonal projection r of p into the plane H with normal a, together with its rotation s counterclockwise $\pi/2$ radians around that axis, serve as a convenient orthogonal coordinate system in H.

It remains to construct s, and this is traditionally done with the aid of the cross product $a \times r$ since (see (14.1)) it is a vector perpendicular to a and r. Hence, assuming without loss that a is normalized, we now know that $a \times r$ is in the plane H and perpendicular to r and of the same length as r. Of the two vectors in H that have this property, it also happens to be the one obtained from r by a $(\pi/2)$-rotation that appears counterclockwise when looking down on H from the side that the vector a points into. (Just try it out.)

The calculations can be further simplified. The map

$$r \mapsto a \times r$$

is linear and, by inspection, $a \times a = 0$. Since a is normalized by assumption, we compute

$$r = p - (a^{\mathrm{t}}p)a,$$

hence

$$a \times r = a \times p.$$

So, altogether

$$Rp = (a^{\mathrm{t}}p)a + \cos(\theta)(p - (a^{\mathrm{t}}p)a) + \sin(\theta)(a \times p) = \cos(\theta)p + (1 - \cos(\theta))(a^{\mathrm{t}}p)a + \sin(\theta)(a \times p).$$

This is the formula that is most efficient for the calculation of Rp. However, if the matrix for $R = R\,\mathrm{id}_3$ (with respect to the natural basis) is wanted, we read it off as

$$R = \cos(\theta)\,\mathrm{id}_3 + (1 - \cos(\theta))[a][a]^{\mathrm{t}} + \sin(\theta)(a\times),$$

with

$$a\times := \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix}$$

the matrix for the linear map $r \mapsto a \times r$.

Markov Chains

Recall from page 100 our example of a random walk on some graph. There we were interested in the matrices M^k, $k = 1, 2, 3, \ldots$, with the entries of the square matrix M all nonnegative and all entries in any particular row adding up to 1. In other words, $M \geq 0$ and $Me = e$, with

$$e := (1, 1, \ldots, 1).$$

In particular, $1 \in \mathrm{spec}(M)$. Further, since $\|M\|_\infty = 1$, we conclude from (13.1) that $\rho(M) \leq 1$. Hence, 1 is an absolutely largest eigenvalue for M. Assume, in addition, that M is irreducible. This is certainly guaranteed if $M > 0$. Then, by the Perron-Frobenius theory, 1 is a nondefective eigenvalue of M, and is the unique absolutely largest eigenvalue. By (11.10)Theorem, this implies that M is convergent. In fact, since 1 is a nondefective simple eigenvalue of M with corresponding eigenvector e, there is a basis $V = [e, W]$, with W a basis for $\mathrm{ran}(M - \mathrm{id})$, hence

$$MV = [e, MW] = V\,\mathrm{diag}(1, B)$$

for some B with $\rho(B) < 1$. Therefore,

$$M^k V = V\,\mathrm{diag}(1, B^k) \xrightarrow[k\to\infty]{} V\,\mathrm{diag}(1, 0).$$

In other words,

$$\lim_{k\to\infty} M^k = eu^{\mathrm{t}},$$

with $M^{\mathrm{t}}u = u$, i.e., u is an eigenvector of M^{t} belonging to the eigenvalue 1. In particular, all rows of M^k converge to this particular nonnegative vector whose entries sum to 1.

An example from CAGD

In Computer-Aided Geometric Design, one uses repeated corner-cutting to refine a given polygon into a smooth curve of approximately the same shape. The best-known example is the **Chaikin algorithm**. This algorithm consists in applying repeatedly, until satisfied, the following step:

input: the vertices $x_1, x_2, \ldots, x_n, x_{n+1} := x_1 \in \mathbb{R}^2$ of a closed polygon.
for $j = 1 : n$**, do:** $y_{2j-1} \leftarrow (3x_j + x_{j+1})/4$; $y_{2j} \leftarrow (x_j + 3x_{j+1})/4$; **enddo**
output: the vertices $y_1, y_2, \ldots, y_{2n}, y_{2n+1} := y_1 \in \mathbb{R}^2$ of a closed polygon that is inscribed into the input polygon.

In other words,

$$[y_1, \ldots, y_{2n}] = [x_1, \ldots, x_n]C_n,$$

with C_n the $n \times (2n)$-matrix

$$C_n := \begin{bmatrix} 3 & 1 & 0 & 0 & 0 & 0 & \cdots & 1 & 3 \\ 1 & 3 & 3 & 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & 3 & 3 & 1 & \cdots & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 3 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ldots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 3 & 1 \end{bmatrix} /4.$$

It is possible to show that, as $k \to \infty$, the polygon with vertex sequence

$$[x_1^{(k)}, \ldots, x_{2^k n}^{(k)}] := [x_1, \ldots, x_n]C_n C_{2n} \cdots C_{2^k n}$$

converges to a smooth curve, namely the curve

$$t \mapsto \sum_j x_j B_2(t - j),$$

with B_2 a certain smooth piecewise quadratic function, a so-called quadratic B-spline (whatever that may be).

Here, we consider the following much simpler and more radical corner-cutting:

$$[y_1, \ldots, y_n] = [x_1, \ldots, x_n]A,$$

with

$$\text{(14.3)} \qquad A := \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 1 \\ 1 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 1 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ldots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 1 & 1 \end{bmatrix} /2.$$

In other words, the new polygon is obtained from the old by choosing as the new vertices the midpoints of the edges of the old.

Simple examples, hand-drawn, quickly indicate that the sequence of polygons, with vertex sequence

$$[x_1^{(k)}, \ldots, x_n^{(k)}] := [x_1, \ldots, x_n]A^k$$

seem to shrink eventually into a point. Here is the analysis that this is, in fact, the case, with that limiting point equal to the average, $\sum_j x_j/n$, of the original vertices.

(i) The matrix A, defined in (14.3), is a **circulant**, meaning that each row is obtainable from its predecessor by shifting everything one to the right, with the right-most entry in the previous row becoming the left-most entry of the current row. All such matrices have eigenvectors of the form

$$u_\lambda := (\lambda^1, \lambda^2, \ldots, \lambda^n),$$

with the scalar λ chosen so that $\lambda^n = 1$, hence $\lambda^{n+1} = \lambda$. For our A, we compute

$$Au_\lambda = (\lambda^n + \lambda^1, \lambda^1 + \lambda^2, \ldots, \lambda^{n-1} + \lambda^n)/2.$$

Hence, if $\lambda^n = 1$, then

$$Au_\lambda = \frac{1+\lambda}{2\lambda} u_\lambda.$$

(ii) The equation $\lambda^n = 1$ has exactly n distinct solutions, namely the n **roots of unity**

$$\lambda_j := \exp(2\pi \mathrm{i} j/n) = \omega^j, \quad j = 1{:}n.$$

Here,

$$\omega := \omega_n := \exp(2\pi\mathrm{i}/n)$$

is a **primitive nth root of unity**. Note that

$$\overline{\omega} = 1/\omega.$$

Let

$$V = [v_1, \ldots, v_n] := [u_{\lambda_1}, \ldots, u_{\lambda_n}]$$

be the column map whose jth column is the eigenvector

$$v_j := (\omega^j, \omega^{2j}, \ldots, \omega^{nj})$$

of A, with corresponding eigenvalue

$$\mu_j := \frac{1+\lambda_j}{2\lambda_j} = (\omega^{-j} + 1)/2, \quad j = 1{:}n.$$

Since these eigenvalues are distinct, V is 1-1 (by (10.9)Lemma), hence V is a basis for $\mathbb{C}^n$. In particular,

$$A = V \operatorname{diag}(\ldots, \mu_j, \ldots) V^{-1}.$$

(iii) It follows that

$$A^k = V \operatorname{diag}(\ldots, \mu_j^k, \ldots) V^{-1} \xrightarrow[k\to\infty]{} V \operatorname{diag}(0, \ldots, 0, 1) V^{-1}$$

since $|\mu_j| < 1$ for $j < n$, while $\mu_n = 1$. Hence

$$\lim_{k\to\infty} A^k = v_n V^{-1}(n,:).$$

(iv) In order to compute $V^{-1}(n,:)$, we compute $V^{\mathrm{c}}V$ (recalling that $\overline{\omega^r} = \omega^{-1}$):

$$(V^{\mathrm{c}}V)(j,k) = {v_j}^{\mathrm{c}} v_k = \sum_{r=1}^n \omega^{-rj}\ \omega^{rk} = \sum_{r=1}^n \omega^{(k-j)r}.$$

That last sum is a geometric series, of the form $\sum_{r=1}^n \nu^r$ with $\nu := \omega^{k-j}$, hence equals n in case $k = j$, and otherwise $\nu \neq 1$ and the sum equals $(\nu^{n+1} - \nu)/(\nu - 1) = 0$ since $\nu^n = 1$, hence $\nu^{n+1} - \nu = 0$. It follows that

$$V^{\mathrm{c}}V = n\, \mathrm{id}_n,$$

i.e., $V/\sqrt{n}$ is *unitary*, i.e., an o.n. basis for $\mathbb{C}^n$. In particular, $V^{-1} = V^{\mathrm{c}}/n$, therefore

$$V^{-1}(n,:) = {v_n}^{\mathrm{c}}/n.$$

(v) It follows that

$$\lim_{k\to\infty} A^k = (1/n) v_n {v_n}^{\mathrm{c}},$$

with

$$v_n = (1, 1, \ldots, 1).$$

Consequently,

$$\lim_{k\to\infty} [\ldots, x_j^{(k)}, \ldots] \ = \ \sum_j x_j/n\ {v_n}^{\mathrm{c}} \ = \ [\ldots, \sum_j x_j/n, \ldots],$$

i.e., the rank-one matrix all of whose columns equal the average $\sum_j x_j/n$ of the vertices of the polygon we started out with.

Tridiagonal Toeplitz matrix

Circulants are a special case of *Toeplitz* matrices, i.e., of matrices that are constant along diagonals. Precisely, the matrix A is **Toeplitz** if

$$A(i,j) = a_{i-j}, \quad \forall i,j,$$

for some sequence $(\ldots, a_{-2}, a_{-1}, a_0, a_1, a_2, \ldots)$ of appropriate domain. Circulants are special in that the determining sequence a for them is periodic, i.e., $a_{i+n} = a_i$, all i, if A is of order n.

Consider now the case of a **tridiagonal** Toeplitz matrix A. For such a matrix, only the (main) diagonal and the two next-to-main diagonals are (perhaps) nonzero; all other entries are zero. This means that only a_{-1}, a_0, a_1 are, perhaps, nonzero, while $a_i = 0$ for $|i| > 1$. If also $a_{-1} = a_1 \neq 0$, then the circulant trick, employed in the preceding section, still works, i.e., we can fashion some eigenvectors from vectors of the form $u_\lambda = (\lambda^1, \ldots, \lambda^n)$. Indeed, now

$$(Au_\lambda)_j = \begin{cases} a_0\lambda + a_1\lambda^2 & \text{for } j = 1; \\ a_1\lambda^{j-1} + a_0\lambda^j + a_1\lambda^{j+1} & \text{for } j = 2{:}n{-}1; \\ a_1\lambda^{n-1} + a_0\lambda^n & \text{for } j = n. \end{cases}$$

Hence,

$$Au_\lambda = (a_1/\lambda + a_0 + a_1\lambda)u_\lambda - a_1(e_1 + \lambda^{n+1}e_n).$$

At first glance, this doesn't look too hopeful since we are after eigenvectors. However, recall that, for a unimodular λ, i.e., for $\lambda = \exp \mathrm{i}\varphi$ for some real φ, we have $1/\lambda = \overline{\lambda}$, hence

$$Au_{\overline{\lambda}} = (a_1/\lambda + a_0 + a_1\lambda)u_{\overline{\lambda}} - a_1(e_1 + \overline{\lambda^{n+1}}e_n).$$

It follows that, by choosing λ as an $(n+1)$st root of unity, i.e.,

$$\lambda = \lambda_j := \exp(2\pi \mathrm{i} j/(n+1)), \quad j = 1{:}n,$$

and setting

$$v_j := (u_\lambda - u_{\overline{\lambda}})/(2\mathrm{i}) = (\sin(2\pi kj/(n+1)) : k = 1{:}n),$$

we obtain

$$Av_j = \mu_j v_j$$

with

$$\mu_j := a_0 + a_1(\lambda_j + \overline{\lambda_j}) = a_0 + 2a_1\cos(2\pi j/(n+1)).$$

Since we assumed that $a_1 \neq 0$, these n numbers μ_j are pairwise distinct, hence $V =: [v_1, \ldots, v_n]$ is 1-1 by (10.9)Lemma, hence a basis for $\mathbb{C}^n$. In fact, since V maps $\mathbb{R}^n$ to $\mathbb{R}^n$, V is a basis for $\mathbb{R}^n$. Hence if both a_0 and a_1 are real, then also each μ_j is real and then, A is diagonable even over $\mathbb{F} = \mathbb{R}$.

Linear Programming

In Linear Programming, one seeks a minimizer for a *linear* **cost function**

$$x \mapsto c^{\mathrm{t}}x$$

on the set

$$F := \{x \in \mathbb{R}^n : Ax \leq b\}$$

of all n-vectors x that satisfy the m **linear constraint**s

$$A(i,:)^{\mathrm{t}}x \leq b_i, \quad i = 1{:}m,$$

with $c \in \mathbb{R}^n$, $A \in \mathbb{F}^{m\times n}$, $b \in \mathbb{R}^m$ given. Here and below, for $y, z \in \mathbb{R}^m$,

$$y \leq z \;:=\; z - y \in \mathbb{R}^m_+ := \{u \in \mathbb{R}^m : 0 \leq u_j, j = 1{:}m\}.$$

The set F, also called the **feasible set**, is the intersection of m halfspaces, i.e., sets of the form

$$H(a,b) := \{x \in \mathbb{R}^n : a^{\mathrm{t}}x \le b\}.$$

Such a **halfspace** consists of all the points that lie on that side of the corresponding **hyperplane**

$$h(a,b) := \{x \in \mathbb{R}^n : a^{\mathrm{t}}x = b\}$$

that the **normal** a of the hyperplane points away from; see (2.4)Figure, or (14.4)Figure.

Here is a simple example: Minimize

$$2x_1 + x_2$$

over all $x \in \mathbb{R}^2$ for which

$$x_2 \ge -2, \quad 3x_1 - x_2 \le 5, \quad x_1 + x_2 \le 3$$

$$x_1 - x_2 \ge -3, \quad 3x_1 + x_2 \ge -5.$$

In matrix notation, and more uniformly written, this is the set of all $x \in \mathbb{R}^2$ for which $Ax \le b$ with

$$A := \begin{bmatrix} 0 & -1 \\ 3 & -1 \\ 1 & 1 \\ -1 & 1 \\ -3 & -1 \end{bmatrix}, \quad b := \begin{bmatrix} 2 \\ 5 \\ 3 \\ 3 \\ 5 \end{bmatrix}.$$

In this simple setting, you can visualize the set F by drawing each of the hyperplanes $h(A_{i:}, b_i)$ along with a vector pointing in the same direction as its normal vector, $A_{i:}$; the set F lies on the side that the normal vector points away from; see (14.4)Figure.

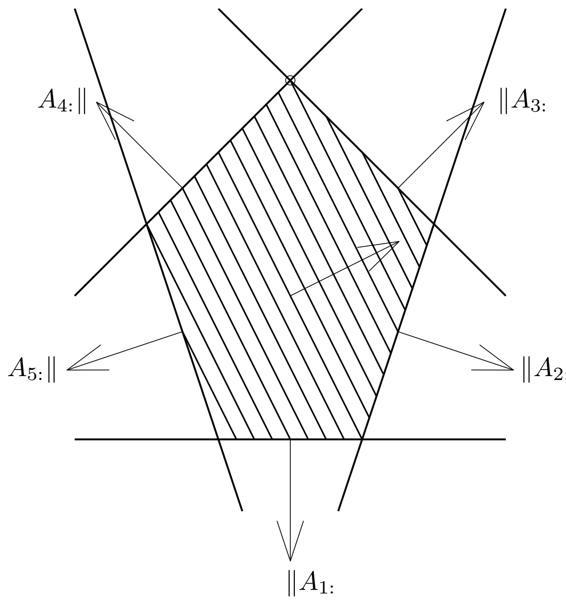

(14.4) Figure. The feasible set for five linear constraints in the plane, as filled out by some level lines of the cost function. Since the gradient of the cost function is shown as well, the location of the minimizer is clear.

□

In order to provide a handier description for F, one introduces the so-called **slack variables**

$$y := b - Ax;$$

earlier, we called this the *residual*. With their aid, we can describe F as

$$F = \{x \in \mathbb{R}^n : \exists y \in \mathbb{R}_+^m \text{ s.t. } (x, y, -1) \in \operatorname{null}[A,\ \mathrm{id}_m, b]\},$$

and use elimination to obtain a concise description of $\operatorname{null}[A,\ \mathrm{id}_m, b]$.

For this, assume that A is 1-1. Then, each column of A is bound, hence is also bound in $[A,\ \mathrm{id}, b]$. Therefore, after n steps of the (3.2)Elimination Algorithm applied to $[A,\ \mathrm{id}, b]$, we will arrive at a matrix B, with the same nullspace as $[A,\ \mathrm{id}, b]$, and an n-vector $\mathtt{nbs}$ (with $\mathtt{nbs}(k)$ the row used as pivot row for the kth unknown or column, all k), for which

$$B(\mathtt{nbs}, 1{:}n)$$

is upper triangular with nonzero diagonals while, with $\mathtt{bs}$ the rows not yet used as pivot rows,

$$B(\mathtt{bs}, 1{:}n) = 0.$$

Further, since the next m columns of $[A,\ \mathrm{id}_m, b]$ have nonzero entries in these pivot rows $\mathtt{nbs}$ only in columns $n + \mathtt{nbs}$, the other columns, i.e., columns $n + \mathtt{bs}$, will remain entirely unchanged. It follows that

$$B(\mathtt{bs}, n + \mathtt{bs}) = \ \mathrm{id}_{m-n}.$$

Therefore, after dividing each of the n pivot rows by their pivot element and then using each pivot row to eliminate its unknown also from all other pivot rows, we will arrive at a matrix, still called B, for which now

$$B([\mathtt{nbs}, \mathtt{bs}], \mathtt{bound}) = \ \mathrm{id}_m$$

with

$$\mathtt{bound} := [1{:}n, n + \mathtt{bs}]$$

the bound columns of $[A,\ \mathrm{id}, b]$.

For our particular example, the matrix B will be reached after just two steps:

$$[A,\ \mathrm{id}, b] = \begin{bmatrix} 0 & -1 & 1 & 0 & 0 & 0 & 0 & 2 \\ 3 & -1 & 0 & 1 & 0 & 0 & 0 & 5 \\ 1 & 1 & 0 & 0 & 1 & 0 & 0 & 3 \\ -1 & 1 & 0 & 0 & 0 & 1 & 0 & 3 \\ -3 & -1 & 0 & 0 & 0 & 0 & 1 & 5 \end{bmatrix} \to \begin{bmatrix} 0 & -1 & 1 & 0 & 0 & 0 & 0 & 2 \\ 0 & -4 & 0 & 1 & -3 & 0 & 0 & -4 \\ 1 & 1 & 0 & 0 & 1 & 0 & 0 & 3 \\ 0 & 2 & 0 & 0 & 1 & 1 & 0 & 6 \\ 0 & 2 & 0 & 0 & 3 & 0 & 1 & 14 \end{bmatrix}$$

$$\to \begin{bmatrix} 0 & 0 & 1 & 0 & 1/2 & 1/2 & 0 & 5 \\ 0 & 0 & 0 & 1 & -1 & 2 & 0 & 8 \\ 1 & 0 & 0 & 0 & 1/2 & -1/2 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1/2 & 1/2 & 0 & 3 \\ 0 & 0 & 0 & 0 & 2 & -1 & 1 & 8 \end{bmatrix} =: B,$$

with $\mathtt{nbs} = [3, 4]$, and $\mathtt{bs} = [1, 2, 5]$.

It follows that

$$\mathtt{free} := [n + \mathtt{nbs}, n + m + 1]$$

gives the free columns of $[A,\ \mathrm{id}, b]$. In particular, we can freely choose $y_{\mathtt{nbs}}$, i.e., the slack variables associated with the n pivot rows, and, once they are chosen, then x as well as the bound slack variables, $y_{\mathtt{bs}}$, are uniquely determined by the requirement that $(x, y, -1) \in \operatorname{null} B$.

This suggests eliminating x altogether, i.e., using the pivot rows $B(\mathtt{nbs}, :)$ to give

$$x = B(\mathtt{nbs}, \mathtt{end}) - B(\mathtt{nbs}, n + \mathtt{nbs})y_{\mathtt{nbs}},$$

(with `end` being `MATLAB`'s convenient notation for the final row or column index) and, with that, rewrite the cost function in terms of $y_{\texttt{nbs}}$:

$$y_{\texttt{nbs}} \mapsto c^{t}B(\texttt{nbs},\texttt{end}) - c^{t}B(\texttt{nbs}, n+\texttt{nbs})y_{\texttt{nbs}}.$$

Correspondingly, we simplify the working array B in the following two ways:

(i) We append the row $B(m+1,:) := c^{t}B(\texttt{nbs},:)$.

(ii) Then, we drop entirely the n rows `nbs` (storing those rows perhaps in some other place against the possibility that we need to compute x from $y_{\texttt{nbs}}$ at some later date), and also drop the first n columns.

In our example, this leaves us with the following, smaller, array B:

$$B = \begin{bmatrix} 1 & 0 & 1/2 & 1/2 & 0 & 5 \\ 0 & 1 & -1 & 2 & 0 & 8 \\ 0 & 0 & 2 & -1 & 1 & 8 \\ 0 & 0 & 3/2 & -1/2 & 0 & 3 \end{bmatrix}, \quad \texttt{bs} = [1,2,5], \quad \texttt{nbs} = [3,4].$$

This change of independent variables, from x to $y_{\texttt{nbs}}$, turns the n hyperplanes $h(A_{k:}, b(k))$, $k \in \texttt{nbs}$, into coordinate planes; see (14.5)Figure. In particular, the choice $y_{\texttt{nbs}} = 0$ places us at the (unique) point of intersection of these n hyperplanes. In our example, that point is $x = (0,3)$, and it is marked in (14.4)Figure, and functions as the origin in (14.5)Figure(a).

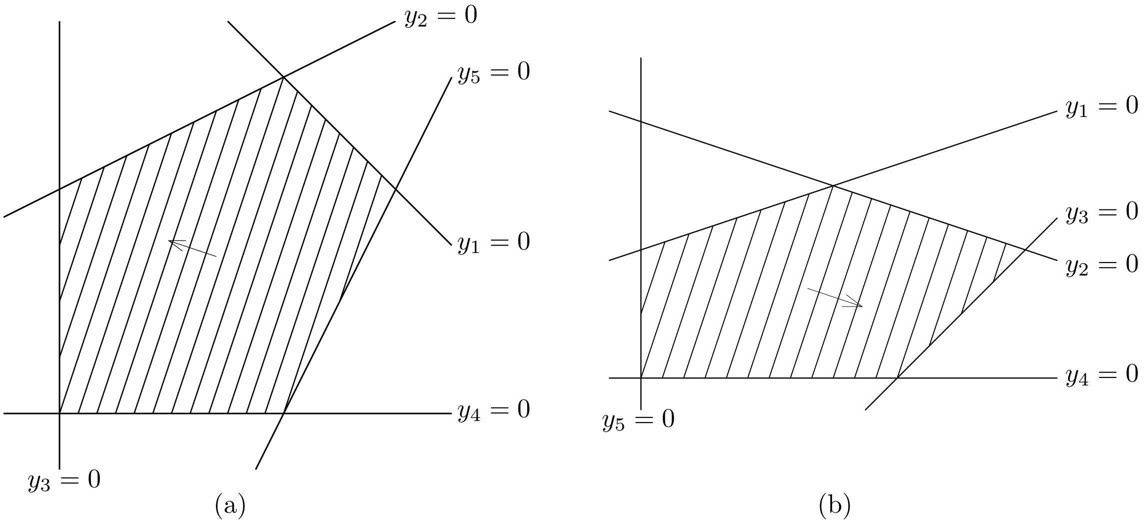

(14.5) Figure. The feasible set for five linear constraints in the plane, as filled out by some level lines of the cost function, viewed in terms of the (nonbasic) variables (a) y_3, y_4; (b) y_5, y_4. From the latter, the minimizing vertex will be reached in one step of the Simplex Method.

In terms of this B as just constructed, and with

$$m' := m - n = \#\texttt{bs},$$

our minimization problem now reads: *Minimize the cost function*

$$y_{\texttt{nbs}} \mapsto B(\texttt{end},\texttt{end}) - B(\texttt{end},\texttt{nbs})^{t} y_{\texttt{nbs}} \tag{14.6}$$

over all $y_{\texttt{nbs}} \in \mathbb{R}^{n}_{+}$ *for which*

$$B(\texttt{bs},\texttt{nbs})\, y_{\texttt{nbs}} \le B(\texttt{bs},\texttt{end}),$$

i.e., for which

$$y_{\mathtt{bs}} := B(\mathtt{bs},\mathtt{end}) - B(\mathtt{bs},\mathtt{nbs})y_{\mathtt{nbs}} \in \mathbb{R}_+^{m'}.$$

This is the form in which linear programming problems are usually stated, and from which most textbooks start their discussion of such problems.

Note how easily accessible various relevant information now is.

(i) $B(\mathtt{end},\mathtt{end})$ tells us the value of the cost function at the current point, $y_{\mathtt{nbs}} = 0$.

(ii) For any $k \in \mathtt{nbs}$, the entry $B(\mathtt{end},k)$ tells us how the cost function would change if we were to change the value of the nonbasic variable y_k in the only way permitted, i.e., from 0 to something positive. Such a move would lower the cost function if and only if $B(\mathtt{end},k) > 0$.

(iii) Our current point, $y_{\mathtt{nbs}} = 0$, is feasible if and only if $B(1{:}m',\mathtt{end}) \ge 0$.

(iv) If we were to change the nonbasic variable y_k from zero to something positive, then the basic variable $y_{\mathtt{bs}(i)}$ would change, from $B(i,\mathtt{end})$ to $B(i,\mathtt{end}) - B(i,k)y_k$. Hence, assuming $B(i,\mathtt{end}) > 0$ and $B(i,k) > 0$, we could change y_k only to $B(i,\mathtt{end})/B(i,k)$ before the $\mathtt{bs}(i)$th constraint would be violated.

In our example (have a look at (14.5)Figure(a)), we already observed that our current point, $y_{\mathtt{nbs}} = 0$, is, indeed, feasible. But we notice that, while $B(\mathtt{end},4) < 0$, hence any feasible change of y_4 would only increase the cost function (14.6), $B(\mathtt{end},3)$ is positive, hence know that we can further decrease the cost function (14.6) by increasing y_3. Such a change is limited by concerns for the positivity of y_1 and y_5. As for y_1, we would reach $y_1 = 0$ when $y_3 = 5/(1/2) = 10$, while, for y_5, we would reach $y_5 = 0$ when $y_3 = 8/2 = 4$. We take the smaller change and thereby end up at a new vector y, with $y_4 = 0 = y_5$, i.e., are now at the intersection of the constraints 4 and 5, with the cost further reduced by $(3/2)4 = 6$, to -3.

In other words, after this change, y_4 and y_5 are now the nonbasic variables. In order to have our B tell us about this new situation, and since $5 = \mathtt{bs}(3)$, we merely divide its 3rd row by $B(3,3)$, then use the row to eliminate y_3 from all other rows of B. This leads to

$$B = \begin{bmatrix} 1 & 0 & 0 & 3/4 & -1/4 & 3 \\ 0 & 1 & 0 & 3/2 & 1/2 & 12 \\ 0 & 0 & 1 & -1/2 & 1/2 & 4 \\ 0 & 0 & 0 & 1/2 & -3/4 & -3 \end{bmatrix}, \quad \mathtt{bs} = [1,2,3], \quad \mathtt{nbs} = [5,4].$$

In particular, we readily see that the cost at $y_4 = 0 = y_5$ is, indeed, -3. We also see (see also (14.5)Figure(b)) that it is possible to reduce the cost further by changing y_4, from 0 to something positive. Such a change would reduce $y_1 = 3$ by $(3/4)y_4$ and would reduce $y_2 = 12$ by $(3/2)y_4$. Hence, this change is limited to the smaller of $3/(3/4) = 4$ and $12/(3/2) = 8$, i.e., to the change $y_4 = 4$ that makes $y_1 = 0$.

We carry out this exchange, of y_4 into $\mathtt{bs}$ and y_1 into $\mathtt{nbs}$, by dividing $B(1,:)$ by $B(1,4)$ and then using that row to eliminate y_4 from all other rows, to get the following B:

$$B = \begin{bmatrix} 4/3 & 0 & 0 & 1 & -1/3 & 4 \\ -2 & 1 & 0 & 0 & 1 & 6 \\ 2/3 & 0 & 1 & 0 & 1/3 & 6 \\ -1/3 & 0 & 0 & 0 & -2/3 & -4 \end{bmatrix}, \quad \mathtt{bs} = [4,2,3], \quad \mathtt{nbs} = [5,1].$$

In particular, now $B(\mathtt{end},\mathtt{nbs}) \le 0$, showing that no further improvement is possible, hence -4 is the minimum of the cost function on the feasible set. At this point, $y_{3:4} = (6,4)$, hence, from the rows used as pivot rows to eliminate x, we find that, in terms of x, our optimal point is $x = (0,3) - (1/2)\begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}(6,4) = -(1,2)$, and, indeed, $c^{\mathrm{t}}x = (2,1)^{\mathrm{t}}(-1,-2) = -4$.

The steps just carried out for our example are the standard steps of the **Simplex Method**. In this method (as proposed by Dantzig), one examines the value of the cost function only at a **vertex**, i.e., at the unique intersection of n of the constraint hyperplanes, i.e., at a point corresponding to $y_{\mathtt{nbs}} = 0$ for some choice of $\mathtt{nbs}$. Assuming that vertex feasible, one checks whether $B(\mathtt{end},\mathtt{nbs}) \le 0$. If it is, then one knows that one is at the minimum. Otherwise, one moves to a neighboring vertex at which the cost is less by exchanging some y_k for which $B(\mathtt{end}, k) > 0$ with some $y_{\mathtt{bs}(i)}$ with i chosen as the minimizer for $B(i, \mathtt{end})/B(i, k)$ over all i with $B(i, k) > 0$. This exchange is carried out by just one full elimination step applied to B, by dividing $B(i, :)$ by $B(i, k)$ and then using this row to eliminate y_k from all other rows, and then updating the sequences `bs` and `nbs`.

This update step is one full elimination step (sometimes called a **(Gauss-)Jordan** step in order to distinguish it from the **Gauss** step, in which the unknown is eliminated only from the rows not yet used as pivot rows).

Since all the information contained in the columns $B(:, \mathtt{bs})$ is readily derivable from `bs` and `nbs`, one usually doesn't bother to carry these columns along. This makes the updating of the matrix $B(:, \mathtt{nbs})$ a bit more mysterious.

Finally, there are the following points to consider:

unbounded feasible set If, for some $k \in \mathtt{nbs}$, $B(\mathtt{end}, k)$ is the only positive entry in its column, then increasing y_k will strictly decrease the cost *and* increase all basic variables. Hence, if $y_{\mathtt{nbs}} = 0$ is a feasible point, then we can make the cost function on the feasible set as close to $-\infty$ as we wish. In our example, this would be the case if we dropped constraints 1 and 5. Without these constraints, in our very first Simplex step, we could have increased y_3 without bound and so driven the cost to $-\infty$.

finding a feasible point In our example, we were fortunate in that the very first vertex we focused on was feasible. However, if it is not, then one can use the very Simplex Method to obtain a feasible point, simply by introducing an additional variable, y_0, which is added to each infeasible row, and then using the Simplex Method to minimize the cost function

$$y \mapsto y_0.$$

In this, the variable y_0 starts off nonbasic, i.e., $y_0 = 0$, and, then, as an extraordinary first step, we would exchange y_0 for the most negative basic variable, and then proceed until the minimum of this auxiliary cost function is reached. If it is positive, then we now know that *the feasible set is empty*. Otherwise, the minimum is zero (i.e., y_0 is again a nonbasic variable) and can simply be dropped now since the point corresponding to the remaining nonbasic variables being zero is feasible.

Note that, in this way, the Simplex Method can be used to solve any finite set of linear inequalities in the sense of either providing a point satisfying them all or else proving that none exists.

convergence in finitely many steps If we can guarantee that, at each step, we strictly decrease the cost, then we must reach the vertex with minimal cost in finitely many steps since, after all, there are only finitely many vertices. A complete argument has to deal with the fact that the cost may not always strictly decrease because the current point may lie on more than just n of the constraint hyperplanes.

Approximation by broken lines

Flats: points, vectors, barycentric coordinates, differentiation

In CAGD and Computer Graphics, Linear Algebra is mainly used to change one's point of view, that is, to change coordinate systems. In this, even the familiar 3-space, $\mathbb{R}^3$, is often treated as an 'affine space' or 'flat' rather than a vector space, in order to deal simply with useful maps other than linear maps, namely the affine maps.

For example, the **translation**

$$\mathbb{R}^3 \to \mathbb{R}^3 : p \mapsto p + v$$

of $\mathbb{R}^3$ by the vector v is not a linear map. Nevertheless, it can be represented by a matrix, using the following trick. Embed $\mathbb{R}^3$ into $\mathbb{R}^4$ by the 1-1 map

$$\mathbb{R}^3 \to \mathbb{R}^4 : x \mapsto (x, 1).$$

The image of $\mathbb{R}^3$ under this map is the 'flat'

$$F := \mathbb{R}^3 \times 1 = \{(x,1) : x \in \mathbb{R}^3\}.$$

Consider the linear map on $\mathbb{R}^4$ given by

$$T_v := \begin{bmatrix} \mathrm{id}_3 & v \\ 0 & 1 \end{bmatrix}.$$

Then, for any $x \in \mathbb{R}^3$,

$$T_v(x,1) = (\,\mathrm{id}_3 x\; + v, 0^{\mathrm{t}} x + 1) = (x+v, 1).$$

In other words, the linear map T_v carries F into itself in such a way that the point $p = (x,1)$ is carried to its 'translate' $(p+v,1) = p + (v,0)$.

Let, now, $\widehat{A} \in \mathbb{R}^{4\times 4}$ be an arbitrary linear map on $\mathbb{R}^4$ subject only to the condition that it map F into itself. Breaking up $\widehat{A}$ in the same way as we did T_v, i.e.,

$$\widehat{A} =: \begin{bmatrix} A_0 & v \\ [u]^{\mathrm{t}} & t \end{bmatrix},$$

we get

$$\widehat{A}\,(x,1) = (A_0 x\; + v, u^{\mathrm{t}} x + t),$$

hence want $u^{\mathrm{t}} x + t = 1$ for all $x \in \mathbb{R}^3$, and this holds if and only if $u = 0$ and $t = 1$, i.e.,

$$\widehat{A} =: \begin{bmatrix} A_0 & v \\ 0 & 1 \end{bmatrix}$$

is the most general such map.

Look at its action as concerns the general point $p = (x,1)$ in F: After subtracting $p_0 := (0,1)$ from p, we obtain a vector in the linear subspace

$$F{-}F = \{p - q : p, q \in F\} = \mathrm{ran}[e_1, e_2, e_3].$$

Since $\widehat{A}$ maps F into itself, it also maps this subspace into itself, hence $\widehat{A}(p - p_0) = (A_0 \tau_{p-p_0}, 0)$ is again in this subspace and, after adding to it the element $\widehat{A}p_0 = (v,1) \in F$, we finally obtain $\widehat{A}p$ as $\widehat{A}(p-p_0) + \widehat{A}p_0$.

Three-dimensional plots in `MATLAB` show, in fact, the orthogonal projection onto the (x,y)-plane *after* an affine transformation of $\mathbb{R}^3$ that makes the center of the plotting volume the origin and a rotation that moves a line, specified by azimuth and elevation, to the z-axis. This affine map is recorded in a matrix of order 4, obtainable by the command `view`, and also changeable by that command, but, fortunately, in down-to-earth terms like *azimuth* and *elevation*, or *viewing angle*.

□

Consider now the question which weighted sums

$$\sum_{j=0}^{r} p_j \alpha_j$$

of $p_j \in F$ are again in F. Apparently, all that is required is that

$$\sum_{j=0}^{r} \alpha_j = 1.$$

Such a weighted sum is called an **affine combination**. Thus, as far as the set F is concerned, these are the only linear combinations allowed. Note that such an affine sum can always be rewritten as

$$p_0 + \sum_{j=1}^{r} (p_j - p_0)\alpha_j,$$

where now the weights α_j, $j = 1{:}r$, are *arbitrary.* In other words, an affine sum on F is obtained by adding to some point in F an arbitrary weighted sum of elements in the vector space $F{-}F$.

An **affine map** on F is any map from F to F that preserves affine combinations, i.e., for which

$$A(p_0 + \sum_j (p_j - p_0)\alpha_j) = Ap_0 + \sum_j (Ap_j - Ap_0)\alpha_j$$

for all $p_j \in F$, $\alpha_j \in \mathbb{R}$. It follows that the map on $F{-}F$ defined by

$$A_0 : F{-}F \to F{-}F : p - q \mapsto Ap - Aq$$

must be well-defined and linear, hence A is necessarily the restriction to F of some linear map $\widehat{A}$ on $\mathbb{R}^4$ that carries F into itself and therefore also carries the linear subspace $F{-}F$ into itself.

The main pay-off, in CAGD and in Computer Graphics, of these considerations is the fact that one can represent the composition of affine maps by the product of the corresponding matrices.

This concrete example has led to the following abstract definition of a flat, whose notational conventions strongly reflect the concrete example. It should be easy for you to verify that the standard example is, indeed, a flat in the sense of this abstract definition.

(14.7) Definition: A **flat** or **affine space** or **linear manifold** is a nonempty set F of **point**s, a vector space T of **translation**s, and a map

(14.8) $$\varphi : F \times \mathrm{T} \to F : (p, \tau) \mapsto \tau(p) =: p + \tau$$

satisfying the following:

(a) $\forall\{(p,\tau) \in F \times \mathrm{T}\}\;\; p + \tau = p \quad \Longleftrightarrow \quad \tau = 0.$

(b) $\forall\{\tau, \sigma \in \mathrm{T}\}\;\; (\cdot + \tau) + \sigma = \cdot + (\tau + \sigma).$

(c) $\exists\{p_0 \in F\}\;\; \varphi(p_0, \cdot)$ is onto.

Translations are also called **vector**s since (like 'vehicles' or 'conveyors', words that have the same Latin root as 'vector') they carry points to points.

Condition (a) ensures the uniqueness of the solution of the equation $p{+}? = q$ whose existence (see the proof of (3) below) is guaranteed by (c).

Condition (b) by itself is already satisfied, for arbitrary F and T, by, e.g., $\varphi : (p, \tau) \mapsto p$.

Condition (c) is needed to be certain that T is rich enough. (a) + (b) is already satisfied, e.g., by $\mathrm{T} = \{0\}$, $\varphi(\cdot, 0) = \mathrm{id}$. As we will see in a moment, (a)+(b)+(c) together imply that $\varphi(p, \cdot)$ is onto for every $p \in F$. In other words, there is nothing special about the p_0 that appears in (c). In fact, the notion of a flat was developed explicitly as a set that, in contrast to a vector space which has an origin, does not have a distinguished point.

Consequences

(1) $\varphi(\cdot, 0) = \mathrm{id}$ (by (a)).

(2) For any $\tau \in \mathrm{T}$, $\varphi(\cdot, \tau)$ is invertible; its inverse is $\varphi(\cdot, -\tau)$ (by (1) and (b)). The corresponding abbreviation

$$p - \tau := p + (-\tau)$$

is helpful and standard.

(3) $\forall\{p, q \in F\}\ \exists!\{\tau \in \mathrm{T}\}\ \ p + \tau = q$. This unique τ is correspondingly denoted

$$q - p.$$

Proof: If $p+\tau = q = p+\sigma$, then, by (2) and (b), $p = q+(-\sigma) = (p+\tau)+(-\sigma) = p+(\tau-\sigma)$, therefore, by (1), $\tau - \sigma = 0$, showing the *uniqueness* of the solution to $p+? = q$, regardless of p and q. The *existence* of a solution is, offhand, only guaranteed, by (c), for $p = p_0$. However, with the invertibility of $\varphi(p_0, \cdot) : \mathrm{T} \to F$ thus established, hence with $p - p_0$ and $q - p_0$ well-defined, we have $q = p_0 + (q - p_0)$ and $p = p_0 + (p - p_0)$, hence $p_0 = p - (p - p_0)$, therefore

$$q = p - (p - p_0) + (q - p_0),$$

showing that the equation $p+? = q$ has a solution (namely the vector $(q - p_0) - (p - p_0)$). □

(4) Note that (3) provides a 1-1 correspondence (in many different ways) between F and T. Specifically, for any particular $o \in F$,

$$F \to \mathrm{T} : p \mapsto p - o$$

is an invertible map, as is its inverse,

$$\mathrm{T} \to F : \tau \mapsto o + \tau.$$

However, the wish to avoid such an arbitrary choice of an 'origin' o in F provided the impetus to define the concept of flat in the first place. The **dimension of a flat** is, by definition, the dimension of the associated vector space of translations. Also, since the primary focus is usually the flat, F, it is very convenient to write its vector space of translations as

$$F{-}F.$$

(5) The discussion so far has only made use of the additive structure of T. Multiplication by scalars provides additional structure. Thus, for arbitrary $Q \subset F$, the **affine hull** of Q is, by definition,

$$\flat(Q) := q + \mathrm{span}(Q - q),$$

with the right side certainly independent of the choice of $q \in Q$, by (4). The affine hull of Q is, itself, a flat, with $\mathrm{span}(Q - q)$ the vector space of its translations.

(6) In particular, the affine hull of a finite subset Q of F is

$$\flat(Q) = q_0 + \mathrm{ran}[q - q_0 : q \in Q\backslash q_0], \quad q_0 \in Q.$$

Let

$$q_0 + \sum_{q \neq q_0} (q - q_0)\alpha_q$$

be one of its elements. In order to avoid singling out $q_0 \in Q$, it is customary to write instead

$$\sum_q q\alpha_q, \quad \text{with} \ \ \alpha_{q_0} := 1 - \sum_{q \neq q_0} \alpha_q.$$

This makes $\flat(Q)$ the set of all **affine combination**s

$$\sum_{q\in Q} q\alpha_q, \quad \sum_q \alpha_q = 1,$$

of the elements of Q. The affine hull $\flat(q_0,\ldots,q_r)$ of a sequence $q_0,\ldots,q_r$ in F is defined analogously. But I prefer to work here with the set Q in order to stress the point of view that, in a flat, all points are of equal importance.

A special case is the straight line through $p \ne q$, i.e.,

$$\flat(p,q) = p + \mathrm{I\!R}(q-p) = q + \mathrm{I\!R}(p-q) = \{(1-\alpha)p + \alpha q : \alpha \in \mathrm{I\!R}\}.$$

(7) The finite set $Q \subset F$ is called **affinely independent** in case, for some (hence for every) $o \in Q$, $[q - o : q \in Q\backslash o]$ is 1-1. In that case, each $p \in \flat(Q)$ can be written in exactly one way as an affine combination

$$p =: \sum_q q\ell_q(p), \quad \sum_q \ell_q(p) = 1,$$

of the $q \in Q$. Indeed, in that case, for any particular $o \in Q$, $V_o := [q - o : q \in Q\backslash o]$ is a basis for the vector space of translations on $\flat(Q)$, hence, for all $p \in \flat(Q)$,

$$p = o + (p-o) = o + V_oV_o^{-1}(p-o) = \sum_{q\in Q} q\ell_q(p),$$

with

$$(\ell_q(p) : q \in Q\backslash o) := V_o^{-1}(p-o), \quad \ell_o(p) := 1 - \sum_{q\ne o}\ell_q(p).$$

The 'affine' vector $\ell(p) = (\ell_q(p) : q \in Q) \in \mathrm{I\!R}^Q$ constitutes the **barycentric coordinates of** p **with respect to** Q.

It follows that, for arbitrary $p_i \in \flat(Q)$ and arbitrary $\alpha_i \in \mathrm{I\!R}$ with $\sum_i \alpha_i = 1$, we have

$$\sum_i \alpha_i p_i = \sum_i \alpha_i \sum_q \lambda_q(p_i)q = \sum_q (\sum_i \alpha_i\lambda_q(p_i))q,$$

with

$$\sum_i \alpha_i(\sum_q \lambda_q(p_i)) = \sum_i \alpha_i = 1.$$

Hence, by the uniqueness of the barycentric coordinates, the map

$$\lambda : \flat(Q) \to \mathrm{I\!R}^Q : p \mapsto (\lambda_q(p) : q \in Q)$$

is **affine**, meaning that

$$\lambda(\sum_i \alpha_i p_i) = \sum_i \alpha_i\lambda(p_i).$$

It is also 1-1, of course, and so is, for our flat $\flat(Q)$, what a coordinate map is for a vector space, namely a convenient structure-preserving numerical representation of the flat.

It follows that, with $f_0 : Q \to G$ an arbitrary map on Q into some flat G, the map

$$f : \flat(Q) \to G : \sum_{q\in Q}\lambda_q(p)q \mapsto \sum_{q\in Q}\lambda_q(p)f_0(q)$$

is affine. Hence, if $A : f \to G$ is an affine map that agrees with f_0 on Q, then it must equal f.

(8) Let the $r+1$-subset Q of the r-dimensional flat F be affinely independent. Then, for any $o \in Q$, $[q-o : q \in Q\backslash o]$ is a basis for $F{-}F$, and the scalar-valued map

$$\ell_o : F \to \mathbb{R} : p \mapsto \ell_o(p)$$

is a **linear polynomial** on F. Some people prefer to call it an **affine polynomial** since, after all, it is not a *linear* map. However, the adjective 'linear' is used here in the sense of 'degree ≤ 1', in distinction to quadratic, cubic, and higher-degree polynomials. A description for the latter can be obtained directly from the ℓ_q, $q \in Q$, as follows. The column map

$$[\ell_\alpha := \prod_{q\in Q} (\ell_q)^{\alpha(q)} : \alpha \in \mathbb{Z}_+^Q, |\alpha| = k]$$

into $\mathbb{R}^F$ is a basis for the (scalar-valued) polynomials of degree $\le k$ on F.

(9) An affine combination with nonnegative weights is called a **convex combination**. The weights being affine, hence summing to 1, they must also be no bigger than 1. The set

$$[q \,..\, q] := \{(1-\alpha)p + \alpha q : \alpha \in [0 \,..\, 1]\}$$

of all convex combinations of the two points p, q is called the **interval with endpoints** p**,** q. The set

$$\sigma_Q := \{\sum_{q\in Q} q\alpha_q : \alpha \in [0 \,..\, 1]^Q, \sum_q \alpha_q = 1\}$$

of all convex combinations of points in the finite set Q is called the **simplex with vertex set** Q in case Q is affinely independent.

(10) Flats are the proper setting for *differentiation*. Assume that the flat F is finite-dimensional. Then there are many ways to introduce a vector norm on the corresponding vector space $F{-}F$ of translations, hence a notion of convergence, but which vector sequences converge and which don't is independent of the choice of that norm. This leads in a natural way to convergence on F: *The point sequence* $(p_n : n \in \mathbb{N})$ *in* F **converges** *to* $p \in F$ *exactly when* $\lim_{n\to\infty} \|p_n - p\| = 0$. Again, this characterization of convergence does not depend on the particular vector norm on $F{-}F$ chosen.

With this, the function $f : F \to G$, on the finite-dimensional flat F to the finite-dimensional flat G, is **differentiable at** $p \in F$ in case the limit

$$D_\tau f(p) := \lim_{h\searrow 0} (f(p + h\tau) - f(p))/h$$

exists for every $\tau \in F{-}F$. In that case, $D_\tau f(p)$ is called the **derivative of** f **at** p **in the direction** τ.

Notice that $D_\tau f(p)$ is a *vector*, in $G{-}G$. It tells us the direction into which $f(p)$ gets translated as we translate p to $p+\tau$. Further, its magnitude gives an indication of the size of the change as a function of the size of the change in p. Exactly,

$$f(p + h\tau) = f(p) + hD_\tau f(p) + o(\|\tau\|h), \quad h \ge 0.$$

In particular, if f is differentiable at p, then

$$Df(p) : F{-}F \to G{-}G : \tau \mapsto D_\tau f(p)$$

is a well-defined map, from $F{-}F$ to $G{-}G$. This map is positively homogeneous, i.e.,

$$D_{h\tau} f(p) = hD_\tau f(p), \quad h \ge 0.$$

If this map $Df(p)$ is linear, it is called the **derivative of f at p**. Note that then

$$f(p+\tau) = f(p) + Df(p)\tau + o(\|\tau\|), \quad \tau \in F{-}F.$$

If V is any particular basis for $F{-}F$ and W is any particular basis for $G{-}G$, then the matrix

$$Jf(p) := W^{-1}\, Df(p)V$$

is the **Jacobian** of f at p. Its (i,j) entry tells us how much $f(p+\tau)$ moves in the direction of w_i because of a unit change in τ in the direction of v_j. More precisely, if $\tau = V\alpha$, then $Df(p)\tau = W\, Jf(p)\alpha$.

A practical high-point of these considerations is the **chain rule**, i.e., the observation that if $g : G \to H$ is a 'uniformly' differentiable map, then their composition gf, is differentiable, and

$$D(gf)(p) = Dg(f(p))Df(p).$$

grad, div, and curl

15. Optimization and quadratic forms

Minimization

We are interested in *minimizing* a given function

$$f : \operatorname{dom} f \subset \mathbb{R}^n \to \mathbb{R},$$

i.e., we are looking for $x \in \operatorname{dom} f$ so that

$$\forall y \in \operatorname{dom} f \quad f(x) \le f(y).$$

Any such x is called a **minimizer for** f; in symbols:

$$x \in \operatorname{argmin} f.$$

The discussion applies, of course, also to finding some $x \in \operatorname{argmax} f$, i.e., finding a **maximizer** for f, since $x \in \operatorname{argmax} f$ iff $x \in \operatorname{argmin}(-f)$.

Finding minimizers is, in general, an impossible problem since one cannot tell whether or not $x \in \operatorname{argmin} f$ except by checking *every* $y \in \operatorname{dom} f$ to make certain that, indeed, $f(x) \le f(y)$. However, if f is a 'smooth' function, then one can in principle check whether, at least, x is a **local minimizer**, i.e., whether $f(x) \le f(y)$ for all 'nearby' y, by checking whether the **gradient**

$$Df(x) = (D_i f(x) : i = 1{:}n)$$

of f at x is zero. Here, $D_i f = \partial f / \partial x_i$ is the derivative of f with respect to its ith argument.

To be sure, the vanishing of the gradient of f at x is only a *necessary* condition for x to be a minimizer for f, since the gradient of a (smooth) function must also vanish at any local *maximum*, and may vanish at points that are neither local minima nor local maxima but are, perhaps, only saddle points. By definition, any point x for which $Df(x) = 0$ is a **critical point** for f.

At a critical point, f is locally flat. This means that, in the Taylor expansion

$$f(x+h) = f(x) + (Df(x))^{\mathrm{t}} h + h^{\mathrm{t}} (D^2 f(x)/2) h + \text{h.o.t.}(h)$$

for f at x, the linear term, $(Df(x))^{\mathrm{t}} h$, is zero. Thus, if the matrix

$$H := D^2 f(x) = (D_i D_j f(x) : i, j = 1{:}n)$$

of second derivatives of f is 1-1, then x is a local minimizer (maximizer) for f if and only if 0 is a minimizer (maximizer) for the **quadratic form**

$$\mathbb{R}^n \to \mathbb{R} : h \mapsto h^{\mathrm{t}} H h$$

associated with the **Hessian** $H = D^2 f(x)$ for f at x.

If all second derivatives of f are continuous, then also $D_i D_j f = D_j D_i f$, hence the Hessian is real symmetric, therefore

$$H^{\mathrm{t}} = H.$$

However, in the contrary case, one simply defines H to be

$$H := (D^2 f(x) + (D^2 f(x))^{\mathrm{t}})/2,$$

thus making it real symmetric while, still,

$$h^{\mathrm{t}} H h = h^{\mathrm{t}} D^2 f(x) h, \quad \forall h \in \mathbb{R}^n.$$

In any case, it follows that *quadratic forms model the behavior of a smooth function 'near' a critical point.* The importance of minimization of real-valued functions is the prime motivation for the study of quadratic forms, to which we now turn.

Quadratic forms

Each $A \in \mathbb{R}^{n\times n}$ gives rise to a quadratic form, via

$$q_A : \mathbb{R}^n \to \mathbb{R} : x \mapsto x^{\mathrm{t}} Ax.$$

However, as we already observed, the quadratic form 'sees' only the **symmetric part**

$$(A + A^{\mathrm{t}})/2$$

of A, i.e.,

$$\forall x \in \mathbb{R}^n \quad x^{\mathrm{t}}Ax = x^{\mathrm{t}}(A + A^{\mathrm{t}})/2\ x.$$

For this reason, in discussions of the quadratic form q_A, we will always assume that A is real symmetric.

The Taylor expansion for q_A is very simple. One computes

$$q_A(x + h) = (x + h)^{\mathrm{t}}A(x + h) = x^{\mathrm{t}}Ax + x^{\mathrm{t}}Ah + h^{\mathrm{t}}Ax + h^{\mathrm{t}}Ah = q_A(x) + 2(Ax)^{\mathrm{t}}h + h^{\mathrm{t}}Ah,$$

using the fact that $A^{\mathrm{t}} = A$, thus $h^{\mathrm{t}}Ax = x^{\mathrm{t}}Ah = (Ax)^{\mathrm{t}}h$, hence

$$Dq_A(x) = 2Ax, \quad D^2q_A(x) = 2A.$$

It follows that, for any 1-1 A, 0 is the only critical point of q_A. The sought-for classification of critical points of smooth functions has led to the following classification of quadratic forms:

$$A \text{ is } \begin{array}{c}\textbf{positive}\\ \textbf{positive semi-}\\ \textbf{negative semi-}\\ \textbf{negative}\end{array}\textbf{definite} := \begin{array}{ll}\forall x \ne 0 & x^{\mathrm{t}}Ax > 0\\ \forall x & x^{\mathrm{t}}Ax \ge 0\\ \forall x & x^{\mathrm{t}}Ax \le 0\\ \forall x \ne 0 & x^{\mathrm{t}}Ax < 0\end{array} \iff 0 \text{ is } \begin{array}{l}\textbf{the unique minimizer}\\ \textbf{a minimizer}\\ \textbf{a maximizer}\\ \textbf{the unique maximizer}\end{array} \text{ for } q_A.$$

If none of these conditions obtains, i.e., if there exist x and y so that $x^{\mathrm{t}}Ax < 0 < y^{\mathrm{t}}Ay$, then q_A is called **indefinite** and, in this case, 0 is a **saddle point** for q_A.

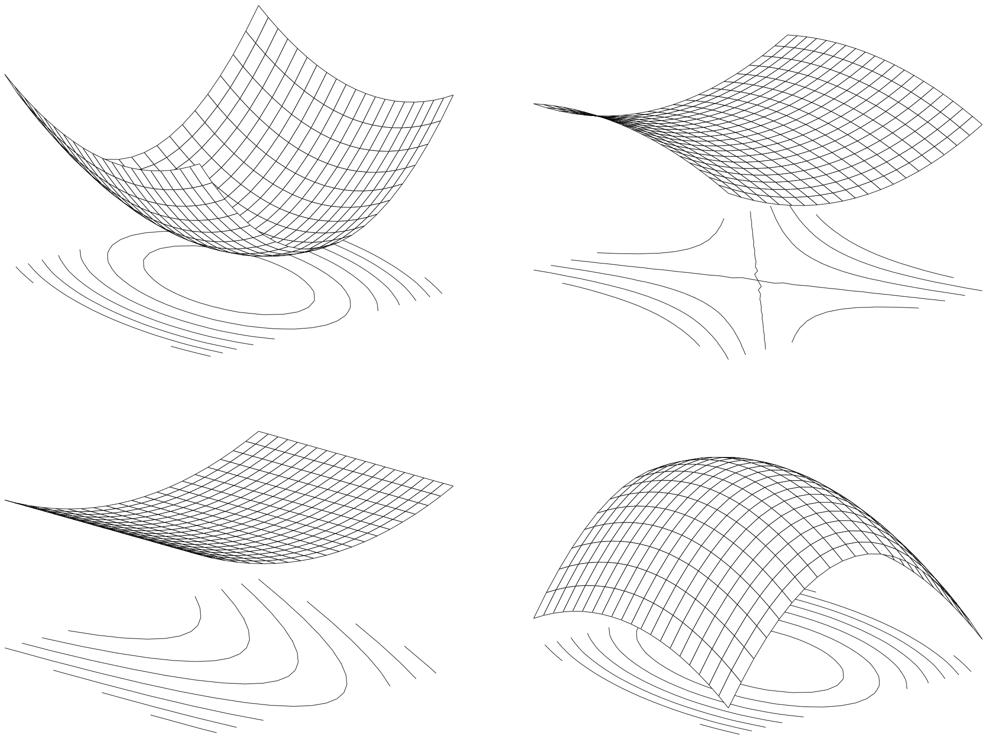

(15.1) Figure. Local behavior near a critical point.

(15.1)Figure shows three quadratic forms near their unique critical point. One is a minimizer, another is a saddle point, and the last one is a maximizer. Also shown is a quadratic form with a whole straight line of critical points. The figure (generated by the `MATLAB` command `meshc`) also shows some **contour lines** or **level lines**, i.e., lines in the domain $\mathbb{R}^2$ along which the function is constant. The contour plots are characteristic: Near an extreme point, be it a maximum or a minimum, the level lines are ellipses, with the extreme point their center, while near a saddle point, the level lines are hyperbolas, with the extreme point their center and with two level lines actually crossing at the saddle point.

There is an intermediate case between these two, also shown in (15.1)Figure, in which the level lines are parabolas and, correspondingly, there is a whole line of critical points. In this case, the quadratic form is semidefinite. Note, however, that the *definition* of semidefiniteness does not exclude the possibility that the quadratic form is actually definite.

Since, near any critical point x, a smooth f behaves like its quadratic term $h \mapsto h^{\mathrm{t}}(D^2 f(x)/2)h$, we can be sure that a contour plot for f near an extremum would approximately look like concentric ellipses while, near a saddle point, it would look approximately like concentric hyperbolas.

These two patterns turn out to be the only two possible ones for definite quadratic forms on $\mathbb{R}^2$. On $\mathbb{R}^n$, there are only $\lceil (n+1)/2 \rceil$ possible distinct patterns, as follows from the fact that, *for every quadratic form q_A, there are o.n. coordinate systems U for which*

$$q_A(x) = \sum_{i=1}^{n} d_i \; (U^{\mathrm{c}} x)_i^2 .$$

15.1 For each of the following three functions on $\mathbb{R}^2$, compute the Hessian $D^2 f(0)$ at 0 and use it to determine whether 0 is a (local) maximum, minimum, or neither. (In an effort to make the derivation of the Hessians simple, I have made the problems so simple that you could tell by inspection what kind of critical point $0 = (0,0) \in \mathbb{R}^2$ is; nevertheless, give your answer based on the spectrum of the Hessian.)

(a) $f(x,y) = (x-y)\sin(x+y)$

(b) $f(x,y) = (x+y)\sin(x+y)$

(b) $f(x,y) = (x+y)\cos(x+y)$.

Reduction of a quadratic form to a sum of squares

Consider the effects of a *change of basis.* Let $V \in \mathbb{R}^n$ be a basis for $\mathbb{R}^n$ and consider the map

$$f := q_A \circ V.$$

We have $f(x) = (Vx)^{\mathrm{t}} A V x = x^{\mathrm{t}}(V^{\mathrm{t}} A V)x$, hence

$$q_A \circ V = q_{V^{\mathrm{t}} A V}.$$

This makes it interesting to look for bases V for which $V^{\mathrm{t}} A V$ is as simple as possible. Matrices A and B for which $B = V^{\mathrm{t}} A V$ are said to be **congruent** to each other. Note that congruent matrices are not necessarily similar; in particular, their spectra can be different. However, by Sylvester's Law of Inertia (see (15.9) below), congruent *hermitian* matrices have the same number of positive, of zero, and of negative, eigenvalues. This is not too surprising in view of the following *reduction to a sum of squares* which is possible for any quadratic form.

(15.2) Proposition: Every quadratic form q_A on $\mathbb{R}^n$ can be written in the form

$$q_A(x) = \sum_{j=1}^{n} d_j ({u_j}^{\mathrm{t}} x)^2,$$

for some suitable o.n. basis $U = [u_1, u_2, \ldots, u_n]$ for which $U^{\mathrm{t}} A U = \mathrm{diag}(d_1, \ldots, d_n) \in \mathbb{R}^n$.

Proof: Since A is hermitian, there exists, by (12.2)Corollary, some o.n. basis $U = [u_1, u_2, \ldots, u_n]$ for $\mathbb{F}^n$ for which $U^{\mathrm{t}}AU = \mathrm{diag}(d_1, d_2, \ldots, d_n) \in \mathbb{R}^{n\times n}$. Now use the facts that $U^{\mathrm{t}}U = \mathrm{id}_n$ and therefore $q_A(x) = q_{U^{\mathrm{t}}AU}(U^{\mathrm{t}}x)$ to obtain for $q_A(x)$ the displayed expression. □

What about the classification introduced earlier, into positive or negative (semidefinite)? The proposition permits us to visualize $q_A(x)$ as a weighted sum of squares (with real weights $d_1, d_2, \ldots, d_n$) and $U^{\mathrm{t}}x$ an arbitrary n-vector (since U is a basis), hence permits us to conclude that q_A is definite if and only if all the d_j are strictly of one sign, semidefinite if and only if all the d_j are of one sign (with zero possible), and indefinite if and only if there are both positive and negative d_j.

`MATLAB` readily provides these numbers d_j by the command `eig(`A`)`.

Consider specifically the case $n = 2$ for which we earlier provided some pictures. Assume without loss that $d_1 \le d_2$. If $0 < d_1$, then A is positive definite and, correspondingly, the contour line

$$c_r := \{x \in \mathbb{R}^2 : q_A(x) = r\} = \{x \in \mathbb{R}^2 : d_1({u_1}^{\mathrm{t}}x)^2 + d_2({u_2}^{\mathrm{t}}x)^2 = r\}$$

for $r > 0$ is an ellipse, with axes parallel to u_1 and u_2. If $0 = d_1 < d_2$, then these ellipses turn into parabolas and, in an extreme case, into straight lines. Similarly, if $d_2 < 0$, then the contour line

$$c_r := \{x \in \mathbb{R}^2 : q_A(x) = r\} = \{x \in \mathbb{R}^2 : d_1({u_1}^{\mathrm{t}}x)^2 + d_2({u_2}^{\mathrm{t}}x)^2 = r\}$$

for $r < 0$ is an ellipse, with axes parallel to u_1 and u_2. Finally, if $d_1 < 0 < d_2$, then, for any r, the contour line

$$c_r := \{x \in \mathbb{R}^2 : q_A(x) = r\} = \{x \in \mathbb{R}^2 : d_1({u_1}^{\mathrm{t}}x)^2 + d_2({u_2}^{\mathrm{t}}x)^2 = r\}$$

is a hyperbola, with axes parallel to u_1 and u_2.

Note that such an o.n. basis U is Cartesian, i.e., its columns are orthogonal to each other (and are normalized). This means that we can visualize the change of basis, from the natural basis to the o.n. basis U, as a rigid motion, involving nothing more than rotations and reflections.

Rayleigh quotient

This section is devoted to the proof and exploitation of the following remarkable

Fact: The eigenvectors of a hermitian matrix A are the critical points of the corresponding **Rayleigh quotient**

$$R_A(x) := \langle Ax, x\rangle / \langle x, x\rangle,$$

and $R_A(x) = \mu$ in case $Ax = \mu x$.

This fact has many important consequences concerning how the eigenvalues of a hermitian matrix depend on that matrix, i.e., how the eigenvalues change when the entries of the matrix are changed, by round-off or for other reasons.

This perhaps surprising connection has the following intuitive explanation: Suppose that $Ax \notin \mathrm{ran}[x]$. Then the error $h := Ax - R_A(x)x$ in the least-squares approximation to Ax from $\mathrm{ran}[x]$ is not zero, and is perpendicular to $\mathrm{ran}[x]$. Consequently, $\langle Ax, h\rangle = \langle h, h\rangle > 0$, and therefore the value

$$\langle A(x + th), x + th\rangle = \langle Ax, x\rangle + 2t\langle Ax, h\rangle + t^2\langle Ah, h\rangle$$

of the numerator of $R_A(x + th)$ grows linearly for positive t, while its denominator

$$\langle x + th, x + th\rangle = \langle x, x\rangle + t^2\langle h, h\rangle$$

grows only quadratically, i.e., much less fast for t near zero. It follows that, in this situation, $R_A(x+th) > R_A(x)$ for all 'small' positive t, hence x cannot be a critical point for R_A. – To put it differently, for any critical point x for R_A, we necessarily have $Ax \in \mathrm{ran}[x]$, therefore $Ax = R_A(x)x$. Of course, that makes any such x an eigenvector with corresponding eigenvalue $R_A(x)$. □

Next, recall from (12.2) that a hermitian matrix is unitarily similar to a real diagonal matrix. This means that we may assume, after some reordering if necessary, that

$$A = UDU^{c}$$

with U unitary and with $\mathrm{M} = \mathrm{diag}(\mu_1, \ldots, \mu_n)$ where

$$\mu_1 \le \mu_2 \le \cdots \le \mu_n.$$

At times, we will write, more explicitly,

$$\mu_j(A)$$

to denote the jth eigenvalue of the hermitian matrix A in this ordering. Note that there may be coincidences here, i.e., $\mu_j(A)$ is the jth smallest eigenvalue of A *counting multiplicities*. Note also that, in contrast to the singular values (and in contrast to most books), we have put here the eigenvalues in *increasing* order.

Now recall that a unitary basis has the advantage that it preserves angles and lengths since $\langle Ux, Uy\rangle = \langle x, y\rangle$ for any orthonormal U. Thus

$$\langle Ax, x\rangle = \langle U\mathrm{M}U^{c}x, x\rangle = \langle \mathrm{M}(U^{c}x), U^{c}x\rangle,$$

and $\langle x, x\rangle = \langle U^{c}x, U^{c}x\rangle$. Therefore

$$R_A(x) = \langle Ax, x\rangle/\langle x, x\rangle = \langle \mathrm{M}(U^{c}x), U^{c}x\rangle/\langle U^{c}x, U^{c}x\rangle = R_{\mathrm{M}}(U^{c}x).$$

This implies that

$$\begin{matrix}\max_x \\ \min_x\end{matrix} R_A(x) = \begin{matrix}\max_y \\ \min_y\end{matrix} R_{\mathrm{M}}(y).$$

On the other hand, since M is diagonal, $\langle \mathrm{M}y, y\rangle = \sum_j \mu_j |y_j|^2$, therefore

$$R_{\mathrm{M}}(y) = \sum_j \mu_j |y_j|^2 / \sum_j |y_j|^2,$$

and this shows that

$$\min_x R_A(x) = \min_y R_{\mathrm{M}}(y) = \mu_1, \qquad \max_x R_A(x) = \max_y R_{\mathrm{M}}(y) = \mu_n.$$

This is **Rayleigh's Principle**. It characterizes the extreme eigenvalues of a hermitian matrix. The intermediate eigenvalues are the solution of more subtle extremum problems. This is the content of the **Courant-Fischer minimax Theorem** and the **?.?. maximin Theorem**. It seems most efficient to combine both in the following

(15.3) MMM (or, **maximinimaxi**) **Theorem:** Let A be a hermitian matrix of order n, hence $A = U\mathrm{M}U^{c}$ for some unitary U and some real diagonal matrix $\mathrm{M} = \mathrm{diag}(\cdots, \mu_j, \ldots)$ with $\mu_1 \le \cdots \le \mu_n$. Then, for $j = 1{:}n$,

$$\max_{\dim G < j} \min_{x \perp G} R_A(x) = \mu_j = \min_{j \le \dim H} \max_{x \in H} R_A(x),$$

with G and H otherwise arbitrary linear subspaces.

Proof: If $\dim G < j \le \dim H$, then one can find $y \in H\backslash 0$ with $y \perp G$ (since, with V a basis for G and W a basis for H, this amounts to finding a nontrivial solution to the equation $V^{c}W? = 0$, and this system is homogeneous with more unknowns than equations). Therefore

$$\min_{x \perp G} R_A(x) \le R_A(y) \le \max_{x \in H} R_A(x).$$

Hence,

$$\max_{\dim G < j} \min_{x \perp G} R_A(x) \le \min_{j \le \dim H} \max_{x \in H} R_A(x).$$

On the other hand, for $G = \mathrm{ran}[u_1, \ldots, u_{j-1}]$ and $H = \mathrm{ran}[u_1, \ldots, u_j]$,

$$\min_{x \perp G} R_A(x) = \mu_j(A) = \max_{x \in H} R_A(x).$$

$\square$

The MMM theorem has various useful (and immediate) corollaries.

(15.4) Interlacing Theorem: If the matrix B is obtained from the hermitian matrix A by crossing out the kth row and column (i.e., $B = A(I, I)$ with $I := (1{:}k-1, k+1{:}n)$), then

$$\mu_j(A) \le \mu_j(B) \le \mu_{j+1}(A), \quad j < n.$$

Proof: It is sufficient to consider the case $k = n$, since we can always achieve this situation by interchanging rows k and n, and columns k and n, of A, and this will not change $\operatorname{spec}(A)$. Let $J : \mathbb{F}^{n-1} \to \mathbb{F}^n : x \mapsto (x, 0)$. Then $R_B(x) = R_A(Jx)$ and $\operatorname{ran} J = \operatorname{ran}[e_n]\perp$, therefore also $J(G\perp) = (JG + \operatorname{ran}[e_n])\perp$ and $\{JG + \operatorname{ran}[e_n] : \dim G < j, G \subset \mathbb{F}^{n-1}\} \subset \{\tilde G : \dim \tilde G < j+1, \tilde G \subset \mathbb{F}^n\}$. Hence

$$\mu_j(B) = \max_{\dim G<j}\min_{x\perp G} R_A(Jx) = \max_{\dim G<j}\min_{y\perp JG+\operatorname{ran}[e_n]} R_A(y) \le \max_{\dim \tilde G<j+1}\min_{y\perp \tilde G} R_A(y) = \mu_{j+1}(A).$$

Also, since $\{JH : j \le \dim H, H \subset \mathbb{F}^{n-1}\} \subset \{\tilde H : j \le \dim \tilde H, \tilde H \subset \mathbb{F}^n\}$,

$$\mu_j(B) = \min_{j\le\dim H}\max_{x\in H} R_A(Jx) = \min_{j\le\dim H}\max_{y\in JH} R_A(y) \ge \min_{j\le\dim \tilde H}\max_{y\in\tilde H} R_A(y) = \mu_j(A).$$

□

(15.5) Corollary: If $A = \begin{bmatrix} B & C \\ D & E \end{bmatrix} \in \mathbb{F}^{n\times n}$ is hermitian, and $B \in \mathbb{F}^{r\times r}$, then at least r eigenvalues of A must be $\le \max\operatorname{spec}(B)$ and at least r eigenvalues of A must be $\ge \min\operatorname{spec}(B)$.

In particular, if the spectrum of B is negative and the spectrum of E is positive, then A has exactly r negative, and $n - r$ positive, eigenvalues.

A different, simpler, application of the MMM theorem is based on the following observation: If

$$f(t) \le g(t) \quad \forall t,$$

then this inequality persists if we take on both sides the maximum or minimum over the same set T, i.e., then

$$\max_{t\in T} f(t) \le \max_{t\in T} g(t), \qquad \min_{t\in T} f(t) \le \min_{t\in T} g(t).$$

It even persists if we further take the minimum or maximum over the same family $\mathbf{T}$ of subsets T, e.g., then also

$$\max_{T\in\mathbf{T}}\min_{t\in T} f(t) \le \max_{T\in\mathbf{T}}\min_{t\in T} g(t).$$

Consequently,

(15.6) Corollary: If A, B are hermitian, and $R_A(x) \le R_B(x) + c$ for some constant c and all x, then

$$\mu_j(A) \le \mu_j(B) + c, \quad \forall j.$$

This gives

(15.7) Weyl's inequalities: If $A = B + C$, with A, B, C hermitian, then

$$\mu_j(B) + \mu_1(C) \le \mu_j(A) \le \mu_j(B) + \mu_n(C), \quad \forall j.$$

Proof: Since $\mu_1(C) \le R_C(x) \le \mu_n(C)$ (by Rayleigh's principle), while $R_B(x) + R_C(x) = R_A(x)$, the preceding corollary provides the proof. □

A typical *application of Weyl's Inequalities* is the observation that, for $A = BB^c + C \in \mathbb{F}^{n\times n}$ with $B \in \mathbb{F}^{n\times k}$ and A hermitian (hence also C hermitian), $\mu_1(C) \le \mu_j(A) \le \mu_n(C)$ for all $j < (n-k)$, since $\operatorname{rank} BB^c \le \operatorname{rank} B \le k$, hence $\mu_j(BB^c)$ must be zero for $j < (n-k)$.

Since $C = A - B$, Weyl's inequalities imply that

$$|\mu_j(A) - \mu_j(B)| \le \max\{|\mu_1(A-B)|, |\mu_n(A-B)|\} = \rho(A-B).$$

Therefore, with the substitutions $A \leftarrow A + E$, $B \leftarrow A$, we obtain

(15.8) max-norm Wielandt-Hoffman: If A and E are both hermitian, then

$$\max_j |\mu_j(A+E) - \mu_j(A)| \le \max_j |\mu_j(E)|.$$

A corresponding statement involving 2-norms is valid but much harder to prove.

Finally, a totally different application of the MMM Theorem is

(15.9) Sylvester's Law of Inertia: Any two *congruent* hermitian matrices have the same number of positive, zero, and negative eigenvalues.

Proof: It is sufficient to prove that if $B = V^c A V$ for some hermitian A and some invertible V, then $\mu_j(A) > 0$ implies $\mu_j(B) > 0$. For this, we observe that, by the MMM Theorem, $\mu_j(A) > 0$ implies that R_A is positive somewhere on every j-dimensional subspace, while (also by the MMM Theorem), for some j-dimensional subspace H,

$$\mu_j(B) \;=\; \max_{x\in H} R_B(x) = \max_{x\in H} R_A(Vx) R_{V^cV}(x),$$

and this is necessarily positive, since $\dim VH = j$ and $R_{V^cV}(x) = \|Vx\|^2/\|x\|^2$ is positive for any $x \ne 0$. □

It follows that we don't have to diagonalize the real symmetric matrix A (as we did in the proof of (15.2)Proposition) in order to find out whether or not A or the corresponding quadratic form q_A is definite. Assuming that A is invertible, hence has no zero eigenvalue, it is sufficient to use Gauss elimination without pivoting to obtain the factorization $A = LDL^c$, with L unit lower triangular. By Sylvester's Law of Inertia, the number of positive (negative) eigenvalues of A equals the number of positive (negative) diagonal entries of D.

This fact can be used to locate the eigenvalues of a real symmetric matrix by *bisection*. For, the number of positive (negative) diagonal entries in the diagonal matrix D_μ obtained in the factorization $L_\mu D_\mu L_\mu{}^c$ for $(A - \mu\,\mathrm{id})$ tells us the number of eigenvalues of A to the right (left) of μ, hence makes it easy to locate and refine intervals that contain just one eigenvalue of A.

16. More on determinants

In this chapter only, n-vectors will be denoted by lower-case boldface roman letters; for example,

$$\mathbf{a} = (a_1, \ldots, a_n) \in \mathbb{F}^n.$$

Determinants are often brought into courses such as this quite unnecessarily. But when they are useful, they are remarkably so. The use of determinants is a bit bewildering to the beginner, particularly if confronted with the classical definition as a sum of signed products of matrix entries.

I find it more intuitive to follow Weierstrass and begin with a few important properties of the determinant, from which all else follows, including that classical definition (which is practically useless anyway).

As to the many determinant identities available, in the end I have almost always managed with just one nontrivial one, viz. *Sylvester's determinant identity*, and this is nothing but Gauss elimination; see the end of this chapter. The only other one I have used at times is the *Cauchy-Binet formula*.

Definition and basic properties

The determinant is a map,

$$\det : \mathbb{F}^{n\times n} \to \mathbb{F} : A \mapsto \det A,$$

with various properties. The first one in the following list is perhaps the most important one.

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

(ii) $\det(\,\mathrm{id}) = 1$.

Consequently, for any invertible A,

$$1 = \det(\,\mathrm{id}) = \det(AA^{-1}) = \det(A)\det(A^{-1}).$$

Hence,

(iii) *If A is invertible, then* $\det A \neq 0$ *and,* $\det(A^{-1}) = 1/\det(A)$.

While the determinant is defined as a map on matrices, it is very useful to think of $\det(A) = \det[\mathbf{a}_1, \ldots, \mathbf{a}_n]$ as a function of the columns $\mathbf{a}_1, \ldots, \mathbf{a}_n$ of A. The next two properties are in those terms:

(iv) $\mathbf{x} \mapsto \det[\ldots, \mathbf{a}_{j-1}, \mathbf{x}, \mathbf{a}_{j+1}, \ldots]$ *is linear*, i.e., for any n-vectors $\mathbf{x}$ and $\mathbf{y}$ and any scalar α (and arbitrary n-vectors $\mathbf{a}_i$),

$$\det[\ldots, \mathbf{a}_{j-1}, \mathbf{x} + \alpha\mathbf{y}, \mathbf{a}_{j+1}, \ldots] = \det[\ldots, \mathbf{a}_{j-1}, \mathbf{x}, \mathbf{a}_{j+1}, \ldots] + \alpha\det[\ldots, \mathbf{a}_{j-1}, \mathbf{y}, \mathbf{a}_{j+1}, \ldots].$$

(v) *The determinant is an* **alternating** *form*, i.e.,

$$\det[\ldots, \mathbf{a}_i, \ldots, \mathbf{a}_j, \ldots] = -\det[\ldots, \mathbf{a}_j, \ldots, \mathbf{a}_i, \ldots].$$

In words: Interchanging two columns changes the sign of the determinant (and nothing else).

It can be shown (see below) that (ii) + (iv) + (v) implies (i) (and anything else you may wish to prove about determinants). Here are some basic consequences first.

(vi) Since 0 is the only scalar α with the property that $\alpha = -\alpha$, it follows from (v) that $\det(A) = 0$ *if two columns of A are the same.*

(vii) *Adding a multiple of one column of A to another column of A doesn't change the determinant.*

Indeed, using first (iv) and then the consequence (vi) of (v), we compute

$$\det[\ldots,\mathbf{a}_i,\ldots,\mathbf{a}_j+\alpha\mathbf{a}_i,\ldots] = \det[\ldots,\mathbf{a}_i,\ldots,\mathbf{a}_j,\ldots] + \alpha\det[\ldots,\mathbf{a}_i,\ldots,\mathbf{a}_i,\ldots] = \det[\ldots,\mathbf{a}_i,\ldots,\mathbf{a}_j,\ldots].$$

Here comes a very important use of (vii): Assume that $\mathbf{b} = A\mathbf{x}$ and consider $\det[\ldots,\mathbf{a}_{j-1},\mathbf{b},\mathbf{a}_{j+1},\ldots]$. Since $\mathbf{b} = x_1\mathbf{a}_1 + \cdots + x_n\mathbf{a}_n$, subtraction of x_i times column i from column j, i.e., subtraction of $x_i\mathbf{a}_i$ from $\mathbf{b}$ here, for each $i \ne j$ is, by (vii), guaranteed not to change the determinant, yet changes the jth column to $x_j\mathbf{a}_j$; then, pulling out that scalar factor x_j (permitted by (iv)), leaves us finally with $x_j \det A$. This proves

(viii) *If* $\mathbf{b} = A\mathbf{x}$, *then* $\det[\ldots,\mathbf{a}_{j-1},\mathbf{b},\mathbf{a}_{j+1},\ldots] = x_j \det A$.

Hence, if $\det A \ne 0$, then $\mathbf{b} = A\mathbf{x}$ implies

$$x_j = \det[\ldots,\mathbf{a}_{j-1},\mathbf{b},\mathbf{a}_{j+1},\ldots]/\det(A), \qquad j = 1,\ldots,n.$$

This is **Cramer's rule**.

In particular, if $\det(A) \ne 0$, then $A\mathbf{x} = \mathbf{0}$ implies that $x_j = 0$ for all j, i.e., then A is 1-1, hence invertible (since A is square). This gives the converse to (iii), i.e.,

(ix) *If* $\det(A) \ne 0$, *then* A *is invertible.*

In old-fashioned mathematics, a matrix was called **singular** if its determinant is 0. So, (iii) and (ix) combined say that a matrix is nonsingular iff it is invertible.

The suggestion that one actually construct the solution to $A? = y$ by Cramer's rule is ridiculous under ordinary circumstances since, even for a linear system with just two unknowns, it is more efficient to use Gauss elimination. On the other hand, if the solution is to be constructed *symbolically* (in a symbol-manipulating system such as `Maple` or `Mathematica`), then Cramer's rule is preferred to Gauss elimination since it treats all unknowns equally. In particular, the number of operations needed to obtain a particular unknown is the same for all unknowns.

We have proved all these facts (except (i)) about determinants from certain postulates (namely (ii), (iv), (v)) without ever saying how to *compute* $\det(A)$. Now, it is the actual formulas for $\det(A)$ that have given determinants such a bad name. Here is the standard one, which (see below) can be derived from (ii), (iv), (v), in the process of proving (i):

(x) *If* $A = (a_{ij} : i,j = 1,\ldots,n)$, *then*

$$\det(A) = \sum_{\sigma\in\mathbb{S}_n} (-1)^\sigma \prod_{j=1}^n a_{\sigma(j),j}.$$

Here, $\sigma \in \mathbb{S}_n$ is shorthand for: σ is a **permutation of the first** n **integers**, i.e.,

$$\sigma = (\sigma(1),\sigma(2),\ldots,\sigma(n)),$$

where $\sigma(j) \in \{1,2,\ldots,n\}$ for all j, and $\sigma(i) \ne \sigma(j)$ if $i \ne j$. In other words, σ is a 1-1 and onto map from $\{1,\ldots,n\}$ to $\{1,\ldots,n\}$. This is bad enough, but I still have to explain the mysterious $(-1)^\sigma$. This number is 1 or -1 depending on whether the **parity** of σ is even or odd. Now, this parity can be determined in at least two equivalent ways:

(a) keep making interchanges until you end up with the sequence $(1,2,\ldots,n)$; the parity of the number of steps it took is the parity of σ (note the implied assertion that this parity will not depend on how you went about this, i.e., the number of steps taken may differ, but the parity never will; if it takes me an even number of steps, it will take you an even number of steps.)

(b) count the number of pairs that are out of order; its parity is the parity of σ.

Here is a simple example: $\sigma = (3,1,4,2)$ has the pairs $(3,1)$, $(3,2)$, and $(4,2)$ out of order, hence $(-1)^\sigma = -1$. Equivalently, the following sequence of 3 interchanges gets me from σ to $(1,2,3,4)$:

$$\begin{gathered}(3,1,4,2)\\(3,1,2,4)\\(1,3,2,4)\\(1,2,3,4)\end{gathered}$$

Therefore, again, $(-1)^\sigma = -1$.

Now, fortunately, we don't really ever have to use this stunning formula (x) in calculations, nor is it physically possible to use it for n much larger than 8 or 10. For $n = 1,2,3$, one can derive from it explicit rules for computing $\det(A)$:

$$\det\,[\,a\,] = a, \quad \det\begin{bmatrix} a & b \\ c & d \end{bmatrix} = ad - bc, \quad \det\begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} = aei + bfg + cdh - (ceg + afh + bdi);$$

the last one can be remembered easily by the following mnemonic:

$$\begin{matrix} a & b & c & a & b \\ d & e & f & d & e \\ g & h & i & g & h \end{matrix}$$

For $n > 3$, this mnemonic *does not work*, and one would not usually make use of (x), but use instead (i) and the following immediate consequence of (x):

(xi) *The determinant of a triangular matrix equals the product of its diagonal entries.*

Indeed, when A is upper triangular, then $a_{ij} = 0$ whenever $i > j$. Now, if $\sigma(j) > j$ for some j, then the factor $a_{\sigma(j),j}$ in the corresponding summand $(-1)^\sigma \prod_{j=1}^n a_{\sigma(j),j}$ is zero. This means that the only possibly nonzero summands correspond to σ with $\sigma(j) \le j$ for all j, and there is only one permutation that manages that, the **identity permutation** $(1,2,\ldots,n)$, and its parity is even (since it takes no interchanges). Therefore, the formula in (x) gives $\det A = a_{11}\cdots a_{nn}$ in this case. – The proof for a lower triangular matrix is analogous; else, use (xiii) below.

Consequently, if $A = LU$ with L unit triangular and U upper triangular, then

$$\det A = \det U = u_{11}\cdots u_{nn}.$$

If, more generally, $A = PLU$, with P some permutation matrix, then

$$\det A = \det(P) u_{11}\cdots u_{nn},$$

i.e.,

(xii) $\det A$ *is the product of the pivots used in elimination, times* $(-1)^i$, *with* i *the number of row interchanges made.*

Since, by elimination, any $A \in \mathbb{F}^n$ can be factored as $A = PLU$, with P a permutation matrix, L unit lower triangular, and U upper triangular, (xii) provides the standard way to compute determinants.

Note that, then, $A^{\mathrm{t}} = U^{\mathrm{t}} L^{\mathrm{t}} P^{\mathrm{t}}$, with U^{t} lower triangular, L^{t} unit upper triangular, and P^{t} the inverse of P, hence

(xiii) $\det A^{\mathrm{t}} = \det A$.

This can also be proved directly from (x). Note that this converts all our statements about the determinant in terms of *columns* to the corresponding statements in terms of *rows*.

(xiv) "expansion by minors":

Since, by (iv), the determinant is slotwise linear, and $\mathbf{x} = x_1\mathbf{e}_1 + x_2\mathbf{e}_2 + \cdots + x_n\mathbf{e}_n$, we obtain

$$\det[\ldots, \mathbf{a}_{j-1}, \mathbf{x}, \mathbf{a}_{j+1}, \ldots] = x_1 C_{1j} + x_2 C_{2j} + \cdots + x_n C_{nj}, \tag{16.1}$$

with

$$C_{ij} := \det[\ldots, \mathbf{a}_{j-1}, \mathbf{e}_i, \mathbf{a}_{j+1}, \ldots]$$

the socalled **cofactor** of a_{ij}. With the choice $\mathbf{x} = \mathbf{a}_k$, this implies

$$a_{1k}C_{1k} + a_{2k}C_{2k} + \cdots + a_{nk}C_{nk} = \det[\ldots, \mathbf{a}_{j-1}, \mathbf{a}_k, \mathbf{a}_{j+1}, \ldots] = \begin{cases} \det A & \text{if } k = j; \\ 0 & \text{otherwise.} \end{cases}$$

The case $k = j$ gives the **expansion by minors** for $\det A$ (and justifies the name 'cofactor' for C_{ij}). The case $k \neq j$ is justified by (vi). In other words, with

$$\operatorname{adj} A := \begin{bmatrix} C_{11} & C_{21} & \cdots & C_{n1} \\ C_{12} & C_{22} & \cdots & C_{n2} \\ \vdots & \vdots & \cdots & \vdots \\ C_{1n} & C_{2n} & \cdots & C_{nn} \end{bmatrix}$$

the socalled **adjugate** of A (note that the subscripts appear reversed), we have

$$\operatorname{adj}(A)\, A = (\det A)\ \mathrm{id}.$$

For an *invertible* A, this implies that

$$A^{-1} = (\operatorname{adj} A)/\det A.$$

The expansion by minors is useful since, as follows from (x), the cofactor C_{ij} equals $(-1)^{i+j}$ times the determinant of the matrix $A(\mathbf{n}\backslash i|\mathbf{n}\backslash j)$ obtained from A by removing row i and column j, i.e.,

$$C_{ij} = (-1)^{i+j} \det \begin{bmatrix} \cdots & \cdots & \cdots & \cdots \\ \cdots & a_{i-1,j-1} & a_{i-1,j+1} & \cdots \\ \cdots & a_{i+1,j-1} & a_{i+1,j+1} & \cdots \\ \cdots & \cdots & \cdots & \cdots \end{bmatrix},$$

and this is a determinant of order $n-1$, and so, if $n-1 > 1$, can itself be expanded along some column (or row).

As a practical matter, for $[\mathbf{a}, \mathbf{b}, \mathbf{c}] := A \in \mathbb{R}^3$, the formula $\operatorname{adj}(A)\, A = (\det A)$ id implies that

$$(\mathbf{a} \times \mathbf{b})^{\mathrm{t}}\mathbf{c} = \det[\mathbf{a}, \mathbf{b}, \mathbf{c}],$$

with

$$\mathbf{a} \times \mathbf{b} := (a_2b_3 - a_3b_2, a_3b_1 - a_1b_3, a_1b_2 - a_2b_1)$$

the **cross product** of $\mathbf{a}$ with $\mathbf{b}$. In particular, $\mathbf{a} \times \mathbf{b}$ is perpendicular to both $\mathbf{a}$ and $\mathbf{b}$. Also, if $[\mathbf{a}, \mathbf{b}]$ is o.n., then so is $[\mathbf{a}, \mathbf{b}, \mathbf{a} \times \mathbf{b}]$ but, in addition, $\det[\mathbf{a}, \mathbf{b}, \mathbf{a} \times \mathbf{b}] = 1$, i.e., $[\mathbf{a}, \mathbf{b}, \mathbf{a} \times \mathbf{b}]$ provides a right-handed cartesian coordinate system for $\mathbb{R}^3$.

(xv) $\det A$ *is the n-dimensional (signed) volume of the parallelepiped*

$$\{A\mathbf{x} : 0 \le x_i \le 1, \text{all } i\}$$

spanned by the columns of A.

For $n > 3$, this is a *definition*, while, for $n \le 3$, one works it out (see below). This is a very useful *geometric* way of thinking about determinants. Also, it has made determinants indispensable in the *definition* of multivariate integration and the handling therein of changes of variable.

Since $\det(AB) = \det(A)\det(B)$, it follows that *the linear transformation* $T : \mathbb{F}^n \to \mathbb{F}^n : \mathbf{x} \mapsto A\mathbf{x}$ *changes volumes by a factor of* $\det(A)$, meaning that, for any set M in the domain of T,

$$\mathrm{vol}_n(T(M)) = \det(A)\,\mathrm{vol}_n(M).$$

As an example, consider $\det[\mathbf{a}, \mathbf{b}]$, with $\mathbf{a}$, $\mathbf{b}$ vectors in the plane linearly independent, and assume, wlog, that $a_1 \ne 0$. By (iv), $\det[\mathbf{a}, \mathbf{b}] = \det[\mathbf{a}, \tilde{\mathbf{b}}]$, with $\tilde{\mathbf{b}} := \mathbf{b} - (b_1/a_1)\mathbf{a}$ having its first component equal to zero, and so, again by (iv), $\det[\mathbf{a}, \mathbf{b}] = \det[\tilde{\mathbf{a}}, \tilde{\mathbf{b}}]$, with $\tilde{\mathbf{a}} := \mathbf{a} - (a_2/\tilde{b}_2)\tilde{\mathbf{b}}$ having its second component equal to zero. Therefore, $\det[\mathbf{a}, \mathbf{b}] = \tilde{a}_1\tilde{b}_2 = \pm\|\tilde{\mathbf{a}}\|\|\tilde{\mathbf{b}}\|$ equals $\pm$ the area of the rectangle spanned by $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{b}}$. However, following the derivation of $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{b}}$ graphically, we see, by matching congruent triangles, that the rectangle spanned by $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{b}}$ has the same area as the parallelepiped spanned by $\mathbf{a}$ and $\tilde{\mathbf{b}}$, and, therefore, as the parallelepiped spanned by $\mathbf{a}$ and $\mathbf{b}$. Thus, up to sign, $\det[\mathbf{a}, \mathbf{b}]$ is the area of the parallelepiped spanned by $\mathbf{a}$ and $\mathbf{b}$.

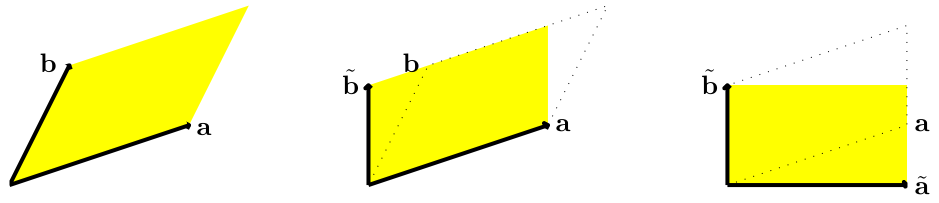

Here, finally, for the record, is a *proof* that (ii) + (iv) + (v) implies (i), hence everything else we have been deriving so far. Let A and B be arbitrary matrices (of order n). Then the linearity (iv) implies that

$$\det(BA) = \det[B\mathbf{a}_1, B\mathbf{a}_2, \ldots, B\mathbf{a}_n] = \det[\ldots, \sum_i \mathbf{b}_i a_{ij}, \ldots] = \sum_{\sigma\in\{1,\ldots,n\}^n} \det[\mathbf{b}_{\sigma(1)}, \ldots, \mathbf{b}_{\sigma(n)}] \prod_j a_{\sigma(j),j}.$$

By the consequence (vi) of the alternation property (v), most of these summands are zero. Only those determinants $\det[\mathbf{b}_{\sigma(1)}, \ldots, \mathbf{b}_{\sigma(n)}]$ for which all the entries of σ are different are *not* automatically zero. But that are exactly all the $\sigma \in \mathbb{S}_\mathbf{n}$, i.e., the permutations of the first n integers. Further, for such σ,

$$\det[\mathbf{b}_{\sigma(1)}, \ldots, \mathbf{b}_{\sigma(n)}] = (-)^\sigma \det(B)$$

by the alternation property (v), with $(-)^\sigma = \pm 1$ depending on whether it takes an even or an odd number of interchanges to change σ into a strictly increasing sequence. (We discussed this earlier; the only tricky part remaining here is an argument that shows the parity of such number of needed interchanges to be independent of how one goes about making the interchanges. The clue to the proof is the simple observation that any one interchange is bound to change the number of sequence entries out of order by an *odd* amount.) Thus

$$\det(BA) = \det(B) \sum_{\sigma\in\mathbb{S}_\mathbf{n}} (-)^\sigma \prod_j a_{\sigma(j),j}.$$

Since $\mathrm{id}A = A$ while, by the defining property (ii), $\det(\mathrm{id}) = 1$, the formula (x) follows and, with that, $\det(BA) = \det(B)\det(A)$ for arbitrary B and A. On the other hand, starting with the formula in (x) as a definition, one readily verifies that det so defined satisfies the three properties (ii) ($\det(\mathrm{id}) = 1$), (iv) (multilinear), and (v) (alternating) claimed for it. In other words, there actually is such a function (necessarily given by (x)).

Sylvester

Here, for the record, is a proof and statement of Sylvester's Determinant Identity. For it, the following notation will be useful: If $\mathbf{i} = (i_1, \ldots, i_r)$ and $\mathbf{j} = (j_1, \ldots, j_s)$ are suitable integer sequences, then $A(\mathbf{i}, \mathbf{j}) = A(\mathbf{i}|\mathbf{j})$ is the $r \times s$-matrix whose (p, q) entry is $A(i_p, j_q)$, $p = 1, \ldots, r$, $q = 1, \ldots, s$. This is just as in `MATLAB` except for the vertical bar used here at times, for emphasis and in order to list, on either side of it, a sequence without having to encase it in parentheses. Also, it will be handy to denote by : the entire sequence $1{:}n$, and by $\backslash\mathbf{i}$ the sequence obtained from $1{:}n$ by removing from it the entries of $\mathbf{i}$. Thus, as in `MATLAB`, $A(:|j) = A(:, j)$ is the jth column of A. Finally, $A(\mathbf{i}) := A(\mathbf{i}|\mathbf{i})$.

With $\mathbf{k} := 1{:}k$, consider the matrix B with entries

$$B(i, j) := \det A(\mathbf{k}, i|\mathbf{k}, j).$$

On expanding (see property (xiv)) $\det A(\mathbf{k}, i|\mathbf{k}, j)$ by entries of the last row,

$$B(i, j) = A(i, j) \det A(\mathbf{k}) - \sum_{r \le k} A(i, r)(-)^{k-r} \det A(\mathbf{k}|(\mathbf{k}\backslash r), j).$$

This shows that

$$B(:, j) \in A(:, j) \det A(\mathbf{k}) + \operatorname{span} A(:|\mathbf{k}),$$

while, directly, $B(i, j) = 0$ for $i \in \mathbf{k}$ since then $\det A(\mathbf{k}, i|\mathbf{k}, j)$ has two rows the same.

In the same way,

$$B(i, :) \in A(i, :) \det A(\mathbf{k}) + \operatorname{span} A(\mathbf{k}|:),$$

while, directly, $B(i, j) = 0$ for $j \in \mathbf{k}$. Thus, if $\det A(\mathbf{k}) \neq 0$, then, for $i > k$,

$$B(i, :)/\det A(\mathbf{k})$$

provides the ith row of the matrix obtained from A after k steps of Gauss elimination (without pivoting). This provides the following useful

(16.2) Determinantal expressions for LU factors and Schur complement: The matrix $S := B/\det A(\mathbf{k})$ contains the **Schur complement** $S(\backslash\mathbf{k})$ in A of the **pivot block** $A(\mathbf{k})$. Further,

$$B(k+1, k+1) = \det A(\mathbf{k+1})/\det A(\mathbf{k})$$

is the pivot for the $k+1$st elimination step, hence, for $i > k$,

$$L(i, k+1) = B(i, k+1)/B(k+1, k+1) = \det A(\mathbf{k}, i|\mathbf{k+1})/\det A(\mathbf{k+1})$$

is the $(i, k{+}1)$ entry of the resulting unit lower triangular left factor of A and, correspondingly,

$$U(k+1, i) = B(k+1, i)/B(k+1, k+1) = \det A(\mathbf{k+1}|\mathbf{k}, i)/\det A(\mathbf{k+1})$$

is the $(k+1, i)$ entry of the resulting unit upper triangular right factor of A.

Since such row elimination is done by elementary matrices with determinant equal to 1, it follows that

$$\det A = \det A(\mathbf{k}) \det S(\backslash\mathbf{k}).$$

Since, for any $\#\mathbf{i} = \#\mathbf{j}$, $B(\mathbf{i}, \mathbf{j})$ depends only on the square matrix $A(\mathbf{k}, \mathbf{i}|\mathbf{k}, \mathbf{j})$, this implies

Sylvester's determinant identity. *If*

$$S(i,j) := \det A(\mathbf{k},i|\mathbf{k},j)/\det A(\mathbf{k}), \quad \forall i,j,$$

then

$$\det S(\mathbf{i}|\mathbf{j}) = \det A(\mathbf{k},\mathbf{i}|\mathbf{k},\mathbf{j})/\det A(\mathbf{k}).$$

Cauchy-Binet

Cauchy-Binet formula. $\det(BA)(\mathbf{i}|\mathbf{j}) = \sum_{\#\mathbf{h}=\#\mathbf{i}} \det B(\mathbf{i}|\mathbf{h}) \ \det A(\mathbf{h}|\mathbf{j})$.

Even the special case $\#\mathbf{i} = \#A$ of this, i.e., the most important determinant property (i),

$$\det(BA) = \det B \det A,$$

Binet and Cauchy were the first to prove. Not surprisingly, the proof of the formula follows our earlier proof of that identity.

Proof: Since $(BA)(\mathbf{i}|\mathbf{j}) = B(\mathbf{i}|:)A(:|\mathbf{j})$, it is sufficient to consider the case $B, A^{\mathrm{t}} \in \mathbb{F}^{m\times n}$ for some m and n. If $m > n$, then B cannot be onto, hence BA must fail to be invertible, while the sum is empty, hence has value 0. It is therefore sufficient to consider the case $m \le n$.

For this, using the linearity of the determinant in each slot,

$$\begin{aligned}
\det BA &= \det[BA(:,1),\ldots,BA(:,m)] \\
&= \sum_{h_1}\cdots\sum_{h_m} \det[B(:,h_1)A(h_1,1),\ldots,B(:,h_m)A(h_m,m)] \\
&= \sum_{h_1}\cdots\sum_{h_m} \det[B(:,h_1),\ldots,B(:,h_m)]\, A(h_1,1)\cdots A(h_m,m) \\
&= \sum_{h_1<\ldots<h_m} \det B(:|\mathbf{h}) \sum_{\sigma\in\mathbb{S}_m} (-1)^\sigma A(h_{\sigma(1)},1)\cdots A(h_{\sigma(m)},m) = \sum_{h_1<\ldots<h_m} \det B(:|\mathbf{h}) \det A(\mathbf{h}|:).
\end{aligned}$$

□

16.1 Prove: *For any $A \in \mathbb{F}^{n\times n+1}$, the vector $((-1)^k \det A(:,\backslash k) : k = 1{:}n+1)$ is in* $\operatorname{null} A$.

16.2 Let $A \in \mathbb{Z}^{n\times n}$, i.e., a matrix of order n with integer entries, and assume that A is invertible. Prove: $A^{-1} \in \mathbb{Z}^{n\times n}$ *if and only if* $|\det A| = 1$. (Hint: Use Cramer's Rule to prove that such A maps $\mathbb{Z}^n$ onto itself in case $\det A = \pm 1$.)

16.3 Let R be a ring (see Backgrounder). Prove the following claim, of use in ideal theory: *If $Ax = 0$ for $A \in R^{n\times n}$ and $x \in R^n$, then $x_i \det(A) = 0$ for all i.*

16.4 Use the previous homework to prove the following (see the Backgrounder on rings for background): *If R is a commutative ring with identity, $s_1,\ldots,s_n \in R$, $F := [s_1,\ldots,s_n](R^n)$ and H is an ideal in R for which $F \subset HF := \{hf : h \in H, f \in F\}$, then, for some $h \in H$, $(1-h)F = 0$.*

16.5 Prove that the elementary matrix $A := \mathrm{id}_n - qr^{\mathrm{t}}$ has a factorization $A = LDU$ with L unit lower triangular, D diagonal, and U unit upper triangular provided the numbers

$$p_i := 1 - \sum_{j\le i} q_j r_j$$

are non zero for $i < n$, and verify that then $D = \operatorname{diag}(p_i/p_{i-1} : i = 1{:}n)$ and

$$L(i,j) = -q_i r_j / p_j = U(j,i), \quad i > j.$$

17. Background

A nonempty finite subset of $\mathbb{R}$ contains a maximal element

Let m be an arbitrary element of the set M in question; there is at least one, by assumption. Then the algorithm

$$\textbf{for } r \in M \textbf{ do: if } r > m\textbf{, } m \leftarrow r\textbf{, od}$$

produces the maximal element, m, after finitely many steps.

Since a bounded subset of $\mathbb{Z}$ necessarily has only finitely many elements, it follows that *a nonempty bounded subset of $\mathbb{Z}$ contains a maximal element.* This latter claim is used several times in these notes.

Also, note the corollary that *a bounded function into the integers takes on its maximal value*: its range then contains a maximal element and any preimage of that maximal element will do.

A nonempty bounded subset of $\mathbb{R}$ has a least upper bound

Let M be a subset of $\mathbb{R}$. Then, as the example of the open interval $(0\,..\,1)$ readily shows, such M need not have a maximal (or, rightmost) element. However, if the set

$$\{r \in \mathbb{R} : m \le r, \forall m \in M\}$$

of **upper bound**s for M is not empty, then this set has a smallest (or, leftmost) element. This smallest element is called the **least upper bound**, or the **supremum**, for M and is correspondingly denoted

$$\sup M.$$

The existence of a least upper bound for any real set M that has an upper bound is part of our understanding or definition of the set $\mathbb{R}$. What if M has no upper bound? Then some would say that $\sup M = \infty$. What if M is empty? Then, offhand, $\sup M$ is not defined. On the other hand, since $M \subset N \Longrightarrow \sup M \le \sup N$, some would, consistent with this, define $\sup\{\} = \inf \mathbb{R} = -\infty$.

One also considers the set

$$\{r \in \mathbb{R} : r \le m, \forall m \in M\}$$

of all **lower bound**s of the set M and understands that this set, if nonempty, has a largest (or, right-most) element. This element is called the **greatest lower bound**, or **infimum**, of M, and is denoted

$$\inf M.$$

What if M has no lower bound? Then some would say that $\inf M = -\infty$.

Note that

$$-\sup M = \inf(-M).$$

Complex numbers

A complex number is of the form

$$z = a + \mathrm{i}b,$$

with a and b real numbers, called, respectively, the **real part of** z and the **imaginary part of** z, and i the **imaginary unit**, i.e.,

$$\mathrm{i} := \sqrt{-1}.$$

Actually, there are two complex numbers whose square is -1. We denote the other one by $-\mathrm{i}$. Be aware that, in parts of Engineering, the symbol j is used instead of i.

> `MATLAB` works internally with (double precision) complex numbers. Both variables `i` and `j` in `MATLAB` are initialized to the value i. □

One adds complex numbers by adding separately their real and imaginary parts. One multiplies two complex numbers by multiplying out and rearranging, mindful of the fact that $\mathrm{i}^2 = -1$. Thus,

$$(a + \mathrm{i}b)(c + \mathrm{i}d) = ac + a\mathrm{i}d + b\mathrm{i}c - bd = (ac - bd) + \mathrm{i}(ad + bc).$$

Note that both addition and multiplication of complex numbers is commutative. Further, the product of $z = a + \mathrm{i}b$ with its **complex conjugate**

$$\overline{z} := a - \mathrm{i}b$$

is the nonnegative number

$$z\overline{z} = a^2 + b^2,$$

and its (nonnegative) squareroot is called the **absolute value** or **modulus** of z and denoted by

$$|z| := \sqrt{z\overline{z}}.$$

For $z \neq 0$, we have $|z| \neq 0$, hence $\overline{z}/|z|^2 = a/|z|^2 - \mathrm{i}b/|z|^2$ is a well-defined complex number. It is the **reciprocal** of z since $z\overline{z}/|z|^2 = 1$, of use for *division* by z. Note that, for any two complex numbers z and ζ,

$$|z\zeta| = |z||\zeta|.$$

It is very useful to visualize complex numbers as points in the so called **complex plane**, i.e., to identify the complex number $a + \mathrm{i}b$ with the point (a, b) in $\mathbb{R}^2$. With this identification, its absolute value corresponds to the (Euclidean) distance of the corresponding point from the origin. The sum of two complex numbers corresponds to the vector sum of their corresponding points. The product of two complex numbers is most easily visualized in terms of the **polar form**

$$z = a + \mathrm{i}b = r \exp(\mathrm{i}\varphi),$$

with $r \geq 0$, hence $r = |z|$ its *modulus*, and $\varphi \in \mathbb{R}$ is called its **argument**. Indeed, for any real φ, $\exp(\mathrm{i}\varphi) = \cos(\varphi) + \mathrm{i}\sin(\varphi)$ has absolute value 1, and φ is the angle (in radians) that the vector (a, b) makes with the positive real axis. Note that, for $z \neq 0$, the argument, φ, is only defined up to a multiple of 2π, while, for $z = 0$, the argument is arbitrary. If now also $\zeta = \alpha + \mathrm{i}\beta = |\zeta| \exp(\mathrm{i}\psi)$, then, by the law of exponents,

$$z\zeta = |z| \exp(\mathrm{i}\varphi)|\zeta| \exp(\mathrm{i}\psi) = |z||\zeta| \exp(\mathrm{i}(\varphi + \psi)).$$

Thus, as already noted, the absolute value of the product is the product of the absolute values of the factors, while the argument of a product is the sum of the arguments of the factors.

For example, in as much as the argument of $\overline{z}$ is the negative of the argument of z, the argument of the product $z\overline{z}$ is necessarily 0. As another example, if $z = a + \mathrm{i}b$ is of modulus 1, then z lies on the unit circle in the complex plane, and so does any power z^k of z. In fact, then $z = \exp(\mathrm{i}\varphi)$ for some real number φ, and therefore $z^k = \exp(\mathrm{i}(k\varphi))$. Hence, the sequence $z^0, z^1, z^2, \cdots$ appears as a sequence of points on the unit circle, equally spaced around that circle, never accumulating anywhere unless $\varphi = 0$, i.e., unless $z = 1$.

(17.1) Lemma: Let z be a complex number of modulus 1. Then the sequence $z^0, z^1, z^2, \ldots$ of powers of z lies on the unit circle, but fails to converge except when $z = 1$.

Groups, rings, and fields

A **semigroup** (F, op) is a set F and an **operation** op on F, i.e., a map $op : F \times F \to F : (f, g) \mapsto fg$ that is **associative**, meaning that

$$\forall\{f, g, h \in F\}\ (fg)h = f(gh).$$

The semigroup is **commutative** if

$$\forall\{f, g \in F\}\ fg = gf.$$

The prime example of a semigroup is the set $(M \to M) = M^M$ of all maps on some set M, with map composition as the operation, or any of its subsets H that are **closed under** the operation, i.e., satisfy $HH := \{gh : g, h \in H\} \subset H$. M^M is commutative only if $\#M = 1$.

A **group** (G, op) is a semigroup (necessarily nonempty) whose operation is a **group operation**, meaning that, in addition to associativity, it has the following properties:

(g.1) there exists a **left neutral element** and a **right neutral element**, i.e., an $e_l, e_r \in G$ (necessarily $e_l = e_r$, hence unique, denoted by e and called the **neutral element**) such that

$$\forall\{g \in G\}\ e_l g = g = g e_r;$$

(g.2) every $g \in G$ has a **left inverse** and a **right inverse**, i.e., $f, h \in G$ so that $fg = e = gh$ (and, necessarily, these are unique and coincide, leading to the notation $f = g^{-1} = h$).

G is said to be 'a group under the operation op'.

A group G is called **Abelian** if it is commutative.

If also H is a group, then a **homomorphism** from G to H is any map $\varphi : G \to H$ that 'respects the group structure', i.e., for which

$$\forall\{f, g \in G\}\ \varphi(fg) = \varphi(f)\varphi(g).$$

The prime example of a group is the collection of all invertible maps on some set, with map composition the group operation. The most important special case of these is $\mathbb{S}_n$, called the **symmetric group of order** n and consisting of all permutations of order n, i.e., of all invertible maps on $\underline{n} = \{1, 2, \ldots, n\}$. Any finite group G can be **represented by** a subgroup of $\mathbb{S}_n$ for some n in the sense that there is a **group monomorphism** $\varphi : G \to \mathbb{S}_n$, i.e., a 1-1 homomorphism from G to $\mathbb{S}_n$.

Here are some specific examples:

(i) $(\mathbb{Z}, +)$, i.e., the integers under addition; note that, for each $n \in \mathbb{Z}$, the map $n : \mathbb{Z} \to \mathbb{Z} : m \mapsto m + n$ is, indeed, invertible, with $-n : \mathbb{Z} \to \mathbb{Z} : m \mapsto m - n$ its inverse.

(ii) $(\mathbb{Q}\backslash 0, *)$, i.e., the nonzero rationals under multiplication; note that, for each $q \in \mathbb{Q}\backslash 0$, the map $q : \mathbb{Q}\backslash 0 \to \mathbb{Q}\backslash 0 : p \mapsto pq$ is, indeed, invertible, with $q^{-1} : \mathbb{Q}\backslash 0 \to \mathbb{Q}\backslash 0 : p \mapsto p/q$ its inverse.

(iii) The collection of all rigid motions that carry an equilateral triangle to itself. It can be thought of as $\mathbb{S}_3$ since each such motion, being rigid, must permute the vertices and is completely determined once we know what it does to the vertices.

17.1 Prove that, for $M = \{1, 2\}$, the semigroup M^M is not commutative.

17.2 Verify all the parenthetical claims made in the above definition of a group.

17.3 Give an example of a nonabelian group.

A **ring** $R = (R, +, *)$ is a set R (necessarily nonempty) with two operations, $(f, g) \mapsto f + g$ and $(f, g) \mapsto f * g =: fg$, called addition and multiplication respectively, such that

(r.1) $(R, +)$ is an Abelian group, with neutral element usually denoted 0;

(r.2) $(R, *)$ is a semigroup;

(r.3) (distributive laws): for every $f \in R$, the maps $R \to R : g \mapsto fg$ and $R \to R : g \mapsto gf$ are homomorphisms of the group $(R, +)$, i.e., $f(g + h) = fg + fh$ and $(g + h)f = gf + hf$.

A **field** is a ring $(R, +, *)$ for which $(R\backslash 0, *)$ is a group.

If multiplication in the ring R is commutative, i.e., $fg = gf$ for all $f, g \in R$, then R is called commutative.

If the ring R has a neutral element for multiplication, i.e., an element e so that $eg = g = ge$ for all $g \neq 0$, then it has exactly one such, and it is usually denoted by 1. In that case, R is called a **ring with identity**. Any field is a ring with identity.

Both $\mathbb{R}$ and $\mathbb{C}$ are commutative fields. The prime example of a ring is the set $\Pi(\mathbb{F}^d)$ of all polynomials in d (real or complex) variables with (real or complex) coefficients, with pointwise addition and multiplication the ring operations. It is a commutative ring with identity. It has given the major impetus to the study of (two-sided) **ideal**s, i.e., of nonempty subsets S of a ring R closed under addition, and containing both SR and RS, i.e., closed also under left or right multiplication by any element of the ring. This makes S a subring of R, i.e., a ring in its own right, but not all subrings are ideals.

Let R be a commutative ring. Then the set $[s_1, \ldots, s_r](R^r) = \{s_1 g_1 + \cdots + s_r g_r : (g_1, \ldots, g_r) \in R^r)\}$ is an ideal, the ideal **generated by** $(s_1, \ldots, s_r)$. Such an ideal is called **finitely generated**. A ring R is called **Noetherian** if all its ideals are finitely generated. **Hilbert's Basis Theorem** famously states that $\Pi(\mathbb{F}^d)$ is Noetherian.

17.4 Verify that, for any $s_1, \ldots, s_n$ in the commutative ring R, $[s_1, \ldots, s_n](R^n)$ is an ideal.

The ring of univariate polynomials

$\Pi = \Pi(\mathbb{F})$ is, by definition, the set of univariate polynomials, i.e., the collection of all maps

$$p : \mathbb{F} \to \mathbb{F} : z \mapsto \widehat{p}_0 + \widehat{p}_1 z + \widehat{p}_2 z^2 + \cdots + \widehat{p}_d z^d,$$

with $\widehat{p}_0, \ldots, \widehat{p}_d \in \mathbb{F}$ and some $d \in \mathbb{Z}_+$. If $\widehat{p}_d \neq 0$, then d is the **degree** of p, i.e.,

$$d = \deg p := \max\{j : \widehat{p}_j \neq 0\}.$$

This leaves the degree of the zero polynomial, $0 : \mathbb{F} \to \mathbb{F} : z \mapsto 0$, undefined. It is customary to set

$$\deg 0 := -1.$$

As already mentioned, Π is a ring under pointwise addition and multiplication. More than that, Π is a **principal ideal domain**, meaning that any of its ideals is generated by just one element. Indeed, if I is an ideal, then it contains an element p of smallest possible nonnegative degree. If f is any element of Π, then, by the Euclidean algorithm (see below), we can find $q, r \in \Pi$

so that $f = qp + r$ and $\deg r < \deg p$. If now $f \in I$, then also $r = f - qp \in I$ and $\deg r < \deg p$ hence, by the minimality of $\deg p$, r must be 0. In other words,

$$I = \Pi p := \{qp : q \in \Pi\}.$$

To be sure, already $\Pi(\mathbb{F}^2)$ fails to be a principal ideal domain.

It is simple algebra (see, e.g., the discussion of Horner's method below) that the set

$$Z(p) := \{z \in \mathbb{F} : p(z) = 0\}$$

of zeros of $p \in \Pi$ contains at most $\deg p$ elements. It is the **Fundamental Theorem of Algebra** that $\#Z(p) = \deg p$, counting multiplicities, in case $\mathbb{F} = \mathbb{C}$. More explicitly, this theorem says that, with $d := \deg p$,

$$p = c(\cdot - z_1)\cdots(\cdot - z_d)$$

for some nonzero constant c and some $z \in \mathbb{C}^d$.

It is in this sense that $\mathbb{C}$ is said to be **algebraically closed** while $\mathbb{R}$ is not. E.g., the real polynomial $()^2 + 1$ has no real zeros. It is remarkable that, by adjoining one, hence the other, of the 'imaginary' zeros of $()^2 + 1$, i.e., $\mathrm{i} = \sqrt{-1}$, appropriately to $\mathbb{R}$, i.e., by forming $\mathbb{C} = \mathbb{R} + \mathrm{i}\mathbb{R}$, we obtain enough additional scalars so that now, even if we consider polynomials with complex coefficients, all nonconstant polynomials have a full complement of zeros (counting multiplicities).

Convergence of a scalar sequence

A subset Z of $\mathbb{C}$ is said to be **bounded** if it lies in some ball

$$B_r := \{z \in \mathbb{C} : |z| < r\}$$

of (finite) radius r. Equivalently, Z is bounded if, for some r, $|\zeta| < r$ for all $\zeta \in \mathrm{Z}$. In either case, the number r is called a **bound for** Z.

In particular, we say that the scalar sequence $(\zeta_1, \zeta_2, \ldots)$ is **bounded** if the set $\{\zeta_m : m \in \mathbb{N}\}$ is bounded. For example, the sequence $(1, 2, 3, \ldots)$ is not bounded.

(17.2) Lemma: The sequence $(\zeta^1, \zeta^2, \zeta^3, \cdots)$ is bounded if and only if $|\zeta| \le 1$. Here, ζ^k denotes the kth power of the scalar ζ.

Proof: Assume that $|\zeta| > 1$. I claim that, for all m,

$$|\zeta^m| - 1 > (|\zeta| - 1)m. \tag{17.3}$$

This is certainly true for $m = 1$. Assume it correct for $m = k$. Then

$$|\zeta^{k+1}| - 1 = (|\zeta^{k+1}| - |\zeta^k|) + (|\zeta^k| - 1).$$

The first term on the right-hand side gives

$$|\zeta^{k+1}| - |\zeta^k| = (|\zeta| - 1)|\zeta|^{k-1} > |\zeta| - 1,$$

since $|\zeta| > 1$, while, for the second term, $|\zeta^k| - 1 > (|\zeta| - 1)k$ by induction hypothesis. Consequently,

$$|\zeta^{k+1}| - 1 > (|\zeta| - 1) + (|\zeta| - 1)k = (|\zeta| - 1)(k + 1),$$

showing that (17.3) also holds for $m = k + 1$.

In particular, for any given c, choosing m to be any natural number bigger than $c/(|\zeta| - 1)$, we have $|\zeta^m| > c$. We conclude that the sequence $(\zeta^1, \zeta^2, \zeta^3, \cdots)$ is unbounded when $|\zeta| > 1$.

Assume that $|\zeta| \le 1$. Then, for any m, $|\zeta^m| = |\zeta|^m \le 1^m = 1$, hence the sequence $(\zeta^1, \zeta^2, \zeta^3, \cdots)$ is not only bounded, it lies entirely in the **unit disk**

$$B_1^- := \{z \in \mathbb{C} : |z| \le 1\}.$$

A sequence $(\zeta_1, \zeta_2, \zeta_3, \cdots)$ of (real or complex) scalars is said to **converge to the scalar** ζ, in symbols:

$$\lim_{m\to\infty} \zeta_m = \zeta,$$

if, for all $\varepsilon > 0$, there is some m_ε so that, for all $m > m_\varepsilon$, $|\zeta - \zeta_m| < \varepsilon$.

Assuming without loss the scalars to be complex, we can profitably visualize this definition as saying the following: Whatever small circle $\{z \in \mathbb{C} : |z - \zeta| = \varepsilon\}$ of radius ε we draw around the point ζ, *all* the terms of the sequence except the first few are inside that circle.

(17.4) Lemma: A convergent sequence is bounded.

Proof: If $\lim_{m\to\infty} \zeta_m = \zeta$, then there is some m_0 so that, for all $m > m_0$, $|\zeta - \zeta_m| < 1$. Therefore, for all m,

$$|\zeta_m| \le r := |\zeta| + 1 + \max\{|\zeta_k| : k = 1{:}m_0\}.$$

Note that r is indeed a well-defined nonnegative number, since a *finite* set of real numbers always has a largest element. □

(17.5) Lemma: The sequence $(\zeta^1, \zeta^2, \zeta^3, \cdots)$ is convergent if and only if either $|\zeta| < 1$ or else $\zeta = 1$. In the former case, $\lim_{m\to\infty} \zeta^m = 0$, while in the latter case $\lim_{m\to\infty} \zeta^m = 1$.

Proof: Since the sequence is not even bounded when $|\zeta| > 1$, it cannot be convergent in that case. We already noted that it cannot be convergent when $|\zeta| = 1$ unless $\zeta = 1$, and in that case $\zeta^m = 1$ for all m, hence also $\lim_{m\to\infty} \zeta^m = 1$.

This leaves the case $|\zeta| < 1$. Then either $|\zeta| = 0$, in which case $\zeta^m = 0$ for all m, hence also $\lim_{m\to\infty} \zeta^m = 0$. Else, $0 < |\zeta| < 1$, therefore $1/\zeta$ is a well-defined complex number of modulus greater than one, hence, as we showed earlier, $1/|\zeta^m| = |(1/\zeta)^m|$ grows monotonely to infinity as $m \to \infty$. But this says that $|\zeta^m|$ decreases monotonely to 0. In other words, $\lim_{m\to\infty} \zeta^m = 0$. □

Horner, or: How to divide a polynomial by a linear factor

Recall that, given the polynomial p and one of its roots, μ, the polynomial $q := p/(\cdot - \mu)$ can be constructed by **synthetic division**. This process is also known as **nested multiplication** or **Horner's scheme** as it is used, more generally, to evaluate a polynomial efficiently. Here are the details, for a polynomial of degree ≤ 3.

If $p(t) = a_0 + a_1 t + a_2 t^2 + a_3 t^3$, and z is any scalar, then

$$p(z) = \underbrace{a_0 + z\,\underbrace{(a_1 + z\,\underbrace{(a_2 + z\,\underbrace{a_3}_{=:b_3})}_{=:b_2})}_{=a_1+zb_2=:b_1}}_{a_0+zb_1=:b_0}.$$

In other words, we write such a polynomial in **nested form** and then evaluate from the inside out. Each step is of the form

$$b_j := a_j + zb_{j+1}, \tag{17.6}$$

it involves one multiplication and one addition. The last number calculated is b_0; it is the value of p at z. There are 3 such steps for our cubic polynomial (the definition $b_3 := a_3$ requires no calculation!). So, for a polynomial of degree n, we would use n multiplications and n additions.

Now, not only is b_0 of interest, since it equals $p(z)$; the other b_j are also useful since

$$p(t) = b_0 + (t - z)(b_1 + b_2 t + b_3 t^2).$$

We verify this by multiplying out and rearranging terms according to powers of t. This gives

$$\begin{array}{rcccccccc} b_0 + (t-z)(b_1 + b_2 t + b_3 t^2) & = & b_0 & + & b_1 t & + & b_2 t^2 & + & b_3 t^3 \\ & & -zb_1 & - & zb_2 t & - & zb_3 t^2 & & \\ & = & b_0 - zb_1 & + & (b_1 - zb_2)t & + & (b_2 - zb_3)t^2 & + & b_3 t^3 \\ & = & a_0 & + & a_1 t & + & a_2 t^2 & + & a_3 t^3 \end{array}$$

The last equality holds since, by (17.6),

$$b_j - zb_{j+1} = a_j$$

for $j < 3$ while $b_3 = a_3$ by definition.

(17.7) Nested Multiplication (aka Horner): To evaluate the polynomial $p(t) = a_0 + a_1 t + \cdots + a_k t^k$ at the point z, compute the sequence $(b_0, b_1, \ldots, b_k)$ by the prescription

$$b_j := \begin{cases} a_j & \text{if } j = k; \\ a_j + zb_{j+1} & \text{if } j < k. \end{cases}$$

Then $p(t) = b_0 + (t - z)q(t)$, with

$$q(t) := b_1 + b_2 t + \cdots + b_k t^{k-1}.$$

In particular, if z is a root of p (hence $b_0 = 0$), then

$$q(t) = p(t)/(t - z).$$

Since $p(t) = (t - z)q(t)$, it follows that $\deg q < \deg p$. This provides another proof (see (3.22)) for the *easy* part of the *Fundamental Theorem of Algebra*, namely that a polynomial of degree k has at most k roots.

Euclid's Algorithm

Horner's method is a special case of **Euclid's Algorithm** which constructs, for given polynomials f and $p \neq 0$, (unique) polynomials q and r with $\deg r < \deg p$ so that

$$f = pq + r.$$

For variety, here is a nonstandard discussion of this algorithm, in terms of elimination.

Assume that

$$p(t) = \widehat{p}_0 + \widehat{p}_1 t + \cdots + \widehat{p}_d t^d, \quad \widehat{p}_d \neq 0,$$

and

$$f(t) = \widehat{f}_0 + \widehat{f}_1 t + \cdots + \widehat{f}_n t^n$$

for some $n \ge d$, there being nothing to prove otherwise. Then we seek a polynomial

$$q(t) = \widehat{q}_0 + \widehat{q}_1 t + \cdots + \widehat{q}_{n-d} t^{n-d}$$

for which

$$r := f - pq$$

has degree $< d$. This amounts to the square upper triangular linear system

$$\begin{array}{cccccccccc}
\widehat{p}_d\widehat{q}_0 & + & \widehat{p}_{d-1}\widehat{q}_1 & + & \cdots & & & & = & \widehat{f}_d \\
& & \widehat{p}_d\widehat{q}_1 & + & \widehat{p}_{d-1}\widehat{q}_2 & + \cdots & & & = & \widehat{f}_{d+1} \\
& & & \ddots & & & & & \vdots & \\
& & & & & \widehat{p}_d\widehat{q}_{n-d-1} & + & \widehat{p}_{d-1}\widehat{q}_{n-d} & = & \widehat{f}_{n-1} \\
& & & & & & & \widehat{p}_d\widehat{q}_{n-d} & = & \widehat{f}_n
\end{array}$$

for the unknown coefficients $\widehat{q}_0, \ldots, \widehat{q}_{n-d}$ which can be uniquely solved by back substitution since its diagonal entries all equal $\widehat{p}_d \neq 0$.

A real continuous function on a compact set in $\mathbb{R}^n$ has a maximum

This basis result of Analysis is referred to in these notes several times. It goes beyond the scope of these notes.

Here is the phrasing of this result that is most suited for these notes.

(17.8) Theorem: An upper semicontinuous real-valued function f on a closed and bounded set M in $X := \mathbb{R}^n$ has a maximum, i.e.,

$$\sup f(M) = f(m)$$

for some $m \in M$.

In particular, $\sup f(M) < \infty$.

A subset M of X is **closed** if $m = \lim_n x_n$ and $x_n \in M$, all n, implies that $m \in M$.

A subset M of X is **bounded** if $\sup \|M\| < \infty$.

A function $f : M \subset X \to \mathbb{R}$ is **continuous at** m if $\lim_n x_n = m$ implies that $\lim f(x_n) = f(m)$. The function is **continuous** if it is continuous at every point of its domain.

A function $f : M \subset X \to \mathbb{R}$ is **upper semicontinuous at** m if $\lim_n x_n = m$ implies that $\lim f(x_{\mu(n)}) \ge f(m)$ for every convergent subsequence $n \mapsto f(x_{\mu(n)})$ of $n \mapsto f(x_n)$.

Let $b := \sup f(M)$. Then, for each $r < b$, the set

$$M_r := \{m \in M : f(m) \ge r\}$$

is not empty. Also, M_r is closed, by the upper semicontinuity of f and bounded. Also, M_r is decreasing as r increases. This implies (by the Heine-Borel Theorem) that $\cap_r M_r$ is not empty. But, for any $m \in \cap_r M_r$, $f(m) \ge r$ for all $r < b$, hence $f(m) \ge b = \sup f(M)$, therefore $f(m) = \sup f(M)$.

The theorem is also valid if X is any finite-dimensional normed vector space. For, with V any basis for X, we can write $f = gV^{-1}$ with $g := FV$ upper semicontinuous on $V^{-1}M$ and $\sup f(M) = \sup g(V^{-1}M) = g(h)$ for some $h \in V^{-1}M$, and so $m := Vh$ does the job for f.

Rough index for these notes

