

Honours Linear Algebra

Course Compendium TATA53

Jonathan Nilsson



Last updated May 11, 2025

Preface

This text grew out of the lecture notes that I wrote in spring 2024 when I gave the course TATA53 "Honours Linear Algebra" at Linköping University.

The presentation in these notes assumes a good understanding of the concepts and techniques of basic linear algebra: Vectors, bases, linear systems, inner products, matrices, linear maps, eigenvalues and so on. In this text we generalize the concepts from a first linear algebra course to vector spaces over an arbitrary field, with a special focus on the complex numbers \mathbb{C} . We also cover number of additional topics not usually found in a first linear algebra course, notably matrix-factorizations, Jordan normal form, Perron-Frobenius theory, singular values, and a brief introduction to multilinear algebra and tensors. The course is largely theoretical in nature, but we also illustrate many applications of linear algebra, including dynamical systems such as predator-prey models, image and video compression, ranking models, principal component analysis, and neural networks.

The text also contains a number of **problems** for each main section of the text, a subset of these are suggested exercises for the course. Almost all problems have a **hint** in a separate section at the end of this text, these are intended to get you started if you get stuck, but you are encouraged to try and attack the problem without looking at the hints first. There is also a section of **answers**, sometimes including short proofs, for most problems.

Thanks to Göran Bergqvist whose lecture notes were useful when preparing the lectures. I also appreciate feedback from Mats Aigner, Magnus Herberthson, and Axel Hultman for suggestions regarding this text. Thanks also to the students of TATA53 in spring 2024 and 2025 who found several errors and typos in these notes.

Contents

1	Vector spaces	7
1.1	Motivation	7
1.2	Definition of vector spaces	7
1.3	Basis and dimension	8
1.4	Subspaces	8
1.5	Direct sum	9
1.6	Affine subsets	11
1.7	Linear maps	12
2	Matrices	15
2.1	A basis for the matrix space	15
2.2	Echelon forms	17
2.3	Elementary matrices	17
2.4	LU-decomposition	18
2.5	Cholesky-factorization	20
2.6	Determinants	21
2.7	Block matrices	22
3	Introductory spectral theory	25
3.1	Eigenvalues and eigenvectors	25
3.2	Invariant subspaces	28
3.3	Matrix polynomials	29
4	Jordan normal form	33
4.1	Properties of matrices on Jordan form	33
4.2	Structure theory for nilpotent operators	35
4.3	Jordan chains and Jordanization	40
4.4	Matrix functions and applications	44
5	Inner product spaces	51
5.1	Inner products	51
5.2	Norms	53
5.3	Orthonormal bases and projections	55
5.4	Gram-Schmidt	56
5.5	QR-decomposition	59
5.6	Self-adjoint, unitary, and normal operators	62
5.7	Spectral theorem for normal operators	63
5.8	Positive definite operators	65
5.9	Least squares	67
6	Perron-Frobenius theory	71
6.1	Positive and non-negative matrices	71
6.2	Perron's theorem	72
6.3	Primitive matrices and the power method	74
6.4	Frobenius' theorem for irreducible matrices	76
6.5	Connection to graphs	77
6.6	Ranking models	79

6.7	Markov chains	82
7	Singular values	85
7.1	Geometrical intuition	85
7.2	Singular values and singular vectors	86
7.3	Singular values and matrix-norms	92
7.4	SVD as a sum - Schmidt-decomposition	92
7.5	Low rank approximation	94
7.6	Moore-Penrose pseudo inverse	98
7.7	Condition number	100
7.8	Polar decomposition	100
7.9	Data analysis with SVD	101
7.10	Principal component analysis	107
7.11	Total least squares	113
8	Multilinear algebra	115
8.1	Introduction	115
8.2	Tensor products of vector spaces	115
8.3	Dual spaces	119
8.4	Higher order tensors	122
8.5	Tensor products of linear maps	125
8.6	Kronecker product	126
8.7	Tensors in neural networks	129
9	Appendix	133
9.1	Fields	133
9.2	Complexification and realification	136
9.3	Basis in infinite dimension	137
10	Problems	139
11	Hints	167
12	Answers	183

Notation

For this course you will need to absorb material from various sources, and authors prefer different notations, something one has to get used to in more advanced math courses. Here is a list of common notations that appear in this and other texts on linear algebra.

$\mathbb{R}, \mathbb{C}, \mathbb{Q}$	The field of real, complex, and rational numbers respectively.
\mathbb{Z}_p or $\mathbb{Z}/p\mathbb{Z}$ or $\mathbb{Z}/(p)$	The field of integers modulo some prime p .
$\mathbb{F}, \mathbb{K}, k$	Common notation for an arbitrary field (like the ones above).
$u, \mathbf{v}, \bar{w}, \vec{x}$	Common notations for vectors.
\mathcal{P}_n or \mathbb{P}_n	Polynomials of degree $\leq n$, with coefficients in some field.
$\mathcal{C}[a, b]$	Set of continuous functions $[a, b] \rightarrow \mathbb{R}$
$\mathcal{C}^n[a, b]$	Set of n times continuously differentiable functions $[a, b] \rightarrow \mathbb{R}$
$A, B, C, D, E, M, N, X, Y$	Capital letters commonly used for matrices.
$\lambda, \mu, \alpha, \beta, a, b, c, s, t, r, z$	Lower-case and Greek letters commonly used for scalars.
$(1, 2, 3)$ or $(1, 2, 3)^T$	Vectors in \mathbb{R}^3 , some authors always prefer writing them as columns.
$\text{Mat}_{m \times n}(\mathbb{C}), M_{m \times n}(\mathbb{C})$	The set of $m \times n$ matrices with complex coefficients.
$\text{Mat}_n(\mathbb{C}), M_n(\mathbb{C})$	Shorthand for the above when $m = n$ (the matrix is square)
$\text{Mat}_{m \times n}, M_{m \times n}, \text{Mat}_n, M_n$	Same as above, the field is supposed to be understood by the context.
$\text{diag}(d_1, d_2, d_3)$	The diagonal matrix with d_1, d_2, d_3 on the diagonal.
$(a_{ij})_{i,j}, (a_{ij})$	The matrix with the number a_{ij} in position (i, j)
A_{ij}	The element of the matrix A at position (i, j)
A_i	The i 'th column of the matrix A
δ_{ij}	The Kronecker delta function, 1 if $i = j$, otherwise 0.
$e_{ij}, e_{i,j}, E_{i,j}, E_{ij}$	The unit-matrix with a single 1 in position (i, j) and zeroes elsewhere.
e_i	Standard basis vector in \mathbb{C}^n (or \mathbb{F}^n) with a single 1 in position i .
$\text{span}(v_1, v_2, v_3), [v_1, v_2, v_3]$	The span of v_1, v_2, v_3 (the set of all linear combinations of the vectors).
$\ker(F), N(F)$	The kernel, or nullspace of the linear map (or matrix) F .
$\text{Im}(F), \text{Ran}(F), V(F)$	The image or range of a linear map (or matrix) F .
A^T, A^t	The transpose of a matrix A .
$A^*, \bar{A}^T, A^H, A^\dagger$	The conjugate transpose of a matrix A .
$\begin{pmatrix} 2 \\ 3 \end{pmatrix}_{\mathbf{e}}$ or $\mathbf{e} \begin{pmatrix} 2 \\ 3 \end{pmatrix}$	$2e_1 + 3e_2$, the vector with coordinates $(2, 3)$ in basis $\mathbf{e} = (e_1, e_2)$
$v_{\mathbf{e}}$	The coordinate vector of v with respect to the basis \mathbf{e}
$F_{\mathbf{e}}$ or $[F]_{\mathbf{e}}$	The matrix of $F : V \rightarrow V$ with respect to a basis \mathbf{e} for V .
$[F]_{\mathbf{f}}$ or $_{\mathbf{f}}[F]_{\mathbf{e}}$	The matrix of $F : V \rightarrow W$ with respect to a bases \mathbf{e} for V , and \mathbf{f} for W .
$[F]$	the matrix of F with respect to some basis understood by the context.
$\sigma(F)$ or $\text{Spec}(F)$	The spectrum of F , the set of eigenvalues.
$p_A(\lambda)$ or just p_A	$\det(A - \lambda I)$, the characteristic polynomial (also called the secular polynomial).
$m_A(\lambda)$ or just m_A	The minimal polynomial of A
$\text{rnk}(F), \text{rank}(F)$	The rank of a matrix (or a linear map) F .
E_λ	The eigenspace for the eigenvalue λ
\tilde{E}_λ	The generalized eigenspace for the eigenvalue λ
$u \bullet v, u \cdot v, \bar{v}^T u$	The dot product of u , and v , the standard inner product on \mathbb{R}^n or \mathbb{C}^n
$(u, v), (u v), \langle u, v \rangle$	common notations for various inner products of u and v
$\ v\ $ or $ v $	The norm of v or the length of v .
$\ A\ _F$	The Frobenius norm of a matrix A (root of sum of squared absolute values of entries)
U^\perp	The orthogonal complement to a subspace U (with respect to some inner product)
$P_U(v), \text{proj}_U(v), v_{\ U}, v_U$	The projection of v onto the subspace U
$A > 0$	The matrix A is positive (all entries are positive)
$\deg p(t), \text{deg } p$	The degree of a polynomial p
m_λ, g_λ	The algebraic and the geometric multiplicity of an eigenvalue λ .
A^+	The Moore-Penrose pseudo-inverse of A .

Chapter 1

Vector spaces

1.1 Motivation

In a first linear algebra course, we typically think of a **vector** as something that looks like $v = (1, 2, -3)$, a triple of real numbers. However, it turns out that the concepts, tools, and techniques of linear algebra (such as linear systems, linear maps, matrices, eigenvalues, etc) are useful in a much broader context, and they can be used to solve problems seemingly unrelated to the space \mathbb{R}^n . In fact, it turns out that all we need is any type of objects that we can add and multiply by scalars in a coherent way.

1.2 Definition of vector spaces

Definition 1.2.1. A **vector space** over a field^a \mathbb{F} is a set V together with an "addition" operation $V \times V \rightarrow V$ and a "scalar multiplication" $\mathbb{F} \times V \rightarrow V$ that satisfy the following **vector space axioms** for all $u, v, w \in V$ and all $\lambda, \mu \in \mathbb{F}$:

$$(V1) \quad u + v = v + u$$

$$(V2) \quad (u + v) + w = u + (v + w)$$

$$(V3) \quad \text{There is an element } 0 \in V \text{ satisfying } 0 + v = v \text{ for all } v \in V$$

$$(V4) \quad \text{For every } v \in V \text{ there is an element } -v \in V \text{ such that } v + (-v) = 0$$

$$(V5) \quad 1 \cdot v = v$$

$$(V6) \quad \lambda \cdot (\mu \cdot v) = (\lambda\mu) \cdot v$$

$$(V7) \quad \lambda \cdot (u + v) = (\lambda \cdot u) + (\lambda \cdot v)$$

$$(V8) \quad (\lambda + \mu) \cdot v = (\lambda \cdot v) + (\mu \cdot v)$$

^aIn this course the focus will be on the fields \mathbb{R} or \mathbb{C} , but we will also consider some examples and applications where \mathbb{F} is a finite field. A proper definition of what a field is can be found in the appendix. A discussion of finite fields can be found in the appendix.

A vector space over \mathbb{F} is also called an \mathbb{F} -vector space. The elements of \mathbb{F} are called **scalars**, while the elements of the vector space V are called **vectors**. We typically write λv instead of $\lambda \cdot v$. A prototypical example of a real vector space is \mathbb{R}^2 , the set of all pairs (x, y) of real numbers, where addition is defined as $(x_1, y_1) + (x_2, y_2) := (x_1 + x_2, y_1 + y_2)$ and $\lambda \cdot (x, y) = (\lambda x, \lambda y)$. The same construction works if we replace \mathbb{R} by \mathbb{C} or any other field. There are however many other examples:

- \mathbb{C}^n , the set of complex n -tuples with coordinate-wise sum and scalar multiplication is a complex vector space.
- The set \mathcal{P}_n of polynomials with complex coefficients and degree $\leq n$ is a complex vector space.
- The set $\text{Mat}_{m \times n}(\mathbb{R})$ of real $m \times n$ -matrices (with the usual ways of adding matrices and multiplying matrices by real numbers) is a real vector space.

- The set $\mathcal{C}(\mathbb{R})$ of continuous functions $f : \mathbb{R} \rightarrow \mathbb{R}$ is a real vector space.
- The set of infinite sequences of complex numbers (a_1, a_2, \dots) is a complex vector space (like \mathbb{C}^n where $n = \infty$)
- The set of solutions to the differential equation $y''(x) + y'(x) - 6y(x) = 0$ is a real vector space.
- The field \mathbb{Z}_2 of two elements is the set $\{0, 1\}$ where addition and multiplication is defined like on the real numbers except that $1 + 1 := 0$. The set of triples of 0's and 1's form a vector space over \mathbb{Z}_2 .

1.3 Basis and dimension

For reference we recap some of the important concepts familiar from a first linear algebra course. Let V be a vector space over a field \mathcal{F} , all vectors below belong to V and all scalars lie in \mathcal{F} . A **linear combination** of vectors $v_1, \dots, v_n \in V$ is a vector of form $\lambda_1 v_1 + \dots + \lambda_n v_n$, where the coefficients λ_i are scalars. The set of all such linear combination is called their **span**, we write

$$\text{span}(v_1, \dots, v_n) = \{\lambda_1 v_1 + \lambda_2 v_2 + \dots + \lambda_n v_n \mid \lambda_i \in \mathbb{F}\}.$$

If $V = \text{span}(v_1, \dots, v_n)$ we say that the vectors span or generate V , this means that every vector of V can be expressed as a linear combination of v_1, \dots, v_n .

On the other hand, the vectors v_1, \dots, v_n are called **linearly independent** if

$$\lambda_1 v_1 + \dots + \lambda_n v_n = 0 \implies \lambda_i = 0 \forall i.$$

In other words, the only linear combination of the vectors that is zero is the trivial combination.

An ordered set of vectors $\mathcal{B} = (v_1, \dots, v_n)$ that both spans V and is linearly independent is called a **basis** for V , and we define the **dimension** $\dim V = n$, the number of basis vectors. The basis-conditions guarantee (and are equivalent to) the fact that every vector in V has a unique expression as a linear combination of the basis vectors. If $v = \lambda_1 v_1 + \dots + \lambda_n v_n$ we say that $(\lambda_1, \dots, \lambda_n)$ is the **coordinate vector** of v with respect to the basis $\mathcal{B} = (v_1, \dots, v_n)$. Different choices of bases produces different coordinate-vectors, so if we want to explicitly state what basis we refer to we can write this as $v = (\lambda_1, \dots, \lambda_n)_{\mathcal{B}}$.

If we are working with a fixed basis, then summing and rescaling vectors in V corresponds to making the same operations on the respective coordinate-vectors, so with the basis in hand we can forget V and instead do all computations in the vector space \mathbb{F}^n , the set of n -tuples of elements from the field.

Example 1.3.1. Consider again some of the vector spaces discussed above.

- In \mathbb{C}^3 , a basis is $((1, 0, 0), (0, 1, 0), (0, 0, 1))$. The dimension is 3.
- In \mathcal{P}_3 , a basis is $(1, x, x^2, x^3)$. The dimension is 4.
- A basis for $\text{Mat}_{2 \times 2}(\mathbb{R})$ is $\left(\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}\right)$. The dimension is 4.
- The set of continuous functions is infinite-dimensional (and it's hard to write down a basis^a)
- The solution space to $y''(x) + y'(x) - 6y(x) = 0$ is 2-dimensional with basis (e^{-3x}, e^{2x}) .

^aIn fact, the axiom of choice is required.

△

1.4 Subspaces

Definition 1.4.1. Let V be a vector space. A nonempty subset $U \subset V$ is called **subspace** of V if it is closed under taking sums and products by scalars:

- $u_1, u_2 \in U \implies u_1 + u_2 \in U$
- $u \in U, \lambda \in \mathbb{F} \implies \lambda \cdot u \in U$

The definition is equivalent to saying that U itself is a vector space with the addition and scalar action inherited from V .

Example 1.4.2. Let $V = \text{Mat}_3(\mathbb{R})$, the real vector space of 3×3 -matrices. Let U be the subset of V consisting of *skew-symmetric* matrices:

$$U = \{A \in V \mid A^T = -A\}.$$

We show that U is a subspace of V by verifying the conditions of Definition 1.4.1. Clearly U is nonempty (for example, the zero-matrix is clearly skew-symmetric). Assume that A and B are skew-symmetric. Then we have

$$(A + B)^T = A^T + B^T = (-A) + (-B) = -(A + B)$$

which shows that $A + B$ is skew symmetric too and $A + B \in U$. For the second condition, let A be skew-symmetric and let $\lambda \in \mathbb{R}$. Then

$$(\lambda A)^T = \lambda(A^T) = \lambda(-A) = -(\lambda A)$$

which shows that $\lambda A \in U$. Therefore U is a subspace of V . Since U is a vector space in its own right we can find a basis for it. Solving the equation $A = -A^T$ shows that U is three-dimensional with basis

$$\left(\begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{pmatrix} \right).$$

△

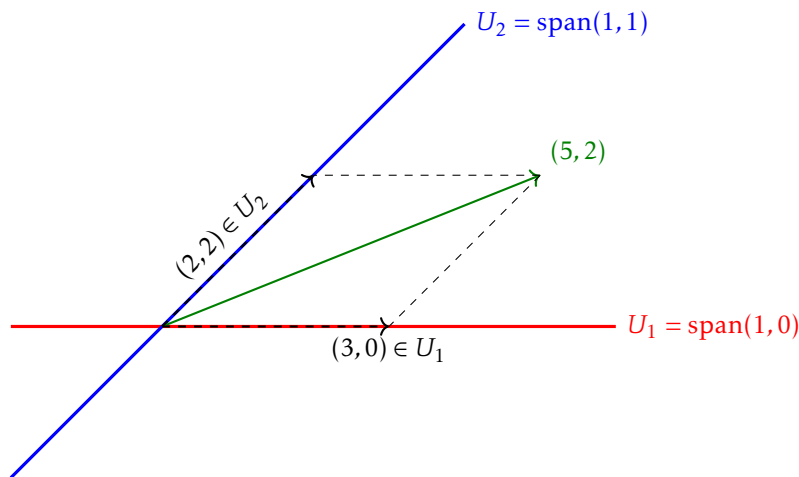
1.5 Direct sum

A basic way to understand an object is to break it down into smaller pieces and analyze them separately.

Definition 1.5.1. Let U_1 and U_2 be two subspaces of a vector space V . If every vector $v \in V$ has a *unique representation* $v = u_1 + u_2$ where $u_1 \in U_1$ and $u_2 \in U_2$, we say V is the (internal) **direct sum** of U_1 and U_2 , and we write

$$V = U_1 \oplus U_2.$$

Example 1.5.2. Let $U_1 = \text{span}((1,0))$ and^a $U_2 = \text{span}((1,1))$ be two lines in \mathbb{R}^2 . Then $U_1 \oplus U_2 = \mathbb{R}^2$; every vector $v \in \mathbb{R}^2$ has a unique representation as $u_1 + u_2$ with $u_1 \in U_1$ and $u_2 \in U_2$. For example, $(5,2) = (3,0) + (2,2)$, or more generally, $(x,y) = (x-y,0) + (y,y)$.



^aWhen we are taking the span of a single vector we shall sometimes skip the parentheses and write it as $\text{span}(1,0)$.

The subspaces in the direct sum need not be 1-dimensional: △

Example 1.5.3. Let $\pi = \{(x, y, z) \in \mathbb{R}^3 \mid x + 2y + 3z = 0\}$ and $\ell = \text{span}((1, 1, 1))$. Then

$$\mathbb{R}^3 = \pi \oplus \ell.$$

On the other hand, \mathbb{R}^3 can never be the direct sum of two planes. For example, with △

$$U_1 : z = 0 \text{ and } U_2 : x + y + z = 0,$$

each vector in \mathbb{R}^3 can be expressed as a sum $u_1 + u_2$ of vectors from the two subspaces, but the expression is not unique, for example there are two ways (and infinitely many more ways) of writing the vector $(1, 2, 3)$ as a sum $u_1 + u_2$ of vectors from each respective subspace:

$$(1, 5, 0) + (0, -3, 3) = (2, 4, 0) + (-1, -2, 3).$$

A vector space can be expressed as a direct sum in many different way, the decomposition one chooses depends on what one wants to achieve.

Note that if $V = U_1 \oplus U_2$ we have a natural way to project a vector onto each subspace:

$$P_{U_1}(v) = u_1 \text{ and } P_{U_2}(v) = u_2 \text{ where } v = u_1 + u_2.$$

Because of the uniqueness of the expression $v = u_1 + u_2$ these functions are well defined. Note that this is different from *orthogonal projection*¹, for example, in the graphical example above we have $P_{U_2}(5, 2) = (2, 2)$.

Theorem 1.5.4. Let U_1 and U_2 be subspaces of a vector space V . Then $V = U_1 \oplus U_2$ if and only if $U_1 + U_2 = V$ and $U_1 \cap U_2 = \{0\}$.

Proof. The first condition $U_1 + U_2 = V$ is clearly equivalent to being able to express each $v \in V$ as $v = u_1 + u_2$. We show that the second condition is equivalent to uniqueness of the expression: assume that $U_1 \cap U_2$ contains some nonzero vector u_0 . Then for any vector v with representation $v = u_1 + u_2$, we can also write $v = (u_1 + u_0) + (u_2 - u_0)$, so no vector has a unique representation. Conversely, if a vector v has two distinct representations $u_1 + u_2 = v = u'_1 + u'_2$, then subtracting we get $(u_1 - u'_1) = (u'_2 - u_2)$. This vector clearly lies in U_1 since the left hand does, and it lies in U_2 since the right side does. Moreover the vector is nonzero since we assumed the representations were different. Thus we have a nonzero vector that lies in both U_1 and U_2 , so $U \cap U' \neq \{0\}$. □

Note that if (u_1, \dots, u_n) is a basis for U and if (u'_1, \dots, u'_m) is a basis for U' , then $(u_1, \dots, u_n, u'_1, \dots, u'_m)$ is a basis for $V = U \oplus U'$. In particular $\dim(U \oplus U') = \dim(U) + \dim(U')$.

Example 1.5.5. Consider the vector space $V = \text{Mat}_n(\mathbb{R})$, and let S be the subspace consisting of symmetric matrices, and let S' be the subspace consisting of skew-symmetric matrices. Let us prove that

$$V = S \oplus S'.$$

According to the theorem above, it suffices to prove that $S + S' = V$ and that $S \cap S' = \{0\}$. Starting with the latter, assume that a matrix A lies in $S \cap S'$. Then A is both symmetric and skew symmetric, so $A = A^T = -A$, so $2A = 0$ and therefore $A = 0$. Thus S and S' intersects only in the zero matrix.

For the other part we need to show that any square matrix can be expressed as the sum of a symmetric and a skew-symmetric matrix. For this we note that $B = A + A^T$ is symmetric (since $B^T = (A + A^T)^T = A^T + A = B$), and that $C = (A - A^T)$ is skew symmetric (since $C^T = (A - A^T)^T = A^T - A = -C$). But then we can express A as a sum a symmetric and a skew-symmetric matrix as follows:

$$A = \frac{1}{2}(A + A^T) + \frac{1}{2}(A - A^T).$$

We can also talk about direct sum of more than two components in an analogous way: If U_1, \dots, U_n are subspaces of V , we say that $V = U_1 \oplus \dots \oplus U_n$ if and only if each vector $v \in V$ has a unique expression △

$$v = u_1 + \dots + u_n \text{ where } u_i \in U_i.$$

¹Indeed, in a general vector space we don't even have the concept of orthogonality before defining an inner product. Note also that it doesn't make sense to project on a subspace U_1 in this sense without specifying what we choose as U_2 .

²Here $U_1 + U_2 := \{u_1 + u_2 \mid u_1 \in U_1, u_2 \in U_2\}$.

External direct sum

There is also an analogous construction called the **external direct sum**. Starting with two vector spaces we can construct a larger vector space that has the two vector spaces as direct summands in the above sense (roughly speaking).

Definition 1.5.6. Let V and W be vector spaces over the same field \mathbb{F} . We define $V \oplus W$ to be the set $V \times W$ which consists of all pairs (v, w) where $v \in V$ and $w \in W$, and we define the sum and scalar action on such pairs in the natural way:

$$(v, w) + (v', w') := (v + v', w + w') \text{ and } \lambda \cdot (v, w) := (\lambda \cdot v, \lambda \cdot w).$$

Under these operations $V \oplus W$ becomes a vector space, called the (external) **direct sum** of V and W .

Note that technically V and W are not subspaces (and not even a subsets) of $V \oplus W$, since the latter object consists of elements of form (v, w) while V and W do not. However, if we identify V with pairs of form $(v, 0)$, and W with pairs $(0, w)$, then $V \oplus W$ is the internal direct sum of V and W . If (v_1, \dots, v_m) is a basis for V and if (w_1, \dots, w_n) is a basis for W then

$$((v_1, 0), \dots, (v_m, 0), (0, w_1), \dots, (0, w_n))$$

is a basis for $V \oplus W$.

Example 1.5.7. The vector space $V = \text{Mat}_{2 \times 2}(\mathbb{R}) \oplus \mathcal{C}(\mathbb{R})$ consists of objects of form $(A, f(x))$, where A is a 2×2 -matrix and $f : \mathbb{R} \rightarrow \mathbb{R}$ is a continuous function. Here's an example of a linear combination in V :

$$\frac{1}{2} \left(\begin{pmatrix} 1 & 2 \\ 4 & 3 \end{pmatrix}, e^x \right) - \frac{1}{2} \left(I, e^{-x} \right) = \left(\begin{pmatrix} 0 & 1 \\ 2 & 1 \end{pmatrix}, \sinh x \right).$$

△

External direct sums of more than two vector spaces can also be defined in the natural way.

1.6 Affine subsets

If u is a vector of a subspace, we know that $0 \cdot u \in U$. This means that the zero-vector always belongs to a subspace. An affine subset is something that looks like a subspace (think a line or a plane), but it may be shifted away from the origin.

Definition 1.6.1. Let U be a subspace of V . For each $v \in V$ we define

$$v + U := \{v + u \mid u \in U\}$$

and call this an **affine subset** (or affine subspace) parallel to U .

Intuitively, this is just the subspace U shifted by a vector v . For example, the line

$$\{(2 + 3t, 1 + 5t) \in \mathbb{R}^2 \mid t \in \mathbb{R}\}$$

is an affine subset of \mathbb{R}^2 . It can be written as $(2, 1) + U$ where $U = \text{span}(3, 5)$. Similarly, the plane $x + 2y + 3z = 5$ is an affine subspace of \mathbb{R}^3 , it can be written $(0, 1, 1) + U$ where U is the plane $x + 2y + 3z = 0$.

Note that $v \in U$ if and only if $v + U = U$. Also note that an affine subset typically is not closed under addition or scalar multiplication.

It turns out that the set of all affine subsets is itself a vector space in a natural way.

Definition 1.6.2. Let U be a subspace of V . We define the **quotient space**

$$V/U := \{v + U \mid v \in V\}.$$

This is a vector space in the natural way, addition of two affine subsets is defined as

$$(v + U) + (v' + U) := (v + v') + U$$

and multiplication of scalars is defined by

$$\lambda(v + U) := (\lambda v) + U.$$

So intuitively, to add two affine subsets, pick a vector in each one, add them, and then take the affine subset which the sum lies in. It turns out this operation is well defined. For example. Let $U = \text{span}(1, 1)$. Then \mathbb{R}^2/U is the set of lines in \mathbb{R}^2 with slope 1.

If (u_1, \dots, u_m) is a basis for U and if we extend it to a basis (u_1, \dots, u_n) for V , it is easy to see that a basis for V/U is given by the affine subsets

$$(u_{m+1} + U, \dots, u_n + U).$$

In particular $\dim(V/U) = \dim(V) - \dim(U)$. Note however that V/U is not itself a subset of V .

Note also that $0 + U = u + U$ for $u \in U$ - this means that in V/U , we can't tell the difference between different elements of U , so a good way to think about V/U is that we take V and then we "make all the elements of U equal".

1.7 Linear maps

Definition 1.7.1. Let V and W be \mathbb{F} -vector spaces. A map $F : V \rightarrow W$ is called **linear** if

$$F(u + v) = F(u) + F(v) \quad \text{and} \quad F(\lambda v) = \lambda F(v)$$

holds for all $u, v \in V$ and $\lambda \in \mathbb{F}$.

Linear maps are sometimes called linear transformations, homomorphisms, or operators if $V = W$.

A linear map is completely determined by its action on the basis vectors, because if $\mathcal{B} = (v_1, \dots, v_n)$ is a basis for V , a linear map F satisfies

$$F(\lambda_1 v_1 + \dots + \lambda_n v_n) = \lambda_1 F(v_1) + \dots + \lambda_n F(v_n).$$

In fact, if pick a basis $\mathcal{B}' = (w_1, \dots, w_m)$ of W and express the vectors $F(v_i)$ in this basis, and put these coordinate-vectors $F(v_i)_{\mathcal{B}'}$ as columns in a matrix, we get the matrix of F with respect to the two bases \mathcal{B} and \mathcal{B}' :

$$[F]_{\mathcal{B}', \mathcal{B}} = \begin{pmatrix} | & | & \cdots & | \\ F(v_1) & F(v_2) & \cdots & F(v_n) \\ | & | & \cdots & | \end{pmatrix}.$$

Example 1.7.2. Consider the map $D : \mathcal{P}_3 \rightarrow \mathcal{P}_2$ defined by taking derivative, $D(p(x)) = p'(x)$. This map is linear because of the familiar rules proved in a first calculus course:

$$\frac{d}{dx}(f(x) + g(x)) = \frac{d}{dx}f(x) + \frac{d}{dx}g(x) \quad \text{and} \quad \frac{d}{dx}(\lambda f(x)) = \lambda \frac{d}{dx}f(x).$$

Let $\mathcal{B} = (1, x, x^2, x^3)$ and $\mathcal{B}' = (1, x, x^2)$ be the standard bases \mathcal{P}_3 and \mathcal{P}_2 respectively. To find the matrix for D with respect to these bases, we evaluate D on the basis vectors in \mathcal{B} and express them in the basis \mathcal{B}' :

$$D(1) = 0 = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}_{\mathcal{B}'}, \quad D(x) = 1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}_{\mathcal{B}'}, \quad D(x^2) = 2x = \begin{pmatrix} 0 \\ 2 \\ 0 \end{pmatrix}_{\mathcal{B}'}, \quad D(x^3) = 3x^2 = \begin{pmatrix} 0 \\ 0 \\ 3 \end{pmatrix}_{\mathcal{B}'}. .$$

So the matrix for D with respect to \mathcal{B} and \mathcal{B}' is

$$[D]_{\mathcal{B}',\mathcal{B}} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix}.$$

△

There are two important subspaces associated to a linear map:

Definition 1.7.3. Let $F : V \rightarrow W$ be linear. We define

$$\ker(F) = \{v \in V \mid F(v) = 0\} \quad \text{and} \quad \text{Im}(F) = \{F(v) \mid v \in V\}$$

and call these the **kernel** and **image** of F respectively.

Note that $\ker(F)$ is a subset of V and $\text{Im}(F)$ is a subset of W , in fact it is not hard to prove that they are *subspaces*. The kernel is also called the nullspace of F , and the image is sometimes called the range of F .

We recall the important *dimension theorem*, also called the *rank-nullity theorem*:

Theorem 1.7.4. Let $F : V \rightarrow W$ be linear. Then

$$\dim \ker(F) + \dim \text{Im}(F) = \dim V.$$

The theorem holds even for infinite-dimensional vector spaces if we define $5 + \infty = \infty$ and so on. The dimension of the image $\dim \text{Im}(F)$ is also called the **rank** of the linear map.

Recall also that if $F : V \rightarrow W$ is a map, an **inverse** of F is a map in the other direction $G : W \rightarrow V$ such that

$$G(F(v)) = v \text{ for all } v \in V \quad \text{and} \quad F(G(w)) = w \quad \text{for all } v \in V \text{ and } w \in W.$$

This is commonly expressed as $G \circ F = \text{id}_V$ and³ $F \circ G = \text{id}_W$.

Direct sum of linear maps

Let $F : V \rightarrow V$ and $G : W \rightarrow W$ be linear maps. Then we get a corresponding linear map

$$F \oplus G : V \oplus W \rightarrow V \oplus W \quad \text{defined by} \quad (F \oplus G)(v, w) = (F(v), G(w)).$$

Note that if A is the matrix for F with respect to a given basis \mathcal{B}_V of V , and if B is the matrix of G with respect to a given basis \mathcal{B}_W of W , then the matrix of $F \oplus G$ with respect to the corresponding basis $\mathcal{B}_V \cup \mathcal{B}_W$ of $V \oplus W$ is of block form

$$\left(\begin{array}{c|c} A & 0 \\ \hline 0 & B \end{array} \right).$$

Linear maps and quotient spaces

Let U be a subspace of BV . We have a natural linear projection map

$$\pi : V \rightarrow V/U \quad \text{defined by} \quad \pi(v) = v + U.$$

On the other hand, suppose $F : V \rightarrow W$ is a linear map, and suppose that U is a subspace of V that lies inside the kernel of F : $F(U) = \{0\}$. Then we can construct a corresponding linear map

$$\tilde{F} : V/U \rightarrow W \text{ defined by } F(v + U) := F(v).$$

The condition $F(U) = \{0\}$ guarantees that the map is well defined.

³Here $\text{id}_V : V \rightarrow V$ is the identity map on V mapping each v to itself.

Chapter 2

Matrices

2.1 A basis for the matrix space

Let e_{ij} be the $m \times n$ -matrix which has a single 1 in position (i, j) and zeroes elsewhere¹. The set of all these matrices clearly form a basis for $\text{Mat}_{m \times n}(\mathbb{F})$. Note that if e_{ij} and e_{kl} are such matrices (of compatible sizes), then

$$e_{ij}e_{kl} = \begin{cases} e_{il} & \text{if } j = k \\ 0 & \text{otherwise} \end{cases}.$$

This is typically expressed more compactly as $e_{ij}e_{kl} = \delta_{jk}e_{il}$, where δ_{jk} is the **Kronecker-delta** function, it is 1 if $j = k$ and zero otherwise.

We write $(a_{ij})_{ij}$ or just (a_{ij}) for the matrix that has element a_{ij} in position (i, j) , in other words $(a_{ij})_{ij} = \sum_{i,j} a_{ij}e_{ij}$.

Let $A = (a_{ij})_{ij}$ be a matrix. Recall that the **transpose** of A is $A^T = (a_{ji})_{ij}$. For complex matrices we also define the **Hermitian conjugate** of A as $A^* = (\overline{a_{ji}})_{ij}$, this is just the conjugate-transpose of A .

A matrix $A = (a_{ij})_{ij}$ is called...

- Diagonal if $a_{ij} = 0$ whenever $i \neq j$
- Upper triangular if $a_{ij} = 0$ whenever $i > j$
(strictly upper triangular if $a_{ii} = 0$ also)
- Lower triangular if $a_{ij} = 0$ whenever $j > i$
(strictly lower triangular if $a_{ii} = 0$ also)
- Symmetric if $a_{ij} = a_{ji}$
(skew-symmetric if $a_{ij} = -a_{ji}$)
- Hermitian if $a_{ij} = \overline{a_{ji}}$
(skew-Hermitian if $a_{ij} = -\overline{a_{ji}}$)

Note that a matrix needs to be square in order to be symmetric/skew-symmetric/Hermitian, but the other concepts apply for any size of matrix.² Note also that Hermitian and symmetric has the same meaning when the matrix is real.

From the product-rule for matrices e_{ij} it follows that matrix multiplication can be expressed like this: If $A = (a_{ij})_{ij}$ is an $m \times n$ -matrix and $B = (b_{ij})_{ij}$ is an $n \times k$ -matrix, then

$$AB = \left(\sum_{r=1}^n a_{ir}b_{rj} \right)_{ij},$$

this is just an algebraic way of expressing the familiar rule that the element in position (i, j) in the product is the scalar-product of row number i in A and column number j in B .

¹Note that just writing e_{ij} is not completely clear since the size of the matrix is not specified, the 3×3 -matrix e_{12} is different from the 2×2 -matrix e_{12} . However, the format of the matrix is usually obvious from the context.

²Although when we say diagonal matrix without qualification, we usually mean a square matrix.

Note that $e_{12}e_{23} = e_{13}$ while $e_{23}e_{12} = 0$, so in general $AB \neq BA$ when A and B are matrices. We say that two $n \times n$ -matrices A and B **commute** if $AB = BA$.

Usually we shall not differentiate between square matrices and linear operators. For example, if A is an $m \times n$ -matrix, $\ker(A)$ is the set of vectors $X \in \mathbb{C}^n$ satisfying $AX = 0$ (when X is written as a column), and similarly $\text{Im}(A)$ is the set $\{AX \mid X \in \mathbb{C}^m\}$, which is the same as the span of the columns of A . The rank of A , defined as $\dim \text{Im}(A)$, can therefore be characterized as the maximum number of linear independent columns of A .

Recall that if a linear operator $F : V \rightarrow V$ has matrix A with respect to one basis, and matrix B with respect to a different basis, then $A = SBS^{-1}$, where S has the new basis vectors as columns expressed in the old basis. We say that two square matrices A and B are **similar** if there is a matrix S such that $A = SBS^{-1}$. Then two matrices are similar if they represent the same linear map $V \rightarrow V$ with respect to different choices of basis.

The **trace** of an $n \times n$ -matrix A is the sum of the diagonal entries:

$$\text{tr}(A) := \sum_{k=1}^n a_{kk}.$$

For example we have $\text{tr} \begin{pmatrix} 2+i & 3 \\ 4i & 3-4i \end{pmatrix} = 5 - 3i$.

Now let $A = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix}$ and $B = \begin{pmatrix} 1 & 1 \\ 1 & 3 \end{pmatrix}$. Then we have

$$\text{tr}(AB) = \text{tr} \begin{pmatrix} 2 & 3 \\ 2 & 5 \end{pmatrix} = 7 \text{ and } \text{tr}(BA) = \text{tr} \begin{pmatrix} 1 & 2 & 1 \\ 1 & 3 & 2 \\ 1 & 4 & 3 \end{pmatrix} = 7.$$

So $\text{tr}(AB) = 7 = \text{tr}(BA)$. This is no accident:

Proposition 2.1.1. *If both products AB and BA are defined, we have*

$$\text{tr}(AB) = \text{tr}(BA)$$

The proof is left as an exercise. An important corollary is that similar matrices have the same trace:

$$\text{tr}((S^{-1}B)S) = \text{tr}(S(S^{-1}B)) = \text{tr}(IB) = \text{tr}(B).$$

This lets us define the trace of a linear operator, as it is independent of the choice of basis. A further consequence of this is that if A is diagonalizable with $A = SDS^{-1}$, where D has the eigenvalues of A on the diagonal, then $\text{tr}(A) = \text{tr}(D)$ which is the *sum of the eigenvalues of A including multiplicities*.³

A square matrix⁴ is called **nilpotent** if $N^d = 0$ for some d , the minimal such d is called the nilpotency-degree of N . Nilpotent matrices will be important later when we investigate the Jordan normal form. The prototypical example of a nilpotent matrix is the matrix N below, we have:

$$N = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad N^2 = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad N^3 = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad N^4 = 0.$$

If we think of N as a linear operator, it acts on the standard basis like so:

$$e_4 \mapsto e_3 \mapsto e_2 \mapsto e_1 \mapsto 0.$$

The nilpotency-degree of N is 4, and it is clear that $N^m = 0$ for all $m \geq 4$. Note that if A is similar to N and $N^d = 0$, then $A^d = 0$ too, so the nilpotency-degree is basis independent.

³This also true for non-diagonalizable operators, we will prove this later.

⁴Nilpotency can be defined the exact same way when $N : V \rightarrow V$ is a linear operator.

2.2 Echelon forms

When solving a linear system by Gaussian elimination, we use row operations to reduce the coefficient matrix to a form suitable for writing down the solutions. The matrix $A \in \text{Mat}_{5 \times 6}(\mathbb{R})$ below is in *row echelon form*. With some further row operations we can reduce it to B which is in *reduced row echelon form*, the encircled elements are called *pivots*.

$$A = \begin{pmatrix} \textcircled{1} & 1 & 9 & 2 & 1 & 8 \\ 0 & \textcircled{2} & 4 & 1 & 2 & 5 \\ 0 & 0 & 0 & \textcircled{4} & 1 & 9 \\ 0 & 0 & 0 & 0 & \textcircled{3} & 3 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad B = \begin{pmatrix} \textcircled{1} & 0 & 7 & 0 & 0 & 2 \\ 0 & \textcircled{1} & 2 & 0 & 0 & 1 \\ 0 & 0 & 0 & \textcircled{1} & 0 & 2 \\ 0 & 0 & 0 & 0 & \textcircled{1} & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Definition 2.2.1. A matrix of $\text{Mat}_{m \times n}(F)$ is said to be in **row echelon form** (REF) if the first nonzero element of each row is to the *left* of the first nonzero elements in all the rows below, and if all zero-rows are at the bottom. The first nonzero elements of each row in the REF are called **pivots**. The matrix is said to be in **reduced row echelon form** (RREF) if additionally, all pivots are 1, and the pivots have zeros above them.

We know that any matrix can be reduced to REF and RREF by row operations. The RREF of a matrix is unique, and the number of the pivots in the RREF (and in the REF) is the **rank** of the matrix.

From the RREF we can immediately solve the corresponding linear system. For example, to solve $BX = 0$ for the matrix B above, we introduce parameters $x_3 = s$ and $x_6 = t$ for the variables corresponding to non-pivot columns, then from the RREF we immediately see that the set of solutions⁵ is

$$(x_1, x_2, x_3, x_4, x_5, x_6) = (-7s - 2t, -2s - t, s, -2t, -t, t) \quad s, t \in \mathbb{R}.$$

Since the matrices A and B are in fact row equivalent, the equation $AX = 0$ has the same solutions.

In the above example, we considered a matrix in $\text{Mat}_{5 \times 6}(\mathbb{R})$, but note that matrices, row operations, REF, and RREF makes sense over any field \mathbb{F} .

2.3 Elementary matrices

Row operations on a matrix can be performed by multiplying the matrix from the *left* by an **elementary matrix**. This is best explained by looking at some concrete examples:

$$E_1 A = \begin{pmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 1 & -1 \\ 2 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 1 & -1 \\ 0 & 1 & -2 & 3 \\ 1 & 1 & 1 & 1 \end{pmatrix}.$$

Here we note that multiplying A by the matrix E_1 on the left has the same effect as performing the row operation of *adding (-2) times the first row to the second row of A* .

Another row operation is to multiply one of the rows of A by a nonzero scalar λ , this can be achieved by multiplying by another type of matrix on the left, for example

$$E_2 A = \begin{pmatrix} 3 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 1 & -1 \\ 2 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} = \begin{pmatrix} 3 & 0 & 3 & -3 \\ 2 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix}.$$

The last type of row-operation is switching two rows:

$$E_3 A = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 & 1 & -1 \\ 2 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 1 & -1 \\ 1 & 1 & 1 & 1 \\ 2 & 1 & 0 & 1 \end{pmatrix}.$$

⁵Here one can ask if we should write $s, t \in \mathbb{C}$ instead, we still get solutions to the system for such s, t . In this case, the matrix was real, so we assumed that we were working over the real numbers.

In general, an elementary matrix is an $m \times m$ -matrix of one of the three forms below (empty positions are zeros). Multiplying such a matrix by an $m \times n$ -matrix A from the left has the effect of making a row operation on A .

Elementary matrix	Corresponding row operation
$\begin{pmatrix} 1 & & & & \\ & 1 & & & \\ & & 1 & & \\ & & \lambda & \ddots & \\ & & & & 1 \end{pmatrix} = I + \lambda e_{ij}$ <p>(λ in position (i, j))</p>	Add λ times row j to row i
$\begin{pmatrix} 1 & & & & \\ & \ddots & & & \\ & & \lambda & & \\ & & & \ddots & \\ & & & & 1 \end{pmatrix} = I + (\lambda - 1)e_{ii}$ <p>(identity except $\lambda \neq 0$ on position (i, i))</p>	Multiply row i by a nonzero scalar λ
$\begin{pmatrix} 1 & & & & \\ & \ddots & & & \\ & & 1 & & \\ & & & \ddots & \\ & & & & 1 \end{pmatrix}$ $= I - e_{ii} - e_{jj} + e_{ij} + e_{ji}$ <p>(I but with rows i and j switched)</p>	Switching rows i and j

Note that multiplying a matrix by an elementary matrix from the *right side* instead has the effect of performing a corresponding column-operation. This however is less useful, for example if we perform a column-operation on a linear system it no longer has the same solutions.

2.4 LU-decomposition

Definition 2.4.1. An **LU-decomposition** of an $m \times n$ matrix A is a factorization

$$A = LU$$

where L is a lower-triangular $m \times m$ -matrix, and U is an upper triangular $m \times n$ -matrix.

An LU-decomposition of a matrix A can typically be obtained by reducing A to row echelon form (REF) and keeping track of the elementary matrices corresponding to the row operations.

Example 2.4.2. We shall find an LU-decomposition $A = LU$ of the matrix A below. We start by row-reducing A to row echelon form U :

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 3 & 3 & -1 \\ -2 & 2 & -1 & -5 \end{pmatrix} \sim \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 2 & 2 & -2 \\ 0 & 4 & 1 & -3 \end{pmatrix} \sim \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 2 & 2 & -2 \\ 0 & 0 & -3 & 1 \end{pmatrix} = U.$$

We performed three row operations:

- Add (-1) times the first row to the second row
- Add (2) times the first row to the third row

- Add (-2) times the second row to the the third row

These row operations correspond to left multiplication by these elementary matrices:

$$E_1 = \begin{pmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad E_2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 2 & 0 & 1 \end{pmatrix} \quad E_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -2 & 1 \end{pmatrix}.$$

Thus we have $E_3(E_2(E_1A)) = (E_3E_2E_1)A = U$ so $A = (E_3E_2E_1)^{-1}U = (E_1^{-1}E_2^{-1}E_3^{-1})U = LU$, where

$$L = E_1^{-1}E_2^{-1}E_3^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -2 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 2 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ -2 & 2 & 1 \end{pmatrix}$$

We conclude that an LU-decomposition is given by

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 3 & 3 & -1 \\ -2 & 2 & -1 & -5 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ -2 & 2 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 2 & 2 & -2 \\ 0 & 0 & -3 & 1 \end{pmatrix} = LU.$$

△

Some questions remain - when do LU-decompositions exist, are they unique when they do, and what are they good for?

The reason that the method in the example works is that typically we can reduce a matrix to REF simply by adding multiples of rows to rows below them - the corresponding elementary matrices will be lower triangular, and then their inverses and their product L is also lower triangular. The only problem is if we need to switch two rows to reach an echelon form of A . In this case we can first perform a sequence of row switches in A by left-multiplying by elementary row-switching matrices. Let P be the product of these matrices⁶ Then we proceed with the LU-decomposition as usual to obtain a factorization⁷ $PA = LU$. We have proved the following theorem:

Theorem 2.4.3. *Each $m \times n$ matrix A admits a decomposition $PA = LU$ where*

- L is lower triangular $m \times m$ matrix with ones on the diagonal
- P is a permutation matrix of size $m \times m$
- U is an upper triangular $m \times n$ matrix (a row echelon form of A)

Is such a decomposition unique? Well more or less. First, consider our example above and factor the matrix U as DU' like so:

$$A = LU = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ -2 & 2 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 2 & 2 & -2 \\ 0 & 0 & -3 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ -2 & 2 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & -3 \end{pmatrix} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & -1 \\ 0 & 0 & 1 & -\frac{1}{3} \end{pmatrix} = LDU'.$$

This is called an **LDU-decomposition** of A , and it can clearly be found from the LU-decomposition as above.

Definition 2.4.4. An LDU-decomposition of A is a factorization $A = LDU$ where L is lower triangular $m \times m$ with ones on the diagonal, D is $m \times m$ and diagonal, and U is a row echelon form of A with ones as pivots.

Here we note that $A = (LD)U = L(DU)$ are two different LU-decompositions of A (unless $D = I$). So in general the LU-decomposition is not unique, however, with some additional conditions it is.

Proposition 2.4.5. *If an invertible $n \times n$ matrix A admits an LU-decomposition, then it is unique if we require that L has ones on the diagonal. It follows that it has a unique LDU-decomposition too.*

⁶A product of row-switching elementary matrices is called a *permutation matrix* - these can be characterized as having a single 1 in each row and each column, with zeros elsewhere.

⁷Equivalently, $A = P^{-1}LU$. Since P^{-1} is also a permutation matrix, some books prefer to write the factorization as $A = PLU$

Proof. Assume $L_1U_1 = L_2U_2$ are two LU-decompositions of A . Then with $L = L_1^{-1}L_2$ we have $U_1 = LU_2$, where U_1 and U_2 are both in row echelon form. But note that since A is invertible, any echelon form must be upper triangular with nonzero elements on the diagonal, but then LU can only itself be upper triangular if $L = I$, which means $L_1 = L_2$, and then since this matrix is invertible we also get $U_1 = U_2$. \square

What are LU-decompositions good for? Imagine that we have a large system of linear equations $Ax = b$ (let's say A is $n \times n$, while $b, x \in \mathbb{R}^n$). Assume we want to solve this system for many different right sides b , and perhaps at different times (or perhaps one b is used to calculate the next recursively). Then if we have a factorization $A = LU$, we have

$$Ax = b \Leftrightarrow L(Ux) = b \Leftrightarrow Ly = b \text{ and } Ux = y.$$

So instead of solving $Ax = b$ directly we can solve the two systems $Ly = b$ and then $Ux = y$. These systems are triangular, so they are fast to solve by back-substitution. If this is done by a computer on a very large matrix (say $n = 10^4$), the speed-increase is significant.⁸

2.5 Cholesky-factorization

Definition 2.5.1. A **Cholesky-factorization** of a real or complex square matrix A is of form

$$A = CC^*$$

where C is a lower triangular matrix.

Since CC^* is always Hermitian, so must A be in order to admit such a factorization. In case an invertible Hermitian matrix A has a decomposition $A = LDU$, we note that

$LDU = A = A^* = U^*D^*L^*$ are two LDU-decompositions, so by uniqueness $L = U^*$ and $D = D^*$ so $A = LDL^*$ and D is real. Moreover, assume that D has positive entries

$$D = \begin{pmatrix} d_1 & & & \\ & d_2 & & \\ & & \ddots & \\ & & & d_n \end{pmatrix} \text{ and define } \sqrt{D} = \begin{pmatrix} \sqrt{d_1} & & & \\ & \sqrt{d_2} & & \\ & & \ddots & \\ & & & \sqrt{d_n} \end{pmatrix}.$$

The reason for the notation is that $(\sqrt{D})^2 = D$. Now take $C = L\sqrt{D}$. Then C is lower triangular, and

$$CC^* = L\sqrt{D}\sqrt{D}^*L^* = L\sqrt{D}\sqrt{D}L^* = LDL^* = A,$$

so $A = CC^*$ is the⁹ Cholesky-factorization of A .

The Cholesky-factorization can be used when solving linear systems with a Hermitian coefficient matrix - algorithmically it is twice as efficient as using the LU-decomposition.

Example 2.5.2. Let's find the Cholesky-factorization of $A = \begin{pmatrix} 2 & 4 \\ 4 & 12 \end{pmatrix}$. We can reduce it to REF by a single row-operation:

$$EA = \begin{pmatrix} 1 & 0 \\ -2 & 1 \end{pmatrix} \begin{pmatrix} 2 & 4 \\ 4 & 12 \end{pmatrix} = \begin{pmatrix} 2 & 4 \\ 0 & 4 \end{pmatrix} \text{ so } A = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix} \begin{pmatrix} 2 & 4 \\ 0 & 4 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix} \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix} \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} = LDU$$

is the LDU-decomposition of A . Now take

$$C = L\sqrt{D} = \begin{pmatrix} 1 & 0 \\ -2 & 1 \end{pmatrix} \begin{pmatrix} \sqrt{2} & 0 \\ 0 & 2 \end{pmatrix} = \begin{pmatrix} \sqrt{2} & 0 \\ 2\sqrt{2} & 2 \end{pmatrix},$$

⁸But we need to find the LU-decomposition first, so for a single right side b it is not an effective method.

⁹One can show that the Cholesky-factorization exists for Hermitian matrices A that are positive semidefinite (meaning that all its eigenvalues are ≥ 0). If the matrix is positive definite the Cholesky-factorization is additionally unique.

then it is easy to verify that

$$CC^* = \begin{pmatrix} \sqrt{2} & 0 \\ 2\sqrt{2} & 2 \end{pmatrix} \begin{pmatrix} \sqrt{2} & 2\sqrt{2} \\ 0 & 2 \end{pmatrix} = \begin{pmatrix} 2 & 4 \\ 4 & 12 \end{pmatrix} = A$$

is the Cholesky-decomposition of A .

△

2.6 Determinants

Recall from a first linear algebra course that if $A \in \text{Mat}_n(\mathbb{R})$, the **determinant** of A is a real number that could be calculated by several methods:

- Sarrus' rule for 2×2 and 3×3 matrices
- Expansion along rows or columns
- Make row or column operations until the matrix is upper- or lower-triangular, then the determinant is the product of the diagonal entries¹⁰
- As a sum over the permutation group S_n (maybe not in a first course)

The determinant also has a number of important properties:

- $\det(I) = 1$
- $\det(A) = \det(A^T)$
- $\det(AB) = \det(A)\det(B)$
- $\det(A) \neq 0 \Leftrightarrow A^{-1}$ exists and each system $AX = 0$ has a unique solution

It follows from the second point that that $\det(S^{-1}AS) = \det(A)$, so the determinant is basis-independent and can be defined for any linear map. If a linear map is diagonalizable, its determinant is therefore the same as the determinant of $D = \text{diag}(\lambda_1, \dots, \lambda_n)$, in other words, $\det(A) = \lambda_1 \cdots \lambda_n$, the determinant is the product of all eigenvalues counting multiplicities. This also holds for non-diagonalizable maps which will see later.

For now we just remark that all these rules and properties above work the same over any field \mathbb{F} . When $\mathbb{F} = \mathbb{C}$ we also note that $\det(\overline{A}) = \overline{\det(A)}$, and therefore $\det(A^*) = \overline{\det(A)}$.

Example 2.6.1. The determinant of the linear map $\mathbb{C}^2 \rightarrow \mathbb{C}^2$ with standard matrix $A = \begin{pmatrix} 1+i & 1 \\ i & 3 \end{pmatrix}$ is $\det(A) = 3(1+i) - i = 3 + 2i \neq 0$, so the map is invertible.

△

Example 2.6.2. Let $A = \begin{bmatrix} \spadesuit & \heartsuit \\ \clubsuit & \diamonds \end{bmatrix} \in \text{Mat}_2(F)$ where F is the field from Example 9.1.9 of the appendix. Then

$$\det(A) = \spadesuit \cdot \diamonds - \clubsuit \cdot \heartsuit = \heartsuit - \clubsuit = \heartsuit + \clubsuit = \heartsuit.$$

The zero-element of F is \clubsuit , so since $\det(A) \neq \clubsuit$, the matrix A is invertible in $\text{Mat}_2(F)$.

△

¹⁰Remember that changing two rows negates the determinant, and multiplying a row or column by a λ changes the determinant by a factor λ^n

2.7 Block matrices

It is sometimes useful to split a matrix into **blocks**, and consider it as a matrix where the coefficients are themselves matrices.

For example, the matrix

$$X = \left(\begin{array}{cc|cc} 3 & 2 & 3 & 0 \\ 1 & 1 & 0 & 3 \\ \hline 0 & 0 & 6 & 4 \\ 0 & 0 & 2 & 2 \end{array} \right) \in \text{Mat}_4(\mathbb{R})$$

can be considered as a **block matrix** in $\text{Mat}_2(\text{Mat}_2(\mathbb{R}))$, a 2×2 -matrix where the coefficient themselves are 2×2 -matrices¹¹: For the matrix X above we have

$$X = \left(\begin{array}{c|c} A & 3I \\ \hline 0 & 2A \end{array} \right) \quad \text{where} \quad A = \begin{pmatrix} 3 & 2 \\ 1 & 1 \end{pmatrix}.$$

Block-matrices multiply according to normal matrix multiplication as long as the blocks are compatible in the sense that all products and sums of blocks are defined. For example, we have

$$X^2 = \left(\begin{array}{c|c} A & 3I \\ \hline 0 & 2A \end{array} \right) \left(\begin{array}{c|c} A & 3I \\ \hline 0 & 2A \end{array} \right) = \left(\begin{array}{cc|cc} A \cdot A + 3I \cdot 0 & A \cdot 3I + 3I \cdot 2A \\ \hline 0 \cdot A + 2A \cdot 0 & 0 \cdot 3I + 2A \cdot 2A \end{array} \right) = \left(\begin{array}{c|c} A^2 & 9A \\ \hline 0 & 4A^2 \end{array} \right) = \left(\begin{array}{cc|cc} 11 & 8 & 27 & 18 \\ 4 & 3 & 9 & 9 \\ \hline 0 & 0 & 44 & 32 \\ 0 & 0 & 16 & 12 \end{array} \right).$$

It is also possible to express the inverse of X as a block matrix: we claim that the inverse of X above can be expressed in block-form as $X^{-1} = \left(\begin{array}{c|c} A^{-1} & -\frac{3}{2}A^{-2} \\ \hline 0 & \frac{1}{2}A^{-1} \end{array} \right)$, this is easy to verify by multiplying these matrices in block form to show that $X \cdot X^{-1} = \left(\begin{array}{c|c} I & 0 \\ \hline 0 & I \end{array} \right) = I$.

A common example of block matrices are the **block-diagonal** matrices. For example, if A, B, C are square matrices we say that the matrix $\text{diag}(A, B, C)$ is a block-diagonal matrix. It is easy to exponentiate such block-diagonal matrices:

$$D = \text{diag}(A, B, C) := \begin{pmatrix} A & & \\ & B & \\ & & C \end{pmatrix} \implies D^n = \begin{pmatrix} A & & \\ & B & \\ & & C \end{pmatrix}^n = \begin{pmatrix} A^n & & \\ & B^n & \\ & & C^n \end{pmatrix}.$$

Note that this works even when A, B, C are square matrices of different size; this will be useful in Chapter 4 where we investigate Jordan forms of matrices.

We illustrate another useful example where the blocks have different size, this comes from computer graphics where one wants to represent affine-transformations: linear transformations composed with translations.

Example 2.7.1. An affine map $F : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a map of form $F(v) = Av + w$ where A is a 2×2 -matrix and w is a fixed vector written as a 2×1 -matrix. So F is a linear map composed with a translation, note that F is not a linear map when $w \neq 0$. In this example, let $A = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$ and $w = \begin{pmatrix} 3 \\ 5 \end{pmatrix}$ such that F is a counter-clockwise quarter-rotation around the origin followed by a translation by the vector $\begin{pmatrix} 3 \\ 5 \end{pmatrix}$.

All our methods and results from linear algebra are for linear maps, so in order for this theory to be useful we would like to represent F as some linear operator. To do this, let

$$M = \left(\begin{array}{c|c} A & w \\ \hline 0 & 1 \end{array} \right) = \left(\begin{array}{cc|c} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} & 3 \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 5 \\ \hline 0 & 0 & 1 \end{array} \right).$$

Now let us decide to always embed vectors $v \in \mathbb{R}^2$ into \mathbb{R}^3 by setting the third coordinate to 1: for

¹¹There is a technical difference between $\text{Mat}_2(\text{Mat}_2(\mathbb{R}))$ and $\text{Mat}_4(\mathbb{R})$, but ignoring parentheses there is a natural identification.

$v \in \mathbb{R}^2$ define $\bar{v} := \begin{pmatrix} v \\ 1 \end{pmatrix} \in \mathbb{R}^3$. Then by computing the block-matrix product

$$M\bar{v} = \left(\begin{array}{c|c} A & w \\ \hline 0 & 1 \end{array} \right) \begin{pmatrix} v \\ 1 \end{pmatrix} = \begin{pmatrix} Av + w \\ 1 \end{pmatrix} = \overline{F(v)}$$

we see that the affine map F is represented by the linear map M when vectors are embedded via $v \leftrightarrow \bar{v}$. Using this correspondence we can use the standard tools of linear algebra to study affine maps, for example, composition of affine maps will correspond to products of the corresponding 3×3 -matrices. The analogous construction works for affine maps on \mathbb{R}^3 , or more generally on \mathbb{F}^n .

△

Chapter 3

Introductory spectral theory

3.1 Eigenvalues and eigenvectors

We recap the theory of eigenvalues and eigenvectors from a first linear algebra course. In this chapter though, unless otherwise stated, all vector spaces are defined over an arbitrary field of scalars \mathbb{F} .

Definition 3.1.1. Let $F : V \rightarrow V$ be a linear map on a vector space over a field \mathbb{F} . If

$$F(v) = \lambda v \text{ for some } \lambda \in \mathbb{F} \text{ and some nonzero } v \in V,$$

we say that λ is an **eigenvalue** for F , and v is a corresponding **eigenvector**.

When $\dim V < \infty$, the map F may be described by a matrix A after a basis is picked, so we shall also speak of eigenvalues and eigenvectors of *matrices*, the condition in the definition then looks like $Av = \lambda v$.

The eigenvalues typically give important information about the geometric nature of a map as the following example illustrates.

Example 3.1.2. Let $P : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be given by $[P] = \frac{1}{14} \begin{pmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \\ 3 & 6 & 9 \end{pmatrix}$ in the standard basis. Then by the standard methods of linear algebra we may find its eigenvalues and eigenvectors. It turns out that P has two eigenvalues: 0 and 1. Every nonzero vector in the plane $x + 2y + 3z = 0$ is an eigenvector for the eigenvalue 0, and every nonzero vector on the line $t(1, 2, 3)$ is an eigenvector for the eigenvalue 1. From this information we can deduce that geometrically, F is the orthogonal projection onto the line $t(1, 2, 3)$.

△

The definition of eigenvalues also makes sense for operators on infinite-dimensional vector spaces, but note that our standard method for finding the eigenvalues does not work here.

Example 3.1.3. Let $C^\infty(\mathbb{R})$ be the real vector space of infinitely differentiable functions $\mathbb{R} \rightarrow \mathbb{R}$, and let D be the linear operator on this space that takes the derivative, $D(f(x)) = f'(x)$. Then every $\lambda \in \mathbb{R}$ is an eigenvalue for D , the set of corresponding eigenvectors are nonzero multiples of $e^{\lambda x}$.

△

The **spectrum** of a linear operator F is the set of eigenvalues, and it is written $\sigma(F)$. The spectrum gives important qualitative information about the operator. In the above examples we have $\sigma(P) = \{0, 1\}$ and $\sigma(D) = \mathbb{R}$. We shall focus on the case when the dimension of the vector space is finite, then there is a concrete standard method for finding the eigenvalues and eigenvectors: fix a basis for V and consider the matrix A of the linear operator. Then λ is an eigenvalue and a nonzero v is an eigenvector if and only if $Av = \lambda v$, or equivalently $(A - \lambda I)v = 0$. This matrix-equation has a nontrivial solution v if and only if $\det(A - \lambda I) = 0$. Solving this equation¹ gives us the eigenvalues, and for each eigenvalue λ we

¹Two remarks: When the field $\mathbb{F} = \mathbb{C}$, solving a polynomial equation of degree ≥ 5 algebraically is not feasible in general, but there are good numeric methods for finding a good approximations, so this is not a problem in applications. For other \mathbb{F} such as finite fields, the polynomial $\det(A - \lambda I)$ may not have any solutions $\lambda \in \mathbb{F}$, and hence may not have any eigenvalues.

can then solve the linear system $(A - \lambda I)v = 0$ to find the corresponding eigenvectors.

For $\lambda \in \mathbb{F}$, we define the corresponding **eigenspace**, to be $E_\lambda := \ker(A - \lambda I)$. This is the subspace of V consisting of all vectors² v satisfying $Av = \lambda v$. Note that $\ker(A - \lambda I) = \{0\}$ when λ is not an eigenvalue.

Now let A be a matrix representing a linear operator on a finite-dimensional vector space over a field \mathbb{F} . The **characteristic polynomial**³ for A is defined as

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

Then $p_A(\lambda) = 0$ if and only if $\lambda \in \mathbb{F}$ is an eigenvalue of A . Note that the coefficients of p_A lie in \mathbb{F} , and that $\deg p_A = \dim V$. A quick calculation

$$\det(SAS^{-1} - \lambda I) = \det(S(A - \lambda I)S^{-1}) = \det(S)\det(A - \lambda I)\det(S^{-1}) = \det(SS^{-1})\det(A - \lambda I) = \det(A - \lambda I)$$

shows that the characteristic polynomial is the same regardless of the choice of basis in our matrix-representation and we can therefore speak about $p_F(\lambda)$, the characteristic polynomial of a linear map $F : V \rightarrow V$ without specifying a basis.

The **algebraic multiplicity** m_λ of an eigenvalue λ is the multiplicity of λ as a zero in $p_A(t)$, in other words we can factor the characteristic polynomial as

$$p_A(t) = (t - \lambda)^{m_\lambda} q(t) \quad \text{where} \quad q(\lambda) \neq 0.$$

On the other hand, the **geometric multiplicity** g_λ of λ is defined as the dimension of the λ -eigenspace:

$$g_\lambda = \dim E_\lambda = \dim \ker(A - \lambda I).$$

Typically we expect the geometric and algebraic multiplicities to coincide for each eigenvalue of an operator. This was the case for the operator $P : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ from Example 3.1.2 above where we had $g_0 = 2 = m_0$ and $g_1 = 1 = m_1$. But there are exceptions:

Example 3.1.4. Let $A = \begin{pmatrix} 3 & 5 \\ 0 & 3 \end{pmatrix}$. The characteristic polynomial is $p_A(t) = (t-3)^2$, so $\lambda = 3$ is an eigenvalue with algebraic multiplicity 2 since $t = 3$ is a double zero for $p_A(t)$. But solving $Av = 3v$ we notice that the only eigenvectors are multiples of $(1, 0)$, so $\ker(A - 3I) = \text{span}((1, 0))$ is 1-dimensional, so the geometric multiplicity of $\lambda = 3$ is 1. In other words, $1 = g_3 < m_3 = 2$.

△

Proposition 3.1.5. *Let λ be an eigenvalue of an operator $F : V \rightarrow V$ where V is finite-dimensional. Then the geometric multiplicity of λ is less than the algebraic multiplicity of λ :*

$$g_\lambda \leq m_\lambda.$$

Proof. Let V be n -dimensional and assume the geometric multiplicity of λ is $g_\lambda = m$. Then we can pick a basis of m vectors (v_1, \dots, v_m) in the eigenspace E_λ and extend this to a basis $\mathcal{B} = (v_1, \dots, v_n)$ of V . With respect to this latter basis, the matrix of F has block form

$$A = [F]_{\mathcal{B}} = \left(\begin{array}{ccc|cc} \lambda & & & & \\ & \ddots & & & \\ & & \lambda & & B \\ \hline & & & 0 & C \end{array} \right)$$

where the top-left block is of size $m \times m$.

Since the characteristic polynomial is independent of the choice of basis we have

$$p_F(t) = p_A(t) = \det(A - tI) = (\lambda - t)^m \det(C - tI) = (\lambda - t)^m p_C(t),$$

which was obtained by expanding the determinant along each of the first m columns. This shows that $(\lambda - t)^m$ divides $p_F(t)$, so the algebraic multiplicity of λ is at least m . This completes the proof. □

²Note that $\ker(A - \lambda I)$ also contains the zero-vector, so technically it consists of all eigenvectors of eigenvalue λ and the zero vector.

³Also called the secular polynomial. Some books define it instead as $\ker(\lambda I - A)$, this is the same up to a sign change. Sometimes a different variable is used in the polynomial, for example $p_A(t) = \det(A - tI)$. The variable-name is unimportant, $p_A(\lambda)$ and $p_A(t)$ is the same *function*. Note that $p_A(t) \in \mathbb{F}[t]$; the characteristic polynomial has coefficients in the field \mathbb{F} .

Note however that if λ is an eigenvalue, g_λ is at least 1 since there is always at least one eigenvector.

If a matrix A can be factored as $A = SDS^{-1}$ where D is a diagonal matrix, we say that A is **diagonalizable**. This is equivalent to saying that there is a basis for V consisting of eigenvectors for A . This means that the characteristic polynomial factors completely into linear factors, and that the algebraic and geometric multiplicities agree for all eigenvalues. Concretely, let D be a diagonal matrix with the eigenvalues $\lambda_1, \dots, \lambda_n$ on the diagonal (repeating according to algebraic multiplicities), and pick a basis v_1, \dots, v_n for V consisting of eigenvectors ordered in the same order as the eigenvalues, such that $Av_i = \lambda_i v_i$. Let $S = (v_1 \cdots v_n)$ be the matrix with these eigenvectors as columns. Then $S^{-1}AS = D$, or equivalently $SDS^{-1} = A$. A factorization like this is called a diagonalization of A , it corresponds to making a change of basis so that the new basis vectors are eigenvectors. This just means that geometrically, any diagonalizable linear map just stretches vectors with different factors along a number of axes.

One main difference between real and complex vector spaces when it comes to spectral theory is:

Theorem 3.1.6. *Each linear operator on a finite-dimensional complex vector space has an eigenvalue.*

Proof. By the fundamental theorem of algebra, every nonconstant polynomial with complex coefficients has a zero in \mathbb{C} . A zero of the characteristic polynomial is an eigenvalue. \square

Example 3.1.7. Let's diagonalize the linear map $F : \mathbb{C}^2 \rightarrow \mathbb{C}^2$ with matrix $[F] = A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$. Since $p_A(\lambda) = \lambda^2 + 1$, the eigenvalues are $\pm i$. To find the eigenvectors for $\lambda = i$ we solve the system $(A - iI)x = 0$:

$$\begin{cases} -ix_1 & - x_2 = 0 \\ x_1 & -ix_2 = 0 \end{cases} \Leftrightarrow x_1 - ix_2 = 0 \Leftrightarrow (x_1, x_2) = t(i, 1) \text{ where } t \in \mathbb{C}.$$

Similarly we get $(A + iI)x = 0 \Leftrightarrow x = t(1, i)$. We put the eigenvalues in a matrix D and corresponding eigenvectors as columns in a matrix S . We now have a diagonalization of A :

$$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} = A = SDS^{-1} = \begin{pmatrix} i & 1 \\ 1 & i \end{pmatrix} \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} \frac{1}{2} \begin{pmatrix} -i & 1 \\ 1 & -i \end{pmatrix}.$$

Note that the same matrix A can be viewed as the matrix of a linear map $G : \mathbb{R}^2 \rightarrow \mathbb{R}^2$. This map G does *not* have any eigenvalues, since $p_A(\lambda) = \lambda^2 + 1$ has no zeroes in \mathbb{R} . Geometrically the map G corresponds to a rotation a quarter of a turn counter-clockwise in \mathbb{R}^2 .

\triangle

The same algorithms work for attempting to diagonalize a matrix over an arbitrary field:

Example 3.1.8. Let $A = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix} \in \text{Mat}_2(\mathbb{Z}_2)$ be a matrix with coefficients in the finite field \mathbb{Z}_2 (see Example 9.1.5 in the appendix for details), it represents a linear map $(\mathbb{Z}_2)^2 \rightarrow (\mathbb{Z}_2)^2$. The standard method gives $p_A(t) = (t - 1)t$ so both 0 and 1 is an eigenvalue (all elements of the field). The standard computation (but over \mathbb{Z}^2) shows that $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ is an eigenvector of eigenvalue 1 and that $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ is an eigenvector of eigenvalue 0. So A is diagonalizable as $A = SDS^{-1}$ where

$$S = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad D = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

We remark that there is an alternative way to find the eigenvectors without finding the characteristic polynomial or solving any matrix equations: Since $(\mathbb{Z}_2)^2$ only has four elements we can multiply A by each of the three nonzero vectors and see which, if any, are eigenvectors. We immediately get

$$A \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad A \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad A \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

so $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ is an eigenvector with eigenvalue 1 and that $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ is an eigenvector with eigenvalue 0.

△

The following facts about eigenvectors will be used in later chapters:

Proposition 3.1.9. *Eigenvectors corresponding to different eigenvalues are linearly independent.*

Proof. We prove the statement by contradiction. Let A be a matrix representing an operator, and let v_1, \dots, v_n be eigenvectors with $Av_i = \lambda_i v_i$. Suppose that the eigenvalues λ_i are distinct and that they satisfy linear dependence relation

$$\alpha_1 v_1 + \dots + \alpha_n v_n = 0,$$

where we may assume that all $\alpha_i \neq 0$ (otherwise just remove the corresponding vectors v_i).

We may also assume that n is minimal, such that there is no dependence relation of eigenvectors of A with fewer vectors involved. Then we apply $(A - \lambda_1 I)$ to the relation and obtain

$$(A - \lambda_1 I)(\alpha_1 v_1 + \dots + \alpha_n v_n) = \alpha_1(A - \lambda_1 I)v_1 + \dots + \alpha_n(A - \lambda_1 I)v_n = \alpha_1(\lambda_1 - \lambda_1)v_1 + \dots + \alpha_n(\lambda_n - \lambda_1)v_n = 0.$$

But we have now produced another nontrivial linear dependence but with exactly one less term since the coefficient of v_1 is now zero. This contradicts the minimality of n , and finishes the proof. The argument still works if we replace the matrix A by an arbitrary linear operator F . □

3.2 Invariant subspaces

Let $F : V \rightarrow V$ be a linear operator, and let U be a subspace of V . We say that the subspace U is **F -invariant**, or invariant under F , if

$$u \in U \Rightarrow F(u) \in U.$$

In other words, we "stay in the subspace" if we apply F to a vector in the subspace. Since matrices describe linear maps we shall also talk about invariant subspaces for matrices.

Let $F : V \rightarrow W$ be a linear map, and let $U \subset V$ be a subspace. The **restriction** $F|_U$ of F to U is the map

$$F|_U : U \rightarrow W \quad \text{where } F|_U(u) = F(u),$$

in other words, $F|_U$ is the same map as F but with domain restricted to U .

But if $F : V \rightarrow V$ is an operator and $U \subset V$ is a subspace that is *invariant* under a map F , then we can its restriction as a map from U to itself⁴.

An eigenspace for a linear map F is clearly an F -invariant subspace, and so is a sum of eigenspaces. But there are also other examples as illustrated in the following example:

Example 3.2.1. Let $F : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be the rotation a quarter of a turn in the positive direction around the axis $(1, 2, 2)$. Then $U_1 = \text{span}(1, 2, 2)$ is a one-dimensional F -invariant subspace, and the restriction $F|_{U_1} : U_1 \rightarrow U_1$ is the identity map since every vector in U_1 is mapped to itself by F . In the basis $((1, 2, 2))$ of U_1 , the matrix of $F|_{U_1}$ is the 1×1 -matrix (1) .

Similarly, the rotation-plane $U_2 : x + 2y + 2z = 0$ is a two-dimensional F -invariant subspace since every vector of U_2 is mapped to another vector in U_2 when it is rotated around $(1, 2, 2)$. If we pick $((2, -2, 1), (2, 1, -2))$ as a basis for U_2 , the matrix of $F|_{U_2}$ is $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$. Note that U_2 contains no eigenvectors of F . Here we have

$$\mathbb{R}^3 = U_1 \oplus U_2,$$

a nice decomposition of \mathbb{R}^3 with respect to our map F , our vector space has decomposed into two F -invariant components.

△

Invariant subspaces also make sense for maps on infinite-dimensional vector spaces:

Example 3.2.2. Let \mathcal{P} be the vector space of all polynomials with real coefficients. The differential operator $D : \mathcal{P} \rightarrow \mathcal{P}$ is given by $D(p(x)) = p'(x)$. Then \mathcal{P}_n , the polynomials of degree $\leq n$, is a D -

⁴Technically, to get a map $U \rightarrow U$ we should compose $F|_U$ with the projection onto the subspace U , and form $P_U \circ F|_U$, but it is standard to just write $F|_U$ when the context is understood.

| invariant subspace for each n .

△

We know that many maps can be represented by diagonal matrices. What about maps that can't? Before we prove the Jordan theorem in Chapter 4 we need the following preliminary result:

Theorem 3.2.3. *Let $F : V \rightarrow V$ be linear operator on a finite-dimensional complex vector space V . Then there exists a basis for V for which the matrix of F is upper triangular. The elements on the diagonal in this triangular matrix are the eigenvalues for F , counting multiplicities.*

Proof. We proceed by induction on $\dim V$, the statement is trivial if $\dim V = 1$ since every 1×1 -matrix is upper triangular. Let $\dim V = n$ and assume the statement is true for all vector spaces of dimension $< n$. By Theorem 3.1.6 we know that F has an eigenvalue, so assume $F(u_1) = \lambda u_1$. Pick any $(n - 1)$ -dimensional subspace U that doesn't contain u_1 ; then $V = \text{span}(u_1) \oplus U$. Pick a basis (u_2, \dots, u_n) of U . Then with respect to the basis (u_1, \dots, u_n) of V , the matrix of F looks like

$$[F] = \left(\begin{array}{c|ccc} \lambda & b_2 & \cdots & b_n \\ \hline 0 & & & \\ \vdots & & A & \\ 0 & & & \end{array} \right)$$

where A is the matrix for the linear map $\pi_U(F|_U) : U \rightarrow U$, the projection onto U of the restriction $F|_U$. By the induction hypothesis, for this operator there exists a basis for U for which A becomes upper triangular. Together with the first basis vector u_1 , we get a basis for V for which the matrix for F is upper triangular. Let $T = (t_{ij})$ be this upper triangular matrix, then since the determinant of a triangular matrix is the product of the diagonal elements, we have $\det(T - \lambda I) = (t_{11} - \lambda) \cdots (t_{nn} - \lambda)$ is the characteristic polynomial of T , and therefore of F , and its roots t_{11}, \dots, t_{nn} are the diagonal elements of T , so these are the eigenvalues of F . □

Corollary 3.2.4. *Let A be a square matrix with complex coefficients with eigenvalues $\lambda_1, \dots, \lambda_n$ (repeating according to algebraic multiplicity). Then the trace of A is the sum of the eigenvalues, and the determinant of A is the product of the eigenvalues:*

$$\text{tr}(A) = \lambda_1 + \dots + \lambda_n \quad \text{and} \quad \det(A) = \lambda_1 \cdots \lambda_n.$$

Proof. We know that trace and determinant is independent of the choice of basis, so use Theorem 3.2.3 to find a basis for which the matrix of A is upper triangular with the eigenvalues on the diagonal. Then the result follows. □

This result is usually proved in a first course only for diagonalizable operators, now we see that it holds in general for operators on complex vector spaces.

3.3 Matrix polynomials

There is a natural way to "plug in" a square matrix into a polynomial.

Definition 3.3.1. Let $p(t) = a_n t^n + \dots + a_1 t + a_0$ be a polynomial with coefficients $a_i \in \mathbb{F}$, and let A be a square matrix with coefficients in the same field \mathbb{F} . Then we define

$$p(A) = a_n A^n + \dots + a_1 A + a_0 I.$$

Example 3.3.2. Let $p(t) = t^7 + 12t^4 - t^3 + 2t^2 + 5t + 3$ and $N = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}$. Then

$$p(N) = N^7 + 12N^4 - N^3 + 2N^2 + 5N + 3I = 2N^2 + 5N + 3I = \begin{pmatrix} 3 & 5 & 2 \\ 0 & 3 & 5 \\ 0 & 0 & 3 \end{pmatrix}$$

here the calculation was made easy by the fact that N is nilpotent, $N^k = 0$ for $k > 2$, so the first terms disappeared.

△

Note that $p(A)$ is a square matrix of the same size as A , and that $P(A)$ commutes with A since each term does. Moreover, if $B = SAS^{-1}$, we get $p(B) = Sp(A)S^{-1}$. For this reason it also makes sense to define $p(F)$ where $F : V \rightarrow V$ is an operator of an \mathbb{F} -vector space.

For a given $n \times n$ -matrix A , can we always find a nonzero polynomial $p(t)$ such that $p(A) = 0$? Yes - consider the matrices $I, A, A^2, \dots, A^{n^2}$. These are $n^2 + 1$ vectors in an n^2 -dimensional vector space $\text{Mat}_n(\mathbb{F})$, therefore they are linearly dependent and there are scalars λ_k such that

$$\sum_{k=0}^{n^2} \lambda_k A^k = 0$$

so with $p(t) = \sum_{k=0}^{n^2} \lambda_k t^k$ we have $p(A) = 0$.

Can we find find such a polynomial of smaller degree?

Example 3.3.3. Let $p(t) = t^2 - 5t - 2$ and $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$. Then

$$p(A) = A^2 - 5A - 2I = \begin{pmatrix} 7 & 10 \\ 15 & 22 \end{pmatrix} - 5 \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix} - 2 \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

So in this case $p(A)$ is the zero matrix. Note that the characteristic polynomial of A is $\det(A - \lambda I) = \lambda^2 - 5\lambda - 2$.

△

The example above shows an example where a matrix satisfies its characteristic equation: $p_A(A) = 0$. This is in fact always true, it's a famous result of linear algebra:

Theorem 3.3.4. (The Cayley-Hamilton theorem) *If A is a square matrix with characteristic polynomial $p_A(\lambda) = \det(A - \lambda I)$, then*

$$p_A(A) = 0.$$

The theorem holds over any field, but we prove it only for the complex numbers below:

Proof. Use Theorem 3.2.3 to find a matrix S and an upper triangular matrix $T = (t_{ij})$ (with the eigenvalues $\lambda_1, \dots, \lambda_n$ of A on the diagonal) such that $S^{-1}AS = T$. We have previously shown that similar matrices have the same characteristic polynomial, so $p_A(\lambda) = p_T(\lambda)$, and

$$p_A(A) = p_T(A) = p_T(STS^{-1}) = Sp_T(T)S^{-1}$$

so it suffices to show that $p_T(T) = 0$ for an upper triangular matrix T .

Now factor $p_T(\lambda)$ completely over \mathbb{C} , we know that the eigenvalues of T are its diagonal entries, so

$$p_T(\lambda) = (t_{11} - \lambda) \cdots (t_{nn} - \lambda).$$

Then $p_T(T) = (t_{11}I - T) \cdots (t_{nn}I - T)$, and by taking successive products from the right we see that $p_T(T)v = 0$ for every vector v , as illustrated below when $n = 3$:

$$\begin{aligned} p_T(T)v &= (t_{11}I - T)(t_{22}I - T)(t_{33}I - T)v = \begin{pmatrix} 0 & * & * \\ 0 & * & * \\ 0 & 0 & * \end{pmatrix} \begin{pmatrix} * & * & * \\ 0 & 0 & * \\ 0 & 0 & * \end{pmatrix} \begin{pmatrix} * & * & * \\ 0 & * & * \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} \\ &= \begin{pmatrix} 0 & * & * \\ 0 & * & * \\ 0 & 0 & * \end{pmatrix} \begin{pmatrix} * & * & * \\ 0 & 0 & * \\ 0 & 0 & * \end{pmatrix} \begin{pmatrix} * \\ * \\ 0 \end{pmatrix} = \begin{pmatrix} 0 & * & * \\ 0 & * & * \\ 0 & 0 & * \end{pmatrix} \begin{pmatrix} * \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} \end{aligned}$$

where in the calculation above, $*$ means some arbitrary number. Each multiplication gives one more zero in the resulting vector. This shows that $p_T(T)v = 0$ for all v , which means that T is the zero matrix. In light of the above remarks, this completes the proof. □

Example 3.3.5. Let $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$. Let us find a polynomial p such that $p(A) = A^{-1}$. Since $p_A(\lambda) = \lambda^2 - 5\lambda - 2$ we know by Cayley-Hamilton that $p_A(A) = 0$, so solving for I we get

$$A^2 - 5A - 2I = 0 \Leftrightarrow I = \frac{1}{2}(A - 5I)A \Leftrightarrow A^{-1} = \frac{1}{2}(A - 5I).$$

This idea works in general whenever A is invertible, because then 0 is not an eigenvalue, and the constant term in p_A is nonzero.

△

The Cayley-Hamilton also provides a quicker way to evaluate a polynomial of high degree at A without computing many matrix products. For example, in our previous example we had $p_A(t) = t^2 - 5t - 2$, so to evaluate $p(A)$ for some complicated polynomial $p(t)$, first use standard polynomial division to find polynomials $q(t)$ and $r(t)$ with $\deg r(t) < 2$ such that

$$p(t) = p_A(t)q(t) + r(t).$$

Then by Cayley-Hamilton we have

$$p(A) = p_A(A)q(A) + r(A) = 0 \cdot q(A) + r(A) = r(A),$$

so we only have to evaluate $r(A)$ where r has degree ≤ 1 .

Minimal polynomial

Let A be an $n \times n$ matrix. Cayley-Hamilton says that there exists a polynomial of degree n (the characteristic polynomial) satisfying $p(A) = 0$, but a polynomial of even lower degree may exist.

Definition 3.3.6. Let A be a square matrix. The **minimal polynomial** of A is the monic polynomial $m_A(t)$ of lowest degree for which $m_A(A) = 0$.

One can also define the minimal polynomial of any operator on a finite-dimensional vector space. The adjective *monic* just means that the coefficient of the highest degree term is 1.

Proposition 3.3.7. *The minimal polynomial exists and is unique.*

Proof. Cayley-Hamilton shows that there exists some monic polynomial p annihilating A (meaning $p(A) = 0$), namely $\pm p_A$. Now if two monic polynomials p_1 and p_2 of the same minimal degree n both annihilates A , then $(p_1 - p_2)(A) = p_1(A) - p_2(A) = 0 - 0 = 0$, so $p_1 - p_2$ also annihilates A , and it has lower degree than n (and can be made monic by dividing by its leading coefficient). This contradicts minimality unless $p_1 = p_2$ which shows uniqueness. □

Proposition 3.3.8. *The minimal polynomial divides any polynomial that annihilates A :*

$$p(A) = 0 \Rightarrow p(t) = m_A(t)q(t).$$

In particular, the minimal polynomial divides the characteristic polynomial.

Proof. Assume $p(A) = 0$. Divide $p(t)$ by $m_A(t)$ using polynomial division. We obtain a polynomial equation $p(t) = q(t)m_A(t) + r(t)$ where $\deg r(t) < \deg m_A$ or $r(t) = 0$. Replacing t by A we get $p(A) = q(A) \cdot m_A(A) + r(A)$ and $0 = q(A) \cdot 0 + r(A)$, so $r(A) = 0$. But then r is the zero-polynomial, otherwise it would be a polynomial of lower degree than m_A that annihilates A , which would contradict the minimality of m_A . We conclude that $p(t) = q(t)m_A(t)$ so $m_A(t)$ divides $p(t)$. □

Proposition 3.3.9. *The characteristic and minimal polynomial have the same zeros:*

$$p_A(\lambda) = 0 \Leftrightarrow m_A(\lambda) = 0.$$

Proof. Any zero of m_A is a zero of p_A since $m_A | p_A$. On the other hand, let λ be a zero of p_A . Then λ is an eigenvalue and there is some eigenvector v with $Av = \lambda v$. Then $0 = m_A(A)v = m_A(\lambda)v$ so $m_A(\lambda)$ is zero too. □

Example 3.3.10. Let us find the minimal polynomial of $A = \begin{pmatrix} 1 & 1 & 1 \\ -4 & 4 & 3 \\ -4 & 1 & 6 \end{pmatrix}$. We compute and factor the characteristic polynomial: $p_A(t) = -(t-3)^2(t-5)$. Now since $m_A(t)$ divides $p_A(t)$ and still has 3 and 5 as zeros, there are only two options: either $m_A(t) = -p_A(t) = (t-3)^2(t-5)$ or $m_A(t) = (t-3)(t-5)$. We test whether the second option annihilates A , but find that $(A-3I)(A-5I)$ is not the zero-matrix. Therefore $m_A(t) = (t-3)^2(t-5)$.

△

If we know the eigenvalues of some complex matrix A , what can be said about the eigenvalues of $p(A)$ where p is some polynomial? The answer is given by the *spectral mapping theorem*.

Theorem 3.3.11. (*Spectral mapping theorem*) Let A be a square complex matrix, and let p be any polynomial. Then we have

$$\sigma(p(A)) = p(\sigma(A)).$$

In other words, λ is an eigenvalue for A if and only if $p(\lambda)$ is an eigenvalue for $p(A)$.

Proof. Note that by $p(\sigma(A))$ we mean the set $\{p(\lambda) \mid \lambda \in \sigma(A)\}$. One direction is easy: if $Av = \lambda v$ and $p(t) = \sum a_k t^k$, then $p(A)v = \sum a_k A^k v = \sum a_k \lambda^k v = p(\lambda)v$. The other direction is left as an exercise. □

Intuitively, the spectrum of a complex matrix A is some finite set of points in \mathbb{C} . The theorem says that if we apply the polynomial p to each of these points we get the spectrum of the operator $p(A)$. The same holds if A is replaced by any operator F on a finite-dimensional complex vector space. In fact, the theorem holds if \mathbb{C} is replaced by any algebraically closed field⁵.

⁵A field \mathbb{F} is called algebraically closed if every polynomial $p(t) \in \mathbb{F}[t]$ has a zero in \mathbb{F} . This is equivalent to saying that p factors completely as $p(t) = c \prod_{i=1}^n (t - \lambda_i)$ for some $c \in \mathbb{F}$ and $\lambda_i \in \mathbb{F}$.

Chapter 4

Jordan normal form

In this chapter, all vector spaces are assumed to be finite-dimensional unless otherwise stated. Our goal of this section is to show that for *any* linear operator on a complex finite-dimensional vector space, there is a basis such that the matrix for the operator has a particular canonical format called the *Jordan form*¹. We shall soon prove this, and we shall discuss the algorithm for *Jordanizing* a matrix, but first, let's investigate the properties of matrices in this form. The definitions and many results of this section applies to any operator on a finite-dimensional \mathbb{F} -vector space, so assume that all matrices have coefficients in some arbitrary field \mathbb{F} unless otherwise stated.

4.1 Properties of matrices on Jordan form

Definition 4.1.1. The **Jordan block** of size n and with eigenvalue λ is defined as the $n \times n$ -matrix

$$J_n(\lambda) = \begin{pmatrix} \lambda & 1 & & & \\ & \lambda & 1 & & \\ & & \lambda & \ddots & \\ & & & \ddots & 1 \\ & & & & \lambda \end{pmatrix}.$$

In other words, $J_n(\lambda)$ has λ 's on the diagonal, and ones on the super-diagonal (one step over the diagonal), and zeros elsewhere.

Example 4.1.2. The matrix

$$J = J_3(5) = \begin{pmatrix} 5 & 1 & 0 \\ 0 & 5 & 1 \\ 0 & 0 & 5 \end{pmatrix}$$

is a Jordan block. We note that $p_J(t) = \det(J - tI) = (5 - t)^3$ so the only eigenvalue is 5, and since

$$J - 5I = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \text{ only multiples of } \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$$

are eigenvectors, so the geometric multiplicity of $t = 5$ is 1 and J is not diagonalizable.

The same goes in general, the characteristic polynomial of $J = J_n(\lambda)$ is $p_J(t) = \det(J - tI) = (\lambda - t)^n$, and the eigenspace $E_\lambda = \ker(J - \lambda I)$ is spanned by the first standard basis vector e_1 . In particular, the geometric multiplicity g_λ of J is 1, so J is not diagonalizable (except for the trivial case $n = 1$).

¹Named after Camille Jordan who first stated what is now known as the Jordan decomposition Theorem in 1870.

Definition 4.1.3. A matrix J is said to be in **Jordan form** if it is a block-diagonal matrix where each block is a Jordan block. In other words,

$$J = \text{diag}(J_{n_1}(\lambda_1), J_{n_2}(\lambda_2), \dots, J_{n_k}(\lambda_k)).$$

Note that the blocks may have different sizes, and that some of the λ_i may coincide.

Recall that direct sums of linear maps correspond to block-diagonal matrices. For this reason another common notation for the Jordan-matrix in the definition is

$$J = J_{n_1}(\lambda_1) \oplus J_{n_2}(\lambda_2) \oplus \dots \oplus J_{n_k}(\lambda_k).$$

Example 4.1.4. The following matrix is in Jordan form, for readability it is common to omit off-diagonal zeros and to draw boxes to indicate the Jordan blocks:

$$J = \begin{pmatrix} 5 & 1 & 0 & 0 & 0 & 0 \\ 0 & 5 & 1 & 0 & 0 & 0 \\ 0 & 0 & 5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 1 & 0 \\ 0 & 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2 \end{pmatrix} = \begin{pmatrix} \boxed{\begin{matrix} 5 & 1 \\ & 5 & 1 \\ & & 5 \end{matrix}} & & & & & \\ & & & & & & \\ & & & & & & \\ & & & \boxed{\begin{matrix} 2 & 1 \\ & 2 \end{matrix}} & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \boxed{2} \end{pmatrix} = J_3(5) \oplus J_2(2) \oplus J_1(2)$$

Here we see that the characteristic polynomial is $p_J(t) = (t - 5)^3(t - 2)^3$ so the eigenvalues are 5 and 2. Each Jordan-block corresponds to an eigenvector, in this example we see that e_1 is an eigenvector with eigenvalue 5, and that e_4 and e_6 both are eigenvectors of eigenvalue 2.

△

Note in particular that any Jordan block is in Jordan form (with a single block), and that any diagonal matrix is in Jordan form (all the blocks have size 1×1).

How can we determine the minimal polynomial for a matrix in Jordan form? Recall that the minimal polynomial of J is the monic polynomial $m_\lambda(t)$ of minimal degree that annihilates the matrix: $m_J(J) = 0$. Let us first consider an example where all the eigenvalues coincide:

Example 4.1.5. For the matrix J below, we have $p_J(t) = (2 - t)^7$, so the minimal polynomial has form $(t - 2)^n$ for some $1 \leq n \leq 7$, we plug in J and compute:

$$J = \begin{pmatrix} \boxed{\begin{matrix} 2 & 1 & 0 \\ 0 & 2 & 1 \\ 0 & 0 & 2 \end{matrix}} & \begin{matrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{matrix} \\ \begin{matrix} 0 & 0 & 0 & 2 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 \end{matrix} \end{pmatrix} \Rightarrow (J - 2I)^n = \begin{pmatrix} \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}^n & & & & & & \\ & & & & & & \\ & & & \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}^n & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}^n \end{pmatrix}$$

and we see that this is the zero-matrix if and only if $n \geq 3$, because then the largest Jordan block is annihilated by the polynomial, so $m_J(t) = (t - 2)^3$.

△

We see that the analogous argument works in general: For each zero λ of $p_J(t)$, the corresponding exponent of $(t - \lambda)$ in $m_J(t)$ is the size of the largest Jordan block of eigenvalue λ .

For example, in Example 4.1.4, the largest Jordan block for eigenvalue 5 had size 3, and for eigenvalue 2, the largest block had size 2. So the minimal polynomial is $m_J(t) = (t - 5)^3(t - 2)^2$. Here it is easy

to verify that this polynomial annihilates J by computing the product of block-diagonal matrices: we have

$$\begin{aligned}
 m_J(J) &= (J - 5I)^3(J - 2I)^2 = \left(\begin{matrix} \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}^3 & & \\ & \begin{bmatrix} -3 & 1 \\ 0 & -3 \end{bmatrix}^3 & \\ & & [-3]^3 \end{matrix} \right) \left(\begin{matrix} \begin{bmatrix} 3 & 1 & 0 \\ 0 & 3 & 1 \\ 0 & 0 & 3 \end{bmatrix}^2 & & \\ & \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}^2 & \\ & & [0]^2 \end{matrix} \right) \\
 &= \left(\begin{array}{c|c|c} 0 & & \\ \hline & * & \\ \hline & & * \end{array} \right) \left(\begin{array}{c|c|c} * & & \\ \hline & 0 & \\ \hline & & 0 \end{array} \right) = \left(\begin{array}{c|c|c} 0 & & \\ \hline & 0 & \\ \hline & & 0 \end{array} \right) = 0.
 \end{aligned}$$

To summarize the above discussion, we collect what we know about a matrix in Jordan form.

- The eigenvalues are the diagonal elements of J , repeating according to algebraic multiplicity. In other words, the algebraic multiplicity of λ is the sum of the sizes of all λ -Jordan blocks in J .
- The geometric multiplicity of λ is the number of λ -Jordan blocks. The top-left position of each Jordan block corresponds to an eigenvector: if a Jordan-block $J_n(\lambda)$ sits in J with top-left corner in position (i, i) , then e_i is an eigenvector of eigenvalue λ .
- If $p_J(t) = \pm(t - \lambda_1)^{m_1} \dots (t - \lambda_k)^{m_k}$, the minimal polynomial is $m_J(t) = (t - \lambda_1)^{m_1} \dots (t - \lambda_k)^{m_k}$, where m_i is the size of the largest λ_i -Jordan block in J .

Algebraic multiplicities, geometric multiplicities, characteristic- and minimal polynomials are invariant under a change of basis, so if we know these invariants for an arbitrary matrix A , and we are trying to write $J = S^{-1}AS$, we get some information about the shape of J .

Given a linear map $F : V \rightarrow V$, the goal of this chapter is to find a basis for V such that the matrix of F is in Jordan form. This can be somewhat complicated, so let us first restrict ourselves to *nilpotent* maps F .

4.2 Structure theory for nilpotent operators

If F is nilpotent, then $F^d = 0$ for some minimal d which is the nilpotency degree of F . Then the minimal polynomial of F is $m_F(t) = t^d$ which shows that zero is the only eigenvalue of F .

We are looking for a special type of basis for V which clearly reveals structure of F :

Definition 4.2.1. Let $F : V \rightarrow V$ be nilpotent. A **string** (for F) in V is a sequence of nonzero vectors (v_1, \dots, v_n) , such that $F(v_i) = v_{i-1}$ and $F(v_1) = 0$, or visually

$$v_n \mapsto \dots \mapsto v_2 \mapsto v_1 \rightarrow 0.$$

Here we call v_n the **first vector**, and we call v_1 the **last vector** of the string, and the **length** of the string is n . We shall sometimes draw simpler arrows and sometimes omit the vectors when drawing strings.

A **string basis** for V is a basis for V which is a union of strings.

In our definition from previous chapters, a basis needs to be ordered. The different strings in the string-basis can be taken in any order, but each string should be ordered as above, from right to left.

Strings are sometimes called *Jordan chains* or just *chains*, I will use the word string when referring to nilpotent operators and reserve the other words for the more general case that we will investigate later.

Example 4.2.2. Suppose that a linear map $F : \mathbb{C}^8 \rightarrow \mathbb{C}^8$ has a string basis (e_1, \dots, e_8) on which F acts as

indicated in the left diagram below. In this basis, the matrix of F will have the form to the right:

$$\begin{array}{l}
 e_3 \rightarrow e_2 \rightarrow e_1 \rightarrow 0 \\
 e_5 \rightarrow e_4 \rightarrow 0 \\
 e_7 \rightarrow e_6 \rightarrow 0 \\
 e_8 \rightarrow 0
 \end{array}
 \quad [F] = \begin{pmatrix}
 \boxed{0 & 1 & 0} & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & \boxed{0 & 1} & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & \boxed{0 & 1} & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & \boxed{0}
 \end{pmatrix}$$

Here we note that $[F] = J_3(0) \oplus J_2(0) \oplus J_2(0) \oplus J_1(0)$ is in Jordan form, and that the nilpotency degree of F is 3, which corresponds to the length of the longest string. This also shows that $m_F(t) = t^3$. We see that each string corresponds to a Jordan block, and that the rightmost vectors of the strings are the eigenvectors with eigenvalue 0.

The images and kernels $\text{Im}(F^k)$ and $\text{ker}(F^k)$ can easily be visualized from the string-diagram. For example, $\text{ker}(F)$ consists of vectors that are mapped to zero, so this subspace is spanned by the rightmost vectors of all the strings: $\text{ker}(F) = \text{span}(e_1, e_4, e_6, e_8)$. Similarly, the subspace $\text{ker}(F^2)$ is spanned by vectors which are mapped to zero when F is applied twice, so this is spanned by all the standard basis vectors except e_3 . Clearly $\text{ker}(F^3) = \mathbb{C}^8$.

Dually, the image of F is spanned by all basis vectors except the first (leftmost) vectors of each string, $\text{Im}(F) = \text{span}(e_1, e_2, e_4, e_6)$, while $\text{Im}(F^2) = \text{span}(e_1)$. Clearly $\text{Im}(F^3) = 0$.

△

Lemma 4.2.3. *If the last vectors of a set of strings are linearly independent (the eigenvectors with eigenvalue 0), then so are the whole strings. More precisely, if $\mathcal{S} = \{v_k^{(i)}\}$ are strings for F with $F(v_1^{(i)}) = 0$ and $F(v_k^{(i)}) = v_{k-1}^{(i)}$, then if the last vectors $\{v_1^{(i)}\}$ is a linearly independent set, then so is \mathcal{S} .*

Proof. Following the same idea as in the proof of Proposition 3.1.9, assume a minimal nontrivial linear combination of string-vectors is zero. Then applying the operator F shifts all the vectors in the linear combination one step to the right, so by applying F until a vector in the linear combination is mapped to zero, we have obtained a linear combination with fewer elements which is zero, contradicting the minimality of our original linear combination. □

Theorem 4.2.4. *For any nilpotent map $F : V \rightarrow V$ there exists a string basis, and with respect to this basis, the matrix $[F]$ is in Jordan form with zeroes on the diagonal.*

Proof. We proceed by induction on $\dim V$, the base case $\dim V = 0$ is obvious.

Now assume $F : V \rightarrow V$ is nilpotent and $\dim V = n$, and that the statement is true for all nilpotent maps on vector spaces of dimension lower than n . Now if F is the zero map the statement is obvious (any basis is a string basis). Otherwise, $0 < \dim \text{ker } F < n$, so by rank-nullity $\text{Im}(F)$ is a proper subspace. Consider the restriction $F|_{\text{Im}(F)} : \text{Im}(F) \rightarrow \text{Im}(F)$, by the induction hypothesis there exists a string basis for this map, denote this basis

$$\{e_j^{(i)} \mid 1 \leq j \leq k_i\}$$

where the top index indicate the different strings, and k_i is the length of string number i , and where $F(e_j^{(i)}) = e_{j-1}^{(i)}$ and $F(e_1^{(i)}) = 0$. Consider the leftmost vectors $e_{k_i}^{(i)}$ in these strings, they belong to $\text{Im}(F)$, so we can extend each string westwards by picking $e_{k_i+1}^{(i)} \in V$ such that $F(e_{k_i+1}^{(i)}) = e_{k_i}^{(i)}$. Now our extended string-vectors are linearly independent, and they span a subspace $U \subset V$ containing $\text{Im}(F)$. Now by construction, $F(U) = \text{Im}(F)$. The subspace U might not be all of V , so we extend U by vectors v_1, \dots, v_s to a basis of V . But $F(v_i)$ belongs to $\text{Im}(F) = F(U)$, so $F(v_i) = F(u_i)$ for some $u_i \in U$. Now take $w_i = v_i - u_i$. Then the w_i still span a complement to U , and we have $F(w_i) = F(v_i) - F(u_i) = 0$, so $w_i \in \text{ker}(F)$. But then take our string basis for U and adjoin the w_i as strings of length 1, then we have a string basis for V . □

If you want to understand the proof, it might be instructive to study how it applies to the map:

$$\begin{array}{c}
 e_2 \mapsto e_1 \mapsto 0 \\
 \swarrow \\
 e_3
 \end{array}$$

Here the image is spanned by e_1 , we extend the only string to a string of length 2: $e_2 \mapsto e_1 \mapsto 0$. Then $U = \text{span}(e_1, e_2)$ is not all of V , so we extend by $v = e_3$, and note that $F(e_3) = e_1 \in F(U)$, and we pick an element w of U that is also mapped to e_1 , let's take $u = e_2$. Then $w = v - u = e_3 - e_2$ is mapped to zero, and is a string of length 1. So in this case our derived string basis would be $(e_1, e_2, e_3 - e_2)$.

Corollary 4.2.5. *Let A be a nilpotent matrix. Then there exists an invertible matrix S and a matrix J in Jordan-form (with zeros on the diagonal) such that $S^{-1}AS = J$. The matrix J is unique up to permutation of the Jordan blocks.*

Proof. Use Theorem 4.2.4 to find a string basis, and let S be the matrix with the string-basis vectors as columns in order string by string from the right side of each string, then clearly $S^{-1}AS$ is in Jordan form. For the uniqueness claim, we note that the sequence of numbers $\dim \ker(A^k) = \dim \ker(J^k)$ will determine uniquely the number of strings and the length of strings in the string basis, and therefore the shape of the Jordan form up to permutation. \square

Since the Jordan form of nilpotent operators is unique we shall speak of *the* Jordan form of a nilpotent operator².

Example 4.2.6. Let us classify nilpotent operators on a 5-dimensional vector space V . The corollary above guarantees that any nilpotent operator is similar to a matrix in Jordan form, so we need only find all Jordan-forms of 5×5 matrices where the diagonal elements are zero. The block-partition of the Jordan matrix corresponds to a partition of the integer 5 as a sum of positive integers, and there are seven such partitions:

$$(5), (4 + 1), (3 + 2), (3 + 1 + 1), (2 + 2 + 1), (2 + 1 + 1 + 1), (1 + 1 + 1 + 1 + 1).$$

Here are the corresponding Jordan forms:

The corollary guarantees that for any nilpotent map F on a 5-dimensional vector space, there exists a basis for which the matrix of $[F]$ is one (and only one) of the matrices above. Alternatively stated, any nilpotent matrix A is similar to exactly one of the matrices above: $A = SJS^{-1}$ where J is one of the seven matrices above.

\triangle

Given a nilpotent operator, how do you actually find a string basis algorithmically? Before going hunting for basis vectors, a good starting point is determining the number of strings and their lengths. This determines the Jordan form of the operator.

²Usually we consider two Jordan forms to be the same if they contain the same Jordan blocks in different order. Note however that while J is unique, there may be several choices for the actual vectors of the string basis (the columns of S).

Lemma 4.2.7. For a nilpotent operator (or matrix) $A : V \rightarrow V$, define $n_k = \dim \ker(A^k)$, and let d_k be the number of strings of length k in any string basis, or equivalently, the number of Jordan blocks of size k in the Jordan form of A . Then for $k \geq 1$,

$$d_k = 2n_k - n_{k-1} - n_{k+1}.$$

Proof. Note $n_0 = 0$ and that $n_k = \dim V$ for all $k \geq m$ where m is the nilpotency-degree of A , and that n_0, n_1, \dots, n_m is a strictly increasing integer sequence until it reaches $\dim V$. Note also that the maximum length of a string is m , so $d_k = 0$ for $k > m$.

Visualize a generic string-diagram and consider what the different kernels $\ker(A^k)$ look like in terms of the strings. Now $n_1 = n_1 - n_0 = \dim \ker(A)$ equals the total number of strings, or in other words, the number of strings of length ≥ 1 . Similarly, $n_2 - n_1$ equals the number of strings of length ≥ 2 . By this line of reasoning, $n_i - n_{i-1}$ equals the number of strings of length $\geq i$. But then the number of strings of length *exactly* k is equal to the number of strings of length $\geq k$ minus the number of strings of length $\geq k + 1$, so

$$d_k = \underbrace{(n_k - n_{k-1})}_{\text{\#strings of length } \geq k} - \underbrace{(n_{k+1} - n_k)}_{\text{\#strings of length } \geq k+1} = 2n_k - n_{k-1} - n_{k+1}.$$

□

Example 4.2.8. Let's say that we have a matrix $A \in \text{Mat}_{11}(\mathbb{C})$ and we have found that $A^4 = 0$, and that $\text{rank}(A) = 7, \text{rank}(A^2) = 3, \text{rank}(A^3) = 1$. Let us determine the Jordan form of A from this information.

Let $n_k = \dim \ker(A^k)$ as in the lemma. By rank-nullity we get $n_0 = 0, n_1 = 11 - 7 = 4, n_2 = 11 - 3 = 8, n_3 = 11 - 1 = 10$, and $n_k = 11$ for $k > 4$. By the lemma, the numbers d_k of strings can be computed as $d_k = 2n_k - n_{k-1} - n_{k+1}$, which gives

$$d_1 = 0, d_2 = 2, d_3 = 1, d_4 = 1,$$

and $d_k = 0$ for $k \geq 4$.

So we conclude that there are 2 strings of length 2, 1 strings of length 3, and 1 chain of length 4, so the Jordan form for A is:

$$S^{-1}AS = J_2(0) \oplus J_2(0) \oplus J_3(0) \oplus J_4(0).$$

△

Now, to find the actual vectors in a string basis, it is easiest to start with the *first vectors of the longest strings*.

Suppose that we know for some nilpotent map A that there are 2 strings of length 3, and no longer strings. Some starting vectors for these strings should lie in $\ker(A^3) \setminus \ker(A^2)$, so we pick two vectors v and w that span a complement⁴ to $\ker(A^3)$ in $\ker(A^2)$. Then we apply A to get the rest of these strings:

$$v \mapsto Av \mapsto A^2v \mapsto 0 \text{ and } w \mapsto Aw \mapsto A^2w \mapsto 0.$$

Now we proceed to find the strings of length 2 by finding their first vectors which should lie in $\ker(A^2) \setminus \ker(A)$. However, we should also be careful not to pick any vector which is linearly dependent the six vectors v, Av, A^2v, w, Aw, A^2w from our previously chosen strings. Proceeding like this eventually produces a string basis.

In light of the discussion above, here is an algorithm⁵ for finding a string basis for a nilpotent matrix (or operator).

³Recall that $A \setminus B$ means the set-difference, it meaning the elements of A that are not in B . Don't confuse it with our notation V/U for quotient-spaces.

⁴In other words, extend a basis of $\ker(A^2)$ to $\ker(A^3)$

⁵Several variations of this can be found in the literature.

Algorithm 4.2.9. To find a string basis for a nilpotent matrix A , do the following:

1. Write down the matrices A^k and find a basis in each subspace $\ker(A^k)$ until $A^m = 0$. Let $n_k = \dim \ker(A^k)$.
2. Find $d_k = 2n_k - n_{k+1} - n_{k-1}$, the number of strings of length k .
3. Sketch a string-diagram and write down the corresponding Jordan matrix J .
4. For each string-length k , from longest to shortest do:
 - (a) Let \mathcal{B} be the set of previously chosen vectors (\mathcal{B} is empty in the first step)
 - (b) Pick linearly independent vectors v_1, \dots, v_{d_k} in $\ker(A^k)$ that are also linearly independent to $\ker(A^{k-1})$ and to the previously picked vectors in \mathcal{B} . These vectors will be first vectors in strings of length k .
 - (c) Compute the rest of the strings $v_i, Av_i, A^2v_i \dots A^{k-1}v_i$ and verify that $A^k v_i = 0$, adjoin all these nonzero vectors to \mathcal{B} .
 - (d) Decrease k by 1 and repeat until $k = 0$.
5. \mathcal{B} now contains a string basis. Order its vectors string by string from right to left, in the same order as your string diagram (and Jordan matrix J). Let S be the matrix with the string-basis vectors as columns in this order.
6. Verify that $SJS^{-1} = A$.

In step 1, things are easier if pick a basis in each successive $\ker(A^{k+1})$ by extending a basis from $\ker(A^k)$. If we know that the matrix S is invertible, it is easier to verify that $SJ = AS$ in the last step, even faster is to just verify that for each column s_i of S , we have either $As_i = 0$ or $As_i = s_{i-1}$.

The algorithm may seem complicated, especially step 4b. However, for small matrices some steps are quite obvious, let's illustrate:

Example 4.2.10. Let $A = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}$. This gives $A^2 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$ and $A^3 = 0$. We see immediately

that

$$\ker(A) = \text{span}(e_2, e_1 - e_4) \text{ and } \ker(A^2) = \text{span}(e_1, e_2, e_4).$$

So $(n_1, n_2, n_3) = (2, 3, 4)$ and $(d_1, d_2, d_3) = (1, 0, 1)$, so we are looking for two strings, one of length 1 and one of length 3. The Jordan form will therefore be $J_3(0) \oplus J_1(0)$.

Following the algorithm we are first looking for one vector v in $\ker(A^3) \setminus \ker(A^2) = \mathbb{C}^4 \setminus \text{span}(e_1, e_2, e_4)$. We pick $v = e_3$ as the first vector of this string. The rest of the string will be $Ae_3 = e_4$ and $Ae_4 = e_2$, and as we expect, this last vector is indeed mapped to 0 by A . So this is our string of length 3. There are no strings of length 2, so we look instead at the last string of length 1. Now we should pick a vector in $\ker(A^1) \setminus \ker(A^0) = \ker(A) \setminus \{0\}$ that is also linearly independent to our previously chosen vectors (e_3, e_4, e_2) . Since $\ker(A) = \text{span}(e_2, e_1 - e_4)$, a natural choice is to pick $e_1 - e_4$, this is indeed mapped to zero.

Now we have our string basis $\mathcal{B} = (e_2, e_4, e_3, e_1 - e_4)$, where the order matches our chosen Jordan form. We put these string basis vectors as columns of a matrix S , and we take J to be the corresponding Jordan matrix we found above:

$$S = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & -1 \end{pmatrix} \quad J = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

Now we can verify that we indeed have $S^{-1}AS = J$.

△

4.3 Jordan chains and Jordanization

Once we know how to deal with nilpotent matrices, the rest of the theory is not hard.

Generalized eigenspaces

Recall that the eigenspace E_λ for a linear operator (matrix) $A : V \rightarrow V$ consists of all vectors v for which $(A - \lambda I)v = 0$. *Generalized eigenvectors* for the eigenvalue λ are vectors that are *eventually* mapped to zero by $(A - \lambda I)$.

Definition 4.3.1. Let $A : V \rightarrow V$ be an operator (or matrix). A nonzero vector v is called a **generalized eigenvector** for the eigenvalue λ if there exists a positive integer n such that $(A - \lambda I)^n v = 0$. The set of all generalized eigenvectors for the eigenvalue λ (and the zero vector) is called the **generalized eigenspace** for λ , we denote it by \tilde{E}_λ :

$$\tilde{E}_\lambda = \{v \in V \mid (A - \lambda I)^n v = 0 \text{ for some } n\}.$$

We note that we have a sequence

$$E_\lambda = \ker(A - \lambda I) \subset \ker((A - \lambda I)^2) \subset \ker((A - \lambda I)^3) \subset \dots$$

so \tilde{E}_λ is the union of all these kernels. However, if V is finite-dimensional, this sequence of kernels eventually stabilizes, so there exists some minimal integer n such that $\ker((A - \lambda I)^n) = \ker((A - \lambda I)^{n+k})$ for all $k \geq 0$. Then $\tilde{E}_\lambda = \ker((A - \lambda I)^n)$, so \tilde{E}_λ is in fact a subspace. Note also that the operator $(A - \lambda I)$ acts nilpotently on the subspace \tilde{E}_λ , in other words:

$$((A - \lambda I)|_{\tilde{E}_\lambda})^n = 0.$$

But as a map $V \rightarrow V$ on the whole space, the operator $(A - \lambda I)$ typically is not nilpotent.

Example 4.3.2. Consider the matrix J on Jordan form below

$$A = \begin{pmatrix} \boxed{5} & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \boxed{5} & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \boxed{5} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \boxed{8} & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \boxed{8} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \boxed{5} & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \boxed{5} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \boxed{3} \end{pmatrix}$$

Here, the eigenvalues are 5, 3 and 8. In the matrix $A - 5I$, the 5's on the diagonal disappear, leaving two nilpotent Jordan blocks. We see that

$$\ker(A - 5I) = \text{span}(e_1, e_6), \ker((A - 5I)^2) = \text{span}(e_1, e_2, e_6, e_7), \ker((A - 5I)^3) = \text{span}(e_1, e_2, e_3, e_6, e_7).$$

But then the sequence clearly stabilizes and the $\ker((A - 5I)^n)$ doesn't change for $n \geq 3$. So the generalized eigenspace $\tilde{E}_5 = \ker((A - 5I)^3) = \text{span}(e_1, e_2, e_3, e_6, e_7)$ is 5-dimensional. Similarly, $\tilde{E}_8 = \text{span}(e_4, e_5)$ and $\tilde{E}_3 = \text{span}(e_8)$. Note that the direct sum of the generalized eigenspaces is the whole space

$$\mathbb{C}^8 = \tilde{E}_5 \oplus \tilde{E}_8 \oplus \tilde{E}_3,$$

and that the dimension of each \tilde{E}_λ is equal to the *algebraic multiplicity* of λ in p_A .

△

Jordan chains

The idea for finding the Jordan form of an arbitrary matrix is to treat each generalized eigenspace separately.

Definition 4.3.3. Let $A : V \rightarrow V$ be a linear operator (or matrix). A **Jordan chain** (for eigenvalue λ) in V is a sequence of nonzero vectors (v_1, \dots, v_n) such that $(A - \lambda I)v_i = v_{i-1}$ and $(A - \lambda I)v_1 = 0$. In other words, a Jordan chain is just a string in the sense of the previous section, but for the operator $(A - \lambda I)$.

A **Jordan basis** for V is a basis which is a union of Jordan chains (possibly with different λ).

So a Jordan basis looks like a string basis but where the chains may correspond to different eigenvalues. Note that the last vector in each string is an *eigenvector* since some $(A - \lambda I)$ maps it to zero. With respect to a Jordan basis, the matrix of the operator is clearly in Jordan form.

We can now generalize our algorithm for finding a string basis to an algorithm for finding a Jordan basis and the corresponding Jordan form, in other words, to **Jordanize** the operator. Given a square matrix A , the algorithm produces a matrix J in Jordan form, and an invertible matrix S (with columns forming a Jordan basis) such that $SJS^{-1} = A$.

Jordanizing a matrix

Algorithm 4.3.4. To Jordanize a square matrix A :

1. Find the characteristic polynomial, and solve $p_A(\lambda) = 0$ to find the eigenvalues.
2. For each eigenvalue λ do:
 - (a) Let \mathcal{B} be the set of previously chosen vectors (start with an empty set).
 - (b) Find the generalized eigenspace \tilde{E}_λ by computing $\ker((A - \lambda I)^k)$ for $k = 1, 2, \dots$ until the sequence stabilizes: $\ker((A - \lambda I)^n) = \ker((A - \lambda I)^{n+1})$, this happens when the dimension of $\dim \ker((A - \lambda I)^n)$ has reached the algebraic multiplicity of λ . Then $\tilde{E}_\lambda = \ker((A - \lambda I)^n)$.
 - (c) Let $N := (A - \lambda I)|_{\tilde{E}_\lambda} : \tilde{E}_\lambda \rightarrow \tilde{E}_\lambda$. This operator is nilpotent.
 - (d) Apply Algorithm 4.2.9 to find a string basis in \tilde{E}_λ for the operator N . This basis is a union of Jordan chains with eigenvalue λ , adjoin all these strings to \mathcal{B} .
3. Now \mathcal{B} should contain a Jordan basis. Let S be the matrix whose columns are the elements of the Jordan chains, order each chain from right to left (starting with the eigenvector). Take J to be the corresponding Jordan matrix with block sizes and eigenvalues corresponding to the ordering of the chains in S .
4. Verify that $S^{-1}AS = J$.

When we have found \tilde{E}_i , one possibility is to pick a basis and construct a smaller matrix for the operator $A|_{\tilde{E}_i} : \tilde{E}_i \rightarrow \tilde{E}_i$. This may simplify finding the string basis, but the basis should then be converted back to be written in the standard basis for the whole space. Alternatively, when following Algorithm 4.2.9, work with vectors sitting inside \tilde{E}_i , but expressed in the basis for the whole space.

Before looking at an example, let's consider why this algorithm actually works.

There are two important pieces missing: what if the Jordan chains we pick inside one generalized eigenspace are linearly dependent with chains from other generalized eigenspaces? And how can we be sure that all the Jordan chains we find in the algorithm actually span the whole space? These concerns are put to rest by the following lemma:

Lemma 4.3.5. Let $A : V \rightarrow V$ be an linear operator (or a matrix) on a finite dimensional complex vector space. Then V is the direct sum of the generalized eigenspaces for the operator:

$$V = \tilde{E}_{\lambda_1} \oplus \tilde{E}_{\lambda_2} \oplus \dots \oplus \tilde{E}_{\lambda_k}$$

where $\sigma(A) = \{\lambda_1, \dots, \lambda_k\}$.

Proof. First we claim that vectors from different generalized eigenspaces are linearly independent. To show this one can combine the proof-techniques used to prove Proposition 3.1.9 (vectors from different eigenspaces are linearly independent), and Lemma 4.2.3 (strings are linearly independent). We omit the details here.

It remains to prove that the eigenspaces span V . Let $\dim V = n$, factor the characteristic polynomial of A over \mathbb{C} :

$$p_A(t) = (-1)^n (t - \lambda_1)^{n_1} \cdots (t - \lambda_k)^{n_k}.$$

Now perform a partial fraction decomposition of the rational function $\frac{1}{p_A(t)}$, and multiply it by $p_A(t)$, we get

$$1 = p_A(t) \cdot \frac{1}{p_A(t)} = p_A(t) \left(\frac{q_1(t)}{(t - \lambda_1)^{n_1}} + \cdots + \frac{q_k(t)}{(t - \lambda_k)^{n_k}} \right) = h_1(t)q_1(t) + \cdots + h_k(t)q_k(t),$$

where we introduced $h_i(t) = \frac{p_A(t)}{(t - \lambda_i)^{n_i}}$ in the last step. Note that the h_i are *polynomials* as the denominators cancel out. Now evaluate the polynomial identity above at the matrix A , define the operator $P_i := h_i(A)q_i(A)$, we get:

$$I = h_1(A)q_1(A) + \cdots + h_k(A)q_k(A) = P_1 + \cdots + P_k.$$

We now claim that P_i maps vectors into the generalized eigenspace \tilde{E}_{λ_i} , in other words, $\text{Im}(P_i) \subset \tilde{E}_{\lambda_i}$. Indeed for any $v \in V$, we have

$$(A - \lambda_i I)^{n_i} \cdot P_i v = (A - \lambda_i I)^{n_i} q_i(A) h_i(A) v = p_A(A) q_i(A) v = 0 \cdot q_i(A) v = 0,$$

which shows that vectors $P_i v$ in the image of P_i are killed by $(A - \lambda_i I)^{n_i}$, so $P_i v \in \tilde{E}_{\lambda_i}$ as claimed.

But then for $v \in V$ we have

$$v = Iv = (P_1 + \cdots + P_k)v = P_1 v + \cdots + P_k v,$$

so we can express any vector as a sum of vectors from each generalized eigenspace, so the \tilde{E}_{λ_i} span V . □

Note that the proof of Lemma 4.3.5 actually provide an explicit way to construct projection maps onto each eigenspace: The map $P_i : V \rightarrow V$ where projects each vector v onto \tilde{E}_{λ_i} . It is not hard to prove that these projections satisfy

$$P_i^2 = P_i, \quad P_i P_j = 0, \quad P_1 + \cdots + P_k = I.$$

The proof also shows that the algebraic multiplicity of λ is equal to the dimension of the corresponding eigenspace:

$$m_\lambda = \dim \tilde{E}_\lambda.$$

This guarantees that the method of Algorithm 4.3.4 works and will produce a Jordan basis.

The results of this section can be summarized in a single theorem:

Theorem 4.3.6. Jordan Theorem. *Let $F : V \rightarrow V$ be an operator on a complex vector space. Then there exists a basis for V with respect to which the matrix of $[F]$ has Jordan form. Equivalently, any complex square matrix A has a factorization $A = SJS^{-1}$. The Jordan form is unique up to permutation of the blocks.*

Proof. Lemma 4.3.5 shows that Algorithm 4.3.4 will produce a Jordan-basis for the space V . For the uniqueness-claim, we note that the blocks of the Jordan form are uniquely determined by the numbers $\dim \ker((A - \lambda I)^k)$, and these are invariant under a change of basis. □

Example 4.3.7. Let us Jordanize the matrix A below. In other words, we shall find an invertible matrix S and a matrix J in Jordan form such that $S^{-1}AS = J$.

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

We find that the characteristic polynomial is $p_A(t) = \det(A - tI) = -(t - 1)^3(t - 3)^2$, so the eigenvalues are 1 and 3.

We start with the eigenvalue 1. We compute powers of $A - I$ and find that

$$A - I = \begin{pmatrix} 2 & -2 & -3 & 3 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 1 & -1 & 0 & 2 & -1 \\ 1 & -1 & 1 & 1 & -2 \\ 0 & 0 & -1 & 1 & 1 \end{pmatrix} \quad (A - I)^2 = \begin{pmatrix} 4 & -4 & -4 & 4 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 4 & -4 & 0 & 4 & -4 \\ 4 & -4 & 0 & 4 & -4 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

We solve $(A - I)v = 0$ and find that

$$\ker(A - I) = \text{span}((1, 0, 1, 0, 1), (1, 1, 0, 0, 0)).$$

Next we solve $(A - I)^2v = 0$ and we find that $\ker((A - I)^2)$ is spanned by one more vector, so we extend our previous kernel and write

$$\ker((A - I)^2) = \text{span}((1, 0, 1, 0, 1), (1, 1, 0, 0, 0), (1, 0, 0, -1, 0)).$$

Now $\ker((A - I)^3)$ cannot be bigger (looking at the algebraic multiplicity of $\lambda = 1$).

We conclude that \tilde{E}_1 is three-dimensional, and the dimensions of the kernels tell us that we are looking for one Jordan chain of length 2, and one of length 1 (see Algorithm 4.2.9 for details). To find the longest chain, we start with a vector in $\ker((A - I)^2) \setminus \ker(A - I)$, let's take $v_2 := (1, 0, 0, -1, 0)$. Then take $v_1 := (A - I)v_2 = (1, 0, 1, 0, 1)$, so now we have a chain $v_2 \mapsto v_1 \mapsto 0$. For the second chain of length 1, we just pick a vector in $\ker(A - I)$ independent from v_1 and v_2 , let's take $v_3 = (1, 1, 0, 0, 0)$. Now we have our chains for the first generalized eigenspace $\tilde{E}_1 = \text{span}(v_1, v_2, v_3)$. With respect to this basis, the restriction of A onto \tilde{E}_1 will have Jordan form $J_2(1) \oplus J_1(1)$.

We move on to the second eigenvalue $\lambda = 3$ and find that

$$A - 3I = \begin{pmatrix} 0 & -2 & -3 & 3 & 1 \\ 0 & -2 & 0 & 0 & 0 \\ 1 & -1 & -2 & 2 & -1 \\ 1 & -1 & 1 & -1 & -2 \\ 0 & 0 & -1 & 1 & -1 \end{pmatrix} \quad (A - 3I)^2 = \begin{pmatrix} 0 & 4 & 8 & -8 & -4 \\ 0 & 4 & 0 & 0 & 0 \\ 0 & 0 & 4 & -4 & 0 \\ 0 & 0 & -4 & 4 & 4 \\ 0 & 0 & 4 & -4 & 0 \end{pmatrix}$$

And there is no reason to look further, we are only missing two vectors. We find a basis for $\ker(A - 3I)$ and extend it to a basis for $\ker((A - 3I)^2)$:

$$\ker(A - 3I) = \text{span}(0, 0, 1, 1, 0) \quad \ker((A - 3I)^2) = \text{span}((0, 0, 1, 1, 0), (1, 0, 0, 0, 0)).$$

So we are looking for a single chain of length 2, we pick its first vector in $\ker((A - 3I)^2) \setminus \ker(A - 3I)$, let's take $v_5 := (1, 0, 0, 0, 0)$, and we compute $v_4 := (A - 3I)v_5 = (0, 0, 1, 1, 0)$.

Now we have our Jordan chains. We collect the data by creating $S = (v_1, v_2, v_3, v_4, v_5)$ with the chain vectors as columns, and we take J to be the corresponding Jordan matrix:

$$S = \begin{pmatrix} 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 \\ 0 & -1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \end{pmatrix} \quad J = \begin{pmatrix} 1 & 1 & & & \\ & 1 & & & \\ & & 1 & & \\ & & & 3 & 1 \\ & & & & 3 \end{pmatrix}$$

Then $S^{-1}AS = J$.

△

What if we consider a vector space over a field \mathbb{F} other than \mathbb{C} ? Well, we need the characteristic polynomial to factor completely over \mathbb{F} for the algorithm to work.

The general result is that if $A : V \rightarrow V$ is an operator on an \mathbb{F} -vector space, and if $p_A(t)$ factors completely over $\mathbb{F}[t]$ into linear factors:

$$p_A(t) = \det(A - tI) = \pm(t - \lambda_1)(t - \lambda_2) \cdots (t - \lambda_k)$$

⁶Here $\mathbb{F}[t]$ is the set of polynomials in a variable t with coefficients in the field \mathbb{F} .

for some $\lambda_i \in \mathbb{F}$, then we can Jordanize A : there exists a basis in V for which the matrix of A has Jordan form.

The previous computation was actually an example of this, we started with a *real* matrix in $\text{Mat}_5(\mathbb{R})$, and we managed to factor $p_A(t)$ completely as a product of *real* linear factors $(t - \lambda)$ where $\lambda \in \mathbb{R}$, therefore the algorithm worked and we obtained a factorization by real matrices $A = SJS^{-1}$.

4.4 Matrix functions and applications

Motivation

So what is the point of knowing how to Jordanize a matrix? Well, first of all it has many theoretical applications - if we want to prove a statement for any linear map $F : V \rightarrow V$ we may always pick a Jordan basis for V , so it suffices to prove the statement for matrices in Jordan form.

But aside from this there are several applications of a more practical nature. In a first course in linear algebra, we could approach a large class of problems via diagonalization:

- Dynamical systems in discrete time, such as predator prey models
- Explicit forms for recursively defined sequences, such as the Fibonacci numbers
- Linear systems of differential equations

Solutions to these problems were found by diagonalizing some constant coefficient-matrix. With the method of this chapter, these problems become tractable even when the coefficient matrix is non-diagonalizable.

Dynamical systems in discrete time

A dynamical system in discrete time is a system $X_{n+1} = AX_n$ where $X_n = (x_n^{(1)}, \dots, x_n^{(m)})^T$ is a set of m sequences of numbers, and A is a constant $m \times m$ -matrix. The system tells you how to go from X_n to X_{n+1} (we can think of this as a number of variables changing in the next "time step").

Now clearly $X_n = A^n X_0$, so to find the explicit form of the solutions we need to find a general formula for A^n . This could be done by the Jordan form. Write $A = SJS^{-1}$, where we can decompose $J = D + N$ into a diagonal matrix D and a nilpotent matrix N (with some ones on the superdiagonal). Then

$$A^n = (SJS^{-1})^n = SJ^nS^{-1} = S(D + N)^nS^{-1} = S(D + N)^nS^{-1}$$

So it suffices to find $(D + N)^n$. But the matrices D and N commute, so the binomial theorem applies:

$$(D + N)^n = \sum_{k=0}^n \binom{n}{k} N^k D^{n-k}.$$

But since N is nilpotent, only the first few terms of this sum are nonzero, and powers of a diagonal matrix is easy to compute.

Hopefully an example will clarify the general method:

Example 4.4.1. Rabbits and foxes are living in forest, the foxes are hunting the rabbits. In year number n there are r_n rabbits and f_n foxes in the forest. From the start there are 40 rabbits and 10 foxes, and we assume that the populations evolve according to the following model:

$$\begin{cases} r_{n+1} = 3r_n - f_n \\ f_{n+1} = r_n + f_n \end{cases}$$

Intuitively this should sort of make sense: If there are no foxes, the rabbit population increase 3-fold each year, but the more foxes there are the smaller the increase. Similarly, without rabbits the foxes can just sustain themselves, but more rabbits means an increase in the fox-population.

Let us find explicit formulas for r_n and f_n . First let $X_n = \begin{pmatrix} r_n \\ f_n \end{pmatrix}$ and $A = \begin{pmatrix} 3 & -1 \\ 1 & 1 \end{pmatrix}$. Then $X_0 = \begin{pmatrix} 40 \\ 10 \end{pmatrix}$ and $X_{n+1} = AX_n$ so $X_n = A^n X_0$.

To compute A^n we first Jordanize A and find that $A = SJS^{-1}$ for $J = \begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix}$ and $S = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$. We write $J = 2I + N$ where $N^2 = 0$. Then

$$J^n = (2I + N)^n = \sum_{k=0}^n \binom{n}{k} (2I)^{n-k} N^k = (2I)^n + n(2I)^{n-1}N = \begin{pmatrix} 2^n & 0 \\ 0 & 2^n \end{pmatrix} + \begin{pmatrix} 0 & n2^{n-1} \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 2^n & n2^{n-1} \\ 0 & 2^n \end{pmatrix}.$$

So the solution is

$$X_n = A^n X_0 = (SJS^{-1})^n X_0 = SJ^n S^{-1} X_0 = \dots = 5 \cdot 2^n \begin{pmatrix} 8 + 3n \\ 2 + 3n \end{pmatrix},$$

in other words, in year n there are $2^n(40 + 15n)$ rabbits and $2^n(10 + 15n)$ foxes. Note that the solution gives qualitative information about what happens to the populations in the long run; in this example we get $\frac{r_n}{f_n} \rightarrow 1$ as $n \rightarrow \infty$, so in the long run there will be roughly as many foxes as rabbits.

△

Obviously a model like the one in the example will not match reality exactly, for example, if there are too many foxes, r_k can turn negative which doesn't make sense. Also, both populations will tend towards infinity here which is not possible in reality. However, such models can provide a good first approximation of a real dynamical system, and they may be useful for reasonably small populations.

Recursively defined sequences

Consider a simple recursively defined sequence $x_{n+1} = 5x_n$ with $x_0 = 3$. Since the next term is obtained by multiplication by 5, the explicit form of this sequence is clearly $x_n = 3 \cdot 5^n$.

The exact same formula actually works when we have several variables.

Example 4.4.2. Let us find an explicit formula for the sequence b_n defined by

$$b_n = 4b_{n-1} - 4b_{n-2}, \text{ where } b_0 = 0, b_1 = 1.$$

We introduce $X_n := \begin{pmatrix} b_{n+1} \\ b_n \end{pmatrix}$, $X_0 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, and $A = \begin{pmatrix} 4 & -4 \\ 1 & 0 \end{pmatrix}$. Then $X_{n+1} = AX_n$, so $X_n = A^n X_0$.

We Jordanize A and find that $A = SJS^{-1}$ with $S = \begin{pmatrix} 2 & 1 \\ 1 & 0 \end{pmatrix}$ and $J = \begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix}$.

But then we can calculate

$$X_n = A^n X_0 = (SJS^{-1})^n X_0 = SJ^n S^{-1} X_0 = S(2I + N)^n S^{-1} X_0.$$

Here we can expand $(2I + N)^n = (2I)^n + n(2I)^{n-1}N = 2^n I + n2^{n-1}N$ by the binomial theorem, as higher powers of N are zero. The expression above simplifies to

$$\begin{aligned} X_n &= S(2^n I + n2^{n-1}N)S^{-1} X_0 = \begin{pmatrix} 2 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 2^n & n2^{n-1} \\ 0 & 2^n \end{pmatrix} \begin{pmatrix} 0 & 1 \\ 1 & -2 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} 2 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 2^n & n2^{n-1} \\ 0 & 2^n \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 2 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} n2^{n-1} \\ 2^n \end{pmatrix} = \begin{pmatrix} (n+1)2^n \\ n2^{n-1} \end{pmatrix}, \end{aligned}$$

And since b_n is the bottom coordinate of X_n we obtain the explicit formula

$$b_n = n2^{n-1}.$$

△

Note that the same technique works if we just want to find a general solution to a recursively defined sequence without given starting values. In the above example, such a formula would look like $b_n = 2^n(C + Dn)$ for arbitrary scalars C, D .

Analytic matrix functions

Before looking at systems of differential equations we need to dig a bit deeper into the topic of matrix functions.

For a square matrix A , can we evaluate $f(A)$ when f is not a polynomial? For many functions the answer is yes, and this will turn out to be a useful thing to do.

Definition 4.4.3. Let $f : \mathbb{C} \rightarrow \mathbb{C}$ be a function whose Maclaurin-series converges for all x :

$$f(x) = f(0) + f'(0)x + \frac{f''(0)x^2}{2!} + \dots = \sum_{k=0}^{\infty} \frac{f^{(k)}(0)x^k}{k!}.$$

For a square matrix A we then define

$$f(A) := f(0)I + f'(0)A + \frac{f''(0)A^2}{2!} + \dots = \sum_{k=0}^{\infty} \frac{f^{(k)}(0)A^k}{k!}.$$

Note that $f(A)$ is an infinite sum of matrices, so here we should ask ourselves what it means for an infinite sum of matrices to converge. We will return for this question later when we talk about inner products and norms, but for now we shall think of it as element-wise convergence: If A_n is a sequence of matrices, we say that $A_n \rightarrow B$ as $n \rightarrow \infty$ if $(A_n)_{ij} \rightarrow B_{ij}$ for each position (i, j) . The condition that the Maclaurin-series converges for all x guarantees that we can evaluate $f(A)$ at any square matrix A . However, we can use the same definition for $f(A)$ for a larger class of functions, for example $\log(A)$ - but then this expression will only make sense for a certain class of matrices⁷.

Example 4.4.4. For the matrices $A = \begin{pmatrix} 4 & 0 \\ 0 & -5 \end{pmatrix}$ and $B = \begin{pmatrix} 0 & 2 & 0 \\ 0 & 0 & 3 \\ 0 & 0 & 0 \end{pmatrix}$, let us compute $e^A, e^B, \cos(A), \cos(B)$.

Recall that

$$e^x = 1 + x + \frac{x^2}{2} + \frac{x^3}{3!} + \dots \text{ and } \cos(x) = 1 - \frac{x^2}{2} + \frac{x^4}{4!} - \dots$$

Therefore

$$e^A = 1 + A + \frac{A^2}{2} + \frac{A^3}{3!} + \dots = \sum_{k=0}^{\infty} \frac{A^k}{k!} = \begin{pmatrix} \sum_{k=0}^{\infty} \frac{4^k}{k!} & 0 \\ 0 & \sum_{k=0}^{\infty} \frac{(-5)^k}{k!} \end{pmatrix} = \begin{pmatrix} e^4 & 0 \\ 0 & e^{-5} \end{pmatrix}.$$

Similarly, we get $\cos(A) = \begin{pmatrix} \cos(2) & 0 \\ 0 & \cos(-3) \end{pmatrix}$.

Since B is nilpotent with $B^3 = 0$, the computation becomes easy as all but the first terms of the series disappear:

$$e^B = I + B + \frac{B^2}{2} = \begin{pmatrix} 1 & 2 & 3 \\ 0 & 1 & 3 \\ 0 & 0 & 1 \end{pmatrix} \text{ and } \cos(B) = I - \frac{B^2}{2} = \begin{pmatrix} 1 & 0 & -3 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

△

Note that equalities of functions in one variable are preserved when evaluating them at a matrix, for example $\sin(2x) = 2\sin(x)\cos(x)$ are two ways of writing the same function, and the Maclaurin-series are preserved when taking sums, so the corresponding equality will hold if we replace x by any square matrix A .

Now we can always Jordanize a matrix as $A = SJS^{-1}$, and by factoring out S and S^{-1} from each term of the Maclaurin-expansion we get $f(A) = f(SJS^{-1}) = Sf(JS^{-1})$, so we only need to be able to evaluate f at a Jordan-matrix. Since the different Jordan-blocks do not interact when taking powers and sums of the matrix, it is in fact enough to consider how to evaluate f at a single Jordan block.

⁷Technically, the spectral radius $\max\{|\lambda| : \lambda \in \sigma(A)\}$ of A should be smaller than the radius of convergence for f , more on this later.

Proposition 4.4.5. Let $f : \mathbb{C} \rightarrow \mathbb{C}$ be a function whose Maclaurin-series convergent at every x as in Definition 4.4.3. For

$$J = J_n(\lambda) = \begin{pmatrix} \lambda & 1 & 0 & \cdots & 0 \\ 0 & \lambda & 1 & \ddots & 0 \\ 0 & 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \lambda & 1 \\ 0 & \cdots & \cdots & 0 & \lambda \end{pmatrix} \text{ we have } f(J) = \begin{pmatrix} \frac{f(\lambda)}{0!} & \frac{f'(\lambda)}{1!} & \frac{f''(\lambda)}{2!} & \cdots & \frac{f^{(n-1)}(\lambda)}{(n-1)!} \\ 0 & \frac{f(\lambda)}{0!} & \frac{f'(\lambda)}{1!} & \vdots & \frac{f^{(n-2)}(\lambda)}{(n-2)!} \\ 0 & 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \frac{f(\lambda)}{0!} & \frac{f'(\lambda)}{1!} \\ 0 & \cdots & \cdots & 0 & \frac{f(\lambda)}{0!} \end{pmatrix},$$

in other words, in $f(J)$, on the k 'th superdiagonal, we have the element $\frac{f^{(k)}(\lambda)}{k!}$.

Proof. Write $J = \lambda I + N$ and in each term of $f(J)$, expand $J^k = (\lambda I + N)^k$ using the binomial theorem. Then collect terms with the same term factor N^j and use the fact that $N^n = 0$, we omit the details. \square

Indeed, the proposition above is another common way to define $f(J)$.

Lemma 4.4.6. If A and B commute, we have

$$e^{A+B} = e^A e^B.$$

Proof. By definition we have

$$e^A e^B = \left(\sum_{i=0}^{\infty} \frac{A^i}{i!} \right) \left(\sum_{j=0}^{\infty} \frac{B^j}{j!} \right) = \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} \frac{A^i B^j}{i! j!}$$

Now reorder the terms so that terms with same total degree are written together. With $k := i + j$, the expression above can be further simplified to

$$= \sum_{k=0}^{\infty} \sum_{i=0}^k \frac{A^i B^{k-i}}{i!(k-i)!} = \sum_{k=0}^{\infty} \frac{1}{k!} \sum_{i=0}^k \binom{k}{i} A^i B^{k-i} = \sum_{k=0}^{\infty} \frac{1}{k!} (A+B)^k = e^{A+B}.$$

In the second to last step we used the fact that the binomial theorem holds for commuting matrices. \square

The numbers in our vectors or matrices may depend on some unknown parameter t . Then we can differentiate such matrices and vectors element-wise, for example:

$$\text{For } A(t) = \begin{pmatrix} t & 3 \\ t^2 & \sin(t) \end{pmatrix} \text{ we write } \frac{d}{dt} A(t) = A'(t) = \begin{pmatrix} 1 & 0 \\ 2t & \cos(t) \end{pmatrix}.$$

Lemma 4.4.7. For any square matrix A we have we have $\frac{d}{dt} e^{tA} = A e^{tA}$.

Proof.

$$\frac{d}{dt} e^{At} = \frac{d}{dt} \sum_{k=0}^{\infty} \frac{(tA)^k}{k!} = \sum_{k=0}^{\infty} \frac{d}{dt} t^k \frac{A^k}{k!} = \sum_{k=0}^{\infty} k t^{k-1} \frac{A^k}{k!}$$

The step where we differentiated term-wise may require some motivation, but since the Maclaurin-series converges anywhere, this step is fine. Now the first term in the sum is zero, so we make a change of variables and introduce $j := k - 1$, the sum above becomes:

$$= \sum_{j=0}^{\infty} (j+1) t^j \frac{A^{(j+1)}}{(j+1)!} = \sum_{j=0}^{\infty} t^j A \frac{A^j}{j!} = A \sum_{j=0}^{\infty} \frac{(tA)^j}{j!} = A e^{tA}.$$

\square

Note that we can easily calculate e^J where J is a Jordan block, since $J = J_n(\lambda) = \lambda I + N$ where N is the nilpotent matrix with ones on the superdiagonal. Since λI and N commute, by Lemma 4.4.6 we have

$$e^J = e^{\lambda I + N} = e^{\lambda I} e^N = (e^{\lambda I}) e^N = e^{\lambda} e^N,$$

and e^N is easy to compute, since $N^n = 0$, we have $e^N = I + N + \frac{N^2}{2} + \dots + \frac{N^{n-1}}{(n-1)!}$.

Since any Jordan matrix is a direct sum of such blocks, we can do the same thing in general in each block: write $J = D + N$ where $D = \text{diag}(\lambda_1, \dots, \lambda_n)$ is the diagonal of J , and N is nilpotent. Then

$$e^J = e^{D+N} = e^D e^N = \text{diag}(e^{\lambda_1}, \dots, e^{\lambda_n}) e^N.$$

Now let A be any square matrix. Jordanize A and write $A = SJS^{-1}$. Now

$$e^A = \sum_{k=0}^{\infty} \frac{A^k}{k!} = \sum_{k=0}^{\infty} \frac{(SJS^{-1})^k}{k!} = \sum_{k=0}^{\infty} \frac{S J^k S^{-1}}{k!} = S \left(\sum_{k=0}^{\infty} \frac{J^k}{k!} \right) S^{-1} = S e^J S^{-1},$$

which we can calculate with the previous method.

So with this method we can explicitly find e^A for any square matrix A .

Example 4.4.8. Let us compute e^A for $A = \begin{pmatrix} 5 & -1 \\ 1 & 3 \end{pmatrix}$. We first Jordanize A . We get $p_A(t) = \det(A - tI) = (t - 4)^2$ so the only eigenvalue is 4. We have $A - 4I = \begin{pmatrix} 1 & -1 \\ 0 & -1 \end{pmatrix}$, and $(A - 4I)^2 = 0$, so $\ker((A - 4I)^2) = \mathbb{C}^2$ and $\ker(A - 4I) = \text{span}(1, 1)$. So we will get a single chain of length 2, and as first vector we pick any vector not in $\ker(A - 4I)$, let's take $v_2 = (1, 0)$, and then we take $v_1 = (A - 4I)v_2 = (1, 1)$. We put v_1 and v_2 as columns in a matrix S and take J as the corresponding Jordan matrix, we then have $A = SJS^{-1}$ where

$$S = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \quad J = \begin{pmatrix} 4 & 1 \\ 0 & 4 \end{pmatrix} = \begin{pmatrix} 4 & 0 \\ 0 & 4 \end{pmatrix} + \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} = 4I + N \quad S^{-1} = \begin{pmatrix} 0 & 1 \\ 1 & -1 \end{pmatrix}.$$

Now we can evaluate

$$e^A = e^{SJS^{-1}} = S e^J S^{-1} = S e^{4I+N} S^{-1} = S e^J S^{-1} = S e^{4I} e^N S^{-1} = S e^4 I (I + N) S^{-1} = S e^4 (I + N) S^{-1}$$

$$e^4 \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} S^{-1} = e^4 \begin{pmatrix} 1 & 2 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ 1 & -1 \end{pmatrix} = e^4 \begin{pmatrix} 2 & -1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 2e^4 & -e^4 \\ e^4 & 0 \end{pmatrix}$$

△

We now expand on this example by solving a related system of differential equations:

Example 4.4.9. Let us find the general solutions to the linear system of differential equations

$$\begin{cases} x_1'(t) = 5x_1(t) - x_2(t) \\ x_2'(t) = x_1(t) + 3x_2(t), \end{cases}$$

and let us then in particular find the solution which also satisfies the boundary condition $x_1(0) = 3$, $x_2(0) = 5$.

Let $X(t) = \begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix}$ and let $A = \begin{pmatrix} 5 & -1 \\ 1 & 3 \end{pmatrix}$. Then the system can be written simply as $X'(t) = AX(t)$.

Now we claim that $X(t) = e^{tA}C$ is a solution for every 2×1 matrix C . Indeed, for $X(t) = e^{tA}C$ we have

$$X'(t) = \frac{d}{dt} e^{tA} C = A e^{tA} C = AX(t).$$

So to write down the general solution we only need to compute e^{tA} for the matrix A above. This exact matrix A was Jordanized in the previous problem, we had $A = SJS^{-1}$ where $S = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ $J =$

$$\begin{pmatrix} 4 & 1 \\ 0 & 4 \end{pmatrix} = 4I + N.$$

So

$$e^{tJ} = e^{4tI+tN} = e^{4tI} e^{tN} = e^{4t} I (I + tN) = e^{4t} \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} e^{4t} & t e^{4t} \\ 0 & e^{4t} \end{pmatrix}$$

so

$$e^{tA} = Se^{tJ}S^{-1} = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} e^{4t} & te^{4t} \\ 0 & e^{4t} \end{pmatrix} \begin{pmatrix} 0 & 1 \\ 1 & -1 \end{pmatrix} = e^{4t} \begin{pmatrix} t+1 & -t \\ t & 1-t \end{pmatrix}.$$

So an arbitrary solution can be written

$$X(t) = e^{tA}C = e^{4t} \begin{pmatrix} t+1 & -t \\ t & 1-t \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = e^{4t} \begin{pmatrix} c_1(t+1) - c_2t \\ c_1t + c_2(1-t) \end{pmatrix} = e^{4t} \begin{pmatrix} (c_1 - c_2)t + c_1 \\ (c_1 - c_2)t + c_2 \end{pmatrix}.$$

In particular, taking $C = \begin{pmatrix} 3 \\ 5 \end{pmatrix}$, we get $X(0) = e^{0A}C = IC = C = \begin{pmatrix} 3 \\ 5 \end{pmatrix}$, which is the solution satisfying our boundary condition, explicitly this solution is

$$\begin{cases} x_1(t) = (-2t + 3)e^{4t} \\ x_2(t) = (-2t + 5)e^{4t} \end{cases}.$$

△

The method in the above example works in general. We state the result as a proposition:

Proposition 4.4.10. *An $n \times n$ system of linear differential equations*

$$\begin{cases} x_1'(t) = a_{11}x_1(t) + a_{12}x_2(t) + \cdots + a_{1n}x_n(t) \\ x_2'(t) = a_{21}x_1(t) + a_{22}x_2(t) + \cdots + a_{2n}x_n(t) \\ \quad \quad \quad \vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \\ x_n'(t) = a_{n1}x_1(t) + a_{n2}x_2(t) + \cdots + a_{nn}x_n(t) \end{cases}$$

can be written as $X'(t) = AX(t)$ where $A = (a_{ij})$ and $X(t) = (x_1(t), \dots, x_n(t))^T$.

The general solution to this system is $X(t) = e^{tA}C$ where $C = (c_1, \dots, c_n)^T$ is an arbitrary vector. We note that $X(0) = C$.

Chapter 5

Inner product spaces

5.1 Inner products

In a vector space we only have the operations of adding vectors and multiplying vectors by scalars. Although we usually visualize a 2-dimensional vector space as a plane, we haven't defined the concept of distances and angles between vectors yet. For more abstract vector spaces such as the vector space $\text{Mat}_n(\mathbb{C})$ of matrices, or $\mathcal{P}(\mathbb{R})$ of polynomials, it is not even clear how we should define lengths and angles. So in order to be able to do some geometry in the vector space setting we need to introduce these concepts. It turns out that all you need is an *inner product*.

Definition 5.1.1. Let V be a complex vector space. A function $\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{C}$ is called an **inner product** or a **scalar product** on V if for all $u, v, w \in V$ and $\lambda, \mu \in \mathbb{C}$ we have:

1. $\langle \lambda u + \mu v, w \rangle = \lambda \langle u, w \rangle + \mu \langle v, w \rangle$
2. $\langle u, v \rangle = \overline{\langle v, u \rangle}$
3. $\langle v, v \rangle \geq 0$ with equality if and only if $v = 0$.

A vector space equipped with an inner product is called a (complex) **inner product space**. If we replace \mathbb{C} by \mathbb{R} we get the definition of a *real inner product space*, also called a *Euclidean space*.

A few remarks:

- By applying the second axiom we get $\langle v, v \rangle = \overline{\langle v, v \rangle}$ which shows that $\langle v, v \rangle$ is a real number, and the third condition makes sense¹.
- Sometimes we want to talk about different inner products on the same space, other common notations for the inner product of two vectors include (u, v) , $\langle u, v \rangle_1$, $(u|v)$, $u \bullet v$, etcetera.
- Axioms (1) and (2) imply that inner products are conjugate-linear in the second argument:

$$\langle u, \lambda v \rangle = \bar{\lambda} \langle u, v \rangle.$$

So inner products are linear in the first argument, and satisfies half of the axioms for being linear in the second argument, for this reason inner products are called *sesqui-linear forms*, "one and a half"-linear forms.

- In some contexts, such as in quantum mechanics, but *not in this course*, axiom (1) is replaced by linearity in the second argument:

$$\langle w, \lambda u + \mu v \rangle = \lambda \langle w, u \rangle + \mu \langle w, v \rangle,$$

this is just a convention and the theory will be analogous.

¹In general, when for $z \in \mathbb{C}$ we write $z > 0$, we mean that z is *real and positive*.

- In some branches of physics, like in Lorenzian geometry, Pseudo-Riemannian geometry and when studying Minkowski spacetime, axiom (3) is relaxed to

$$\langle u, v \rangle = 0 \forall u \Rightarrow v = 0,$$

which allows $\langle v, v \rangle < 0$ for some vectors.

We have a number of *standard inner products* on some of our familiar vector spaces:

Example 5.1.2. The following functions are inner products on the respective vector space:

- On \mathbb{C}^n we have the inner product $\langle u, v \rangle = \langle (x_1, \dots, x_n), (y_1, \dots, y_n) \rangle = x_1 \overline{y_1} + \dots + x_n \overline{y_n}$. This is also called the *dot-product* and can also be written $u \bullet v$, it is defined the same way on \mathbb{R}^n .
- On $\text{Mat}_{m \times n}(\mathbb{C})$, we have the *Frobenius inner product* $\langle A, B \rangle = \text{tr}(AB^*)$. Although this looks different than the dot-product, upon closer inspection it really is the same thing: element-wise products of entries of A with the conjugate of the corresponding elements of B : $\text{tr}(AB^*) = \sum_{i,j} a_{i,j} \overline{b_{i,j}}$.
- On $\mathcal{C}[a, b]$, the space of continuous functions from $[a, b]$ to \mathbb{C} (or to \mathbb{R}), we have $\langle f(x), g(x) \rangle = \int_a^b f(x) \overline{g(x)} dx$. The same works on subspaces, such as the polynomials.

It is straight forward to verify that the axioms of Definition 5.1.1 hold for these inner products.

△

We shall assume our vector spaces are equipped with these standard inner products unless otherwise stated.

Definition 5.1.3. Let V be an inner product space with inner product $\langle \cdot, \cdot \rangle$ and let $\|\cdot\|$ be the norm derived from the inner product: $\|v\| := \sqrt{\langle v, v \rangle}$. Then for $u, v \in V$:

- We define the **length** of v , also called the **norm** of v , to be $\|v\| = \sqrt{\langle v, v \rangle}$.
- We define the **distance** between u and v to be $\|u - v\|$.
- We define the **angle** between u and v to be $\arccos\left(\frac{\langle u, v \rangle}{\|u\| \|v\|}\right)$ (when V is a real inner product space).
- We say that u and v are **orthogonal** if $\langle u, v \rangle = 0$.

Example 5.1.4. The space $\mathcal{P}(\mathbb{R})$ of polynomials with real coefficients becomes a real inner product space when equipped with the inner product

$$\langle p(x), q(x) \rangle := \int_0^1 p(x)q(x)dx.$$

Let us find the angle θ between the polynomials 1 and x with respect to this inner product.

We have

$$\langle 1, 1 \rangle = \int_0^1 1 dx = [x]_0^1 = 1, \quad \langle 1, x \rangle = \int_0^1 x dx = \left[\frac{x^2}{2}\right]_0^1 = \frac{1}{2}, \quad \langle x, x \rangle = \int_0^1 x^2 dx = \left[\frac{x^3}{3}\right]_0^1 = \frac{1}{3},$$

So

$$\theta = \arccos\left(\frac{\langle 1, x \rangle}{\|1\| \cdot \|x\|}\right) = \arccos\left(\frac{\langle 1, x \rangle}{\sqrt{\langle 1, 1 \rangle} \cdot \sqrt{\langle x, x \rangle}}\right) = \arccos\left(\frac{\frac{1}{2}}{1 \cdot \frac{1}{\sqrt{3}}}\right) = \arccos\left(\frac{\sqrt{3}}{2}\right) = \frac{\pi}{6}.$$

△

We summarize a number of direct consequences of the definitions:

Proposition 5.1.5. Let V be an inner product space with inner product $\langle \cdot, \cdot \rangle$ and let $\|v\| := \sqrt{\langle v, v \rangle}$ be the norm derived from the inner product. Then for $u, v \in V$ and $\lambda \in \mathbb{C}$ we have:

1. $\|\lambda v\| = |\lambda| \cdot \|v\|$.

2. $\|u + v\| \leq \|u\| + \|v\|$ (the triangle inequality)
3. $\|v\| \geq 0$ with equality if and only if $v = 0$.
4. $|\langle u, v \rangle| \leq \|u\| \cdot \|v\|$ with equality if and only if $u\|v$ (the Cauchy-Schwartz inequality)
5. $\|u + v\|^2 + \|u - v\|^2 = 2\|u\|^2 + 2\|v\|^2$ (the parallelogram identity)
6. $4\langle u, v \rangle = \begin{cases} \|u + v\|^2 - \|u - v\|^2 & \text{if } V \text{ is real.} \\ \|u + v\|^2 - \|u - v\|^2 + i\|u + iv\|^2 - i\|u - iv\|^2 & \text{if } V \text{ is complex.} \end{cases}$

We leave the proofs as an exercise.

Note that the Cauchy-Schwarz inequality guarantees that the definition of angles in Definition 5.1.3 makes sense. The last formulas in (6) are called *polarization identities*, these show that the inner product can be calculated from the norm.

5.2 Norms

In some contexts, we want to define norms without necessarily having a corresponding inner product. It turns out that the first three properties of Proposition 5.1.5 are the key to make norms useful as a concept.

Definition 5.2.1. Let V be a complex vector space. A **norm** on V is a function $\|\cdot\| : V \rightarrow \mathbb{R}$ such that for $u, v \in V$ and $\lambda \in \mathbb{C}$ we have:

1. $\|\lambda v\| = |\lambda| \cdot \|v\|$
2. $\|u + v\| \leq \|u\| + \|v\|$
3. $\|v\| \geq 0$ with equality if and only if $v = 0$

The definition is the same if V is a real vector space, except that $\lambda \in \mathbb{R}$. A vector space equipped with a norm is called a **normed space**. We still define the length of v as $\|v\|$ and the distance between u and v as $\|u - v\|$ in any normed space.

Norms appear frequently in functional analysis and physics when investigating some important infinite-dimensional classes of vector spaces such as Banach-spaces² and Hilbert spaces³.

Proposition 5.1.5 shows that every inner product gives rise to a norm, for example, the Frobenius-norm on $\text{Mat}_n(\mathbb{C})$ is $\|A\|_F = \sqrt{\sum |a_{ij}|^2}$. But the opposite is *not true*, there are norms that do not come from any inner product. Here are a few examples:

- On \mathbb{C}^n the **maximum norm** is defined as

$$\|(x_1, \dots, x_n)\|_{\max} = \|(x_1, \dots, x_n)\|_{\infty} = \max\{|x_1|, \dots, |x_n|\}.$$

- On \mathbb{C}^n the **Manhattan norm** or the **taxicab norm** is defined as

$$\|(x_1, \dots, x_n)\|_{\text{Mh}} = |x_1| + \dots + |x_n|.$$

- More generally, for $p \geq 1$, the **p-norm** on \mathbb{C}^n is defined as

$$\|(x_1, \dots, x_n)\|_p = (|x_1|^p + \dots + |x_n|^p)^{\frac{1}{p}}.$$

For $p \neq 2$ this does not correspond to an inner product.

²A Banach space is a complete normed vector space. "Complete" means that every Cauchy-sequence v_n converges to some v with respect to the norm: $\|v_n - v\| \rightarrow 0$.

³A Hilbert space is an inner product space, which is complete with respect to the norm induced from the inner product.

- The **spectral norm** on $\text{Mat}_{m \times n} \mathbb{C}$ is defined⁴ as

$$\|A\|_{\sigma} = \max\{\sqrt{\lambda} \mid \lambda \in \sigma(A^*A)\}.$$

- The **operator norm** of a linear operator between inner product spaces $F : V \rightarrow W$ is defined as

$$\|F\|_{\text{op}} = \max\{\|F(v)\| \mid \|v\| = 1\}.$$

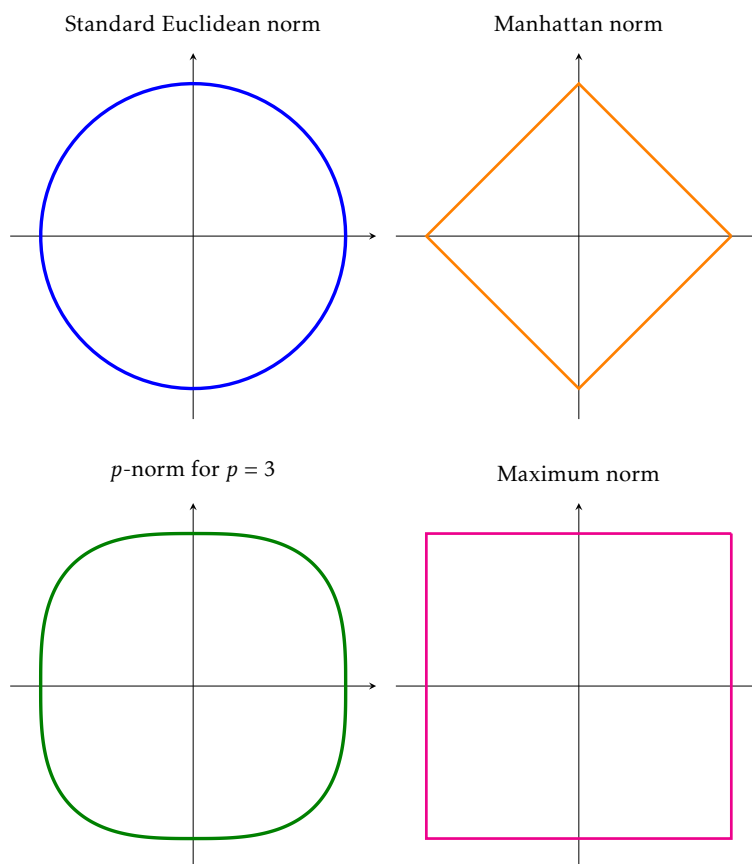
- The **supremum norm**⁵ on $\mathcal{C}[a, b]$ is defined as

$$\|f(x)\|_{\text{sup}} = \max\{|f(x)|; x \in [a, b]\}.$$

These all satisfy the conditions in Definition 5.2.1. In fact many of these are the p -norm $\|\cdot\|_p$ in disguise: The standard norm is $\|\cdot\|_2$, the Manhattan norm is $\|\cdot\|_1$, and the maximum norm is $\|\cdot\|_{\infty}$ in the sense that $\|v\|_{\text{max}} = \lim_{p \rightarrow \infty} \|v\|_p$.

Why should we consider different norms, isn't the distance between two points objective? Well, depending on the context it will make sense to measure the distance in different ways. For example, if you are on Manhattan, the streets form an orthogonal grid, and the distance you need to travel to go between the points is the sum of the vertical and the horizontal distances between the points (since you cannot drive diagonally through buildings). For this reason the Manhattan norm is appropriate in this context. In another setting we might be able to move vertically and horizontally independently with the same speed, such as a king on a chessboard - in this setting the maximum norm is the appropriate way to measure distance.

The point is, different norms give rise to different geometries on the vector space. For example, here is how the unit circle $\{v \in \mathbb{R}^2 : \|v\| = 1\}$ looks for different choices of norms $\|\cdot\|$ on \mathbb{R}^2 :



What norms can be derived from inner products? The answer is given in:

⁴For square matrices A satisfying $A^*A = AA^*$, this is the same as the *spectral radius* of A , the largest absolute value of the eigenvalues of A . We will discuss this in more detail later.

⁵The analogous definition works on any vector space of continuous functions from some compact subset $X \subset \mathbb{C}^n$ to \mathbb{C} .

Theorem 5.2.2. Let V be a vector space equipped with a norm $\|\cdot\|$. There exists an inner product $\langle \cdot, \cdot \rangle$ on V such that $\|v\|^2 = \langle v, v \rangle$ if and only if the norm satisfies the parallelogram identity:

$$\|u + v\|^2 + \|u - v\|^2 = 2\|u\|^2 + 2\|v\|^2 \text{ for all } u, v \in V.$$

Proof. The idea of the proof is given by the last point of Proposition 5.1.5: The only option is to define $\langle u, v \rangle = \frac{1}{4}(\|u + v\|^2 - \|u - v\|^2)$ (when V is a real vector space), and then show that this function is an inner product if and only if the parallelogram law holds. It is quick to check that axioms (1) and (2) of Definition 5.1.1 follows from the norm axioms, but showing that $\langle \cdot, \cdot \rangle$ is linear in the first argument takes some work. See the literature for details. \square

As soon as we have a norm on a vector space we can talk about the concept of convergence in V with respect to the norm.

Definition 5.2.3. Let $\|\cdot\|$ be a norm on V . We say that a sequence of vectors v_1, v_2, \dots in V **converges** to some vector $v \in V$ with respect to the norm $\|\cdot\|$ if and only if

$$\|v_n - v\| \rightarrow 0 \text{ as } n \rightarrow \infty.$$

This just means that for each $\epsilon > 0$ there exists N such that $\|v_n - v\| < \epsilon$ for $n > N$.

Here we might ask ourselves if it is possible that the sequence v_n converges with respect to one norm and diverges with respect to another. When V is finite-dimensional such problems do not arise:

Proposition 5.2.4. If V is finite-dimensional and $v_n \rightarrow v$ with respect to some norm, then $v_n \rightarrow v$ with respect to every norm.

This lets us choose norms freely when trying show that a sequence of vectors converges.

5.3 Orthonormal bases and projections

In this section V is a real or complex inner product space of finite dimension unless otherwise stated, we write $\langle \cdot, \cdot \rangle$ for the inner product and $\|\cdot\|$ for the norm derived from the inner product.

Definition 5.3.1. A basis (e_1, \dots, e_n) for V is called an **orthonormal basis** or an **ON-basis**, if the basis vectors have length 1 and are pairwise orthogonal:

$$\langle e_i, e_j \rangle = \delta_{ij}.$$

If we drop the condition that the lengths have to be 1, we call the basis an **orthogonal basis**.

The same definition works in infinite-dimensional spaces.

Orthonormal bases make many calculations easier: if (e_1, \dots, e_n) is an ON-basis, we have

$$v = \langle v, e_1 \rangle e_1 + \dots + \langle v, e_n \rangle e_n.$$

We know that every subspace of a vector space has a complement; for $U \subset V$ we can always find another subspace U' such that $U \oplus U' = V$. In an inner product space there is a canonical choice of the complement U^\perp :

Definition 5.3.2. For a subspace $U \subset V$ we define its *orthogonal complement* to U as

$$U^\perp := \{v \in V \mid \langle v, u \rangle = 0 \text{ for all } u \in U\}.$$

So U^\perp is the set of vectors that are orthogonal to every vector in U . It is not too hard to show that $V = U \oplus U^\perp$ and that $(U^\perp)^\perp = U$ when V is finite-dimensional.

The canonical choice of complement also gives canonical choices for projection maps onto subspaces:

Definition 5.3.3. If U is a subspace of V we define the map $P_U : V \rightarrow V$ by $P_U(v) = u$ where $v = u + u'$ is the unique expression of v as a sum of $u \in U$ and $u' \in U^\perp$.

If $U = \text{span}(u)$ is one-dimensional, P_U can be calculated explicitly by the familiar projection formula^a

$$P_U(v) = P_u(v) = \frac{\langle v, u \rangle}{\langle u, u \rangle} u,$$

Otherwise, if U has higher dimension, we can pick an orthogonal basis (u_1, \dots, u_m) of U , and then we explicitly have

$$P_U(v) = P_{u_1}(v) + \dots + P_{u_m}(v) = \frac{\langle v, u_1 \rangle}{\langle u_1, u_1 \rangle} u_1 + \dots + \frac{\langle v, u_m \rangle}{\langle u_m, u_m \rangle} u_m.$$

We call $P_U(v)$ the **(orthogonal) projection** of v onto the subspace U (or onto the vector u if $U = \text{span}(u)$).

^aNote the order of the vectors in $\langle v, u \rangle$, for complex vector spaces these cannot be switched.

It is not hard to show that $P_U^2 = P_U$ and that $P_U(v)$ is the vector in U with minimal distance to v . Note that it is important that the basis u_1, \dots, u_n of the subspace we are projecting onto are indeed pairwise orthogonal, otherwise the result will be wrong.

We also remark that $P_v(u) = P_{\lambda v}(u)$ whenever λ is a nonzero complex number so the length of the vector we project onto is irrelevant, and when projecting onto a subspace it is enough to project onto the vectors of an orthogonal basis of the subspace.

5.4 Gram-Schmidt

The Gram-Schmidt process is an algorithm for converting a basis for a finite-dimensional inner product space (or a subspace) into an orthonormal basis for the same space.

Theorem 5.4.1. (Gram-Schmidt) Let v_1, \dots, v_n be linearly independent vectors in an inner product space V , and let $U = \text{span}(v_1, \dots, v_n)$.

$$\begin{aligned} e_1 &= v_1 && = v_1 \\ e_2 &= v_2 - P_{e_1}(v_2) && = v_2 - \frac{\langle v_2, e_1 \rangle}{\langle e_1, e_1 \rangle} e_1 \\ e_3 &= v_3 - P_{e_1}(v_3) - P_{e_2}(v_3) && = v_3 - \frac{\langle v_3, e_1 \rangle}{\langle e_1, e_1 \rangle} e_1 - \frac{\langle v_3, e_2 \rangle}{\langle e_2, e_2 \rangle} e_2 \\ &\vdots && \vdots \\ e_n &= v_n - P_{e_1}(v_n) - \dots - P_{e_{n-1}}(v_n) && = v_n - \frac{\langle v_n, e_1 \rangle}{\langle e_1, e_1 \rangle} e_1 - \dots - \frac{\langle v_n, e_{n-1} \rangle}{\langle e_{n-1}, e_{n-1} \rangle} e_{n-1} \end{aligned}$$

Then (e_1, \dots, e_n) is an orthogonal basis for U . If we normalize it and define $f_i := \frac{1}{\|e_i\|} e_i$, the vectors (f_1, \dots, f_n) will be an orthonormal basis for U .

Proof. Since each e_k is defined as $v_k - u$ for some u in the span of the previous vectors e_i , we get $\text{span}(v_1, \dots, v_k) = \text{span}(e_1, \dots, e_k)$ for each $1 \leq k \leq n$. So we need only verify that the vectors e_i and e_k are orthogonal for $i \neq k$. Without loss of generality we can assume $i < k$. We do this by induction. Assume that the vectors e_i are pairwise orthogonal for $i < k$. We shall prove that e_k is orthogonal to each of these vectors e_i . By the definition of e_k in the Gram-Schmidt process we have

$$\begin{aligned} \langle e_k, e_i \rangle &= \left\langle v_k - \frac{\langle v_k, e_1 \rangle}{\langle e_1, e_1 \rangle} e_1 - \dots - \frac{\langle v_k, e_{k-1} \rangle}{\langle e_{k-1}, e_{k-1} \rangle} e_{k-1}, e_i \right\rangle = \left\langle v_k - \sum_{j=0}^{k-1} \frac{\langle v_k, e_j \rangle}{\langle e_j, e_j \rangle} e_j, e_i \right\rangle \\ &= \langle v_k, e_i \rangle - \sum_{j=0}^{k-1} \frac{\langle v_k, e_j \rangle}{\langle e_j, e_j \rangle} \langle e_j, e_i \rangle = \langle v_k, e_i \rangle - \frac{\langle v_k, e_i \rangle}{\langle e_i, e_i \rangle} \langle e_i, e_i \rangle = \langle v_k, e_i \rangle - \langle v_k, e_i \rangle = 0. \end{aligned}$$

□

It is also possible to normalize each e_i at each step in Gram-Schmidt. This produces the ON-basis f_i directly, but the calculations of the projections will typically involve square roots even if we start with integer-vectors.

We also remark that the Gram-Schmidt process still can be applied if the set of vectors is linearly dependent, then some e_i will be zero, but if we remove these we will end up with an ON-basis of $\text{span}(v_1, \dots, v_n)$.

Example 5.4.2. Consider the vector space $\mathcal{C}[0, 1]$ of real-valued continuous functions on the unit interval with the standard inner product $\langle f, g \rangle = \int_0^1 f(x)g(x)dx$.

In this example we shall use Gram-Schmidt to construct an ON-basis for \mathcal{P}_1 , the subspace of polynomials of degree ≤ 1 (with domains restricted to $[0, 1]$). We shall then use this basis to find the function in \mathcal{P}_1 closest to $f(x) = x^2$.

Since $\mathcal{P}_1 = \text{span}(1, x)$, we apply Gram-Schmidt to convert the basis $(v_1, v_2) = (1, x)$ of \mathcal{P}_1 to an orthogonal basis (e_1, e_2) of \mathcal{P}_1 . By the algorithm we should take $e_1 = v_1 = 1$ and we note that $\|e_1\|^2 = \langle e_1, e_1 \rangle = 1$, so we get

$$e_2 = v_2 - P_{e_1}(v_2) = v_2 - \frac{\langle v_2, e_1 \rangle}{\langle e_1, e_1 \rangle} e_1 = x - \frac{\int_0^1 x \cdot 1 dx}{\int_0^1 1 \cdot 1 dx} 1 = x - \frac{1}{2}$$

and we note that $\|e_2\|^2 = \langle e_2, e_2 \rangle = \int_0^1 (x - \frac{1}{2})^2 dx = \frac{1}{12}$. So we conclude that

$$(e_1, e_2) = \left(1, x - \frac{1}{2}\right)$$


is an orthogonal basis for \mathcal{P}_1 . The second basis vector does not have length 1 though, so if we need an ON-basis we should also normalize. So with $f_1 = e_1$, and $f_2 = \frac{1}{\|e_2\|} e_2$ we get an ON-basis for \mathcal{P}_1 :

$$(f_1, f_2) = \left(1, \sqrt{3}\left(x - \frac{1}{2}\right)\right).$$

Now the function $g(x)$ in \mathcal{P}_1 closest to $f(x) = x^2$ is the projection of x^2 onto \mathcal{P}_1 , and this can be found by projecting onto each of our basis vectors f_1, f_2 in our ON-basis. To avoid square roots in our calculations we might as well project onto e_1 and e_2 instead:

$$g(x) = P_{\mathcal{P}_1}(x^2) = P_{f_1}(x^2) + P_{f_2}(x^2) = \frac{\langle x^2, 1 \rangle}{\langle 1, 1 \rangle} \cdot 1 + \frac{\langle x^2, (x - \frac{1}{2}) \rangle}{\langle (x - \frac{1}{2}), (x - \frac{1}{2}) \rangle} \cdot (x - \frac{1}{2}) = \frac{1}{3} + (x - \frac{1}{2}) = x - \frac{1}{6}.$$

So $g(x) = x - \frac{1}{6}$ is the function in \mathcal{P}_1 closest to x^2 . Concretely, by our definition of the inner product, this means that $\|x^2 - (x - \frac{1}{6})\|^2 = \int_0^1 (x^2 - (x - \frac{1}{6}))^2 dx$ is minimized, which we can think of as minimizing the area between the graphs of $x - \frac{1}{6}$ and x^2 on $[0, 1]^a$.



So in a sense $x - \frac{1}{2}$ is the "best approximation" of x^2 by a line if we only care about how the functions behave on the unit interval. If wanted to find the best approximation for another interval, we could perform the analogous computation for a different choice of inner product.

^aTechnically the area under the squared difference of the functions is being minimized.

△

The idea of the example still works if we replace polynomials by another class of functions. Projection onto the subspace produces an approximation of a given function as a linear combination of functions from this class. For example, in *Fourier-analysis*, we approximate periodic functions by *trigonometric functions*, as illustrated by the following example.

Example 5.4.3. Consider the space of all real 2π -periodic functions equipped with the inner product^a

$$\langle f(x), g(x) \rangle = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x)g(x)dx.$$

Let $\mathcal{F}_n = \text{span}(\sin(x), \sin(2x), \dots, \sin(nx))$. Let us find the function $g(x) \in \mathcal{F}_3$ which best approximates the *square wave function* $f(x)$, which is defined as $\text{sgn}(x)$ on $[-\pi, \pi)$ and is 2π -periodic.

One can check that $\{\sin(kx)\}_{k \in \mathbb{N}}$ is actually an orthonormal set of functions with respect to our inner product, so we obtain our approximation immediately as


$$g(x) = P_{\mathcal{F}_3}(f(x)) = \langle f(x), \sin(x) \rangle \sin(x) + \langle f(x), \sin(2x) \rangle \sin(2x) + \langle f(x), \sin(3x) \rangle \sin(3x)$$

For general k we calculate $\langle f(x), \sin(kx) \rangle$. Since $f(x) \sin(kx)$ is an even function and $f(x) = 1$ for $x \in (0, \pi)$ we get

$$\begin{aligned} \langle f(x), \sin(kx) \rangle &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \sin(kx) dx = \frac{2}{\pi} \int_0^{\pi} \sin(kx) dx = \frac{2}{\pi} \left[\frac{-\cos(kx)}{k} \right]_0^{\pi} \\ &= \frac{2(1 + \cos(k\pi))}{k\pi} = \frac{2(1 - (-1)^k)}{k\pi} = \begin{cases} \frac{4}{k\pi} & \text{for odd } k \\ 0 & \text{for even } k. \end{cases} \end{aligned}$$

This shows that

$$g(x) = \langle f(x), \sin(x) \rangle \sin(x) + \langle f(x), \sin(2x) \rangle \sin(2x) + \langle f(x), \sin(3x) \rangle \sin(3x) = \frac{4}{\pi} \sin(x) + \frac{4}{3\pi} \sin(3x).$$



Our calculation actually directly tells us the projection of the square wave function onto \mathcal{F}_{2m+1} for arbitrary m :

$$P_{\mathcal{F}_{2m+1}}(f(x)) = \sum_{k=0}^m \frac{4}{k\pi} \sin((2k+1)x).$$

Higher m gives better approximations of the function.

^aSome technical restrictions of the whole space of functions is needed in order for the integral in the inner product to exist, let us assume that all functions are continuous except for finitely many points on $[-\pi, \pi)$. We shall also consider two functions in the space to be equal if they differ for only finitely many points on that interval (otherwise the inner product would not be positive definite).

△

Approximating functions by trigonometric functions has plentiful applications in signal-processing, audio-compression, etcetera.

5.5 QR-decomposition

Definition 5.5.1. A *QR-factorization* or *QR-decomposition* of a matrix A is a factorization of form

$$A = QR$$

R is an upper-triangular matrix and Q is a square matrix satisfying $Q^*Q = I$.

The condition $Q^*Q = I$ is equivalent to saying that the columns for Q form an ON-basis with respect to the standard inner product on \mathbb{C}^n , such matrices are called *unitary*, we will return to these soon.

If A is an invertible matrix, A has a *unique* QR factorization if we require that the diagonal entries of R are real positive.

The QR-factorization can be obtained by applying Gram-Schmidt⁶ to the columns of the matrix as the following example demonstrates:

Example 5.5.2. Let us find the QR-factorization of the matrix

$$A = \begin{pmatrix} 1 & 3 & 1 \\ 1 & 1 & 2 \\ 0 & 2 & 2 \end{pmatrix}.$$

We are seeking a matrix Q satisfying $QQ^* = I$, and an upper triangular matrix R such that $A = QR$.

We apply Gram-Schmidt to the columns of A with respect to the standard inner product on \mathbb{C}^3 , call the columns v_1, v_2, v_3 so that $A = (v_1 \ v_2 \ v_3)$.

We obtain $e_1 = v_1 = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}$, for which we have $\|e_1\|^2 = \langle e_1, e_1 \rangle = 2$, so we take $f_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}$.

Next we get

$$e_2 = v_2 - \frac{\langle v_2, e_1 \rangle}{\langle e_1, e_1 \rangle} e_1 = v_2 - 2e_1 = \begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix} - 2 \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ -1 \\ 2 \end{pmatrix},$$

for which $\|e_2\|^2 = \langle e_2, e_2 \rangle = 6$, so we take so we take $f_2 = \frac{1}{\sqrt{6}} \begin{pmatrix} 1 \\ -1 \\ 2 \end{pmatrix}$. And finally

$$e_3 = v_3 - \frac{\langle v_3, e_1 \rangle}{\langle e_1, e_1 \rangle} e_1 - \frac{\langle v_3, e_2 \rangle}{\langle e_2, e_2 \rangle} e_2 = v_3 - \frac{3}{2} e_1 - \frac{1}{2} e_2 = \begin{pmatrix} 1 \\ 2 \\ 2 \end{pmatrix} - \frac{3}{2} \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} - \frac{1}{2} \begin{pmatrix} 1 \\ -1 \\ 2 \end{pmatrix} = \begin{pmatrix} -1 \\ 1 \\ 1 \end{pmatrix},$$

and we have $\|e_3\|^2 = \langle e_3, e_3 \rangle = 3$ so we take $f_3 = \frac{1}{\sqrt{3}} \begin{pmatrix} -1 \\ 1 \\ 1 \end{pmatrix}$.

The steps in Gram-Schmidt give us a way to express our original vectors v_1, v_2, v_3 as linear combinations of our new basis vectors f_1, f_2, f_3 . We had:

⁶However, the Gram-Schmidt algorithm is unstable numerically, so modern computer algorithms use different methods to find QR-factorizations.

$$\begin{cases} v_1 = e_1 & = \sqrt{2}f_1 \\ v_2 = 2e_1 + e_2 & = 2\sqrt{2}f_1 + \sqrt{6}f_2 \\ v_3 = \frac{3}{2}e_1 + \frac{1}{2}e_2 + e_3 & = \frac{3\sqrt{2}}{2}f_1 + \frac{\sqrt{6}}{2}f_2 + \sqrt{3}f_3 \end{cases}$$

which can be expressed in matrix form

$$\begin{aligned} A &= \begin{pmatrix} 1 & 3 & 1 \\ 1 & 1 & 2 \\ 0 & 2 & 2 \end{pmatrix} = \begin{pmatrix} | & | & | \\ v_1 & v_2 & v_3 \\ | & | & | \end{pmatrix} = \begin{pmatrix} | & | & | \\ f_1 & f_2 & f_3 \\ | & | & | \end{pmatrix} \begin{pmatrix} \sqrt{2} & 2\sqrt{2} & \frac{3\sqrt{2}}{2} \\ 0 & \sqrt{6} & \frac{\sqrt{6}}{2} \\ 0 & 0 & \sqrt{3} \end{pmatrix} \\ &= \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{6}} & -\frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{6}} & \frac{1}{\sqrt{3}} \\ 0 & \frac{2}{\sqrt{6}} & \frac{1}{\sqrt{3}} \end{pmatrix} \begin{pmatrix} \sqrt{2} & 2\sqrt{2} & \frac{3\sqrt{2}}{2} \\ 0 & \sqrt{6} & \frac{\sqrt{6}}{2} \\ 0 & 0 & \sqrt{3} \end{pmatrix} = QR. \end{aligned}$$

△

In the example we obtained our matrix R by keeping track of the coefficients relating our original vectors v_i to our new vectors f_i . However, note that if $A = QR$ is a QR-decomposition, we have

$$R = IR = Q^*QR = Q^*A,$$

so we can just perform Gram-Schmidt to obtain the matrix Q where the new ON-basis are columns, and then use it to compute $R = Q^*A$. If we followed Gram-Schmidt correctly, R will be upper triangular.

The method in the example works in general for finding a QR-factorization. Although the columns of an $m \times n$ -matrix are vectors of \mathbb{C}^m , they may not span \mathbb{C}^m , but we can always use Gram-Schmidt to obtain an ON-basis for their span⁷, and then we extend this to an ON-basis of the full space \mathbb{C}^m , then we take Q as the square matrix with this ON-basis as columns, and $R = Q^*A$ will be upper triangular. We summarize:

Proposition 5.5.3. Any matrix $A \in \text{Mat}_{m \times n}(\mathbb{C})$ has a QR-factorization $A = QR$ with Q unitary ($Q^*Q = I$) and R upper triangular. If $m \geq n$ we can express this in block form as

$$A = QR = (Q_1 \mid Q_2) \begin{pmatrix} R_1 \\ 0 \end{pmatrix} = Q_1 R_1,$$

where $Q \in \text{Mat}_m(\mathbb{C})$ is a unitary and R is upper triangular. In the block-decomposition, Q_1 contains first n columns of Q , and R_1 is of shape $n \times n$.

The more compact version $A = Q_1 R_1$ is sometimes called a **thin** QR-decomposition (but note that Q_1 is typically not unitary since it is not square).

Example 5.5.4. Let $A = \begin{pmatrix} 1 & 5 \\ 2 & 7 \\ 2 & 4 \end{pmatrix}$. We apply Gram-Schmidt to the columns and get a matrix $Q_1 = \frac{1}{3} \begin{pmatrix} 1 & 2 \\ 2 & 1 \\ 2 & -2 \end{pmatrix}$ whose columns are orthonormal and span the same space as the columns of A . We then take $R_1 = Q_1^* A = \begin{pmatrix} 3 & 9 \\ 0 & 3 \end{pmatrix}$. We now have the thin QR-factorization

$$A = Q_1 R_1 = \frac{1}{3} \begin{pmatrix} 1 & 2 \\ 2 & 1 \\ 2 & -2 \end{pmatrix} \begin{pmatrix} 3 & 9 \\ 0 & 3 \end{pmatrix}.$$

To get the full QR-factorization we extend the basis obtained via Gram-Schmidt to an ON-basis for

⁷In the G-S-algorithm, some resulting vectors may be zero, we drop these when constructing the basis.

\mathbb{C}^3 , and adjoin this last column to Q_1 . Then $R = Q^*A = \begin{pmatrix} R_1 \\ 0 \end{pmatrix}$. We then get a full QR-factorization:

$$A = QR = (Q_1 \mid Q_2) \begin{pmatrix} R_1 \\ 0 \end{pmatrix} = \frac{1}{3} \begin{pmatrix} 1 & 2 & 2 \\ 2 & 1 & -2 \\ 2 & -2 & 1 \end{pmatrix} \begin{pmatrix} 3 & 9 \\ 0 & 3 \\ 0 & 0 \end{pmatrix}.$$

△

Applications of the QR-decomposition

In the last section we saw how we could solve minimization-problems via Gram-Schmidt: The vector in v closest to a subspace U is $P_U(v)$, and this can be computed by constructing an ON-basis for U via Gram-Schmidt, and then projecting v onto each of these basis vectors. The QR-factorization gives a matrix version of this.

We illustrate by an example:

Example 5.5.5. Let us use the QR-factorization to find the vector in $\mathbb{U} = \text{span}((1, 2, 2), (5, 7, 4))$ closest to $(1, 2, 3)$ (with respect to the standard norm). We put the vectors generating U as columns in a matrix A - this matrix is the same as in the last example and we already have its QR-factorization.

For the minimization we are trying to find x_1, x_2 such that $x_1(1, 2, 2) + x_2(5, 7, 4) - (1, 2, 3)$ has minimal length, this length can be expressed in matrix form as:

$$\begin{aligned} \left\| \begin{pmatrix} 1 & 5 \\ 2 & 7 \\ 2 & 4 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} - \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} \right\| &= \|AX - Y\| = \|QRX - Y\| = \|Q^*QRX - Q^*Y\| = \|RX - Q^*Y\| \\ &= \left\| \begin{pmatrix} R_1 \\ 0 \end{pmatrix} X - \frac{1}{3} \begin{pmatrix} 11 \\ -2 \\ 1 \end{pmatrix} \right\| = \left\| \begin{pmatrix} R_1 X - \frac{1}{3} \begin{pmatrix} 11 \\ -2 \end{pmatrix} \\ -\frac{1}{3} \end{pmatrix} \right\|, \end{aligned}$$

where in one step we used that $\|u - v\| = \|Q(u - v)\|$ which works since $QQ^* = I$. We conclude from the calculation that the distance is minimized for $R_1 X = \frac{1}{3} \begin{pmatrix} 11 \\ -2 \end{pmatrix}$ which lets us find $X = \frac{1}{3} R_1^{-1} \begin{pmatrix} 11 \\ -2 \end{pmatrix} = \frac{1}{9} \begin{pmatrix} 17 \\ -2 \end{pmatrix}$, which gives $P_U((1, 2, 3)) = \frac{1}{9}(7, 20, 26)$.

△

For a computer there are efficient methods for finding QR-factorizations, so for large systems the above technique is useful for solving minimization problems.

Another useful application for QR-factorization in computer algebra is the *QR-algorithm*, an algorithm that is used to find eigenvalues and eigenvectors of a matrix. For large matrices A , our usual method of computing the characteristic polynomials and finding their roots is not tenable, instead the following works: Find a QR-factorization of our matrix, flip Q and R , and repeat. More precisely, start with $A_0 = A$, and for each k , QR-factorize $A_k = Q_k R_k$ and take $A_{k+1} = R_k Q_k$. Then the sequence R_k converges to an upper triangular matrix with the eigenvalues of A on the diagonal, and Q_k converges to a matrix with the corresponding eigenvectors as columns. The proof of this lies beyond the scope of this course.

Another application of QR-factorizations is that it simplifies some matrix products, later we shall need to compute products of form A^*A for non-square matrices A . If $A = QR$ we get $A^*A = (QR)^*(QR) = R^*Q^*QR = R^*IR = R^*R$ which is faster to compute since R is upper triangular. The calculation also shows that R^*R is an LU-factorization of A^*A .

5.6 Self-adjoint, unitary, and normal operators

Note that if $\mathcal{B} = (e_1, \dots, e_n)$ is an ON-basis for an inner product space V , then we have

$$\begin{aligned} \left\langle \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}_{\mathcal{B}}, \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix}_{\mathcal{B}} \right\rangle &= \langle x_1 e_1 + \dots + x_n e_n, y_1 e_1 + \dots + y_n e_n \rangle = \sum_{i,j} \langle x_i e_i, y_j e_j \rangle \\ &= \sum_{i,j} x_i \overline{y_j} \langle e_i, e_j \rangle = \sum_{i,j} x_i \overline{y_j} \delta_{ij} = \sum_i x_i \overline{y_i} = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \bullet \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix}. \end{aligned}$$

In other words, if we choose an ON-basis, the inner product of two vectors corresponds to the standard dot-product of their coordinate vectors in \mathbb{C}^n .

We also remark that if we write vectors $X, Y \in \mathbb{C}^n$ as column-matrices ($X, Y \in \text{Mat}_{n \times 1}(\mathbb{C})$), the dot-product can be expressed in matrix form as a matrix product

$$X \bullet Y = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \bullet \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} = \sum_{i=1}^n x_i \overline{y_i} = (\overline{y_1} \ \dots \ \overline{y_n}) \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} = Y^* X,$$

note the flipped order: $X \bullet Y = Y^* X$.

Linear functionals and adjoints

Since an inner product is linear in the first argument, if we fix a vector $w \in V$, we get a linear function $\alpha : V \rightarrow \mathbb{C}$ (a linear functional) by defining $\alpha(v) = \langle v, w \rangle$, also written $\alpha = \langle \cdot, w \rangle$. The next theorem says that every linear functional has this form when V is finite-dimensional.

Proposition 5.6.1. Riesz representation theorem. *Let V be finite dimensional inner product space and let $\alpha : V \rightarrow \mathbb{C}$ be linear. Then there exists a unique vector $v_\alpha \in V$ such that*

$$\alpha(u) = \langle u, v_\alpha \rangle \text{ for all } u \in V.$$

Proof. Let (e_1, \dots, e_n) be an orthonormal basis for V , and define $v_\alpha := \overline{\alpha(e_1)} e_1 + \dots + \overline{\alpha(e_n)} e_n$. Then for any $u \in V$ we have

$$\langle u, v_\alpha \rangle = \langle u, \overline{\alpha(e_1)} e_1 + \dots + \overline{\alpha(e_n)} e_n \rangle = \alpha(e_1) \langle u, e_1 \rangle + \dots + \alpha(e_n) \langle u, e_n \rangle = \alpha(\langle u, e_1 \rangle e_1 + \dots + \langle u, e_n \rangle e_n) = \alpha(u),$$

which shows the existence of such a v_α . For uniqueness, we note that if v'_α also satisfies $\alpha(u) = \langle u, v'_\alpha \rangle$, then for all u we have

$$0 = \alpha(u) - \alpha(u) = \langle u, v_\alpha \rangle - \langle u, v'_\alpha \rangle = \langle u, v_\alpha - v'_\alpha \rangle.$$

But this shows that $v_\alpha - v'_\alpha = 0$ by positive-definiteness of the inner product. \square

A similar theorem also applies in a more general context⁸.

Proposition 5.6.2. *Let $F : V \rightarrow W$ be a linear map between finite-dimensional inner product spaces (real or complex). Then there exists a unique linear map $F^* : W \rightarrow V$ called the **adjoint** of F that satisfies*

$$\langle F(v), w \rangle = \langle v, F^*(w) \rangle \text{ for all } v \in V \text{ and } w \in W.$$

If we pick orthonormal bases for V and W , we have $[F^] = [F]^*$, in other words, the matrix of the adjoint map is the Hermitian conjugate of the matrix for the original map.*

Proof. For each fixed $w \in W$, let $\alpha_w : V \rightarrow \mathbb{C}$ be defined by $\alpha_w(v) = \langle F(v), w \rangle$. This map is linear, so by Proposition 5.6.1 there exists a unique element, v_{α_w} such that

$$\langle F(v), w \rangle = \alpha_w(v) = \langle v, v_{\alpha_w} \rangle,$$

so define a map $F^* : W \rightarrow V$ by $F^*(w) = v_{\alpha_w}$, then it is not too hard to verify that this map is linear. \square

This just means that we can talk about adjoints of linear maps without specifying bases.

⁸In an infinite-dimensional Hilbert space, the theorem says that every bounded linear functional α can be represented as $\alpha = \langle \cdot, v \rangle$ for some fixed v .

Definition 5.6.3. Let $F : V \rightarrow V$ be an operator on an inner product space V .
 F is called **self-adjoint** if $F = F^*$.

F is called **unitary** if $F \circ F^* = \text{id}_V = F^* \circ F$.
 A square matrix $A \in \text{Mat}_n(\mathbb{C})$ is called **unitary** if $AA^* = I = A^*A$.

F is called **normal** if $F \circ F^* = F^* \circ F$.
 A square matrix in $A \in \text{Mat}_n(\mathbb{C})$ is called **normal** if $AA^* = A^*A$.

Some remarks are in order.

First, we see that F is self-adjoint if and only if the matrix of F with respect to an ON-basis is Hermitian.

We see that a matrix A is unitary if and only if it is square and $A^{-1} = A^*$, this also shows that a square matrix is unitary if and only if $A^*A = I$. Moreover, columns of a unitary matrix forms an ON-basis with respect to the standard inner product. To see this we denote the columns of A by A_1, \dots, A_n , and note that the block matrix product shows that

$$A^*A = \begin{pmatrix} A_1^* \\ \vdots \\ A_n^* \end{pmatrix} (A_1 \ \cdots \ A_n) = (A_i^*A_j)_{ij} = (A_j \bullet A_i)_{ij}$$

so $A^*A = I$ if and only if the columns of A (and equivalently the rows of A) form an orthonormal basis in \mathbb{C}^n .

If A is a real matrix, it is unitary if and only if it is "orthogonal"⁹: $A^T A = I = AA^T$.

If F is unitary we also note that

$$\|F(v)\|^2 = \langle F(v), F(v) \rangle = \langle v, F^*(F(v)) \rangle = \langle v, \text{id}(v) \rangle = \langle v, v \rangle = \|v\|^2,$$

which means that $\|F(v)\| = \|v\|$ for all v , and consequently $\|F(u) - F(v)\| = \|u - v\|$ so applying F doesn't change distances in the vector space, and F is called an **isometry**. Intuitively we should think of unitary operators as a sort of rotation or reflection on \mathbb{C}^n .

We also remark that both self-adjoint and unitary operators are normal.

Theorem 5.6.4. (Schur's Theorem) Let $F : V \rightarrow V$ be an operator on a finite dimensional complex inner product space. Then there exists an orthonormal basis for V with respect to which the matrix of F is upper triangular. Equivalently, any square complex matrix A is **unitarily equivalent** to an upper triangular matrix T : there exists a unitary matrix U such that $A = UTU^*$.

Proof. This follows directly from a modification of the proof of Theorem 3.2.3. We just make sure to normalize the eigenvectors we pick, and in the induction step, instead of picking an arbitrary complement to the line spanned by an eigenvector v , we pick its orthogonal complement $\text{span}(v)^\perp$. \square

5.7 Spectral theorem for normal operators

We recall the real version of the spectral theorem from a first linear algebra course, it says that *symmetric* matrices are *orthogonally diagonalizable*:

Theorem 5.7.1. (Real spectral theorem) Let $A \in \text{Mat}_n(\mathbb{R})$. There exists an orthonormal basis for \mathbb{R}^n consisting of eigenvectors for A if and only if A is *symmetric*. This means that there exists a factorization $A = SDS^T$ where D is diagonal and S is orthogonal ($SS^T = I$).

We may now prove the more general complex spectral theorem which says that the *normal* operators are precisely the ones that are orthogonally diagonalizable:

⁹"Orthonormal matrix", or "ON-matrix" in some texts.

Theorem 5.7.2. (Complex spectral theorem) Let $F : V \rightarrow V$ be an operator on a finite dimensional complex inner product space. Then there exists an orthonormal basis for V consisting of eigenvectors for F if and only if F is normal.

Equivalently, if A is a square complex matrix there exists unitary U and diagonal D such that $A = UDU^*$ if and only if A is normal.

Proof. We prove the statement in matrix form. Suppose first that $A = UDU^*$ as in the theorem. Then

$$AA^* = (UDU^*)(UD^*U^*) = UDD^*U^* = UD^*DU^* = (UD^*U^*)(UDU^*) = A^*A,$$

so A is normal.

In the other direction, suppose A is normal. We use Schur's Theorem 5.6.4 to write $A = UTU^*$, and we will show that the upper triangular matrix T is in fact diagonal. Since A is normal, so is T :

$$TT^* = (U^*AU)(UA^*U^*) = UAA^*U^* = UA^*AU^* = (U^*A^*U)(UAU^*) = T^*T.$$

We claim that normal upper triangular matrices are always diagonal, we prove this by induction. The statement is trivially true for 1×1 -matrices. Let $T = (t_{ij})_{ij}$. Then explicitly the normality $TT^* = T^*T$ looks like

$$\begin{pmatrix} t_{11} & t_{12} & \cdots & t_{1n} \\ 0 & t_{22} & \cdots & t_{2n} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & t_{nn} \end{pmatrix} \begin{pmatrix} \overline{t_{11}} & 0 & \cdots & 0 \\ \overline{t_{12}} & \overline{t_{22}} & \cdots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ \overline{t_{1n}} & \overline{t_{2n}} & \cdots & \overline{t_{nn}} \end{pmatrix} = \begin{pmatrix} \overline{t_{11}} & 0 & \cdots & 0 \\ \overline{t_{12}} & \overline{t_{22}} & \cdots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ \overline{t_{1n}} & \overline{t_{2n}} & \cdots & \overline{t_{nn}} \end{pmatrix} \begin{pmatrix} t_{11} & t_{12} & \cdots & t_{1n} \\ 0 & t_{22} & \cdots & t_{2n} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & t_{nn} \end{pmatrix}$$

We compare position $(1, 1)$ in both sides of the product and get

$$|t_{11}|^2 + |t_{12}|^2 + \cdots + |t_{1n}|^2 = |t_{11}|^2$$

which shows that $t_{12} = t_{13} = \cdots = t_{1n} = 0$. So T has block form $\begin{pmatrix} t_{11} & | & 0 \\ 0 & | & T' \end{pmatrix}$, where T' is normal since T is, and since T' is smaller than T , the inductive hypothesis gives that T' is diagonal which shows that T is diagonal. \square

To actually find the matrix U we follow the standard diagonalization algorithm, eigenvectors corresponding to different eigenvalues will then be orthogonal, but when the geometric multiplicity is greater than one we need to make sure to pick an orthogonal basis for each eigenspace, and also normalize all eigenvectors before putting them as columns in the matrix U .

Corollary 5.7.3. Self-adjoint operators, Hermitian, skew-Hermitian, and unitary matrices, as well as all complex multiples of such matrices and operators are orthogonally diagonalizable.

Proof. It is easy to verify that all the operators listed are normal, they commute with their adjoints. \square

We remark that if A is a normal matrix whose eigenvalues are all real, we get $D = D^*$ above, which implies

$$A^* = (UDU^*)^* = UD^*U^* = UDU^* = A,$$

and A is not only normal but Hermitian.

If $A = UDU^*$, and if we write u_i for the columns of U such that $U = (u_1 \cdots u_n)$, then we have:

$$A = UDU^* = (u_1 \cdots u_n) \begin{pmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{pmatrix} \begin{pmatrix} u_1^* \\ \vdots \\ u_n^* \end{pmatrix} = (u_1 \cdots u_n) \begin{pmatrix} \lambda_1 u_1^* \\ \vdots \\ \lambda_n u_n^* \end{pmatrix} = \lambda_1 u_1 u_1^* + \cdots + \lambda_n u_n u_n^*.$$

So here we have written A as a linear combination of $n \times n$ -matrices of form $u_i u_i^*$. Such matrices clearly have rank 1 and we can see that they are in fact all orthonormal with respect to the Frobenius-inner product:

$$\langle u_i u_i^*, u_j u_j^* \rangle_F = \text{tr}(u_i u_i^* (u_j u_j^*)^*) = \text{tr}(u_i u_i^* u_j u_j^*) = \text{tr}(\underbrace{(u_i u_i^* u_j u_j^*)}_{n \times n})$$

$$= \underbrace{\text{tr}(u_j^*(u_i u_i^* u_j))}_{1 \times 1} = \text{tr}((u_j^* u_i)(u_i^* u_j)) = \text{tr}((u_i \bullet u_j)(u_j \bullet u_i)) = \text{tr}(\delta_{ij} \delta_{ji}) = \delta_{ij}.$$

We will return to such decomposition of matrices when we talk about singular values and the Schmidt-decomposition, but then in the more general context of rectangular matrices.

Example 5.7.4. Let $A = \begin{pmatrix} 3 & 2 \\ -2 & 3 \end{pmatrix}$. One easily verifies that $A^*A = AA^*$, so A is normal. The standard-method shows that $\begin{pmatrix} 1 \\ i \end{pmatrix}$ is an eigenvector with eigenvalue $3 + 2i$ and that $\begin{pmatrix} 1 \\ -i \end{pmatrix}$ is an eigenvector with eigenvalue $3 - 2i$, these two eigenvectors are indeed orthogonal with respect to the standard inner product on \mathbb{C}^2 : $\begin{pmatrix} 1 \\ i \end{pmatrix} \bullet \begin{pmatrix} 1 \\ -i \end{pmatrix} = 1 \cdot \bar{1} + i \cdot \overline{-i} = 0$. We normalize them and put them as columns in a matrix U , which will then be unitary, and we take D as the diagonal matrix with the eigenvalues in the same order: with

$$U = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ i & -i \end{pmatrix} \quad \text{and} \quad D = \begin{pmatrix} 3+2i & 0 \\ 0 & 3-2i \end{pmatrix}$$

we have $A = UDU^*$.

If we want to express A as a linear combination of Frobenius-orthonormal matrices as discussed above, this looks like

$$\begin{aligned} A = UDU^* &= \begin{pmatrix} u_1 & u_2 \end{pmatrix} \begin{pmatrix} 3+2i & 0 \\ 0 & 3-2i \end{pmatrix} \begin{pmatrix} u_1^* \\ u_2^* \end{pmatrix} = (3+2i)u_1 u_1^* + (3-2i)u_2 u_2^* \\ &= (3+2i) \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix} \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \end{pmatrix} + (3-2i) \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -i \end{pmatrix} \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i \end{pmatrix} = (3+2i) \frac{1}{2} \begin{pmatrix} 1 & -i \\ i & 1 \end{pmatrix} + (3-2i) \frac{1}{2} \begin{pmatrix} 1 & i \\ -i & 1 \end{pmatrix}. \end{aligned}$$

△

5.8 Positive definite operators

Definition 5.8.1. Let $F : V \rightarrow V$ be a self-adjoint^a operator on a complex inner product space V . F is called...

- **Positive definite** if $\langle F(v), v \rangle > 0$ for all nonzero $v \in V$.
- **Positive semi-definite** if $\langle F(v), v \rangle \geq 0$ for all nonzero $v \in V$.
- **Negative definite** if $\langle F(v), v \rangle < 0$ for all nonzero $v \in V$.
- **Negative semi-definite** if $\langle F(v), v \rangle \leq 0$ for all nonzero $v \in V$.

These concept are defined analogously for matrices, for example: A Hermitian matrix A is called **positive definite** if $AX \bullet X > 0$, or equivalently

$$X^*AX > 0 \text{ for all nonzero columns } X.$$

^aThis condition is not necessary for complex matrices, every positive (semi-)definite matrix will be Hermitian, the proof of this is left as an exercise.

So we note that F is positive definite if and only if its matrix $[F]$ with respect to an ON-basis is positive definite.

Example 5.8.2. Let $A = \begin{pmatrix} 3 & 0 \\ 0 & 5 \end{pmatrix}$ and $B = \begin{pmatrix} 6 & 0 \\ 0 & 0 \end{pmatrix}$ and $C = \begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix}$, and let $X = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$ be an arbitrary nonzero vector.

$$\text{Then } X^*AX = \begin{pmatrix} \bar{x}_1 & \bar{x}_2 \end{pmatrix} \begin{pmatrix} 3 & 0 \\ 0 & 5 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = 3|x_1|^2 + 5|x_2|^2 > 0 \text{ so } A \text{ is positive definite.}$$

For B we have $X^*BX = (\overline{x_1} \ \overline{x_2}) \begin{pmatrix} 6 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = 6|x_1|^2 \geq 0$, but it can be zero for $X \neq 0$, take for example $X = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$. So B is positive semi-definite.

For C we have $X^*CX = (\overline{x_1} \ \overline{x_2}) \begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = 2|x_1|^2 + \overline{x_1}x_2 + |x_2|^2$. This may be non-real number (take for example $X = \begin{pmatrix} 1 \\ i \end{pmatrix}$), so it is neither ≥ 0 nor ≤ 0 for all X , and C is neither positive or negative (semi)-definite.

△

Suppose that A is positive definite and that A is unitarily equivalent to a matrix B via $B = U^*AU$. Then for $X \neq 0$ we have

$$X^*BX = X^*U^*AUX = (UX)^*A(UX) > 0,$$

and since X is arbitrary nonzero and U is bijective, $Y = UX$ is also arbitrary nonzero. This shows that B is positive definite too.

In particular, since A is necessarily Hermitian, by the spectral theorem it is unitarily equivalent to a diagonal matrix (with the eigenvalues of A on the diagonal), and for a diagonal matrix $D = \text{diag}(\lambda_1, \dots, \lambda_n)$ we have

$$X^*DX = \lambda_1|x_1|^2 + \dots + \lambda_n|x_n|^2$$

which is clearly positive for all nonzero X if and only if all $\lambda_i > 0$. We conclude that a Hermitian matrix A is positive definite if and only if all its eigenvalues are positive (and positive semi-definite if all its eigenvalues are ≥ 0).

We note that if we write $X, Y \in \mathbb{C}^n$ as columns-matrices, $\langle X, Y \rangle := X^*AY$ defines an inner product on \mathbb{C}^n if and only if A is positive definite. The sesqui-linearity and conjugate-symmetry follows directly, and positive-definiteness of the inner product $\langle X, X \rangle > 0$ is equivalent to A being positive definite.

The following result provides a useful condition for testing whether a Hermitian matrix is positive definite:

Theorem 5.8.3. (Sylvester’s criterion) *Let $A \in \text{Mat}_n(\mathbb{C})$ be a Hermitian matrix ($A = A^*$). The **principal minor** of size $m \times m$ in A is the determinant of the matrix obtained from removing all but the first m rows and columns from A .*

The Hermitian matrix A is positive definite if and only if all its principal minors are positive, and A is positive semi-definite if all its principal minors are ≥ 0 .

We don’t prove this here, but we illustrate with an example:

Example 5.8.4. The matrix $A = \begin{pmatrix} 1 & i & 1 \\ -i & 2 & 1 \\ 1 & 1 & 5 \end{pmatrix}$ is Hermitian. Its three principal minors of size 1, 2, 3 are:

$$|1| = 1, \quad \begin{vmatrix} 1 & i \\ -i & 2 \end{vmatrix} = 1, \quad \begin{vmatrix} 1 & i & 1 \\ -i & 2 & 1 \\ 1 & 1 & 5 \end{vmatrix} = 2.$$

Since *all* of these are positive, by Sylvester’s criterion A is positive definite.

△

Square roots

In calculus, for $x \in \mathbb{R}$ we normally define \sqrt{x} as the unique *non-negative* real number y satisfying $y^2 = x$. So for $x < 0$, the square root \sqrt{x} is normally *undefined*. For $x > 0$ there are two real numbers y satisfying $y^2 = x$, but only one of them is positive.

For matrices the situation is analogous if we replace “positive” by “positive definite”:

Definition 5.8.5. Let A be a positive semi-definite matrix. Then A is necessarily Hermitian, and by the spectral theorem we may write $A = UDU^*$ for $D = \text{diag}(\lambda_1, \dots, \lambda_n)$ and since A is positive semi-definite, so is the matrix D , so $\lambda_i \geq 0$. We define

$$\sqrt{D} = \text{diag}(\sqrt{\lambda_1}, \dots, \sqrt{\lambda_n}) \quad \text{and} \quad \sqrt{A} = U\sqrt{D}U^*.$$

So we leave \sqrt{A} *undefined* when A is not positive semi-definite, even if there may exist matrices which squares to A .

Proposition 5.8.6. \sqrt{A} is well defined for positive semi-definite A . \sqrt{A} is the unique positive semi-definite matrix whose square is A .

Proof. We have $\sqrt{A}^2 = U\sqrt{D}U^*U\sqrt{D}U^* = U\sqrt{D}\sqrt{D}U^* = UDU^* = A$, and \sqrt{A} is positive semi-definite since all its eigenvalues are ≥ 0 .

For the uniqueness claim, suppose B is any positive semi-definite matrix for which $B^2 = A$. Since B is positive semi-definite it is Hermitian, and we can use the spectral theorem to write $B = \tilde{U}\tilde{D}\tilde{U}^*$ where $\tilde{D} = \text{diag}(\mu_1, \dots, \mu_n)$ has the eigenvalues of B on the diagonal. Then $A = B^2 = \tilde{U}\tilde{D}^2\tilde{U}^*$ so $A\tilde{U} = \tilde{U}\tilde{D}^2$ which shows that the columns of \tilde{U} are orthonormal eigenvectors for A with eigenvalues μ_i^2 , this shows that $\mu_i^2 = \lambda_i$. Since A is positive semi-definite and Hermitian, all its eigenvalues λ_i are real and ≥ 0 , which shows that $\mu_i = \pm\sqrt{\lambda_i}$, but since B is required to be positive semi-definite, only the plus-signs can occur. But then \sqrt{A} and B acts the same way on a basis for the vector space (the columns of \tilde{U}), so $B = \sqrt{A}$. \square

Example 5.8.7. Let $A = \begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix}$. By Sylvester's criterion, A is positive definite and it has a square root \sqrt{A} . Let us compute it:

We find the eigenvalues of A and an orthonormal basis of eigenvectors, this gives us

$$A = UDU^* \quad \text{where} \quad U = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad \text{and} \quad D = \begin{pmatrix} 4 & 0 \\ 0 & 2 \end{pmatrix}.$$

So we have

$$\sqrt{A} = U\sqrt{D}U^* = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} 2 & 0 \\ 0 & \sqrt{2} \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 2 + \sqrt{2} & 2 - \sqrt{2} \\ 2 - \sqrt{2} & 2 + \sqrt{2} \end{pmatrix}.$$

Note that each of the four matrices $B = U \begin{pmatrix} \pm 2 & 0 \\ 0 & \pm\sqrt{2} \end{pmatrix} U^*$ satisfies $B^2 = A$, but only one of them is positive definite.

\triangle

5.9 Least squares

Definition 5.9.1. Consider a system $AX = b$ where $A \in \text{Mat}_{m \times n}(\mathbb{C})$, and where $X \in \mathbb{C}^n$ and $b \in \mathbb{C}^m$ are written as column-matrices. Such a system may not have a solution (this typically happens when $m > n$). A **least square solution** to $AX = b$ is a vector X for which $\|AX - b\|$ is minimal, where $\|\cdot\|$ is the standard norm on \mathbb{C}^m .

Write $A = (A_1 \ \dots \ A_n)$ and $X = (x_1, \dots, x_n)^T$. We then note that $AX = x_1A_1 + \dots + x_nA_n \in \text{Im}(A)$, so $\|AX - b\|$ is minimal when $X = P_{\text{Im}(A)}(b)$ in which case $b - AX$ is orthogonal to $\text{Im}(A) = \text{span}(A_1, \dots, A_n)$. This orthogonality can be expressed in terms of the *normal equations*

$$(b - AX) \bullet A_i = 0 \ \forall i \Leftrightarrow A_i^*(AX - b) = 0 \ \forall i \Leftrightarrow \begin{pmatrix} A_1^* \\ \vdots \\ A_n^* \end{pmatrix} (AX - b) = \begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix} \Leftrightarrow A^*(AX - b) = 0 \Leftrightarrow A^*AX = A^*b.$$

The $n \times n$ -matrix A^*A is typically invertible, and then the least square solution is unique:

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

This gives an explicit way to find the vector in $\text{Im}(A)$ closest to b without applying Gram-Schmidt to the columns of A .

Proposition 5.9.2. *The least square solution to the $m \times n$ system $AX = b$ is unique if and only if the columns of A are linearly independent.*

Proof. We claim that $\ker(A^*A) = \ker(A)$, the inclusion \supset is obvious, and on the other hand:

$$(A^*A)X = 0 \Rightarrow X^*A^*AX = 0 \Rightarrow X^*(A^*A)X = 0 \Rightarrow (AX)^*(AX) = 0 \Rightarrow \|AX\|^2 = 0 \Rightarrow AX = 0$$

so $\ker(A^*A) \subset \ker(A)$ too, and the kernels are equal. Both A and A^*A have n columns, so by the rank-nullity theorem

$$\text{rank}(A^*A) = n - \dim(\ker(A^*A)) = n - \dim(\ker(A)) = \text{rank}(A).$$

Thus the $n \times n$ -matrix A^*A is invertible if and only if $\text{rank}(A^*A) = n$ if and only if $\text{rank}(A) = n$ which is the same as the columns of A being linearly independent. \square

Definition 5.9.3. Let A be an $m \times n$ matrix with linearly independent columns. The **Moore-Penrose pseudo-inverse** of A is defined as the $n \times m$ -matrix

$$A^+ := (A^*A)^{-1}A^*.$$

By the discussion above, the least square solution to the system $AX = b$ is then $X = A^+b$, so A^+ is "almost" an inverse to the non-square matrix A , but note that typically $AA^+ \neq I$.

Proposition 5.9.4. *Let $A \in \text{Mat}_{m \times n}(\mathbb{C})$ with linearly independent columns, and let $A^+ := (A^*A)^{-1}A^*$. Then*

$$A^+A = I_n \text{ and } AA^+ = P_{\text{Im}(A)}.$$

Proof. $A^+A = ((A^*A)^{-1}A^*)A = (A^*A)^{-1}(A^*A) = I_n$. For the second equality, we note that by the arguments above, for each $b \in \mathbb{C}^m$, we have $P_{\text{Im}(A)}(b) = AX = AA^+b$, so $P_{\text{Im}(A)} = AA^+$. \square

Later we will use singular values to also define the Moore-Penrose pseudo-inverse for arbitrary matrices A (when the columns possibly are linearly dependent).

We give an example where we use these this pseudo-inverse to solve a least-square problem:

Example 5.9.5. Let us use the Moore-Penrose pseudo inverse to find the least square solution to the linear system

$$\begin{cases} x + y &= 1 \\ x &= 1 \\ ix + 2y &= 0. \end{cases}$$

The system can be written as $AX = b$ where $A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \\ i & 2 \end{pmatrix}$ and $X = \begin{pmatrix} x \\ y \end{pmatrix}$ and $b = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}$. The pseudo inverse is

$$A^+ = (A^*A)^{-1}A^* = \begin{pmatrix} 3 & 1 - 2i \\ 1 + 2i & 5 \end{pmatrix}^{-1} A^* = \frac{1}{10} \begin{pmatrix} 5 & -1 + 2i \\ -1 - 2i & 3 \end{pmatrix} A^* = \frac{1}{10} \begin{pmatrix} 4 + 2i & 5 & -2 - i \\ 2 - 2i & -1 - 2i & 4 + i \end{pmatrix}$$

so the least square solution to our system is given by

$$\begin{pmatrix} x \\ y \end{pmatrix} = A^+b = \frac{1}{10} \begin{pmatrix} 9 + 2i \\ 1 - 4i \end{pmatrix}.$$

△

In this example it was not really important to find the matrix A^+ , it would have been faster to multiply the matrices in the following order to find X :

$$X = A^+b = ((A^*A)^{-1}A^*)b = (A^*A)^{-1}(A^*b).$$

However, knowing the pseudo-inverse A^+ is useful if we need to find least square solutions $AX = b$ for several different right sides b .

Here is another example illustrating how one can fit a curve to some given data:

Example 5.9.6. Let us find the centered ellipse of form $c_1x^2 + c_2y^2 = 1$ which in the least square sense best approximates the data:

$$\begin{array}{c|cccc} x & 0 & 2 & 1 & -1 \\ \hline y & -1 & 0 & 1 & 1 \end{array}$$

We insert the four datapoints (x, y) in our equation $c_1x^2 + c_2y^2 = 1$ and obtain four linear equations in the variables c_1 and c_2 :

$$\begin{pmatrix} 0 & 1 \\ 4 & 0 \\ 1 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \Leftrightarrow AX = b.$$

This system is inconsistent, but the two columns of A are independent, so we can find its least square solution^a $X = A^+b$.

Here the Moore-Penrose pseudo-inverse is given by


$$A^+ = (A^*A)^{-1}A^* = \begin{pmatrix} 18 & 2 \\ 2 & 3 \end{pmatrix}^{-1} \begin{pmatrix} 0 & 4 & 1 & 1 \\ 1 & 0 & 1 & 1 \end{pmatrix} = \frac{1}{52} \begin{pmatrix} 3 & -2 \\ -2 & 18 \end{pmatrix} \begin{pmatrix} 0 & 4 & 1 & 1 \\ 1 & 0 & 1 & 1 \end{pmatrix} = \frac{1}{50} \begin{pmatrix} -2 & 12 & 1 & 1 \\ 18 & -8 & 16 & 16 \end{pmatrix},$$

so the least square solution is

$$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = A^+b = \frac{1}{50} \begin{pmatrix} -2 & 12 & 1 & 1 \\ 18 & -8 & 16 & 16 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} = \frac{1}{25} \begin{pmatrix} 6 \\ 21 \end{pmatrix},$$

so the ellipse we seek is

$$\frac{6x^2}{25} + \frac{21y^2}{25} = 1.$$



"This is what is meant by the phrasing "in the least square sense" in the beginning, but technically this is not *exactly* the same as the ellipse which minimizes the distances from the points to the ellipse.

We also remark that if we know a QR -factorization of A , helps with the computation of A^+b , we have: △

$$A^+ = (A^*A)^{-1}A^* = (R^*Q^*QR)^{-1}R^*Q^* = (R^*R)^{-1}R^*Q^*.$$

Chapter 6

Perron-Frobenius theory


6.1 Positive and non-negative matrices

Introduction

Perron-Frobenius theory is the study of the spectral properties of matrices where the entries are *positive*, or at least *non-negative*. Such matrices appear in a multitude of applications, such as in graph theory, stochastic processes, ranking models, Markov chains, models for economic growth etcetera. As in introductory example, let us consider a 4×4 -matrix with positive entries. We find its eigenvalues¹ and label them in order of decreasing absolute value:

$$A = \begin{pmatrix} 3 & 2 & 4 & 1 \\ 3 & 2 & 4 & 3 \\ 1 & 3 & 2 & 2 \\ 4 & 2 & 2 & 3 \end{pmatrix} \quad \sigma(A) = \{\lambda_1, \lambda_2, \lambda_3, \lambda_4\} \text{ where } \begin{cases} \lambda_1 \approx 10.08 \\ \lambda_2 \approx 0.60 + 1.74i \\ \lambda_3 \approx 0.60 - 1.74i \\ \lambda_4 \approx -1.29 \end{cases}$$

Here is what the spectrum $\sigma(A)$ looks like in the complex plane:



We see that there is a unique eigenvalue with maximal absolute value, and that that eigenvalue is *real and positive*. Here are the corresponding eigenvectors satisfying $Av_i = \lambda_i v_i$:

$$v_1 \approx \begin{pmatrix} 0.47 \\ 0.57 \\ 0.40 \\ 0.54 \end{pmatrix} \quad v_2 \approx \begin{pmatrix} -0.27 + 0.54i \\ -0.031 - 0.14i \\ -0.22 - 0.37i \\ 0.66 \end{pmatrix} \quad v_3 \approx \begin{pmatrix} -0.27 - 0.54i \\ -0.031 + 0.14i \\ -0.22 + 0.37i \\ 0.66 \end{pmatrix} \quad v_4 \approx \begin{pmatrix} -0.25 \\ -0.73 \\ 0.55 \\ 0.31 \end{pmatrix}$$

We see that the eigenvector corresponding to λ_1 was chosen to have real and positive components.

In this section we shall show that this is always the case for square matrices with positive entries. The largest eigenvalue of A determines the growth of $\|A^n v\|$ as $n \rightarrow \infty$ which is important in applications.

Recall that for $z, w \in \mathbb{C}$ we write $z > w$ if and only if $z - w$ is *real and positive*.

¹Found numerically and rounded to 2 decimal places.

Definition 6.1.1. If A and B are complex matrices of the same size we shall write

$$A > B \quad \text{when } a_{ij} > b_{ij} \text{ for all indices } (i, j)$$

$$A \geq B \quad \text{when } a_{ij} \geq b_{ij} \text{ for all indices } (i, j)$$

In particular, if $A > 0$ the matrix A is called **positive**, this means that all the entries of A are positive. If $A \geq 0$ the matrix A is called **non-negative**, this means each entry of A is positive or zero.

The relation \leq is a *partial order*², but note that most matrices are incomparable: neither $A \leq B$ nor $B \leq A$ holds. The naming convention "non-negative" is a bit unfortunate here - it does not mean that $A \not\leq 0$ by the definition.

Note that positivity is not a basis-independent concept, changing basis in \mathbb{C}^n might turn a positive matrix non-positive. In this section we will always consider vectors as columns, and since these are $n \times 1$ -matrices, we can also talk about positive *vectors*.

Example 6.1.2.

$$\begin{pmatrix} 3 & 6 \\ -4 & 5+i \end{pmatrix} > \begin{pmatrix} 1 & -7 \\ -6 & 3+i \end{pmatrix} \quad \text{and} \quad v = \begin{pmatrix} 1 \\ 0 \\ 2 \end{pmatrix} \geq 0.$$

Note that the vector v above satisfies $v \geq 0$ and $v \neq 0$, but it doesn't satisfy $v > 0$.

△

Note that "positive" is a different concept than "positive definite", unfortunately positive definite matrices are called positive in some texts.

Clearly the product of positive matrices is positive, and the product of non-negative matrices are non-negative. If $A \in \text{Mat}_{m \times n}$ is positive and $X \in \text{Mat}_{n \times 1}$ is nonzero and non-negative, then AX is positive.

6.2 Perron's theorem

Let A be an $n \times n$ -matrix. In this section we shall always order its eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$ such that

$$|\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_n|, \text{ and such that } |\lambda_i| = |\lambda_j| \text{ with } \lambda_i \geq 0 \Rightarrow i \leq j,$$

in other words, the eigenvalues are ordered by absolute value with largest values first, and real positive eigenvalues comes first in a tie. For example, if $\sigma(A) = \{-4, 5, 3+4i\}$, the only correct order is $(\lambda_1, \lambda_2, \lambda_3) = (5, 3+4i, -4)$.

We start by summarizing the main result on positive matrices.

Theorem 6.2.1. Perron's theorem. *Let $A > 0$ be a positive matrix. Then the eigenvalue λ_1 of largest absolute value is real and positive. Moreover, λ_1 has algebraic multiplicity 1 and no other eigenvalue has the same absolute value. The corresponding eigenvector can be taken as positive, and no other non-negative eigenvectors for A exist.*

Definition 6.2.2. An eigenvalue λ is called **dominant** if its absolute value is strictly larger than the absolute value of all other eigenvalues of A . When the matrix A is positive, the dominant eigenvalue is called a **Perron eigenvalue**, and a corresponding positive eigenvector is called a **Perron vector**.

By the theorem above, the unique Perron-eigenvalue is real and positive, and the corresponding Perron-vector is unique up to positive multiples.

We prove Theorem 6.2.1 by splitting it into a couple of lemmata.

Lemma 6.2.3. *Let $A > 0$ be a square matrix. Then the eigenvalue λ_1 of maximal absolute value is real and positive, and the corresponding eigenvector v can also be chosen as positive with $v > 0$.*

²meaning that $A \leq A$, and that if $A \leq B$ and $B \leq C$, then $A \leq C$, and that if $A \leq B$ and $B \leq A$ then $A = B$.

Proof. Let λ_1 be an eigenvalue of maximal absolute value, and let $u = (u_1, \dots, u_n)^T$ be a corresponding eigenvector: $Au = \lambda_1 u$. We shall show that $v = u_{\text{abs}} := (|u_1|, \dots, |u_n|) > 0$ is also an eigenvector but with eigenvalue $|\lambda_1|$. This will imply that $\lambda_1 = |\lambda_1|$, and hence $\lambda_1 \geq 0$ since we decided that real eigenvalues comes first when ordering the eigenvalues.

We see that $Av > 0$ and

$$Av = A_{\text{abs}}u_{\text{abs}} \geq (Au)_{\text{abs}} = (\lambda_1 u)_{\text{abs}} = |\lambda_1|u_{\text{abs}} = |\lambda_1|v,$$

where the inequality follows from the triangle inequality in each component of the vector. So $Av \geq |\lambda_1|v$, and we will show that this inequality is in fact an equality. Assume the contrary: that $Av \neq |\lambda_1|v$. Then

$$\underbrace{A}_{>0} \underbrace{(Av - |\lambda_1|v)}_{\geq 0} > 0$$

this means that we can pick ϵ small enough that

$$A(Av - |\lambda_1|v) > \epsilon Av,$$

which implies that $A(Av) > (\epsilon + |\lambda_1|)Av$, and $\frac{1}{\epsilon + |\lambda_1|}AAv > Av$. Let $B = \frac{1}{\epsilon + |\lambda_1|}A$. Then $B > 0$ and $B(Av) > Av$ and by repeatedly multiplying by the positive matrix B we get a chain of positive vectors:

$$Av < BAv < B^2Av < B^3Av < \dots$$

so $B^kAv > Av$ for all $k \geq 1$. We claim that $B^k \rightarrow 0$ entry-wise. To see this, Jordanize B with $B = SJS^{-1}$, where $J = D + N$ with N nilpotent (with say $N^{p+1} = 0$) and D is a diagonal matrix with the eigenvalues of B on the diagonal. Since $|\lambda_1|$ is the largest absolute-value of an eigenvalue of A , the largest absolute-value of an eigenvalue of B is $\frac{|\lambda_1|}{\epsilon + |\lambda_1|} < 1$, so D has diagonal elements of absolute-value < 1 on the diagonal, so $D^k \rightarrow 0$ as $k \rightarrow \infty$. But then via the binomial theorem we have

$$B^k = SJ^kS^{-1} = S(D + N)^kS^{-1} = S(D^k + kD^{k-1}N + \dots + \binom{k}{p}N^pD^{k-p})S^{-1} \rightarrow 0 \text{ when } k \rightarrow \infty,$$

since S, S^{-1} , and p are fixed, and $D^k \rightarrow 0$ (the terms are polynomials in k multiplied by exponentials α^k for $0 < \alpha < 1$).

Taking the limit of $B^kAv > Av$ as $k \rightarrow \infty$ we obtain $0 \geq Av$, contradicting $Av > 0$. Thus our original assumption that $Av \neq |\lambda_1|v$ is false, so $Av = |\lambda_1|v$, and as discussed in the beginning, the ordering of eigenvalues shows that $\lambda_1 = |\lambda_1|$ so $\lambda_1 \geq 0$, and the $v = u_{\text{abs}}$ used in the proof is clearly non-negative.

So we know $\lambda_1 \geq 0$ and $v \geq 0$, finally we need to confirm that these inequalities are strict. But if $\lambda_1 = 0$, then all eigenvalues are zero which is not possible since the sum of eigenvalues is $\text{tr}(A)$ which is positive since $A > 0$, so $\lambda_1 > 0$. But then since $Av = \lambda_1 v$ we have $v = \frac{1}{\lambda_1} \underbrace{Av}_{>0} > 0$. \square

Lemma 6.2.4. For a positive matrix A , the dominant eigenvalue is unique: $|\lambda_k| < |\lambda_1|$ when $\lambda_k \neq \lambda_1$.

Proof. We prove the statement by contradiction. Suppose $\lambda_k \neq \lambda_1$ but $|\lambda_k| = |\lambda_1|$, and let $y = (y_1, \dots, y_n)^T$ be an eigenvector corresponding to the eigenvalue λ_k such that $Ay = \lambda_k y$. Just as in Lemma 6.2.3 it follows that y_{abs} also is an eigenvector with $y_{\text{abs}} > 0$, but with eigenvalue $|\lambda_k|$, so $Ay_{\text{abs}} = |\lambda_k|y_{\text{abs}} = \lambda_1 y_{\text{abs}}$. In other words, for all vector indices j we have

$$\lambda_1 |y_j| = |\lambda_k y_j| = |(Ay)_j| = \left| \sum_k a_{jk} y_k \right| \leq \underbrace{\sum_k a_{jk} |y_k|}_{*} = \sum_k \underbrace{|a_{jk}|}_{a_{jk}} \cdot |y_k| = (Ay_{\text{abs}})_j = \lambda_1 |y_j|.$$

since the left and right sides are equal, the triangle-inequality $*$ in the middle must in fact be an equality, and this happens precisely if all components y_k have the same argument θ when expressed in polar coordinates. Thus for each k we have $e^{-i\theta} y_k > 0$, which gives $x := e^{-i\theta} y > 0$. But then $0 < Ax = e^{-i\theta} Ay = e^{-i\theta} \lambda_k y = \lambda_k x$, so both $x > 0$ and $\lambda_k x > 0$ which implies that λ_k must be real and positive, and since $|\lambda_k| = |\lambda_1|$ this implies that $\lambda_1 = \lambda_k$. This contradiction proves that $\lambda_1 \neq \lambda_k$ with $|\lambda_k| = |\lambda_1|$ is impossible, which is the statement of the lemma. \square

Lemma 6.2.5. For a positive matrix A , the dominant eigenvalue λ_1 has algebraic multiplicity 1.

Proof. We first show that the geometric and algebraic multiplicity of λ_1 must agree, and then that the geometric multiplicity must be 1. Let J be the Jordan form of A with $A = SJS^{-1}$. Then we have³

$$\frac{1}{\lambda_1^k} \|J^k\|_F = \left\| S^{-1} \frac{1}{\lambda_1^k} A^k S \right\|_F \leq \|S\|_F \cdot \left\| \frac{1}{\lambda_1^k} A^k \right\|_F \cdot \|S^{-1}\|_F. \tag{6.1}$$

Now if the algebraic and geometric multiplicities of λ_1 differ, then J contains a block $J_m(\lambda_1)$ for $m > 1$. We claim that this implies that the left side of Equation (6.1) tends towards infinity as $k \rightarrow \infty$. This follows from Proposition 4.4.5, to illustrate, when $m = 2$, on the diagonal of $\frac{1}{\lambda_1^k} J^k$ we have a block

$$\frac{1}{\lambda_1^k} (J_2(\lambda_1))^k = \frac{1}{\lambda_1^k} \begin{pmatrix} \lambda_1 & 1 \\ 0 & \lambda_1 \end{pmatrix}^k = \frac{1}{\lambda_1^k} \begin{pmatrix} \lambda_1^k & k\lambda_1^{k-1} \\ 0 & \lambda_1^k \end{pmatrix} = \begin{pmatrix} 1 & \frac{k}{\lambda_1} \\ 0 & 1 \end{pmatrix},$$

whose Frobenius-norm tends to infinity since $\frac{k}{\lambda_1} \rightarrow \infty$ as $k \rightarrow \infty$.

Thus the left hand side, and also the right hand side of Equation (6.1) tends towards infinity, which means that $\left\| \frac{1}{\lambda_1^k} A^k \right\|_F \rightarrow \infty$ as $k \rightarrow \infty$. We claim that this is not possible, indeed let $v > 0$ be a Perron vector with $Av = \lambda_1 v$. Then $A^k v = \lambda_1^k v$ so $v = \frac{1}{\lambda_1^k} A^k v$, summing the positive coordinates of v we get

$$\sum_j v_j = \frac{1}{\lambda_1^k} \sum_j (A^k v)_j = \frac{1}{\lambda_1^k} \sum_{i,j} (A^k)_{ji} v_i \geq \frac{1}{\lambda_1^k} \underbrace{\left(\sum_{i,j} (A^k)_{ji} \right)}_{\rightarrow \infty} \min_i v_i.$$

The right hand side tends towards infinity as $k \rightarrow \infty$, but the left hand side is fixed, so this is a contradiction, and our initial assumption that there existed a Jordan block of size > 1 is false, which shows that the algebraic and geometric multiplicities for λ_1 agree.

It remains only to show that the geometric multiplicity of λ_1 is 1. Let v be a Perron vector, and suppose there exists another eigenvector w of eigenvalue λ_1 which is non-parallel to v . Pick an index i for which w_i is nonzero and form $u = v - \frac{v_i}{w_i} w$. Then u is also an eigenvector of eigenvalue λ_1 . As in Lemma 6.2.3 get that u_{abs} is also an eigenvector for eigenvalue λ_1 , and $u_{\text{abs}} \geq 0$. But then since $Au = \lambda_1 u$ we get $u = \frac{1}{\lambda_1} Au > 0$ since $A > 0$ and $u \geq 0$ is not zero. But by the definition of u , we see that $u_i = 0$ so u can not be positive, this contradiction shows that there can not exist two linearly independent eigenvectors of eigenvalue λ_1 . This finishes the proof of the lemma. \square

The last lemma shows that no eigenvalue of a positive matrix other than λ_1 can have a non-negative eigenvector.

Lemma 6.2.6. *Let $A > 0$. If $Au = \lambda_k u$ with $u \geq 0$, then $k = 1$.*

Proof. Let $Au = \lambda_k u$ with $u \geq 0$, and suppose $\lambda_k \neq \lambda_1$. Since A^T has the same eigenvalues as A , the Perron-eigenvalue for A will also be the Perron-eigenvalue for A^T , so there exists a Perron vector $u' > 0$ for A^T satisfying $A^T u' = \lambda_1 u'$. But then

$$\underbrace{(\lambda_1 - \lambda_k)}_{\neq 0} \underbrace{u^T u'}_{1 \times 1} = u^T (\lambda_1 u') - (\lambda_k u^T) u' = u^T (A^T u') - (Au)^T u' = u^T A^T u' - u^T A^T u' = 0,$$

but this is a contradiction since the left hand side is clearly nonzero since $\lambda_1 \neq \lambda_k$, and $u^T u' > 0$ since $0 \neq u^T \geq 0$ and $u' > 0$. \square

6.3 Primitive matrices and the power method

Now let us start to consider matrices which are non-negative. A weaker version of Perron's theorem still applies to any non-negative matrix A :

³Recall that the Frobenius-norm of a matrix is $\|A\|_F^2 = \sum_{i,j} |a_{ij}|^2$. This norm satisfies $\|AB\|_F \leq \|A\|_F \cdot \|B\|_F$, the proof of this is an exercise (just add terms and use the triangle inequality)

Lemma 6.3.1. *Let $A \geq 0$ be a non-negative square matrix. Then there exists a real eigenvalue $\lambda_1 \geq 0$ for A , such that all other eigenvalues have less or equal absolute value. There exists a corresponding non-negative eigenvector $v \geq 0$ with $Av = \lambda_1 v$.*

Proof. We note that A is a limit of positive matrices: if $(1) = (1)_{ij}$ is the matrix full of ones, then $A = \lim_{n \rightarrow \infty} A + \frac{1}{n}(1)$. Since each matrix $A + \frac{1}{n}(1)$ is positive, it has a unique positive dominant eigenvalue $\lambda_1^{(n)}$ and a corresponding unique positive normalized Perron-vector $v^{(n)} > 0$. Since the roots of the characteristic polynomial depends continuously on the matrix coefficients, $\lambda_1^{(n)}$ converges to λ_1 , an eigenvalue of A greater or equal than the absolute value of all other eigenvalues, and $v^{(n)}$ converges to a corresponding eigenvector v (this last statement really requires checking some additional details that we omit here). Since $\lambda_1^{(n)} > 0$, we get $\lambda_1 \geq 0$, and since $v^{(n)} > 0$ we get $v \geq 0$. \square

When $A \geq 0$ we shall still refer to the $\lambda_1 \geq 0$ from the lemmas as the Perron-eigenvalue, and a corresponding eigenvector $v \geq 0$ as a Perron-vector even if its direction might not be unique.

If we require some additional properties of the matrix, we can say something more specific.

Definition 6.3.2. A matrix $A \geq 0$ is called **primitive** if $A^p > 0$ for some positive integer p .

Example 6.3.3. The matrix A below is primitive with $A^3 > 0$, while the matrix B below is not primitive since no power of B is strictly positive.

$$A = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix} \quad A^3 = \begin{pmatrix} 2 & 2 & 1 \\ 1 & 1 & 1 \\ 2 & 2 & 1 \end{pmatrix} \quad B = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad B^p = \begin{pmatrix} 1 & p \\ 0 & 1 \end{pmatrix}$$

△

For primitive non-negative matrices, the exact same conclusions from Perron's theorem 6.2.1 still hold:

Theorem 6.3.4. *Let $A \geq 0$ is a primitive non-negative matrix. Then there is a unique eigenvalue $\lambda_1 > 0$ of largest absolute value. Moreover, λ_1 has algebraic multiplicity 1 and no other eigenvalue has the same absolute value. The corresponding eigenvector v , a Perron-vector, can be taken as positive, and besides multiples of v , no other eigenvectors for A are non-negative.*

Proof. According to Lemma 6.3.1 there exists $\lambda \geq 0$ and $v \geq 0$ such that $Av = \lambda v$ and such that all other eigenvalues of A have absolute value less or equal to λ . Since A is primitive there exists $k > 0$ such that $A^k > 0$. Then $A^k v = \lambda^k v$, so v is a non-negative eigenvector of the positive matrix A^k , so by Perron's theorem 6.2.1, λ^k is the Perron-eigenvalue of A^k and in fact v must be strictly positive, and no other eigenvalue of A^k can have the same absolute value as λ^k . But by the spectral mapping theorem 3.3.11, the eigenvalues of A^k all have form μ^k where μ is an eigenvalue of A , so this shows that no other eigenvalue with absolute value λ can exist for A . Since the algebraic multiplicity of an eigenvalue μ of A is the same as the algebraic multiplicity of μ^k in A^k , the eigenvalue of λ of A must have multiplicity 1. Finally, if another non-negative vector w was an eigenvector for A , then it would also be a non-negative vector for the positive matrix A^k which we know is not possible by Perron. \square

Finding the Perron-vector

In applications, the $n \times n$ -matrix A can be very large, so the classical approach to finding eigenvalues/eigenvectors is ineffective. Instead we can use the *power method*.

Proposition 6.3.5. The power method.

Let A be an $n \times n$ -matrix which has a unique eigenvalue λ_1 of maximal absolute value, and whose algebraic multiplicity is 1.

Pick a random nonzero vector $w_0 \in \mathbb{C}^n$, and repeatedly multiply by A and normalize:

$$w_{k+1} = \frac{Aw_k}{\|Aw_k\|}.$$

Then with probability 1, the sequence of vectors w_k converges to an eigenvector w for λ_1 .

Proof. We note that w_k can also be calculated directly as $w_k = \frac{A^k w_0}{\|A^k w_0\|}$ (where we normalized afterwards instead of in each step), however in applications $A^k w_0$ will be very large, so for numerical reasons it is better to normalize in each step.

Assume first that A is diagonalizable with eigenvalues $\lambda_1, \dots, \lambda_n$ and corresponding eigenvectors v_1, \dots, v_n . Express our initially chosen vector w_0 from the power method in the eigenbasis: $w_0 = c_1 v_1 + \dots + c_n v_n$. Then with probability 1, the coefficient c_1 will be nonzero and we have

$$\begin{aligned} w_k &= \frac{A^k w}{\|A^k w\|} = \frac{c_1 A^k v_1 + c_2 A^k v_2 + \dots + c_n A^k v_n}{\|c_1 A^k v_1 + c_2 A^k v_2 + \dots + c_n A^k v_n\|} = \frac{c_1 \lambda_1^k v_1 + c_2 \lambda_2^k v_2 + \dots + c_n \lambda_n^k v_n}{\|c_1 \lambda_1^k v_1 + c_2 \lambda_2^k v_2 + \dots + c_n \lambda_n^k v_n\|} \\ &= \frac{c_1 v_1 + c_2 (\frac{\lambda_2}{\lambda_1})^k v_2 + \dots + c_n (\frac{\lambda_n}{\lambda_1})^k v_n}{\|c_1 v_1 + c_2 (\frac{\lambda_2}{\lambda_1})^k v_2 + \dots + c_n (\frac{\lambda_n}{\lambda_1})^k v_n\|} \rightarrow \frac{c_1 v_1}{\|c_1 v_1\|}, \end{aligned}$$

where all $(\frac{\lambda_m}{\lambda_1})^k \rightarrow 0$ since λ_1 is strictly larger than the absolute value of all other eigenvalues. So w_k converges to an eigenvector for A , and since each w_k is non-negative, so is the limit. The same argument works for non-diagonalizable A , we know that the algebraic multiplicity of λ_1 is 1, and working with generalized eigenvectors instead of eigenvectors some terms of form $k(\frac{\lambda_m}{\lambda_1})^k$ may occur in the calculation, but these also approach zero. □

Corollary 6.3.6. *Let $A \geq 0$ be primitive. Then take $w_0 > 0$ as a positive vector in the power method. Then with probability 1, the vectors w_k converges to a Perron vector $v > 0$ for A , and the Perron-eigenvalue can be obtained as $\lambda_1 = \frac{\|Av\|}{\|v\|}$.*

Proof. We know that non-negative primitive matrices has a unique eigenvalue $\lambda_1 > 0$ of algebraic multiplicity 1, so the previous proposition applies. If v is a Perron vector, then $Av = \lambda_1 v$, so $\|Av\| = \|\lambda_1 v\| = |\lambda_1| \|v\|$, so $\frac{\|Av\|}{\|v\|} = |\lambda_1| = \lambda_1$. □

6.4 Frobenius' theorem for irreducible matrices

Even for matrices that are not primitive, some of the results of Perron's theorem may hold. Frobenius found a large class of such matrices called *irreducible matrices*:

Definition 6.4.1. An $n \times n$ -matrix A is called **reducible** if there exists a permutation matrix P (a matrix with a single 1 in each row and each column, and zeros elsewhere), such that $P^T A P$ is block-upper triangular:

$$P^T A P = \left(\begin{array}{c|c} B & C \\ \hline 0 & D \end{array} \right)$$

where in this block form B and D are square matrices, possibly of different size. If A is not reducible it is called **irreducible**.

An $n \times n$ -permutation matrix P acts on \mathbb{C}^n by permuting the standard basis vectors, so we can think of the matrix $P^T A P$ as the matrix obtained from A by relabelling the rows and columns according to the permutation P . So intuitively, an irreducible matrix does not have "too many zeros".

Theorem 6.4.2. Frobenius' theorem.

If $A \geq 0$ is a square non-negative matrix that is irreducible, then the Perron-eigenvalue λ_1 has algebraic multiplicity 1, and we can pick the corresponding eigenvector $v > 0$ strictly positive, and no other non-negative eigenvectors exist.

We omit the proof here.

Here is an example to illustrate the contrast of Perron's and Frobenius' theorems:

Example 6.4.3. Let $A = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. Then $A \geq 0$ and A is irreducible (the only 2×2 permutation matrices P are the identity and A itself, and one can check that neither $I^T A I$ nor $A^T A A$ are block-upper triangular). The spectrum of A is $\sigma(A) = \{1, -1\}$, so both eigenvalues have the same absolute value.

An eigenvector to $\lambda_1 = 1$ is $v = \begin{pmatrix} 1 \\ 1 \end{pmatrix} > 0$ (positive as guaranteed by the theorem), and for $\lambda_2 = -1$ an eigenvector is $\begin{pmatrix} 1 \\ -1 \end{pmatrix}$. Note that the power method doesn't work here: with $w_0 = \begin{pmatrix} a \\ b \end{pmatrix}$ we get $w_k = \begin{pmatrix} a \\ b \end{pmatrix}$ when k is even and $w_k = \begin{pmatrix} b \\ a \end{pmatrix}$ when k is odd, so the sequence of vectors w_k fails to converge in general.

△

The primitivity and irreducibility-condition in the above versions of Perron's theorem may seem confusing, but it they have very concrete interpretations if we view the matrix as a *graph*.

6.5 Connection to graphs

Definition 6.5.1. A **directed graph** $G = (V, E)$ consists of a set $V = (v_1, \dots, v_n)$ of vertices, or nodes, and a set $E \subset V \times V$ of edges. An edge (v_i, v_j) should be visualized as an arrow from node i to node j .

Given a directed graph with n nodes, we define its **adjacency matrix** as the $n \times n$ -matrix A where $a_{ij} = 1$ if there is an edge $v_j \rightarrow v_i$ in the graph (note the reversed index order^a), and $a_{ij} = 0$ otherwise.


Conversely, if A is an $n \times n$ -matrix we define its **associated di-graph** to have n nodes, and where $v_j \rightarrow v_i$ is a directed edge if and only if $a_{ij} \neq 0$.

A directed graph is called **strongly connected** if for each pair of vertices v and v' there is a walk from v to v' (a sequence of directed edges from v to v').

^asome texts prefer to use the transposed matrix as the adjacency-matrix

Example 6.5.2. Here is a matrix and its associated directed graph, the graph is strongly connected.

$$\begin{pmatrix} 0 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 0 \end{pmatrix}$$



△

Graphs come in a number of flavours. Graphs where edges do not have directions correspond to symmetric matrices. A similar correspondence between graphs and matrices works for multi-graphs⁴ and weighted⁵ graphs.

Proposition 6.5.3. A matrix A is irreducible if and only if its associated di-graph G is strongly connected.

Proof. One one hand, if the $n \times n$ -adjacency-matrix has block upper triangular form $A = \begin{pmatrix} B & C \\ 0 & D \end{pmatrix}$ with $B \in \text{Mat}_{m \times m}$, then this corresponds to a partition of the nodes: let $V_1 = \{v_1, \dots, v_m\}$ and $V_2 = \{v_{m+1}, \dots, v_n\}$. Then the matrix B corresponds to the arrows $V_1 \rightarrow V_1$, the matrix D corresponds to arrows $V_2 \rightarrow V_2$, and the matrix C corresponds to arrows $V_2 \rightarrow V_1$. The zero block says that there are no arrows $V_1 \rightarrow V_2$, which means that if we start at a vertex in V_1 we can never reach vertices in V_2 , so the matrix is not


⁴In multi-graphs we allow several (directed) edges between two nodes, and we can define its adjacency matrix by a_{ij} =number of directed edges $v_j \rightarrow v_i$

⁵In a weighted graph each (directed) edge has a corresponding weight (usually a positive number) attached, and we can define its adjacency matrix by a_{ij} =weight of $v_j \rightarrow v_i$ if the edge exists, and zero otherwise

strongly connected. But even if A is not block upper-triangular it might be possible to relabel the vertices so that the adjacency-matrix with respect to the new labelling is reducible. This relabelling of nodes corresponds to permuting both rows and columns in the adjacency-matrix by a permutation matrix P , which means that the new adjacency-matrix will have form $A' = PAP^T$, this sets up the correspondence between irreducible matrices and strongly connected di-graphs. \square

Let us illustrate this correspondence with an example.


Example 6.5.4. If we label the nodes of the left graph G as indicated, its adjacency-matrix is the matrix A :



$$A = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}$$

It is not clear whether this adjacency matrix is irreducible.

However, if we take the same graph but switch the labelling of the nodes 2 and 3, we get the new adjacency matrix:



$$A' = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

where the adjacency-matrix A' is now block upper-triangular. Switching the labelling of nodes 2 and 3 has the effect of switching row 2 by row 3, and column 2 by column 3; in other words, multiplying A from the left and from the right by the permutation matrix $P = P^T$ that performs the row/column switches, and indeed it is easy to check that

$$P^TAP = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} = A'.$$

So A is in fact reducible.

This is also easy to see visually in the graph: the original graph is not strongly connected, because if we start at nodes 1 or 3 we can never reach node 2 or 4 (in the original labelling).

\triangle

The graph-theoretic version of primitive matrices has to do with walks⁶ in the graph.

Proposition 6.5.5. *Let $A \geq 0$ be a non-negative square matrix and let $G = (V, E)$ be its associated di-graph. Then A is primitive with $A^p > 0$ if and only if there exists a walk of length exactly p between any pair of nodes in G .*

Proof. Recall how matrix multiplication is defined in terms of the basis-matrices: $e_{i,j}e_{j,k} = e_{i,k}$. This shows that $(A^p)_{i,j}$ is a linear combination of elements of form $a_{i,i_1}a_{i_1,i_2} \cdots a_{i_{p-1},j}$ and this element is nonzero if and only if $v_j \rightarrow v_{i_{p-1}} \rightarrow \cdots \rightarrow v_2 \rightarrow v_1 \rightarrow v_i$ is a walk of length p in G . \square

Consider for example the graphs associated to the matrices A and B in Example 6.3.3.

⁶A "walk" in a graph is a way of moving around in the graph along the directed edges. A walk can be represented as a sequence of vertices $(v_{i_1}, v_{i_2}, \dots, v_{i_n})$ of vertices such that there is an edge between each consecutive pair of vertices in the list

6.6 Ranking models

Given a network of websites linking to each other, how should we rank websites in order of importance? The naive approach of counting incoming links to each website is problematic since someone might just create hundreds of empty websites linking to their own website. So we want to rank a website higher when it has links from websites that are themselves high-ranked. Since ranking values should then depend on themselves this turns into an eigenvalue problem.


PageRank

*PageRank*⁷ is one of the foundational ideas for Google’s original ranking algorithm. The model says that a website should distribute a portion of its own value to all websites that it links to. In other words, if a given website has ranking value r and it links eight other web-pages, it should contribute a number proportional to $\frac{r}{8}$ to each of those websites.

This same ideas can be used in various context where a directed graph describes the data, such as ranking players in tournament, or modelling traffic.

Let us illustrate the ranking algorithm by an example.

Example 6.6.1. Suppose there are only four websites on the web. They link to each other according to the (strongly connected) graph below:



According to our model, the ranking values $r = (r_1, r_2, r_3, r_4)^T$ should satisfy:

$$\begin{cases} r_1 = k(\frac{1}{2}r_2 + \frac{1}{2}r_3 + \frac{1}{3}r_4) \\ r_2 = k(r_1 + \frac{1}{3}r_4) \\ r_3 = k(\frac{1}{2}r_2 + \frac{1}{3}r_4) \\ r_4 = k(\frac{1}{2}r_3) \end{cases} \Leftrightarrow r = k \begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} & \frac{1}{3} \\ 1 & 0 & 0 & \frac{1}{3} \\ 0 & \frac{1}{2} & 0 & \frac{1}{3} \\ 0 & 0 & \frac{1}{2} & 0 \end{bmatrix} r \Leftrightarrow Ar = \lambda r$$

where in the last step we wrote $k = \frac{1}{\lambda}$ for the proportionality constant and A for the coefficient-matrix.

By construction, all the column sums in A are 1 which means that 1 is an eigenvalue of A^T with eigenvector $v = (1, 1, 1, 1)^T$, and since $A^T > 0$ and $v > 0$, the Perron-eigenvalue of A^T is 1. Since A and A^T have the same eigenvalues, 1 is also a dominant eigenvalue of A , and the ranking vector is the unique (up to scaling) positive eigenvector. We compute the eigenvector for A of eigenvalue 1 and find that the ranking vector is

$$r = (0.32, 0.36, 0.21, 0.11)^T,$$

it is unique up to scaling by a positive constant. Note that node 2 has highest ranking despite not being the most linked to website.

△

We generalize the notions from the example:

⁷“Page” as in “Larry Page”, not “webpage”

Definition 6.6.2. Given a directed graph with nodes (v_1, \dots, v_n) , let n_j be the number of edges starting at v_j . Then the **PageRank-matrix** of the graph is the matrix $A = (a_{ij})$ where

$$a_{ij} = \begin{cases} \frac{1}{n_j} & \text{if there is an edge } v_j \rightarrow v_i \\ 0 & \text{otherwise.} \end{cases}$$

In other words, A is the adjacency-matrix of the graph where each nonzero column is scaled so that its sum is 1.

A corresponding **ranking vector** is a Perron-vector of A . If A is irreducible this vector is unique up to scaling.

Damping factor

Numerically it is beneficial to work with a positive matrix instead of only a non-negative one, in particular we do not have to check our irreducible criterion in Frobenius' theorem for a large matrix if we know that it is positive. One approach is to modify the PageRank matrix by choosing $0 < d < 1$ and replacing the PageRank matrix A of a di-graph by a *dampened* version:

Definition 6.6.3. Let A be the PageRank matrix of a di-graph (the column-normalized adjacency matrix). For $0 \leq d \leq 1$, the **dampened PageRank matrix** is defined as

$$A_d = d \cdot A + (1 - d) \cdot \frac{1}{n}(1),$$

where $(1) = (1)_{ij}$ is the matrix full of ones.

The Perron vector r_d to A_d is a **dampened ranking vector** for the di-graph.

We note that for $d < 1$ the matrix A_d is strictly positive, so Perron's theorem applies. The matrix $\frac{1}{n}(1)$ is the PageRank matrix of the complete graph, and A_d is a weighted average between this and our original PageRank matrix.

Note that $A = A_1$, and that high values for d will give similar rankings for our modified matrix. In Google's original approach a value of $d = 0.85$ was used. Google's ranking matrix is an $n \times n$ -matrix where n is the number of websites indexed by Google, estimated⁸ to be over 10^{10} or 10 billion.

Example 6.6.4. Let us go back to our original PageRank illustration in Example 6.6.1. We modify the original PageRank matrix A with a damping factor of $d = 0.9$. We get

$$A_{0.9} = 0.9A + 0.1 \frac{1}{4}(1) = 0.9 \begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} & \frac{1}{3} \\ 1 & 0 & 0 & \frac{1}{3} \\ 0 & \frac{1}{2} & 0 & \frac{1}{3} \\ 0 & 0 & \frac{1}{2} & 0 \end{bmatrix} + 0.025 \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 0.025 & 0.475 & 0.475 & 0.325 \\ 0.925 & 0.025 & 0.025 & 0.325 \\ 0.025 & 0.475 & 0.025 & 0.325 \\ 0.025 & 0.025 & 0.475 & 0.025 \end{bmatrix}.$$

We compute the corresponding ranking vector $r_{0.9}$ for $A_{0.9}$ and compare it to the ranking vector r of $A = A_1$ from Example 6.6.1:


$$r_{0.9} = \begin{pmatrix} 0.31 \\ 0.35 \\ 0.22 \\ 0.12 \end{pmatrix} \quad r = \begin{pmatrix} 0.32 \\ 0.36 \\ 0.21 \\ 0.11 \end{pmatrix}$$

We see that the ranking values are marginally different, but using $d = 0.9$ didn't affect the ranking order.

△

⁸The ranking algorithm has probably evolved significantly in recent years, the material presented here is from Page's and Brin's original scientific paper "The anatomy of a large-scale hypertextual Web search engine" from 1998, the exact algorithm is less public today.

The following image illustrates how page rank applies to a larger network. The nodes are labelled by their ranking values (normalized so that the values sum to 1000), and the node areas are proportional to the ranking values. A dampening factor of $d = 0.85$ was used.



The random surfer

As observed in Page's and Brin's original paper, the PageRank model can be interpreted in terms of user behaviour. The *random surfer* starts at some website and then repeatedly clicks random links on that website. The ranking-vector then describes the proportion of time spent at each website in the long run.

The irreducibility condition in Frobenius' theorem guarantees that the surfer never permanently gets stuck in a nest of websites without links out of them, so called *spider traps*.

In this context, the dampened version of PageRank corresponds to the possibility of the random surfer getting bored: at each time-step there is a probability of $1 - d$ that the random surfer doesn't click a link and instead just surfs to a completely random website.



Applications of ranking models

PageRank has many applications beyond ranking website, any system that can be described by a weighted directed graph is amenable to the techniques of this section. We list a few applications of PageRank:

- Quantifying scientific impact of authors and journals in Scientific research. Here the graph describe citations of publications.
- In biology, PageRank can be used to measure the importance of different species in the eco-system. The graph studied has species as vertices and arrows could correspond to various relationships between them (such as predator-prey relationships).

- Player ratings in Round Robin tournaments where all players play each other, or where each player just plays a subset of the others. The arrows in the graph describes who won between two players (possibly weighted by the score-difference in their game).
- City planning, when determining the scale of various spaces in a city, PageRank can be used to estimate how much space is needed (for example for parking spaces). The vertices of the graph are spaces in the city, and the arrows comes from measurements or estimations of the traffic, the number of persons or vehicles moving from one space to another.

6.7 Markov chains

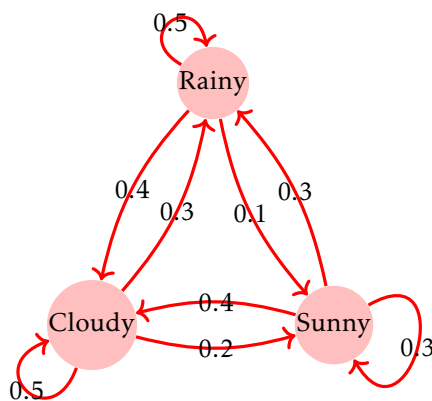
Definition 6.7.1. A **Markov chain**, or a **Markov process**, is a stochastic model in which a system can be in a number of different states (v_1, \dots, v_n) , and given that we are in state j , there is a fixed probability a_{ij} that we transfer to the state i in the next time-step.

A Markov chain corresponds to a **weighted graph** where the vertices are the different states, and for each pair of states v_j and v_i , the "weight" of the arrow $v_j \rightarrow v_i$ is the probability a_{ij} that we will move to the state j from state i . The corresponding matrix $A = (a_{ij})_{ij}$ is called the **transition matrix**^a of the Markov chain. Since the columns of the transition matrix describe probabilities of moving *from* a given state, the columns will sum up to 1. Square matrices where each column consists of non-negative numbers that add up to 1 are called **stochastic matrices**.

^aIt is common to instead call A^T the transition matrix, then the row sums will be 1 instead, this is just a convention.

We illustrate by an example:

Example 6.7.2. Suppose the weather in Linköping changes from day to day according to the following (simplistic) Markov chain model:



For example, the model (somewhat depressingly) says that given that it is rainy today, there is a 50% change that it will be rainy tomorrow too, with 40% chance of clouds tomorrow, and only a 10% chance for sunshine tomorrow.

If we order the states as (Rainy,Cloudy,Sunny), the transition matrix of the system is

$$A = \begin{bmatrix} 0.5 & 0.3 & 0.3 \\ 0.4 & 0.5 & 0.4 \\ 0.1 & 0.2 & 0.3 \end{bmatrix}.$$

Note that the columns add up to 1 since the probability of *some weather* happening tomorrow is 1.

Assuming the weather follows this model, how many days will be sunny on average? We note that given that it is sunny today, represented by the state vector $w = (0, 0, 1)^T$, then the probabilities for to-

tomorrow's weather is described by the vector $Aw = (0.3, 0.4, 0.3)^T$, the probabilities for the day after tomorrow is $A^2w = (0.36, 0.44, 0.20)^T$ and so on. In general, $A^k w$ will describe the weather-probabilities k days after today, and for k very large this vector will correspond to the average probabilities of weathers in the long run. In fact, if the matrix is irreducible with a unique dominant Perron-eigenvalue, with probability 1, the vectors $A^n w$ will converge the Perron-vector for the matrix A , regardless of today's weather w .

In this example, the Perron vector v (scaled so that $\|v\|_1 = 1$) is

$$v = (0.38, 0.44, 0.18)^T,$$

so 18% of days would be sunny on average.

△

Note that in Markov chains, the probabilities for transitions is only dependent on the current state, this is of course not very realistic for predicting weather, as many other factors apply (such as time of year).

We formulate a more general version of our conclusions from the example.

As with the ranking model, since the column sums of the transition matrix is 1, the positive vector $(1, \dots, 1)$ is a Perron vector for A^T with eigenvalue 1, so 1 is also the Perron-eigenvalue for A .

Definition 6.7.3. Let A be an irreducible transition matrix of a Markov chain. Let v be the Perron vector of A normalized such that $\|v\|_1 = 1$ (the sum of the entries is 1). Then v is called the **stationary distribution** or the **steady state distribution** of A , it satisfies $Av = v$ and the entries of v describe the probabilities of being in a given state asymptotically.

Applications of Markov chains

Markov chains appear in several applications in science and technology, we list a few:

- In speech recognition, a given word is more likely to be followed by certain other words. This can be modelled by a Markov chain where words are the nodes of the graph.
- In theoretical physics, Brownian motion and random movement of particles can be modelled by Markov processes.
- In finance, asset pricing can be modelled by viewing the evolution of the market as a random process.
- In neuroscience, the transition of states of a brain (or other neural networks) can be modelled as Markov chains.
- In cryptography, one approach to attack a hidden or unknown system (such as a cryptographic algorithm for which we do not know its source code) is to model how its outputs are affected by its inputs by modelling it as a Markov process.
- In gene-prediction, a subject in computational biology, one use Markov models to predict what regions of DNA correspond to particular functional elements of the organism.


Chapter 7

Singular values

7.1 Geometrical intuition

Singular values can be viewed as a generalization of eigenvalues that exist for *any* complex matrix or linear operator between inner product spaces. Singular values have a multitude applications both in mathematics and science and technology.

Let us start with some geometrical intuition for what singular values are. Let $F : W \rightarrow W'$ be a map between two three-dimensional inner product spaces, and pick any ON-bases in W and in W' . Then, working with coordinate vectors, we can think of F as a 3×3 -matrix A , or as a map $\mathbb{C}^3 \rightarrow \mathbb{C}^3$. Let us assume for now that F is invertible¹, then the unit ball in W will be mapped to an ellipsoid in W' .



Now we pick a new ON-basis $\mathcal{B}' = (u_1, u_2, u_3)$ for W' such that the basis vectors point in the direction of axes of the ellipsoid (but the u_i are normalized), ordered so that longer semi-axes come before short ones, then the semi-axes of the ellipsoid are described by the rescaled vectors $\sigma_1 u_1$, $\sigma_2 u_2$, and $\sigma_3 u_3$. Since these vectors lie on the ellipsoid surface, their inverse image $v_i := F^{-1}(\sigma_i u_i)$ under F lie on the sphere surface, and they will in fact form a new ON-basis $\mathcal{B} = (v_1, v_2, v_3)$ for W . This means that with respect to these two choices of bases, we have

$$\Sigma := [F]_{\mathcal{B}', \mathcal{B}} = \begin{pmatrix} \sigma_1 & 0 & 0 \\ 0 & \sigma_2 & 0 \\ 0 & 0 & \sigma_3 \end{pmatrix},$$

so $[F]$ has diagonal form with real diagonal elements σ_i in decreasing order: $\sigma_1 \geq \sigma_2 \geq \sigma_3$ because of how we ordered our basis for W' . But note that for the map $F : W \rightarrow W'$ we are using *different* bases in the domain W and the codomain W' of the map, despite both being viewed as the same space \mathbb{C}^3 , so this is not a diagonalization of the map (F may in fact not be diagonalizable).

¹if F is not invertible, say $\text{rank}(F) = 2$, then the same thing happens except that the image of the unit circle will be flattened to an ellipse in W' , and we have to extend our ellipse-axis-basis to an ON-basis in all of \mathbb{C}^3 , this corresponds to one of the sigmas being zero in what follows.

Here the real non-negative values σ_i are called the *singular values* of F , they describe the length of the axes of the ellipsoid. The vectors v_1, v_2, v_3 are called *right singular vectors*, they are mapped to multiples of the vectors u_1, u_2, u_3 which are called *left singular vectors*. The left singular vectors u_i point in the axis-directions of the ellipsoid.

If we put these basis vectors as columns in two matrices $V := (v_1 \ v_2 \ v_3)$ and $U = (u_1 \ u_2 \ u_3)$, the equations $F(v_i) = \sigma_i u_i$ can be expressed in matrix form as $[F]V = U\Sigma$, where $[F]$ is the matrix for F with respect to the standard bases, or equivalently

$$[F] = U\Sigma V^*$$

which is called the singular value decomposition of F , note that the matrices U and V are unitary. We also remark that the columns of U are the left singular vectors and the columns of V are the right singular vectors, corresponding to the order of the factorization $A = U\Sigma V^*$.

7.2 Singular values and singular vectors

Let us try to formalize and generalize our introductory example.

Definition 7.2.1. Let A be a complex $m \times n$ -matrix. If

$$Av = \sigma u \quad \text{and} \quad A^*u = \sigma v$$

for some $\sigma \geq 0$, and some columns $v \in \mathbb{C}^n$ and $u \in \mathbb{C}^m$, then σ is called a **singular value**, and v is called a **right singular vector**, and u is called a **left singular vector**.

A **singular value decomposition** or just the **SVD** of A is a factorization

$$A = U\Sigma V^*$$

where U is a unitary $m \times m$ matrix, and V is a unitary $n \times n$ -matrix, and where Σ is a (pseudo) diagonal $m \times n$ matrix

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

with real non-negative elements $\Sigma_{ii} = \sigma_i$ on the diagonal ordered such that

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p \geq 0,$$

where $p = \min(m, n)$. We note that this implies that for $1 \leq i \leq p$ we have

$$Av_i = \sigma_i u_i \quad \text{and} \quad A^*u_i = \sigma_i v_i,$$

so $\sigma_1, \dots, \sigma_p$ are singular values, and the first p columns of U are left singular vectors, and the first p columns of V are right singular vectors.

You may note that there may be singular vectors that do not appear as columns in U and V (consider for example $A = I$), but in fact every left singular vector is a linear combination of the first p columns U , and similarly, every right singular vector is a linear combination of the first p columns of V .

Now suppose that $A = U\Sigma V^*$, and suppose that there are exactly r nonzero singular values. Let $\tilde{\Sigma}$ be the diagonal $r \times r$ -matrix containing these values (the upper left corner of Σ). Let \tilde{V} contain only the first r columns of V with $V = [\tilde{V} | V_0]$, and let \tilde{U} contain the first r columns of U with $U = [\tilde{U} | U_0]$.

Then a block-matrix computation yields

$$A = U\Sigma V^* = \begin{bmatrix} \tilde{U} & | & U_0 \end{bmatrix} \begin{bmatrix} \tilde{\Sigma} & | & 0 \\ \hline 0 & | & 0 \end{bmatrix} \begin{bmatrix} \tilde{V}^* \\ \hline -V_0^* \end{bmatrix} = \begin{bmatrix} \tilde{U}\tilde{\Sigma} & | & 0 \end{bmatrix} \begin{bmatrix} \tilde{V}^* \\ \hline -V_0^* \end{bmatrix} = \tilde{U}\tilde{\Sigma}\tilde{V}^*$$

This factorization on the right is called a *compact SVD* of A , note that \tilde{U} and \tilde{V} are no longer unitary in general since they typically are not square: \tilde{V} is $m \times r$ and \tilde{U} is $n \times r$, while $\tilde{\Sigma}$ is an invertible $r \times r$ -matrix. The number r is the rank of the original matrix A .

We summarize:

Definition 7.2.2. Let A be a complex $m \times n$ -matrix with rank r . A **compact SVD** is a factorization

$$A = \tilde{U} \tilde{\Sigma} \tilde{V}^*$$

where \tilde{U} is an $m \times r$ matrix with orthonormal columns, and \tilde{V} is an $n \times r$ matrix with orthonormal columns, and where $\tilde{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_r)$ is diagonal, and $\sigma_1 \geq \dots \geq \sigma_r > 0$.

We note that $\tilde{\Sigma}$ contains the nonzero singular values and the columns of \tilde{U} and of \tilde{V} are left and right singular vectors corresponding to nonzero singular values.

How do we actually find the SVD? Let us illustrate by an example before looking at the general algorithm:

Example 7.2.3. Let us find the SVD and the compact SVD for the matrix

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

Let us start with the compact SVD. Suppose that $A = \tilde{U} \tilde{\Sigma} \tilde{V}^*$. Then

$$A^*A = (\tilde{U} \tilde{\Sigma} \tilde{V}^*)^*(\tilde{U} \tilde{\Sigma} \tilde{V}^*) = \tilde{V} \tilde{\Sigma}^* \tilde{U}^* \tilde{U} \tilde{\Sigma} \tilde{V}^* = \tilde{V} \tilde{\Sigma}^* \tilde{\Sigma} \tilde{V}^* = \tilde{V} \tilde{\Sigma}^2 \tilde{V}^* = \tilde{V} D \tilde{V}^*.$$

This suggests that we can find \tilde{V} and $\tilde{\Sigma}$ by orthogonally diagonalizing the matrix A^*A . We have

$$A^*A = \begin{pmatrix} 1 & 1 & 1 \\ 1 & -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \end{pmatrix} = \begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix},$$

and since this matrix is Hermitian it has an orthogonal diagonalization. We get $\det(A^*A - \lambda I) = (\lambda - 4)(\lambda - 2)$ so the eigenvalues of A^*A are 4 and 2. We pick $v_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ as a normalized eigenvector for A^*A of eigenvalue 4, and we pick $v_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ 1 \end{pmatrix}$ as a normalized eigenvector of eigenvalue 2. We put these vectors as columns in a matrix \tilde{V} , and we let D be the diagonal matrix with 4 and 2 on the diagonal. We conclude that with

$$\tilde{V} = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} \quad D = \begin{pmatrix} 4 & 0 \\ 0 & 2 \end{pmatrix}$$

we have $A^*A = \tilde{V} D \tilde{V}^*$. Since $D = \tilde{\Sigma}^2$ and we want $\tilde{\Sigma}$ to have non-negative entries, we need to take $\tilde{\Sigma} = \sqrt{D} = \begin{pmatrix} 2 & 0 \\ 0 & \sqrt{2} \end{pmatrix}$, so the singular values are $\sigma_1 = 2$ and $\sigma_2 = \sqrt{2}$. Now we have found \tilde{V} and $\tilde{\Sigma}$. Since $A = \tilde{U} \tilde{\Sigma} \tilde{V}^*$ we get $A \tilde{V} = \tilde{U} \tilde{\Sigma}$ which means

$$A v_1 = \sigma_1 u_1 \quad \text{and} \quad A v_2 = \sigma_2 u_2,$$

so we can find the columns of \tilde{U} :

$$u_1 = \frac{1}{\sigma_1} A v_1 = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \end{pmatrix} \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} \quad u_2 = \frac{1}{\sigma_2} A v_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \end{pmatrix} \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ -1 \\ 0 \end{pmatrix}.$$

We note that u_1 and u_2 have length 1 and are orthogonal to each other, so \tilde{U} has orthonormal columns as desired. So we conclude that a compact SVD of A is given by

$$A = \tilde{U} \tilde{\Sigma} \tilde{V}^* \quad \text{where} \quad \tilde{U} = \begin{pmatrix} \frac{1}{\sqrt{2}} & 0 \\ 0 & -1 \\ \frac{1}{\sqrt{2}} & 0 \end{pmatrix} \quad \tilde{\Sigma} = \begin{pmatrix} 2 & 0 \\ 0 & \sqrt{2} \end{pmatrix} \quad \tilde{V} = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}.$$

Very little work is now needed to find the full (non-compact) SVD of A . We need U to be a unitary 3×3 -matrix, so we need to extend u_1, u_2 to an ON-basis for \mathbb{C}^3 . We pick $u_3 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix}$ and we take U as the matrix with columns u_1, u_2, u_3 . The matrix Σ needs to be of shape 3×2 (the same shape as the original matrix A), so here we need to add a row of zeros to $\tilde{\Sigma}$. The matrix V needs to be a unitary 2×2 -matrix which \tilde{V} already is, so we take $V = \tilde{V}$.

We conclude that an SVD of A is given by

$$A = U\Sigma V^* \quad \text{where} \quad U = \begin{pmatrix} \frac{1}{\sqrt{2}} & 0 & \frac{1}{\sqrt{2}} \\ 0 & -1 & 0 \\ \frac{1}{\sqrt{2}} & 0 & -\frac{1}{\sqrt{2}} \end{pmatrix} \quad \Sigma = \begin{pmatrix} 2 & 0 \\ 0 & \sqrt{2} \\ 0 & 0 \end{pmatrix} \quad V = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}.$$

△

We formulate the method employed in the example as a general algorithm:

Algorithm 7.2.4. To find the singular value decomposition $A = U\Sigma V^*$ and the compact singular value decomposition $A = \tilde{U}\tilde{\Sigma}\tilde{V}^*$ of an $m \times n$ -matrix A do the following:

1. Assume that $m \geq n$ so that the matrix is *tall*, otherwise find the SVD of A^* first, and take the conjugate transpose of that factorization.
2. Find the eigenvalues of A^*A (these are real and non-negative) and order them in decreasing order $\lambda_1 \geq \dots \geq \lambda_n$, the singular values are the square roots of these eigenvalues: $\sigma_i = \sqrt{\lambda_i}$, giving

$$\sigma_1 \geq \dots \geq \sigma_n \geq 0.$$
3. Let r be the number of nonzero singular values σ_i . Then r is the rank of the matrix A .
4. Put the nonzero singular values in decreasing order in an $r \times r$ -matrix $\tilde{\Sigma}$. Let Σ be the corresponding $m \times n$ -matrix (same size as A) obtained by padding $\tilde{\Sigma}$ by zero-rows on the bottom and zero-columns on the right.
5. Find a corresponding orthonormal basis v_i of eigenvectors for A^*A , such that $A^*Av_i = \sigma_i^2 v_i$. Put these orthonormal eigenvectors as columns (in the same order) in an $n \times n$ -matrix $V = (v_1 \dots v_n)$. Let \tilde{V} be the (smaller) $n \times r$ -matrix $\tilde{V} = (v_1 \dots v_r)$ consisting of the first r columns of V (corresponding to nonzero singular values).
6. For each nonzero singular value σ_i , take $u_i := \frac{1}{\sigma_i} Av_i$. These vectors are orthonormal. Let $\tilde{U} = (u_1 \dots u_r)$ be the $m \times r$ -matrix with these vectors as columns.
7. Extend u_1, \dots, u_r to a basis u_1, \dots, u_m for \mathbb{C}^m . Let $U = (u_1 \dots u_m)$, this is the $m \times m$ -matrix obtained from \tilde{U} by adjoining the new basis vectors u_{r+1}, \dots, u_m as columns on the right, then U is unitary.

8. Now

$$\underbrace{A}_{m \times n} = \underbrace{U}_{m \times m} \underbrace{\Sigma}_{m \times n} \underbrace{V^*}_{n \times n}$$

is a singular value decomposition for A . Similarly,

$$\underbrace{A}_{m \times n} = \underbrace{\tilde{U}}_{m \times r} \underbrace{\tilde{\Sigma}}_{r \times r} \underbrace{\tilde{V}^*}_{r \times n}$$

is a compact singular value decomposition for A . This can be verified by multiplying the matrices.

The following result shows that the above algorithm always works:

Theorem 7.2.5. Every complex matrix has a singular value decomposition $A = U\Sigma V^*$ and a corresponding compact singular value decomposition $A = \tilde{U}\tilde{\Sigma}\tilde{V}^*$ where the matrix Σ and $\tilde{\Sigma}$ of singular values is are unique.

Additionally, if the nonzero singular values are all distinct, then the matrices \tilde{U} and \tilde{V} are also unique up to scaling their columns by complex numbers.

Proof. Let A have size $m \times n$. We follow Algorithm 7.2.4, it remains only to verify that all the steps will work for a general matrix A .

First we note that if $A = U\Sigma V^*$, then $A^* = V^{**}\Sigma^*U^* = V\Sigma^T U^*$ where Σ^T is still diagonal with non-negative entries, and where the matrices to the left and right are still unitary, so given a SVD of A we automatically get the SVD for A^* , for this reason we can assume without loss of generality that $m \geq n$, such that the matrix A is "tall" rather than "wide".

Now the square $n \times n$ -matrix A^*A satisfies $(A^*A)^* = A^*A$ so it is Hermitian and therefore normal, so by the spectral theorem there is a corresponding orthonormal basis of eigenvectors for A^*A in \mathbb{C}^n . Moreover, A^*A is always positive semi-definite:

$$\langle A^*AX, X \rangle = \langle AX, AX \rangle = \|AX\|^2 \geq 0 \text{ for all columns } X \in \mathbb{C}^n.$$

Therefore the eigenvalues λ_i of A^*A are indeed non-negative real numbers, and $\sigma_i = \sqrt{\lambda_i}$ is well-defined.

Next we need to verify that the first r columns of U , defined as $u_i = \frac{1}{\sigma_i}Av_i$ are indeed orthonormal vectors in \mathbb{C}^m . Using the fact that v_i are orthonormal eigenvectors of A^*A we get

$$u_j \bullet u_i = \frac{1}{\sigma_i\sigma_j}(Av_j) \bullet (Av_i) = \frac{1}{\sigma_i\sigma_j}(Av_i)^*(Av_j) = \frac{1}{\sigma_i\sigma_j}v_i^*A^*Av_j = \frac{1}{\sigma_i\sigma_j}v_i^*(\sigma_j^2v_j) = \frac{\sigma_j^2}{\sigma_i\sigma_j}v_j \bullet v_i = \frac{\sigma_j^2}{\sigma_i\sigma_j}\delta_{ij} = \delta_{ij},$$

so u_1, \dots, u_r are orthonormal.

We also claimed that the number of nonzero singular values was the rank of A . Indeed, since multiplying by an invertible matrix from the left or right doesn't change the rank, we have

$$\text{rank}(A) = \text{rank}(U\Sigma V^*) = \text{rank}(\Sigma) = r.$$

For the uniqueness-claim, we note that if $A = U\Sigma V^*$, then since $A^*A = V\Sigma^*\Sigma V^*$, the diagonal non-negative matrix $\Sigma^*\Sigma$ necessarily has the eigenvalues of A^*A as on the diagonal, these are non-negative since A^*A is positive semi-definite, so since Σ should have non-negative entries in decreasing order on the diagonal, it is uniquely determined by A . If the singular values are distinct, then each eigenvalue of A^*A has algebraic multiplicity 1, so the columns of \tilde{V} are also uniquely-determined up to multiplication by complex numbers c (where $|c| = 1$ since the columns of V need to be normalized). But then the columns of \tilde{U} are also uniquely determined by \tilde{V} up to a corresponding scaling. \square

We can interpret the singular value decomposition in a basis-free way.

Corollary 7.2.6. *Let $F : W \rightarrow W'$ be a linear map between finite-dimensional inner product spaces. Then for $p = \min(\dim W, \dim W')$, there exists non-negative real numbers $\sigma_1 \geq \dots \geq \sigma_p$, and there exist orthonormal vectors (v_1, \dots, v_p) in W and orthonormal vectors (u_1, \dots, u_p) in W' such that*

$$F(v_i) = \sigma_i u_i \quad \text{and} \quad F^*(u_i) = \sigma_i v_i.$$

In other words, we can always choose an ON-basis B for W and an ON-basis B' for W' such that the matrix for F becomes pseudo-diagonal with non-negative diagonal entries.

Proof. Pick any ON-bases in B_0 in W and pick any ON-basis B'_0 in W' , and let $A = [F]_{B'_0, B_0}$ be the matrix for F with respect to these bases. Let $A = U\Sigma V^*$ be an SVD of A . Then with respect to the basis B_0 , the columns of V are coordinate vectors for an ON-basis $B = (v_1, \dots, v_n)$ for W , and with respect to the basis B'_0 , the columns of U forms an ON-basis $B' = (u_1, \dots, u_m)$ of W' . Then with respect to these bases we have $[F]_{B', B} = \Sigma$ and $[F^*]_{B, B'} = \Sigma^* = \Sigma^T$, so in particular $F(v_i) = \sigma_i u_i$ and $F^*(u_i) = \sigma_i v_i$. \square

Corollary 7.2.7. *For any complex $m \times n$ -matrix A , the matrices A^*A and AA^* always have the same set of nonzero eigenvalues.*

Proof. Let $A = U\Sigma V^*$ be an SVD of A . Then $A^*A = V\Sigma^*\Sigma V^*$ and $AA^* = U\Sigma\Sigma^*U^*$ are orthogonal diagonalizations of the matrices A^*A and AA^* respectively. Thus the nonzero eigenvalues of the $n \times n$ -matrix A^*A are $\sigma_1^2, \dots, \sigma_r^2$, which is the same as the nonzero eigenvalues of the $m \times m$ -matrix AA^* . Note that if $m \neq n$, the algebraic multiplicity of the eigenvalue 0 will differ in A^*A and AA^* . \square

Let us look at a bigger example:

Example 7.2.8. Let us find the SVD and the compact SVD of the matrix

$$A = \frac{1}{6} \begin{pmatrix} 4 & 5 & 2 \\ 0 & 3 & 6 \\ 4 & 5 & 2 \\ 0 & 3 & 6 \end{pmatrix}.$$

We follow Algorithm 7.2.4. The matrix is already tall. We find that

$$A^*A = \frac{1}{9} \begin{pmatrix} 8 & 10 & 4 \\ 10 & 17 & 14 \\ 4 & 14 & 20 \end{pmatrix}.$$

With the standard method we find an orthonormal eigenbasis, the eigenvalues are $(\lambda_1, \lambda_2, \lambda_3) = (4, 1, 0)$ (note that we placed them in decreasing order) and the corresponding orthonormal eigenvectors can be picked as

$$v_1 = \frac{1}{3} \begin{pmatrix} 1 \\ 2 \\ 2 \end{pmatrix}, \quad v_2 = \frac{1}{3} \begin{pmatrix} 2 \\ 1 \\ -2 \end{pmatrix}, \quad v_3 = \frac{1}{3} \begin{pmatrix} 2 \\ -2 \\ 1 \end{pmatrix}.$$

Now the singular values are the roots of the eigenvalues so $(\sigma_1, \sigma_2, \sigma_3) = (2, 1, 0)$. Two of these are nonzero, so the rank of A is $r = 2$. We put the singular values in an 4×3 -matrix Σ (the same shape as our original matrix), and we put our ON-basis v_1, v_2, v_3 as columns in a matrix V . Since the rank is 2, we also make smaller versions of these matrices for our compact SVD, $\tilde{\Sigma}$ will be the 2×2 -matrix containing the two nonzero singular values, and \tilde{V} will consist of the two first columns of V :

$$\Sigma = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad \tilde{\Sigma} = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} \quad V = \frac{1}{3} \begin{pmatrix} 1 & 2 & 2 \\ 2 & 1 & -2 \\ 2 & -2 & 1 \end{pmatrix} \quad \tilde{V} = \frac{1}{3} \begin{pmatrix} 1 & 2 \\ 2 & -2 \end{pmatrix}.$$

Lastly we need to find the basis vectors u_i for \mathbb{C}^4 . We have two nonzero singular values, so we take $u_1 = \frac{1}{\sigma_1}Av_1 = \frac{1}{2}(1, 1, 1, 1)^T$ and $u_2 = \frac{1}{\sigma_2}Av_2 = \frac{1}{2}(1, -1, 1, -1)^T$, these will be orthonormal (indeed we could have just calculated the integer matrix product $(6A)(3v_1)$ and then normalized afterwards). We put u_1 and u_2 as columns in a matrix \tilde{U} . To find the matrix U we extend (u_1, u_2) to an ON-basis for \mathbb{C}^4 with the standard method, and we put these four vectors u_i as columns in a 4×4 -matrix U . For example, we can choose $u_3 = \frac{1}{2}(1, 1, -1, -1)^T$ and $u_4 = \frac{1}{2}(1, -1, -1, 1)^T$, which gives us:

$$U = \frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix} \quad \tilde{U} = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

Now we have our SVD of A :

$$A = U\Sigma V^* = \frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix} \begin{pmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \frac{1}{3} \begin{pmatrix} 1 & 2 & 2 \\ 2 & 1 & -2 \\ 2 & -2 & 1 \end{pmatrix},$$

and the corresponding compact SVD is

$$A = \tilde{U}\tilde{\Sigma}\tilde{V}^* = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} \frac{1}{3} \begin{pmatrix} 1 & 2 & 2 \\ 2 & 1 & -2 \end{pmatrix},$$

note that $V^* = V$ in this example, but this is not the case in general.

△

Let us now verify the original claim in this section: that a linear map transforms the unit ball to a

(hyper)-ellipsoid. Let $F : W \rightarrow W'$ be linear with $\dim W = n$ and $\dim W' = m$, let $r = \text{rank}(F)$ and let $p = \min(m, n)$. Pick ON-bases in W and in W' according to the singular value decomposition, and let us express vectors as column-matrices with respect to these bases. Let $X = (x_1, \dots, x_n)^T \in W$ be a vector on the surface of the unit ball, meaning $\|X\|^2 = 1$.

Then $F(X) = \Sigma X = (\sigma_1 x_1, \sigma_2 x_2, \dots, \sigma_r x_r, 0, \dots, 0) = (y_1, \dots, y_m)$. So $y_i = 0$ for $i > r$ and


$$1 = |x_1|^2 + \dots + |x_n|^2 = \frac{|y_1|^2}{\sigma_1^2} + \dots + \frac{|y_r|^2}{\sigma_r^2} + |x_{r+1}|^2 + \dots + |x_n|^2 \geq \frac{|y_1|^2}{\sigma_1^2} + \dots + \frac{|y_r|^2}{\sigma_r^2}.$$

So the image $Y = (y_1, \dots, y_m)$ of the unit vector X satisfies

$$\frac{|y_1|^2}{\sigma_1^2} + \dots + \frac{|y_r|^2}{\sigma_r^2} \leq 1$$

with equality if and only if the last components of X are zero: $x_{r+1} = \dots = x_n = 0$, and the equations $\frac{|y_1|^2}{\sigma_1^2} + \dots + \frac{|y_r|^2}{\sigma_r^2} \leq 1$ and $y_{r+1} = \dots = y_m = 0$ indeed defines a hyper-ellipsoid in \mathbb{C}^m where the lengths of the semi-axes are $\sigma_1, \dots, \sigma_r$.

Let us look at the geometric interpretation of the SVD of the matrix $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$.



The top half of the image describes the map A with respect to the standard basis, it shows how the unit disk is mapped to an ellipse, and it shows that the colored thick arrows $(1, 0)$ and $(0, 1)$ are mapped to $(1, 0)$ and $(1, 1)$ respectively. The dashed black arrows top left represents our new choice of basis v_1, v_2 , these are mapped by A to directions aligned the semi-axes of the ellipse, so rescaling these to have length one we get our second ON-basis u_1, u_2 . The image shows how the same map A can be viewed as a composition of three linear transformation:

- First rotate the ball by applying the unitary transformation V^* , this maps v_1, v_2 to the standard basis $(1, 0)$ and $(0, 1)$.
- Then stretch the ball with a factor $\sigma_1 \approx 1.62$ along the x -axis and compress it along the y -axis by a factor $\sigma_2 \approx 0.62$. This produces an ellipse aligned with the standard axes.
- Finally rotate the ellipse by applying the unitary transformation U , which maps the standard basis to u_1, u_2 .

The same idea works in general, an $m \times n$ -matrix maps the n -dimensional unit ball into a hyper-ellipsoid in m -dimensional space. If the dimensions of the domain and the codomain differ (the matrix is non-square), then the stretching step also involves embedding or projecting the ball into or onto a space of different dimension, but in the above example, A was a map from \mathbb{C}^2 to \mathbb{C}^2 .

The moral of the story is:

Any linear transformation between inner product spaces can be viewed as a combination of a rotation, a stretch along the coordinate axes, and another rotation. The stretching factors correspond to the singular values.

7.3 Singular values and matrix-norms

Several norms for matrices and linear operators can be expressed in terms of the singular values.

Proposition 7.3.1. *Let A be a complex matrix with singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p \geq 0$.*

- *If A is square, the spectral norm of A is $\|A\|_\sigma = \sigma_1$.*
- *More generally, for any shape of A , the operator norm of A is also $\|A\|_{\text{op}} = \sigma_1$.*
- *The Frobenius norm of A is $\|A\|_F = \sqrt{\sigma_1^2 + \dots + \sigma_p^2}$.*

Additionally, we define the **nuclear norm** of A as

$$\|A\|_{\blacktriangle} = \sigma_1 + \dots + \sigma_p.$$

Proof. Let $A = U\Sigma V^*$ be an SVD of A . The spectral norm was defined as $\max\{\sqrt{\lambda} \mid \lambda \in \sigma(A^*A)\}$, since $A^*A = V\Sigma^*\Sigma V^*$, the singular values are precisely the square roots of the eigenvalues of A^*A so the maximum equals σ_1 .

For the second claim, recall that the operator norm of an operator $F : W \rightarrow W'$ is defined as $\|F\|_{\text{op}} = \max_{\|v\|=1} \|F(v)\|$, and if A is the matrix of F this becomes

$$\begin{aligned} \|A\|_{\text{op}} &= \max_{\|X\|=1} \|AX\| = \max_{\|X\|=1} \|U\Sigma V^*X\| \\ &= \max_{\|Y\|=1} \|\Sigma Y\| = \max_{\|Y\|=1} \sqrt{\sigma_1^2|y_1|^2 + \dots + \sigma_p^2|y_p|^2} \leq \max_{\|Y\|=1} \sqrt{\sigma_1^2|y_1|^2 + \dots + \sigma_1^2|y_p|^2} = \max_{\|Y\|=1} \sqrt{\sigma_1^2} \|Y\| = \sigma_1. \end{aligned}$$

where we took $Y = V^*X$ and used the fact that unitary transformations do not change the norm. When taking $Y = (1, 0, \dots, 0)^T$ we see that the inequality above is an equality and the maximum is attained.

The Frobenius norm of A is defined as

$$\|A\|_F = \sqrt{\langle A, A \rangle_F} = \sqrt{\text{tr}(A^*A)} = \sqrt{\text{tr}(V\Sigma^*\Sigma V^*)} = \sqrt{\text{tr}(\Sigma^*\Sigma)} = \sqrt{\sigma_1^2 + \dots + \sigma_p^2}$$

as claimed. □

7.4 SVD as a sum - Schmidt-decomposition

The singular value decomposition allows us to write any matrix as a *sum of rank 1 matrices*. Matrices of rank one can always be written as the product of a column-matrix and a row-matrix (in that order!) as the following example illustrates.

Example 7.4.1. Consider the following matrix product of a column and a row:

$$CR = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} \begin{pmatrix} 2 & -1 & 1 & 0 & 3 \end{pmatrix} = \begin{pmatrix} 2 & -1 & 1 & 0 & 3 \\ 4 & -2 & 2 & 0 & 6 \\ 6 & -3 & 3 & 0 & 9 \end{pmatrix}$$

We note that every column is parallel with C and every row is parallel with R , so the matrix CR has rank 1. △

Any rank 1 matrix has the property that all the rows are parallel and that all the columns are parallel, and it can therefore be expressed as a product of a column by a row as in this example.

Now let A be an $m \times n$ -matrix of rank r , and let

$$A = U\Sigma V^*$$

be an SVD of A . Let $\sigma_1, \dots, \sigma_r$ be the nonzero singular values in decreasing order, and write u_1, \dots, u_m for the columns of U and write v_1, \dots, v_n for the columns of V . Then viewing A as a product of block matrices we have

$$A = \begin{pmatrix} | & & | \\ u_1 & \cdots & u_m \\ | & & | \end{pmatrix} \begin{pmatrix} \sigma_1 & & & \\ & \ddots & & \\ & & \sigma_r & \\ & & & 0 \\ & & & & \ddots \end{pmatrix} \begin{pmatrix} - & v_1^* & - \\ & \vdots & \\ - & v_n^* & - \end{pmatrix} = \sigma_1 u_1 v_1^* + \cdots + \sigma_r u_r v_r^* = \sum_{i=1}^r \sigma_i u_i v_i^*.$$

Here we have expressed A as a linear combination of rank 1 matrices $u_i v_i^*$, this is called the *Schmidt decomposition*² of A .

Definition 7.4.2. A **Schmidt decomposition** of a complex $m \times n$ -matrix A is an expression

$$A = \sigma_1 S_1 + \sigma_2 S_2 + \cdots + \sigma_r S_r$$

where the σ_i are real positive numbers and the $m \times n$ -matrices S_i have rank 1 and are orthonormal with respect to the Frobenius inner product.

If $A = U\Sigma V^*$ is an SVD of A , then a Schmidt decomposition is given explicitly by

$$A = \sum_{i=1}^r \sigma_i u_i v_i^*,$$

where σ_i are the singular values in decreasing order, u_i are the columns of U , and v_i are the columns of V , and $r = \text{rank}(A)$.

The matrices $u_i v_i^*$ clearly have rank 1, and they are orthonormal (see the remarks after Corollary 5.7.3), but they do not form an ON-basis for $\text{Mat}_{m \times n}(\mathbb{C})$, because $r \leq \min(m, n) < m \cdot n$. When all the nonzero singular values are different, the Schmidt-decomposition is unique.

Example 7.4.3. Let us consider the compact SVD we found for the matrix in Example 7.2.8. We had found that

$$\frac{1}{6} \begin{pmatrix} 4 & 5 & 2 \\ 0 & 3 & 6 \\ 4 & 5 & 2 \\ 0 & 3 & 6 \end{pmatrix} = A = \tilde{U} \tilde{\Sigma} \tilde{V}^* = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} \frac{1}{3} \begin{pmatrix} 1 & 2 & 2 \\ 2 & 1 & -2 \end{pmatrix}.$$

Thus the Schmidt-decomposition of A is

$$A = \sigma_1 u_1 v_1^* + \sigma_2 u_2 v_2^* = 2 \cdot \frac{1}{6} \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \begin{pmatrix} 1 & 2 & 2 \end{pmatrix} + 1 \cdot \frac{1}{6} \begin{pmatrix} 1 \\ -1 \\ 1 \\ -1 \end{pmatrix} \begin{pmatrix} 2 & 1 & -2 \end{pmatrix}$$

²Schmidt decomposition can be defined more generally as an expression of a vector $w \in V_1 \otimes V_2$ in the *tensor product* of two vector spaces. However, a linear map $F : V \rightarrow W$ can be viewed as an element of the tensor product $W \otimes V^*$, which gives the connection to the Schmidt decomposition defined here. Tensor products will be discussed later when we talk about multilinear algebra.

$$= 2 \cdot \frac{1}{6} \begin{pmatrix} 1 & 2 & 2 \\ 1 & 2 & 2 \\ 1 & 2 & 2 \\ 1 & 2 & 2 \end{pmatrix} + 1 \cdot \frac{1}{6} \begin{pmatrix} 2 & 1 & -2 \\ -2 & -1 & 2 \\ 2 & 1 & -2 \\ -2 & -1 & 2 \end{pmatrix} = 2S_1 + 1S_2.$$

△

7.5 Low rank approximation

Let A be a complex $m \times n$ -matrix. The Schmidt decomposition gives

$$A = \sum_{i=1}^r \sigma_i u_i v_i^* = \sigma_1 u_1 v_1^* + \sigma_2 u_2 v_2^* + \dots + \sigma_r u_r v_r^*,$$

where r is the rank of A . Since $\|u_i v_i^*\|_F = 1$ and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r \geq 0$, the terms in this sum gets smaller and smaller, so we obtain a good approximation of A by simply cutting of the tail of the sum.

Definition 7.5.1. Let A be an $m \times n$ -matrix and let $A = \sum_{i=1}^r \sigma_i u_i v_i^*$ be its SVD as a sum (its Schmidt-decomposition). For $0 \leq k \leq r$, the corresponding **low rank approximation** of rank k is the matrix

$$A_{(k)} := \sum_{i=1}^k \sigma_i u_i v_i^*.$$

Equivalently, if $A = U \Sigma V^*$ is an SVD of A , we obtain $A_{(k)} = U \Sigma_{(k)} V^*$, where in $\Sigma_{(k)}$ we keep only the first k singular values on the diagonal, and replace the others by zero. This is also called a **truncated SVD** of A .

Then for $k \leq r$, the matrix $A_{(k)}$ clearly has rank k : since U and V^* are invertible we have

$$\text{rank}(A_{(k)}) = \text{rank}(U \Sigma_{(k)} V^*) = \text{rank}(\Sigma_{(k)}) = k.$$

Let us work out a small example first.

Example 7.5.2. Let $A = \begin{pmatrix} 13 & 1 & 9 & -3 \\ 11 & 5 & 15 & 9 \\ 5 & 17 & 3 & 15 \end{pmatrix}$. Then $\text{rank}(A) = 3$, and the standard algorithm produces the SVD

$$A = U \Sigma V^* = \frac{1}{3} \begin{pmatrix} 1 & 2 & 2 \\ 2 & 1 & -2 \\ 2 & -2 & 1 \end{pmatrix} \cdot \begin{pmatrix} 30 & 0 & 0 & 0 \\ 0 & 18 & 0 & 0 \\ 0 & 0 & 6 & 0 \end{pmatrix} \cdot \frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}$$

Expressing this as a sum of rank one matrices we get

$$A = \sum_{i=1}^3 \sigma_i u_i v_i^* = \begin{pmatrix} 5 & 5 & 5 & 5 \\ 10 & 10 & 10 & 10 \\ 10 & 10 & 10 & 10 \end{pmatrix} + \begin{pmatrix} 6 & -6 & 6 & -6 \\ 3 & -3 & 3 & -3 \\ -6 & 6 & -6 & 6 \end{pmatrix} + \begin{pmatrix} 2 & 2 & -2 & -2 \\ -2 & -2 & 2 & 2 \\ 1 & 1 & -1 & -1 \end{pmatrix}.$$

Using only the first matrix term gives us the best rank 1 approximation of A , and using only the first two terms gives us the best rank 2 approximation of A :

$$A_{(1)} = \sigma_1 u_1 v_1^* = \begin{pmatrix} 5 & 5 & 5 & 5 \\ 10 & 10 & 10 & 10 \\ 10 & 10 & 10 & 10 \end{pmatrix} \quad A_{(2)} = \sigma_1 u_1 v_1^* + \sigma_2 u_2 v_2^* = \begin{pmatrix} 11 & -1 & 11 & -1 \\ 13 & 7 & 13 & 7 \\ 4 & 16 & 4 & 16 \end{pmatrix}.$$

Since $\text{rank}(A) = 3$ we have $A_{(3)} = A$.

△

Low rank approximation can be viewed as a form of data-compression. Let us consider how effective the compression is. Let A be an $m \times n$ -matrix. Then the matrix itself contains $m \cdot n$ numbers. In order

to store the information about the rank k approximation we need to store the numbers in $\sum_{i=1}^k \sigma_i u_i v_i^*$. Since each u_i is a matrix of form $m \times 1$ and v_i is of form $n \times 1$, we need to store $(m + n + 1)$ numbers³ for each i , so in total $k(m + n + 1)$.

So for example, if A is a 1000×1000 -matrix, and we retain the first 20 singular values when forming the approximation, then we need 10^6 numbers for the original matrix, and $20 \cdot (1000 + 1000 + 1) = 40020$ numbers for the approximation, corresponding to a data compression factor rate of $\frac{10^5}{4002} \approx 25$, so we only need 4% of the data for this approximation. How good the approximation is depends on how large the 20 first singular values are compared to the rest.

Exactly what k to choose of course depends on what the data looks like.

Eckart-Young-Mirsky theorem

The Eckart-Young-Mirsky theorem says that the low rank approximation $A_{(k)}$ defined above in fact is the *best*⁴ approximation both with respect to the operator norm and with respect to the Frobenius norm. It is one of the most used theorems in applications.

Theorem 7.5.3. Eckart-Young-Mirsky theorem

Let A be an $m \times n$ -matrix of rank r with SVD given by $A = \sum_{i=1}^r \sigma_i u_i v_i^*$. For each $k \leq r$ define the corresponding truncated SVD

$$A_{(k)} := \sum_{i=1}^k \sigma_i u_i v_i^*.$$

Then among the matrices of rank $\leq k$, the matrix $A_{(k)}$ is as close as possible to A both with respect to the operator norm and with respect to the Frobenius norm, more precisely:

The minimum of $\{\|A - X\|_{\text{op}} \mid \text{rank}(X) \leq k\}$ is attained for $X = A_{(k)}$.

The minimum of $\{\|A - X\|_F \mid \text{rank}(X) \leq k\}$ is attained for $X = A_{(k)}$.

Proof. For $k = r$ we have $A_{(k)} = A$ and the theorem holds trivially, so assume that $k < r$. We first prove the statement for the operator-norm, recall that $\|A\|_{\text{op}} = \sigma_1$, the largest singular value of A .

We have

$$\|A - A_{(k)}\|_{\text{op}} = \left\| \sum_{i=1}^r \sigma_i u_i v_i^* - \sum_{i=1}^k \sigma_i u_i v_i^* \right\|_{\text{op}} = \left\| \sum_{i=k+1}^r \sigma_i u_i v_i^* \right\|_{\text{op}} = \sigma_{k+1}.$$

So it remains to show that no other approximation of rank k or less can be better. So let B be any $m \times n$ -matrix of rank $\leq k$. Since $\dim \text{span}(v_1, \dots, v_{k+1}) = k+1$ and $\dim \ker(B) \geq n - k$, the sum of dimensions of these two subspaces of \mathbb{C}^n is $> n$, so they intersect non-trivially, and we can pick a vector w of length 1 in the intersection: take

$$w = \sum_{j=1}^{k+1} c_j v_j \in \text{span}(v_1, \dots, v_{k+1}) \cap \ker(B) \quad \text{with} \quad 1 = \|w\| = \sqrt{\sum_{j=1}^{k+1} |c_j|^2}.$$

Then since $\|A - B\|_{\text{op}}$ is the maximum length of a vector $(A - B)x$ when $\|x\| = 1$, in particular we have

$$\begin{aligned} \|A - B\|_{\text{op}} &\geq \|(A - B)w\| = \left\| Aw - \underbrace{Bw}_{=0} \right\| = \|Aw\| = \left\| \sum_{i=1}^r \sigma_i u_i v_i^* \sum_{j=1}^{k+1} c_j v_j \right\| \\ &= \left\| \sum_{i=1}^r \sum_{j=1}^{k+1} c_j \sigma_i u_i \underbrace{(v_i^* v_j)}_{\delta_{ij}} \right\| = \left\| \sum_{j=1}^{k+1} c_j \sigma_j u_j \right\| \geq \sigma_{k+1} \left\| \sum_{j=1}^{k+1} c_j u_j \right\| = \sigma_{k+1} \sqrt{\sum_{j=1}^{k+1} |c_j|^2} = \sigma_{k+1}, \end{aligned}$$

where we used that $w \in \ker(B)$, and that $\{u_i\}$ and $\{v_j\}$ are both orthonormal sets.

³since u_i and v_i are normalized we could possibly get away with storing $m + n - 1$ numbers, but the difference is marginal.

⁴There may in fact be several choices that are equally as good as $A_{(k)}$ when some singular values coincide.

This shows that any matrix of rank $\leq k$ has at least distance σ_{k+1} to A with respect to the operator norm, so since $\|A - A_{(k)}\|_{\text{op}} = \sigma_{k+1}$ we have completed the proof in the operator norm case.

Now let us consider the Frobenius norm case. The error when approximating A by $A_{(k)}$ is:

$$\|A - A_{(k)}\|_F = \left\| \sum_{i=k+1}^r \sigma_i u_i v_i^* \right\|_F = \sqrt{\sigma_{k+1}^2 + \dots + \sigma_r^2},$$

so it remains to show that if B is an arbitrary matrix of rank $\leq k$, then $\|A - B\|_F \geq \sqrt{\sigma_{k+1}^2 + \dots + \sigma_r^2}$.

In general, let us write $\sigma_i(C)$ for the i 'th singular value of any matrix C , we also define $\sigma_i(C) = 0$ for all integers $i > \text{rank}(C)$. By just σ_i we mean $\sigma_i(A)$, the i 'th singular value of our original matrix.

Suppose first that $A = A' + A''$ is an arbitrary expression of A as a sum of two matrices. Then for $i, j \geq 1$ we have:

$$\begin{aligned} \sigma_i(A') + \sigma_j(A'') &\underbrace{=} \sigma_1(A' - A'_{(i-1)}) + \sigma_1(A'' - A''_{(j-1)}) \underbrace{\geq} \sigma_1(A' + A'' - A'_{(i-1)} - A''_{(j-1)}) \\ &\underbrace{=} \sigma_1(A - (A'_{(i-1)} + A''_{(j-1)})) \underbrace{\geq} \sigma_1(A - A_{(i+j-2)}) \underbrace{=} \sigma_{i+j-1}(A). \end{aligned}$$

(1) (2) (3) (4)

In step (1) and (4) we used the fact that when subtracting the components corresponding to the first $(i - 1)$ singular values from a matrix C , the largest remaining singular value is $\sigma_i(C)$. Step (2) follows from the triangle inequality for the operator norm: $\sigma_1(X + Y) \leq \sigma_1(X) + \sigma_1(Y)$. In step (3) we used the fact that

$$\text{rank}(A'_{(i-1)} + A''_{(j-1)}) \leq \text{rank}(A'_{(i-1)}) + \text{rank}(A''_{(j-1)}) = (i - 1) + (j - 1) = \text{rank}(A_{(i+j-2)}),$$

so (3) follows from the fact that $A_{(i+j-2)}$ is the *best* rank $i + j - 2$ approximation for A with respect to the operator norm as we showed in the first part of this proof.

So we have shown that

$$\sigma_i(A') + \sigma_j(A'') \geq \sigma_{i+j-1} \quad \text{whenever} \quad A = A' + A''.$$

Now let B be an arbitrary $m \times n$ -matrix of rank $\leq k$. Take $A' = A - B$ and $A'' = B$ and $j = k + 1$ in our inequality above. Then for all $i \geq 1$ we obtain

$$\sigma_i(A - B) = \sigma_i(A - B) + \underbrace{\sigma_{k+1}(B)}_{=0} \geq \sigma_{i+(k+1)-1}(A) = \sigma_{k+i}(A),$$

where $\sigma_{k+1}(B) = 0$ since we assumed that B has rank $\leq k$. But then

$$\|A - B\|_F = \sqrt{\sum_{i=1}^r (\sigma_i(A - B))^2} \geq \sqrt{\sum_{i=1}^r (\sigma_{k+i}(A))^2} = \sqrt{\sum_{i=k+1}^r \sigma_i^2} = \|A - A_{(k)}\|_F.$$

This proves that any approximation B of rank $\leq k$ of A is equal or worse than the approximation $A_{(k)}$, which proves the theorem in the Frobenius norm case. □

We note that if $\sigma_k > \sigma_{k+1}$, the best approximation of rank k is unique, but if $\sigma_k = \sigma_{k+1}$ there will be several matrices with the same minimal distance to A since the matrix $A_{(k)}$ actually depends on how we order the k 'th and $(k + 1)$ 'th columns in U and in V in the SVD.

Image compression

Let us look at an example where we compress an image using its low rank approximations. A grayscale image can be represented as an $m \times n$ -matrix where the image is m pixels high and n pixels wide. Each pixel brightness can be encoded as a real number, a simple standard is to use integers between 0 and 255 where 0 means black and 255 means white.


Let us use the following image of a macaw for demonstration, the image is of size 540×900 . We start by converting it to grayscale and call that matrix A , the rank of A is 540.



We compute the SVD of A :

$$A = U\Sigma V^* = \sum_{i=1}^{540} \sigma_i u_i v_i^*.$$

From this we can construct the low rank approximations of rank 5, 20, and 100 by just using the first 5, 20, and 100 terms of this sum respectively. This produces:



Rank 5
Compression 1.5%



Rank 20
Compression 5.9%




Rank 100
Compression 29.7%

The compression percentages measure how much data needs to be stored compared to the original black and white image, for example, for the 100-rank matrix we need to store $100 \cdot (540 + 900 + 1) = 144100$ numbers while the original gray image needs $540 \cdot 900 = 539460$.

The colored image can be approximated with the same method, since each color can be represented as a triple (r, g, b) , write A_r, A_g, A_b for the three matrices that keeps track of the red, blue, and green

pixel-components respectively. Then we perform the low rank approximation on each color channel separately and merge the resulting images. For rank 100 this looks like:



Red channel

Green channel


Blue channel

Merged

In applications, finding the full SVD and then truncating it may be time-consuming. However, there are efficient algorithms that find just the first k singular values, and the corresponding first k columns in the matrices U and V in the SVD.

Let us plot the 20 first singular values of the matrix A , the grayscale version of the macaw:

Decay of Singular Values



This graph is typical for what the magnitudes of the singular values tend to look like in collected data. The first singular value is very large, the following few singular values are relatively large, and then they approach zero quite quickly. Looking at such a graph may help us decide how to pick the cut-off point k in choosing an approximation $A_{(k)}$.

7.6 Moore-Penrose pseudo inverse

In the context of least squares, when A had *linearly independent columns* we defined $A^+ = (A^*A)^{-1}A^*$, and showed that this matrix in a sense was the best possible generalization of an inverse when A was not square, see Definition 5.9.3 and the following results. In this section we will generalize this to matrices of arbitrary size using the SVD.

Definition 7.6.1. Let A be an $m \times n$ -matrix with compact SVD given by $A = \tilde{U}\tilde{\Sigma}\tilde{V}^*$. We define the **Moore-Penrose pseudo inverse** of A to be the $n \times m$ -matrix

$$A^+ := \tilde{V}\tilde{\Sigma}^{-1}\tilde{U}^*.$$

Note that if $\tilde{\Sigma}$ is the $r \times r$ -matrix $\tilde{\Sigma} = \begin{pmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{pmatrix}$, then $\tilde{\Sigma}^{-1} = \begin{pmatrix} \frac{1}{\sigma_1} & & \\ & \ddots & \\ & & \frac{1}{\sigma_r} \end{pmatrix}$. We also remark that

$A^+ = \tilde{U}\tilde{\Sigma}^{-1}\tilde{V}^*$ is not an SVD-expression of A^+ since the diagonal entries of $\tilde{\Sigma}^{-1}$ are in fact increasing, not decreasing. If we need to find the SVD of A^+ we should reverse the ordering of the diagonal elements of $\tilde{\Sigma}^{-1}$, and correspondingly we should also reverse the order of the columns in \tilde{U} and in \tilde{V} . When A has linearly independent columns this coincides with our previous definition of A^+ .

There is an alternative characterization of A^+ in terms of the so called *Moore-Penrose conditions*:

Proposition 7.6.2. *The Moore-Penrose pseudo inverse A^+ satisfies:*

1. A^+A and AA^+ are both Hermitian.
2. $AA^+A = A$.
3. $A^+AA^+ = A^+$.

These are called the Moore-Penrose conditions.

The proof is left as an exercise. Indeed, one can prove that there is a unique matrix satisfying the Moore Penrose conditions, so this is an alternative definition of A^+ .

Example 7.6.3. Let us find the Moore-Penrose inverse of the matrix $A = \frac{1}{6} \begin{pmatrix} 4 & 5 & 2 \\ 0 & 3 & 6 \\ 4 & 5 & 2 \\ 0 & 3 & 6 \end{pmatrix}$. The columns of

A are linearly dependent, so our previous method of finding $(A^*A)^{-1}A^*$ does not work since A^*A does not have an inverse. In Example 7.2.8 we found that a compact SVD of A is given by

$$A = \tilde{U}\tilde{\Sigma}\tilde{V}^* \quad \text{where} \quad \tilde{U} = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \\ 1 & -1 \end{pmatrix} \quad \tilde{\Sigma} = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} \quad \tilde{V} = \frac{1}{3} \begin{pmatrix} 1 & 2 \\ 2 & 1 \\ 2 & -2 \end{pmatrix}.$$

Here we have $\tilde{\Sigma}^{-1} = \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 1 \end{pmatrix}$, so we get

$$A^+ = \tilde{V}\tilde{\Sigma}^{-1}\tilde{U}^* = \frac{1}{3} \begin{pmatrix} 1 & 2 \\ 2 & 1 \\ 2 & -2 \end{pmatrix} \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 1 \end{pmatrix} \frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \end{pmatrix} = \frac{1}{12} \begin{pmatrix} 5 & -3 & 5 & -3 \\ 4 & 0 & 4 & 0 \\ -2 & 6 & -2 & 6 \end{pmatrix}.$$

Now neither A^+A nor AA^+ is the identity, because the rank of A is 2 while the two matrices A^+A nor AA^+ are of size 3×3 and 5×5 respectively. Nevertheless, we can think of A^+ as being as close as possible to being an inverse of A .

△

The square matrices A^+A and AA^+ are in a sense *almost* the identity matrices, here is how we can think of these maps geometrically:

Proposition 7.6.4. *Geometrically,*

- AA^+ is the orthogonal projection onto $\text{Im}(A)$, the space spanned by the columns of A .
- A^+A is the orthogonal projection onto $\text{Im}(A^*)$, the space spanned by the conjugates of rows of A .

Here is how the matrix A^+ relates to least square solutions:

Proposition 7.6.5. *Given an arbitrary linear system $Ax = b$, the vector A^+b is the shortest least square solution to $Ax = b$, in other words, $y = A^+b$ is the vector of smallest length that satisfies $A^*Ay = A^*b$.*

We leave the proofs of these statements as exercises.

7.7 Condition number

Consider a linear system of equations $Ax = b$ where A is an invertible square $n \times n$ -matrix. It has unique solution $x = A^{-1}b$.

In practical applications, this system may come from noisy data, so suppose that we add a small measurement or rounding error Δb in the right side b . To what extent does this change the solution to the system?

Let $x + \Delta x$ be the solution to the modified system:

$$A(x + \Delta x) = b + \Delta b.$$

Since $Ax = b$ this gives $A(\Delta x) = \Delta b$ and $\Delta x = A^{-1}(\Delta b)$. Hopefully the error Δx will be small compared to the size of x , we compute the *relative error*:

$$\frac{\|\Delta x\|}{\|x\|} = \frac{\|A^{-1}(\Delta b)\|}{\|x\|} = \frac{\|A^{-1}(\Delta b)\|}{\|b\|} \cdot \frac{\|Ax\|}{\|x\|} \leq \frac{\|A^{-1}\|_{\text{op}} \cdot \|\Delta b\|}{\|b\|} \cdot \frac{\|A\|_{\text{op}} \cdot \|x\|}{\|x\|} = \frac{\sigma_1}{\sigma_n} \frac{\|\Delta b\|}{\|b\|},$$

where we used that $\|A\|_{\text{op}} = \sigma_1$ and $\|A^{-1}\|_{\text{op}} = \frac{1}{\sigma_n}$, where $\sigma_1, \dots, \sigma_n$ are the singular values of A .

This says that when A is fixed, the relative error in the solution is bounded by a constant times the relative error in the right side b . That constant is called the *condition number* of the matrix.

Definition 7.7.1. Let A be an invertible $n \times n$ -matrix with singular values $\sigma_1 \geq \dots \geq \sigma_n > 0$. We define the **condition number** of A to be the quotient of the largest by the smallest singular value:

$$\kappa(A) = \frac{\sigma_1}{\sigma_n}.$$

So we conclude that $\frac{\|\Delta x\|}{\|x\|} \leq \kappa(A) \frac{\|\Delta b\|}{\|b\|}$. Perhaps the matrix A comes from data we collected, and we have some freedom in choosing this matrix. In that case we should try and choose A so that its condition number is as small as possible. Matrices with small condition number are sometimes called *well-conditioned*, and matrices with large condition numbers are sometimes called *ill-conditioned*, although exactly how small condition number is needed depends on the context.

7.8 Polar decomposition

Recall that the polar form of a complex number is

$$z = e^{i\theta} r$$

where $r \geq 0$ is a non-negative real number, and where $e^{i\theta}$ lies on the unit circle in the complex plane. If we think of z as an 1×1 -matrix $[z]$, it corresponds to the a linear map $\mathbb{C} \rightarrow \mathbb{C}$ which multiplies complex numbers by z . The factorization $[z] = [e^{i\theta}][r]$ expresses this map as the product of a rotation by an angle θ , and a scaling by a factor r . It turns out that this concept generalizes well to arbitrary square complex matrices:

Definition 7.8.1. A **polar decomposition** of a square complex matrix A is a factorization

$$A = UP$$

where U is unitary, and P is positive semi-definite.

With the Polar decomposition we can think of an arbitrary square matrix (or a map $V \rightarrow V$) as a composition of a rotation U and stretching P along a number of orthogonal axes.

A polar decomposition can be found via the SVD of A : we have a polar decomposition

$$A = U\Sigma V^* = (UV^*)(V\Sigma V^*) = U'P,$$

where $U' := UV^*$ is unitary since U and V is, and where $P = V\Sigma V^*$ is positive semi-definite since Σ is.

In general, if $A = UP$, we get $A^*A = P^*U^*UP = P^*P = P^2$ where we used that a positive semi-definite matrix is always Hermitian. So $P^2 = A^*A$ and therefore $P = \sqrt{A^*A}$ since P is positive semi-definite. If zero is not a singular value of A , then P is invertible, and we get a unique $U = AP^{-1}$ as well.

We have proven:

Proposition 7.8.2. *Every square complex matrix A has a polar decomposition $A = UP$. The matrix P is unique and is given by $P = \sqrt{A^*A}$. When all the singular values of A are positive (meaning that A is invertible), the matrix U is also unique, and it can then be found as $U = AP^{-1}$.*

In analogy with the complex numbers it is common to define $|A| = \sqrt{A^*A}$, and think of this as the "stretching part" of the matrix A , the polar decomposition then reads $A = U|A|$.

Example 7.8.3. Let us find the polar decomposition of the matrix $A = \frac{1}{\sqrt{5}} \begin{pmatrix} 1 & 5 \\ 7 & -5 \end{pmatrix}$.

We diagonalize $A^*A = \begin{pmatrix} 10 & -6 \\ -6 & 10 \end{pmatrix}$, we see that $v_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix}$ is an eigenvector for the eigenvalue 16, and that $v_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ is an eigenvector for the eigenvalue 4, so with $V = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}$ we have

$$A^*A = V \begin{pmatrix} 16 & 0 \\ 0 & 4 \end{pmatrix} V^* \quad \text{so} \quad P = \sqrt{A^*A} = V \sqrt{\begin{pmatrix} 16 & 0 \\ 0 & 4 \end{pmatrix}} V^* = V \begin{pmatrix} 4 & 0 \\ 0 & 2 \end{pmatrix} V^* = \begin{pmatrix} 3 & -1 \\ -1 & 3 \end{pmatrix},$$

and we see that P is indeed positive definite (it is Hermitian and satisfies Sylvester's criterion). Since P is invertible we get $U = AP^{-1} = \frac{1}{\sqrt{5}} \begin{pmatrix} 1 & 2 \\ 2 & -1 \end{pmatrix}$, which is indeed unitary. So the polar decomposition of A is

$$A = UP = \frac{1}{\sqrt{5}} \begin{pmatrix} 1 & 2 \\ 2 & -1 \end{pmatrix} \begin{pmatrix} 3 & -1 \\ -1 & 3 \end{pmatrix}$$

△

7.9 Data analysis with SVD

In this section we shall work only with matrices whose entries are *real*⁵.

Suppose that we have collected some data in an $m \times n$ -matrix A . We shall think of the n columns of A as corresponding to different variables or measurements, and we think of the m rows as samples or instances of some experiment. For example, the rows could correspond to different people, and the columns could correspond to measurements such as height, weight, age, etcetera. Typically m is a lot larger than n . In the context of data-analysis we shall call a matrix of such form a **data-matrix**.

We are interested in comparing how *similar* two samples or two variables are:

Definition 7.9.1. The **cosine-similarity** of two vectors u and v is defined as


$$\text{cosSim}(u, v) = \frac{u \bullet v}{\|u\| \cdot \|v\|}.$$

So the cosine similarity is just the cosine of the angle between the vectors u and v (even if these may sit in some large-dimensional space). By Cauchy-Schwarz, this number is always between -1 and 1 . When

⁵With some modification, the techniques in this section applies to complex data too, but this complicates our exposition so we skip it here.

the cosine similarity is close to 1 it means that u and v are highly correlated in the sense that their entries u_i and v_i tend to have the same sign, and u_i is large when v_i is large. Negative cosine similarity means that u_i and v_i tends to have opposite signs. Cosine-similarity close to zero means that u and v are uncorrelated in the sense that the size of the entries u_i can not be used to predict the size of the entries v_i .

Example 7.9.2. Let's say we have measured age and weight of five people. We collect the data in a matrix A where the first column is weight and the second is age, and where the rows correspond to five people (Ada, Bob, Cam, Deb, Eli). We plot the data in a diagram.



Here the cosine-similarity between the two columns A_1 and A_2 of A is

$$\text{cosSim}(A_1, A_2) = \frac{(66, 83, 51, 70, 79) \bullet (31, 40, 18, 65, 59)}{\|(66, 83, 51, 70, 79)\| \cdot \|(31, 40, 18, 65, 59)\|} \approx 0.955$$

which tells us that higher age tends to imply higher weight.

We can also compare two individuals:

$$\text{cosSim}(\text{Bob}, \text{Deb}) = \frac{(83, 40) \bullet (70, 65)}{\|(83, 40)\| \cdot \|(70, 65)\|} \approx 0.953.$$

Note that the cosine-similarity only compares the "direction" or angle of the data-points, not the relative size. For example, Bob is closer to Cam than to Deb with respect to cosine-similarity, but with respect to the euclidean distance Bob is closer to Deb than to Cam.

△

In general, if A is an $m \times n$ -matrix, we should think of a similar picture, but we should try and visualize m points in some higher dimensional space \mathbb{R}^n (think of \mathbb{R}^3 and imagine their are even more directions).

Now our goal is to find some new orthogonal coordinate system in \mathbb{R}^n which *best explains the variation* in the data. First, let us pick a vector v of length 1 in \mathbb{R}^n , and consider the axis-line through the origin in the direction of v .

Write x_i for the i 'th row of A , such that each x_i corresponds to one of the points in \mathbb{R}^n . Since $\|v\| = 1$, the projection of x_i on the axis through v is $P_v(x_i) = (x_i \bullet v)v$ and the squared length of this projection is simply $\|P_v(x_i)\|^2 = (x_i \bullet v)^2$, and therefore by the Pythagorean theorem, the squared distance from x_i to the axis through v is $\|x_i - P_v(x_i)\|^2 = \|x_i\|^2 - (x_i \bullet v)^2$. So the sum of the squared distance of all the data-points to the axis through v is

$$d_v = \sum_{i=1}^m \|x_i\|^2 - (x_i \bullet v)^2.$$

Since $\sum_{i=1}^m \|x_i\|^2$ is constant, d_v is minimized when $\sum_{i=1}^m (x_i \bullet v)^2$ is maximized, this latter quantity is the sum of the squared lengths of the projections of all data-points onto the v -axis.

But now we note that this quantity is exactly $\|Av\|^2$:

$$\|Av\|^2 = \left\| \begin{pmatrix} x_1 \\ \vdots \\ x_m \end{pmatrix} v \right\|^2 = \left\| \begin{pmatrix} x_1 \bullet v \\ \vdots \\ x_m \bullet v \end{pmatrix} \right\|^2 = \sum_{i=1}^m (x_i \bullet v)^2.$$

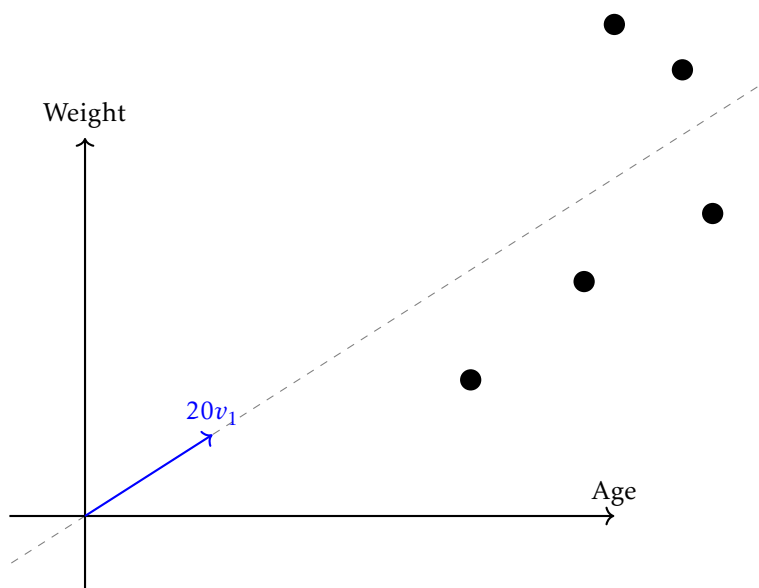
So to minimize the distance from the v -axis to the data, we should pick the unit vector v such that $\|Av\|$ is maximized. But by the properties of the operator norm we know that

$$\|Av\| \leq \|A\|_{\text{op}} \cdot \|v\| = \sigma_1 \|v\| = \sigma_1,$$

with equality if and only if v is parallel to the first right singular vector of A . We conclude:

Proposition 7.9.3. *Let A be an $m \times n$ data-matrix and let v be a unit vector in \mathbb{R}^n , and consider the axis through the origin in the v -direction. The sum of the squared distances of the data-points to this axis is minimized when v is a right singular vector for the largest singular value of A . So in order to pick an axis which minimizes the squared distance to the data, the axis direction v should be chosen as an eigenvector of $A^T A$ corresponding to its largest eigenvalue.*

Example 7.9.4. Following up on our weight-age example above we can now plot the direction of the first right singular vector $v_1 = (0.843, 0.538)$, but let's rescale this vector so that it is visible in the diagram). We see that it indeed seems to minimize the (squared) distances to the data.



△

Now suppose that we have a singular value decomposition of A given by⁶ $A = U\Sigma V^T$.

In general, given our first axis-direction v_1 of \mathbb{R}^n along which the data varies the most, we can project all the data points onto the orthogonal complement $\text{span}(v_1)^\perp$, this is an $n - 1$ -dimensional hyperplane, and we can use the same method to find the axis along which the projected data varies the most.

The projection of the data-point x_i (still viewed as a row matrix) onto $\text{span}(v_1)^\perp$ is $x_i - (x_i \bullet v_1)v_1^T$, doing this with all rows x_i of A simultaneously gives the new (projected) data-matrix

$$A - (Av_1)v_1^T = A - \sigma_1 u_1 v_1^T.$$

⁶Note that $V^* = V^T$ since all matrices have real entries in this section. We could also use the compact version of the SVD here.

Here we have subtracted the rank 1-approximation $A_{(1)}$ from A , so the largest remaining singular value is σ_2 , and the projected data varies most in the direction v_2 .

Continuing this, we see that the right singular vectors of A are precisely the directions along which the data varies the most.

So the right singular vectors v_1, \dots, v_n corresponding to the singular values values of A in decreasing order is the canonical choice for a basis of our new coordinate system. These vectors are the eigenvectors of $A^T A$ in order by decreasing eigenvalues, so in fact we do not find the full SVD to find this coordinate system, the eigenvalues and eigenvectors of $A^T A$ suffice.

The new basis vectors v_i still span the space \mathbb{R}^n in which the data-points sit, so we can think of each v_i as a linear combination of the n variables. Any of the data-points x (a row of A) can be expressed in this basis as

$$x = (x \bullet v_1)v_1 + \dots + (x \bullet v_n)v_n,$$

so $x \bullet v_i$ is the coordinate of x along the v_i -axis, so these inner products tell us how well each data-point aligns with that direction.

On the other hand, let us consider the left singular vectors u_i from the SVD. These form an ON-basis for \mathbb{R}^m in which the columns of A sit (each columns corresponding to all measurements of a fixed variable), and u_i is parallel to Av_i . So we can think of the vectors u_i as as linear combinations of the different variable-measurements. Any column y of A can be expressed in the basis u_i as

$$y = (y \bullet u_1)u_1 + \dots + (y \bullet u_n)u_n,$$

so $y \bullet u_i$ is the coordinate of y along the u_i -axis, so these inner products tell us how well each variable aligns with that direction.

Here the idea is that we can perform a **dimensionality reduction** of the data. If we just keep the first k major directions along which the data varies we get a good approximation of the data. This corresponds to replacing the data-matrix A by its low rank approximation $A_{(k)}$. We summarize our conclusions:

Proposition 7.9.5. *Let A be an $m \times n$ data-matrix of rank r , and let $A = U\Sigma V^T$ be a singular decomposition of A . Then the right singular vectors $v_1, \dots, v_r \in \mathbb{R}^n$ form an ON-basis for the row-space of A , and for each $1 \leq k \leq r$, the span $V_k = \text{span}(v_1, \dots, v_k)$ is the k -dimensional subspace of \mathbb{R}^n which lies closest to the data-points in the sense that the sum of the square distances from the data-points to V_k is minimal.*

Dually, the vectors $u_1, \dots, u_r \in \mathbb{R}^m$ is an ON-basis for the column-space of A , and for each $1 \leq k \leq r$, the span $U_k = \text{span}(u_1, \dots, u_k)$ is a k -dimensional subspace of \mathbb{R}^m which lies closest to the n variable-vectors in the sense that the sum of the square distances from the columns of A to U_k is minimal.

Here we can think of each singular value as detecting some "feature" of the data. Each right singular vector v is some linear combination of the variables and corresponds to such a feature, the corresponding left singular vector tells us how each data point aligns with this feature, and the size of the singular value determines how much of the variance in the data is explained by that feature. To understand this better we should look at a concrete example.

Movie ratings example

Suppose that we have a matrix A of movie ratings, where users have rated a set of movies by 1-5 stars.

		Interstellar				
		Mean Girls				
		The Matrix				
		Borat				
		Ghost Busters				
Alex	(5	2	4	2	4
Bert		5	1	4	1	4
Cleo		4	5	2	4	5
Dany		3	5	3	5	5
Elle		5	2	5	2	3
Fred		5	3	4	2	4
)					

The singular value decomposition, rounded to two decimals of the movie matrix is

$$A = U\Sigma V^T$$

$$U = \begin{pmatrix} 0.39 & 0.26 & 0.12 & 0.17 & -0.46 \\ 0.36 & 0.45 & 0.41 & 0.52 & 0.15 \\ 0.44 & -0.49 & 0.48 & -0.38 & -0.35 \\ 0.45 & -0.57 & -0.42 & 0.49 & 0.22 \\ 0.39 & 0.37 & -0.64 & -0.30 & -0.29 \\ 0.41 & 0.16 & 0.09 & -0.46 & 0.71 \end{pmatrix} \quad \Sigma = \begin{pmatrix} 20.1 & & & & \\ & 5.60 & & & \\ & & 1.57 & & \\ & & & 0.96 & \\ & & & & 0.35 \end{pmatrix} \quad V = \begin{pmatrix} 0.54 & 0.45 & 0.33 & -0.42 & -0.47 \\ 0.38 & -0.56 & -0.05 & -0.61 & 0.41 \\ 0.44 & 0.47 & -0.67 & 0.14 & 0.36 \\ 0.34 & -0.50 & -0.41 & 0.28 & -0.63 \\ 0.51 & -0.13 & 0.53 & 0.60 & 0.30 \end{pmatrix}$$

Let us try and interpret this data. The entries of Σ are the singular values of A , and each singular value corresponds to some *latent feature*, and the sizes of the singular values indicate the importance of that feature in explaining the variance in the data. In this case, the first singular value $\sigma_1 = 20.1$ will correspond to the feature *average movie popularity*, while $\sigma_2 = 5.60$ will correspond to *movie genre* on a scale between sci-fi and comedy⁷.

The columns v_i of V form an ON-basis for the 5-dimensional *movie-space*, we have one standard basis vector for each movie, and each users movie preferences correspond to a vector in this space. We can also think of each of the new basis vectors v_i as a formal linear combination of the five movies. The first right singular vector

$$v_1 = (0.54, 0.38, 0.44, 0.34, 0.51)^T$$

will explain the greatest variance in the movie ratings data, and its entries will correspond to the average ratings of the movies. So from v_1 we see that the most popular movie on average was *Interstellar* closely followed by *Ghost Busters*. The second right singular vector

$$v_2 = (0.45, -0.56, 0.47, -0.50, -0.13)^T$$

will explain the greatest variance in the ratings data *which is not explained by the average popularity*. In this context, the entries of v_2 will correspond to the movie genre, on a scale ranging from sci-fi to comedy. We see that *Interstellar* and *The Matrix* have high positive coefficients indicating sci-fi, while *Mean Girls* and *Borat* have high-negative scores indicating comedies, while *Ghost Busters* has a score close to zero, indicating that it might be a sci-fi/comedy mix. The remaining columns of V correspond to other latent feature that are not explained by the two first factors, they are a lot harder to understand intuitively.

Let's move on to the columns of U , these form an ON-basis for the six-dimensional *user-space* in which we have one standard basis vector for each user, so each vector u_i can be viewed as a formal linear combination of the users (movie watchers). The set of user ratings for each fixed movie is also a point in this space. The vectors u_i will tell us how important each latent feature is to explaining that users movie-ratings. So

$$u_1 = (0.39, 0.36, 0.44, 0.45, 0.39, 0.41)^T$$

tells us the average movie ratings for the users, in particular, the fourth user Dany gave the highest average ratings. Some users may be predisposed to giving higher ratings on average, but if we subtract all such effects, the greatest factor explaining user ratings is explained by the second left singular vector

$$u_2 = (0.26, 0.45, -0.49, -0.57, 0.37, 0.16)^T.$$

The entries of u_2 correspond to user preferences when it comes to the second latent feature, the movie genre. The two negative values in the middle tells us that Cleo and Dany strongly prefer comedies over science fiction, while the others prefer sci-fi (with Fred being relatively neutral when it comes to genres).

The two first singular values are quite a bit larger than the others. This tells us that most of the variance in the ratings data can be explained by the two first latent features, and we shouldn't lose too

⁷In general it is very hard to deduce exactly what properties a certain latent feature corresponds to.

much information if we replace A by its rank 2 approximation $A_{(2)}$. So we retain only the first two columns of U and of V , and we retain only the two first singular values. With


$$\hat{U} = \begin{pmatrix} 0.39 & 0.26 \\ 0.36 & 0.45 \\ 0.44 & -0.49 \\ 0.45 & -0.57 \\ 0.39 & 0.37 \\ 0.41 & 0.16 \end{pmatrix} \quad \hat{\Sigma} = \begin{pmatrix} 20.1 & \\ & 5.60 \end{pmatrix} \quad \hat{V} = \begin{pmatrix} 0.54 & 0.45 \\ 0.38 & -0.56 \\ 0.44 & 0.47 \\ 0.34 & -0.50 \\ 0.51 & -0.13 \end{pmatrix}$$

we get

$$\hat{U}\hat{\Sigma}\hat{V}^* = \begin{pmatrix} 4.93 & 2.18 & 4.16 & 1.93 & 3.85 \\ 5.02 & 1.31 & 4.34 & 1.16 & 3.35 \\ 3.54 & 4.86 & 2.59 & 4.33 & 4.86 \\ 3.45 & 5.22 & 2.47 & 4.65 & 5.04 \\ 5.16 & 1.82 & 4.41 & 1.60 & 3.74 \\ 4.89 & 2.63 & 4.06 & 2.34 & 4.12 \end{pmatrix} \approx \begin{pmatrix} 5 & 2 & 4 & 2 & 4 \\ 5 & 1 & 4 & 1 & 4 \\ 4 & 5 & 2 & 4 & 5 \\ 3 & 5 & 3 & 5 & 5 \\ 5 & 2 & 5 & 2 & 3 \\ 5 & 3 & 4 & 2 & 4 \end{pmatrix} = A$$

This is the matrix of rank 2 which best approximates the original data, it only takes overall popularity and genre-preference in account, and we see that it is quite a good approximation to A .


Replacing data by a lower rank approximation can eliminate noise and simplify further data analysis. It is also useful for data-visualization purposes, in this example the approximation above lets us visualize the data in a two-dimensional space. Retaining only the two first right singular vectors and projecting the users onto the two-dimensional space $\text{span}(v_1, v_2)$ we can visualize this projected data in a 2d-diagram:



So in this diagram, the coordinates of each point is the length of its projection onto v_1 and on v_2 respectively.

We can clearly identify two clusters, the people who prefer sci-fi and the people who prefer comedy.

We can make a similar diagram in movie-space, retaining only the two first left singular vectors and projecting the movies onto the two-dimensional space $\text{span}(u_1, u_2)$ we can visualize this projected data in a two-dimensional diagram:



In this diagram, the coordinates of each point is the length of its projection onto u_1 and on u_2 respectively.

Here we see how similar movies cluster together corresponding to the sci-fi- and comedy-generes, with one movie in the middle between the two clusters.

Suppose now that a seventh user Gina absolutely loved Interstellar and hated Borat, but haven't watched any of the other movies. We can represent her movie ratings by a vector $v = (5, 3, 3, 1, 3)$ in movie-space (where we put an average rating of 3 for unseen movies). What movies should we recommend to Gina? To estimate Gina's preferences we project her (incomplete) rating vector onto $\text{span}(v_1, v_2)$ and obtain $v = (7.02, 1.06) = 7.02v_1 + 1.06v_2$, where the coefficients tells us how well Gina aligns with the average popularity and with the genre-features. So if we plot her in the user-diagram above we see that she is close to Alex, Burt, Elle, and Fred and will probably prefer the same movies as these users. If we move into movie-space, then corresponding combination of left singular vectors is $(7.02, 1.06) = 7.02u_1 + 1.06u_2$ which we could plot in the movie-diagram above. This point lies very close to Interstellar and The Matrix with very high cosine-similarity to both, and quite far from the other movies, so we should probably recommend Gina to watch The Matrix which she didn't see yet.

7.10 Principal component analysis

Principal component analysis can be thought of a rebranding of the SVD but in the context of a normalized data-matrix.

Say that we have any data in an $m \times n$ -matrix B with each of the m rows corresponding to a sample or experiment, and each column corresponds to some variable or measurement. Let z_i be the i 'th row of B . The *mean* of the data-matrix is the average of the rows:

$$c = \frac{1}{m} \sum_{i=1}^m z_i.$$

We center our data around the origin by defining $x_i := z_i - c$, and we let A be the normalized data-matrix with rows x_i . Then the sum in each column of A is zero.

Definition 7.10.1. The **covariance matrix** of the normalized data is

$$\text{cov}(A) = \frac{1}{m-1} A^T A$$

its element at position (i, j) measure the correlation between variable i and j . The directions of the eigenvectors of the covariance matrix are called the **principal axes** or **principal components** of the data-matrix A ,


By replacing our data-points by their projections onto the first k principal components, we perform what is called a **dimensionality reduction**.

The normalization factor $\frac{1}{m-1}$ makes the covariance independent of the number of measurements, but note that the covariance matrix still has the same eigenvectors as $A^T A$, and in some text $A^T A$ is even called the covariance matrix.

The principal components correspond to eigenvectors of the covariance matrix, are these eigenvectors are always sorted in order of decreasing eigenvalues. The principal components are precisely the same as the right singular vectors of the normalized data-matrix A , and they describe the direction in which the data varies the most.

Dimensionality reduction corresponds to finding the subspace of dimension k closest to the data, which is equivalent to saying that the projection of the data onto this subspace is as "spread out" as possible. This essentially reduces the number of variables in our data, which simplifies analysis and visualization.

Example 7.10.2. Suppose that we have some two dimensional data consisting of the 10 points (x, y) below. We have already normalized our data so that the sum of all x -coordinates and the sum of all y -coordinates is zero. The data can be plotted as points in the plane:



Now the covariance matrix is $\frac{1}{10-1} A^T A = \frac{1}{9} \begin{pmatrix} 106 & 41 \\ 41 & 38 \end{pmatrix}$.


To find the principal components we compute the (compact) SVD of A , we get:

$$\tilde{U} = \begin{pmatrix} -0.19 & -0.53 \\ 0.20 & 0.01 \\ 0.52 & 0.14 \\ -0.24 & 0.29 \\ 0.36 & -0.18 \\ -0.40 & 0.49 \\ 0.32 & 0.12 \\ -0.04 & 0.31 \\ -0.08 & -0.42 \\ -0.44 & -0.23 \end{pmatrix} \quad \tilde{\Sigma} = \begin{pmatrix} 11.19 & 0.00 \\ 0.00 & 4.33 \end{pmatrix} \quad \tilde{V} = \begin{pmatrix} 0.91 & -0.43 \\ 0.43 & 0.91 \end{pmatrix}$$

Thus the principal axes are spanned by the two right singular vectors v_1 and v_2 , the columns of \tilde{V} . Since the covariance matrix is $\frac{1}{9}A^T A = \tilde{V} \left(\frac{1}{\sqrt{9}} \Sigma \right)^2 \tilde{V}^T$, the square roots of the eigenvalues of the covariance matrix corresponding to the principal components are $\frac{11.19}{3} \approx 3.73$ and $\frac{4.33}{3} \approx 1.44$.

We can think of PCA as fitting an ellipse (or hyper-ellipsoid in general) to our data, the semi-axes of the ellipse are the eigenvectors of the covariance matrix and the corresponding square roots of the eigenvalues are the sizes of the axes.

We plot the right singular vectors $v_1 = (0.91, 0.43)$ and $v_2 = (-0.43, 0.91)$, these form an ON-basis in the two-dimensional space where the data points sit. Then the first vector v_1 will describe the direction in which the data varies most, this just means the line in the v_1 direction lies closest to the data points. If we project the data points onto the space orthogonal to v_1 , the second vector v_2 describes the direction orthogonal to v_1 in which the projected data varies most since our data was two-dimensional in this example, there are no other directions the data can vary. The singular values represent how great the data variance is in the respective directions.



The size of the ellipse corresponds to 1 standard deviation of the data. If the points were collected from a normal distribution, we would expect this ellipse to cover around 40% of the points.


△

Eigenfaces example

Let us go through an example where we use SVD to analyze pictures of faces.


In this example, 70 pictures of employees at the mathematics department at Linköping university

is used as a data-set. Each picture is aligned using a computer algorithm such that they have the same relative size and such that the eyes, noses and mouths aligned. The pictures are then converted to grayscale 150×150 images, so that each picture is a 150×150 -matrix with integer entries on the interval $[0, 255]$. These matrices are then vectorized row by row so that each picture becomes a vector $z_i \in \mathbb{R}^{150 \cdot 150} = \mathbb{R}^{22500}$. We then compute the mean face:

$$c = \frac{1}{70} \sum_{i=1}^{70} z_i =$$


So this is how we look at MAI on average⁸.

Now we subtract the average face from each face: $x_i := z_i - c$, and we put these new faces x_i as rows in a large matrix A , this matrix has size 70×22500 .



Original → Aligned → Resized and decolorized → Mean subtracted


We use a computer algorithm find a compact SVD⁹ of A . Since $\text{rank}(A) = 70$ we get $A = \tilde{U} \tilde{\Sigma} \tilde{V}^T$ where \tilde{U} and $\tilde{\Sigma}$ are both 70×70 , and \tilde{V} is of size 22500×70 .

Here the columns v_i of \tilde{V} are called **eigenfaces**. The eigenfaces form an orthonormal basis for the 70-dimensional face-space, the subspace of \mathbb{R}^{22500} spanned by the faces, but note that (luckily) each vector v_i is not a face of someone at MAI, but rather a linear combination of such faces. Moreover, v_1 explains the largest in the variation of the faces, while v_2 explains the most variation in the directions orthogonal to the v_1 -axis and so on.

⁸Here the contrast was increased to make the face in the image more visible. Indeed, some face-vectors in this example will even have negative elements, so all images in this section are transformed such that their lowest pixel value is 0 and their highest is 255.

⁹In fact, we don't really need the matrices \tilde{U} and $\tilde{\Sigma}$ for the PCA, so it would suffice to diagonalize $A^T A = V D V^T$.

Here are the first 18 right singular vectors v_1, \dots, v_{18} , the directions of the principal component axes:



We can think of these as the best possible faces for representing our data-set, any face at MAI is can be quite well estimated as a liner combination of these first 18 eigenfaces.


Since the v_i form an ON-basis in face-space, any normalized face x in our dataset has a unique expression

$$x = (x \bullet v_1)v_1 + \dots + (x \bullet v_{70})v_{70},$$

and by projecting only onto the first k principal components we get an approximation of the face.


In fact we can project any image (such as a face not belonging to our data-set) onto the principal components to find the linear combination of faces from MAI that best approximates the given image.

Here are the projections of a normalized face x in our dataset and its projection onto the first k components for different values of k :



Just as we did in the movie-example, let us plot each of the 70 faces from the data-set onto the first

two principal components, this gives us a representation of how varied our faces are at MAI:



Further applications of data analysis with SVD

Here follows a short list of further applications where similar techniques can be used in data analysis via SVD.

- *Ancestry analysis*: By analyzing occurrences of genes in the genomes of a number of people or animals, one can detect what effects genes have on physical traits of the organism, such as what combinations of genes may increase the risk of some disease. Here we perform SVD on a matrix A where the columns correspond to genes, the rows to people, and the entries to the number of occurrences of a gene in a person's genome.
- *Latent Semantic Analysis*: Given a large collection of documents we can set up a matrix A which counts the occurrences of a set of words within the documents. Projecting this large data-set on a low-dimensional space using SVD we can look for clusters which correspond to different "topics" which we can identify. Interestingly, we can perform this topic classification without actually knowing the language in which the texts are written.
- *Recommender systems*: Like in our movie example above, SVD can be used to recommend movies, music, articles, products, etcetera based on previous user data.
- *Signal processing*: By just looking at the first principal components we can remove noise and redundancy in a signal.
- *Climate data analysis*: By looking at large collections of temperature, precipitation, and wind data we use SVD to remove noise and identify major trends.
- *Word embeddings*: Similar to latent semantic analysis, but instead we construct a matrix A where rows and columns both correspond to words in a large data-set, and where the matrix entries describe *co-occurrence*, how often the two words occur together, possibly weighted by the distance between the two words in the texts. With the SVD we can embed the words in a lower-dimensional space in such a way that the meaning of the words somewhat corresponds to the locations of the words in this space. Embeddings such as these are important in generative AI.

An example of latent semantic analysis can be found in the exercises.

7.11 Total least squares

With the techniques of the previous sections it is now easy to solve the **total least squares** problem:

Given m points x_1, \dots, x_m in \mathbb{R}^n , we want to find the k -dimensional affine subspace closest to the points (in the sense that the sum of the squares of distances to the affine subspace is as small as possible).

Let c be the average, or the center of mass, of the given points:

$$c = \frac{1}{m} \sum_{i=1}^m x_i.$$

Now let $y_i := x_i - c$, and insert these normalized vectors as rows in an $m \times n$ matrix A , this because of our normalization this new matrix will have columns which sum to zero. Find the eigenvectors v_1, \dots, v_n of $A^T A$ (ordered by decreasing eigenvalue as usual), in other words, v_i are the right singular vectors of A .

Then we know that for each $0 \leq k \leq n$, the k -dimensional subspace V_k that lies closest to the points y_i is spanned by the k first right singular vectors: $V_k = \text{span}(v_1, \dots, v_k)$.

So by translating back the whole system such that it is centered at c , we can conclude that $c + V_k$ is the affine subspace of dimension k that lies closest to the original points x_i .

We summarize:

Proposition 7.11.1. Let $x_1, \dots, x_m \in \mathbb{R}^n$. Let $c = \frac{1}{m} \sum_{i=1}^m x_i$ be the center of the points, let $y_i = x_i - c$ and let A be the matrix with rows y_i . Let v_1, \dots, v_n be the right singular vectors of A (obtained as eigenvectors of $A^T A$ ordered by decreasing eigenvalues). Then for each $0 \leq k \leq n$,

$$S = c + \text{span}(v_1, \dots, v_k)$$

is the k -dimensional affine subspace that lies closest to the points x_1, \dots, x_m , in the sense that $\sum_{i=1}^m d(x_i, S)^2$ is minimized, where $d(x_i, S)$ is the distance between x_i and S :

$$d(x_i, S) := \min_{s \in S} \|x_i - s\|.$$

When $n = 2$ and $k = 1$ the total least square problem reduces to finding a line that is closest to a number of points in the plane. Here the regular least square method minimizes the sum of squares of *vertical distance* of the points to the line, while the total least square method minimizes the sum of squares of the *minimal distance* to the line. So we can think of the regular least square method as assuming that only the y -coordinates have some errors deviating from the line, while the total least square method assumes there may also be some error in the first coordinate.

Example 7.11.2. Let us compare our two least square methods for finding the line closest to the three points

$$(1, 5), (1, 2), (4, 2).$$

With our old least square method we look for a line $y = kx + m$ of form, and inserting our three points (x, y) we obtain three linear equations in k and m :

$$\begin{cases} 1k + m = 5 \\ 1k + m = 2 \\ 4k + m = 2 \end{cases} \Leftrightarrow \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 4 & 1 \end{pmatrix} \begin{pmatrix} k \\ m \end{pmatrix} = \begin{pmatrix} 5 \\ 2 \\ 2 \end{pmatrix} \Leftrightarrow AX = b.$$

This system has no solution, but since the columns of A are independent the least square solution is obtained by a standard calculation as

$$\begin{pmatrix} k \\ m \end{pmatrix} = X = (A^* A)^{-1} A^* b = \frac{1}{2} \begin{pmatrix} -1 \\ 8 \end{pmatrix},$$

so we obtain the line

$$\ell_1 : y = -\frac{1}{2}x + 4.$$

We know that these values of k and m minimizes

$$\|AX - b\|^2 = \left\| \begin{pmatrix} (kx_1 + m) - y_1 \\ (kx_2 + m) - y_2 \\ (kx_3 + m) - y_3 \end{pmatrix} \right\|^2 = \left\| \begin{pmatrix} (1k + m) - 5 \\ (1k + m) - 2 \\ (4k + m) - 2 \end{pmatrix} \right\|^2,$$

which is the sum of squares of the vertical distances from the three points to the line.


Moving on, let us find a total least square solution. We calculate the mean in each column and find out that $c = \frac{1}{3}((1, 5) + (1, 2) + (4, 2)) = (2, 3)$ is the centre of mass of the three points. We subtract this mean from each data-point and obtain the centered points $(-1, 2)$, $(-1, -1)$, and $(2, -1)$. We put these as rows into a matrix B , and we compute $B^T B$:

$$B = \begin{pmatrix} -1 & 2 \\ -1 & -1 \\ 2 & -1 \end{pmatrix} \quad \text{which gives} \quad B^T B = \begin{pmatrix} 6 & -3 \\ -3 & 6 \end{pmatrix}.$$

Here we see that $(1, 1)$ and $(1, -1)$ are eigenvectors of $B^T B$ corresponding to the eigenvalues 3 and 9 respectively. So $v = (1, -1)$ corresponds to the largest singular value, and this is the direction in which we have to pick our 1-dimensional affine subspace according to the total least square method. Thus we get our second line as $c + tv = (2, 3) + t(1, -1)$ where $t \in \mathbb{R}$, as an equation this becomes

$$\ell_2: y = -x + 5.$$

Finally, we draw our two lines and indicate what squared distances are minimized:



△

Chapter 8

Multilinear algebra

8.1 Introduction

In mathematics we say that two objects are *isomorphic* if they fundamentally are the same in some sense even though they may look different. In linear algebra this concept is defined as follows:

Definition 8.1.1. Two vector spaces U and V over the same field \mathbb{F} are called **isomorphic** if there exists a linear bijective (invertible) map

$$\varphi : U \rightarrow V.$$

If such a map φ exists we write $U \simeq V$, and we call φ an **isomorphism** from U to V .

In general, finding out what objects are isomorphic can be a challenging problem, but in linear algebra this is trivial:

Proposition 8.1.2. Two vector spaces U and V are isomorphic if and only if they have the same dimension:

$$U \simeq V \iff \dim U = \dim V.$$

Proof. We prove the proposition in the finite-dimensional case. If $\varphi : U \rightarrow V$ is an isomorphism, then since φ is bijective, so the rank-nullity theorem gives

$$\dim(U) = \dim(\text{Im}(\varphi)) + \dim(\ker(\varphi)) = \dim(V) + \dim\{0\} = \dim(V).$$

On the other hand, if U and V have the same dimension n , pick a basis (u_1, \dots, u_n) of U and a basis (v_1, \dots, v_n) of V and define $\varphi : U \rightarrow V$ requiring $\varphi(u_i) = v_i$. This defines the linear map φ explicitly as:

$$\varphi(c_1 u_1 + \dots + c_n u_n) = c_1 v_1 + \dots + c_n v_n.$$

This map is obviously bijective. □

8.2 Tensor products of vector spaces

Tensor products of vector spaces

Just like the (external) direct sum $U \oplus V$, the *tensor product* $U \otimes V$ is a way of combining two vector spaces into a larger one. The tensor product encodes the property of *bilinearity* into a vector space.

Definition 8.2.1. Let U, V, W be vector spaces over the same field \mathbb{F} . A map $b : U \times V \rightarrow W$ is called **bilinear** if

$$b(\lambda u + \lambda' u', v) = \lambda b(u, v) + \lambda' b(u', v) \quad \text{and} \quad b(u, \lambda v + \lambda' v') = \lambda b(u, v) + \lambda' b(u, v')$$

holds for all $u, u' \in U$, $v, v' \in V$, and scalars λ, λ' . In other words, for every fixed $u \in U$, the map $b(u, -) : V \rightarrow W$ is linear, and similarly, for fixed v the map $b(-, v) : U \rightarrow W$ is linear.


For example, each inner product $\langle -, - \rangle$ on a real vector space V is a bilinear map $V \times V \rightarrow \mathbb{R}$.

Note that in the definition above, $U \times V$ is *not* a vector space, but simply consists of all pairs (u, v) where $u \in U$ and $v \in V$ with no addition or scalar multiplications defined. The tensor product $U \otimes V$ is a *vector space* in which we can embed $U \times V$, and which is universal in the sense that every bilinear map $U \times V \rightarrow W$ factors through $U \otimes V$:

Definition 8.2.2. Definition of the tensor product via universal property.

Let U and V be vector spaces over the same field \mathbb{F} . The **tensor product** $U \otimes V$ of U and V is an \mathbb{F} -vector space together with a bilinear map $\tau : U \times V \rightarrow U \otimes V$ such that the following **universal property** holds:

For every bilinear map $b : U \times V \rightarrow W$, there exists a *unique* linear map $B : U \otimes V \rightarrow W$ such that $B \circ \tau = b$ as in the following diagram:



Usually the field \mathbb{F} is understood by the context, otherwise we can write $U \otimes_{\mathbb{F}} V$ to specify which field we are taking the tensor product over.

This means that bilinear maps from $U \times V$ to W are in one to one correspondence with linear maps, but from $U \otimes V$ to W . So the bilinearity property is encoded into this new vector space.

The problem with this purely abstract definition is that we do not even know if such an object $U \otimes V$ with the required properties even exists yet, and we don't know what its elements "look like". Luckily $U \otimes V$ does exist, here is an equivalent definition that actually constructs the required object:

Definition 8.2.3. Definition of the tensor product via construction.

Let U and V be \mathbb{F} -vector spaces. Let $\mathcal{F} = \mathcal{F}(U \times V)$ be the free vector space generated by $U \times V$. This is a huge vector space in which $U \times V$ is a basis: every element (u, v) of $U \times V$ is a basis vector. This just means that the vectors of \mathcal{F} looks like finite linear combinations of such basis elements with no other relations. Now let S be the subspace of \mathcal{F} spanned by all elements of forms

$$\begin{aligned}
 (u + u', v) - (u, v) - (u', v), & \quad (u, v + v') - (u, v) - (u, v'), \\
 \lambda(u, v) - (\lambda u, v), & \quad \lambda(u, v) - (u, \lambda v),
 \end{aligned}$$

for all $u, u' \in U$, $v, v' \in V$, and all scalars $\lambda \in \mathbb{F}$.

We define the **tensor product** $U \otimes V$ to be the quotient space

$$U \otimes V = \mathcal{F}(U \times V)/S,$$

and we let $\tau : U \times V \rightarrow U \otimes V$ be the map $(u, v) \mapsto (u, v) + S$. We shall write just $u \otimes v$ for $\tau(u, v)$.

Elements of $U \otimes V$ are sometimes called **tensors**^a. Elements of $U \otimes V$ which can be written $u \otimes v$ are called **pure tensors**, every element of $U \otimes V$ is a linear combination pure tensors.

^aBut note that this is a very common word and unfortunately it is used differently depending on the context.

Now let us think about what the vectors of $U \otimes V$ are more concretely. First, from the definition we see that every element of $U \otimes V$ is a linear combination of elements of form $u \otimes v$ where $u \in U$ and $v \in V$.

Since the elements of $U \otimes V$ lie in the quotient space $\mathcal{F}(U \times V)/S$, we can add or subtract elements of S without changing the affine subset. For example, in $U \otimes V$ we have

$$\begin{aligned}
 (3u - 2u') \otimes v &= (3u - 2u', v) + S = (3u - 2u', v) - \underbrace{\left((3u - 2u', v) - 3(u, v) + 2(u', v) \right)}_{\in S} + S \\
 &= 3(u, v) - 2(u', v) + S = 3(u \otimes v) - 2(u' \otimes v).
 \end{aligned}$$

The point is that since we are taking a quotient by S , we are effectively turning every element of S to the zero vector in $U \otimes V$. So the following relations will hold in $U \otimes V$:

Proposition 8.2.4. *The tensor product $U \otimes V$ is spanned by elements of form $u \otimes v$ where $u \in U$ and $v \in V$, such elements satisfy the relations:*

- $(u + u') \otimes v = (u \otimes v) + (u' \otimes v)$
- $u \otimes (v + v') = (u \otimes v) + (u \otimes v')$
- $(\lambda u) \otimes v = \lambda(u \otimes v) = u \otimes (\lambda v)$

Proof. Each equation differ only by an element of S , so since $U \otimes V$ is a quotient by S , each equality holds in $U \otimes V$. □

This shows that the symbols $u \otimes v$ are linear in each component separately, and the third property shows that we are free to "move scalars" through the tensor product sign, and we don't really need to put all the parentheses here because the meaning is clear.

From now on we shall write just $\lambda u \otimes v + \lambda' u' \otimes v'$ instead of $\lambda(u \otimes v) + \lambda'(u' \otimes v')$, the tensor product implicitly comes before addition.

Example 8.2.5. The space $\mathbb{R}^3 \otimes \mathbb{R}^2$ is spanned by elements of form $u \otimes v$ where $u \in \mathbb{R}^3$ and $v \in \mathbb{R}^2$. We can make a simplification like:

$$\begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ 6 \end{pmatrix} + \begin{pmatrix} 1 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 2 \\ 4 \end{pmatrix} = 3 \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} + 2 \begin{pmatrix} 1 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} = \begin{pmatrix} 3 \\ 6 \\ 9 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} + \begin{pmatrix} 2 \\ 2 \\ 2 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} = \begin{pmatrix} 5 \\ 8 \\ 11 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix}.$$

However, an expression like

$$\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

can never be simplified further to just one term $u \otimes v$.

△

The easiest way to think about about $U \otimes V$ when the vector spaces U and V are finite-dimensional is probably the following:

Proposition 8.2.6. *Let (e_1, \dots, e_m) be a basis for U and let (f_1, \dots, f_n) be a basis for V . Then*

$$\{e_i \otimes f_j \mid 1 \leq i \leq m, 1 \leq j \leq n\}$$

is a basis for $U \otimes V$. In particular we have

$$\dim(U \otimes V) = \dim(U) \cdot \dim(V).$$

Proof. Let (e_1, \dots, e_m) be a basis for U and let (f_1, \dots, f_n) be basis for V . Then for $u \otimes v \in U \otimes V$ we have

$$u \otimes v = \left(\sum_{i=1}^m a_i e_i \otimes \sum_{j=1}^n b_j f_j \right) = \sum_{i,j} a_i b_j (e_i \otimes f_j).$$

Since $U \otimes V$ is spanned by such elements $u \otimes v$, we have shown that each element of $U \otimes V$ is a linear combination of elements of form $e_i \otimes f_j$, so $\{e_i \otimes f_j \mid 1 \leq i \leq m, 1 \leq j \leq n\}$ spans $U \otimes V$. One can use the universal property of the tensor product to show that the elements $e_i \otimes f_j$ are linearly independent, we omit this part here.

Since there are $m \cdot n$ elements of form $e_i \otimes f_j$ in the basis we get

$$\dim(U \otimes V) = m \cdot n = \dim(U) \cdot \dim(V).$$

□

The same statement is true when U and V are infinite-dimensional: If \mathcal{B}_U is a basis for U and if \mathcal{B}_V is a basis for V , then

$$\mathcal{B}_{U \otimes V} = \{e \otimes f \mid e \in \mathcal{B}_U, f \in \mathcal{B}_V\}$$

is a basis for $U \otimes V$.

Recall that for the direct sum, we can view both U and V as subspaces of $U \oplus V$ in a natural way.

Tensor products don't work quite like that, instead we can think of $U \otimes V$ as being obtained by attaching a copy of V at every direction in U or vice versa: for each fixed $u \in U$ we have a subspace

$$U \otimes V \supset u \otimes V = \{u \otimes v \mid v \in V\} \simeq V.$$

Proposition 8.2.7. *The tensor product is associative, and it distributes over direct sums; for vector spaces U, V, W over the same field there are canonical isomorphisms:*

$$(U \otimes V) \otimes W \simeq U \otimes (V \otimes W)$$

$$U \otimes (V \oplus W) \simeq (U \otimes V) \oplus (U \otimes W)$$

$$(U \oplus V) \otimes W \simeq (U \otimes W) \oplus (V \otimes W).$$

Proof. The word "canonical" here means that there is a natural way to pick the map between the two vector spaces. For the first statement the canonical isomorphism is defined by $(u \otimes v) \otimes w \mapsto u \otimes (v \otimes w)$, it is just a reordering of the parenthesis.

For the second statement, the isomorphism is given by $u \otimes (v, w) \mapsto (u \otimes v, u \otimes w)$ and in the third it is given by $(u, v) \otimes w \mapsto (u \otimes w, v \otimes w)$. All of these are bijective.

Alternatively, for the finite-dimensional case it suffices to show that the spaces on each side have the same dimension. This follows from the fact that $\dim(A \oplus B) = \dim A + \dim B$ and that $\dim(A \otimes B) = \dim A \cdot \dim B$. □

Because of these natural isomorphisms we do not really need to differentiate between the spaces on the left and the right side in the proposition, for example, we can write simply $U \otimes V \otimes W$ without parentheses and think of it as a vector space spanned by elements of form

$$u \otimes v \otimes w \quad \text{for } u \in U, v \in V, w \in W.$$

So far we have defined tensor products and their elements in a rather abstract way. How can we represent elements of a tensor products in a way suitable for calculations? The answer is *multi-dimensional arrays*.

Consider a tensor product $U \otimes V$. If we pick bases (u_1, \dots, u_m) of U and (v_1, \dots, v_n) of V we can write each element x of the tensor product $U \otimes V$ as

$$x = \sum_{i=1}^m \sum_{j=1}^n \alpha_{ij} u_i \otimes v_j.$$

Then with respect to these bases can represent x as the $m \times n$ -matrix $[x] = (\alpha_{ij})_{ij}$.

Analogously, if we have yet another space W with basis (w_1, \dots, w_p) , we can express a "3-tensor" y in $U \otimes V \otimes W$ as

$$y = \sum_{i=1}^m \sum_{j=1}^n \sum_{k=1}^p \alpha_{ijk} u_i \otimes v_j \otimes w_k,$$

where the coefficients α_{ijk} now have three indices. So with respect to these three bases, we can represent y as a three-dimensional array containing the element α_{ijk} in position (i, j, k) , we can think of these as $m \times n \times p$ -matrices stacked p times, or as "shoe boxes" of (length,width,height) equal to (m, n, p) respectively.

In general, vectors in $V_1 \otimes \dots \otimes V_N$ (a tensor products of N vector spaces) are harder to visualize, but with respect to fixed bases in all the vector spaces we can think of these vectors as N -dimensional arrays with coefficients $\alpha_{i_1, i_2, \dots, i_N}$ with N indices.

Just as with matrices however, these are just representations of the tensors, they depend on the choice of basis of all the vector spaces. Other choices of bases will yield different arrays of numbers.

Example 8.2.8. In the space $\mathbb{R}^3 \otimes \mathbb{R}^2$, a basis is given by the six vectors

$$e_1 \otimes e_1, \quad e_1 \otimes e_2, \quad e_2 \otimes e_1, \quad e_2 \otimes e_2, \quad e_3 \otimes e_1, \quad e_3 \otimes e_2,$$

where (e_1, e_2, e_3) on the left side is the standard basis of \mathbb{R}^3 , and where (e_1, e_2) on the right side is the standard basis of \mathbb{R}^2 .

The tensor

$$x = \begin{pmatrix} 5 \\ 0 \\ 3 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} + \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

can be expressed in this standard basis as

$$x = 5e_1 \otimes e_1 + 11e_1 \otimes e_2 + e_2 \otimes e_2 + 3e_3 \otimes e_1 + 7e_3 \otimes e_2.$$

So with respect to these bases we can represent x as a matrix

$$[x] = \begin{pmatrix} 5 & 11 \\ 0 & 1 \\ 3 & 7 \end{pmatrix}.$$

△

The analogous idea works for 3-tensors:

Example 8.2.9. In the space $\mathbb{R}^2 \otimes \mathbb{R}^3 \otimes \mathbb{R}^2$, let us use a more compact notation and write $e_{ijk} = e_i \otimes e_j \otimes e_k$. Then with respect to this basis, the tensor

$$y = \begin{pmatrix} 1 \\ 2 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ -1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 2 \\ -1 \end{pmatrix} + 3 \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

can be written

$$\begin{aligned} y &= 6e_{111} - 3e_{112} - 2e_{121} + e_{122} + 12e_{211} - 6e_{212} - 4e_{221} + 2e_{222} \\ &\quad + 3e_{211} + 3e_{212} + 3e_{221} + 3e_{222} + 3e_{231} + 3e_{232} \\ &= 6e_{111} - 3e_{112} - 2e_{121} + e_{122} + 15e_{211} - 3e_{212} - e_{221} + 5e_{222} + 3e_{231} + 3e_{232} \end{aligned}$$

So we can visualize $[y]$ as the $2 \times 3 \times 2$ array

Adding two tensors represented in the same basis corresponds to adding their coordinate arrays, but note that there is no well defined unique way of multiplying tensor-arrays like how we can multiply matrices.

Just like the matrix of a linear maps depends on the choice of bases, the same tensor y above will have different 3d-array representations if we change basis in any of the three component spaces of $\mathbb{R}^2 \otimes \mathbb{R}^3 \otimes \mathbb{R}^2$.

△

8.3 Dual spaces

Definition 8.3.1. Let V be a vector space over a field \mathbb{F} . The corresponding **dual space** V^* is the vector space of all linear maps from V to \mathbb{F} :

$$V^* = \{f : V \rightarrow \mathbb{F} \mid f \text{ is linear}\} = \text{Hom}(V, \mathbb{F}).$$

Such maps are also called **linear functionals** on V .

V^* is a vector space where addition and scalar multiplication is defined in the natural way: for $f, g \in V^*$ and $\lambda \in \mathbb{F}$ we have

$$(f + g)(v) = f(v) + g(v) \quad \text{and} \quad (\lambda f)(v) = \lambda f(v).$$

We have previously defined both the Hermitian conjugate A^* of a matrix A , and more generally the adjoint F^* of a linear map F . However, there shouldn't be at risk of mixing these concepts up, because in this new definition, V and V^* are vector space and not a linear maps.

Example 8.3.2. For example, consider an evaluation e on the vector space \mathcal{P} of polynomials defined by $e(p(x)) = p(5)$. Then $e \in \mathcal{P}^*$.

Another example: the function g taking the trace of 2×2 matrices $g(A) = \text{tr}(A)$ is also an element of the dual space, $g \in \text{Mat}_2(\mathbb{R})^*$.

△

Given a basis for a finite-dimensional vector space, we get a corresponding basis for its dual.

Proposition 8.3.3. Let V be a finite-dimensional vector space over a field \mathbb{F} , and let (e_1, \dots, e_n) be a basis for V . For $1 \leq i \leq n$ we define elements $e_i^* \in V^*$, where

$$e_i^* : V \rightarrow \mathbb{F} \quad \text{is defined by} \quad e_i^*(e_j) = \delta_{ij} = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

Then (e_1^*, \dots, e_n^*) is a basis for V^* which is called the **dual basis** corresponding to the basis (e_1, \dots, e_n) of V .

Proof. Assume that e_1^*, \dots, e_n^* are linearly dependent with $0 = \sum_{i=1}^n c_i e_i^*$. Both sides of this equality is a map $V \rightarrow \mathbb{F}$, and in particular we can apply it to a basis element e_j which gives

$$0 = \sum_{i=1}^n c_i e_i^*(e_j) = \sum_{i=1}^n c_i \delta_{ij} = c_j$$

for every j , so all c_j are zero, and the elements e_i^* are linearly independent.

To prove that the elements e_i^* span V^* , let $f \in V^*$. We claim that

$$f = f(e_1)e_1^* + \dots + f(e_n)e_n^*,$$

this is clearly true, because applying both sides to e_j we get

$$f(e_j) = f(e_1)e_1^*(e_j) + \dots + f(e_n)e_n^*(e_j) = f(e_1)\delta_{1j} + \dots + f(e_n)\delta_{nj} = f(e_j)\delta_{jj} = f(e_j).$$

So the two linear maps f and $f(e_1)e_1^* + \dots + f(e_n)e_n^*$ are equal on a basis for V , and thus they are the same map. So $f = f(e_1)e_1^* + \dots + f(e_n)e_n^*$ holds for any f , which shows that the elements e_i^* span V^* . □

Note that in general we can not express elements of V^* as v^* for individual elements $v \in V$. Only when v is an element of a specified basis of V we define v^* , and then the function v^* actually depends on how the other vectors of the basis were chosen.

Example 8.3.4. Consider the space \mathcal{P}_2 of real polynomials of degree ≤ 2 . The standard basis of \mathcal{P}_2 is $(e_0, e_1, e_2) = (1, x, x^2)$. The corresponding dual basis of \mathcal{P}_2^* is (e_0^*, e_1^*, e_2^*) where explicitly

$$e_0^*(a + bx + cx^2) = e_0^*(ae_0 + be_1 + ce_2) = ae_0^*(e_0) + be_0^*(e_1) + ce_0^*(e_2) = a$$

$$e_1^*(a + bx + cx^2) = e_1^*(ae_0 + be_1 + ce_2) = ae_1^*(e_0) + be_1^*(e_1) + ce_1^*(e_2) = b$$

$$e_2^*(a + bx + cx^2) = e_2^*(ae_0 + be_1 + ce_2) = ae_2^*(e_0) + be_2^*(e_1) + ce_2^*(e_2) = c$$

so each e_i^* acts on a polynomial by picking out the corresponding coefficient with respect to the basis $(1, x, x^2)$ of \mathcal{P}_2 .

Any linear map $\mathcal{P}_2 \rightarrow \mathbb{R}$ can be expressed in this basis (e_0^*, e_1^*, e_2^*) . For example, consider the evaluation map $f : \mathcal{P}_2 \rightarrow \mathbb{R}$ defined by $f(p(x)) = p(2)$. Then

$$f = e_0^* + 2e_1^* + 4e_2^*$$

which can be verified by evaluating both sides at an arbitrary polynomial $a + bx + cx^2$.

△

Note that the proof above shows that $V \simeq V^*$ when V is finite-dimensional. This is no longer true in the infinite-dimensional case as the following example illustrates:

Example 8.3.5. Let \mathcal{P} be the space of polynomials. Let $(1, x, x^2, \dots) = (e_0, e_1, e_2, \dots)$ be the standard basis for this space and for $i \in \mathbb{N}$, define $e_i^*(e_j) = \delta_{ij}$ as before. Then each e_i^* lies in \mathcal{P}^* , but these elements no longer span \mathcal{P} , since each $f \in V^*$ is not a finite linear combination of the vectors e_i^* .

In particular, define $f \in \mathcal{P}$ by $f(p(x)) = p(1)$. Then f is not expressible as a finite linear combination of e_i^* since f has the property that $f(e_n) = 1$ for all n , but no finite linear combination of e_i^* can have this same property.

△

When V is finite-dimensional however, it turns out that there is a natural way to identify a vector space V with its *double dual* V^{**} .

Theorem 8.3.6. The double dual isomorphism theorem.

Let V be a finite-dimensional vector space over a field \mathbb{F} . Then the spaces V and V^{**} are canonically isomorphic.

Proof. Here V^{**} means the dual space of the dual space of V , in other words $V^{**} = (V^*)^*$. The isomorphism $\varphi : V \rightarrow V^{**}$ is defined by $\varphi(v) = \Phi_v$ where each $\Phi_v \in V^{**}$ is a map $\Phi_v : V^* \rightarrow \mathbb{F}$ defined by

$$\Phi_v(f) = f(v) \quad \text{for each } f \in V^*.$$

We omit the proof that this is indeed an isomorphism, some parts of the proof is left for the exercises. □

We can now combine tensors and duals to show an important correspondence which lets us think of linear maps as tensors (without necessarily picking a basis in each space).

Proposition 8.3.7. Let V and W be finite-dimensional vector spaces over \mathbb{F} and let $\text{Hom}_{\mathbb{F}}(V, W)$ be the vector space of linear maps from V to W , then there is a canonical isomorphism

$$\text{Hom}_{\mathbb{F}}(V, W) \simeq W \otimes V^*.$$

Proof. Each pure tensor $w \otimes f \in W \otimes V^*$ corresponds to a linear map $F : V \rightarrow W$ defined by $F(v) = f(v)w$, and linear combinations of such pure tensors give the corresponding maps. This is indeed an isomorphism, we leave it to the exercises to verify this.

Alternatively we can proceed more concretely: pick a basis (e_1, \dots, e_n) of V and a basis (f_1, \dots, f_m) in W , then we know that the tensor product $W \otimes V^*$ has basis consisting of elements of form $f_i \otimes e_j^*$ and we can define our isomorphism $\varphi : W \otimes V^* \rightarrow \text{Hom}_{\mathbb{F}}(V, W)$ explicitly by

$$\varphi\left(\sum_{i,j} a_{ij} f_i \otimes e_j^*\right)(v) = \sum_{i,j} a_{ij} e_j^*(v) f_i,$$

taking $v = e_k$, for each basis vector $e_k \in V$ we see that the matrix of this map with respect to our chosen bases is precisely (a_{ij}) , and each basis matrix e_j^* with a single 1 in position (i, j) is the image under φ of the basis elements $f_i \otimes e_j^*$ of $W \otimes V^*$. □

In the special case that V is a finite-dimensional inner product space, there is actually a natural isomorphism between V and V^* given by

$$\varphi : V \rightarrow V^* \quad \text{where} \quad \varphi(v) = \langle -, v \rangle,$$

the Riesz representation theorem 5.6.1 guarantees that this map is bijective; each $f \in V^*$ has form $\langle \cdot, v \rangle$ for some v .

Given a linear map $F : V \rightarrow W$ we get a corresponding map $F^* : W^* \rightarrow V^*$ given by *pre-composition* by F given by

$$F^* : g \mapsto g \circ F \quad \begin{array}{ccc} V & \xrightarrow{F} & W \\ g \circ F \downarrow & & \swarrow g \\ \mathbb{F} & & \end{array}$$

8.4 Higher order tensors

We can form tensor product of more than two vector spaces:

Definition 8.4.1. Let V_1, \dots, V_n and W be vector spaces over the same field \mathbb{F} . A map

$$f : V_1 \times \dots \times V_n \rightarrow W$$

is called **multilinear** if it is linear when fixing all but one of the components; for each index $1 \leq i \leq n$ we have

$$\begin{aligned} f(v_1, \dots, v_{i-1}, \lambda v_i + \mu v'_i, v_{i+1}, \dots, v_n) \\ = \lambda f(v_1, \dots, v_{i-1}, v_i, v_{i+1}, \dots, v_n) + \mu f(v_1, \dots, v_{i-1}, v'_i, v_{i+1}, \dots, v_n). \end{aligned}$$

The tensor product

$$\bigotimes_{i=1}^n V_i := V_1 \otimes \dots \otimes V_n$$

is called a **tensor product of order n**. It is a vector space together with a multilinear embedding $\tau : V_1 \times \dots \times V_n \rightarrow V_1 \otimes \dots \otimes V_n$ such that for every multilinear map f as above, there exists a unique linear map $F : V_1 \otimes \dots \otimes V_n \rightarrow W$ such that $f = F \circ \tau$.

The tensor product $V_1 \otimes \dots \otimes V_n$ is spanned by elements of form $v_1 \otimes \dots \otimes v_n$ where $v_i \in V_i$, such elements are linear in each component:

$$\begin{aligned} v_1 \otimes \dots \otimes v_{i-1} \otimes (\lambda v_i + \mu v'_i) \otimes v_{i+1} \otimes \dots \otimes v_n \\ = \lambda (v_1 \otimes \dots \otimes v_{i-1} \otimes v_i \otimes v_{i+1} \otimes \dots \otimes v_n) + \mu (v_1 \otimes \dots \otimes v_{i-1} \otimes v'_i \otimes v_{i+1} \otimes \dots \otimes v_n). \end{aligned}$$

Elements of $V_1 \otimes \dots \otimes V_n$ are called **tensors of order n**. Elements that can be written $v_1 \otimes \dots \otimes v_n$ with $v_i \in V_i$ are called **pure tensors**.

The tensor product $V_1 \otimes \dots \otimes V_n$ is essentially the same as what we get if we iteratively take the tensor product of two vector spaces at a time:

$$V_1 \otimes \dots \otimes V_n \simeq (\dots (((V_1 \otimes V_2) \otimes V_3) \otimes \dots) \otimes V_n),$$

the parentheses is the only difference.

A special case is when all the vector spaces V_i are all the same, then we write

$$V^{\otimes n} := \underbrace{V \otimes V \otimes \dots \otimes V}_{n \text{ copies}}$$

and call this the **n-fold** tensor product on V .

Tensors products of both a space and its dual is an important object of study both in theory and applications:

Definition 8.4.2. A (p, q) -tensor is an element of

$$\underbrace{V \otimes V \otimes \cdots \otimes V}_p \otimes \underbrace{V^* \otimes V^* \otimes \cdots \otimes V^*}_q.$$

One should be aware that this definition varies in different books and in different fields, even within mathematics. It is common to define a (p, q) -tensor as *multilinear map*

$$T : \underbrace{V^* \times V^* \times \cdots \times V^*}_p \times \underbrace{V \times V \times \cdots \times V}_q \rightarrow \mathbb{F}$$

or equivalently as a *linear map*

$$T : \underbrace{V^* \otimes V^* \otimes \cdots \otimes V^*}_p \otimes \underbrace{V \otimes V \otimes \cdots \otimes V}_q \rightarrow \mathbb{F}.$$

This is essentially the same as our definition above, because via the canonical isomorphisms

$$\text{Hom}(V, W) \simeq W \otimes V^*, \quad V \simeq \mathbb{F} \otimes V, \quad (V \otimes V)^* \simeq V^* \otimes V^*, \quad V \simeq V^{**},$$

we get a canonical isomorphism

$$\begin{aligned} & \text{Hom}(\underbrace{V^* \otimes V^* \otimes \cdots \otimes V^*}_p \otimes \underbrace{V \otimes V \otimes \cdots \otimes V}_q, \mathbb{F}) \\ & \simeq \mathbb{F} \otimes \left(\underbrace{V^* \otimes V^* \otimes \cdots \otimes V^*}_p \otimes \underbrace{V \otimes V \otimes \cdots \otimes V}_q \right)^* \\ & \simeq \underbrace{V \otimes V \otimes \cdots \otimes V}_p \otimes \underbrace{V^* \otimes V^* \otimes \cdots \otimes V^*}_q. \end{aligned}$$

In texts more geared towards applications, a (p, q) -tensor is sometimes defined a multidimensional array

$$T_{j_1, \dots, j_q}^{i_1, \dots, i_p},$$

or perhaps more precisely as linear association of such an array to each choice of basis of V , subject to certain transformation laws which are somewhat non-intuitive and hard to write down explicitly.

However, this is also equivalent to our definition; if we choose a basis (e_1, \dots, e_n) of V we get a corresponding dual basis (e_1^*, \dots, e_n^*) of V^* , then $\underbrace{V \otimes V \otimes \cdots \otimes V}_p \otimes \underbrace{V^* \otimes V^* \otimes \cdots \otimes V^*}_q$ has a basis consisting

of the $n^{(p+q)}$ elements

$$e_{i_1} \otimes e_{i_2} \otimes \cdots \otimes e_{i_p} \otimes e_{j_1}^* \otimes e_{j_2}^* \otimes \cdots \otimes e_{j_q}^*$$

where all indices ranges from 1 to $n = \dim V$. So with respect to this basis, we can represent each tensor as a $(p + q)$ -dimensional array $T_{j_1, \dots, j_q}^{i_1, \dots, i_p}$ containing the coefficients of the (p, q) -tensor with respect to our chosen basis.

Many objects of linear algebra can be viewed as tensors, let us look at (p, q) -tensors on a space V for various choices of p and q :

Example 8.4.3.

- $(1, 0)$ -tensors are just vectors in V .
- $(0, 1)$ -tensors are just functionals $V \rightarrow \mathbb{F}$.
- $(1, 1)$ -tensors are linear operators on V since $V \otimes V^* \simeq \text{Hom}(V, V)$.
- $(0, 2)$ -tensors are bilinear forms, since $V^* \otimes V^* \simeq \text{Hom}(V \otimes V, \mathbb{F})$.
- $(1, 2)$ -tensors are bilinear multiplications on V since $\text{Hom}(V \otimes V, V) \simeq V \otimes V^* \otimes V^*$. Such a multiplication is a rule that takes two vectors as input and produces a vector. For example, the cross

product is an example of a $(1, 2)$ -tensor when $V = \mathbb{R}^3$, since the cross product is a bilinear map $\times : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$.

- Matrix-multiplication can be viewed as a $(1, 2)$ -tensor for $V = \text{Mat}_{m \times n}(\mathbb{R})$, or, when $m = n$ we can view it as a $(3, 3)$ -tensor for $V = \mathbb{R}^n$ (since $\text{Mat}_n(\mathbb{R}) \simeq \mathbb{R}^n \otimes (\mathbb{R}^n)^*$).
- $(0, n)$ -tensors take n vectors and produce a number in a multilinear way. For example, the determinant corresponds to a $(0, n)$ -tensor where $n = \dim V$.
- $(0, 0)$ -tensors are usually *defined* as the scalars \mathbb{F} .

△

Contraction of tensors

Elements of V^* take vectors of V and produces numbers, so there is a natural way of "cancelling" factors V and V^* pairwise in a (p, q) -tensor, this is called a *contraction* of the tensor.

Definition 8.4.4. Let $V \otimes \dots \otimes V \otimes V^* \otimes V^* = V^{\otimes p} \otimes (V^*)^{\otimes q}$ be a (p, q) -tensor product. For each pairs of indices $1 \leq i \leq p$ and $1 \leq j \leq q$, we define a corresponding **(i, j) -contraction** as the linear map

$$C_{ij} : V^{\otimes p} \otimes (V^*)^{\otimes q} \rightarrow V^{\otimes(p-1)} \otimes (V^*)^{\otimes(q-1)}$$

defined on pure tensors by

$$C_{ij}(v_1 \otimes \dots \otimes v_p \otimes f_1 \otimes \dots \otimes f_q) = f_j(v_i)(v_1 \otimes \dots \otimes v_{i-1} \otimes v_{i+1} \otimes \dots \otimes v_p \otimes f_1 \otimes f_{j-1} \otimes f_{j+1} \otimes f_q),$$

in other words, we evaluate the j 'th V^* -component at the i 'th V -component and multiply the resulting scalar by the $(p - 1, q - 1)$ -tensor obtained by removing these components.

Example 8.4.5. Let $V = \mathbb{R}^2$, let $a \in V^*$ be the functional taking the average $a(x, y) = \frac{x+y}{2}$, and let $p_1, p_2 \in V^*$ be the projection onto the first and second component respectively: $p_1(x, y) = x$ and $p_2(x, y) = y$. Consider the $(3, 2)$ -tensor $w \in V \otimes V \otimes V \otimes V^* \otimes V^*$ where

$$w = 2 \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} \otimes a \otimes p_1 + 3 \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ 4 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} \otimes p_2 \otimes a.$$

The $(3, 1)$ -contraction cancel out the "middle" components: $V \otimes V \otimes (V \otimes V^*) \otimes V^*$ of w like so:

$$\begin{aligned} C_{3,1}(w) &= 2a\left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}\right) \cdot \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ 1 \end{pmatrix} \otimes p_1 + 3p_2\left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}\right) \cdot \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ 4 \end{pmatrix} \otimes a \\ &= 3 \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ 1 \end{pmatrix} \otimes p_1 + 6 \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ 4 \end{pmatrix} \otimes a. \end{aligned}$$

Another example: the $(1, 1)$ -contraction of w is

$$C_{1,1}(w) = 2a\left(\underbrace{\begin{pmatrix} 1 \\ -1 \end{pmatrix}}_0\right) \cdot \begin{pmatrix} 3 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} \otimes p_1 + 3p_2\left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) \cdot \begin{pmatrix} 3 \\ 4 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} \otimes a = 3 \begin{pmatrix} 3 \\ 4 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} \otimes a.$$

△

An example of tensor contraction is composition of linear maps on a vector space¹, here we represent two linear maps as $(1, 1)$ -tensors $f, g \in V \otimes V^*$. Then the composition $f \circ g$ corresponds the contraction of $f \otimes g \in V \otimes V^* \otimes V \otimes V^*$ obtained by contracting the middle two vector spaces.

If we identify V and V^* via an inner product, we can also perform contractions of $V \otimes V \otimes \dots \otimes V$ the same way. Tensor contractions are useful in applications such as neural networks.

¹Or think of it as matrix multiplication.

8.5 Tensor products of linear maps

There is a natural way to form the tensor product of two linear maps:

Definition 8.5.1. Let $F : U \rightarrow V$ and $G : U' \rightarrow V'$ be two linear maps between vector spaces, all over the same field.

Then we define the **tensor product** of F and G as the linear map defined by

$$F \otimes G : U \otimes U' \rightarrow V \otimes V' \quad \text{defined by} \quad (F \otimes G)(u \otimes u') = F(u) \otimes G(u').$$

So we just apply each map to each component separately.

The following properties of tensor products of linear maps are easy to prove:

Proposition 8.5.2. Let F, G, H, K be linear maps and let λ be a scalar. The tensor products of linear maps satisfies:

1. $(\lambda F) \otimes G = \lambda(F \otimes G) = F \otimes (\lambda G)$
2. $(F + G) \otimes H = F \otimes H + G \otimes H$
3. $F \otimes (G + H) = F \otimes G + F \otimes H$
4. $(F \otimes G) \circ (H \otimes K) = (F \circ H) \otimes (G \circ K)$

whenever the domains and co-domains of F, G, H, K are such that the corresponding operations of each statement are defined.

Proof. Apply each equality to a pure tensor, for example, (2) follows from

$$\begin{aligned} ((F + G) \otimes H)(u \otimes v) &= (F + G)(u) \otimes H(v) = (F(u) + G(u)) \otimes H(v) \\ &= F(u) \otimes H(v) + G(u) \otimes H(v) = (F \otimes H + G \otimes H)(u \otimes v), \end{aligned}$$

since the two linear maps are equal on pure tensors and such tensors span the domain, the maps are equal. \square

The tensor product of linear map should be compared to the direct sum of linear maps defined previously. The direct sum $U \oplus V$ of vector spaces consists of pairs (u, v) with component-wise addition and scalar multiplication. Recall that we have previously defined the direct sum of two linear maps: If $F : U \rightarrow V$ and $G : U' \rightarrow V'$ as

$$F \oplus G : U \oplus U' \rightarrow V \oplus V' \quad \text{where} \quad (F \oplus G)(u, u') = (F(u), G(u')).$$

The following proposition tells us how the direct sum and tensor products of linear maps interact.

Proposition 8.5.3. Let F, G, H be linear maps between vector spaces over the same field. Under the natural identification between domains and co-domains in 8.2.7, we have

$$\begin{aligned} (F \otimes G) \otimes H &= F \otimes (G \otimes H) \\ F \otimes (G \oplus H) &= (F \otimes G) \oplus (F \otimes H) \\ (F \oplus G) \otimes H &= (F \otimes H) \oplus (G \otimes H) \end{aligned}$$

Proof. Let

$$F : U_1 \rightarrow V_1 \quad G : U_2 \rightarrow V_2 \quad H : U_3 \rightarrow V_3.$$

There is a slight annoyance here. The map $(F \otimes G) \otimes H$ is technically different² from $F \otimes (G \otimes H)$, since the maps have different domains and co-domains: for example, $(F \otimes G) \otimes H$ is defined on $(U_1 \otimes U_2) \otimes U_3$ while $F \otimes (G \otimes H)$ is defined on $U_1 \otimes (U_2 \otimes U_3)$. But if we identify the domains and co-domains and think of them both as being spanned by elements $w_1 \otimes w_2 \otimes w_3$ without parenthesis, both maps send the element $u_1 \otimes u_2 \otimes u_3$ to $F(u_1) \otimes G(u_2) \otimes H(u_3)$ for all $u_i \in U_i$. An analogous argument works for the other two statements. \square

Note that different choices of bases in $U_1, U_2, U_3, V_1, V_2, V_3$ will give different matrix representations of the maps in the proposition.

²Hence the phrasing "under the natural identification" in the proposition.

Example 8.5.4. Define three linear maps F, G, H as follows:

$$F : \mathcal{C}^\infty(\mathbb{R}) \rightarrow \mathcal{C}^\infty(\mathbb{R}) \quad \text{where} \quad F(f(x)) = f'(x),$$

$$G : \text{Mat}_2(\mathbb{R}) \rightarrow \mathbb{R} \quad \text{where} \quad G(A) = \text{tr}(A),$$

$$H : \mathbb{R} \rightarrow \mathbb{R}^2 \quad \text{where} \quad H(x) = (x, x).$$

Then $T = F \otimes (G \oplus H)$ is a map

$$T : \mathcal{C}^\infty(\mathbb{R}) \otimes (\text{Mat}_2(\mathbb{R}) \oplus \mathbb{R}) \rightarrow \mathcal{C}^\infty(\mathbb{R}) \otimes (\mathbb{R} \oplus \mathbb{R}^2).$$

For example we may evaluate:

$$T\left(\sin(x) \otimes \left(\begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}, 7\right)\right) = \cos(x) \otimes (5, (7, 7)).$$

If we compare this with the map $S = (F \otimes G) \oplus (F \otimes H)$ we see that they have the same effect but that the arguments are "packaged" differently.

△

The tensor product gives an alternative and easier way to define the **complexification** of vector spaces and maps that we discussed in Section 9.2.

Example 8.5.5. Let V be a real vector space. The complex numbers \mathbb{C} is also a real vector space (of dimension 2). So we can define

$$V^{\mathbb{C}} = \mathbb{C} \otimes_{\mathbb{R}} V.$$

Initially this is a real vector space, but we can turn it into a complex vector space in the natural way by defining multiplication on the left component: for $\lambda \in \mathbb{C}$ we define $\lambda(\alpha \otimes v) := (\lambda\alpha) \otimes v$. Then $V^{\mathbb{C}}$ becomes a complex vector space. Since we can move *real* numbers through the tensor product sign we can express every element of $V^{\mathbb{C}}$ as $1 \otimes v + i \otimes w$ for $v, w \in V$.

Moreover, if $F : V \rightarrow W$ is a map between real vector space, then $\text{id} \otimes F$ is the corresponding complexification of the map F that in the previous section this was written $F^{\mathbb{C}}$.

△

8.6 Kronecker product

What is the matrix of the linear map $F \otimes G$? So far we have defined the tensor product $F \otimes G$ of linear maps in a basis-free way. But given choices of bases for all involved vector spaces, can the matrix of $F \otimes G$ be expressed in terms of the matrices of F and of G ?

The answer is yes, but to see this, let us start with defining the Kronecker-product of matrices:

Definition 8.6.1. Let $A = (a_{ij})$ be an $m \times n$ matrix, and let $B = (b_{ij})$ be an $m' \times n'$ -matrix. The **Kronecker product** of A and B is defined as the block matrix

$$A \otimes B = (a_{ij}B) = \begin{pmatrix} a_{11}B & a_{12}B & \cdots & a_{1n}B \\ a_{21}B & a_{22}B & \cdots & a_{2n}B \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1}B & a_{m2}B & \cdots & a_{mn}B \end{pmatrix}.$$

note that $A \otimes B$ is an $(mm') \times (nn')$ -matrix.

So we obtain $A \otimes B$ by multiplying each entry of A by the matrix B .

Example 8.6.2. Let $A = \begin{pmatrix} 1 & -2 \\ 3 & 0 \end{pmatrix}$ and $B = \begin{pmatrix} 1 & 2 & 3 \\ -1 & 0 & 1 \end{pmatrix}$.

Then

$$A \otimes B = \begin{pmatrix} 1 \begin{pmatrix} 1 & 2 & 3 \\ -1 & 0 & 1 \end{pmatrix} & -2 \begin{pmatrix} 1 & 2 & 3 \\ -1 & 0 & 1 \end{pmatrix} \\ 3 \begin{pmatrix} 1 & 2 & 3 \\ -1 & 0 & 1 \end{pmatrix} & 0 \begin{pmatrix} 1 & 2 & 3 \\ -1 & 0 & 1 \end{pmatrix} \end{pmatrix} = \begin{pmatrix} 1 & 2 & 3 & -2 & -4 & -6 \\ -1 & 0 & 1 & 2 & 0 & -2 \\ 3 & 6 & 9 & 0 & 0 & 0 \\ -3 & 0 & 3 & 0 & 0 & 0 \end{pmatrix}.$$

The order of A and B is not irrelevant here, compare with

$$B \otimes A = \begin{pmatrix} 1 \begin{pmatrix} 1 & -2 \\ 3 & 0 \end{pmatrix} & 2 \begin{pmatrix} 1 & -2 \\ 3 & 0 \end{pmatrix} & 3 \begin{pmatrix} 1 & -2 \\ 3 & 0 \end{pmatrix} \\ -1 \begin{pmatrix} 1 & -2 \\ 3 & 0 \end{pmatrix} & 0 \begin{pmatrix} 1 & -2 \\ 3 & 0 \end{pmatrix} & 1 \begin{pmatrix} 1 & -2 \\ 3 & 0 \end{pmatrix} \end{pmatrix} = \begin{pmatrix} 1 & -2 & 2 & -4 & 3 & -6 \\ 3 & 0 & 6 & 0 & 9 & 0 \\ -1 & 2 & 0 & 0 & 1 & -2 \\ -3 & 0 & 0 & 0 & 3 & 0 \end{pmatrix}.$$

Note that $A \otimes B \neq B \otimes A$, but the matrices are of the same size and they contain the same *entries* (all combinations of products of entries $a_{ij}b_{kl}$ from the two matrices).

△

Let us investigate some properties of such Kronecker-products.

Proposition 8.6.3. *The Kronecker product satisfies the mixed product property:*

$$(A \otimes B)(C \otimes D) = (AC) \otimes (BD)$$

whenever AC and BD are defined.

Proof. Thinking of $(A \otimes B)(C \otimes D) = (a_{ij}B)_{ij}(c_{ij}D)_{ij}$ as a block-matrix product, all matrix-product of blocks will be BD , so

$$(A \otimes B)(C \otimes D) = (a_{ij}B)(c_{kl}D) = \left(\sum_k a_{ik}c_{kj}BD \right)_{ij} = \left((AC)_{ij}BD \right)_{ij} = (AC) \otimes (BD).$$

□

It may seem confusing that we use the same symbol \otimes for the Kronecker-product of two matrices as for the tensor product of linear maps defined above. But the two are directly related:

Proposition 8.6.4. *Let $F : U \rightarrow V$ and $G : U' \rightarrow V'$ be linear maps between finite-dimensional vector spaces. Pick bases $\mathcal{B}_U = (u_1, \dots, u_m)$ of U and $\mathcal{B}_{U'} = (u'_1, \dots, u'_{m'})$ of U' , and pick bases $\mathcal{B}_V = (v_1, \dots, v_n)$ of V and $\mathcal{B}_{V'} = (v'_1, \dots, v'_{n'})$ of V' . Let A be the matrix of F with respect to the bases \mathcal{B}_U and \mathcal{B}_V , and let B be the matrix of G with respect to the bases $\mathcal{B}_{U'}$ and $\mathcal{B}_{V'}$. Define bases for the tensor products of vector spaces*

$$\begin{aligned} \mathcal{B}_{U \otimes U'} &= (u_1 \otimes u'_1, u_1 \otimes u'_2, \dots, u_1 \otimes u'_{m'}; u_2 \otimes u'_1, u_2 \otimes u'_2, \dots, u_2 \otimes u'_{m'}; \dots), \\ \mathcal{B}_{V \otimes V'} &= (v_1 \otimes v'_1, v_1 \otimes v'_2, \dots, v_1 \otimes v'_{n'}; v_2 \otimes v'_1, v_2 \otimes v'_2, \dots, v_2 \otimes v'_{n'}; \dots), \end{aligned}$$

Then the Kronecker product $A \otimes B$ is the matrix of $F \otimes G$ with respect to the bases $\mathcal{B}_{U \otimes U'}$ and $\mathcal{B}_{V \otimes V'}$ above.

Proof. With respect to the bases above, the coordinate-matrices (columns) of vectors in $U \otimes U'$ satisfy

$$[u \otimes u']_{\mathcal{B}_{U \otimes U'}} = [u]_{\mathcal{B}_U} \otimes [u']_{\mathcal{B}_{U'}}, \quad \text{for all } u \in U \text{ and } u' \in U',$$

where the order in which we enumerated the basis in the proposition was used. Note here that the right side is the Kronecker-product of two column matrices, so both sides is a column of height mm' . The analogous rule $[v \otimes v']_{\mathcal{B}_{V \otimes V'}} = [v]_{\mathcal{B}_V} \otimes [v']_{\mathcal{B}_{V'}}$ holds for all $v \in V$ and $v' \in V'$. So with respect to the 6 different bases above we have

$$\begin{aligned} [F \otimes G]_{\mathcal{B}_{V \otimes V'}, \mathcal{B}_{U \otimes U'}} [u \otimes u']_{\mathcal{B}_{U \otimes U'}} &= [(F \otimes G)(u \otimes u')]_{\mathcal{B}_{V \otimes V'}} = [F(u) \otimes G(u')]_{\mathcal{B}_{V \otimes V'}} \\ &= [F(u)]_{\mathcal{B}_V} \otimes [G(u')]_{\mathcal{B}_{V'}} = [F]_{\mathcal{B}_V, \mathcal{B}_U} [u]_{\mathcal{B}_U} \otimes [G]_{\mathcal{B}_{V'}, \mathcal{B}_{U'}} [u']_{\mathcal{B}_{U'}} = A [u]_{\mathcal{B}_U} \otimes B [u']_{\mathcal{B}_{U'}} \\ &= (A \otimes B) ([u]_{\mathcal{B}_U} \otimes [u']_{\mathcal{B}_{U'}}) = (A \otimes B) [u \otimes u']_{\mathcal{B}_{U \otimes U'}} \end{aligned}$$

Therefore $[F \otimes G]_{\mathcal{B}_{V \otimes V'}} = A \otimes B$.

□

We shouldn't mix up the Kronecker product of column vectors with the tensor products of vectors. For example, the tensor product $\begin{pmatrix} 1 \\ 2 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ 4 \end{pmatrix}$ is an element of $\mathbb{R}^2 \otimes \mathbb{R}^2$, while the Kronecker product of the two matrices is a 4×1 -matrix $(3, 4, 6, 8)^T$.

Using the correspondence between linear maps and Kronecker product it is easy to prove some additional properties of the Kronecker product:

Proposition 8.6.5. *The Kronecker product satisfies*

1. $(A \otimes B)(C \otimes D) = (AC) \otimes (BD)$
2. $(\lambda A) \otimes B = \lambda(A \otimes B) = A \otimes (\lambda B)$
3. $(A + B) \otimes C = A \otimes C + B \otimes C$
4. $A \otimes (B + C) = A \otimes B + A \otimes C$
5. $(A \otimes B) \otimes C = A \otimes (B \otimes C)$
6. $(A \otimes B)^T = A^T \otimes B^T$
7. $\overline{A \otimes B} = \overline{A} \otimes \overline{B}$
8. $(A \otimes B)^* = A^* \otimes B^*$
9. $I_m \otimes I_n = I_{mn}$

for all matrices A, B, C and scalars λ where the corresponding operations are defined.

Proof. We already proved 1. Properties 2-5 follows from the corresponding properties for tensor products of linear maps. Number 6-9 are easy to check and are left as an exercise. \square

Spectral properties

How are eigenvalues of F , G , and $F \otimes G$ related? Let $F : U \rightarrow U$ and $G : V \rightarrow V$ be linear maps, and assume that

$$F(u) = \lambda u \quad \text{and} \quad G(v) = \mu v.$$

Then we get

$$(F \otimes G)(u \otimes v) = F(u) \otimes G(v) = (\lambda u) \otimes (\mu v) = \lambda \mu (u \otimes v).$$

This shows that if u and v are eigenvectors of F and G respectively, then $u \otimes v$ is an eigenvector of $F \otimes G$ where the eigenvalue is the product of the corresponding eigenvalues of F and of G .

In fact there can be no other eigenvalues of $F \otimes G$:

Proposition 8.6.6. *Let $F : U \rightarrow U$ and $G : V \rightarrow V$ be linear maps. The spectrum of $F \otimes G$ is*

$$\sigma(F \otimes G) = \{\lambda \mu \mid \lambda \in \sigma(F), \mu \in \sigma(G)\}.$$

Moreover, let $\lambda_1, \dots, \lambda_m$ be the eigenvalues of F repeated according to algebraic multiplicity, and let μ_1, \dots, μ_n be the eigenvalues of G repeated according to algebraic multiplicity.

Then the algebraic multiplicity of an eigenvalue γ for $F \otimes G$ is the number of ways to write $\gamma = \lambda_i \mu_j$ where for $1 \leq i \leq m$ and $1 \leq j \leq n$.

If F and G are both diagonalizable, then so is $F \otimes G$.

Proof. Let $m = \dim U$ and $n = \dim V$. Pick two Jordan-bases of U and of V such that $[F] = A$ and $[G] = B$ where A and B both are on Jordan form. Let A have the elements $(\lambda_1, \dots, \lambda_m)$ on the diagonal; these elements are the eigenvalues of A repeated according to algebraic multiplicity. Similarly, let B have (μ_1, \dots, μ_n) on the diagonal. Then with respect to these bases, the matrix of the operator $F \otimes G$ will be the upper triangular Kronecker product $A \otimes B = (a_{ij}B)$, and the diagonal blocks of this matrix will be $(\lambda_1 B, \lambda_2 B, \dots, \lambda_m B)$. This means that the diagonal of $A \otimes B$ will be

$$(\lambda_1 \mu_1, \dots, \lambda_1 \mu_n; \lambda_2 \mu_1, \dots; \lambda_2 \mu_n; \dots \lambda_m \mu_1, \dots, \lambda_m \mu_n).$$

Since $A \otimes B$ is upper triangular, its eigenvalues are the diagonal elements repeated according to algebraic multiplicity, so the statement in the proposition follows. \square

8.7 Tensors in neural networks

Basic setup


Let's start with a brief and informal description of a very simple neural network setup. The idea is to simulate a brain consisting of *neurons* that are interconnected, each connection between neurons has a certain *weight*, the "strength" of the connection. When a neuron *activates* it fires a signal to all neurons it connects to, with signal strength proportional to the weighted sum of its input signals, and if a neuron receives activation signals larger than some threshold, it too activates.

The basic setup for modelling is to form a number of layers of neurons where each neuron connects to every neuron in the next layer. The leftmost layer is called the *input layer* and the rightmost layer is called the *output layer*, in between these there can exist any number of *hidden layers* of various size. All signals travel from left to right.

We start with an input, a sequence of real numbers for the input neurons. Then we use the weights to compute the sum of the inputs to each neuron in the second layer. If the value at a neuron doesn't reach the threshold we set its value to zero. Then we repeat this process for each layer until we reach the output layer.

As a whole, the network transforms input vectors to output vectors, and by tuning the weights (strength of the neuron connections) we can get the network to perform quite advanced functions. To tune the weights one can *train* the network by running the network on inputs where we know the desired output (training data), and at each step we tune the weights a tiny bit so that the output better matches the inputs. This is done using the usual chain rule of multivariate calculus.

Let us look at a concrete example: Here is a network:




In this example, zero is the threshold for neuron activation. Let's run the network for the input vector (1, 2, 3). We input these values at the left nodes and compute the second layer by forming linear combinations of the inputs according to the weights, then the values at the middle layer becomes:

$$(0.7, -0.1, 0.6, -0.2).$$

Here the second and fourth node didn't reach our threshold 0 so we set these values to 0, so our node values at the middle layer is (0.7, 0, 0.6, 0).

We repeat the process to compute the outputs

$$(0.23, 0.10).$$



Here we observe that the computation we performed when going from one layer to the next was really just matrix-multiplication by the matrix whose coefficients are the weights of the network, followed by a cutoff-function. The matrices in this example are:

$$A_1 = \begin{pmatrix} 0.2 & 0.1 & 0.1 \\ 0.1 & 0.2 & -0.2 \\ -0.5 & 0.4 & 0.1 \\ 0.1 & 0.0 & -0.1 \end{pmatrix} \quad \text{and} \quad A_2 = \begin{pmatrix} 0.5 & -0.3 & -0.2 & -0.2 \\ -0.2 & -0.3 & 0.4 & -0.2 \end{pmatrix}.$$

Several variations of this setup are possible. Naturally, layers can have any size and there can be any number of layers between the input and output. It is also common to add a vector b , a "bias" to each layer after forming the linear combinations, the parameters of the vectors b can also be tuned in the training process. Additionally, the non-linear threshold function σ can vary, this is the function that was applied to the vector after taking the linear combinations. In our example we used $\sigma(x) = \max(0, x)$ for each entry.

Let's summarize by writing down a general definition of a basic neural network in this setup:

Definition 8.7.1. A **basic neural network** is a function $F : \mathbb{R}^N \rightarrow \mathbb{R}^M$ which is a composition of functions of form $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ with $f(v) = \sigma(Av + b)$, where A is an $m \times n$ -matrix, $b \in \mathbb{R}^m$, and where $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a nonlinear function, and $\sigma(v) := (\sigma(v_1), \dots, \sigma(v_m))$.

So a neural network is really just a composition of linear (or affine) maps with some cut-off function applied after each step. Common choices for σ includes the *ReLU-function* $\sigma(x) = \max(0, x)$ or a sigmoid function $\sigma(x) = (1 + e^{-x})^{-1}$, a smooth function from \mathbb{R} to $[0, 1]$.

To make things more explicit, let F is a composition of q functions

$$F = f_q \circ \dots \circ f_1$$

so that the network has $q + 1$ layers. Let (n_0, n_1, \dots, n_q) be the number of neurons in each layer. Then for each $1 \leq i \leq q$ we have a function

$$f_i : \mathbb{R}^{n_{i-1}} \rightarrow \mathbb{R}^{n_i} \quad \text{given by} \quad f_i(v) = \sigma_i(A_i v + b_i),$$

where $A_i \in \text{Mat}_{n_i \times n_{i-1}}(\mathbb{R})$ and $b_i \in \mathbb{R}^{n_i}$, and where σ_i is a nonlinear functions (usually defined coordinate-wise, and usually the same throughout the network).

Multidimensional data

Now our input at the leftmost layer may be some two-dimensional data such as grayscale image, or some three-dimensional data like a color-image, or some four-dimensional data such as a color-video. We could vectorize our data to a vector in \mathbb{R}^n for some very large n , but then we loose information about the topology of the data³. To keep the data-structure it is then natural to represent the data being forwarded throughout the network as *tensors*.

³Meaning the information about what pixels are close to each other, it may be important to keep this structure if we apply some transformations such as filters to our data throughout the network.

For example, say that we input an $n \times n$ -pixel grayscale image. This image is a tensor $v \in \mathbb{R}^n \otimes \mathbb{R}^n$. Assume for simplicity that all layers have the same size so that the signal in each layer correspond to a tensor $\mathbb{R}^n \otimes \mathbb{R}^n$. Then when moving from each layer to the next, we should apply a linear map $A : \mathbb{R}^n \otimes \mathbb{R}^n \rightarrow \mathbb{R}^n \otimes \mathbb{R}^n$, but we know that

$$\text{Hom}(\mathbb{R}^n \otimes \mathbb{R}^n, \mathbb{R}^n \otimes \mathbb{R}^n) \simeq (\mathbb{R}^n)^* \otimes (\mathbb{R}^n)^* \otimes \mathbb{R}^n \otimes \mathbb{R}^n \simeq \mathbb{R}^n \otimes \mathbb{R}^n \otimes \mathbb{R}^n \otimes \mathbb{R}^n = (\mathbb{R}^n)^{\otimes 4}$$

so each transition map A between layers can be viewed as a tensor of order 4, or more concretely as a 4-dimensional array. We also add a constant tensor $b \in \mathbb{R}^n \otimes \mathbb{R}^n$ in each layer and apply a threshold-function σ .

We may of course want to use different number of layers of different sizes, but as a whole, the network can be described as a collection of tensors. Transferring higher dimensional data through a network requires higher order tensors.

Filtering images

If we have an image represented as a matrix we can apply a *filter* to it, this is done picking a smaller matrix K called a *kernel*⁴, and sliding it across every position of the image. At each position (i, j) of our filter we compute the dot-product of the kernel by the pixels it covers, and we put the result at the corresponding position (i, j) of our output matrix $F_K(A)$. For example, let us take the image matrix A and the kernel K_y below to produce the output $F_{K_y}(A)$:

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 \end{pmatrix} \quad K_y = \begin{pmatrix} 1 & -1 \\ 1 & -1 \end{pmatrix} \quad \rightarrow \quad F_{K_y}(A) = \begin{pmatrix} 0 & 1 & -1 & 0 \\ 1 & 1 & -1 & -1 \\ 2 & 0 & 0 & -2 \\ 2 & 0 & 0 & -2 \end{pmatrix}$$

We see that if K_y is over a block of four equal pixels it will produce 0, and in fact, if K_y is over a block where the two top pixels are equal and the two bottom pixels are equal, it will also produce zero. So K_y really only detects *vertical lines segments* in the image, when the filter is over such lines it will produce, nonzero values, and larger positive or negative outputs correspond to sharper lines.

The filter F_{K_x} with kernel $K_x = \begin{pmatrix} 1 & 1 \\ -1 & -1 \end{pmatrix}$ detects horizontal lines instead.

Similarly, consider the same image but with a different filter:

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 \end{pmatrix} \quad K_a = \frac{1}{4} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \rightsquigarrow F_{K_a}(A) = \begin{pmatrix} 1 & 0.75 & 0.75 & 1 \\ 0.75 & 0.25 & 0.25 & 0.75 \\ 0.5 & 0 & 0 & 0.5 \\ 0.5 & 0 & 0 & 0.5 \end{pmatrix}$$

Here the filter F_{K_a} takes the average value of the four pixels it covers, so the result is a blurred version of the original image⁵.

So different filters can detect different *features* of the image. Each filter is a 4-tensor, a linear map that transforms a matrices to a matrices (2-tensors to 2-tensors). For example, the three filters F_x, F_y, F_a described above are all linear maps $\mathbb{R}^5 \otimes \mathbb{R}^5 \rightarrow \mathbb{R}^4 \otimes \mathbb{R}^4$, so we can view each of them as a 4-tensors, elements of $\mathbb{R}^5 \otimes \mathbb{R}^5 \rightarrow \mathbb{R}^4 \otimes \mathbb{R}^4$.

We can also construct a linear map that extracts all the three features of our image at once and puts the result in a tensor 3-tensor. Viewing each of the maps $F_{K_x}, F_{K_y}, F_{K_a}$ above as 4-tensors, we can form

$$F = F_{K_x} \otimes e_1 + F_{K_y} \otimes e_2 + F_{K_a} \otimes e_3,$$

a 5-tensor that produces an three images of horizontal lines, vertical lines, and a blurred version of the image, and stacks these three images together in a 3-tensor.

If we instead start with a 5×5 color-image, an order 3 tensor of dimensions $5 \times 5 \times 3$, then filters can act differently on each color-layers, or they could even take linear combinations of the pixels from

⁴The meaning here is different from the usual kernel $\ker(A)$ of linear algebra.

⁵You may note that we lose one pixel in height and width when applying a filter, but this can be circumvented by padding the original matrix, there are several variations of this.

different layers, so the kernels also needs to be an order 3 tensors of size $d \times d' \times 3$ where $d \times d'$ is the size of the "sliding window". For example, the kernel $K = K_a \otimes (e_1 + e_2 + e_3)$ is a 3-tensor with dimensions $2 \times 2 \times 3$. The corresponding map F_K then takes the average brightness in each layer separately, so the image of F_K is a $5 \times 5 \times 3$ -image where each layer is obtained by blurring each color channel separately, this resulting tensor correspond to the blurred version of the original image. Here the filter F_K maps 3-tensors to 3-tensors, so F_K can be viewed as a 6-tensor of dimensions $5 \times 5 \times 3 \times 4 \times 4 \times 3$.

Convolutional neural networks

In a **Convolutional neural network** (CNN), our data being transferred through the network is represented as 3-tensors. In each step we apply a number of filters to the previous tensor, where the number of filters is the *depth* of the input tensor. For example, say that our input is a 5×5 -pixel color-image represented as a tensor X of dimensions $5 \times 5 \times 3$. Then to move to the next layer of our network we have a number of filters F_1, \dots, F_p where each corresponding kernel K_i is a $d \times d \times 3$ -tensor. Then the full transition map from the first layer to the second is $F = F_1 \otimes e_1 + F_2 \otimes e_2 + \dots + F_p \otimes e_p$, which is a 6-tensor. Usually such operations are combined with non-linear threshold functions σ , and *pooling layers* that reduce the dimension of the tensor being transmitted.

So a convolutional neural network can essentially be described by a certain set of 6-tensors corresponding to the linear transformation between layers, a set of 3-tensors corresponding to biases we add at each layer, and a set of nonlinear functions σ .

Thanks to this architecture of CNNs they can detect hierarchy's of spatial relationships within an image which turns out to be useful for image recognition. Several variations of the setup for CNNs have been studied.

Chapter 9

Appendix

9.1 Fields

Definition

A field is a number system where you can add, subtract, multiply, and divide numbers with each other to produce new numbers. The numbers themselves may be any type of objects as long as they are subject to a set of rules, these rules guarantee that the numbers interact similarly to how the real numbers do.

Definition 9.1.1. A field is a set \mathbb{F} equipped with two operations $\mathbb{F} \times \mathbb{F} \rightarrow \mathbb{F}$ written like normal addition and multiplication, such that the following **field-axioms** hold for all $a, b, c \in \mathbb{F}$:

(F1) $a + b = b + a$ *(Commutativity of addition)*

(F2) $(a + b) + c = a + (b + c)$ *(Associativity of addition)*

(F3) There exists $0 \in \mathbb{F}$ such that $a + 0 = a$ for all $a \in \mathbb{F}$ *(Additive identity)*

(F4) For all $a \in \mathbb{F}$ there exists $(-a) \in \mathbb{F}$ such that $a + (-a) = 0$ *(Additive inverses)*

(F5) $a \cdot b = b \cdot a$ *(Commutativity of multiplication)*

(F6) $(a \cdot b) \cdot c = a \cdot (b \cdot c)$ *(Associativity of multiplication)*

(F7) There exists an element $1 \in \mathbb{F}$ such that $a \cdot 1 = a$ for all $a \in \mathbb{F}$ *(Multiplicative identity)*

(F8) For all $a \neq 0$ in \mathbb{F} there exists an $a^{-1} \in \mathbb{F}$ such that $a \cdot a^{-1} = 1$ *(Multiplicative inverses)*

(F9) $a \cdot (b + c) = (a \cdot b) + (a \cdot c)$ *(Distributivity of multiplication over addition)*

Additionally we require that $0 \neq 1$.

Some convention of notation: as usual we write just ab instead of $a \cdot b$, and $a + bc$ means $a + (bc)$. We can always subtract elements of a field, by $a - b$ we mean $a + (-b)$. We can also divide by nonzero elements, by $\frac{a}{b}$ we mean $a \cdot b^{-1}$ whenever $b \neq 0$. By a^n we mean $a \cdot a \cdots a$ as usual, the product of n copies of a . If all the axioms except possibly (F5) and (F8) hold, then F is called a **ring**.

Example 9.1.2.

- \mathbb{Q} , \mathbb{R} , and \mathbb{C} - the rational, real, and complex numbers are all examples of fields.
- The integers \mathbb{Z} is not a field since all numbers except ± 1 have no multiplicative inverse in \mathbb{Z} .
- The set of 2×2 -matrices $\text{Mat}_n(\mathbb{C})$ is not a field, because not all matrices commute, and because some nonzero matrices do not have inverses.
- The set of real rational functions $\mathbb{R}(x) = \{\frac{p(x)}{q(x)} \mid p, q \in \mathbb{R}[x], q \neq 0\}$ is a field.

△

Integers modulo n

Fix a positive integer n and let \mathbb{Z}_n be the set of integers modulo n . As a set we can write

$$\mathbb{Z}_n = \{0, 1, 2, \dots, n - 1\},$$

and these numbers can be added and multiplied in the standard way, except that we reduce the result modulo n .¹

Proposition 9.1.3. *Then for any $n > 1$, \mathbb{Z}_n is a ring. \mathbb{Z}_n is a field if and only if n is a prime number.*

Proof. All field axioms except (F8) follows from the corresponding properties in \mathbb{Z} . If p is prime and $m \in \mathbb{Z}_p$ with $1 \leq m \leq p - 1$, then the greatest common divisor of m and p is 1, which means that we can express $1 = a \cdot m + b \cdot p$ for some integers a and b . But then in \mathbb{Z}_p this equality says that $a \cdot m = 1$, so $m^{-1} = a$. This shows that every nonzero element in \mathbb{Z}_p has a multiplicative inverse. On the other hand, if n is not prime it has a factorization $n = ab$ with $1 < a, b < n$. But then $a \cdot b = 0$ in \mathbb{Z}_p , so in \mathbb{Z}_p we have

$$b = 1 \cdot b = (a^{-1} \cdot a) \cdot b = a^{-1} \cdot (a \cdot b) = a^{-1} \cdot 0 = 0,$$

which is a contradiction, so \mathbb{Z}_n is not a field if n is composite. □

In a finite field, the addition and multiplication can be completely expressed by its addition- and multiplication-tables.

Example 9.1.4. The finite field \mathbb{Z}_5 consists of the elements 0, 1, 2, 3, 4 where addition and multiplication are defined by the following two tables:

+	0	1	2	3	4	0	1	2	3	4	
0	0	1	2	3	4	·	0	1	2	3	4
1	1	2	3	4	0	0	0	0	0	0	
2	2	3	4	0	1	1	0	1	2	3	4
3	3	4	0	1	2	2	0	2	4	1	3
4	4	0	1	2	3	3	0	3	1	4	2
						4	0	4	3	2	1

Here is an example of a calculation in \mathbb{Z}_5 :

$$2 - 4 + \frac{2}{3} = 2 + 1 + 2 \cdot 3^{-1} = 3 + 2 \cdot 2 = 3 + 4 = 2.$$

Here we used that in \mathbb{Z}_5 we have $-4 = 1$ and that $3^{-1} = 2$ since $3 \cdot 2 = 1$. △

Since finite fields allow exact computation on a computer, they are useful in applications such as cryptography, signal processing, and error correcting codes.

Example 9.1.5. A common example in applications is the finite field \mathbb{Z}_2 consists of the elements 0, 1 where addition and multiplication are defined by the following two tables:

+	0	1	0	1	
0	0	1	·	0	1
1	1	0	0	0	0
			1	0	1

Multiplication, addition, and subtraction are easy to compute in \mathbb{Z}_p , we just do the corresponding operation in \mathbb{Z} and reduce modulo p . Division however can be more tricky computationally. Since $\frac{a}{b} = a \cdot b^{-1}$, this boils down to finding the multiplicative inverse of elements in \mathbb{Z}_p . For low primes p we can find b^{-1} by trial and error, by computing $a \cdot b \pmod p$ for $a = 1, 2, 3, \dots$ until we find a product that is 1 modulo p . △

For general p we can use the Euclidean algorithm, we illustrate with an example:

¹More formally, one can define $\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$, the quotient of the integers by the subset $n\mathbb{Z}$ of integer multiples of n . The elements of \mathbb{Z}_n are then technically subsets of \mathbb{Z} of form $\{a + kn \mid k \in \mathbb{Z}\}$, more details can be found in a book on introductory abstract algebra.

Example 9.1.6. Let us find 12^{-1} in \mathbb{Z}_{41} . By dividing 41 by 12 we see that $41 = 12 \cdot 3 + 5$, and by dividing 12 by 5 we see that $12 = 5 \cdot 2 + 2$, and dividing 5 by 2 we get $5 = 2 \cdot 2 + 1$. We can use these three equalities to write 1 as an integer combination of 41 and 12:

$$1 = 5 - 2 \cdot 2 = 5 - 2 \cdot (12 - 5 \cdot 2) = 5 \cdot 5 - 2 \cdot 12 = 5 \cdot (41 - 12 \cdot 3) - 2 \cdot 12 = 5 \cdot 41 - 17 \cdot 12.$$

But this means that in \mathbb{Z}_{41} we have $1 = (-17) \cdot 12 = 24 \cdot 12$ so $12^{-1} = 24$ in \mathbb{Z}_{41} .

△

In general, to find the inverse of b in \mathbb{Z}_p , start by dividing p by b , and then repeatedly divide the divisor by the remainder until the remainder is 1. Then use the obtained equalities backwards to express $1 = x \cdot b + y \cdot p$ for integers x, y . Then in \mathbb{Z}_p , the inverse of b is x modulo p .

The following famous theorem is often useful when doing calculations in \mathbb{Z}_p :

Theorem 9.1.7. (*Fermat's little theorem*)

When x is any integer and p is a prime we have

$$x^p \equiv x \pmod{p}.$$

In other words, we have $a^p = a$ for all $a \in \mathbb{Z}_p$.

Example 9.1.8. Let us give an example of how to solve a system of linear equation where all coefficients and variables are elements of \mathbb{Z}_7 . We can proceed with Gaussian elimination as usual if we just remember to perform all arithmetic in the field \mathbb{Z}_7 . We have

$$\begin{cases} 3x + y + z = 1 \\ 2x + 5y + z = 2 \end{cases} \Leftrightarrow \begin{cases} 3x + y + z = 1 \\ 2y + 5z = 1 \end{cases} \Leftrightarrow \begin{cases} 3x + 2z = 4 \\ 2y + 5z = 1 \end{cases} \Leftrightarrow \begin{cases} x + 3z = 6 \\ y + 6z = 4 \end{cases}$$

where we used the following steps: first we added 4 times the first equation to the second, then we added 3 times the second equation to the first, then we multiplied the first equation by $3^{-1} = 5$ and the second by $2^{-1} = 4$. We now see that for each value of $t \in \mathbb{Z}_5$, setting $z = t$ gives us $x = 6 - 3t = 6 + 4t$ and $y = 4 - 6t = 4 + t$, so the set of solutions is

$$(x, y, z) \in \{(6 + 4t, 4 + t, t) \mid t \in \mathbb{Z}_7\} = \{(6, 4, 0), (3, 5, 1), (0, 6, 2), (4, 0, 3), (1, 1, 4), (5, 2, 5), (2, 3, 6)\}.$$

Note that even though we had one free parameter we only got finitely many solutions, since \mathbb{Z}_7 is a finite field.

△

Abstract fields

The elements of a field F need not be "numbers" in the traditional sense as the following example illustrates.

Example 9.1.9. There is a field F with four elements $F = \{\heartsuit, \spadesuit, \clubsuit, \diamondsuit\}$ where the operations are determined by the following tables:

+		\clubsuit	\heartsuit	\spadesuit	\diamondsuit
\clubsuit		\clubsuit	\heartsuit	\spadesuit	\diamondsuit
\heartsuit		\heartsuit	\clubsuit	\diamondsuit	\spadesuit
\spadesuit		\spadesuit	\diamondsuit	\clubsuit	\heartsuit
\diamondsuit		\diamondsuit	\spadesuit	\heartsuit	\clubsuit

·		\clubsuit	\heartsuit	\spadesuit	\diamondsuit
\clubsuit		\clubsuit	\clubsuit	\clubsuit	\clubsuit
\heartsuit		\heartsuit	\heartsuit	\spadesuit	\diamondsuit
\spadesuit		\spadesuit	\spadesuit	\diamondsuit	\heartsuit
\diamondsuit		\diamondsuit	\diamondsuit	\heartsuit	\spadesuit

Here one can verify that field axioms (F1)-(F9) hold.

△

There is a nice classification of the finite fields:

Proposition 9.1.10. For each integer $q = p^n$ where p is a prime and n is a positive integer, there exists a field $\text{GF}(q)$ with q elements. This field is unique up to isomorphism.

The fields \mathbb{Z}_p correspond to $n = 1$ in the above proposition.

9.2 Complexification and realification

In this section of the appendix we discuss the connection between real and complex vector spaces and maps.

In Example 3.1.7 we saw how a matrix A could represent both a linear map between real vector spaces $\mathbb{R}^2 \rightarrow \mathbb{R}^2$, and a linear map between two complex vector spaces $\mathbb{C}^2 \rightarrow \mathbb{C}^2$. Let's try to formalize this idea.

Complexification

Every complex number can be written uniquely as $a + bi$ with $a, b \in \mathbb{R}$. In other words, the complex numbers \mathbb{C} is a real vector space of dimension 2 with basis $(1, i)$.

Similarly, every vector in \mathbb{C}^2 can be written uniquely as $v + iv'$ where $v, v' \in \mathbb{R}^2$. For example:

$$\begin{pmatrix} 3 + 2i \\ 5 - i \end{pmatrix} = \begin{pmatrix} 3 \\ 5 \end{pmatrix} + i \begin{pmatrix} 2 \\ -1 \end{pmatrix}.$$

Following this idea we construct the *complexification* of a real vector space by considering formal pairs of real vectors.

Definition 9.2.1. Let V be a real vector space. We define the **complexification** $V^{\mathbb{C}}$ of V to consist of all formal sums $v + iv'$ where $v, v' \in V$, where addition of such objects is defined in the natural way:

$$(v_1 + iv'_1) + (v_2 + iv'_2) := (v_1 + v_2) + i(v'_1 + v'_2),$$

and where multiplication by a complex number $a + bi$ on such an object is defined via:

$$(a + bi) \cdot (v + iv') := (av - bv') + i(av' + bv).$$

Then $V^{\mathbb{C}}$ is a complex vector space, it satisfies all the vector space axioms.

For example, the complexification of \mathbb{R}^2 is \mathbb{C}^2 , but the construction above is more general and works for any real vector space V .

Note that a basis for V is still a basis for $V^{\mathbb{C}}$, but complex coefficients are allowed in the latter space².

Now, if $F : V \rightarrow W$ is a linear map between real vector spaces, we can define a map $F^{\mathbb{C}} : V^{\mathbb{C}} \rightarrow W^{\mathbb{C}}$ by

$$F^{\mathbb{C}}(v + iv') := F(v) + iF(v').$$

This construction has the effect witnessed before: if $[F] = A$ with respect to some choice of bases in V and W , then $[F^{\mathbb{C}}] = A$ too with respect to the same bases in $V^{\mathbb{C}}$ and $W^{\mathbb{C}}$.

Realification

Is the opposite construction possible? Can we from a complex vector space construct a corresponding real one? Yes, if V is a complex vector space, let $V_{\mathbb{R}}$ be the same set as V , where addition is defined the same way, and where multiplication by a scalar λ is defined the same way as in V when $\lambda \in \mathbb{R}$, but where it is *undefined* when λ is not real. Intuitively, just take the same vector space V but "forget" how to multiply vectors by non-real complex numbers³. This construction is called the **decomplexification** or **realification** of V . If $F : V \rightarrow W$ is a linear map between complex vector spaces, then we get a corresponding linear map between real vector spaces $F_{\mathbb{R}} : V_{\mathbb{R}} \rightarrow W_{\mathbb{R}}$, where $F_{\mathbb{R}}(v) = F(v)$.

Note that if (e_1, \dots, e_n) is a basis for the complex vector space V , it means that every vector $v \in V$ can be expressed uniquely as

$$v = \lambda_1 e_1 + \lambda_2 e_2 + \dots + \lambda_n e_n = (a_1 + ib_1)e_1 + (a_1 + ib_1)e_2 + \dots + (a_n + ib_n)e_n.$$

This shows that $(e_1, ie_1, e_2, ie_2, \dots, e_n, ie_n)$ is a basis for the real vector space $V_{\mathbb{R}}$ since every v can be expressed uniquely as a linear combination of these vectors with *real* coefficients. Note that $\dim V_{\mathbb{R}} = 2 \cdot \dim V$.

²Indeed, this is an alternative way to define complexification. If (e_1, \dots, e_n) is a basis for a real vector space V , then let $V^{\mathbb{C}}$ be the complex vector space with the same basis. It consists of all *complex* linear combinations of the basis vectors.

³In category theory, the function that maps each complex vector space V to $V_{\mathbb{R}}$ is called the *forgetful functor*.

9.3 Basis in infinite dimension

The majority of this text focuses on finite-dimensional vector spaces. In this section we summarize some basic definitions and results for the infinite-dimensional case, and we prove the famous theorem that every vector space has a basis (assuming that the axiom of choice holds).

Span, linear dependence, and bases revisited

The concept of span, linear dependence, and basis are usually defined a bit different for infinite-dimensional vector spaces. In this section, let V be a vector space over a field \mathbb{F} which is not necessarily finite-dimensional.

Definition 9.3.1. Let \mathcal{B} be a subset of a vector space V .

We say that \mathcal{B} span (or generate) V if every vector $v \in V$ is a *finite* linear combination of the elements of \mathcal{B} :

$$v = \lambda_1 b_1 + \dots + \lambda_n b_n \text{ where } b_i \in \mathcal{B}, \text{ and } \lambda_i \in \mathbb{F}.$$

Defining $\text{span}(\mathcal{B})$ as the set of all *finite* linear combinations, this is equivalent to the usual condition $V = \text{span}(\mathcal{B})$.

On the other hand we say that \mathcal{B} is a linearly independent set if the only *finite* linear combination of \mathcal{B} that is zero is the trivial combination: for any finite subset $\{b_1, \dots, b_n\} \subset \mathcal{B}$ we have

$$\lambda_1 b_1 + \dots + \lambda_n b_n = 0 \Rightarrow \lambda_1 = \dots = \lambda_n = 0.$$

If \mathcal{B} is a linearly independent set that span V we say that \mathcal{B} is a **basis** for V . This is equivalent to saying that every element of V can be expressed in a unique way as a *finite* linear combination of elements from \mathcal{B} .

So the difference here is that we do not consider bases as ordered, and we always require that linear combinations are finite.

Example 9.3.2. Let $\mathcal{P}(\mathbb{R})$ be the vector space of *all* polynomials with real coefficients, and take $\mathcal{B} = (1, x, x^2, \dots)$. Then \mathcal{B} is a basis for $\mathcal{P}(\mathbb{R})$, since every polynomial is a finite linear combination of such elements (since a polynomial cannot have infinitely many terms).

On the other hand, let \mathbb{R}^∞ be the space of all infinite sequences of real numbers

$$\mathbb{R}^\infty = \{(x_1, x_2, x_3, \dots) \mid x_i \in \mathbb{R}\}$$

and let $\mathcal{B} = \{e_1, e_2, e_3, \dots\}$ where $e_k \in \mathbb{R}^\infty$ has a single 1 in position k and zeroes elsewhere. Then \mathcal{B} is a linearly independent set, but it is *not* a basis for \mathbb{R}^∞ as it does not span it. For example, the vector $(1, 1, 1, \dots) \in \mathbb{R}^\infty$ is not a *finite* linear combination of elements of \mathcal{B} .

We shall prove that \mathbb{R}^∞ in fact has a basis, but it not possible to write it down explicitly.

△

Zorn's lemma

The proof that every vector space has a basis depends on the axiom of choice, that essentially says that given an infinite collection of sets it is always possible to pick one element from each set. This may sound obvious, but it is not provable within the standard axiom system. The axiom of choice is equivalent to Zorn's lemma:

Lemma 9.3.3.

(Zorn's lemma)

If P is a partially ordered set such that every chain has an upper bound, then P contains a maximal element.

Let us try and decipher the words of the Lemma:

A partially ordered set means set P with a relation \leq such that for all $x, y, z \in P$ we have

- $x \leq x$ (reflexivity)
- $x \leq y$ and $y \leq x$ implies $x = y$ (anti-symmetry)

- $x \leq y$ and $y \leq z$ implies $x \leq z$ (transitivity)

If we would have added the condition that for any $x, y \in P$ we had either $x \leq y$ or $y \leq x$, then P would be a **totally ordered set**.

A chain in P is a totally ordered subset of $C \subset P$ (not necessarily finite), we can think of it as a set of elements satisfying

$$\cdots \leq x \leq y \leq z \leq w \leq \cdots.$$

An upper bound for a chain $C \subset P$ is an element $m \in P$ such that $m \geq c$ for all $c \in C$.

A maximal element of P is an element $M \in P$ such that the only element $p \in P$ satisfying $p \geq M$ is M itself.

The basis theorem

We now proceed to prove that every vector space has a basis (under the assumption that Zorn's lemma holds).

Theorem 9.3.4. (*The basis theorem*)

Zorn's lemma implies that every vector space has a basis.

Proof. Let V be a vector space. If $V = \{0\}$ the empty set \emptyset is a basis for V , so assume $V \neq \{0\}$.

Let P be the set of all linearly independent subsets of V . P is a partially ordered set via the subset-relation $A \leq B \Leftrightarrow A \subset B$.

We show that the condition of Zorn's lemma holds, that every chain has an upper bound. Let C be a non-empty chain in P (a set of subsets of V totally ordered under inclusion), and let $m = \bigcup_{c \in C} c$ be the union of the chain C . Then clearly m contains all elements of C as a set, so to prove that m is an upper bound for C it suffices to show that it is still a linearly independent set.

Assume for contradiction that m is linearly dependent. Then there exists a finite linear combination that is zero:

$$\sum_{i=1}^n \lambda_i v_i = 0 \text{ for some } v_i \in m = \bigcup_{c \in C} c.$$

But then since m is the union of C , for each $1 \leq i \leq n$ there is a subset $c_i \in C$ such that $v_i \in c_i$, and since C is totally ordered, the finite set $\{c_1, \dots, c_n\}$ has a maximum element c' . But then all $v_i \in c'$, and then $\sum_{i=1}^n \lambda_i v_i = 0$ is a finite-zero combination of vectors of c' , which is a contradiction since $c' \in C$ was assumed to be a linearly independent set. This contradiction proves that in fact every chain has an upper bound, and Zorn's lemma applies to the partially ordered set P : there exists a maximal element M of P .

We now prove that this maximal element M is in fact a basis for V . By definition of P , the element M is linearly independent so it remains to show that M spans V . Assume for contradiction that $w \in V$ can not be written as a finite linear combination of M . Then $M' = M \cup \{w\}$ is a linearly independent set, so $M' \in P$ and $M \leq M'$ with $M \neq M'$, which contradicts the maximality of M . This proves that our assumption that M was linearly independent was false, so M is a basis for V . \square

This result has theoretical interest, but note that the proof is non-constructive; given a vector space such as \mathbb{R}^∞ the proof gives no hint on how to actually construct a basis for it.

Chapter 10

Problems

Vector spaces, subspaces, direct sum, quotients

P 1.1. Use the vector space axioms to prove that $v + v = 2 \cdot v$ holds in any complex vector space.

P 1.2. Consider $\mathbb{R}_+ =]0, \infty[$, the set of positive real numbers. We define an addition $+$ on \mathbb{R}_+ by

$$x + y := xy.$$

We define a multiplication of real numbers λ on \mathbb{R}_+ by

$$\lambda \bullet x := x^\lambda.$$

Under this addition and scalar multiplication, \mathbb{R}_+ is in fact a vector space.

- What is the zero element (the additive identity) of the vector space V ?
- What is -5 ? (the additive inverse of the vector 5)?
- Verify that the vector space axiom $(\lambda + \mu) \bullet v = \lambda \bullet v + \mu \bullet v$ holds in V .
- Verify that $\lambda \bullet (\mu \bullet v) = (\lambda\mu) \bullet v$.

P 1.3. Which of the following subsets of \mathbb{R}^2 are subspaces? Which satisfy the additive property, and which satisfy the homogeneity property?

- $S_1 = \{(x, y) \in \mathbb{R}^2 \mid 2x + y = 3\}$
- $S_2 = \{(x, y) \in \mathbb{R}^2 \mid x = y\}$
- $S_3 = \{(x, y) \in \mathbb{R}^2 \mid xy = 0\}$
- $S_4 = \{(x, y) \in \mathbb{R}^2 \mid x \geq 0 \text{ and } y \geq 0\}$
- $S_5 = \{(x, y) \in \mathbb{R}^2 \mid x \in \mathbb{Z} \text{ and } y \in \mathbb{Z}\}$
- $S_6 = \mathbb{R}^2$
- $S_7 = \emptyset$

P 1.4. In \mathbb{R}^3 , let U be the plane $x + 2y + 3z = 0$ and let U' be the line spanned by $(1, 1, 1)$. Then $\mathbb{R}^3 = U \oplus U'$. Find the projection of $v = (1, 4, 1)$ on U with respect to this direct sum.

P 1.5. Show that $S = \{p(x) \in \mathcal{P}_3 \mid p(2) = 0\}$ is a subspace of \mathcal{P}_3 . Also find a basis for it.

P 1.6. Let S and S' be any two subspaces of a vector space V . Which of the following statements are true? For false statements, give a counterexample. For true statements, give a short proof.

- The intersection $S \cap S' := \{v \in V \mid v \in S \text{ and } v \in S'\}$ is a subspace.

b) The union $S \cup S' := \{v \in V \mid v \in S \text{ or } v \in S'\}$ is a subspace.

c) The sum $S + S' := \{u + v \mid u \in S, v \in S'\}$ is a subspace.

P 1.7. Let U_1, U_2, U_3 be three subspaces of V such that $U_1 + U_2 + U_3 = V$ and $U_1 \cap U_2 = U_1 \cap U_3 = U_2 \cap U_3 = \{0\}$. Show that we do not necessarily have $V = U_1 \oplus U_2 \oplus U_3$.

P 1.8. Let \mathcal{F} be the vector space of all functions $\mathbb{R} \rightarrow \mathbb{R}$. Let \mathfrak{e} be the set of all even functions, and let \mathfrak{o} be the set of all odd functions in \mathcal{F} .

a) Show that \mathfrak{e} and \mathfrak{o} are subspaces of \mathcal{F} .

b) Show that $\mathcal{F} = \mathfrak{e} \oplus \mathfrak{o}$.

c) Find the projection of e^x onto the subspace \mathfrak{e} .

d) Find the projection of $f(x) = \sin(x)x^{10} \arctan(x)$ onto \mathfrak{o} .

P 1.9. Which of the following maps are linear?

a) $F : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ where $F(x, y) = (y + x, 2x - 1)$.

b) $I : \mathcal{P}_n \rightarrow \mathcal{P}_{n+1}$ where $I(p(x)) = \int_0^1 p(x) dx$

c) $G : \text{Mat}_{n \times n}(\mathbb{C}) \rightarrow \mathbb{C}$ where $G(A) = \text{tr}(A)$

d) $H : \text{Mat}_{n \times n}(\mathbb{C}) \rightarrow \text{Mat}_{n \times n}(\mathbb{C})$ where $H(A) = A^T + 3A$

e) $T : \mathcal{C}(\mathbb{R}) \rightarrow \mathcal{C}(\mathbb{R})$ where $T(f(x)) = f(x + 1)$

f) $C : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ where $C(v) = v \times (1, 2, 3)$.

P 1.10. Let $T : \text{Mat}_n(\mathbb{C}) \rightarrow \text{Mat}_n(\mathbb{C})$ where $T(A) = A - A^T$. Show that T is linear and determine $\ker(T)$ and $\text{Im}(T)$.

P 1.11. Show that a linear map F is injective if and only if $\ker(F) = \{0\}$.

P 1.12. Let V be the vector spaces of all infinite sequences (a_1, a_2, \dots) where $a_i \in \mathbb{R}$; the sum and scalar action is defined coordinate-wise. Let $F : V \rightarrow V$ be the right-shifting operator $F(a_1, a_2, \dots) = (0, a_1, a_2, \dots)$, this is a linear map. Is F injective? Surjective? Does F have an inverse?

P 1.13. Let \mathcal{P}_4 be the space of polynomials with real coefficients and of degree ≤ 4 . Define a linear map $F : \mathcal{P}_4 \rightarrow \mathcal{P}_4$ by $F(p(x)) = p(x + 1)$. Find the matrix of F with respect to the standard basis of \mathcal{P}_4 . Also find the inverse of F .

P 1.14. In Lie theory an important object of study is \mathfrak{sl}_2 , the set of complex matrices with trace zero:

$$\mathfrak{sl}_2 = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{Mat}_{2 \times 2}(\mathbb{C}) \mid a + d = 0 \right\},$$

this is a subspace of $\text{Mat}_{2 \times 2}(\mathbb{C})$ with basis $(X, H, Y) = \left(\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} \right)$. Define two maps $F, G : \mathfrak{sl}_2 \rightarrow \mathfrak{sl}_2$ by

$$F(A) = HA - AH \text{ and } G(A) = XA - AX.$$

Find the matrices of F and of G with respect to the basis (X, H, Y) .

P 1.15. Let U be a subspace of a vector space V . Let $v + U$ and $w + U$ be two affine subsets. Prove first that $v + U = (v + u') + U$ whenever $u' \in U$. Use this to prove that $v + U = w + U$ whenever $v - w \in U$.

P 1.16. Let $V = \mathbb{R}^2$ and let U be the subspace spanned by $(1, 1)$. Let $A = (2, 3) + U$, $B = (1, 0) + U$, and $C = (4, 5) + U$ be affine subsets.

a) Are any of A, B, C equal?

- b) Sketch the affine subsets A, B , and $A + B$ in a picture of \mathbb{R}^2 .
- c) Show that A is a basis for V/U and express B in this basis.

P 1.17. Let $U \subset \mathbb{R}^3$ be the subspace spanned by $(1, 2, 3)$. Then the pair $(A, B) = ((1, 1, 0) + U, (0, 1, 1) + U)$ is a basis for \mathbb{R}^3/U . Express the vector $C = (1, 1, 1) + U$ of V/U in this basis.

P 1.18. Let $V = \mathbb{R}^4/\ell$ where $\ell = \text{span}(3, 2, 1, 2)$. Is

$$(e_1, e_2, e_3) = ((1, 1, 0, 0) + \ell, (0, 1, 1, 0) + \ell, (0, 0, 1, 1) + \ell)$$

a basis for V ?

P 1.19. Let $F : V \rightarrow V$ be a linear map, and let $U \subset V$ be a subspace. Show that if $F(u) = 0$ for all $u \in U$, then the map defined by $\tilde{F} : V/U \rightarrow V/U$ with $F(v + U) = F(v) + U$ is well-defined and linear. How does the matrix of F and \tilde{F} look?

P 1.20. Write $\text{Aspan}(p_1, \dots, p_n)$ for the **affine span** of vectors (or think of them as points) p_1, \dots, p_n , defined as the smallest affine subset containing all vectors p_1, \dots, p_n . Find a geometric description of each affine span below, and express it as $v + U$ for suitable vector v and subspace U .

- a) $\text{Aspan}((1, 2), (3, 4))$ in \mathbb{R}^2 .
- b) $\text{Aspan}((1, 0, 0), (0, 1, 0), (0, 0, 1))$ in \mathbb{R}^3 .

P 1.21. A map $F : V \rightarrow V$ is called an **affine map** if $F(v) = G(v) + w$ where G is a linear map and w is a fixed translation-vector.

- a) Show that the composition of affine maps is affine.
- b) Show that the map that reflects points of \mathbb{R}^2 in the line $(1 + t, t)$ is an affine map.

P 1.22. Let $X = \{1, 2, 3, 4, 5\}$ be a finite set, and let $\mathcal{P}(X)$ be the set of all subsets of X . Define an addition on $\mathcal{P}(X)$ by the symmetric difference operator

$$S_1 + S_2 := S_1 \Delta S_2 = (S_1 \cup S_2) \setminus (S_1 \cap S_2).$$

- a) How should scalar multiplication be defined in order for $\mathcal{P}(X)$ to become a vector space over \mathbb{Z}_2 ? What is the additive identity in $\mathcal{P}(X)$?
- b) Compute $\{1, 3, 5\} + \{1, 2, 3\}$
- c) Compute $-\{1, 3, 5\}$
- d) Find a basis for $\mathcal{P}(X)$
- e) Are the vectors $\{1, 3, 4\}, \{1, 2\}, \{1, 4, 5\}, \{2, 3, 4\} \in \mathcal{P}(X)$ linearly dependent?

P 1.23. Let V and W be vector spaces over \mathbb{Q} , and let $F : V \rightarrow W$ be a map satisfying $F(u + v) = F(u) + F(v)$ for all $u, v \in V$. Show that F is linear.

P 1.24. In this problem we consider vector spaces and matrices over the field \mathbb{Z}_3 , it consists of only three elements $\{0, 1, 2\}$ which can be added and multiplied modulo 3 as indicated in these tables:

+	0	1	2
0	0	1	2
1	1	2	0
2	2	0	1

·	0	1	2
0	0	0	0
1	0	1	2
2	0	2	1

- a) Compute $2(1, 2, 1, 1) + (2, 2, 0, 1)$

b) Solve the linear system $\begin{pmatrix} 1 & 2 \\ 2 & 2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$

c) Compute $\det \begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix}$

d) Find $\begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix}^{-1}$

e) How many elements does $\text{Mat}_{2 \times 2}(\mathbb{Z}_3)$ have?

P 1.25. Solve the following system of equations over the field \mathbb{Z}_5 :

$$(a) \begin{cases} x + 2y = 4 \\ 2x + y = 1 \end{cases} \quad (b) \begin{cases} x + 2y = 4 \\ 3x + y = 1 \end{cases} \quad (c) \begin{cases} x + 2y = 4 \\ 3x + y = 2 \end{cases}$$

P 1.26. Are the following vectors of $(\mathbb{Z}_3)^4$ linearly independent?

$$v_1 = \begin{pmatrix} 1 \\ 1 \\ 2 \\ 1 \end{pmatrix}, \quad v_2 = \begin{pmatrix} 1 \\ 0 \\ 1 \\ 2 \end{pmatrix}, \quad v_3 = \begin{pmatrix} 0 \\ 2 \\ 2 \\ 1 \end{pmatrix}$$

P 1.27. If F is a field, the vector space F^n consists of all n -tuples (x_1, x_2, \dots, x_n) with coefficients $x_i \in F$. For $P \in F^n$ and $v \in F^n$ we can construct the *line* in F^n through P with direction v ; it is defined as

$$\ell : \{P + tv \mid t \in F\}.$$

When $F = \mathbb{R}$ this is just the standard parameter form of the line. For other fields F the definition still makes sense even though it may be harder to visualize the line geometrically.

Now let $F = \mathbb{Z}_3$, and $n = 2$: Let ℓ_1 be the line through $P = (1, 2)$ in the direction $v = (1, 1)$. Let ℓ_2 be the line that goes through the points $Q = (1, 0)$ and $R = (2, 2)$. Do the lines intersect?

Matrices, echelon forms, LU-factorization

P 2.1. Let A and B be Hermitian matrices of the same size. Which of the matrices below are guaranteed to be Hermitian? Give a proof or a counterexample to each.

$$A + B \quad AB \quad \lambda A \quad A^T \quad AB^* + BA^*$$

P 2.2. Any complex matrix can be written uniquely as $A + Bi$ where A and B are real matrices. Prove that $A + Bi$ is Hermitian if and only if A is symmetric and B is skew-symmetric.

P 2.3. Prove that $\text{tr}(AB) = \text{tr}(BA)$ whenever both matrix products are defined. Recall that the trace of a matrix is the sum of its diagonal entries.

P 2.4.

a) Show that if an $n \times n$ -matrix A has eigenvalues $\lambda_1, \dots, \lambda_n$ (all different), then

$$\text{tr}(A) = \lambda_1 + \dots + \lambda_n.$$

b) Suppose that A is a 3×3 -matrix with 3 different eigenvalues $\lambda_1, \lambda_2, \lambda_3$. Find a formula for $\text{tr}(A^n)$.

c) Suppose that A is 3×3 and $\text{tr}(A) = 1$, $\text{tr}(A^2) = 6$, and $\text{tr}(A^3) = 10$. Find the eigenvalues of A .

P 2.5. The matrix P below is an example of what is called a *permutation matrix*. Compute P^n for each integer n .

$$P = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$$

P 2.6. A **permutation matrix** is a matrix which has a single 1 in each column and in each row.

- Prove that a permutation matrix must be square.
- Prove that the product of two permutation matrices is a permutation matrix.
- Prove that if P is a permutation matrix, then $P^n = I$ for some $n > 0$.
- We define the *order* of a permutation matrix P as the minimal positive n for which $P^n = I$. Find an example of an $n \times n$ -permutation matrix whose order is greater than its size n .

P 2.7. Let $A \in \text{Mat}_n(\mathbb{R})$ be a square matrix, and let $\mathcal{C}_A := \{B \in \text{Mat}_n(\mathbb{R}) \mid AB = BA\}$ be the **commutant** of A , the set of matrices that commute with A . Show that \mathcal{C}_A is a subspace of $\text{Mat}_n(\mathbb{R})$. Then find the commutant of $A = \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$.

P 2.8. Let $\mathcal{C} = \{A \in \text{Mat}_n(\mathbb{C}) \mid AB = BA \text{ for all } B \in \text{Mat}_n(\mathbb{C})\}$ be the set of matrices that commute with every other matrix. Describe the set \mathcal{C} explicitly.

P 2.9. Let $A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ where the coefficient lie in the finite field \mathbb{Z}_2 . Compute A^n for every integer $n \geq 0$.

P 2.10. In algebra, an object x satisfying $x \cdot x = x$ is called an **idempotent**. For example, there are exactly two idempotents in \mathbb{R} : 1 and 0. Show that there are infinitely many idempotents in $\text{Mat}_2(\mathbb{R})$.

P 2.11. Let A be a real diagonalizable 5×5 -matrix. What values are possible for the dimension of the commutant $\dim \mathcal{C}_A$?

P 2.12. Recall that matrix N is called nilpotent if $N^d = 0$ for some d . Let N and M be nilpotent matrices that commute. Show that NM and $N + M$ are both nilpotent. Is the statement still true if the matrices do not commute?

P 2.13. Let $F : V \rightarrow V$ be nonzero operator satisfying $\text{Im}(F) \subset \ker(F)$. Show that F is nilpotent, and find its nilpotency degree.

P 2.14. Show that if $N \in \text{Mat}_n(\mathbb{C})$ is nilpotent, then its nilpotency degree is $\leq n$.

P 2.15. Show that $I + N$ is invertible whenever N is nilpotent.

P 2.16. Let $M = \left(\begin{array}{c|c} 2I & 0 \\ \hline A & 3I \end{array} \right)$ be the blockmatrix of size $2n \times 2n$, where each of the blocks has size $n \times n$ and A is some given matrix. Show that $M^{-1} = \frac{1}{6} \left(\begin{array}{c|c} 3I & 0 \\ \hline -A & 2I \end{array} \right)$.

P 2.17. Find a general formula for the inverse of the block matrix $M = \left(\begin{array}{c|c} A & B \\ \hline 0 & C \end{array} \right)$, where each block has size $n \times n$, and where A and C are invertible.

P 2.18. For the matrix below, find a row echelon form and find $\text{rank}(A)$. Also find the reduced row echelon form of A , and use your result to find the nullspace $\ker(A)$.

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

P 2.19. For the matrix C below, find the reduced row echelon form, and use it to solve the linear system $CX = 0$.

$$C = \begin{pmatrix} 1 & 2i & 1+i & i \\ 2i & -4 & 1 & 3i \end{pmatrix}$$

P 2.20. A is a 3×5 matrix with $\ker(A) = \{(-2s + 3t - r, s, t, -r, 2r) \mid s, t, r \in \mathbb{R}\}$. Find the reduced row echelon form of A .

P 2.21. Find the reduced row echelon form of the matrix $A \in \text{Mat}_{3 \times 4}(\mathbb{Z}_3)$ below, and use your result to solve $AX = 0$.

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

P 2.22. Later, when determining Jordan canonical form of a matrix, we will have to find bases for various subspaces related to the matrix.

$$\text{Let } A = \begin{pmatrix} 0 & 1 & 0 & 2 & -1 \\ -1 & 1 & 1 & -1 & 0 \\ 1 & 0 & -1 & 4 & -1 \\ 0 & 0 & 0 & 0 & 0 \\ -1 & 1 & 1 & -2 & 0 \end{pmatrix}$$

- Find a basis for $\text{Im}(A)$
- Find a basis for $\ker(A)$
- Find a basis for $\ker(A) \cap \text{Im}(A)$

P 2.23. Give an example of a matrix A such that $\ker(A)$ and $\text{Im}(A)$ have a nontrivial intersection, and neither is a subset of the other

P 2.24. Find the inverse of each elementary matrix below:

$$E_1 = \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad E_2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 5 \end{pmatrix} \quad E_3 = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Then describe a general formula for the inverse of an elementary matrix.

P 2.25. Row and column operations can be achieved by matrix multiplication. Let A be a 3×3 matrix. What matrix should we multiply A by, and from what side, to have the effect of

- Adding two times the first row of A to the third row of A
- Multiplying the middle column of A by 3
- Switching the first two rows in A

P 2.26. Write $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$ as a product of elementary matrices. Then do the same for A^{-1} .

P 2.27. Recall the Gauss-Jordan method for finding the inverse of a square matrix A : Write down the block-matrix $[A|I]$ and do row operations on this matrix until it has form $[I|B]$, then $A^{-1} = B$. Prove that the Gauss-Jordan method works.

P 2.28. Let

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

Find the LU-decomposition and the LDU-decomposition of A .

P 2.29. Let

$$A = \begin{pmatrix} 1 & -1 & 1 & 1 \\ -2 & 5 & 0 & -1 \\ 3 & 3 & 1 & 2 \end{pmatrix}.$$

Find the LU-decomposition and the LDU-decomposition of A .

P 2.30. The matrix A below does not admit an LU-decomposition. Find instead a decomposition $PA = LU$ where L is lower triangular, U is upper triangular, and P is a permutation matrix.

$$A = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 2 \\ 1 & 2 & 3 \end{pmatrix}$$

P 2.31. Let $A = \begin{pmatrix} 4 & 3 \\ 5 & 1 \end{pmatrix} \in \text{Mat}_2(\mathbb{Z}_7)$, where all the coefficients lie in the finite field \mathbb{Z}_7 of integers modulo 7. Find the LU- and the LDU-decomposition of A .

P 2.32. Let U be a row echelon form of A . Prove that if the columns of A satisfy a linear dependence relation

$$\lambda_1 A_1 + \lambda_2 A_2 + \cdots + \lambda_n A_n = 0$$

then the columns of U satisfy the same relation:

$$\lambda_1 U_1 + \lambda_2 U_2 + \cdots + \lambda_n U_n = 0.$$

Conclude that the dimension of the span of the columns of A is equal to the number of pivots in U .

P 2.33. We say that a matrix A is in (reduced) **column echelon form** if and only if A^T is in (reduced) row echelon form. Describe the set of matrices that are both in reduced row echelon form and reduced column echelon form.

P 2.34. Solving a linear system by row operations involves multiplying and adding numbers (and making a few divisions too to figure out what row operations to make, and to solve the final diagonal system, but let's ignore all these). For a computer, addition is a lot faster than multiplication, so the number of multiplications is the limiting factor when solving a linear system.

- How many multiplications are required to solve a generic system $AX = b$ where A is a 4×4 -matrix with the standard Gaussian elimination algorithm? (here generic means that no unexpected zeros appear when performing the row operations)
- Assume now that $A = LU$ is an LU-factorization of the matrix above. The system $AX = b$ can then be written $L(UX) = b$, and we can solve it by solving the two triangular systems $LY = b$ and then $UX = Y$. How many multiplications are required in total?
- Assume that in (b) the matrix A is an $n \times n$ -matrix. Determine the number of multiplications needed.

P 2.35. Find the Cholesky-factorization for each of the matrices below:

$$A = \begin{pmatrix} 1 & 2 \\ 2 & 6 \end{pmatrix} \quad B = \begin{pmatrix} 1 & 2-i \\ 2+i & 9 \end{pmatrix}.$$

P 2.36. Find the Cholesky-factorization of the matrix A below.

$$A = \begin{pmatrix} 9 & 3 & -3 \\ 3 & 2 & 1 \\ -3 & 1 & 10 \end{pmatrix}$$

P 2.37. Prove that if A admits a Cholesky-factorization, then A is Hermitian. Also prove that the reverse implication does not hold.

Introductory spectral theory

P 3.1. Find the spectrum $\sigma(F)$, and the dimension of the corresponding eigenspaces for each of the linear maps $F : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ described below:

- Projection onto a line
- Reflection in a plane
- Rotation around an axis
- The identity map
- A nilpotent operator of nilpotency-degree 3

P 3.2. Find all eigenvalues and eigenvectors of the linear map on \mathbb{C}^2 given by the matrix $A = \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}$. Also diagonalize A , in other words, find matrices $S, D \in \text{Mat}_2(\mathbb{C})$ such that D is diagonal and $A = SDS^{-1}$.

P 3.3. Let \mathcal{P} be the space of all polynomials with real coefficients. Find all eigenvectors and eigenvalues of the operator $F : \mathcal{P} \rightarrow \mathcal{P}$ defined by $F(p(x)) = xp'(x)$. What is the spectrum $\sigma(F)$?

P 3.4. Show that if A is a matrix with real entries, that has a complex eigenvalue λ , then the complex conjugate $\bar{\lambda}$ is an eigenvalue too.

P 3.5. We know that a linear operator on \mathbb{C}^2 with matrix $A \in \text{Mat}_2(\mathbb{R})$ has $2 + 3i$ as an eigenvalue with corresponding eigenvector $(1, 1 + i)$. Find the matrix A .

P 3.6. Let $R : \text{Mat}_2(\mathbb{C}) \rightarrow \text{Mat}_2(\mathbb{C})$ be the linear map that rotates matrices a quarter of a turn clockwise like so:

$$R \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} c & a \\ d & b \end{pmatrix}.$$

- Find all eigenvalues and eigenvectors of R . Is R diagonalizable?
- Do the same for the corresponding map $\text{Mat}_3(\mathbb{C}) \rightarrow \text{Mat}_3(\mathbb{C})$.

P 3.7. Let V be a three-dimensional complex vector space with basis (e_1, e_2, e_3) , and let $P : V \rightarrow V$ be the linear map that permutes the basis vectors cyclically: $P(e_1) = e_2$, $P(e_2) = e_3$, $P(e_3) = e_1$. Show that P is diagonalizable and find a new basis of V consisting of eigenvectors of P .

P 3.8. Let $A = \begin{pmatrix} 1 & 2 & 3 & 4 \end{pmatrix}$. Find all eigenvalues and eigenvectors of the 4×4 -matrix $A^T A$ without writing down the matrix first.

P 3.9. How many elements does the set $\text{Mat}_{2 \times 2}(\mathbb{Z}_2)$ have? How many of them are invertible?

P 3.10. Find all eigenvalues and eigenvectors of each of the two matrices

$$A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix} \in \text{Mat}_2(\mathbb{Z}_5) \quad A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix} \in \text{Mat}_2(\mathbb{Z}_{11}).$$

Note that the matrices are different as the coefficients belong to different fields.

P 3.11. Let F be the field of four elements from Example 9.1.9 of the appendix.

- Compute the following matrix product where the coefficients lie in F : $\begin{bmatrix} \clubsuit & \heartsuit \\ \spadesuit & \diamondsuit \end{bmatrix} \cdot \begin{bmatrix} \heartsuit \\ \spadesuit \end{bmatrix}$.
- Find all eigenvalues of $A = \begin{bmatrix} \clubsuit & \heartsuit \\ \spadesuit & \diamondsuit \end{bmatrix}$.

P 3.12. An integer sequence a_n is defined recursively by $a_0 = 4, a_1 = 6$ and

$$a_n = 2a_{n-1} - 2a_{n-2} \quad \text{for } n \geq 2.$$

Find an explicit expression for a_n .

P 3.13. Compute $p(A)$ and $p(B)$ where $p(t) = t^4 + 2t^2 - 5t + 3$ and $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ and $B = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}$.

P 3.14. Let $A = \begin{pmatrix} 1 & 2 \\ 3 & 2 \end{pmatrix}$. Find a polynomial $p(t)$ for which $p(A) = 0$ by computing I, A, A^2, \dots until these matrices become linearly dependent in $\text{Mat}_n(\mathbb{R})$.

P 3.15. The Cayley-Hamilton says that $p_A(A) = 0$ where $p_A(t) = \det(A - tI)$. A famous "fake proof" of the theorem goes like this:

$$p_A(A) = \det(A - A \cdot I) = \det(0) = 0.$$

This proof is incorrect, why?

P 3.16. Verify that the Cayley-Hamilton theorem holds for $A = \begin{pmatrix} -2 & 5 & 0 \\ 0 & -2 & 0 \\ 0 & 0 & 3 \end{pmatrix}$.

P 3.17. Let $p(t) = t^2 - 4t + 3$ and $q(t) = (t - 1)^2(t - 2)^2(t - 3)^2$. Compute $p(A)$ and $q(A)$ where $A = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$.

P 3.18. We know that $p_A(\lambda) = \lambda^2 + (-3 - i)\lambda + 2 + 2i$. Find all eigenvalues of $B = A^2 + 3A - 5I$.

P 3.19. We know that a matrix A has characteristic polynomial $p_A(t) = -t^5 - 2t^4 - t^3$. What are the possible expressions for the minimal polynomial $m_A(t)$?

P 3.20. We know that a matrix A has minimal polynomial $t^2 - 1$. Simplify $A^3 + 2A^2 + 2A$.

P 3.21. Find the minimal polynomial of each matrix below.

$$\text{a) } A = \begin{pmatrix} 4 & -1 \\ 1 & 2 \end{pmatrix} \quad \text{b) } B = \begin{pmatrix} 0 & 1 & 2 & 3 \\ 0 & 4 & 5 & 6 \\ 0 & 0 & 7 & 8 \\ 0 & 0 & 0 & 9 \end{pmatrix} \quad \text{c) } C = \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 2 \end{pmatrix}$$

P 3.22. Let $T : \text{Mat}_n(\mathbb{R}) \rightarrow \text{Mat}_n(\mathbb{R})$ where T takes the transpose of a matrix: $T(A) = A^T$. Find the minimal polynomial of T .

P 3.23. Let $R : \text{Mat}_n(\mathbb{C}) \rightarrow \text{Mat}_n(\mathbb{C})$ be the linear operator that rotates a matrix M a quarter of a turn counter-clockwise: $R\left(\begin{bmatrix} \square \\ \square \\ \square \\ \square \end{bmatrix}\right) = \begin{bmatrix} \square \\ \square \\ \square \\ \square \end{bmatrix}$. Find the minimal polynomial of R .

Jordan canonical form


P 4.1. Which of the following matrices are in Jordan normal form?

$$\text{a) } \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 1 \\ 0 & 0 & 3 \end{pmatrix} \quad \text{b) } \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix} \quad \text{c) } \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}$$

P 4.2. For each Jordan-matrix below, determine the algebraic and the geometric multiplicity of each eigenvalue. Also find the characteristic polynomial and the minimal polynomial.

$$\text{a) } \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & -3 \end{pmatrix} \quad \text{b) } \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad \text{c) } \begin{pmatrix} 5 & 1 & 0 & 0 & 0 \\ 0 & 5 & 1 & 0 & 0 \\ 0 & 0 & 5 & 0 & 0 \\ 0 & 0 & 0 & 3 & 1 \\ 0 & 0 & 0 & 0 & 3 \end{pmatrix}$$


- P 4.3.** There are exactly 5^9 matrices in $\text{Mat}_3(\mathbb{Z}_5)$. How many of them are in Jordan form?
- P 4.4.** For a matrix A we know that its characteristic polynomial is $\det(A - \lambda I) = (\lambda - 3)^5(\lambda + 1)^3$. Find the number the possible Jordan forms of A (up to permutation of the blocks).
- P 4.5.** How many *different* Jordan forms of nilpotent 6×6 -matrices are there?
- P 4.6.** Find two square matrices A and B with *different* Jordan forms such that
- A and B has the same characteristic and minimal polynomials.
 - A and B has the same characteristic and minimal polynomials, and the same dimension of all the eigenspaces.
- P 4.7.** Show that for any square matrix A , the trace $\text{tr}(A)$ is the sum of the eigenvalues and $\det(A)$ is the product of the eigenvalues (counting multiplicities)
- P 4.8.** Prove that for any nilpotent map $N : \mathbb{C}^n \rightarrow \mathbb{C}^n$ we have $N^n = 0$.
- P 4.9.** Find all matrices that commute with the Jordan block $J = \begin{pmatrix} 5 & 1 & 0 \\ 0 & 5 & 1 \\ 0 & 0 & 5 \end{pmatrix}$. Generalize your result to matrices that commute with an arbitrary Jordan block.
- P 4.10.** For a certain nilpotent operator A we have $p_A(t) = t^7$, $m_A(t) = t^4$, and we know that the geometric multiplicity of the eigenvalue 0 is 3. Determine the Jordan form of A .
- P 4.11.** We know that a nilpotent linear map F on a 14-dimensional vector space has a string basis that looks like the diagram below, where each dot represents a vector in the string basis.



Find the dimension of:

- a) $\ker(F)$ b) $\text{Im}(F)$ c) $\ker(F^3)$ d) $\text{Im}(F^2)$ e) $\ker(F) \cap \text{Im}(F^2)$

- P 4.12.** Suppose $F : \mathbb{R}^6 \rightarrow \mathbb{R}^6$ is defined by the diagram below. Without using any matrices, find the characteristic and minimal polynomials for F . Also find a string basis for F .



- P 4.13.** For the nilpotent matrix N below, find a matrix S and a matrix J in Jordan form such that $SJS^{-1} = N$.

$$N = \begin{pmatrix} 1 & -1 & 0 \\ 1 & -1 & 0 \\ 1 & -1 & 0 \end{pmatrix}$$

- P 4.14.** For the nilpotent matrix M below, find a matrix S and a matrix J in Jordan form such that $SJS^{-1} = M$.

$$M = \begin{pmatrix} 0 & 0 & -1 \\ 2 & -1 & -3 \\ -2 & 1 & 1 \end{pmatrix}$$

P 4.15. Jordanize the matrix A below. In other words, find a matrix J in Jordan form and an invertible matrix S such that $S^{-1}AS = J$.

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

P 4.16. Jordanize the matrix A below. In other words, find a matrix J in Jordan form and an invertible matrix S such that $S^{-1}AS = J$.

$$A = \begin{pmatrix} 1 & 0 & 1 & 0 \\ 3 & 1 & -2 & 1 \\ -2 & 0 & 4 & 0 \\ 3 & -1 & -2 & 3 \end{pmatrix}$$

P 4.17. Jordanize the matrix A below. Its only eigenvalue is 2.

$$A = \begin{pmatrix} 1 & 1 & 5 & -1 & 7 \\ -1 & 3 & 9 & 4 & 10 \\ 0 & 0 & 0 & -2 & -2 \\ 0 & 0 & 1 & 3 & 1 \\ 0 & 0 & 1 & 1 & 3 \end{pmatrix}.$$

P 4.18. Jordanize the matrix

$$A = \begin{pmatrix} 9+i & 9 \\ -4 & -3+i \end{pmatrix}.$$

P 4.19. Find a non-diagonalizable matrix 2×2 -matrix A for which $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ is an eigenvector of eigenvalue 3.

P 4.20. Let $F: \mathcal{P}_3 \rightarrow \mathcal{P}_3$ be the shifting operator $F(p(x)) = p(x+1)$. Find the Jordan form of F .

P 4.21. Prove that a square matrix A is invertible if and only if 0 is not an eigenvalue of A .

P 4.22. Prove that if two matrices A and B have the same Jordan form J , then A and B are similar.

P 4.23. Let U be the subspace of $\mathcal{C}(\mathbb{R})$ generated by $(e^x, xe^x, x^2e^x, \sin(x), \cos(x))$, and let D be complexification (see Section 9.2 of the appendix) of the differentiation operator acting on U . Find the Jordan form of D .

P 4.24. Prove that if A is an operator whose minimal polynomial has simple zeros (the multiplicity of each zero is 1), then A is diagonalizable.

P 4.25. Prove that for $n \geq 2$, the Jordan block $J_n(0)$ does not have a square root: no matrix X satisfies $X^2 = J_n(0)$.

P 4.26. Suppose a matrix A has the Jordan form J below. Find the Jordan form of A^n for each $n > 0$.

$$J = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 0 & 2 \end{pmatrix}$$

P 4.27. Find the Jordan form of $\begin{pmatrix} 2 & 0 & 1 \\ 1 & 1 & 2 \\ 0 & 0 & 2 \end{pmatrix} \in \text{Mat}_3(\mathbb{Z}_3)$.

P 4.28. Find a matrix in $\text{Mat}_2(\mathbb{Z}_2)$ which does not admit a Jordan-decomposition, meaning that $A \neq SJS^{-1}$ for any $S, J \in \text{Mat}_2(\mathbb{Z}_2)$.

P 4.29. Compute A^n for the Jordan matrix

$$A = \begin{pmatrix} 2 & 1 & 0 & 0 & 0 \\ 0 & 2 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & -6 & 1 \\ 0 & 0 & 0 & 0 & -6 \end{pmatrix}.$$

P 4.30. Compute e^A and e^B for each of the matrices below.

$$A = \begin{pmatrix} -3 & 0 & 0 \\ 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad B = \begin{pmatrix} 0 & 1 & 2 \\ 0 & 0 & 3 \\ 0 & 0 & 0 \end{pmatrix}$$

P 4.31. Compute e^J where J is the matrix below.

$$J = \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix}$$

P 4.32. Suppose that v is an eigenvector of A with eigenvalue λ . What can be said about eigenvalues and vectors of the matrix e^A ?

P 4.33. Compute $\sin(A_k)$ and $\cos(A_k)$ for $k = 1, 2, 3$ where

$$A_1 = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad A_2 = \begin{pmatrix} 2 & 1 & 0 & 0 \\ 0 & 2 & 1 & 0 \\ 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 2 \end{pmatrix} \quad A_3 = \begin{pmatrix} \frac{\pi}{3} & 1 & 0 & 0 \\ 0 & \frac{\pi}{3} & 0 & 0 \\ 0 & 0 & \frac{\pi}{2} & 0 \\ 0 & 0 & 0 & \pi \end{pmatrix}.$$

P 4.34. Prove that $\frac{d}{dt} \sin(At) = A \cos(At)$ and that $\frac{d}{dt} \cos(At) = -A \sin(At)$. Find a differential equation that can be solved using this fact.

P 4.35. Show that if $\text{tr}(A) = 0$, then $\det(e^A) = 1$.

P 4.36. A discrete dynamical system evolves according to the model

$$\begin{cases} a_{n+1} = -2a_n + b_n \\ b_{n+1} = -a_n - 4b_n \end{cases} \quad \text{where } a_0 = 2 \text{ and } b_0 = 0.$$

Find explicit expressions for a_n and b_n and determine the limit of $\frac{a_n}{b_n}$ as $n \rightarrow \infty$.

P 4.37. Jordanize the matrix below by finding S and J such that $SJS^{-1} = A$.

$$A = \begin{pmatrix} 1 & 1 & 0 \\ -1 & 3 & 0 \\ -4 & 7 & -1 \end{pmatrix}$$

P 4.38. Solve the initial value problem

$$\begin{cases} x_1'(t) = x_1(t) + x_2(t) \\ x_2'(t) = -x_1(t) + 3x_2(t) \\ x_3'(t) = -4x_1(t) + 7x_2(t) - x_3(t) \end{cases} \quad x_1(0) = 0, \quad x_2(0) = 1, \quad x_3(0) = 0.$$

Inner products and norms

P 5.1. Let $\langle \cdot, \cdot \rangle$ be the standard inner product on \mathbb{C}^2 : $\langle (x_1, x_2), (y_1, y_2) \rangle = x_1 \bar{y}_1 + x_2 \bar{y}_2$. Find:

- a) $\langle (i, 1+i), (3, 1+2i) \rangle$ b) $\|(1+i, 3)\|$ c) All vectors orthogonal to $(1, 1+i)$

P 5.2. Find the length of $p(x) = x^2 + x + 1$ in the real vector space of polynomials equipped with the inner product $\langle p(x), q(x) \rangle = \int_0^1 p(x)q(x)dx$.

P 5.3. Consider a complex inner product space V . Show that the inner product function is conjugate-linear in the second argument:

$$\langle u, \lambda v + \mu w \rangle = \bar{\lambda} \langle u, v \rangle + \bar{\mu} \langle u, w \rangle.$$

P 5.4. Which of the following rules define inner products on the given vector space? In case it is not an inner product, give an example of an axiom that is being violated.

- a) $\langle (x_1, y_1), (x_2, y_2) \rangle = 2x_1x_2 + y_1y_2$ on \mathbb{R}^2 .
 b) $\langle (x_1, y_1), (x_2, y_2) \rangle = x_1x_2 + x_1y_2 + y_1y_2$ on \mathbb{R}^2 .
 c) $\langle (x_1, y_1), (x_2, y_2) \rangle = x_1y_1 + y_1y_2$ on \mathbb{C}^2 .
 d) $\langle (x_1, y_1), (x_2, y_2) \rangle = x_1\bar{x}_2$ on \mathbb{C}^2 .
 e) $\langle x, y \rangle = |xy|$ on \mathbb{R}^2 .
 f) $\langle f(x), g(x) \rangle = \int_0^1 f(x)g(x)dx$ on \mathbb{R} (the real vector space of continuous functions $\mathbb{R} \rightarrow \mathbb{R}$).
 g) $\langle A, B \rangle = \text{tr}(A+B)$ on $\text{Mat}_{n \times n}(\mathbb{R})$.
 h) $\langle A, B \rangle = \text{tr}(AB)$ on $\text{Mat}_{n \times n}(\mathbb{R})$.

P 5.5. Let $\langle \cdot, \cdot \rangle$ be an inner product on a real vector space V . Show that $(u|v) := 2\langle u, v \rangle$ defines a new inner product on V . Inner products are used to define lengths and angles in V - how are lengths and angles changed by using $(\cdot|\cdot)$ instead of $\langle \cdot, \cdot \rangle$ as the inner product on V ?

P 5.6. Does there exist an inner product $\langle \cdot, \cdot \rangle$ on \mathbb{C}^2 for which $u = (1, 1)$ and $v = (1, i)$ are orthogonal? Answer the same question for $u' = (1, 1+i)$ and $v' = (1-i, 2)$?

P 5.7. Let V be an inner product space. Show that the norm arising from the inner product satisfies the parallelogram law:

$$\|u+v\|^2 + \|u-v\|^2 = 2\|u\|^2 + 2\|v\|^2.$$

P 5.8. Let V be an inner product space. Show that the norm arising from the inner product is indeed a norm in the sense of Definition 5.2.1: It satisfies absolute homogeneity, the triangle inequality, and positivity:

$$\|\lambda v\| = |\lambda| \cdot \|v\|, \quad \|u\| + \|v\| \geq \|u+v\|, \quad \|v\| \geq 0 \text{ with equality only for } v=0.$$

P 5.9. Show that the maximum norm on \mathbb{R}^2 can not be derived from an inner product.

P 5.10. Let V be a complex vector space. Then the set of all functions $V \times V \rightarrow \mathbb{C}$ is also a vector space that we call \mathcal{F} . Let \mathcal{I} be the set of all inner products on V . Is \mathcal{I} a subspace of \mathcal{F} ?

P 5.11. Find the distance between $(2, 1, 3)$ and $(4, -2, 6)$ with respect to:

- a) The standard norm
 b) The maximum norm
 c) The Manhattan norm

d) The p -norm for $p = 3$

P 5.12. For the matrix $A = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$, find the Frobenius norm $\|A\|_F$ and the spectral norm $\|A\|_\sigma$.

P 5.13. Prove that an equivalent definition of the operator norm of $F : V \rightarrow W$ is

$$\|F\|_{op} = \max_{v \neq 0} \left\{ \frac{\|F(v)\|}{\|v\|} \right\}.$$

P 5.14. Let $C = \{(\cos^{\frac{3}{2}}(t), \sin^{\frac{3}{2}}(t)) \mid t \in \mathbb{R}\}$. Find p such that C is the unit circle with respect to the p -norm.

P 5.15. Show that the definition of the p -norm does in fact not give a norm for $0 < p < 1$.

P 5.16. Two norms $\|\cdot\|_1$ and $\|\cdot\|_2$ are called *equivalent* if there exists positive constants C and D such that

$$C\|v\|_2 \leq \|v\|_1 \leq D\|v\|_2$$

holds for all $v \in V$. Show that on \mathbb{R}^2 , the maximum norm is equivalent to the standard norm, find the maximal possible C and the minimal possible D .

P 5.17. Show that if any two norms $\|\cdot\|_1$ and $\|\cdot\|_2$ are equivalent (as defined in the previous problem), then $v_n \rightarrow v$ with respect to $\|\cdot\|_1$ if and only if $v_n \rightarrow v$ with respect to $\|\cdot\|_2$.

P 5.18. The familiar vector product on \mathbb{R}^3 is given by

$$(x_1, y_1, z_1) \times (x_2, y_2, z_2) = (y_1 z_2 - y_2 z_1, -x_1 z_2 + x_2 z_1, x_1 y_2 - x_2 y_1).$$

With this rule we know that $u \times v$ is orthogonal to both u and v . Does the same property hold if we use the same definition of a vector product on \mathbb{C}^3 ? Give a proof or a counter-example.

P 5.19. On \mathbb{R}^3 , define $((x_1, x_2, x_3) \mid (y_1, y_2, y_3)) := x_1 y_1 + x_2 y_2 - x_3 y_3$. Show that $(\cdot \mid \cdot)$ is not an inner product, and describe the set of all $v \in \mathbb{R}^3$ which has length zero (meaning $(v \mid v) = 0$)

P 5.20. Let U be a subspace of a finite-dimensional inner product space V . Recall that

$$U^\perp := \{v \in V \mid \langle v, u \rangle = 0 \text{ for all } u \in U\}.$$

a) Show that U^\perp is a subspace of V .

b) Show that $V^\perp = \{0\}$, in other words, the only vector orthogonal to all of V is the zero vector.

c) Show that $U \cap U^\perp = \{0\}$.

P 5.21. In an inner product space V , for $u, v \in V$ with $u \neq 0$, define

$$P_u(v) = \frac{\langle v, u \rangle}{\langle u, u \rangle} u.$$

a) Show that $P_u(v) = P_{\lambda u}(v)$ for complex $\lambda \neq 0$

b) Show that $P_u : V \rightarrow V$ is a linear map

c) Show that $\text{Im}(P_u) = \text{span}(u)$ and that $\ker(P_u) = \text{Im}(P_u)^\perp$

d) Show that $v - P_u(v)$ is orthogonal to u .

e) Show that $P_u^2 = P_u$

P 5.22. Prove that if (e_1, \dots, e_n) is an orthonormal basis in an inner product space, then

$$v = \langle v, e_1 \rangle e_1 + \dots + \langle v, e_n \rangle e_n.$$

P 5.23. Prove the *Pythagorean theorem* in the inner product space setting: If u is orthogonal to v we have

$$\|u + v\|^2 = \|u\|^2 + \|v\|^2.$$

P 5.24. Let $U \subset V$ be a subspace, and fix $v \in V$. Show $u = P_U(v)$ is the vector in U closest to v .

P 5.25. Consider \mathbb{C}^3 with the standard inner product, and let

$$U = \text{span}((1, i, 0), (0, 2i, 1))$$

be a subspace. Use the Gram-Schmidt process to find an ON-basis for U , and extend this basis to an ON-basis for \mathbb{C}^3 .

P 5.26. Find the function $g(x) \in \mathcal{P}_1$ that best approximates e^x with respect to the standard inner product on $\mathcal{C}[0, 1]$:

$$\langle f, g \rangle = \int_0^1 f(x)g(x)dx.$$

P 5.27. Consider the space \mathcal{F} of 2π -periodic real valued continuous functions equipped with the inner product

$$\langle f, g \rangle := \frac{1}{\pi} \int_{-\pi}^{\pi} f(x)g(x)dx.$$

Show that

$$\left\{ \frac{1}{\sqrt{2}}, \sin(x), \cos(x), \sin(2x), \cos(2x), \sin(3x), \cos(3x), \dots \right\}$$

is an orthonormal set of functions in \mathcal{F} , in other words, show that these functions are pairwise orthogonal, and that each function has length 1 with respect to our given inner product.

P 5.28. Let $f(x)$ be a *triangular wave function* with period 2π which equals $|x|$ on $[-\pi, \pi]$. Find the function of form $g(x) = a + b\cos(x) + c\sin(x)$ that best approximates f (with respect to the inner product from the previous problem).

P 5.29. Find a QR-factorization of the matrix $A = \begin{pmatrix} 1 & 3 \\ 1 & 1 \end{pmatrix}$. In other words, find a matrix Q which is unitary ($Q^*Q = I$), and a matrix R which is upper triangular, such that $QR = A$.

P 5.30. Find a QR-factorization of the matrix $A = \begin{pmatrix} 1 & 2 & 1 \\ 1 & 3 & 2 \\ 1 & 1 & 3 \end{pmatrix}$.

P 5.31. Find QR-factorizations of $A = \begin{pmatrix} 1 & 2 & 2 \\ 1 & 4 & 1 \end{pmatrix}$ and of $B = \begin{pmatrix} 1 & 3 \\ 2 & 2 \\ 3 & 1 \end{pmatrix}$.

P 5.32. Let $F : V \rightarrow W$ be a linear operator between inner product spaces. Recall that the adjoint of F is the map $F^* : W \rightarrow V$ where $F^*(w)$ is defined as the unique vector in V such that $\langle F(v), w \rangle = \langle v, F^*(w) \rangle$ for all $v \in V$. Show that F^* is a linear map (without using the fact that $[F^*] = [F]^*$).

P 5.33. Show $(G \circ F)^* = F^* \circ G^*$ and that $(F^*)^* = F$.

P 5.34. Let $F : V \rightarrow W$, and pick ON-bases in V and in W . Prove that $[F^*] = [F]^*$, in other words, the matrix of the adjoint of F is the Hermitian conjugate of the matrix for F with respect to the same ON-bases for V and for W .

P 5.35. Let operators F, G, H be given by the following matrices with respect to an ON-basis:

$$[F] = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix} \quad [G] = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix} \quad [H] = \frac{1}{5} \begin{pmatrix} 3 & 4 \\ -4 & 3 \end{pmatrix}$$

Which of the operators F, G, H are...

- a) Self-adjoint? b) Unitary? c) Normal?

P 5.36. Find an operator that is normal, but that is neither unitary nor self-adjoint.

P 5.37. Determine for what $a, b \in \mathbb{C}$ the matrix $A = \frac{1}{5} \begin{pmatrix} 3 & a \\ b & 3 \end{pmatrix}$ is...

- a) Self adjoint b) Unitary c) Normal

Then answer the same question in the special case when $a, b \in \mathbb{R}$.

P 5.38. Show that the spectral radius of a unitary operator is always 1.

P 5.39. Let $F : V \rightarrow V$ be unitary. Show that F preserves inner products: $\langle F(u), F(v) \rangle = \langle u, v \rangle$ for all $u, v \in V$.

P 5.40. Prove that self-adjoint maps are normal and that unitary maps are normal.

P 5.41. Let $F : V \rightarrow W$ be a linear map between inner product spaces. Show that:

- a) $\ker(F^*) = \text{Im}(F)^\perp$
 b) $\text{Im}(F^*) = \ker(F)^\perp$

P 5.42. Prove that compositions of unitary operators are unitary.

P 5.43. Let A be a normal matrix: $AA^* = A^*A$.

- a) Show that $A + \lambda I$ is normal for all λ .
 b) Show that $\ker(A) = \ker(A^*)$.
 c) Show that if $Av = \lambda v$ then $A^*v = \bar{\lambda}v$.
 d) Show that if u and v are eigenvectors of A corresponding to different eigenvalues, then $u \perp v$.

P 5.44. Recall that a matrix A is called positive definite if $X^*AX > 0$ for all nonzero column-matrices $X \in \mathbb{C}^n$.

- a) Show that a Hermitian matrix need not be positive definite.
 b) Prove that if a matrix is positive definite then it must be Hermitian.
 c) Find a non-Hermitian real matrix A such that $X^*AX > 0$ for all nonzero $X \in \mathbb{R}^n$.
 d) Why does the result in (c) not contradict the result in (b)?

P 5.45. For what $a, b, c \in \mathbb{C}$ is the matrix $A = \begin{pmatrix} 2 & 1+i & 1 \\ a & b & 1 \\ 1 & 1 & c \end{pmatrix}$ positive definite?

P 5.46. Prove Sylvester's criterion for diagonal matrices: If D is diagonal it is positive definite if and only if all principal minors are positive.

P 5.47. Prove one of the implications in Sylvester's criterion: If one principal minor of a matrix is negative, then the matrix is not positive definite.

P 5.48. Write the matrix $A = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$ as a linear combination of two matrices which are orthonormal with respect to the Frobenius inner product.

P 5.49. For square matrices A and B we define

$$A < B \iff B - A \text{ is positive semi-definite.}$$

Prove that $<$ is a partial order relation; in other words, show that

- $A < A$ for all A
- If $A < B$ and $B < C$, then $A < C$
- If $A < B$ and $B < A$, then $A = B$

P 5.50. Find the square root of the matrix $A = \frac{1}{5} \begin{pmatrix} 17 & 6 \\ 6 & 8 \end{pmatrix}$.

P 5.51. Show that if A is invertible, then the Moore-Penrose the same as the usual inverse of A , in other words $A^+ = A^{-1}$.

P 5.52. Find the least square solution to each of the following systems by first finding the Moore-Penrose pseudo-inverse to the coefficient matrix.

$$\text{a) } \begin{cases} x + y = 1 \\ x - y = 2 \\ 2x + y = 3 \end{cases} \quad \text{b) } \begin{cases} x + y = 1 \\ x - y = 0 \\ 2x + y = i \end{cases}$$

P 5.53. The *four Moore-Penrose conditions* can be used to define A^+ for general A . If A is an $m \times n$ -matrix, A^+ is the $n \times m$ matrix satisfying:

$$AA^+A = A \quad A^+AA^+ = A^+ \quad AA^+ \text{ is Hermitian} \quad A^+A \text{ is Hermitian.}$$

- Show that if A has linearly independent columns, then $A^+ := (A^*A)^{-1}A^*$ satisfies the conditions.
- Find a formula for a matrix A^+ satisfying the four conditions when A has linearly independent rows.
- Consider the matrix $A = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$. Find a matrix A^+ satisfying the four conditions.

Perron-Frobenius theory and applications

P 6.1. Let A, B, C be $n \times n$ -matrices, and let X be an $n \times 1$ -matrix.

- Show that if $A \geq B$ and $B \geq C$, then $A \geq C$.
- Show that if $A > B$ and $C > 0$ then $CA > CB$.
- Show that if $A > 0$ and $X \geq 0$ and $X \neq 0$, then $AX > 0$.

P 6.2. Note that *positive* and *positive definite* are two completely different properties of matrices.

- Find a Hermitian matrix which is positive but not positive definite.
- Find a square matrix which is positive definite but not positive.

P 6.3. *Positivity* of square matrices is not a property that can be defined in a basis-free way. Find a matrix A which is not positive, and an invertible matrix S such that $S^{-1}AS$ is positive.

P 6.4. Show that the Frobenius norm of matrices satisfies $\|AB\|_F \leq \|A\|_F \cdot \|B\|_F$ whenever the sizes of A and B match so that AB is defined. This fact was used in the proof of Perron's theorem.

P 6.5. Let $A = \begin{pmatrix} 3 & 2 \\ 2 & 3 \end{pmatrix}$. Find the Perron-eigenvalue and a corresponding Perron-vector $v > 0$ of A .

P 6.6. Show that if A is a diagonalizable matrix with dominant eigenvalue λ_1 (meaning that no other eigenvalue has absolute value $|\lambda_1|$), then for any vector v there exists a positive constant C such that $\|A^k v\| \leq C|\lambda_1|^k$. (The statement is in fact true also when A is not diagonalizable)

- P 6.7.** Make a brute-force proof of the Perron theorem for *stochastic* 2×2 -matrices (stochastic means that the sum in each *row* is 1): show that if A is stochastic of size 2×2 and $A > 0$, then there is a unique eigenvalue of maximal absolute value, and there is a corresponding eigenvector $v > 0$.
- P 6.8.** Make a brute-force proof of the Perron theorem for *stochastic* 3×3 -matrices
- P 6.9.** Prove that if $A > 0$ is stochastic, then the Perron-eigenvalue is $\lambda_1 = 1$.
- P 6.10.** Show that Perron's theorem doesn't hold if we only assume that A is non-negative by finding a square matrix $A \geq 0$ for which the the geometric multiplicity of the dominant eigenvalue has geometric multiplicity > 1 . Also find an example of a matrix $B \geq 0$ that has two different eigenvalues of the same maximal absolute values.
- P 6.11.** A solved Sudoku-puzzle can be viewed as a positive 9×9 -matrix A . Find its Perron-eigenvalue and a corresponding Perron-vector.
- P 6.12.** The *Gershgorin circle theorem* says the following about the eigenvalues of a complex $n \times n$ matrix A . Let s_k be the sum of absolute values of the the non-diagonal elements of the k 'th row of A :

$$s_k = \sum_{i \neq k} |a_{ki}|.$$

Let G_k be the corresponding *Gershgorin disc* in the complex plane:

$$G_k := \{z \in \mathbb{C} : |z - a_{kk}| \leq s_k\}.$$


Then

$$\sigma(A) \subset \bigcup_{k=1}^n G_k,$$

in other words, every eigenvalue lies inside one of the Gershgorin discs. Draw the Gershgorin discs for the matrix A below and verify the statement of the theorem by finding $\sigma(A)$ with the help of a computer program.

$$A = \begin{pmatrix} -4 & 1-i & 1+i \\ 1 & 3 & 1 \\ 1 & -i & 4+2i \end{pmatrix}.$$


- P 6.13.** Prove the Gershgorin circle theorem which was stated in the previous problem (see the hint to get started).
- P 6.14.** Let $A = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$. Find the Perron-eigenvalue. Also, let $v = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and compute $A^k v$ for $k = 1, 2, 3, 4, 5$. Compute the ratio the quotient of the first coordinates of $A^5 v$ by the first coordinate of $A^4 v$ and compare it to the Perron-eigenvalue.
- P 6.15.** Determine whether or not the matrix associated to the graph below is irreducible. Either prove that it is irreducible, or find a relabelling of nodes such that that the adjacency-matrix is block upper triangular.



P 6.16. Which of the matrices below are primitive? Which are irreducible?

a) $A = \begin{pmatrix} 1 & 0 & 3 \\ 4 & 5 & 6 \\ 7 & 0 & 9 \end{pmatrix}$ b) $B = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$ c) $C = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$

P 6.17. Determine whether the adjacency-matrix of the following graph is irreducible and primitive.



P 6.18. Show that if $A \geq 0$ is irreducible and has a nonzero diagonal element $a_{kk} > 0$, then A is primitive.

P 6.19. Show that the graph associated to an $n \times n$ -matrix $A \geq 0$ is strongly connected if and only if


$$A + A^2 + A^3 + \dots + A^{n-1} > 0.$$

P 6.20. Find a 4×4 -matrix with 8 edges which is *not primitive*.

P 6.21. Prove that if A has a unique dominant eigenvalue: $|\lambda_1| > |\lambda_2| \geq |\lambda_3| \geq \dots$, where the algebraic multiplicity of λ_1 is 1, then for all vectors v we have $\|A^k v\| \leq \|\lambda_1^k v\|$. Also show that this is false if the uniqueness condition is removed.

P 6.22. Prove that if A is primitive with $\lambda_1 \leq 1$, then the limit $\lim_{k \rightarrow \infty} A^k$ exists. What can be said about the limit when $\lambda_1 < 1$?

P 6.23. Find the PageRank of the following graph, use the dampening factor $d = 0.7$.



Answer with the ranking-vector scaled so that the sum of its entries is 1.

P 6.24. In a strongly connected graph, suppose that every every node has exactly 5 incoming edges and 5 outgoing edges. What can be said about the PageRank of the different nodes?

P 6.25. Implement the power method in a program and use it to find the PageRank of the graph with the (primitive) adjacency-matrix A below, normalize the ranking vector so that the sum of the components is 1. Which node is ranked highest?

$$A = \begin{pmatrix} 1 & 0 & 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 1 & 0 & 0 \end{pmatrix}$$

P 6.26. Suppose that you are either happy or sad every day. If you are happy today, 90% of the time you will be happy tomorrow too. If you are sad today there is a 50% probability that tomorrow will be a sad day

too. Model this as a Markov process and determine the proportion of days that you will be happy on average.

- P 6.27.** Write a program that simulates the Markov chain in Example 6.7.2 (rainy, cloudy, and sunny days). Start with sunny weather and jump between states according to the given probabilities 1000 times, and count the number of days with rain/sun/clouds. Compare the result with the Perron vector in the example.
- P 6.28.** *The drunkards walk* is a Markov chain in which a person walks between 5 bars lined up on a street. If they are currently at bar number k they will move to either bar $k + 1$ or $k - 1$ with probability $\frac{1}{2}$ each at the next hour (except at bars 1 and 5, then there is only one bar to move to the next hour). Determine at what bar the drunkard will spend most time by finding the stationary distribution for this Markov chain.

Singular values and applications

- P 7.1.** Find the singular values and the singular vectors of $A = \begin{pmatrix} -3 & 0 & 0 & 0 \\ 0 & 1+i & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 3+4i \end{pmatrix}$.
- P 7.2.** Find the SVD and the compact SVD of $A = \begin{pmatrix} 1 & -2 \\ 4 & 2 \end{pmatrix}$.
- P 7.3.** Find the singular value decomposition and the compact singular value decomposition of $A = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$.
- P 7.4.** Find the SVD of $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$.
- P 7.5.** Find the SVD and the compact SVD of $A = \begin{pmatrix} 1 & -1 \\ 1 & 2 \\ 2 & 1 \end{pmatrix}$.
- P 7.6.** Find the SVD of $A = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix}$.
- P 7.7.** Find the compact SVD of $A = \frac{1}{6} \begin{pmatrix} 2 & 4 & 5 \\ -2 & 2 & 1 \\ 4 & 2 & 4 \\ 0 & 0 & 0 \\ 6 & 0 & 3 \end{pmatrix}$.
- P 7.8.** Find the singular values and the singular vectors for the matrix $A = \begin{pmatrix} 0 & 2 & 0 \\ 0 & 0 & 3 \\ 0 & 0 & 0 \end{pmatrix}$ without doing any calculations.
- P 7.9.** Let \mathcal{P}_1 be equipped with the inner product $\langle p, q \rangle = \int_0^1 p(x)q(x)dx$, and let $D : \mathcal{P}_1 \rightarrow \mathcal{P}_1$ be the differentiation operator $D(p(x)) = p'(x)$. Find the singular values of D .
- P 7.10.** Find the singular values of a nilpotent Jordan block $J_n(0)$. Generalize your result to an arbitrary nilpotent matrix in Jordan form.
- P 7.11.** Show that if λ is an eigenvalue of A , then $|\lambda| \leq \sigma_1$.
- P 7.12.** Prove that a square matrix is invertible if and only if it doesn't have 0 as a singular value.
- P 7.13.** Prove that an $m \times n$ -matrix A is an isometry ($\|Av\| = \|v\|$ for all v) if and only if all singular values of A are 1.

P 7.14. Show that if A is Hermitian, then the singular values of A are the absolute values of the eigenvalues of A .

P 7.15. Suppose an $n \times n$ -matrix A has singular values $\sigma_1, \dots, \sigma_n$. For $\lambda \in \mathbb{C}$, determine what the singular values of λA are.

P 7.16. Let A be an invertible matrix with singular value decomposition $A = U\Sigma V^*$. Then $A^{-1} = (V^*)^{-1}\Sigma^{-1}U^{-1} = V\Sigma^{-1}U^*$, but this is in fact *not* an SVD of A in general, why? And how do we find an actual SVD of A^{-1} using the SVD of A ?

P 7.17. Show that for $m \times n$ -matrices A and column-matrices $v \in \mathbb{C}^n$ we have $\|Av\| \leq \|A\|_{\text{op}} \cdot \|v\|$, where the two non-operator norms are the standard norms on \mathbb{C}^m and \mathbb{C}^n respectively.

P 7.18. Show that the operator-norm is sub-multiplicative: whenever AB is defined we have

$$\|AB\|_{\text{op}} \leq \|A\|_{\text{op}} \cdot \|B\|_{\text{op}}.$$

Also determine under what condition equality holds.

P 7.19. Show the following inequalities for the operator norm, Frobenius norm, and the nuclear norm of an arbitrary complex matrix A :

$$\|A\|_{\text{op}} \leq \|A\|_F \leq \|A\|_{\clubsuit}.$$

Also determine when equality holds.

P 7.20. We have found that for $A = \frac{1}{21} \begin{pmatrix} -24 & 62 & 26 \\ 24 & 64 & 58 \\ 42 & -14 & 91 \end{pmatrix}$, a singular value decomposition is given by $A = U\Sigma V^*$

where

$$U = \frac{1}{3} \begin{pmatrix} 1 & 2 & 2 \\ 2 & 1 & -2 \\ 2 & -2 & 1 \end{pmatrix} \quad \Sigma = \begin{pmatrix} 6 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad V = \frac{1}{7} \begin{pmatrix} 2 & -3 & -6 \\ 3 & 6 & -2 \\ 6 & -2 & 3 \end{pmatrix}.$$

Find $A_{(1)}$ and $A_{(2)}$, the best approximations of A by a rank 1 and rank 2 matrix respectively. Also determine how good the approximations are by computing $\|A - A_{(i)}\|$ for both the operator norm and for the Frobenius norm.

P 7.21. We want to approximate a grayscale image of size 3000×2000 by an approximation of rank k such that we need to store about 10% of the data. How should k be chosen (approximately)?

P 7.22. Say that we want to replace a matrix A by an approximation, and we can choose either B or C (we want the error to be as small as possible). Show that our choice may depend on what norm we use. In other words, find three matrices A, B, C such that B is closer to A than C is with respect to the operator norm, while C is closer to A than B is with respect to the Frobenius norm.

P 7.23. Find a matrix such that it has a unique rank 1 optimal approximation, while there are two equally good optimal rank 2 approximations. This exercise shows that $A_{(k)}$ is not necessarily unique.

P 7.24. Prove that the matrices $A_i := u_i v_i^*$ appearing in the low rank approximation of A are orthonormal with respect to the Frobenius inner product.

P 7.25. Prove that the Eckart-Young-Mirsky theorem in fact also holds for the nuclear norm:

$$\text{The minimum } \min_{\text{rank}(X) \leq k} \|A - X\|_{\clubsuit} \text{ is attained for } X = A_{(k)}.$$

P 7.26. Find the Moore-Penrose pseudo-inverse of $A = \begin{pmatrix} 3 & 6 \\ 4 & 8 \end{pmatrix}$.

P 7.27. Find the Moore-Penrose pseudo-inverse of $A = \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \end{pmatrix}$ by the general method that uses the SVD. Then compute A^+A and AA^+ and determine whether A^+ is a left or a right inverse to A .

- P 7.28.** Find a general formula for the Moore-Penrose pseudo inverse A^+ when $\text{rank}(A) = 1$.
- P 7.29.** Prove that if A has linearly independent columns, the old definition of $A^+ = (A^*A)^{-1}A^*$ coincides with the more general definition.
- P 7.30.** Prove that $(A^*)^+ = (A^+)^*$.
- P 7.31.** Prove that the Moore-Penrose pseudo-inverse A^+ of an arbitrary $m \times n$ -matrix A indeed satisfies:
- A^+A and AA^+ are both Hermitian.
 - $AA^+A = A$.
 - $A^+AA^+ = A^+$.

as stated in the text.

- P 7.32.** Prove that geometrically, AA^+ is the orthogonal projection onto $\text{Im}(A)$.
- P 7.33.** Under what condition is $A^+A = I$? Under what condition is $AA^+ = I$?
- P 7.34.** Prove that A^+b is a least square solution to the linear system $Ax = b$.
- P 7.35.** Prove that A^+b is the least square solution to $Ax = b$ of *shortest length*; in other words, among the solutions x to $A^*Ax = A^*b$, we have that $\|x\|$ is minimal for $x = A^+b$.

P 7.36. Find the condition number of $A = \begin{pmatrix} 1+i & & & \\ & -3 & & \\ & & 3+4i & \\ & & & \pi \end{pmatrix}$.

P 7.37. For a certain 2×2 -matrix A we know that $|\det(A)| = 12$ and $\|A\|_F = 5$. Find the condition number $\kappa(A)$.

P 7.38. Let $A = \begin{pmatrix} 3 & \alpha \\ \alpha & 3 \end{pmatrix}$ where $\alpha \in \mathbb{R}$ is a constant. Find the condition number $\kappa(A)$ for each value of α .

P 7.39. Consider a system $Ax = b$ where A is an invertible $n \times n$ -matrix. In what directions should we pick b and Δb to get the worst relative error Δx in the solution? Answer in terms of the singular vectors of A , you can assume that all the singular values are different.

P 7.40. Find a polar factorization of $A = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$.

P 7.41. Find a polar factorization of the map F on \mathbb{C}^3 where $F(x, y, z) = (3z, 5x, 7y)$, in other words, express F as a composition of a unitary and a Positive definite operator.

P 7.42. Find a polar factorization of $A = \frac{1}{\sqrt{2}} \begin{pmatrix} 2+i & -2+i \\ 2-i & 2+i \end{pmatrix}$

P 7.43. Find an example of a nonzero matrix which has infinitely many different polar factorizations.

P 7.44. Show that any square matrix has a "flipped" polar factorization $A = PU$ where P is positive semi-definite and U is unitary.

P 7.45. In *topic classification* we look at frequencies of words across a set of documents in order to identify a number of "topics", this is useful for sorting books with similar contents, and dually, for identifying connections between words. There are many variations and subtleties, but here is a minimal example: We go through seven magazines and count the number of occurrences of the words

adventure, culinary, finance, politics

We collect our data in a matrix, where the columns correspond to our four words above, and where

each row is a magazine:

$$A = \begin{pmatrix} 5 & 3 & 0 & 1 \\ 1 & 0 & 5 & 8 \\ 1 & 0 & 7 & 9 \\ 6 & 4 & 1 & 2 \\ 0 & 0 & 3 & 4 \\ 3 & 3 & 1 & 1 \\ 1 & 1 & 5 & 8 \end{pmatrix}.$$

Write a script that finds the SVD of A . Then project each of the four words onto the first principal components and plot the results. Then do the same in the with the seven magazines. Try and interpret what the two principal components correspond to.

P 7.46. Use the total least squares method to find the 1-dimensional affine subspace (line) that in the least squares sense lies closest to the points

$$(4, 4), (3, 1), (2, -1), (-1, 0).$$

Also find the regular least square solution (the line that minimizes the vertical distances to the points).

P 7.47. Consider the following four points in \mathbb{R}^3 :

$$(2, 2, 0), (2, 3, 8), (-1, 0, 4), (1, 3, 0)$$

Find the line in \mathbb{R}^3 that best approximates the points in a total least square sense. Also find the plane in \mathbb{R}^3 that best approximates the points in a total least square sense.

P 7.48. Show that minimizing the square distances to an affine subspace is not the same as just minimizing the distances. In other words, find an example where you explicitly can show that different subspaces S minimizes the distances to a given set of points, such that the two expressions

$$D = \sum_{i=1}^m d(x_i, S)^2 \quad \text{and} \quad D' = \sum_{i=1}^m d(x_i, S)$$

are minimized for different choices of affine subspace S .

Multilinear algebra

P 8.1. Express the vector $v \in \mathbb{R}^2 \otimes \mathbb{R}^2$ below in the basis $(e_1 \otimes e_1, e_1 \otimes e_2, e_2 \otimes e_1, e_2 \otimes e_2)$. Can v be expressed as a pure tensor $v = u \otimes w$?

$$v = 2 \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} - 3 \begin{pmatrix} -2 \\ 2 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ 1 \end{pmatrix}.$$

P 8.2. Can the vector $w \in \mathbb{R}^2 \otimes \mathbb{R}^2 \otimes \mathbb{R}^2$ below be expressed with fewer terms?

$$w = \begin{pmatrix} 2 \\ -2 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ 3 \end{pmatrix} \otimes \begin{pmatrix} 2 \\ 1 \end{pmatrix} - \begin{pmatrix} -3 \\ 3 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 4 \\ 2 \end{pmatrix} + \begin{pmatrix} 4 \\ -4 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} -6 \\ -3 \end{pmatrix}$$

P 8.3. Find an example of a 2-tensor which is a linear combination $w = \sum u_i \otimes v_i$ with more than one term and where none of the u_i are parallel and where none of the v_i are parallel, but such that the tensor w can still be rewritten as a pure tensor $w = w_1 \otimes w_2$.

P 8.4. A tensor $v \in \mathbb{C}^2 \otimes \mathbb{C}^3 \otimes \mathbb{C}^2$ can be visualized as the three dimensional array below, where the front and back matrix-slabs correspond to the the two basis vectors of the third component \mathbb{C}^2 of the space. Show

that v is a pure tensor by expressing it as $v = v_1 \otimes v_2 \otimes v_3$ explicitly.

2	4i	-6i
i	-2	3
2i	-4	6
-1	-2i	3i

P 8.5. Fill in all the missing values \star such that $v \in \mathbb{C}^2 \otimes \mathbb{C}^3 \otimes \mathbb{C}^2$ visualized below becomes a pure tensor, and express the tensor as $u \otimes v \otimes w$.

-2	-3	\star
\star	\star	\star
\star	\star	-4
\star	9	\star

P 8.6. Let $t \in \mathbb{R}^5 \otimes \mathbb{R}^5 \otimes \mathbb{R}^5$, and think of v as a 3d-array, a cube of size $5 \times 5 \times 5$. With respect to the standard basis, all coefficients in the tensor $[t]$ is either 1 or 0, and they form an alternating pattern with coefficients switching between 0 and 1 when an index increases. For example, $[t]_{111} = 1$, $[t]_{112} = 0$, $[t]_{212} = 1$ and so on. Express t as a sum of pure tensors with as few terms as possible, and prove that your expression has minimal number of terms.

P 8.7. Let $\mathbb{R}[x, y]$ be the vector space of polynomials in two variables. Is either of the two spaces

$$\mathbb{R}[x] \oplus \mathbb{R}[y] \quad \text{or} \quad \mathbb{R}[x] \otimes \mathbb{R}[y]$$

isomorphic to $\mathbb{R}[x, y]$? Provide an explicit isomorphism between $\mathbb{R}[x, y]$ and one of the two spaces.

P 8.8. We know that $\dim(W \otimes V^{\otimes 3} \otimes W) = 5400$. Find the dimension of V and of W .

P 8.9. Prove that if V is a real vector space, then

$$\mathbb{R}^2 \otimes V \simeq V \oplus V,$$

and provide an explicit isomorphism.

P 8.10. In the category of finite-dimensional vector spaces, the tensor-hom adjunction says that for any three vector spaces U, V, W we have

$$\text{Hom}(U \otimes V, W) \simeq \text{Hom}(U, \text{Hom}(V, W)).$$

Prove this statement.

P 8.11. Find an *explicit* isomorphism for the previous problem.

P 8.12. Let $e_1 = (3, 1)$ and $e_2 = (1, 1)$ be a basis for $V = \mathbb{R}^2$, and let (e_1^*, e_2^*) be the dual basis for V^* .

- a) Evaluate $e_1^*((7, 4))$ and $e_2^*((7, 4))$.
- b) Evaluate $(2e_1^* - 3e_2^*)((5, 3))$.
- c) The linear map $g \in V^*$ is defined by $g(x, y) = x + y$. Express g in our basis (e_1^*, e_2^*) of V^* .

P 8.13. Let \mathcal{P}_3 be the space of polynomials of degree ≤ 3 with standard basis $(e_0, e_1, e_2, e_3) = (1, x, x^2, x^3)$. Express the linear functional map

$$f : \mathcal{P}_3 \rightarrow \mathbb{R} \quad \text{defined by} \quad f(p(x)) = p(-1) + p'(1)$$

in the dual basis e_i^* .

P 8.14. Referring to the double dual isomorphism theorem 8.3.6:

- a) Show that the map Φ_v is linear for each v .
- b) Show that the map Φ is linear.
- c) Show that Φ is injective.

P 8.15. Let $\mathbb{F} = \mathbb{Z}_2$, the field of two elements. How many elements does the \mathbb{F} -vector space $V = \mathbb{F}^4 \otimes (\mathbb{F} \oplus \mathbb{F}^2)$ have?

P 8.16. Find the dimension of the space of (p, q) -tensors on V in terms of the dimension of V .

P 8.17. Let $w \in \mathbb{R}^2 \otimes \mathbb{R}^2 \otimes (\mathbb{R}^2)^*$ where

$$w = \begin{pmatrix} 1 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ -1 \end{pmatrix} \otimes f + \begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 1 \end{pmatrix} \otimes g,$$

where $f, g \in (\mathbb{R}^2)^*$ with $f(x, y) = 2x + y$ and $g(x, y) = x - y$. Find all possible contractions of the tensor w .

P 8.18. Let F be a linear map $\mathbb{R}^2 \rightarrow \mathbb{R}^2$ which with respect to the standard basis in \mathbb{R}^2 has matrix $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$.

Express F as an element of $\mathbb{R}^2 \otimes (\mathbb{R}^2)^*$, and use this to compute the $(1, 1)$ -contraction of F (here the image of the contraction is an element of \mathbb{R} , the field itself, think of $V \otimes V^*$ as $\mathbb{R} \simeq V \otimes V^*$ before contracting)

P 8.19. A linear map F corresponds to a tensor $w \in V \otimes V^*$. Prove that for the contraction $C_{1,1} : V \otimes V^* \rightarrow \mathbb{F}$, we have $C_{1,1}(w) = \text{tr}(F)$.

P 8.20. In a first linear algebra course we discussed the the vector product, also called the cross product of vectors $u \times v$. Express this multiplication as a $(1, 2)$ -tensor on \mathbb{R}^3 , in other words, express \times as an element

$$\times = t \in \mathbb{R}^3 \otimes (\mathbb{R}^3)^* \otimes (\mathbb{R}^3)^*$$

with respect to the standard basis e_1, e_2, e_3 of \mathbb{R}^3 such that if $t = \sum a_{ijk} e_i \otimes e_j^* \otimes e_k^*$, we have $u \times v = \sum a_{ijk} e_i \cdot e_j^*(u) e_k^*(v)$.

P 8.21. Prove that if F and G are invertible operators, then so is their tensor product $F \otimes G$.

P 8.22. Let $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$ and $B = \begin{pmatrix} 3 & 1 \\ 0 & 2 \end{pmatrix}$. Compute the Kronecker products:

- a) $I \otimes A$
- b) $B \otimes I$
- c) $A \otimes B$
- d) $A \otimes A$

P 8.23. Let S and R be linear operators on \mathbb{R}^2 where S is reflection in the line $x + y = 0$, and where R is rotation counter-clockwise by the angle $\frac{\pi}{6}$. Find the matrix of $S \otimes R$ with respect to the basis

$$(e_1 \otimes e_1, e_1 \otimes e_2, e_2 \otimes e_1, e_2 \otimes e_2) \quad \text{of} \quad \mathbb{R}^2 \otimes \mathbb{R}^2.$$

P 8.24. The (row)-vectorization of an $m \times n$ matrix $A = \begin{pmatrix} r_1 \\ \vdots \\ r_m \end{pmatrix}$ with rows r_i is the vector

$$\text{vec}(A) = (r_1 \quad r_2 \quad \cdots \quad r_m) \in \text{Mat}_{m \times n} \simeq \mathbb{R}^{mn}.$$

Express $\text{vec}(A)$ algebraically using the Kronecker products and the rows r_i .

P 8.25. Let $F : \mathbb{C}^2 \rightarrow \mathbb{C}^2$ and $G : \mathbb{C}^3 \rightarrow \mathbb{C}^3$ have eigenvalues $\sigma(F) = \{2, 4\}$ and $\sigma(G) = \{2, 3, 6\}$. Find all eigenvalues of $F \otimes G$, and find the dimension of each corresponding eigenspace.

P 8.26. Let $F : V \rightarrow V$ and $G : W \rightarrow W$ be diagonalizable linear maps on finite dimensional vector spaces. Show that

$$\text{tr}(F \otimes G) = \text{tr}(F) \cdot \text{tr}(G).$$

P 8.27. Show that For Kronecker products we have $A \otimes B = 0$ if and only if $A = 0$ or $B = 0$.

P 8.28. Prove that for square diagonalizable matrices A and B , we have $\text{rank}(A \otimes B) = \text{rank}(A) \cdot \text{rank}(B)$.

P 8.29. Let A be an $m \times m$ -matrix and let B be an $n \times n$ -matrix. Show that $A \otimes I_n$ commutes with $I_m \otimes B$.

P 8.30. Prove that if λ is an eigenvalue of an $m \times m$ -matrix A and if μ is an eigenvalue of an $n \times n$ -matrix B , then $\lambda + \mu$ is an eigenvalue of the matrix

$$A \otimes I_n + I_m \otimes B.$$

P 8.31. Use the previous problem to find a polynomial with integer coefficients which has $\sqrt{2} + \sqrt{3}$ as a root.

P 8.32. Use Kronecker products to show that the algebraic numbers are closed under addition and multiplication, and that they thereform form a subfield of \mathbb{C} . A number α is called algebraic if it is the zero of some polynomial with integer coefficients: $p(\alpha) = 0$ for some $p \in \mathbb{Z}[x]$.

P 8.33. Let A and B be square matrices. Prove that the Kronecker products $A \otimes B$ and $B \otimes A$ are *permutation similar*: there is a permutation matrix P (a matrix with a single 1 in each column and in each row and zeroes elsewhere), such that

$$B \otimes A = P(A \otimes B)P^T.$$

Conclude that $A \otimes B$ and $B \otimes A$ have the same Jordan form.

P 8.34. Prove that the Kronecker product of unitary matrices is unitary.

P 8.35. Let A and B be matrices of arbitrary size. Show that $\|A \otimes B\|_{\text{op}} = \|A\|_{\text{op}} \cdot \|B\|_{\text{op}}$.

P 8.36. Prove that if $A \sim B$ and $C \sim D$, then $A \otimes C \sim B \otimes D$, where \sim means matrix similarity.

P 8.37. Prove that for matrices of arbitrary size, we have $\text{rank}(A \otimes B) = \text{rank}(A) \cdot \text{rank}(B)$.

P 8.38. Prove that if $U' \subset U$ and $V \subset V'$ are subspaces, then

$$U' \otimes V' := \text{span}\{u' \otimes v' \in U \otimes V \mid u' \in U', v' \in V'\}$$

is a subspace of $U \otimes V$. Does any subspace of $U \otimes V$ have this form?

P 8.39. Let A be an $m \times n$ -matrix and let B be an $m' \times n'$ -matrix. Show that

$$\ker(A \otimes B) = \ker(A) \otimes \mathbb{R}^{n'} + \mathbb{R}^n \otimes \ker(B),$$

in other words, $(A \otimes B)(u \otimes v) = 0$ if and only if $u \in \ker(A)$ or $v \in \ker(B)$.

P 8.40. In a neural network you discover that you can speed it up significantly by using the identity function $\sigma(x) = x$ instead of $\max(0, x)$. Why is this a bad idea?

Appendix

Fields

P 9.1. None of the following algebraic structures are fields. Why?

- $\text{SL}_n(\mathbb{R}) = \{A \in \text{Mat}_{n \times n}(\mathbb{R}) \mid \det(A) = 1\}$ with the usual sums and products of matrices.
- $\mathbb{R}_+ = [0, \infty)$, the nonnegative real numbers with the usual addition and multiplication.
- \mathbb{R}^2 with the usual addition, and element-wise product $(a, b) \cdot (c, d) := (ac, bd)$.
- \mathbb{Z}_{15} , the integers modulo 15 with the usual multiplication and addition modulo 6.
- \mathbb{R}^3 with the usual addition, and with vector-product as the multiplication.

P 9.2. The complex numbers are numbers of form $a + bi$ where the normal algebraic rules hold with the addition that $i^2 = -1$. Similarly, the algebra of *Quaternions* \mathbb{H} consists of numbers $a + bi + cj + dk$ where $a, b, c, d \in \mathbb{R}$ and where the "new numbers" i, j, k multiply according to the following table:

\cdot	i	j	k
i	-1	k	$-j$
j	$-k$	-1	i
k	j	$-i$	-1

Do the quaternions \mathbb{H} form a field?

P 9.3. The erroneous simplification $(a + b)^n = a^n + b^n$ is called "The freshman's dream" since it is a common error among first year students. Prove that over the field \mathbb{Z}_p where p is a prime, the dream comes true!

P 9.4. Show that the following properties hold in any field F :

- a) If $a + b = a + c$ then $b = c$
- b) $0 \cdot a = 0$ for each a .
- c) $-a = (-1) \cdot a$ for each a
- d) There are no nontrivial *zero-divisors*: If $a \cdot b = 0$, then $a = 0$ or $b = 0$.
- e) The cancellation law holds: If $a \neq 0$ and $ab = ac$, then $b = c$.
- f) The element 0 in field axiom (F4) is *unique* - there can not be two elements with this property.
- g) The element 1 of (F7) is also unique.
- h) The right distributive law $(a + b)c = ac + bc$ holds.

P 9.5. We required that $0 \neq 1$ in the field axioms. Show that if we drop this requirement and we have $0 = 1$ in a field F , then F has only 1 element.

P 9.6. Compute in \mathbb{Z}_7 :

- a) Compute $3 + 5 \cdot 4 - 6$
- b) Compute $\frac{3}{4}$
- c) Compute 6^{11}
- d) Solve the equation $x^2 = 2$
- e) Solve the equation $x^2 = 5$

P 9.7. Find the multiplicative inverse of 7 in \mathbb{Z}_{53} .

P 9.8. Use Fermat's little theorem 9.1.7 to find an explicit formula for a^{-1} in \mathbb{Z}_p when $a \neq 0$.

P 9.9. The *fundamental theorem of algebra* says that over the field \mathbb{C} , every non-constant polynomial $p(x) \in \mathbb{C}[x]$ has a zero in \mathbb{C} . Show that the fundamental theorem of algebra does not hold if \mathbb{C} is replaced by the field \mathbb{Z}_p .

P 9.10. Solve *all* second degree equations over \mathbb{Z}_2 . In other words, for each polynomial $p(t) \in \mathbb{Z}_2[t]$ of degree 2, find its zeroes.

P 9.11. Elliptic curves over finite fields are useful tools in cryptography. The solutions to the equation

$$y^2 = x^3 - x + 1$$

form what is called an *elliptic curve*. Find all solutions to the equation over the field \mathbb{Z}_5 , meaning all pairs $x, y \in \mathbb{Z}_5$ that satisfy the equation.

P 9.12. Let $F = \left\{ \begin{pmatrix} a & b \\ -b & a \end{pmatrix} \mid a, b \in \mathbb{R} \right\}$ equipped with the standard addition and products of matrices. Show that F is a field.

P 9.13. Let F be the field of four elements from Example 9.1.9. We had $F = \{\heartsuit, \diamondsuit, \spadesuit, \clubsuit\}$ with:

+	\clubsuit	\heartsuit	\spadesuit	\diamondsuit
\clubsuit	\clubsuit	\heartsuit	\spadesuit	\diamondsuit
\heartsuit	\heartsuit	\clubsuit	\diamondsuit	\spadesuit
\spadesuit	\spadesuit	\diamondsuit	\clubsuit	\heartsuit
\diamondsuit	\diamondsuit	\spadesuit	\heartsuit	\clubsuit

·	\clubsuit	\heartsuit	\spadesuit	\diamondsuit
\clubsuit	\clubsuit	\clubsuit	\clubsuit	\clubsuit
\heartsuit	\clubsuit	\heartsuit	\spadesuit	\diamondsuit
\spadesuit	\clubsuit	\spadesuit	\diamondsuit	\heartsuit
\diamondsuit	\clubsuit	\diamondsuit	\heartsuit	\spadesuit

- a) Compute $\diamondsuit \cdot (\diamondsuit + ((\diamondsuit + \clubsuit) \cdot \spadesuit))$.
- b) What is the additive identity element 0 of F ?
- c) What is the multiplicative identity 1 of F ?
- d) Find the additive inverses $-\clubsuit$, $-\heartsuit$, $-\spadesuit$, och $-\diamondsuit$ in F .
- e) Find the multiplicative inverses \heartsuit^{-1} , \spadesuit^{-1} , och \diamondsuit^{-1} in F .
- f) Solve the equation $\spadesuit(x + \heartsuit) = \diamondsuit$.
- g) Solve the equation $x^2 + \heartsuit = \spadesuit$.

P 9.14. The *characteristic* $\text{char}(F)$ of a field F is the smallest positive integer k such that in F we have $\underbrace{1 + 1 + \dots + 1}_k = 0$.

- 0. If no such integer exists we write $\text{char}(F) = 0$.
 - a) What is the characteristic of the fields $\mathbb{R}, \mathbb{C}, \mathbb{Q}, \mathbb{Z}_5, \mathbb{Z}_3(x)$, and the field $K = \{\heartsuit, \diamondsuit, \spadesuit, \clubsuit\}$ from the previous problem.
 - b) Prove that any finite field F has positive characteristic: $\text{char}(F) > 0$.
 - c) Prove that the characteristic of any field F is either zero or a prime number.

P 9.15. We say that a field F is *algebraically closed* if every nonconstant polynomial with coefficients in F has a root in F :

$$p(t) \in \mathbb{F}[t] \text{ with } p(t) \notin \mathbb{F} \Rightarrow \exists \alpha \in F : p(\alpha) = 0.$$

Which of the fields $\mathbb{Q}, \mathbb{R}, \mathbb{C}$ and \mathbb{Z}_7 are algebraically closed?

Chapter 11

Hints

H 1.1. The equality looks obvious, but note that $+$ and \cdot can be defined in some non-standard way. Start from the left hand side and simplify it using the axioms. Start with the axiom that says that $1 \cdot v = v$.

H 1.2. The zero element $\mathbf{0}$ is the unique element of \mathbb{R}_+ satisfying $\mathbf{0} + v = v$ for all $v \in \mathbb{R}_+$.

H 1.3. Recall that a (real) subspace of V is a non-empty subset $S \subset V$ such that:

- $u, v \in S \Rightarrow u + v \in S$ (additivity)
- $\lambda \in \mathbb{R}, v \in V \Rightarrow \lambda \cdot v \in S$ (homogeneity).

To show that S is a subspace, these conditions have to be tested for all vectors and scalars. To show that S is not a subspace, it suffices to find a single counter-example to the conditions.

H 1.4. The projection is not orthogonal. Express $v = u + u'$ with $u \in U$ and $u' \in U'$.

H 1.5. For the basis, note that if a polynomial $p(x)$ lies in S , so does all its multiples $q(x)p(x)$.

H 1.6.

H 1.7. Look for a counterexample in $V = \mathbb{R}^2$

H 1.8. A function f is even if $f(-x) = f(x)$ for all x and odd if $f(-x) = -f(x)$. Show these properties are preserved when taking sums of functions or products of functions by scalars. To show that each f can be expressed as $e(x) + o(x)$, note that $f(x) + f(-x)$ is always even and $f(x) - f(-x)$ is odd.

H 1.9. For each map F , check whether $F(u + v) = F(u) + F(v)$ and $F(\lambda v) = \lambda F(v)$ holds for all vectors u, v in the domain of F and all scalars λ .

H 1.10.

H 1.11. Recall that the definition of F being injective is that $F(u) = F(v)$ only when $u = v$.

H 1.12. Recall the definitions. A map $F : V \rightarrow V$ is injective if $\ker(F) = 0$, surjective if $\text{Im}(F) = V$, and has an inverse G if both $G \circ F$ and $F \circ G$ is the identity map on V .

H 1.13. Consider where the standard basis $(1, x, x^2, x^3, x^4)$ is mapped, the images form the columns of $[F]$. For the inverse, figure it out without matrices.

H 1.14. The matrices of F and G will be 3×3 where the columns are the images of the basis vectors expressed in the given basis.

H 1.15. Recall that $v + U = \{v + u \mid u \in U\}$.

H 1.16.

- Two affine subsets are equal $v + U = w + U$ if and only if $v - w \in U$.
- Each subset is a line in \mathbb{R}^2 through the given point and in the direction $(1, 1)$.

- c) The coordinate for B in the basis (A) is the single number λ such that $\lambda A = B$, or in other words $\lambda(2, 3) + U = (0, 1) + U$.

H 1.17. Visually, elements of V/U are all the lines in direction $(1, 1, 1)$. The coordinates (λ_1, λ_2) you seek must satisfy $\lambda_1 A + \lambda_2 B = C$, these can be found by solving $\lambda_1(1, 1, 0) + \lambda_2(0, 1, 1) + \lambda_3(1, 2, 3) = (1, 1, 1)$.

H 1.18. e_1, e_2, e_3 are linearly dependent if $\lambda_1 e_1 + \lambda_2 e_2 + \lambda_3 e_3 = 0 + \ell$ (the right side is the zero vector in V). This means that

$$\lambda_1(1, 1, 0, 0) + \lambda_2(0, 1, 1, 0) + \lambda_3(0, 0, 1, 1) = \lambda_4(3, 2, 1, 2),$$

solve this linear system.

H 1.19. If $v - w \in U$, then $v + U = w + U$, so we must have $\tilde{F}(v) = \tilde{F}(w)$ in order for \tilde{F} to be well-defined. The second question is vague, since it depends on what bases we pick, but consider a bases for (u_1, \dots, u_m) of U and extend it to a basis $(u_1, \dots, u_m, v_1, \dots, v_n)$ of V , then a natural choice of basis for V/U is $(v_1 + U, \dots, v_n + U)$, consider how the matrices look with respect to these bases.

H 1.20. Note that if p_1, p_2 lies in an affine subset $v + U$, then $p_1 - p_2 \in U$, so U is spanned by all differences between the points.

H 1.21. For (b), to find the translation vector w , consider where the origin is mapped.

H 1.22. Note that $v = -v$ when the field is \mathbb{Z}_2 .

H 1.23. It remains to show that the property $F(\lambda v) = \lambda F(v)$ holds for all $\lambda \in \mathbb{Q}$. Start by showing that it holds for $\lambda \in \mathbb{Z}$.

H 1.24. Basically everything you can do over a real vector space you can do the same way over \mathbb{Z}_3 . Recall that the only scalars are 0, 1, 2 though.

H 1.25. Proceed with row operations as usual, just remember that all operations and variables should be over the field \mathbb{Z}_5 . For (a), start with adding 3 times the first row to the second row, this will remove x from the second equation.

H 1.26. The vectors are linearly independent if the equation $\lambda_1 v_1 + \lambda_2 v_2 + \lambda_3 v_3 = 0$ (where $\lambda_i \in \mathbb{Z}_3$) only has the trivial solution $\lambda_1 = \lambda_2 = \lambda_3 = 0$. Use the same algorithm as you would in \mathbb{R}^4 but remember that all operations should now be calculated over \mathbb{Z}_3 instead of over \mathbb{R} .

H 1.27. You can use the same method that you would over \mathbb{R} . Another way is to actually write out all points that lies on each line.

H 2.1. Recall that $A^* = \overline{A}^T$, the conjugate transpose of A . A matrix is Hermitian if $A^* = A$.

H 2.2. That A is symmetric means that $A^T = A$ and that B is skew-symmetric means that $B^T = -B$.

H 2.3. Let A be $m \times n$ and let B be $n \times m$, use the sum-version of matrix multiplication: $(AB)_{ij} = \sum_{k=1}^n a_{ik} b_{kj}$ - what are the diagonal elements of AB ?

H 2.4.

a) Since A is diagonalizable, $A = SDS^{-1}$, where $D = \text{diag}(\lambda_1, \dots, \lambda_n)$. Apply the previous problem to prove that $\text{tr}(A) = \text{tr}(D)$.

b) $(S^{-1}AS)^n = D^n$

c) (the eigenvalues are integers)

H 2.5. Compute P^2 and P^3

H 2.6.

a) Count the 1's of the matrix first row by row, then column by column.

b) Think of the two matrices as performing permutations of the basis vectors.

- c) Same hint as above.
- d) You need at least a 5×5 -matrix.

- H 2.7. For example, let $X = \begin{pmatrix} x_1 & x_2 \\ x_3 & x_4 \end{pmatrix}$ and solve the linear system $AX = XA$
- H 2.8. Suppose that $A \in \mathbb{C}$, and compute $e_{ij}A$ and Ae_{ij} for different choices of (i, j) (where e_{ij} is a matrix with a single 1 in position (i, j) and zeroes elsewhere), what does this say about the entries of A ?
- H 2.9. Start by computing some low powers of A^n . Remember to perform all calculations in $\text{Mat}_2(\mathbb{Z}_2)$ when computing the powers.
- H 2.10. $E = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$ is an idempotent.
- H 2.11. Diagonalize A so that $S^{-1}AS = D$. Then A commutes with B , if and only if D commutes with $S^{-1}BS$, so $\dim \mathcal{C}_A = \dim \mathcal{C}_D$, so it suffices to investigate what matrices commute with a given diagonal matrix. Consider first the case where all eigenvalues are distinct, then the case where all eigenvalues coincide, and finally the case where some of them coincide.
- H 2.12. Assume that $M^m = 0$ and $N^n = 0$, what can be said about $(AB)^{m+n}$ and $(M+N)^{m+n}$? For $M+N$, think about the binomial theorem. For the case when A and B do not commute, look for a 2×2 counterexample.
- H 2.13. What can be said about $F(F(v))$?
- H 2.14. Consider the sequence of subspaces $\ker(N) \subset \ker(N^2) \subset \dots$. Show that the dimensions of these subspaces must be *strictly* increasing. (What would it mean if $\ker(N^k) = \ker(N^{k+1})$ for some k ?)
- H 2.15. You can write down $(I+N)^{-1}$ explicitly, think of the formula for geometric sums and its proof.
- H 2.16. Multiply M by M^{-1} in block form and show that the product is the identity matrix.
- H 2.17. Make an ansatz $M^{-1} = \left(\begin{array}{c|c} X_1 & X_2 \\ \hline X_3 & X_4 \end{array} \right)$ and multiply by M . Both MM^{-1} and $M^{-1}M$ should equal the identity matrix $\left(\begin{array}{c|c} I & 0 \\ \hline 0 & I \end{array} \right)$.
- H 2.18. The rank is the number of pivots in the REF (or RREF). $\ker(A)$ is the set of solutions to $AX = 0$, which is the same as the set of solutions to $\tilde{A}X = 0$ where \tilde{A} is the reduced row echelon form of A .
- H 2.19. Row operations works exactly the same over \mathbb{C} . Recall that a complex fraction can be simplified by multiplying numerator and denominator by the complex conjugate of the denominator.
- H 2.20. Remember that the parameters r, s, t correspond to non-pivot columns, find the rows of the RREF from bottom to top.
- H 2.21. Remember that over \mathbb{Z}_3 there are only three scalars: $0, 1, 2$, so any row operations can only involve these. Start by multiplying the first row by 2 to get a one top left. Remember that the solutions X should also be vectors in $(\mathbb{Z}_3)^4$.
- H 2.22. $\text{Im}(A)$ is spanned by the columns of A , to remove linearly dependent columns, find the echelon form of A , and remove columns of A corresponding to non-pivot columns in the echelon form. $\ker(A)$ consists of all solutions X to $AX = 0$, such X can easily be found from the echelon form. The parametrization of the solutions yields a basis. There are sometimes quick ways to see what the intersection is, but recall the standard algorithm: If u_1, \dots, u_m is a basis for U and v_1, \dots, v_n is a basis for V , then a vector w lies in the intersection if $\lambda_1 u_1 + \dots + \lambda_m u_m = w = \lambda_{m+1} v_1 + \dots + \lambda_{m+n} v_n$ has a solution, in other words there exists a $(m+n)$ -tuple $(\lambda_1, \dots, \lambda_{m+n})$ satisfying this equation. One can find all such λ_i by row-operating on the matrix which has all the $m+n$ vectors as columns.

- H 2.23. For example, consider a linear map on \mathbb{R}^4 such that $e_1 \mapsto e_2 \mapsto e_3 \mapsto 0$ and $e_4 \mapsto 0$. What are the three subspaces?
- H 2.24. Think of the matrices as performing row operations - what is the opposite of a given row operation?
- H 2.25. Multiplying A by elementary matrices from the **left** corresponds to doing row operations on A . Multiplication on the right corresponds to column operations.
- H 2.26. Reduce A to the identity matrix by multiplying by elementary matrices on the left (each multiplication corresponding to a row operation). If $E_3E_2E_1A = I$, then $A = (E_3E_2E_1)^{-1} = E_1^{-1}E_2^{-1}E_3^{-1}$. The factorization of the inverse is trivial once you have the factorization of A .
- H 2.27. Row operations on $[A|I]$ can be realized as multiplying the block matrix by elementary matrices on the left. Consider what happens to the left and to the right part.
- H 2.28. By multiplying A on the left by $E = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}$, we reduce A to echelon form: $EA = U$. Then $A = E^{-1}U$ is the LU decomposition.
- H 2.29. If $E_3E_2E_1A = U$ where U is in row echelon form, then $A = (E_1^{-1}E_2^{-1}E_3^{-1})U = LU$ gives the LU -decomposition.
- H 2.30. After the first two standard row operations, note that rows 2 and 3 will be in the wrong order. So start by multiplying A by a permutation matrix P that switches rows 2 and 3, then PA admits an LU -factorization, proceed as usual.
- H 2.31. Follow the standard method but remember to perform all operations in \mathbb{Z}_7 . To reduce A to echelon form you need to add $-\frac{5}{4}$ of the first row to the second, what is $-5 \cdot 4^{-1}$ in \mathbb{Z}_7 ?
- H 2.32. Such a dependence relation can be written $Av = 0$ where v is a column vector of the coefficients λ_i .
- H 2.33. Think of how a matrix in column echelon form looks, and think what happens when reducing it to row echelon forms
- H 2.34. First use the element in position $(1, 1)$ to get zeros in positions $(2, 1)$ and $(3, 1)$ and $(4, 1)$. This takes 4 multiplications each. Then use the element in position $(2, 2)$ to get zeros in positions $(3, 2)$ and $(4, 2)$, this takes 3 more multiplications each...
- H 2.35. Recall that a Cholesky-factorization of A is $A = CC^*$ where C is lower triangular. First find the LDU-factorization $A = LDU$, then D is diagonal with positive entries, so find a diagonal matrix \tilde{D} such that $\tilde{D}^2 = D$ and let $C = L\tilde{D}$.
- H 2.36. Same hint as the previous problem
- H 2.37. For the first part it is enough to show that CC^* is Hermitian. For the second part it suffices to find a Hermitian matrix that can not be factored as CC^* - look for the smallest possible counter-example.
- H 3.1. Think about where vectors on the lines/planes/axes are mapped to.
- H 3.2.
- H 3.3. Compute $F(x^n)$.
- H 3.4. We know that $Av = \lambda v$, take the complex conjugate. What can be said about the complex conjugate of a matrix product?
- H 3.5. Use the previous problem to obtain S, D such that $A = SDS^{-1}$.
- H 3.6. Assume that v is an eigenvector - what can be said about $R^4(v)$?
- H 3.7. Find the eigenvalues and eigenvectors of P , it may help to introduce some new notation for the eigenvalues.
- H 3.8. Consider $(A^T A)X$ when $X \in \mathbb{R}^4$ parallel with $(1, 2, 3, 4)$. Then consider the case when $AX = 0$.

- H 3.9. Each element of the matrix is either 0 or 1. A 2×2 matrix is invertible if and only if each row is nonzero and the rows are parallel.
- H 3.10. Compute the characteristic polynomials $p_A(t) \in \mathbb{Z}_5[t]$ and $p_B(t) \in \mathbb{Z}_{11}[t]$, and find their zeroes in the respective field. Eigenvectors can be found by the standard method.
- H 3.11. Look up how additive/multiplicative identities and inverses are defined in the field axioms (F1)-(F9).
- H 3.12. Write $X_n := \begin{pmatrix} a_{n+1} \\ a_n \end{pmatrix}$. Then $X_0 = \begin{pmatrix} 6 \\ 4 \end{pmatrix}$ and $X_{n+1} = AX_n$ where $A = \begin{pmatrix} 2 & -2 \\ 1 & 0 \end{pmatrix}$. It follows that $X_n = A^n X_0$. This becomes easy to evaluate if X_0 is written as a linear combination of eigenvectors for A .
- H 3.13. Find A^2 and A^4 , the definitions says that $p(A) = A^4 + 2A^2 - 5A + 3I$. Do the same for B .
- H 3.14. The matrices (vectors) I, A, A^2 are linearly dependent. To find the dependence relation, express these three matrices in the standard basis $\mathbf{e} = (e_{11}, e_{12}, e_{21}, e_{22})$, and use the standard method.
- H 3.15. Take a small (say 2×2) matrix A and write down the matrix $A - tI$, what happens when we try to replace t by A ?
- H 3.16. $p_A(t) = -(t+2)^2(t-3)$, inserting A directly into this factorized form makes the computation trivial.
- H 3.17. Factor $p(t)$ and note that $p(t)$ is a factor in $q(t)$.
- H 3.18. Find the eigenvalues of A by factoring p_A , then use the spectral mapping theorem.
- H 3.19. $m(t)$ is monic and divides $p(t)$, and each root of $p(t)$ is a root of $m(t)$.
- H 3.20. Divide $t^3 + 2t + 2$ by $t^2 - 1$
- H 3.21. For (a) the characteristic polynomial is $(t-3)^2$, so the only possible minimal polynomials are $m(t) = t-3$ or $m(t) = (t-3)^2$, remember that $m(A)$ should be the zero matrix. In (b) and (c) the matrix is triangular so the characteristic polynomial is easy to find.
- H 3.22. The minimal polynomial should annihilate T , meaning that $m_T(T)$ is the zero-operator. Use the fact that matrices satisfy $(A^T)^T = A$.
- H 3.23. What does the map R^4 do?
- H 4.1. Try visualizing the Jordan-blocks on the diagonal. Each block should have the same value one the diagonal and ones on the super-diagonal
- H 4.2. The matrices are all on Jordan-form, so the algebraic multiplicity of λ is the number of occurrences of λ on the diagonal. Each Jordan block with λ 's on the diagonal corresponds to a single eigenvector for that eigenvalue, so the number of Jordan blocks corresponding to the eigenvalue λ is the geometric multiplicity of λ
- H 4.3. Consider each of the possible block-configurations separately.
- H 4.4. The eigenvalues 3 and -1 can be treated independently. For the eigenvalue -1 there are three possibilities of the sizes of Jordan blocks: $(3), (2, 1), (1, 1, 1)$ corresponding to integer partitions of 3.
- H 4.5. Different Jordan forms correspond to decreasing sequences that add up to 6
- H 4.6. You need at least size 4×4 for (a), and 7×7 for (b). Take all eigenvalues as the same. Look for A and B with different partitions of Jordan blocks. Remember that for a given eigenvalue, the minimal polynomial determines the size of the largest Jordan block, and the geometric multiplicity is the number of Jordan blocks.
- H 4.7. We know that any matrix A can be Jordanized: $A = SJS^{-1}$ where J is in Jordan form.
- H 4.8. What eigenvalues can N have? Jordanize N , with $SNS^{-1} = J$.
- H 4.9. Try the 2×2 and the 3×3 case first.

- H 4.10. The given data decoded says that the matrix is of size 7×7 with zeros on the diagonal, the largest Jordan block should be of size 4, and there should be three Jordan blocks.
- H 4.11. Remember, a vector represented by a dot is in $\ker(F)$ if it is mapped to zero, it is in $\ker(F^2)$ if it lands in zero after moving along the arrows in one or two steps, and so on. The image consists of dots that are pointed to by some arrow, etc.
- H 4.12. F is clearly nilpotent, so zero is the only eigenvalue. For what n is F^n first zero? To get started on the string basis, consider where $e_1 - e_4$ is mapped.
- H 4.13. Find $\ker(N)$ and $\ker(N^2)$, think about what the chains should look like (or follow the standard algorithm).
- H 4.14. Same hint as the last problem.
- H 4.15. Follow the standard algorithm, it is easy in this case to find the characteristic polynomial.
- H 4.16. In this case, finding the characteristic and minimal polynomials will give enough information to determine the Jordan form. Find a basis for each eigenspace and extend to string bases.
- H 4.17. Since 2 is the only eigenvalue, the matrix $A - 2I$ will be nilpotent. A string basis for $A - 2I$ will be a Jordan basis for A , so just follow Algorithm 4.2.9.
- H 4.18. Proceed as usual, the only difference in this case is that the characteristic polynomial have non-real roots.
- H 4.19. You know the Jordan form and the first column of the matrix S , complete the matrix S in any way, and use that $SJS^{-1} = A$.
- H 4.20. The only eigenvalue is clearly 1. So the operator $F - \text{id}$ is nilpotent.
- H 4.21. Without loss of generality you can assume that A is in Jordan form.
- H 4.22. Remember that A and B are called similar if there exists T such that $TAT^{-1} = B$. Combine that $A = S_1JS_1^{-1}$ and that $B = S_2JS_2^{-1}$.
- H 4.23. Either find the matrix form of D relative to the given basis for U and go for there. Or think first, perhaps you can guess the eigenvalues and (generalized) eigenvectors.
- H 4.24. Pick a new basis for which the matrix of the operator is in Jordan form, then the same statement is simple to prove.
- H 4.25. Suppose $X^2 = J_n(0)$ where $X = SJS^{-1}$ is the Jordan form of X , consider what the matrix J must look like.
- H 4.26. Consider the two Jordan blocks separately. Note that while A^n is similar to J^n , this latter matrix may not be in Jordan form, in particular A^2 is not in Jordan form.
- H 4.27. Follow the standard algorithm, but remember that all numbers, polynomial coefficients, and matrix elements are in \mathbb{Z}_3 . In particular, the matrices S and J should also belong to $\text{Mat}_3(\mathbb{Z}_3)$.
- H 4.28. The only way such a matrix can fail to have a Jordan-decomposition is if the characteristic polynomial doesn't factor completely over \mathbb{Z}_2 . The characteristic polynomial has degree 2 and there are only four polynomials of degree 2 in $\mathbb{Z}_2[x]$, namely x^2 , $x^2 + 1$, $x^2 + x$, $x^2 + x + 1$. Only one of them can not be factored, find a matrix which has this characteristic polynomial.
- H 4.29. You can consider each of the two Jordan blocks separately, write each block as $J = \lambda I + N$ and use the binomial theorem.
- H 4.30. Recall the definition: $e^A := \sum_{k=0}^{\infty} \frac{A^k}{k!}$. In A , consider the two diagonal positions of e^A separately. B is nilpotent so all but the first terms of the sum disappear.
- H 4.31. Write $J = D + N$ where D is diagonal and N is nilpotent. Verify that N and D commute, so that Show that $e^{D+N} = e^D e^N$. Compute the right side by considering the blocks separately.

- H 4.32. $Av = \lambda v$, so $e^A v = (I + A + \frac{A^2}{2} + \dots)v = \dots$
- H 4.33. Recall that $\sin(x) = x - \frac{x^3}{3!} + \frac{x^5}{5!} - \dots$ and $\cos(x) = 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \dots$. A_1 is nilpotent, use Proposition 4.4.5 for A_2 and A_3 and consider the Jordan blocks separately.
- H 4.34. Follow the proof of the fact that $\frac{d}{dt} e^{At} = A e^{At}$. For a differential equation, if $X = \sin(At)C$, what is $X''(t)$?
- H 4.35. Trace and determinant is basis-independent, so you can assume that A is a Jordan matrix without loss of generality. Consider the diagonal elements of e^A .
- H 4.36. Follow the example with rabbits/foxes. Start by introducing $X_n = \begin{pmatrix} a_n \\ b_n \end{pmatrix}$ and expressing the system in matrix form, then Jordanize the coefficient matrix.
- H 4.37. The eigenvalues are 2 and -1 , so study the kernels of $(A - 2I)^k$ and of $(A + I)$ to find (generalized) eigenvectors forming Jordan chains
- H 4.38. The system can be written $X'(t) = AX(t)$ where A is the same as in the previous problem!
- H 5.1. Recall that the norm is defined via the inner product $\|v\| := \sqrt{\langle v, v \rangle}$, and that vectors are said to be orthogonal if their inner product is zero.
- H 5.2. $\|x^2 + x + 1\|^2 = \langle x^2 + x + 1, x^2 + x + 1 \rangle$
- H 5.3. Use symmetry, linearity in the first argument, and symmetry again.
- H 5.4.
- H 5.5. The axioms for an inner product are easy to verify. In the first case IP-space, $\|v\|_1 = \sqrt{\langle v, v \rangle}$, in the second it is $\|v\|_2 = \sqrt{\langle v, v \rangle}$. Recall that the angle between two vectors u and v of a real IP-space is defined as the number θ between 0 and π satisfying $\|u\| \cdot \|v\| \cos(\theta) = \langle u, v \rangle$.
- H 5.6. Is it possible to define an inner product such that (u, v) is an ON-basis?
- H 5.7. Expand the left side using the fact that $\|u \pm v\|^2 = \langle u \pm v, u \pm v \rangle$.
- H 5.8. For the triangle inequality, expand $\|u + v\|^2 = \langle u + v, u + v \rangle$ using sesqui-linearity, you will have to use Cauchy-Schwarz in one step.
- H 5.9. Recall that $\|(x, y)\|_{\max} = \max\{|x|, |y|\}$. It is enough to show that the norm does not satisfy the parallelogram law.
- H 5.10. \mathcal{I} is clearly a subset of \mathcal{F} . Is the sum of two inner products an inner product? Is a scalar times an inner product still an inner product?
- H 5.11. The distance is defined as $\|(2, 1, 3) - (4, -2, 6)\| = \|(2, 1, -4)\|$, so compute this for the various definitions of the norm.
- H 5.12. $\|A\|_F = \sqrt{\langle A, A \rangle_F}$ where $\langle A, B \rangle_F = \text{tr}(AB^*)$. $\|A\|_\sigma = \max\{|\lambda| \mid \lambda \in \sigma(A)\}$.
- H 5.13. Normalize v and write $v = \|v\| \left(\frac{1}{\|v\|} v\right) \|v\|$.
- H 5.14. Choose p such that the identity $\cos^2(t) + \sin^2(t) = 1$ appear in the norm calculation.
- H 5.15. Find an example of vectors violating the triangle inequality.
- H 5.16. Look at the pictures of the unit circles in the different norms. How should the circles be scaled to contain one another?
- H 5.17. Recall that a sequence of vectors v_n converges to a vector v with respect to a norm $\|\cdot\|$ if and only if $\|v_n - v\| \rightarrow 0$.
- H 5.18. Look for a counterexample.

- H 5.19. The form is symmetric and bilinear, but not positive definite.
- H 5.20. For the first part, verify that for $u_1, u_2 \in U^\perp$ and $\lambda \in \mathbb{C}$ we have $u_1 + u_2 \in U^\perp$ and $\lambda u_1 \in U^\perp$, use linearity in the first argument of the inner product. For the second part, use the positive-definiteness.
- H 5.21. The statements follow directly from the definition. They show that we can define *projections* in arbitrary inner product spaces.
- H 5.22. Show that the difference $w := v - (\langle v, e_1 \rangle e_1 + \dots + \langle v, e_n \rangle e_n)$ is zero by showing that its inner product with each of the basis vectors e_i is zero.
- H 5.23. The left hand side is $\langle u + v, u + v \rangle$.
- H 5.24. $\|v - u\|^2 = \|(v - P_U(v)) + (P_U(v) - u)\|^2$, show that $v - P_U(v) \in U^\perp$ and $P_U(v) - u \in U$, and apply Pythagorean theorem.
- H 5.25. Call the spanning vectors for u_1 and u_2 . Replace the second vector u_2 by itself minus its projection on u_1 , $u'_2 := u_2 - \frac{\langle u_2, u_1 \rangle}{\langle u_1, u_1 \rangle} u_1$. Then u_1 and u'_2 are orthogonal and span U . Normalize to get an ON-basis. Solve a linear system to get the last basis vector.
- H 5.26. Project e^x onto \mathcal{P}_1 , use the basis from Example 5.4.2.
- H 5.27. If f is odd, $\int_{-a}^a f(x) dx = 0$. If g is even, $\int_{-a}^a g(x) dx = 2 \int_0^a g(x) dx$. Eulers formulas may also help.
- H 5.28. The set $\{1, \cos(x), \sin(x)\}$ is an orthogonal set of vectors according to the last problem, so project f onto each of these vectors.
- H 5.29. Applying Gram-Schmidt to the columns of A to get an ON-basis of \mathbb{C}^2 , put these vectors as columns in Q and find $R = Q^*QR = Q^*A$.
- H 5.30. Apply Gram-Schmidt to the columns.
- H 5.31. Use Gram-Schmidt to find orthogonal vectors spanning the same set as the columns. To obtain a unitary square matrix Q , extend this basis to an ON-basis.
- H 5.32. To show homogeneity $F(\lambda w) = \lambda F(w)$ we can calculate $\langle v, F^*(\lambda w) \rangle = \langle F(v), \lambda w \rangle = \bar{\lambda} \langle F(v), w \rangle = \bar{\lambda} \langle v, F^*(w) \rangle = \langle v, \lambda F^*(w) \rangle$, so $F^*(\lambda w) = \lambda F^*(w)$. Additivity can be proved similarly.
- H 5.33. For the first part, $\langle u, (F^* \circ G^*)(v) \rangle = \langle u, F^*(G^*(v)) \rangle = \langle F(u), G^*(v) \rangle = \dots$
- H 5.34. Let $\{v_i\}$ and $\{w_j\}$ be the ON-bases, then $\langle v_i, F^*(w_j) \rangle = \langle F(v_i), w_j \rangle$
- H 5.35. Since the matrices are given with respect to an ON-basis: Self-adjoint maps correspond to symmetric matrices, unitary maps corresponds to matrices whose columns form an ON-basis, and normal maps have matrices satisfying $NN^* = N^*N$.
- H 5.36. Try a scaled rotation-matrix.
- H 5.37. Self adjoint: $A^* = A$, Unitary: $AA^* = I$, Normal: $AA^* = A^*A$.
- H 5.38. Show that if λ is an eigenvalue, then $|\lambda| = 1$.
- H 5.39. $\langle F(u), F(v) \rangle = \langle u, F^*(F(v)) \rangle \dots$
- H 5.40. This follows easily from the definitions
- H 5.41. For (a), the key step is $F^*(w) = 0 \Leftrightarrow \langle v, F^*(w) \rangle = 0 \forall v \in V$.
- H 5.42.
- H 5.43.
- Show that $(A + \lambda I)(A + \lambda I)^* = (A + \lambda I)^*(A + \lambda I)$ by expanding both sides.
 - It suffices to show that $\|A^*v\|^2 = \langle A^*v, A^*v \rangle = \langle Av, Av \rangle = \|Av\|^2$.

- c) Combine parts (a) and (b): $Av = \lambda v \Leftrightarrow v \in \ker(A - \lambda I) \Leftrightarrow v \in \ker((A - \lambda I)^*) = \ker(A^* - \bar{\lambda}I) \Leftrightarrow A^*v = \bar{\lambda}v$.
- d) Assume $Au = \lambda u$ and $Av = \mu v$. Expand both sides of the equality $\langle Au, v \rangle = \langle u, A^*v \rangle$ using part (c) and subtract.

H 5.44.

- a) Finding one counterexample suffices.
- b) Write $A = \frac{A+A^*}{2} + i\frac{i(A^*-A)}{2} = B + Ci$, and show that $X^*AX > 0$ for all nonzero complex $X \in \mathbb{C}^n$ implies that $C = 0$.
- c) Try an upper triangular 2×2 -matrix.
- d) In (b) we allow X to have complex entries...

H 5.45. A must be Hermitian, then use Sylvester's criterion.

H 5.46. Show that all principal minors are positive if and only if all the diagonal elements of D are positive.

H 5.47. If the $k \times k$ principal minor is negative, consider columns X on block form $\begin{pmatrix} X' \\ 0 \end{pmatrix}$ where X' is of size $k \times 1$.

H 5.48. A is normal and can be orthogonally diagonalized by the spectral theorem $A = UDU^* = \begin{pmatrix} u_1 & u_2 \end{pmatrix} \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} \begin{pmatrix} u_1^* \\ u_2^* \end{pmatrix} = \lambda_1 u_1 u_1^* + \lambda_2 u_2 u_2^*$, where the matrices $u_1 u_1^*$ and $u_2 u_2^*$ are orthonormal.

H 5.49. For the second part, start by showing that the sum of positive semi-definite matrices are positive semi-definite.

H 5.50. A is symmetric hence normal, use the standard method to diagonalize A as $A = UDU^*$ with U unitary. Then $\sqrt{A} := U\sqrt{D}U^*$

H 5.51. Recall that $A^+ := (A^*A)^{-1}A^*$, simplify this when A is invertible

H 5.52. Recall that the least square solution to the system $AX = b$ is given by A^+b where $A^+ = (A^*A)^{-1}A^*$.

H 5.53. For (b), note that if A has linearly dependent rows, then A^* has linearly independent columns. For (c), A^+ must be 2×2 , make an ansatz.

H 6.1. Note that $A > B \Leftrightarrow A - B > 0$. Also note that $X \geq 0$ and $X \neq 0$ does not imply $X > 0$.

H 6.2. Look for 2×2 -matrices, remember Sylvester's criterion for positive-definiteness.

H 6.3. Look for a 2×2 -example.

H 6.4. First prove that $\|AB\|_F = \|A\| \cdot \|B\|$ when A is a column and B is a row. Then for the general case, write A as a block matrix of columns and write B as a block matrix of rows, multiply and use the triangle inequality.

H 6.5. Since $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ is positive and it is an eigenvector, it must be a Perron-vector corresponding to the largest eigenvalue.

H 6.6. Express v as a linear combination of eigenvectors, apply A and use the triangle inequality.

H 6.7. The matrix has form $A = \begin{pmatrix} a & 1-b \\ 1-a & b \end{pmatrix}$ for some fixed $a, b \in (0, 1)$. Show explicitly that 1 is an eigenvalue, and show that the other eigenvalue is smaller. Recall that $\text{tr}(A)$ is the sum of the eigenvalues.

H 6.8. Apply the same method as in the previous problem.

H 6.9. Show first that A and A^T have the same spectrum.

H 6.10. Consider 2×2 -counterexamples

H 6.11. In a solved Sudoku puzzle, each row and each column contains all of the the integers 1 – 9 in some order. What is $A \cdot (1, 1, \dots, 1)^T$?

H 6.12. According to the definition, the circle centers are the diagonal elements of the matrix, and the radius is the sum of absolute values of the other elements at each row.

H 6.13. We need to show that any eigenvalue λ lies inside one Gershgorin disc. If $Av = \lambda v$, then coordinate-wise this says that for each k we have

$$\sum_{j=1}^n a_{kj} v_j = \lambda v_k.$$

Pick a k for which $|v_k|$ is maximal, and apply the triangle-inequality.

H 6.14. Find the maximum eigenvalue using the standard algorithm.

H 6.15. Recall that the matrix is irreducible if and only if it is possible to walk between any pair of nodes in the graph.

H 6.16. Draw the associated graphs, are they strongly connected? Is there a p such that there is a walk between any pair of nodes of length exactly p ?

H 6.17. The matrix is clearly irreducible since we can walk between any pair of nodes. Write down the adjacency-matrix and compute its powers. It is also possible to do this without matrices: is there an integer p such that there exists a walk of length p between any pair of nodes?

H 6.18. Consider the graph associated to A .

H 6.19. Recall that $(A^k)_{ij} > 0$ if and only if there is a walk of length exactly k from v_j to v_i in the associated graph.

H 6.20. Start by trying to draw such a graph, think about a square.

H 6.21. Decompose v as a sum of (generalized) eigenvectors: $v = a_1 v_1 + \dots + a_n v_n$ and apply A^k . For the second part, consider a Jordan block of size 2.

H 6.22. Consider a Jordanization of A , what can be said about the limit of the Jordan blocks?

H 6.23. The dampened PageRank matrix will be $0.7 \cdot \begin{pmatrix} 0 & 0 & 0 \\ 0.5 & 0 & 0 \\ 0.5 & 1 & 0 \end{pmatrix} + 0.3 \cdot \frac{1}{3} \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}$, find the Perron-vector for this matrix (you probably need to use a computer for this, the eigenvalues are not nice)

H 6.24. Show that the sums in the rows of the PageRank-matrix will be constant.

H 6.25. The PageRank matrix P is A but with normalized columns, apply P repeatedly to an arbitrary vector to find the ranking vector.

H 6.26. Find the Perron-vector of $A = \begin{pmatrix} 0.9 & 0.5 \\ 0.1 & 0.5 \end{pmatrix}$.

H 6.27.

H 6.28. The transition matrix is $A = \begin{pmatrix} & 0.5 & & \\ 1 & & 0.5 & \\ & 0.5 & & 0.5 \\ & & 0.5 & 1 \end{pmatrix}$ (with zeros removed), find an eigenvector of A with eigenvalue 1.

H 7.1. Perhaps calculation is not necessary, what are the eigenvalues of A^*A ? What vectors satisfy $Av_i = \sigma_i u_i$?

H 7.2. Follow the standard algorithm, start with orthogonally diagonalizing A^*A .

- H 7.3. Follow the algorithm, the matrix $A^*A = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$ is already diagonal, but remember that order of the eigenvalues in Σ should be decreasing.
- H 7.4. This example illustrates that SVD-computation can be tricky even for small matrices. Follow the standard method, since $\det(A^*A) = 1$, the product of the eigenvalues is 1, so introduce some notation α and α^{-1} for the eigenvalues to make it easier to write down the answer. Compare your answer to the figure in the text.
- H 7.5. Follow the standard algorithm, start with orthogonally diagonalizing A^*A .
- H 7.6. Find the SVD of A^* first (then you only have to diagonalize a 2×2 -matrix instead of a 3×3 -matrix.)
- H 7.7. In looking for the eigenvalues of A^*A , recall that the trace is the sum of the eigenvalues. Recall that you only need to find singular vectors for the nonzero eigenvalues in the compact SVD.
- H 7.8. You are looking for vectors satisfying $Av_i = \sigma_i u_i$ (and $A^*u_i = \sigma_i v_i$), consider where the standard basis e_1, e_2, e_3 is mapped.
- H 7.9. Pick an orthonormal basis for \mathcal{P}_1 and consider the SVD of the matrix of D with respect to this basis. It is important that the basis is orthonormal for this to work, 1 and x won't do.
- H 7.10. Let $J = J_n(0)$. Then J^*J will already be diagonal.
- H 7.11. Look at $Av = \lambda v$ and use the SVD form of A .
- H 7.12. Can you write down the inverse explicitly given the SVD of A ?
- H 7.13.
- H 7.14. Use the spectral theorem to write $A = UDU^*$ where D is diagonal with the eigenvalues on the diagonal, then find the eigenvalues of A^*A in terms of this decomposition.
- H 7.15. The singular values of λA are the square roots of the eigenvalues of $(\lambda A)^*(\lambda A) = \lambda^* \lambda A^*A$.
- H 7.16. The problem is that $\Sigma^{-1} = \text{diag}(\frac{1}{\sigma_1}, \dots, \frac{1}{\sigma_n})$, so the diagonal elements are in *increasing* order.
- H 7.17. Pick ON-bases in \mathbb{C}^n and \mathbb{C}^m such that $[A] = \Sigma$. Then for $v = (x_1, \dots, x_n)^T$ we get $\|Av\| = \|\sigma_1 x_1, \dots, \sigma_n x_n\| = \dots \leq \sqrt{\sigma_1^2 \|v\|^2}$ so since $\|v\| = 1$ the result follows. Equality holds if and only if v is an singular vector for the maximal singular value σ_1 .
- H 7.18. Use the characterization that $\|Av\| \leq \sigma_1 \|v\|$ with equality if and only if v is a right singular vector for the first singular value σ_1 .
- H 7.19. Let $\sigma_1, \dots, \sigma_r$ be the singular values of A , then it remains to show that
- $$\sigma_1 \leq \sqrt{\sigma_1^2 + \dots + \sigma_r^2} \leq \sigma_1 + \dots + \sigma_r.$$
- H 7.20. First express the SVD explicitly as a sum $A = \sigma_1 u_1 v_1^* + \sigma_2 u_2 v_2^* + \sigma_3 u_3 v_3^*$. The low rank approximations are obtained by deleting one or two terms from this sum. For computing the norms $\|A - A_{(i)}\|$ you don't need matrix computations, the norms can be expressed in terms of the singular values of $A - A_{(i)}$.
- H 7.21. For the SVD of rank k we need to store k columns, k rows, and k singular values, compared to the $3000 \cdot 2000$ original pixel values.
- H 7.22. Consider diagonal matrices, for such matrices the singular values are the absolute values of the diagonal elements.
- H 7.23. Take for example $A = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$. Its SVD is $A = U\Sigma V^*$ where $U = V = I$ and $\Sigma = A$. How can we form two different optimal approximations of rank 2?

- H 7.24. Recall that $\langle A_i, A_j \rangle_F = \text{tr}(A_i^* A_j)$, simplify this expression using the fact that the vectors u_i are orthonormal, and that the v_i are orthonormal. Also remember that $\text{tr}(AB) = \text{tr}(BA)$ whenever the products AB and BA both are defined.
- H 7.25. Recall that the nuclear norm is the sum of the singular values. Follow the proof of the Eckart-Young-Mirsky theorem in Frobenius norm case; there we showed that when $\text{rank}(B) \leq k$ we have $\sigma_i(A - B) \geq \sigma_{i+k}(A)$.
- H 7.26. Almost no computation is needed: A has rank 1 since $A = \begin{pmatrix} 3 \\ 4 \end{pmatrix} \begin{pmatrix} 1 & 2 \end{pmatrix}$, what does this say about the only nonzero singular value?
- H 7.27. See Example 7.2.3 where the SVD of the same matrix was calculated.
- H 7.28. Any rank 1 matrix A can be expressed as $A = CR$, as a product of a column and a row. What does this say about the first (and only) singular value?
- H 7.29. Plug in the general version of the pseudo-inverse from the SVD into the expression $(A^*A)^{-1}A^*$. Where is the statement about linearly independent columns used?
- H 7.30. Use the definition via the compact SVD. What is $(\tilde{\Sigma}^{-1})^*$?
- H 7.31. Recall that if $A = \tilde{U}\tilde{\Sigma}\tilde{V}^*$ is a compact SVD, then $A^+ := \tilde{V}\tilde{\Sigma}^{-1}\tilde{U}^*$, where $\tilde{\Sigma}^{-1}$ has the inverse of the singular values of A on the diagonal.
- H 7.32. In general, a linear map P on an inner product space is an orthogonal projection if $P^2 = P$ and $P = P^*$, verify these conditions for $P = AA^+$ and also check that $AA^+v = v$ when $v \in \text{Im}(A)$.
- H 7.33. Note that if $A = U\Sigma V^*$ is a full SVD of the $m \times n$ matrix A with $\text{rank}(A) = r$, then $A^+ = V\hat{\Sigma}U^*$ where $\hat{\Sigma}$ is the $n \times m$ matrix where the first diagonal elements are $\frac{1}{\sigma_1}, \dots, \frac{1}{\sigma_r}$ and all other elements are zero. Then $A^+A = V\hat{\Sigma}U^*U\Sigma V^* = V\hat{\Sigma}\Sigma V^*$, determine when this is the identity. (What does the matrix $\hat{\Sigma}\Sigma$ look like?)
- H 7.34. You need to show that $x = A^+b$ satisfies the normal equations $A^*Ax = A^*b$. Use an SVD-expression for A .
- H 7.35. Show that any solution to $A^*Ax = A^*b$ can be written $x = A^+b + w$ where $w \in \ker(A)$, and show that A^+b is orthogonal to w . For the last part it might help to use the fact that $\text{Im}(F) \perp \ker(F^*)$.
- H 7.36. Since A is diagonal, the singular values are the absolute values of the diagonal elements.
- H 7.37. First show that if $A = U\Sigma V^*$, then $|\det(A)| = \det(\Sigma)$, recall also that $\|A\|_F = \sqrt{\sigma_1^2 + \sigma_2^2}$.
- H 7.38. A is Hermitian, so the singular values are the absolute values of the eigenvalues of A (which are real).
- H 7.39. Look back at the line where we showed that $\frac{\|\Delta x\|}{\|x\|} \leq \frac{\|\Delta b\|}{\|b\|}$. When does equality hold?
- H 7.40. Perhaps you can guess what the matrices U and P to pick? Otherwise follow the standard method.
- H 7.41. Write down the standard matrix of F and proceed as usual, it is not too hard to see what an SVD of F is without calculation.
- H 7.42. Use the standard method, start by finding an orthogonal diagonalization of A^*A .
- H 7.43. We know that the polar factorization $A = PU$ is unique when A is invertible, so look for a small example where 0 is a singular value.
- H 7.44. One can explicitly write down such a factorization via the SVD of A , compare with the argument for the regular polar factorization. Or alternatively just take the Hermitian conjugate of a polar factorization of A^* .
- H 7.45. Follow the steps of the movie ratings example in the notes.

H 7.46. For the total least square solution, let subtract the mean c from each data point and put the centered data in a 4×2 -matrix A . The first right singular vector of $A^T A$ is the direction of the affine subspace we seek. For the regular least square method, make an ansatz $y = kx + m$, insert the four points to get four linear equations in k and m , and find the least square solution to this system.

H 7.47. Center the points by subtracting the mean and put the new points as rows in a matrix A . The directions of the line is given by the first right singular vector (eigenvector of $A^T A$ of highest eigenvalue), and the directions spanning the plane are the two first right singular vectors.

H 7.48. Take for example just three points in \mathbb{R}^1 . An affine subspace of dimension 0 is just a fourth point $S = s$.

H 8.1. To express it in the standard basis, for the first term, note that

$$\begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} = (e_1 - e_2) \otimes (e_1 + 2e_2) = e_1 \otimes e_1 + 2e_1 \otimes e_2 - e_2 \otimes e_1 - 2e_2 \otimes e_2,$$

where we used the bilinear properties of tensors. It is possible to express v as a pure tensor, make an ansatz or note that the first component of the two terms of v are parallel.

H 8.2. The first and third components are parallel in the three terms.

H 8.3. You need at least three terms, look for an example where $u \otimes u + v \otimes v + w \otimes w$ is pure in $\mathbb{R}^2 \otimes \mathbb{R}^2$.

H 8.4. You are looking for three vectors $v_1 = \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}$, $v_2 = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix}$, and $v_3 = \begin{pmatrix} c_1 \\ c_2 \\ c_3 \end{pmatrix}$ such that the array element at position (i, j, k) is $a_i b_j c_k$ for all indices of the array. Start by expressing the front matrix as an outer product $\begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \begin{pmatrix} b_1 & b_2 & b_3 \end{pmatrix}$.

H 8.5. Keeping two indices fixed, all the "tube-vectors" should be parallel. You could also make an ansatz

$$\begin{pmatrix} a \\ b \end{pmatrix} \otimes \begin{pmatrix} c \\ d \\ e \end{pmatrix} \otimes \begin{pmatrix} f \\ g \end{pmatrix}.$$

H 8.6. Start by instead considering the tensor s with $[s]_{ijk} = (-1)^{i+j+k}$.

H 8.7. Write down an explicit basis for each space, to see what pairs of bases can be paired up most easily.

H 8.8. Recall that $\dim(U \otimes V) = \dim U \cdot \dim V$. What is a prime factorization of 5400?

H 8.9. Look at the dimension of each side. For an explicit isomorphism, consider that the left side is spanned by elements of form $\begin{pmatrix} a \\ b \end{pmatrix} \otimes v$, and elements on the right has form (v, v') , how can we pair them up?

H 8.10. It suffices to verify that the dimension of each side agree, make an ansatz for the dimensions of U, V, W and compare the sides.

H 8.11. Recall that $\text{Hom}(U, V)$ is just the space of linear maps from U to V . There is a natural or *canonical* choice of an isomorphism between the two spaces. Each $\varphi \in \text{Hom}(U \otimes V, W)$ corresponds to a bilinear map $U \times V \rightarrow W$, so to produce an explicit isomorphism between the given spaces you have to map each such φ to a map from U to $\text{Hom}(V, W)$, think about keeping one index fixed.

H 8.12. Hints:

- Express $(7, 4)$ as a linear combination of e_1 and e_2 , apply the dual basis vectors and remember that $e_i^*(e_j) = \delta_{ij}$.
- Express $(5, 3)$ in e_1 and e_2 , the maps e_i^* are linear, so apply the distributive properties to evaluate the expression.
- Make an ansatz $g = xe_1^* + ye_2^*$, express a general vector (a, b) in the basis e_1, e_2 and apply the ansatz using that $g(a, b) = a + b$.

- H 8.13. Make an ansatz for an arbitrary $p(x) \in \mathbb{P}_3$, apply f and look at the coefficients - which linear combination of the dual basis vectors has the same effect on p ?
- H 8.14. Parts 1 and 2 can be done by explicit calculation. For the last part you may use that for any vector $v \in V$ there exists some linear functional f with $f(v) \neq 0$.
- H 8.15. First find the dimension of V .
- H 8.16. The space of (p, q) -tensors on V is the vector space $V^{\otimes p} \otimes (V^*)^{\otimes q}$, what is its dimension?
- H 8.17. There are two possible contractions, $C_{1,1}$ and $C_{2,1}$, these are both linear maps $\mathbb{R}^2 \otimes \mathbb{R}^2 \otimes (\mathbb{R}^2)^* \otimes \mathbb{R}^2 \rightarrow \mathbb{R}^2$. Here $C_{1,1}$ combines the first and third factor into a scalar, while $C_{2,1}$ combines the second and third to a scalar. Both contractions are linear.
- H 8.18. Write F as a linear combination of $e_1 \otimes e_1^*$, $e_1 \otimes e_2^*$, $e_2 \otimes e_1^*$, $e_2 \otimes e_2^*$. The contraction $C_{1,1}$ maps a pure tensor $v \otimes f$ of $\mathbb{R}^2 \otimes \mathbb{R}^2$ to $f(v) \in \mathbb{R}$.
- H 8.19. Look at the previous exercise. Choose a basis (e_1, \dots, e_n) of V , if F has matrix $A = (a_{ij})$ with respect to this matrix, then F corresponds to the tensor $w = \sum_{i=1}^n \sum_{j=1}^n a_{ij} e_i \otimes e_j^*$. What is the contraction of w ?
- H 8.20. Think of the familiar rule from linear algebra; with respect to an ON-basis (such as the standard basis in \mathbb{R}^3) we have
- $$(x_1, x_2, x_3) \times (y_1, y_2, y_3) = (x_2 y_3 - x_3 y_2, x_3 y_1 - x_1 y_3, x_1 y_2 - x_2 y_1).$$
- Also note that $e_i^*(x_1, x_2, x_3) = x_i$.
- H 8.21. Let $F: V \rightarrow V$ and $G: W \otimes W \rightarrow V \otimes W$, so $F \otimes G: V \otimes W \rightarrow V \otimes W$. The inverse should be a map $V \otimes W \rightarrow V \otimes W$ whose composition (in either order) with $F \otimes G$ becomes the identity on $V \otimes W$, you can write down the inverse explicitly.
- H 8.22. Remember that $AB = (a_{ij}B)$, each 2×2 -block is a multiple of the right matrix in the product.
- H 8.23. Write down the matrices of $[S]$ and of $[R]$ with respect to the standard basis (e_1, e_2) of \mathbb{R}^2 . The matrix of $S \otimes T$ with respect to the basis in the problem is then given by the Kronecker product $[S] \otimes [G]$.
- H 8.24. You need a sum of Kronecker products, first write $\begin{pmatrix} r_1 & 0 & \cdots & 0 \end{pmatrix}$ as a Kronecker product.
- H 8.25. Both F and G are diagonalizable since they have enough distinct eigenvalues. If $F(v) = \lambda v$ and $G(w) = \mu w$, what is $(F \otimes G)(v \otimes w)$?
- H 8.26. The trace of a map is the sum of its eigenvalues.
- H 8.27. One direction is obvious. For the other direction, assume that $a_{ij} \neq 0$ and $b_{kl} \neq 0$, and show that $A \otimes B$ has a nonzero element.
- H 8.28. The rank of a diagonalizable matrix is the number of nonzero eigenvalues, counting multiplicities.
- H 8.29. Write down the product in both orders and use the property $(AB) \otimes (CD) = (A \otimes C)(B \otimes D)$ of the Kronecker product.
- H 8.30. Let $Av = \lambda v$ and $Bw = \mu w$. Multiply $A \otimes I_n + I \otimes B$ by $v \otimes w$ (here, v and w are column-matrices, so their Kronecker product is also a column).
- H 8.31. Use the previous exercise to construct a matrix which has $\sqrt{2} + \sqrt{3}$ as an eigenvalue. The characteristic polynomial of this matrix will have the desired property.
- H 8.32. Given algebraic α and β , you need to show that there is a polynomial p with integer coefficients for which $p(\alpha + \beta) = 0$.
- H 8.33. Let A have size $m \times m$ and let B have size $n \times n$. Then both $A \otimes B$ and $B \otimes A$ have shape $mn \times mn$. By the definition of the Kronecker product, the matrix $A \otimes B$ has the element $a_{ij}b_{kl}$ in position $(i + n(k-1), j + n(l-1))$. Analogously, the matrix $B \otimes A$ has the same element $a_{ij}b_{kl}$ in position $(k + m(i-1), l + n(j-1))$. How should the rows and columns be permuted such that the entries of the two matrices match?

- H 8.34. A square matrix A is unitary if $A^*A = I$.
- H 8.35. Find and SVD of $A \otimes B$ in terms of the SVDs of A and of B
- H 8.36. Recall that $A \sim B$ means that there exists S such that $SAS^{-1} = B$.
- H 8.37. The rank of a matrix is the number of nonzero singular values, counting multiplicities.
- H 8.38. $U' \otimes V'$ is obviously closed under addition and scalar multiplication since it is a span of vectors. For the second questions, what are the possible dimensions of $U' \otimes V'$ when $U = V = \mathbb{R}^2$
- H 8.39. One direction is easy: Apply $A \otimes B$ to $u \otimes v + u' \otimes v'$ where $u \in \ker(A)$ and $v' \in \ker(B)$. For the other direction, use a dimension argument using the rank-nullity theorem. Recall that in general, if S and S' are subspaces of some vector space, we have $\dim(S + S') = \dim(S) + \dim(S') - \dim(S \cap S')$.
- H 8.40. Assume first that all the b_i are zero, what does the composition F look like?
- H 9.1. Consider which of the field axioms (F1)-(F9) hold.
- H 9.2. Consider (F1)-(F9)
- H 9.3. Prove that $(a + b)^p = a^p + b^p$ modulo p , use the binomial theorem.
- H 9.4. Combine the field axioms (F1)-(F9)
- H 9.5. Assume that $0 = 1$ and combine (F3) with (F7)
- H 9.6. For (c), note that $6 = -1$ in \mathbb{Z}_7 . (d) and (e) can be solved with trial and error.
- H 9.7. Use the division algorithm to express $1 = a \cdot 53 + b \cdot 7$ for some integers a and b .
- H 9.8. Fermat's theorem states that $a^p = a$ in \mathbb{Z}_p . Multiply by a^{-1} .
- H 9.9. Use Fermat's little theorem to come up with a polynomial which has no zeroes in \mathbb{Z}_p .
- H 9.10. For a polynomial equation $p(t) = 0$ over \mathbb{Z}_2 there are only two possible values for t so it is enough to compute $p(0)$ and $p(1)$. Additionally, there are only four polynomials of degree 2 in $\mathbb{Z}_2[t]$.
- H 9.11. Make a table where you compute $x^3 - x + 1 \pmod{5}$ for $x = 0, 1, 2, 3, 4$. Make a second table where you compute $y^2 \pmod{5}$ for $y = 0, 1, 2, 3, 4$, compare the tables - for what (x, y) are the two values equal?
- H 9.12. Verify that F is closed under sums and products. Then verify the field axioms, most are obvious, focus on (F5) and (F8).
- H 9.13. Look up how additive/multiplicative identities and inverses are defined in the field axioms (F1)-(F9).
- H 9.14.
- Here, $\mathbb{Z}_3(x)$ is the field of rational functions over \mathbb{Z}_3 : it's elements have form $\frac{p(x)}{q(x)}$ for $p(x), q(x) \in \mathbb{Z}_3[x]$. For the field K of four elements, which element is 0 and 1, the additive and multiplicative identity in K ?
 - Consider the elements $1, 1 + 1, 1 + 1 + 1, \dots \in F$. If F is finite these can not all be different, use the field axioms to prove that one of them has to be zero.
 - Suppose that $\text{char}(F) = m \cdot n$ and let $a = \underbrace{1 + \dots + 1}_m \in F$ and $b = \underbrace{1 + \dots + 1}_n \in F$. What can be said about the element $a \cdot b \in F$?
- H 9.15. To show that a field F is *not* algebraically closed, try and find a non-constant polynomial with coefficients in F that does not have a zero in F . Try polynomials of degree 2.

Chapter 12

Answers

A 1.1. $v+v = 1 \cdot v + 1 \cdot v = (1+1) \cdot v = 2 \cdot v$. The condition that the vector space is complex is not really necessary, the only problematic step is the last one: $1+1=2$, as 2 may not be an element of the field.

A 1.2. The zero element is 1, the additive inverse of 5 is $\frac{1}{5}$

A 1.3.

- a) S_1 is not additive and not homogeneous, it is not a subspace. Geometrically it is a line not passing the origin.
- b) S_2 is both additive and homogeneous, it is a subspace. Geometrically it is a line through the origin.
- c) S_3 is not additive but it is homogeneous, it is not a subspace. Geometrically S_3 is the union of the coordinate axes.
- d) S_4 is additive but not homogeneous, it is not a subspace. Geometrically S_4 is the first quadrant.
- e) S_5 is additive but not homogeneous, it is not a subspace. Geometrically S_5 is a lattice.
- f) S_6 is additive and homogeneous, it is a subspace. By definition every vector space is a subspace of itself, although it is called a *non-proper* subspace.
- g) S_7 is technically both additive and homogeneous since a statement of form $\forall x \in M : P(x)$ is true whenever M is empty. However, it is not a subspace - by definition subspaces are required to be nonempty. Note however that the point set $\{0\}$ is a subspace.

A 1.4. The projection is $u = (-1, 2, -1)$.

A 1.5. One choice of basis is $((x-2), x(x-2), x^2(x-2))$.

A 1.6. $S \cap S'$ and $S + S'$ are both subspaces. $S \cup S'$ is in general not a subspace, as illustrated by S_3 of the previous exercise.

A 1.7. For example, consider the three lines spanned by $(1, 0)$, $(0, 1)$, and $(1, 1)$.

A 1.8.

- a)
- b)
- c) $\cosh(x) = \frac{e^x + e^{-x}}{2}$
- d) 0 (the function is clearly even)

A 1.9. I, G, H, T, C are linear.

A 1.10. The kernel consists of all symmetric matrices, the image of all skew-symmetric matrices.

A 1.11. If $\ker(F)$ contains some nonzero v , then $F(0) = 0 = F(v)$, so F is not injective. On the other hand, if F is not injective with $F(u) = F(v)$ for two different vectors u and v , then by linearity $F(u-v) = F(u) - F(v) = 0$, so $u - v$ is a nonzero vector in the kernel.

A 1.12. F is injective but not surjective since the first coordinate is always zero in the image. F doesn't have an inverse (the function G that left-shifts sequences is a right inverse but not a left-inverse to G)

A 1.13. $[F] = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 3 & 4 \\ 0 & 0 & 1 & 3 & 6 \\ 0 & 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$. $F^{-1}(p(x)) = p(x-1)$.

A 1.14. $[F] = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -2 \end{pmatrix}$ and $[G] = \begin{pmatrix} 0 & -2 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}$

A 1.15.

A 1.16. $A = C$ since $(4, 5) - (2, 3) \in U$. Graphically, A, B , and $A+B$ are lines in \mathbb{R}^2 intersecting the y -axis in 1, -1, and 0 respectively. In the basis (A) , B has coordinates (-1) , since $(-1)A = -1((2, 3) + U) = (-2, -3) + U = (0, 1) + U = B$.

A 1.17. $(\frac{1}{2}, -\frac{1}{2})$.

A 1.18. Since $3e_1 - 1e_2 + 2(0, 0, 1, 1) = 0 + \ell$ the vectors are linearly dependent and is not a basis.

A 1.19. Consider the choice of bases from the hint. If the matrix for \tilde{F} is A , then the matrix for F has block form $\left(\begin{array}{c|c} 0 & B \\ \hline 0 & A \end{array} \right)$ where B is some matrix.

A 1.20.

- The line $(1, 2) + t(1, 1)$ (so $v = (1, 2)$, $U = \text{span}(1, 1)$.)
- The plane $x + y + z = 1$ (so $v = (1, 0, 0)$, $U : x + y + z = 0$.)

A 1.21. For (b): The map is given by $F(v) = G(v) + w$ with $w = (1, -1)$ fixed and $G(x, y) = (y, x)$ linear.

A 1.22.

- The zero vector is the empty set \emptyset since $S + \emptyset = S \Delta \emptyset = S$ for all subsets S . It follows from the axioms that $0 \cdot S = \emptyset$, and $1 \cdot S = S$ must hold, this defines scalar multiplication over \mathbb{Z}_2 completely.
- $\{1, 3, 5\} + \{1, 2, 3\} = \{2, 5\}$
- Compute $-\{1, 3, 5\} = \{1, 3, 5\}$
- The natural choice of basis are the singleton sets $(\{1\}, \{2\}, \{3\}, \{4\}, \{5\})$, any set is a (the natural linear combination of these).
- Yes, since $S_1 + S_2 = S_4$ we have $1 \cdot S_1 + 1 \cdot S_2 + 1 \cdot S_4 = \emptyset$

A 1.23.

A 1.24.

- $(1, 0, 2, 0)$
- $X = \begin{pmatrix} 2 \\ 1 \end{pmatrix}$

c) $2 \cdot 2 - 0 \cdot 1 = 1$

d) $\begin{pmatrix} 2 & 2 \\ 0 & 2 \end{pmatrix}$

e) $3^4 = 81$

A 1.25.

a) $(x, y) = (1, 4)$

b) No solution

c) $(x, y) = (4 + 3t, t)$ where $t \in \mathbb{Z}_5$. In other words, $(x, y) \in \{(4, 0), (2, 1), (0, 2), (3, 3), (1, 4)\}$.

A 1.26. The equation $\lambda_1 v_1 + \lambda_2 v_2 + \lambda_3 v_3 = 0$ becomes a 4×3 linear system over \mathbb{Z}_3 . The standard algorithm shows that the solution is $(\lambda_1, \lambda_2, \lambda_3) = t(1, 2, 1)$ where $t \in \mathbb{Z}_3$. This shows that (for example) $1 \cdot v_1 + 2 \cdot v_2 + 1 \cdot v_3 = 0$ so no, the vectors are not linearly independent.

A 1.27. $\ell_1 : P + tv = (1 + t, 2 + t)$, so taking $t = 0, 1, 2$ (all elements of \mathbb{Z}_3 we get that ℓ_1 consists of three points: $(1, 2), (2, 0), (0, 1)$. The second line has direction $w = \overline{QR} = R - Q = (1, 2)$ so $\ell_2 : Q + tw = (1 + t, 2t)$. Taking $t = 0, 1, 2$ we see that ℓ_2 consists of the three points $(1, 0), (2, 2), (0, 1)$. So ℓ_1 and ℓ_2 intersects in the point $(0, 1) \in (\mathbb{Z}_3)^2$.

A 2.1. $(A + B)^* = A^* + B^* = A + B$ so $A + B$ is always Hermitian when A and B are.

$(AB)^* = B^*A^* = BA$ so AB is only Hermitian when A and B commute. For a counterexample, take

$$A = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \text{ and } B = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

$(\lambda A)^* = \overline{\lambda}A^* = \overline{\lambda}A$ so λA is only Hermitian when $\lambda \in \mathbb{R}$. For a counterexample, take $\lambda = i$ and $A = I$.

$(A^T)^* = (\overline{A^T})^T = (A^*)^T = A^T$ so A^T is always Hermitian when A is.

$(AB^* + BA^*)^* = B^{**}A^* + A^{**}B^* = BA^* + AB^* = AB^* + BA^*$, so $AB^* + BA^*$ is Hermitian when A and B are.

A 2.2. $(A + Bi)^* = A^* + \overline{i}B^* = \overline{A}^T - i\overline{B}^T = A^T - iB^T = A + Bi \Leftrightarrow A = A^T \text{ and } B = -B^T$.

A 2.3. After using the formula and taking the sum of the diagonals in both AB and BA , the trace of both is equal to $\sum_{i,k} a_{ik}b_{ki}$.

A 2.4.

a)

b) $\lambda_1^n + \lambda_2^n + \lambda_3^n$

c) $-1, 1, 2$

A 2.5. $P^3 = I$, which shows that

$$P^n = \begin{cases} I & \text{for } n = 3k \\ P & \text{for } n = 3k + 1 \\ P^2 & \text{for } n = 3k + 2 \end{cases}$$

where $k \in \mathbb{Z}$. Note that since $P^{-1} = P^2$ the formula also holds for negative n .

A 2.6. For the last part: take for example the matrix

$$P = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix}$$

and consider how it acts on basis vectors: $e_1 \rightarrow e_2 \rightarrow e_3 \rightarrow e_1$ and $e_4 \leftrightarrow e_5$. The first cycle becomes the identity in P^n when n is divisible by 3, and the second cycle when n is divisible by 2. The order of P is therefore 6, the least common multiple of 3 and 2.

A 2.7. $\mathcal{C}_A = \text{span}(A, I)$.

A 2.8. $\mathcal{C} = \text{span}(I)$, only multiples of the identity-matrix commute with everything.

A 2.9. With $A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ we get $A^2 = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$ and $A^3 = I$, so thereafter we will cycle between these three matrices:

$$A^n = \begin{cases} I & \text{for } n = 3k \\ A & \text{for } n = 3k + 1 \\ A^2 & \text{for } n = 3k + 2 \end{cases}$$

A 2.10. With E as in the hint, all the matrices of form $S^{-1}ES$ are also idempotents.

A 2.11. $\dim \mathcal{C}_A \in \{5, 7, 9, 11, 13, 17, 25\}$

A 2.12. For the last statement, neither MN nor $M + N$ need be nilpotent if $AB \neq BA$, take for example $M = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ and $N = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$. Then M and N are nilpotent, but $M + N$ and MN are not.

A 2.13. $F(F(v)) = 0$, since by definition, the inner argument $F(v)$ lies in the image, and therefore in the kernel, and thus is mapped to zero. So $F^2 = 0$ and F is nilpotent with nilpotency degree 2.

A 2.14. (If $0 \neq \ker(N^k) = \ker(N^{k+1})$ for some n , then N would act bijectively on $\ker(N^k)$, so $\ker(N^{k'}) = \ker(N^k)$ for all $k' > k$, which contradicts nilpotency).

A 2.15.

A 2.16.

A 2.17. $M^{-1} = \left(\begin{array}{c|c} A^{-1} & -A^{-1}BC^{-1} \\ \hline 0 & C^{-1} \end{array} \right)$

A 2.18.

$$\text{REF} = \begin{pmatrix} 1 & -1 & 1 & 2 & 1 \\ 0 & 0 & 2 & -1 & -1 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad \text{RREF} = \begin{pmatrix} 1 & -1 & 0 & \frac{5}{2} & \frac{3}{2} \\ 0 & 0 & 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Other answers are possible for the REF, but not for the RREF.

$$\ker(A) = \{(s - 5t - 3r, s, t + r, 2t, 2r) \mid s, t, r \in \mathbb{R}\} = \text{span}((1, 1, 0, 0, 0), (-5, 0, 1, 2, 0), (-3, 0, 1, 0, 2))$$

A 2.19. The RREF of C is $\begin{pmatrix} 1 & 2i & 0 & 1 \\ 0 & 0 & 1 & i \end{pmatrix}$. The solutions to $CX = 0$ are

$$(x_1, x_2, x_3, x_4) = (-2is - t, s, -it, t) \text{ where } s, t \in \mathbb{C}.$$

A 2.20. The RREF of A is $\begin{pmatrix} 1 & 2 & -3 & 0 & 1 \\ 0 & 0 & 0 & 1 & \frac{1}{2} \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$

A 2.21.

$$\text{RREF: } \begin{pmatrix} 1 & 2 & 0 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad X = s(1, 1, 0, 0) + t(2, 0, 2, 1) \quad s, t \in \mathbb{Z}_3.$$

Note that since $s, t \in \{0, 1, 2\}$ there are exactly nine solutions to the linear system.

- A 2.22. $((0, -1, 1, 0, -1), (1, 1, 0, 0, 1), (2, -1, 4, 0, -2))$ is one choice of basis for $\text{Im}(A)$
 $((1, 0, 1, 0, 0), (1, 1, 0, 0, 1))$ is one choice of basis for $\text{ker}(A)$
 $((1, 0, 1, 0, 0), (1, 1, 0, 0, 1))$ is one choice of basis for $\text{Im}(A) \cap \text{ker}(A)$ (so in this case $\text{ker}(A) \subset \text{Im}(A)$)

A 2.23. The suggestion from the hint gives: $\text{Im}(A) = \text{span}(e_2, e_3)$, while $\text{ker}(A) = \text{span}(e_3, e_4)$, and their intersection is $\text{span}(e_3)$.

A 2.24.

$$E_1^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad E_2^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \frac{1}{5} \end{pmatrix} \quad E_3^{-1} = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

In general: $(I + \lambda e_{ij})^{-1} = I - \lambda e_{ij}$, $\text{diag}(d_1, d_2, \dots, d_n) = \text{diag}(\frac{1}{d_1}, \frac{1}{d_2}, \dots, \frac{1}{d_n})$, and $P^{-1} = P$ for matrices corresponding to switching two rows of the identity matrix ($P = I - e_{ii} - e_{jj} + e_{ij} + e_{ji} \Rightarrow P^{-1} = P$)

A 2.25.

$$(a) \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 2 & 0 & 1 \end{pmatrix} A \quad (b) A \begin{pmatrix} 1 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (c) \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} A$$

A 2.26.

$$A = \begin{pmatrix} 1 & 0 \\ 3 & 1 \end{pmatrix} \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & -2 \end{pmatrix} \quad A^{-1} = \begin{pmatrix} 1 & 0 \\ 0 & -\frac{1}{2} \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -3 & 1 \end{pmatrix}$$

(other answers are possible, multiply to verify)

A 2.27. If $EA = E_n \cdots E_2 E_1 A = I$ is a product of matrices that reduces A to the identity, this shows that $A = E^{-1}$ and $A = E$. When multiplying the block matrix on the left by E we get $E[A|I] = [EA|EI] = [I|E] = [E|A^{-1}]$.

A 2.28.

$$LU = \begin{pmatrix} 1 & 0 \\ -2 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 & 2 \\ 0 & 3 & 4 \end{pmatrix} \quad LDU = \begin{pmatrix} 1 & 0 \\ -2 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 3 \end{pmatrix} \begin{pmatrix} 1 & 1 & 2 \\ 0 & 1 & \frac{4}{3} \end{pmatrix}$$

A 2.29.

$$LU = \begin{pmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 3 & 2 & 1 \end{pmatrix} \begin{pmatrix} 1 & -1 & 1 & 1 \\ 0 & 3 & 2 & 1 \\ 0 & 0 & -6 & -3 \end{pmatrix} \quad LDU = \begin{pmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 3 & 2 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & -6 \end{pmatrix} \begin{pmatrix} 1 & -1 & 1 & 1 \\ 0 & 1 & \frac{2}{3} & \frac{1}{3} \\ 0 & 0 & 1 & \frac{1}{2} \end{pmatrix}$$

A 2.30. $PA = LU$ where $L = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}$, $P = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$, $U = \begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{pmatrix}$.

A 2.31. We have

$$\begin{pmatrix} 4 & 3 \\ 5 & 1 \end{pmatrix} = LU = \begin{pmatrix} 1 & 0 \\ 3 & 1 \end{pmatrix} \begin{pmatrix} 4 & 3 \\ 0 & 6 \end{pmatrix} = LDU' = \begin{pmatrix} 1 & 0 \\ 3 & 1 \end{pmatrix} \begin{pmatrix} 4 & 0 \\ 0 & 6 \end{pmatrix} \begin{pmatrix} 1 & 6 \\ 0 & 1 \end{pmatrix},$$

where all matrices have coefficients in \mathbb{Z}_7 .

A 2.32. Let E be a product of elementary matrices that reduces A to row echelon form U . Then $Av = 0 \Rightarrow EAv = 0 \Leftrightarrow Uv = 0$.

A 2.33. Such a matrix has block form $\begin{pmatrix} I & 0 \\ 0 & 0 \end{pmatrix}$, a diagonal matrix where the diagonal has a number of ones followed by a number of zeros.

A 2.34.

a) $(3 \cdot 4 + 2 \cdot 3 + 1 \cdot 2) + (3 + 2 + 1) = 26$. The first parenthesis corresponds to the operations to put A in echelon form, the rest corresponds to the back-substitution.

b) $2(3 + 2 + 1) = 12$

c) $2(1+2+\dots+(n-1)) = n(n-1)$ (for the standard Gaussian elimination when A is $n \times n$, the answer is a polynomial of degree 3).

A 2.35. With $C = \begin{pmatrix} 1 & 0 \\ 2 & \sqrt{2} \end{pmatrix}$ we have $A = CC^*$.

With $G = \begin{pmatrix} 1 & 0 \\ 2+i & 2 \end{pmatrix}$ we have $B = GG^*$.

A 2.36. With $C = \begin{pmatrix} 3 & 0 & 0 \\ 1 & 1 & 0 \\ -1 & 2 & \sqrt{5} \end{pmatrix}$ we have $A = CC^*$.

A 2.37. $(CC^*)^* = (C^{**})(C^*) = CC^*$ so CC^* is Hermitian. The 1×1 -matrix (-1) is Hermitian but does clearly not admit a factorization $(-1) = (\lambda)(\lambda)^* = (|\lambda|^2)$.

A 3.1.

- a) $\sigma(F) = \{0, 1\}$, $\dim E_1 = 1$, $\dim E_0 = 2$
- b) $\sigma(F) = \{1, -1\}$, $\dim E_1 = 2$, $\dim E_{-1} = 1$
- c) $\sigma(F) = \{1\}$, $\dim E_1 = 1$ (unless the rotation angle is a multiple of π)
- d) $\sigma(F) = \{1\}$, $E_1 = \mathbb{R}^3$ so $\dim E_1 = 3$.
- e) $\sigma(F) = \{0\}$, $\dim E_0 = 1$.

A 3.2. $D = \begin{pmatrix} 1+i & 0 \\ 0 & 1-i \end{pmatrix}$ and $S = \begin{pmatrix} 1 & 1 \\ i & -i \end{pmatrix}$ works. Other choices may also work.

A 3.3. $\sigma(F) = \mathbb{N}$ since $F(x^n) = nx^n$ for all integers $n \geq 0$. The eigenvectors for eigenvalue k are nonzero multiples of x^k . The operator F is sometimes called the *degree-operator*.

A 3.4. Since A is real we have $\bar{A} = A$. And we can compute

$$A\bar{v} = \bar{A}\bar{v} = \overline{Av} = \overline{\lambda v} = \bar{\lambda}\bar{v}$$

which shows that \bar{v} is an eigenvector of A with eigenvalue $\bar{\lambda}$.

A 3.5. The previous problem shows that $2 - 3i$ is an eigenvalue with eigenvector $(1, 1 - i)$, so

$$A = SDS^{-1} = \begin{pmatrix} 1 & 1 \\ 1+i & 1-i \end{pmatrix} \begin{pmatrix} 2+3i & 0 \\ 0 & 2-3i \end{pmatrix} \frac{1}{2} \begin{pmatrix} 1+i & -i \\ 1-i & i \end{pmatrix} = \begin{pmatrix} -1 & 3 \\ -6 & 5 \end{pmatrix}.$$

A 3.6. In the 2×2 case the eigenvalues are $1, -1, i, -i$ with corresponding eigenvectors

$$\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \quad \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} \quad \begin{pmatrix} 1 & i \\ -i & -1 \end{pmatrix} \quad \begin{pmatrix} 1 & -i \\ i & -1 \end{pmatrix},$$

so the operator is diagonalizable. The same is true in the 3×3 -case, with the geometric multiplicities of the eigenvalues in this case being $g_1 = 3, g_{-1} = 2, g_i = 2, g_{-i} = 2$.

A 3.7. Let $\xi = e^{\frac{2i\pi}{3}} = -\frac{1}{2} + \frac{\sqrt{3}}{2}i$. Then with $D = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \xi & 0 \\ 0 & 0 & \xi^2 \end{pmatrix}$ and $S = \begin{pmatrix} 1 & \xi^2 & 1 \\ 1 & \xi & \xi \\ 1 & 1 & \xi^2 \end{pmatrix}$ we have $SDS^{-1} = [P] =$

$$\begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}.$$

A 3.8. Every nonzero vector in the hyperplane $x + 2y + 3z + 4w = 0$ is an eigenvector with eigenvalue 0. Every vector parallel to $(1, 2, 3, 4)$ is an eigenvector of eigenvalue 30.

A 3.9. $|\text{Mat}_{2 \times 2}(\mathbb{Z}_2)| = 2^4 = 16$. The number of invertible matrices is $|\text{GL}_2(\mathbb{Z}_2)| = 6$.

A 3.10. $p_A(t) = \det(A - \lambda I) = \lambda^2 + 3$ has no zeroes in \mathbb{Z}_5 (as can be seen by plugging in each of the five elements in p_A), so A has no eigenvalues or eigenvectors.

On the other hand, over \mathbb{Z}_{11} we have $p_B(t) = \det(B - \lambda I) = \lambda^2 + 6\lambda + 9 = (\lambda + 3)^2$ which shows that $\lambda = -3 = 8$ is the only eigenvalue of B . For eigenvectors we solve the linear system $Bv = 8v$:

$$(B - 8I)v = 0 \Leftrightarrow (B + 3I)v = 0 \Leftrightarrow \begin{pmatrix} 4 & 2 \\ 3 & 7 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \Leftrightarrow \begin{pmatrix} 2 & 1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \Leftrightarrow v = \begin{pmatrix} x \\ y \end{pmatrix} = t \begin{pmatrix} 1 \\ 9 \end{pmatrix}.$$

So the only eigenvalue is 8 and the corresponding eigenvectors are $t \begin{pmatrix} 1 \\ 9 \end{pmatrix}$ for $t \in \mathbb{Z}_{11}$.

A 3.11.

a) $\begin{bmatrix} \spadesuit \\ \heartsuit \end{bmatrix}$

b)

$$\det(A - \lambda I) = \begin{vmatrix} \clubsuit - \lambda & \heartsuit \\ \spadesuit & \diamond - \lambda \end{vmatrix} = (\lambda - \clubsuit)(\lambda - \diamond) - \heartsuit \spadesuit = \lambda(\lambda + \diamond) + \spadesuit = \lambda^2 + \lambda \diamond + \spadesuit = (\lambda - \heartsuit)(\lambda - \spadesuit)$$

so there are two eigenvalues: \heartsuit and \spadesuit .

A 3.12. With notation as in the hint: the eigenspaces are $E_{1+i} = \text{span}(1 + i, 1)$ and $E_{1-i} = \text{span}(1 - i, 1)$, and $X_0 = (2 - i) \begin{pmatrix} 1 + i \\ 1 \end{pmatrix} + (2 + i) \begin{pmatrix} 1 - i \\ 1 \end{pmatrix}$, so

$$X_n = A^n X_0 = (2 - i)(1 + i)^n \begin{pmatrix} 1 + i \\ 1 \end{pmatrix} + (2 + i)(1 - i)^n \begin{pmatrix} 1 - i \\ 1 \end{pmatrix}.$$

Here the bottom coordinate is a_n , so

$$a_n = (2 - i)(1 + i)^n + (2 + i)(1 - i)^n \text{ for all } n \geq 0.$$

Note that despite how the expression looks, a_n is a real integer for each $n \in \mathbb{N}$.

A 3.13. $p(A) = A^4 + 2A^2 - 5A + 3I = \begin{pmatrix} 1 & 4 \\ 0 & 1 \end{pmatrix} + 2 \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} - 5 \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} + 3 \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 3 \\ 0 & 1 \end{pmatrix}.$

Since $B^2 = 0$ we get $p(B) = B^4 + 2B^2 - 5B + 3I = -5B + 3I = \begin{pmatrix} 3 & 0 & -5 \\ 0 & 3 & -5 \\ 0 & 0 & 3 \end{pmatrix}.$

A 3.14. We have $A^2 - 3A - 4I = 0$ (zero matrix), so $p(t) = t^2 - 3t - 4$ does the job. Note that $p(t)$ is in fact the characteristic polynomial of A .

A 3.15.

A 3.16. $p_A(A) = -(A + 2I)^2(A - 3I) = - \begin{pmatrix} 0 & 5 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 5 \end{pmatrix}^2 \begin{pmatrix} -5 & 5 & 0 \\ 0 & -5 & 0 \\ 0 & 0 & 0 \end{pmatrix} = 0$ (easy when computed as a product of block-matrices)

A 3.17. $p(A) = 0, q(A) = 0$

A 3.18. The eigenvalues of A are $1 + i$ and 2 , so the eigenvalues for B are $2^2 + 3 \cdot 2 - 5 = 5$ and $(1 + i)^2 + 3(1 + i) - 5 = 5i - 2$.

A 3.19. Since $p(t) = -t^3(t + 1)^2$, we have

$$m(t) \in \{t^3(t + 1)^2, t^2(t + 1)^2, t(t + 1)^2, t^3(t + 1), t^2(t + 1), t(t + 1)\}$$

A 4.13. For example, with: $S = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}$ and $J = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$ we have $SJS^{-1} = N$.

A 4.14. For example, with: $S = \begin{pmatrix} 1 & 0 & 1 \\ 2 & 1 & 1 \\ 0 & -1 & 0 \end{pmatrix}$ and $J = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}$ we have $SJS^{-1} = M$.

A 4.15. For example, $S = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ -1 & 2 & 1 \end{pmatrix}$ and $J = \begin{pmatrix} -1 & & \\ & 2 & \\ & & 2 \end{pmatrix}$.

A 4.16. For example, with

$$S = \begin{pmatrix} -1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ -2 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 \end{pmatrix} \quad \text{and} \quad J = \begin{pmatrix} 3 & 0 & 0 & 0 \\ 0 & 2 & 1 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix}$$

we have $SJS^{-1} = A$. Other S are possible, but J is unique up to switching the two blocks.

A 4.17. For example, with S and J as below we have $A = SJS^{-1}$:

$$S = \begin{pmatrix} 1 & 2 & 0 & 4 & 1 \\ 1 & 3 & 0 & 8 & 0 \\ 0 & 0 & 3 & -2 & 1 \\ 0 & 0 & -1 & 1 & 0 \\ 0 & 0 & -2 & 1 & 0 \end{pmatrix} \quad J = \begin{pmatrix} 2 & 1 & 0 & 0 & 0 \\ 0 & 2 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 0 & 2 \end{pmatrix}$$

Other options for S is possible, but not for J (except switching the two Jordan blocks).

A 4.18. For example, with $J = \begin{pmatrix} 3+i & 1 \\ 0 & 3+i \end{pmatrix}$ and $S = \begin{pmatrix} 3 & -1 \\ -2 & 1 \end{pmatrix}$ we have $A = SJS^{-1}$.

A 4.19. For example, $J = \begin{pmatrix} 3 & 1 \\ 0 & 3 \end{pmatrix}$ and $S = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$ gives $A = SJS^{-1} = \begin{pmatrix} 2 & 1 \\ -1 & 4 \end{pmatrix}$.

A 4.20. $[F] = J_4(1)$. $(F - \text{id})(p(x)) = p(x+1) - p(x)$ clearly decreases the degree of a polynomial by exactly 1, so x^3 is a first vector in a Jordan chain of length 4 for the eigenvalue 1, this determines the Jordan form.

A 4.21. The determinant is the product of diagonal in the Jordan-matrix, these are the eigenvalues of A , so $\det(A) = 0 \Leftrightarrow 0 \in \sigma(A)$.

A 4.22. With notation as in the hint, $A = S_1JS_1^{-1} = S_1(S_2^{-1}BS_2)S_1^{-1} = (S_1S_2^{-1})B(S_1S_2^{-1})^{-1}$, so $A = TBT^{-1}$ for $T = S_1S_2^{-1}$.

A 4.23. $J = J_3(1) \oplus J_1(i) \oplus J_1(-i)$. The functions $\cos(x) + i \sin(x) = e^{ix}$ and $\cos(x) - i \sin(x) = e^{-ix}$ are eigenvectors for the eigenvalues $\pm i$, and x^2e^x is the first vector of a Jordan chain of length 3 for the eigenvalue 1.

A 4.24. (A polynomial with simple roots would not be able to annihilate a Jordan block of size > 1 .)

A 4.25.

A 4.26.

$$A^2 \sim \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4 & 1 \\ 0 & 0 & 0 & 0 & 4 \end{pmatrix}, \quad \text{for } n \geq 3 \text{ we have } A^n \sim \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2^n & 1 \\ 0 & 0 & 0 & 0 & 2^n \end{pmatrix}$$

A 4.27. For example, with $J = \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ and $S = \begin{pmatrix} 1 & 2 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$ we have $A = SJS^{-1}$.

A 4.28. $x^2 + x + 1$ is the only irreducible polynomial of degree 2 with coefficients in \mathbb{Z}_2 (note that $x^2 + 1 = (x + 1)(x + 1)$). So take one of the matrices $A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ or $A = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$, these are the only matrices in $\text{Mat}_2(\mathbb{Z}_2)$ with the $p_A(t) = t^2 + t + 1$.

As a remark, if we extend our field of scalars from \mathbb{Z}_2 to the field $\mathbb{F} = \mathbb{Z}_2[x]/\mathbb{Z}_2[x](x^2 + x + 1) = \{0, 1, x, x+1\}$ (a field of four elements), the two matrices above actually do admit a Jordan-decomposition $A = SJS^{-1}$ where $S, J \in \text{Mat}_2(\mathbb{F})$. But constructions like this lie beyond the scope of this course.

A 4.29.

$$A^n = \begin{pmatrix} 2^n & n2^{n-1} & \frac{n(n-1)}{2}2^{n-2} & 0 & 0 \\ 0 & 2^n & n2^{n-1} & 0 & 0 \\ 0 & 0 & 2^n & 0 & 0 \\ 0 & 0 & 0 & (-6)^n & n(-6)^{n-1} \\ 0 & 0 & 0 & 0 & (-6)^n \end{pmatrix} = 2^{n-3} \begin{pmatrix} 8 & 4n & n(n-1) & 0 & 0 \\ 0 & 8 & 4n & 0 & 0 \\ 0 & 0 & 8 & 0 & 0 \\ 0 & 0 & 0 & 8(-3)^n & 4n(-3)^{n-1} \\ 0 & 0 & 0 & 0 & 8(-3)^n \end{pmatrix}$$

A 4.30.

$$e^A = \begin{pmatrix} e^{-3} & 0 & 0 \\ 0 & \sqrt{e} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad e^B = I + B + \frac{B^2}{2} = \begin{pmatrix} 1 & 1 & \frac{7}{2} \\ 0 & 1 & 3 \\ 0 & 0 & 1 \end{pmatrix}$$

A 4.31. $e^J = \begin{pmatrix} e^2 & e^2 & 0 \\ 0 & e^2 & 0 \\ 0 & 0 & e^3 \end{pmatrix} = e^2 \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & e \end{pmatrix}$.

A 4.32. $e^A v = e^\lambda v$ so v is still an eigenvector but with eigenvalue e^λ .

A 4.33.

$$\begin{aligned} \sin(A_1) &= \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} & \sin(A_2) &= \begin{pmatrix} \sin(2) & \cos(2) & -\frac{\sin(2)}{2} & -\frac{\cos(2)}{6} \\ 0 & \sin(2) & \cos(2) & -\frac{\sin(2)}{2} \\ 0 & 0 & \sin(2) & \cos(2) \\ 0 & 0 & 0 & \sin(2) \end{pmatrix} & \sin(A_3) &= \begin{pmatrix} \frac{\sqrt{3}}{2} & \frac{1}{2} & 0 & 0 \\ 0 & \frac{\sqrt{3}}{2} & 0 & 0 \\ 0 & 0 & \frac{\sqrt{2}}{2} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \\ \cos(A_1) &= \begin{pmatrix} 1 & 0 & -\frac{1}{2} & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} & \cos(A_2) &= \begin{pmatrix} \cos(2) & -\sin(2) & -\frac{\cos(2)}{2} & \frac{\sin(2)}{6} \\ 0 & \cos(2) & -\sin(2) & -\frac{\cos(2)}{2} \\ 0 & 0 & \cos(2) & -\sin(2) \\ 0 & 0 & 0 & \cos(2) \end{pmatrix} & \cos(A_3) &= \begin{pmatrix} \frac{1}{2} & -\frac{\sqrt{3}}{2} & 0 & 0 \\ 0 & \frac{1}{2} & 0 & 0 \\ 0 & 0 & \frac{\sqrt{2}}{2} & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \end{aligned}$$

A 4.34. The linear system of differential equations of order two $X''(t) + A^2X(t) = 0$ has the general solution $\sin(At)C + \cos(At)D$.

A 4.35. If A has Jordan form with eigenvalues λ_i on the diagonal, then $\det(e^A) = \prod e^{\lambda_i} = e^{\sum \lambda_i} = e^{\text{tr}A} = e^0 = 1$.

A 4.36. $\begin{cases} a_n = (-3)^{n-1}(2n-6) \\ b_n = (-3)^{n-1}(-2n) \end{cases}$ so $\frac{a_n}{b_n} \rightarrow -1$.

A 4.37. With $v_1 = (1, 1, 1)$, $v_2 = (0, 1, 2)$, and $v_3 = (0, 0, 1)$ we have $(A - 2I)v_1 = 0$, $(A - 2I)v_2 = v_1$, and $(A + I)v_3 = 0$, so (v_1, v_2, v_3) is a Jordan basis:

$$\text{With } S = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 2 & 1 \end{pmatrix} \text{ and } J = \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & -1 \end{pmatrix} \text{ we have } A = SJS^{-1}.$$

A 4.38. With $X_0 = (0, 1, 0)^T$, and with $A = D + J$ from in the previous problem, the solution is $e^{tA}X_0 = Se^{tD}e^{tN}S^{-1}$

$$= \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 2 & 1 \end{pmatrix} \begin{pmatrix} e^{2t} & 0 \\ 0 & e^{2t} & 0 \\ 0 & 0 & e^{-t} \end{pmatrix} \begin{pmatrix} 1 & t & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 1 & -2 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$$

$$= \begin{pmatrix} te^{2t} \\ (t+1)e^{2t} \\ (t+2)e^{2t} - 2e^{-t} \end{pmatrix}$$

A 5.1.

- a) $3 + 2i$
- b) $\sqrt{11}$
- c) $t(i - 1, 1)$ where $t \in \mathbb{C}$

A 5.2. $\|x^2 + x + 1\|^2 = \int_0^1 (x^2 + x + 1)^2 dx = \int_0^1 x^4 + 2x^3 + 3x^2 + 2x + 1 dx = \left[\frac{x^5}{5} + \frac{2x^4}{4} + \frac{3x^3}{3} + \frac{2x^2}{2} + x \right]_0^1 = \frac{1}{5} + \frac{1}{2} + 1 + 1 + 1 = \frac{37}{10}$,
so the sought length is $\sqrt{\frac{37}{10}}$.

A 5.3.

A 5.4. Only the first rule gives an inner product.

A 5.5. $\|v\|_2 = \sqrt{\langle v, v \rangle} = \sqrt{2} \sqrt{\langle v, v \rangle} = \sqrt{2} \|v\|_1$, so in the new norm vectors are a factor of $\sqrt{2}$ longer. Similarly, in the new norm $\cos(\theta) = \frac{\langle u, v \rangle}{\|u\|_2 \|v\|_2} = \frac{2\langle u, v \rangle}{\sqrt{2}\|u\|_1 \sqrt{2}\|v\|_1} = \frac{\langle u, v \rangle}{\|u\|_1 \|v\|_1}$ so angles between vectors are the same with the two inner products.

A 5.6. For the first question the answer is yes: Since (u, v) is a basis for \mathbb{C}^2 we can define an inner product by declaring that $\langle u, v \rangle = 0$, $\langle u, u \rangle = 0$, and $\langle v, v \rangle = 1$.

For the second question the answer is no: Since $v' = (1 - i)u'$ we would then have $0 = \langle v', u' \rangle = \langle (1 + i)u', u' \rangle = (1 + i)\langle u', u' \rangle \Leftrightarrow \langle u', u' \rangle = 0$, but $u' \neq 0$ so this contradicts the inner product axioms.

A 5.7.

A 5.8.

A 5.9. For example, take $u = (1, 0)$ and $v = (0, 1)$. Then the parallelogram law for the maximum norm becomes $2 = 1 + 1 = \|(1, 1)\|^2 + \|(1, -1)\|^2 = 2\|(1, 0)\|^2 + 2\|(0, 1)\|^2 = 2 + 2 = 4$ which does not hold.

A 5.10. No, \mathcal{I} is closed under addition but not scalar multiplication: $(-1) \cdot \langle v, v \rangle \leq 0$ so it is not positive definite.

A 5.11.

- a) $\|(2, 1, -4)\|_2 = \sqrt{2^2 + 1^2 + (-4)^2} = \sqrt{21}$
- b) $\|(2, 1, -4)\|_{\max} = \max\{|2|, |1|, |-4|\} = 4$
- c) $\|(2, 1, -4)\|_{\text{Mh}} = |2| + |1| + |-4| = 7$
- d) $\|(2, 1, -4)\|_3 = (|2|^3 + |1|^3 + |-4|^3)^{\frac{1}{3}} = (73)^{\frac{1}{3}}$

A 5.12. $\|A\|_F = \sqrt{1^2 + (-1)^2 + 1^2 + 1^2} = 2$ and $\|A\|_\sigma = |1 \pm i| = \sqrt{2}$

A 5.13.

A 5.14. $p = \frac{4}{3}$ (In general, $(\cos^{\frac{2}{p}}(t), \sin^{\frac{2}{p}}(t))$ is a parametrization of the unit circle with respect to the p -norm)

A 5.15. Let e_1, e_2 be two standard basis vectors in \mathbb{C}^n . For $0 < p < 1$ we have $\|e_1 + e_2\|_p = (1^p + 1^p)^{\frac{1}{p}} = 2^{\frac{1}{p}} > 2 = 1 + 1 = \|e_1\|_p + \|e_2\|_p$.

A 5.16. $1 \cdot \|v\|_{\max} \leq \|v\| \leq \sqrt{2} \cdot \|v\|_{\max}$. Remark: On a finite dimensional vector space, all norms are in fact equivalent.

A 5.17.

A 5.18. No, for example: $(1, i, 0) \times (1, 1, 1)$ is not orthogonal to $(1, i, 0)$. In fact it is not possible to define a vector product on \mathbb{C}^3 with the same properties as the vector product on \mathbb{R}^3 .

A 5.19. The vectors of length zero corresponds to the surface $z^2 = x^2 + y^2$, a double cone in \mathbb{R}^3 (the "light-cone" in two-dimensional space time in physics).

A 5.20.

A 5.21.

A 5.22.

A 5.23.

A 5.24. By the Pythagorean theorem we get $\|v - u\|^2 = \|v - P_U(v)\|^2 + \|P_U(v) - u\|^2 \geq \|v - P_U(v)\|^2$, so the minimal distance between u and v is $\|v - P_U(v)\|$, and it is attained when $v - P_U(v) = 0$.

A 5.25. Following the method in the hint we obtain

$$(f_1, f_2, f_3) = \left(\frac{1}{\sqrt{2}}(1, i, 0), \frac{1}{\sqrt{3}}(-1, i, 1), \frac{1}{\sqrt{6}}(1, -i, 2) \right)$$

which is an orthonormal basis for \mathbb{C}^3 where (f_1, f_2) is an ON-basis for U .

A 5.26. With the orthogonal basis $(e_1, e_2) = (1, x - \frac{1}{2})$ of \mathcal{P}_1 we get

$$g(x) = P_{\mathcal{P}_1}(e^x) = \frac{\langle e^x, 1 \rangle}{\langle 1, 1 \rangle} 1 + \frac{\langle e^x, x - \frac{1}{2} \rangle}{\langle x - \frac{1}{2}, x - \frac{1}{2} \rangle} (x - \frac{1}{2}) = 6(3 - e)x + 4e - 10.$$

A 5.27.

A 5.28. $g(x) = \frac{\pi}{2} - \frac{4}{\pi} \cos(x)$

A 5.29. With $Q = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ and $R = \sqrt{2} \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$ we have $A = QR$ satisfying all the conditions.

A 5.30.

$$Q = \begin{pmatrix} \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{3}} \\ 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{3}} \\ \frac{-2}{\sqrt{6}} & \frac{1}{\sqrt{6}} & \frac{1}{\sqrt{6}} \end{pmatrix} \quad R = \begin{pmatrix} \frac{3}{\sqrt{3}} & \frac{6}{\sqrt{3}} & \frac{6}{\sqrt{3}} \\ 0 & \frac{2}{\sqrt{2}} & \frac{-1}{\sqrt{2}} \\ 0 & 0 & \frac{3}{\sqrt{6}} \end{pmatrix}.$$

A 5.31. $A = QR = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \cdot \frac{1}{\sqrt{2}} \begin{pmatrix} 2 & 6 & 3 \\ 0 & -2 & 1 \end{pmatrix}$
 $B = Q'R' = \begin{pmatrix} \frac{1}{\sqrt{14}} & \frac{4}{\sqrt{21}} & \frac{1}{\sqrt{6}} \\ \frac{2}{\sqrt{14}} & \frac{1}{\sqrt{21}} & \frac{-2}{\sqrt{6}} \\ \frac{3}{\sqrt{14}} & \frac{-2}{\sqrt{21}} & \frac{1}{\sqrt{6}} \end{pmatrix} \cdot \begin{pmatrix} \sqrt{14} & \frac{10}{\sqrt{14}} \\ 0 & \frac{12}{\sqrt{21}} \\ 0 & 0 \end{pmatrix}.$

A 5.32.

A 5.33.

A 5.34.

A 5.35. a) Self adjoint: G b) Unitary: Only H c) G and H

A 5.36. For example: The operator on \mathbb{C}^2 with matrix $\begin{pmatrix} 0 & 2 \\ -2 & 0 \end{pmatrix}$ is normal but neither self-adjoint nor unitary.

A 5.37. For $a, b \in \mathbb{C}$ we get a) $a = \bar{b}$ b) $|a| = 4, b = -\bar{a}$ c) $|a| = |b|, a - b \in \mathbb{R}$ For $a, b \in \mathbb{R}$ this becomes a) $a = b$
 b) $(a, b) = \pm(4, -4)$ c) $a = \pm b$

A 5.38.

A 5.39.

A 5.40.

A 5.41. For (a): $w \in \ker(F^*) \Leftrightarrow F^*(w) = 0 \Leftrightarrow \langle v, F^*(w) \rangle = 0 \forall v \in V \Leftrightarrow \langle F(v), w \rangle = 0 \forall v \in V \Leftrightarrow w \perp \text{Im}(F) \Leftrightarrow w \in \text{Im}(F)^\perp$. For b, just replace F by F^* and take the complement of both sides in (a).

A 5.42.

A 5.43.

A 5.44.

a) For example, $-I$ is Hermitian but not positive definite.

b) With notation as in the hint we have $X^*AX = X^*BX + iX^*CX$. Since B and C are Hermitian, $X^*BX \in \mathbb{R}$ and $X^*CX \in \mathbb{R}$, so in order for X^*AX to be positive it needs to be real, and thus $0 = C = \frac{i(A^*-A)}{2}$, and $A = A^*$.

c) For example, with $A = \begin{pmatrix} 1 & 2 \\ 0 & 2 \end{pmatrix}$, when $X \in \text{Mat}_{2 \times 1}(\mathbb{R})$ is nonzero we get

$$X^*AX = X^TAX = (x_1 \ x_2) \begin{pmatrix} 1 & 2 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = x_1^2 + 2x_1x_2 + 2x_2^2 = (x_1 + x_2)^2 + x_2^2 > 0.$$

d) Our matrix $A = \begin{pmatrix} 1 & 2 \\ 0 & 2 \end{pmatrix}$ from the last problem is in fact not positive definite, because there exists nonzero complex X such that $X^*AX \not> 0$. For example, $X = \begin{pmatrix} 1 \\ i \end{pmatrix}$ gives $X^*AX = 3 + 2i \not> 0$.

A 5.45. $a = 1 - i, b > 1, c > \frac{b}{2(b-1)}$

A 5.46.

A 5.47.

A 5.48. $A = 3 \cdot \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} + 1 \cdot \frac{1}{2} \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$.

A 5.49.

A 5.50. $\sqrt{A} = \frac{1}{5} \begin{pmatrix} 9 & 2 \\ 2 & 6 \end{pmatrix}$

A 5.51. If A is invertible, then so is A^* , so $A^+ = (A^*A)^{-1}A^* = A^{-1}(A^*)^{-1}A^* = A^{-1}$.

A 5.52. We have $A^+ = \frac{1}{14} \begin{pmatrix} 1 & 5 & 4 \\ 4 & -8 & 2 \end{pmatrix}$ so with $b_1 = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}$ and $b_2 = \begin{pmatrix} 1 \\ 0 \\ i \end{pmatrix}$ we get a) $A^+b_1 = \frac{1}{14} \begin{pmatrix} 23 \\ -6 \end{pmatrix}$ b) $A^+b_2 = \frac{1}{14} \begin{pmatrix} 1+4i \\ 4+2i \end{pmatrix}$

A 5.53.

a) -

b) $A^+ = A^*(AA^*)^{-1}$

c) $A^+ = \frac{1}{4} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$.

A 6.1.

A 6.2. For example:

a) $\begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$ b) $\begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix}$

A 6.3. For example, with $A = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$ and $S = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ we have $S^{-1}AS = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} > 0$.

A 6.4.

A 6.5. $\lambda_1 = 5, v = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$.

A 6.6.

A 6.7.


A 6.8.

A 6.9. Following the hint we see that $0 = \det(A - \lambda I) = \det((A - \lambda I)^T) = \det(A^T - \lambda I)$ so $\sigma(A) = \sigma(A^T)$. Since the rows of A^T have sum 1, we get that $(1, 1, \dots, 1)^T > 0$ is a positive eigenvector with eigenvalue 1 of the positive matrix A^T . By Perron, 1 is the eigenvalue of maximal absolute value in $\sigma(A^T)$, so since $\sigma(A) = \sigma(A^T)$, 1 is also maximal in $\sigma(A)$ (but note that the corresponding eigenvector of A is not $(1, 1, \dots, 1)$).

A 6.10. Take for example $A = 0$ and $B = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. Then every vector is an eigenvalue of A with eigenvalue 0, and $\sigma(B) = \{-1, 1\}$.

A 6.11. $\lambda_1 = 45 = 1 + 2 + \dots + 9$. Since $(1, 1, \dots, 1) > 0$ is a positive eigenvector vector, the eigenvalue 45 is dominant, so there can exist no eigenvalue with larger absolute value. (Note that this holds for any Sudoku-matrix A , but the other eigenvalues might differ).

A 6.12.



The Gershgorin discs are drawn above, the eigenvalues $\sigma(A)$ are marked in the picture:

$$\sigma(A) = \{-4.25 + 0.02i, 3.50 + 1.85i, 3.71 + 0.14i\}$$

A 6.13. Proceeding as in the hint we get

$$\sum_{j=1}^n a_{kj}v_j = \lambda v_k \Leftrightarrow \sum_{j \neq k}^n a_{kj}v_j = (\lambda - a_{kk})v_k$$

so by the triangle inequality we get

$$|(\lambda - a_{kk})v_k| = \left| \sum_{j \neq k}^n a_{kj}v_j \right| \leq \sum_{j \neq k}^n |a_{kj}| |v_j| \leq \sum_{j \neq k}^n |a_{kj}| \cdot |v_k|$$

and since $|v_k| \neq 0$ we have $|\lambda - a_{kk}| \leq \sum_{j \neq k}^n |a_{kj}| = s_k$ which shows that λ lies inside the k 'th Gershgorin disc.

A 6.14. $\lambda_1 = \frac{3+\sqrt{5}}{2} = \phi^2$ (the square of the golden ratio).

$A^k v$ is $\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 2 \\ 3 \end{pmatrix}, \begin{pmatrix} 5 \\ 8 \end{pmatrix}, \begin{pmatrix} 13 \\ 21 \end{pmatrix}, \begin{pmatrix} 34 \\ 55 \end{pmatrix}$, so the requested quotient is $\frac{34}{13} \approx 2.615$, while $\lambda_1 \approx 2.618$. This illustrates a method of approximating the Perron-eigenvalue.

A 6.15. The matrix is not irreducible, you can never escape from the set of nodes numbered 2,3,5. With the relabelling $(1', 2', 3', 4', 5') := (2, 3, 5, 1, 4)$, the adjacency-matrix for the graph becomes block upper triangular

$$\begin{pmatrix} A & B \\ 0 & C \end{pmatrix} = \begin{pmatrix} 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix}$$

with A of size 3×3 and C of size 2×2 .

A 6.16. B and C are irreducible, only C is primitive.

A 6.17. Yes its adjacency-matrix A is both irreducible and primitive with $A^6 > 0$.

A 6.18. Since $A > 0$ is irreducible and $a_{kk} > 0$, the graph associated to A is strongly connected and contains a self-edge $v_k \rightarrow v_k$. Let p be the smallest number such that there exists a walk between any pair of nodes of length $\leq p$. Then for fixed indices i and j there must exist a walk $v_i \rightarrow \dots \rightarrow v_k \rightarrow \dots \rightarrow v_j$ of length $< 2p$, and it can be modified to a walk of length exactly $2p$ by walking along the self edge $v_k \rightarrow v_k$ a sufficient number of times, so $A^{2p} > 0$ and A is primitive.

A 6.19.

A 6.20. For example: $A = \begin{pmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{pmatrix}$. Then in the associated graph, if you start at v_1 and take an even number of steps you will always end up at v_1 or v_3 , and if you take an odd number of steps you always end up at v_2 or v_4 , so for no p there is a walk of length p both from v_1 to v_2 and from v_1 to v_3 .

A 6.21.

A 6.22. If $\lambda_1 < 1$ the spectral radius is < 1 and $A^k \rightarrow 0$.

A 6.23. A ranking vector (normalized so that its sum is 1, and rounded to three decimals) is $v = (0.166, 0.263, 0.570)$.

A 6.24. A Perron-vector will be $r = (1, 1, \dots, 1)$ so all websites will be ranked the same.

A 6.25. A Perron vector is $v = (0.139, 0.167, 0.139, 0.222, 0.139, 0.194)$ (rounded to 3 decimals), so the fourth node is ranked the highest.

A 6.26. A Perron vector for the maximal eigenvalue 1 is $v = \begin{pmatrix} 5 \\ 1 \end{pmatrix}$ so you will be happy 5/6 of the time.

A 6.27.

A 6.28. $v = \frac{1}{8}(1, 2, 2, 2, 1)^T$ is the stationary distribution (a Perron vector for the transition-matrix) so the drunkard will spend twice as much time at the three central bars as at the other two bars.

A 7.1. $(\sigma_1, \sigma_2, \sigma_3, \sigma_4) = (5, 3, \sqrt{2}, 0)$ (note the order). The right singular vectors are $(v_1, v_2, v_3, v_4) = (e_4, e_1, e_2, e_3)$, and corresponding left singular vectors are $(u_1, u_2, u_3, u_4) = (\frac{3+4i}{5}e_4, -e_1, \frac{1+i}{\sqrt{2}}e_2, \frac{3+4i}{5}e_4, e_3)$.

A 7.2. $A = U\Sigma V^*$ where $U = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ and $\Sigma = \begin{pmatrix} 2\sqrt{5} & 0 \\ 0 & \sqrt{5} \end{pmatrix}$ and $V = \frac{1}{\sqrt{5}} \begin{pmatrix} 2 & -1 \\ 1 & 2 \end{pmatrix}$.

A 7.3.

$$A = U\Sigma V^* = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

$$A = \tilde{U}\tilde{\Sigma}\tilde{V}^* = \begin{pmatrix} 1 \\ 0 \end{pmatrix} (1) \begin{pmatrix} 0 & 1 \end{pmatrix}$$

A 7.4. For $\alpha = \frac{3+\sqrt{5}}{2}$ we have $A = U\Sigma V^*$ where

$$V = \frac{1}{\sqrt{\alpha+1}} \begin{pmatrix} 1 & \alpha-1 \\ \alpha-1 & -1 \end{pmatrix} \quad \Sigma = \begin{pmatrix} \sqrt{\alpha} & 0 \\ 0 & \sqrt{\alpha^{-1}} \end{pmatrix} \quad U = \frac{1}{\sqrt{4\alpha-1}} \begin{pmatrix} \alpha & \alpha-1 \\ \alpha-1 & -\alpha \end{pmatrix}$$

A 7.5. $A = U\Sigma V = \tilde{U}\tilde{\Sigma}\tilde{V}^*$ where

$$U = \begin{pmatrix} 0 & -\frac{2}{\sqrt{6}} & -\frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{6}} & -\frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{6}} & \frac{1}{\sqrt{3}} \end{pmatrix} \quad \Sigma = \begin{pmatrix} 3 & 0 \\ 0 & \sqrt{3} \\ 0 & 0 \end{pmatrix} \quad V = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$$

$$\tilde{U} = \begin{pmatrix} 0 & -\frac{2}{\sqrt{6}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{6}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{6}} \end{pmatrix} \quad \tilde{\Sigma} = \begin{pmatrix} 3 & 0 \\ 0 & \sqrt{3} \end{pmatrix} \quad \tilde{V} = V = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}.$$

A 7.6. $A = U\Sigma V$ where

$$U = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \quad \Sigma = \begin{pmatrix} \sqrt{3} & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \quad V = \begin{pmatrix} \frac{1}{\sqrt{6}} & -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{3}} \\ \frac{2}{\sqrt{6}} & 0 & -\frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{6}} & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{3}} \end{pmatrix}.$$

A 7.7. $A = \tilde{U}\tilde{\Sigma}\tilde{V}^*$ where

$$\tilde{U} = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 0 & 1 \\ 1 & 0 \\ 0 & 0 \\ 1 & -1 \end{pmatrix} \quad \tilde{\Sigma} = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} \quad \tilde{V} = \begin{pmatrix} 2 & -2 \\ 1 & 2 \\ 2 & 1 \end{pmatrix}.$$

A 7.8. The right singular vectors are $(v_1, v_2, v_3) = (e_3, e_2, e_1)$, the left singular vectors are $(u_1, u_2, u_3) = (e_2, e_1, e_3)$, and the singular values are $(\sigma_1, \sigma_2, \sigma_3) = (3, 2, 0)$. Then it is easy to verify that $Av_i = \sigma_i u_i$ and $A^* u_i = \sigma_i v_i$.A 7.9. With respect to the ON-basis $(1, 2\sqrt{3}x - \sqrt{3})$ the matrix of D is $[D] = A = \begin{pmatrix} 0 & 2\sqrt{3} \\ 0 & 0 \end{pmatrix}$, so $A^*A = \begin{pmatrix} 0 & 0 \\ 0 & 12 \end{pmatrix}$.The eigenvalues of A^*A are 12 and 0, so the singular values for D are $\sigma_1 = \sqrt{12} = 2\sqrt{3}$ and $\sigma_2 = 0$.A 7.10. For $J = J_n(0)$ we have $J^*J = \text{diag}(0, 1, 1, \dots, 1)$ so the first $n-1$ singular values are 1, and the last one is zero. In general, a nilpotent Jordan matrix J with k Jordan blocks will have k singular values that are zero, and the rest will be one.

A 7.11.

A 7.12. If 0 is a singular value, some unit vector v_i is mapped to $0 \cdot u_i = 0$ so A is not injective. On the other hand, suppose that 0 is not a singular value. Let $A = U\Sigma V^*$. Since each matrix is square we get $A^{-1} = (V^*)^{-1}\Sigma^{-1}U^{-1} = V\Sigma^{-1}U^*$ and Σ is clearly invertible with inverse $\Sigma^{-1} = \text{diag}(\sigma_1^{-1}, \dots, \sigma_n^{-1})$.A 7.13. If $m < n$, A has a nontrivial kernel and can not be an isometry, and zero is a singular value as well. So assume $m \geq n$. On one hand, if all the singular values are 1, then $\Sigma = \begin{pmatrix} I \\ 0 \end{pmatrix}$ so a direct computation

shows that $\|AX\| = \|U\Sigma V^*X\| = \|X\|$. On the other hand, if $\sigma_i \neq 1$ is a singular value, then $Av_i = \sigma_i u_i$ so $\|Av_i\| = |\sigma_i| \|u_i\| = \sigma_i \cdot 1 \neq 1 = \|u_i\|$.

A 7.14. By the spectral theorem for normal operators, we get $A = UDU^*$ with U unitary and where we can order the diagonal elements of D (these are real and non-negative) in decreasing order. We then have $A^*A = (UDU^*)^*UDU^* = UD^*DU^*$, which gives $\Sigma^2 = D^*D$, and since both D^*D and Σ are non-negative, and diagonal we have $\Sigma = \sqrt{D^*D}$.

A 7.15. The singular values are $|\lambda|\sigma_1, \dots, |\lambda|\sigma_n$.

A 7.16. Following the hint, we look for a permutation matrix that reverses the order of the diagonal of Σ^{-1} . Let $P = e_{1,n} + e_{2,n-1} + \dots + e_{n_1}$ (it looks like the identity matrix mirrored horizontally), this is a permutation matrix, so it is unitary, and clearly $P^2 = I$, so we may write

$$A^{-1} = V\Sigma^{-1}U^* = (VP)(P\Sigma^{-1}P)(PU^*)$$

this is an actual SVD of A^{-1} since VP and PU^* are both unitary, and $P\Sigma^{-1}P = \text{diag}(\frac{1}{\sigma_n}, \dots, \frac{1}{\sigma_1})$ is diagonal with the singular values of A^{-1} in decreasing order.

A 7.17.

A 7.18. Let σ be the largest singular value of A and let σ' be the largest singular value for B . Then $\|ABv\| = \|A(Bv)\| \leq \sigma \|Bv\| = \sigma\sigma' \|v\| = \|A\|_{\text{op}} \cdot \|B\|_{\text{op}} \cdot \|v\|$ where both inequalities are equalities if and only if v is a right singular vector for B corresponding to σ' , and if Bv is a right singular vector for A corresponding to σ . In other words so $\|AB\|_{\text{op}} < \sigma\sigma' = \|A\|_{\text{op}} \cdot \|B\|_{\text{op}}$ with equality if and only if a left singular vector for B for the maximal singular value is parallel to a right singular vector for A corresponding to its maximal singular value.

A 7.19. Follow the hint. Equality holds everywhere $\|A\|_{\text{op}} = \|A\|_F = \|A\|_{\bullet}$ if and only if A has at most one nonzero singular value, in other words when $\text{rank}(A) \leq 1$.

A 7.20.

$$A_{(1)} = \frac{1}{21} \begin{pmatrix} 12 & 18 & 36 \\ 24 & 36 & 72 \\ 24 & 36 & 72 \end{pmatrix} \quad A_{(2)} = \frac{1}{21} \begin{pmatrix} -12 & 66 & 20 \\ 12 & 60 & 64 \\ 48 & -12 & 88 \end{pmatrix}$$

$$\|A - A_{(1)}\|_{\text{op}} = 4 \quad \|A - A_{(2)}\|_{\text{op}} = 1 \quad \|A - A_{(1)}\|_F = \sqrt{17} \quad \|A - A_{(2)}\|_F = 1.$$

A 7.21. $\frac{1}{10} = \frac{k(3000+2000+1)}{3000 \cdot 2000}$ gives $k \approx 120$.

A 7.22. For example, take $A = \begin{pmatrix} 7 & 0 \\ 0 & 5 \end{pmatrix}$ and $B = \begin{pmatrix} 4 & 0 \\ 0 & 2 \end{pmatrix}$ and $C = \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix}$. Then

$$\|A - B\|_{\text{op}} = 3 < 5 = \|A - C\|_{\text{op}}$$

so B is a better approximation of A than C with respect to the operator norm, but

$$\|A - B\|_F = \sqrt{18} > \sqrt{17} = \|A - C\|_F$$

so C is a better approximation of A than B with respect to the Frobenius norm.

A 7.23. With $A = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$, its SVD in sum form is $A = 2e_1e_1^* + 1e_2e_2^* + 1e_3e_3^*$ where e_1, e_2, e_3 is the standard

basis of \mathbb{C}^3 written as columns. We get a unique $A_{(1)} = 2e_1e_1^* = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$, but since the second and third singular values coincide we can reorder the second and third term to get another SVD, so we get

that $A_{(2)} = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}$ and $A'_{(2)} = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ are equally good rank 2 approximations of A (with respect to either norm).

A 7.24.

A 7.25.

A 7.26. $A = \begin{pmatrix} 3 \\ 4 \end{pmatrix} \begin{pmatrix} 1 & 2 \end{pmatrix} = 5\sqrt{5} \cdot \frac{1}{5} \begin{pmatrix} 3 \\ 4 \end{pmatrix} \cdot \frac{1}{\sqrt{5}} \begin{pmatrix} 1 & 2 \end{pmatrix} = \sigma_1 u_1 v_1^*$, so $A^+ = \frac{1}{\sigma_1} v_1 u_1^* = \frac{1}{5\sqrt{5}} \cdot \frac{1}{\sqrt{5}} \begin{pmatrix} 1 & 2 \end{pmatrix} \cdot \frac{1}{5} \begin{pmatrix} 3 & 4 \end{pmatrix} = \frac{1}{125} \begin{pmatrix} 3 & 4 \\ 6 & 8 \end{pmatrix}$.

A 7.27.

$$A^+ = \tilde{V} \tilde{\Sigma}^{-1} \tilde{U}^* = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} & 0 & \frac{1}{\sqrt{2}} \\ 0 & -1 & 0 \end{pmatrix} = \frac{1}{4} \begin{pmatrix} 1 & 2 & 1 \\ 1 & -2 & 1 \end{pmatrix}.$$

We note that $A^+A = I$, so A^+ is a left inverse of A , but

$$AA^+ = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 1 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \end{pmatrix} \neq I,$$

but in fact the matrix AA^+ is as close to the identity as possible for a rank 2-matrix, in this case we have $\|AA^+ - I\| = \sqrt{4 \cdot (\frac{1}{2})^2} = 1$.

A 7.28. Write $A = CR$ as a product of a columns and a row. Then with $u = \frac{1}{\|C\|}C$ and $v = \frac{1}{\|R\|}R^*$, and with $\sigma = \|C\| \cdot \|R\|$, the SVD of A is $A = \sigma uv^*$, so $A^+ = \frac{1}{\sigma} v u^* = \frac{1}{\|C\| \cdot \|R\|} \frac{1}{\|R\|} R^* \cdot \frac{1}{\|C\|} C^* = \frac{1}{\|C\|^2 \cdot \|R\|^2} R^* C^* = \frac{1}{\|A\|_F^2} A^*$.

A 7.29.

A 7.30.

A 7.31.

A 7.32. Follow the hint, all the the conditions follow directly from the Moore-Penrose conditions.

A 7.33. If A is of size $m \times n$, then $A^+A = I$ if and only if $\text{rank}(A) = n$, and $AA^+ = I$ if and only if $\text{rank}(A) = m$.

A 7.34.

A 7.35.

A 7.36. $\kappa(A) = \frac{\sigma_1}{\sigma_4} = \frac{5}{\sqrt{2}}$.

A 7.37. From the provided information we get $\sigma_1 \sigma_2 = 12$ and $\sigma_1^2 + \sigma_2^2 = 25$, so $(\sigma_1 + \sigma_2)^2 = 24 + 25 = 49$, so $\sigma_1 + \sigma_2 = 7$, and the only possibility is $\sigma_1 = 4$ and $\sigma_2 = 3$ which gives $\kappa(A) = \frac{\sigma_1}{\sigma_2} = \frac{4}{3}$.

A 7.38. The singular values are $|3 \pm \alpha|$, but which one is largest depends on α . For $\alpha \geq 0$ we get $\kappa(A) = \frac{\alpha+3}{|\alpha-3|}$ and for $\alpha \leq 0$ we get $\kappa(A) = \frac{3-\alpha}{|\alpha+3|}$.

A 7.39. We get the worst relative error when first of all $\|Ax\| = \sigma_1 \|x\|$ which means that $x \|v_1$, which means $b \|u_1$ since $Ax = b$. We also need $\|A^{-1}(\Delta b)\| = \frac{1}{\sigma_n} \Delta b$ which happens when $\Delta b \|u_n$.

A 7.40. For example: $A = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} = UP$

A 7.41. $[A] = UP = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 5 & 0 & 0 \\ 0 & 7 & 0 \\ 0 & 0 & 3 \end{pmatrix}$

A 7.42. $A = UP = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 2 & i \\ -i & 2 \end{pmatrix}$

A 7.43. For example, for $A = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$, for every real $0 \leq \theta < 2\pi$ we have a polar factorization


$$A = UP = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\theta} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

A 7.44. $A = U\Sigma V^* = (U\Sigma U^*)(UV^*) = PU'$ satisfies the conditions.

A 7.45. The compact SVD of the data-matrix A is $A = \tilde{U}\tilde{\Sigma}\tilde{V}^T$ where

$$\tilde{U} = \begin{pmatrix} -0.11 & -0.57 & 0.42 & 0.03 \\ -0.50 & 0.13 & 0.44 & -0.18 \\ -0.60 & 0.19 & -0.15 & 0.68 \\ -0.20 & -0.67 & 0.07 & 0.12 \\ -0.26 & 0.12 & -0.18 & 0.04 \\ -0.12 & -0.39 & -0.74 & -0.11 \\ -0.51 & 0.08 & -0.12 & -0.69 \end{pmatrix} \quad \tilde{\Sigma} = \begin{pmatrix} 18.79 & 0.00 & 0.00 & 0.00 \\ 0.00 & 9.68 & 0.00 & 0.00 \\ 0.00 & 0.00 & 1.18 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.94 \end{pmatrix} \quad \tilde{V} = \begin{pmatrix} -0.20 & -0.79 & 0.43 & 0.39 \\ -0.11 & -0.57 & -0.66 & -0.48 \\ -0.55 & 0.17 & -0.51 & 0.64 \\ -0.80 & 0.16 & 0.33 & -0.47 \end{pmatrix}$$

The projections of the rows and columns of A onto the first two right and left singular vectors respectively gives our two-dimensional data:



Here a reasonable interpretation is that the first principal component (corresponding to the x-axis in both diagrams) has to do with overall word frequency: we see that culinary is the least common word and politics is the most common, in the other diagram the x-coordinates gives some indication of the overall lengths of the magazines. More interesting for the topic-classification is probably the second principal component, in this example, this probably correspond to something like the type of magazine on a scale from "lifestyle" to "newspaper".

A 7.46. Total least squares line: $(2, 1) + t(1, 1)$ or $y = x - 1$.
Regular least squares line: $y = \frac{x}{2}$.

A 7.47. The line that best approximates the data in the total least squares sense is $\ell : (1, 2, 3) + t(0, 0, 1)$. The plane that best approximates the data in the total least squares sense is $\pi : (1, 2, 3) + s(1, 1, 0) + t(0, 0, 1)$ which can also be written $\pi : x - y = -1$.

A 7.48. For example, with the three points $x_1 = 0, x_2 = 0, x_3 = 1$, the affine subspace of dimension 0 is just a point s . Then $D = s^2 + s^2 + (1-s)^2 = 3s^2 - 2s + 1$ is minimized for $s = \frac{1}{3}$ while $D' = |s-1| + 2|s|$ is minimized for $s = 0$. The conclusion should be when minimizing squared distances, larger errors are somewhat exaggerated, so the total least squares method will be a bit biased towards minimizing the larger errors.

A 8.1. $v = 20e_1 \otimes e_1 + 10e_1 \otimes e_2 - 20e_2 \otimes e_1 - 10e_2 \otimes e_2$. To express v as a pure tensor we can rework the original expression and write

$$v = 2 \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} - 3 \begin{pmatrix} -2 \\ 2 \end{pmatrix} \otimes \begin{pmatrix} 3 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 2 \\ 4 \end{pmatrix} + \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 18 \\ 6 \end{pmatrix} = \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 20 \\ 10 \end{pmatrix}.$$

A 8.2.

$$w = 6 \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 2 \\ 1 \end{pmatrix} + 6 \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 2 \\ 1 \end{pmatrix} - 12 \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 2 \\ 1 \end{pmatrix}$$

$$= \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \left(\begin{pmatrix} 6 \\ 6 \end{pmatrix} + \begin{pmatrix} 6 \\ -6 \end{pmatrix} + \begin{pmatrix} -12 \\ 0 \end{pmatrix} \right) \otimes \begin{pmatrix} 2 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ -1 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 2 \\ 1 \end{pmatrix} = 0$$

so w is in fact the zero vector in $\mathbb{R}^2 \otimes \mathbb{R}^2 \otimes \mathbb{R}^2$.

A 8.3. For example

$$\begin{pmatrix} 1 \\ 2 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 2 \end{pmatrix} = 2 \begin{pmatrix} 1 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 1 \end{pmatrix} - \begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} + 2 \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

A 8.4. $v = \begin{pmatrix} 2 \\ i \end{pmatrix} \otimes \begin{pmatrix} i \\ -2 \\ 3 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ -i \end{pmatrix}.$

A 8.5. $v = \begin{pmatrix} 2 \\ 3 \end{pmatrix} \otimes \begin{pmatrix} 2 \\ 3 \\ -2 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ -\frac{1}{2} \end{pmatrix}$

-2	-3	2
-3	$-\frac{9}{2}$	3
4	6	-4
6	9	-6

A 8.6. Let $v = (1, 1, 1, 1, 1)$ and $w = (1, -1, 1, -1, 1)$. Then the given tensor is

$$t = \frac{1}{2} v \otimes v \otimes v + \frac{1}{2} w \otimes w \otimes w.$$

The tensor $w \otimes w \otimes w$ is s from the hint. No expression with fewer terms is possible, keeping the two last indices fixed, a pure tensor would have all "tubes" parallel in \mathbb{R}^5 (a tube is obtained by keeping two indices of $[t]_{ijk}$ fixed).

A 8.7. $\varphi(\sum_{ij} a_{ij} x^i \otimes y^j) := \sum_{ij} a_{ij} x^i y^j$ is an explicit isomorphism $\varphi : \mathbb{R}[x, y] \rightarrow \mathbb{R}[x] \otimes \mathbb{R}[y]$. All three spaces are in fact isomorphic since they all have dimension \aleph_0 , countably infinity. However, the the isomorphism above has a "canonical" structure, writing down the other isomorphism explicitly will be very hard.

A 8.8. $\dim(V) = 6$ and $\dim(W) = 5$.

A 8.9. Both sides have dimension $2 \dim V$. An isomorphism is given by

$$\varphi : V \oplus V \rightarrow \mathbb{R}^2 \otimes V \quad \text{where} \quad \varphi((v, v')) = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes v + \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes v'.$$

A 8.10. The dimension of both sides is $\dim U \cdot \dim V \cdot \dim W$.

A 8.11. Define a map $\Psi : \text{Hom}(U \otimes V, W) \rightarrow \text{Hom}(U, \text{Hom}(V, W))$ by $\Psi(\varphi)(u) = \varphi(u \otimes -) : V \rightarrow W$. Explicitly this can be written

$$(\Psi(\varphi)(u))(v) = \varphi(u \otimes v).$$

This map Ψ is linear and bijective. This same isomorphism is important in many branches of mathematics, it works the same in many different category and is called the "tensor-hom adjunction".

A 8.12.

- a) Since $(7, 4) = \frac{3}{2}(3, 1) + \frac{5}{2}e_2$ we get $e_1^*((7, 4)) = \frac{3}{2}$ and $e_2^*((7, 4)) = \frac{5}{2}$.
- b) $(2e_1^* - 3e_2^*)((5, 3)) = (2e_1^* - 3e_2^*)(e_1 + 2e_2) = 2e_1^*(e_1) + 2e_1^*(2e_2) - 3e_2^*(e_1) - 3e_2^*(2e_2) = 2 + 0 + 0 - 6 = -4$.
- c) $g = 4e_1^* + 2e_2^*$.

A 8.13. $f = e_0^* + 3e_2^* + 2e_3^*$.

A 8.14.

A 8.15. Since $\dim(V) = 4 \cdot (1 + 2) = 12$, the vector space has $2^{12} = 4096$ elements.

A 8.16. $\dim(V^{\otimes p} \otimes (V^*)^{\otimes q}) = (\dim V)^{p+q}$.

A 8.17.

$$C_{1,1}(w) = 3 \begin{pmatrix} 3 \\ -1 \end{pmatrix} + 2 \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 11 \\ -1 \end{pmatrix}, \quad C_{2,1}(w) = 5 \begin{pmatrix} 1 \\ 1 \end{pmatrix} + 0 \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 5 \\ 5 \end{pmatrix}.$$

A 8.18. Under the natural identification between $\text{Hom}(\mathbb{R}^2, \mathbb{R}^2)$ and $\mathbb{R}^2 \otimes \mathbb{R}^2$ we have

$$F = e_1 \otimes e_1^* + 2e_1 \otimes e_2^* + 3e_2 \otimes e_1^* + 4e_2 \otimes e_2^*$$

so $C_{1,1}(F) = e_1^*(e_1) + 2e_2^*(e_1) + 3e_1^*(e_2) + 4e_2^*(e_2) = 1 + 0 + 0 + 4 = 5$.

A 8.19. If F has matrix A with respect to a basis (e_1, \dots, e_n) of V , then F corresponds to the tensor $w = \sum_{i=1}^n \sum_{j=1}^n a_{ij} e_i \otimes e_j^*$, and $C_{1,1}(w) = C_{1,1} \left(\sum_{i=1}^n \sum_{j=1}^n a_{ij} e_i \otimes e_j^* \right) = \sum_{i=1}^n \sum_{j=1}^n a_{ij} C_{1,1}(e_i \otimes e_j^*) = \sum_{i=1}^n \sum_{j=1}^n a_{ij} \delta_{ij} = \sum_{i=1}^n a_{ii} = \text{tr}(A) = \text{tr}(F)$.

A 8.20. For more compact notation, let $e_{ijk} := e_i \otimes e_j^* \otimes e_k^*$. Then the vector product tensor is

$$\times = t = e_{123} - e_{132} + e_{231} - e_{213} + e_{312} - e_{321}.$$

A 8.21. $(F \otimes G)^{-1} = F^{-1} \otimes G^{-1}$.

A 8.22. a) $\begin{pmatrix} 1 & 2 & 0 & 0 \\ 3 & 4 & 0 & 0 \\ 0 & 0 & 1 & 2 \\ 0 & 0 & 3 & 4 \end{pmatrix}$ b) $\begin{pmatrix} 3 & 0 & 1 & 0 \\ 0 & 3 & 0 & 1 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix}$ c) $\begin{pmatrix} 3 & 1 & 6 & 2 \\ 0 & 2 & 0 & 4 \\ 9 & 3 & 12 & 4 \\ 0 & 6 & 0 & 8 \end{pmatrix}$ d) $\begin{pmatrix} 1 & 2 & 2 & 4 \\ 3 & 4 & 6 & 8 \\ 3 & 6 & 4 & 8 \\ 9 & 12 & 12 & 16 \end{pmatrix}$

A 8.23. $[S] = -\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ and $[R] = \frac{1}{2} \begin{pmatrix} \sqrt{3} & -1 \\ 1 & \sqrt{3} \end{pmatrix}$, so the matrix for $S \otimes T$ with respect to the specified basis is

$$[S] \otimes [R] = -\frac{1}{2} \begin{pmatrix} 0 & 0 & \sqrt{3} & -1 \\ 0 & 0 & 1 & \sqrt{3} \\ \sqrt{3} & -1 & 0 & 0 \\ 1 & \sqrt{3} & 0 & 0 \end{pmatrix}.$$

A 8.24. If e_i are the standard basis of \mathbb{R}^m expressed as columns as usual, we get

$$\text{vec}(A) = \sum_{i=1}^m e_i^T \otimes r_i.$$

A 8.25. $F(v) = \lambda v$ and $G(w) = \mu w$ implies $(F \otimes G)(v \otimes w) = (\lambda \mu) v \otimes w$, so pairs of eigenvectors of F and G give eigenvectors of $F \otimes G$ where the new eigenvalue is the product of the two respective eigenvalues. So we get

$$\sigma(F \otimes G) = \{4, 6, 8, 12, 24\},$$

each eigenspace of $F \otimes G$ is 1-dimensional except for E_{12} which is 2-dimensional (because $2 \cdot 6 = 3 \cdot 4$ are two ways of writing 12 as a product of an eigenvalue of F and an eigenvalue of G). Since $F \otimes G$ is a map on the 6-dimensional space $\mathbb{C}^2 \otimes \mathbb{C}^3$ and has six linearly independent eigenvectors, $F \otimes G$ is diagonalizable.

A 8.26. If $(\lambda_1, \dots, \lambda_m)$ are the eigenvalues of F and if (μ_1, \dots, μ_n) are the eigenvalues of G , then the eigenvalues of $F \otimes G$ are $(\lambda_1, \dots, \lambda_m) \otimes (\mu_1, \dots, \mu_n) = (\lambda_1 \mu_1, \lambda_1 \mu_2, \dots, \lambda_m \mu_n)$. So since the trace is the sum of the eigenvalues we get

$$\text{tr}(F)\text{tr}(G) = (\lambda_1 + \dots + \lambda_m)(\mu_1 + \dots + \mu_n) = (\lambda_1 \mu_1 + \lambda_1 \mu_2 + \dots + \lambda_m \mu_n) = \text{tr}(F \otimes G).$$

If some eigenvalues λ_i have multiplicity > 1 we should technically consider the multiset $\{\lambda_1, \dots, \lambda_m\}$ instead, and similar for μ , but the same proof works. The statement is actually true even when F and G are not diagonalizable.

A 8.27.

A 8.28.

A 8.29.

A 8.30. If $Av = \lambda v$ and $Bw = \mu w$, then $(A \otimes I_n + I \otimes B)(v \otimes w) = Av \otimes w + v \otimes Bw = (\lambda v) \otimes w + v \otimes (\mu w) = (\lambda + \mu)(v \otimes w)$, so $\lambda + \mu$ is an eigenvalue of the given operator.

A 8.31. For example, let $A = \begin{pmatrix} 0 & 1 \\ 2 & 0 \end{pmatrix}$ and $B = \begin{pmatrix} 0 & 1 \\ 3 & 0 \end{pmatrix}$. Since $\sqrt{2}$ is an eigenvalue of A and $\sqrt{3}$ is an eigenvalue of B , by the previous exercise, $\sqrt{2} + \sqrt{3}$ will be an eigenvalue of the Kronecker product $A \otimes I + I \otimes B$. We have

$$A \otimes I + I \otimes B = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 3 & 0 & 0 & 1 \\ 2 & 0 & 0 & 1 \\ 0 & 2 & 3 & 0 \end{pmatrix}.$$

We compute the characteristic polynomial $p_{A \otimes B}(t) = t^4 - 10t^2 + 1$, it is easy to verify that $\sqrt{2} + \sqrt{3}$ indeed is a root.

A 8.32. Since α, β are algebraic there exists integer coefficient polynomials p_α and p_β for which $p_\alpha(\alpha) = 0$ and $p_\beta(\beta) = 0$. Take a matrix A for which $\det(A - \lambda I) = p_\alpha$ is the characteristic polynomial and similarly pick B such that $\det(B - \lambda I) = p_\beta$. Such matrices exist (take the companion-matrices). Then by previous exercises $p_{A \otimes B}(\alpha\beta) = 0$ and $p_{A \otimes I + I \otimes B}(\alpha + \beta) = 0$, so both $\alpha\beta$ and $\alpha + \beta$ are algebraic.

A 8.33.

A 8.34. Let U and V be unitary. Then by the properties of the Kronecker product we have $(U \otimes V)(U \otimes V)^* = (U \otimes V)(U^* \otimes V^*) = (UU^*) \otimes (VV^*) = I \otimes I = I$, so $U \otimes V$ is unitary. Note that this still works when U and V are of different size. This gives a way to merge ON-bases to form ON-bases in higher-dimensional space.

A 8.35. If $A = U_1 \Sigma_1 V_1^*$ and $B = U_2 \Sigma_2 V_2^*$, then

$$A \otimes B = (U_1 \otimes U_2)(\Sigma_1 \otimes \Sigma_2)(V_2 \otimes V_1)^*$$

is an SVD of $A \otimes B$ (except that of the singular values in $\Sigma_1 \otimes \Sigma_2$ may be ordered incorrectly). Either way, the operator norm of $A \otimes B$ is its largest singular value, which is the element at position $(1, 1)$ in $\Sigma_1 \otimes \Sigma_2$, this is the product of the largest singular values of A and of B respectively.

A 8.36. If $SAS^{-1} = B$ and $TCT^{-1} = D$, then

$$(S \otimes T)(A \otimes C)(S \otimes T)^{-1} = (S \otimes T)(A \otimes C)(S^{-1} \otimes T^{-1}) = (SAS^{-1}) \otimes (TCT^{-1}) = B \otimes D,$$

so $A \otimes C \sim B \otimes D$.

A 8.37.

A 8.38. To show that not every subspace of $U \otimes V$ has form $U' \otimes V'$, take $U = V = \mathbb{R}^2$, then any subspace has dimension 0, 1, or 2, so $\dim(U' \otimes V') = \dim(U') \cdot \dim(V')$ can never be 3, but $\mathbb{R}^2 \otimes \mathbb{R}^2 \simeq \mathbb{R}^4$ obviously has 3-dimensional subspaces.

A 8.39.

A 8.40. If the b_i are zero, and $\sigma(x) = x$, then the neural network is just a composition of linear functions, which is again linear. So all the middle layers "collapse", and the network $F : \mathbb{R}^N \rightarrow \mathbb{R}^M$ is just a linear map, and it can never be trained to solve non-linear problems. Similarly if the b_i are not zero, F is a composition of affine maps which is still affine.

A 9.1.

- a) For $F = \text{SL}_n(R)$ The addition operation is not closed: it is not a function $F \times F \rightarrow F$. In other words, sums of matrices in F need not lie in F . (F3),(F4),(F5) also do not hold.
- b) (F4) does not hold, no element except 0 has an additive inverse.
- c) (F8) does not hold, nonzero elements of form $(x, 0)$ and $(0, x)$ does not have a multiplicative inverse.
- d) (F8) does not hold, the elements that are divisible by 3 or 5 lacks multiplicative inverse: 0, 3, 5, 6, 9, 10, 12.
- e) (F5) does not hold since $u \times v = -v \times u$. (F6) does not hold, for example $0 = (e_1 \times e_1) \times e_2 \neq e_1 \times (e_1 \times e_2) = e_1 \times e_3 = -2e_2$. (F7) does not hold, because $u \times v$ is orthogonal to v so it can never be v when v is nonzero. (F8) does not hold, it doesn't make sense when 1 is undefined.

A 9.2. No, \mathbb{H} is not a field since multiplication is not commutative.

A 9.3. We have $(a+b)^p = \sum_{k=0}^p \binom{p}{k} a^k b^{p-k} = a^p + \sum_{k=1}^{p-1} \binom{p}{k} a^k b^{p-k} + b^p$ so it is enough to show that $\binom{p}{k}$ is divisible by p for $1 \leq k \leq p-1$, but this is clearly true since $\binom{p}{k} = \frac{p \cdot (p-1) \cdots (p-k+1)}{k!}$ is an integer, and since $k > 1$ the prime p is a factor of the numerator, and it cannot be a factor of the denominator since $k < p$.

A 9.4.

- a) By (F4) there exists an additive inverse $(-a)$ of a . We add it to both sides and use (F2), (F3), (F4):
- $$a + b = a + c \Rightarrow (-a) + (a + b) = (-a) + (a + c) \Leftrightarrow ((-a) + a) + b = ((-a) + a) + c \Leftrightarrow 0 + b = 0 + c \Leftrightarrow b = c.$$
- b) $0 \cdot a = (0 + 0) \cdot a = 0 \cdot a + 0 \cdot a \Leftrightarrow 0 \cdot a + 0 = 0 \cdot a + 0 \cdot a$. Now by the first part we get $0 = 0 \cdot a$.
- c) $0 = 0 \cdot a = (1 + (-1)) \cdot a = 1 \cdot a + (-1) \cdot a = a + (-1) \cdot a$, this shows that $(-1)a$ is the additive inverse to a as stated.
- d) If $a = 0$ we are done. Otherwise a^{-1} exists, and we can multiply the equality with a^{-1} : $a \cdot b = 0 \Rightarrow a^{-1} \cdot (a \cdot b) = (a^{-1} \cdot a) \cdot b = 1 \cdot b = b = 0$. So either $a = 0$ or $b = 0$ (or both).
- e) Suppose that 0 and $0'$ both satisfy the property (F3). Then $0 = 0 + 0' = 0'$ (the left equality by the (F3)-property for $0'$, the right equality by the (F3)-property for 0).
- f) As above, if 1 and $1'$ both satisfy (F7), then $1 = 1 \cdot 1' = 1'$.
- g) We combine (F5) and (F9): $(a+b)c = c(a+b) = ca + cb = ac + bc$.

A 9.5. For any $a \in F$ we have $a = 1 \cdot a = 0 \cdot a = 0$, so 0 is the only element (that $0 \cdot a = 0$ was proved in a previous problem).

A 9.6.

- a) $3 + 5 \cdot 4 - 6 = 3 + 6 + 1 = 3$
- b) $\frac{3}{4} = 3 \cdot 4^{-1} = 3 \cdot 2 = 6$
- c) $6^{11} = (-1)^{11} = -1 = 6$
- d) $x = 3$ or $x = 4$
- e) No solutions

A 9.7. Since $1 = 2 \cdot 53 - 15 \cdot 7$, in \mathbb{Z}_{53} we get $1 = (-15) \cdot 7 = 38 \cdot 7$.

A 9.8. Fermat's little theorem implies that $a^p = a$ so if $a \neq 0$ we have $1 = a^{p-1} = a \cdot a^{p-2}$ and we get $a^{-1} = a^{p-2}$ as an explicit formula for the inverse of a . However, computationally this is not an efficient method since if p and a are large, a^{p-2} is a very big number.

A 9.9. Take for example $p(x) = x^p - x + 1$. By Fermat's little theorem, $p(x) = 1$ for all $x \in \mathbb{Z}_p$, so p is a nonconstant polynomial with no zeros in \mathbb{Z}_p .

- A 9.10. $t^2 = 0$ has solution $t = 0$,
 $t^2 + 1 = 0$ has solution $t = 1$,
 $t^2 + t = 0$ has solutions $t = 0$ and $t = 1$,
 $t^2 + t + 1 = 0$ has no solutions.

A 9.11. The points on the elliptic curve over \mathbb{Z}_5 are

$$(0, 1), (0, 4), (1, 1), (1, 4), (3, 0), (4, 1), (4, 4).$$

A 9.12. We have $\begin{pmatrix} a & b \\ -b & a \end{pmatrix} \begin{pmatrix} c & d \\ -d & c \end{pmatrix} = \begin{pmatrix} x & y \\ -y & x \end{pmatrix}$ where $x = ac - bd$ and $ad + bc$, so the product of matrices in F still lies in F , the same is obviously true for sums. So addition and multiplications are really operations $F \times F \rightarrow F$. All the field axioms except (F8) follow from the fact that $\text{Mat}_2(\mathbb{R})$ is a ring. (F8) is true because the inverse of a nonzero matrix in F can be given explicitly as

$$\begin{pmatrix} a & b \\ -b & a \end{pmatrix}^{-1} = \frac{1}{a^2 + b^2} \begin{pmatrix} a & -b \\ b & a \end{pmatrix} \in F.$$

We remark that F is just the field \mathbb{C} in disguise, the fields correspond to each other via $\mathbb{C} \ni a + bi \leftrightarrow \begin{pmatrix} a & b \\ -b & a \end{pmatrix}$. In other words, the map $\varphi : \mathbb{C} \rightarrow F$ with $\varphi(a + bi) = \begin{pmatrix} a & b \\ -b & a \end{pmatrix}$ is a *field automorphism*.

A 9.13.

- a) ♥
- b) ♣ = 0
- c) ♥ = 1
- d) Each element is its own additive inverse: $-\clubsuit = \clubsuit$, $-\heartsuit = \heartsuit$, $-\spadesuit = \spadesuit$, and $-\diamond = \diamond$, in other words, $-x = x$ in F , this is useful in the calculations below.
- e) ♥⁻¹ = ♥, ♠⁻¹ = ♦, and ♦⁻¹ = ♠
- f) ♠(x + ♥) = ♦ ⇔ x + ♥ = ♠⁻¹ · ♦ ⇔ x = ♠⁻¹ · ♦ - ♥ = ♦ · ♦ + ♥ = ♠ + ♥ = ♦
- g) $x^2 + \heartsuit = \spadesuit \Leftrightarrow x^2 = \spadesuit + \heartsuit \Leftrightarrow x^2 = \diamond \Leftrightarrow x = \spadesuit$

A 9.14.

- a) $\text{char}(\mathbb{R}) = \text{char}(\mathbb{C}) = \text{char}(\mathbb{Q}) = 0$, $\text{char}(\mathbb{Z}_5) = 5$, $\text{char}(\mathbb{Z}_3[x]) = 3$, $\text{char}(K) = 2$.
- b) Let $a_k = \underbrace{1 + \dots + 1}_k \in F$. Since F is finite we must have $a_i = a_j$ for some integers $j > i$. But then $0 = a_j - a_i = \underbrace{a_{j-i}}_{j-i} \in F$ so $\underbrace{1 + \dots + 1}_{j-i} = 0$ which shows that the characteristic is positive and $\leq j - 1$.
- c) Following the hint, assume that $\text{char}(F) = m \cdot n$ for some $m, n > 1$ and define $a_k = \underbrace{1 + \dots + 1}_k \in F$ as before. But then $a_m, a_n \neq 0$ since $m, n < \text{char}(F)$, but $a_m \cdot a_n = a_{mn} = 0$ so $0 = a_m^{-1}(a_m \cdot a_n) = a_n$ which is a contradiction. This shows that if the characteristic of F is not zero, it can not be the product of positive integers, so it is a prime number.

A 9.15. Of the listed fields, only \mathbb{C} is algebraically closed (by the fundamental theorem of algebra). For the others, here are examples of non-constant polynomials which lacks zeroes in the respective field:

$$t^2 - 2 \in \mathbb{Q}[t], \quad t^2 + 1 \in \mathbb{R}[t], \quad t^2 + 4 \in \mathbb{Z}_7[t].$$

Index

- (p, q) -tensor, 123
- $A > 0$ - positive entries, 72
- $A > B$ (entrywise inequality), 72
- $A \geq 0$ - non-negative entries, 72
- $A \otimes B$ - Kronecker product of matrices, 126
- A^* - Hermitian conjugate, 15
- A^+ - pseudo-inverse, 99
- E_λ - eigenspace for λ , 26
- $F \oplus G$ - direct sum of linear maps, 13
- $F \otimes G$ - tensor product of maps, 125
- $F^\mathbb{C}$ - complexification of a linear map, 136
- $F_\mathbb{R}$ - decomplexification of a linear map, 136
- $U \otimes V$ - tensor product of vector spaces, 116
- U^\perp - orthogonal complement of U , 55
- V/U - quotient space, 12
- V^* - dual vector space, 120
- V^{**} - double dual, 121
- $V^\mathbb{C}$ - complexification of V , 136
- $V_\mathbb{R}$ - decomplexification of V , 136
- $\text{Hom}(V, W)$ - space of maps $V \rightarrow W$, 121
- \mathbb{Z}_2 , 8
- \mathbb{Z}_n , 134
- \mathbb{Z}_p , 134
- $\mathcal{C}(\mathbb{R})$, 8
- $\mathcal{C}[a, b]$ - continuous functions on $[a, b]$, 52
- $\mathcal{C}^\infty(\mathbb{R})$, 25
- \mathcal{P}_n , 7
- δ_{ij} , 15
- $\text{Im}(F)$ - image of F , 13
- $\kappa(A)$ - condition number, 100
- $\ker(F)$ - kernel of F , 13
- $\text{Mat}_{m \times n}$, 7
- id_V - identity map on V , 13
- \oplus - direct sum, 9
- $\text{rank}(F)$ - rank of F , 13
- \sqrt{A} - square root of a matrix, 67
- $\text{cosSim}(u, v)$ - cosine similarity, 101
- \tilde{E}_λ - generalized eigenspace for λ , 40
- $|A|$ - stretching part of polar factorization, 101
- e^A - matrix exponential, 46
- e_i^* - dual basis vectors, 120
- e_{ij} , 15
- $f(A)$ - function evaluated at matrix, 46
- $p(A)$ - polynomial evaluated at a matrix, 29
- $p_A(t)$ - characteristic polynomial, 26
- $u \otimes v$ - a pure tensor, 116
- abstract field, 135
- activation (neural network), 129
- adjacency matrix, 77
- adjoint of a linear map, 62
- affine subset, 11
- affine subspace, 11
- algebraic multiplicity, 26
- angle between vectors, 52
- annihilate, 31
- approximation by affine subspace (total least squares), 113
- approximation in function spaces (examples), 58
- associated di-graph (of a matrix), 77
- axiom of choice, 137
- Banach space, 53
- basic neural network, 130
- basis, 8
- Basis theorem (for infinite-dimensional spaces), 138
- bias (neural network), 130
- bilinear map, 115
- block matrix, 22
- block-diagonal matrix, 22
- blur filter, 131
- canonical isomorphism, 118
- Cauchy-Schwarz inequality, 53
- Cayley-Hamilton theorem, 30
- center (data), 107
- center of mass, 113
- chain (in an ordered set), 138
- characteristic polynomial, 26
- Cholesky-factorization, 20
- clusters, 106
- CNN - convolutional neural network, 132
- column matrix, 92
- commute, 16
- compact SVD, 87
- complexification, 136
- complexification (via tensor product), 126
- condition number, 100
- conjugate-linear, 51
- Constructing an ON-basis, 56
- contraction (of tensors), 124
- convergence in a normed space, 55
- convolutional neural network, 132
- coordinates, 8
- correlated, 102
- cosine-similarity, 101
- covariance matrix, 108
- curve fitting with least squares, 69

- dampened PageRank matrix, 80
- dampened ranking vector (PageRank), 80
- damping factor (PageRank), 80
- data compression (with SVD), 94
- data-matrix, 101
- data-points, 103
- decomplexification, 136
- determinant, 21
- di-graph, 77
- diagonal matrix, 15
- diagonalizable, 27
- diagonalization, 27
- Differential equation system, 49
- dimension, 8
- dimension theorem, 13
- dimensionality reduction, 104
- direct sum, 9
- direct sum of linear maps, 13
- directed graph, 77
- distance between vectors, 52
- dominant eigenvalue, 72
- dot-product, 52
- double dual isomorphism theorem, 121
- dual basis, 120
- dual space, 120
- dynamical system (discrete), 44

- Eckart-Young-Mirsky theorem, 95
- edge (of graph), 77
- eigenface, 110
- eigenspace, 26
- eigenvalue, 25
- eigenvector, 25
- element-wise convergence, 46
- elementary matrix, 17
- ellipsoid, 85
- Euclidean space, 51
- external direct sum, 11

- facial recognition (example), 109
- feature (of data), 104
- Fermat's little theorem, 135
- field, 7, 133
- field axioms, 133
- filter (for images), 131
- finite field, 134
- first vector (of a string), 35
- fitting an ellipsoid (to data), 109
- Fourier-analysis, 58
- free vector space, 116
- Frobenius inner product, 52
- Frobenius norm, 53, 92
- Frobenius' theorem, 76
- functional, 62, 120

- generalized eigenspace, 40
- generalized eigenvector, 40
- generate, 8
- geometric multiplicity, 26

- Gram-Schmidt, 56
- graph, 77

- Hermitian conjugate, 15
- Hermitian matrix, 15
- hidden layer (neural network), 129
- Hilbert space, 53
- homomorphism, 12
- hyper-ellipsoid, 91

- ill-conditioned, 100
- image, 13
- image compression example, 96
- incomparable matrices, 72
- inner product, 51
- inner product space, 51
- input layer (neural network), 129
- integers modulo n - \mathbb{Z}_n , 134
- internal direct sum, 9
- invariant subspace, 28
- inverse, 13
- irreducible matrix, 76
- isometry, 63
- isomorphic vector spaces, 115
- isomorphism, 115

- Jordan basis, 41
- Jordan block, 33
- Jordan chain, 41
- Jordan form, 34
- Jordan normal form, 34
- Jordan theorem, 42
- Jordanization algorithm, 41
- Jordanize, 41

- kernel, 13
- kernel (for filters), 131
- Kronecker delta, 15
- Kronecker product, 126
- Kronecker product - properties, 128

- last vector (of a string), 35
- latent feature, 105
- LDU-decomposition, 19
- least square solution, 67
- left singular vector, 86
- line detection filter, 131
- linear combination, 8
- linear functional, 62, 120
- linear map, 12
- linear transformation, 12
- linearly independent, 8
- low rank approximation, 94
- lower triangular matrix, 15
- LU-decomposition, 18

- macaw, 97
- Maclaurin series, 46
- Manhattan norm, 53

- Markov chain, 82
- Markov chains - applications, 83
- Markov process, 82
- matrix function, 45
- matrix of a graph, 77
- maximal element (in an ordered set), 138
- maximum norm, 53
- mean (of data), 107
- minimal polynomial, 31
- mixed product property (of Kronecker product), 127
- modulo, 134
- monic polynomial, 31
- Moore-Penrose conditions, 99
- Moore-Penrose pseudo inverse (general definition), 99
- Moore-Penrose pseudo inverse (independent columns case), 68
- movie ratings example (data analysis), 104
- multi-dimensional arrays, 118
- multi-graph, 77
- multidimensional data (in neural networks), 130
- multilinear map, 122

- n-fold tensor product, 122
- natural isomorphism, 118
- negative definite operator, 65
- negative semi-definite operator, 65
- neural network, 129
- neural network (basic definition), 130
- neuron, 129
- nilpotency-degree, 16
- nilpotent, 16
- node, 77
- non-negative matrix, 72
- norm (general definition), 53
- norm derived from inner product, 52
- normal equations, 67
- normal matrix, 63
- normal operator, 63
- normed space, 53
- nuclear norm, 92
- nullspace, 13

- ON-basis, 55
- ON-matrix, 63
- operator, 12
- operator norm, 54, 92
- order (of tensor product), 122
- order (of tensor), 122
- orthogonal basis, 55
- orthogonal complement of a subspace, 55
- orthogonal matrix, 63
- Orthogonal projection onto a subspace, 56
- Orthogonal projection onto a vector, 56
- orthogonal vectors, 52
- orthonormal basis, 55
- orthonormal matrix, 63
- output layer (neural network), 129

- p-norm, 53
- PageRank, 79
- PageRank (example), 79
- PageRank - applications, 81
- PageRank-matrix, 80
- parallelogram identity, 53
- partial order, 72
- partially ordered set, 137
- PCA, 108
- permutation matrix, 19
- Perron eigenvalue, 72
- Perron vector, 72
- Perron's theorem/, 72
- pivots, 17
- polar decomposition, 100
- polarization, 53
- pooling layer (neural network), 132
- positive definite, 20
- positive definite matrix, 65
- positive definite operator, 65
- positive matrix, 72
- positive semi-definite operator, 65
- positive semidefinite, 20
- positive vector, 72
- power method, 75
- pre-composition, 122
- primitive matrix, 75
- principal axes, 108
- principal component analysis, 108
- principal components, 108
- principal minor, 66
- projection (via direct sum), 10
- Pseudo inverse (independent columns case), 68
- pure tensor, 122
- pure tensors, 116

- QR-algorithm, 61
- QR-decomposition, 59
- QR-decomposition (thin), 60
- quotient space, 12

- random surfer, 81
- range, 13
- rank, 13
- rank-nullity theorem, 13
- ranking algorithm, 79
- ranking models, 79
- ranking vector, 79
- realification, 136
- Recursively defined sequence, 45
- reduced row echelon form, 17
- reducible matrix, 76
- REF - row echelon form, 17
- relative error, 100
- ReLU, 130
- restriction to an invariant subspace, 28
- Riesz representation theorem, 62
- right singular vector, 86
- ring, 133

- root of a matrix, 67
- row echelon form, 17
- row matrix, 92
- RREF - reduced row echelon form, 17

- scalar, 7
- scalar multiplication, 7
- scalar product, 51
- Schmidt decomposition (of a matrix), 93
- Schur's theorem, 63
- secular polynomial, 26
- self-adjoint linear map, 63
- semi-axis, 85
- sesqui-linear, 51
- sigmoid, 130
- similar matrices, 16
- singular value, 86
- singular value decomposition, 86
- singular vector, 86
- skew-Hermitian matrix, 15
- skew-symmetric matrix, 15
- span, 8
- spectral mapping theorem, 32
- spectral norm, 54, 92
- spectral theorem (complex case), 64
- spectral theorem (real case), 63
- spectral theory, 25
- spectrum of a linear operator, 25
- spider trap, 81
- square root of a matrix, 67
- stabilize, 41
- standard deviation (of data), 109
- state (of Markov chain), 82
- stationary distribution, 83
- steady state distribution, 83
- stochastic matrix, 82
- stochastic model, 82
- string (for nilpotent map), 35
- string basis (for a nilpotent map), 35
- string basis algorithm, 39
- strongly connected graph, 77
- subspace, 8
- super-diagonal, 33
- supremum norm, 54
- SVD - algorithm, 88
- SVD - singular value decomposition, 86
- Sylvester's criterion, 66
- symmetric matrix, 15
- System of linear differential equations, 49

- tall matrix, 88
- taxicab norm, 53
- tensor product, 116
- tensor product (construction), 116
- tensor product (of linear maps), 125
- tensors, 116
- topology (of data), 130
- total least squares, 113
- totally ordered set, 138

- trace, 16
- training (neural network), 129
- transition matrix, 82
- transpose, 15
- triangle inequality for norms, 53
- triangular matrix, 15
- truncated SVD, 94

- uncorrelated, 102
- unit circle (for different norms), 54
- unitarily equivalence, 63
- unitary linear map, 63
- unitary matrix, 63
- universal property, 116
- upper bound (in an ordered set), 138
- upper triangular matrix, 15

- variables (in data), 101
- variation (in data), 102
- vector, 7
- vector space, 7
- vector space axioms, 7
- vertex (of graph), 77

- walk (in a graph), 78
- web-pages (ranking models), 79
- weighted graph, 77
- well-conditioned, 100
- wide matrix, 89

- Zorn's lemma, 137