

A (TERSE) INTRODUCTION TO

Linear Algebra

Yitzhak Katznelson

(DRAFT)

Contents

| | | |
|------------|--|-----------|
| I | Vector spaces | 1 |
| 1.1 | Vector spaces | 1 |
| 1.2 | Linear dependence, bases, and dimension | 9 |
| 1.3 | Systems of linear equations. | 14 |
| II | Linear operators and matrices | 23 |
| 2.1 | Linear Operators (maps, transformations) | 23 |
| 2.2 | Operator Multiplication | 27 |
| 2.3 | Matrix multiplication. | 28 |
| 2.4 | Matrices and operators. | 31 |
| 2.5 | Kernel, range, nullity, and rank | 35 |
| ★2.6 | Normed finite dimensional linear spaces | 39 |
| III | Duality of vector spaces | 43 |
| 3.1 | Linear functionals | 43 |
| 3.2 | The adjoint | 47 |
| IV | Determinants | 51 |
| 4.1 | Permutations | 51 |
| 4.2 | Multilinear maps | 53 |
| 4.3 | Alternating n -forms | 56 |
| 4.4 | Determinant of an operator | 58 |
| 4.5 | Determinant of a matrix | 61 |
| V | Invariant subspaces | 65 |
| 5.1 | Invariant subspaces | 65 |
| 5.2 | The minimal polynomial | 69 |

| | | |
|------------|---|------------|
| 5.3 | Reducing. | 75 |
| 5.4 | Semisimple systems. | 81 |
| 5.5 | Nilpotent operators | 83 |
| ★5.6 | The cyclic decomposition | 86 |
| 5.7 | The Jordan canonical form | 88 |
| 5.8 | Functions of an operator | 92 |
| VI | Operators on inner-product spaces | 95 |
| 6.1 | Inner-product spaces | 95 |
| 6.2 | Duality and the adjoint. | 102 |
| 6.3 | Unitary and orthogonal operators | 104 |
| 6.4 | Self-adjoint operators | 105 |
| 6.5 | Normal operators. | 109 |
| 6.6 | Positive operators. | 111 |
| 6.7 | Polar decomposition | 112 |
| VII | Additional topics | 115 |
| 7.1 | Quadratic forms | 115 |
| 7.2 | Positive matrices | 118 |
| 7.3 | Nonnegative matrices | 121 |
| 7.4 | Stochastic matrices. | 126 |
| 7.5 | Representation of finite groups | 129 |
| A | Appendix | 137 |
| A.1 | Equivalence relations — partitions. | 137 |
| A.2 | Maps | 138 |
| A.3 | Groups | 139 |
| ★A.4 | Group actions | 142 |
| A.5 | Fields, Rings, and Algebras | 144 |
| A.6 | Polynomials | 147 |
| | Index | 153 |

Chapter I

Vector spaces

1.1 VECTOR SPACES

The notions of *group* and *field* are defined in the Appendix: A.3 and A.5.1 respectively.

The fields \mathbb{Q} (of rational numbers), \mathbb{R} (of real numbers), and \mathbb{C} (of complex numbers) are familiar, and are the most commonly used. Most of the notions and results we discuss are valid for vector spaces over arbitrary underlying fields. When we do not need to specify the underlying field we denote it by the generic \mathbb{F} and refer to its elements as scalars. Results that require specific fields will be stated explicitly in terms of the appropriate field.

1.1.1 DEFINITION: A *vector space* \mathcal{V} over a *field* \mathbb{F} is an abelian group (the group operation written as addition) and a binary product $(a, v) \mapsto av$ of $\mathbb{F} \times \mathcal{V}$ into \mathcal{V} , satisfying the following conditions:

v-s 1. $1v = v$

v-s 2. $a(bv) = (ab)v,$

v-s 3. $(a + b)v = av + bv, \quad a(v + u) = av + au.$

A *real vector space* is a vector space over the field \mathbb{R} ; A *complex vector space* is one over the field \mathbb{C} .

Vector spaces may have additional geometric structure, such as *inner product*, which we study in Chapter VI, or additional algebraic structure, such as multiplication, which we just mention in passing.

EXAMPLES:

- a.** \mathbb{F}^n , the space of all \mathbb{F} -valued n -tuples (a_1, \dots, a_n) with addition and scalar multiplication defined by

$$\begin{aligned}(a_1, \dots, a_n) + (b_1, \dots, b_n) &= (a_1 + b_1, \dots, a_n + b_n) \\ c(a_1, \dots, a_n) &= (ca_1, \dots, ca_n)\end{aligned}$$

If the underlying field is \mathbb{R} , resp. \mathbb{C} , we denote the space \mathbb{R}^n , resp. \mathbb{C}^n .

We write the n -tuples as rows, as we did here, or as columns. (We sometime write \mathbb{F}_c^n resp. \mathbb{F}_r^n when we want to specify that vectors are written as columns, resp. rows.)

- b.** $\mathcal{M}(n, m; \mathbb{F})$, the space of all \mathbb{F} -valued $n \times m$ matrices, that is, arrays

$$A = \begin{bmatrix} a_{11} & \dots & a_{1m} \\ a_{21} & \dots & a_{2m} \\ \vdots & \dots & \vdots \\ a_{n1} & \dots & a_{nm} \end{bmatrix}$$

with entries from \mathbb{F} . The addition and scalar multiplication are again done entry by entry. As a vector space $\mathcal{M}(n, m; \mathbb{F})$ is virtually identical with \mathbb{F}^{mn} , except that we write the entries in the rectangular array instead of a row or a column.

We write $\mathcal{M}(n; \mathbb{F})$ instead of $\mathcal{M}(n, n; \mathbb{F})$ and when the underlying field is either assumed explicitly, or is arbitrary, we may write simply $\mathcal{M}(n, m)$ or $\mathcal{M}(n)$, as the case may be.

- c.** $\mathbb{F}[x]$, the space¹ of all polynomials $\sum a_n x^n$ with coefficients from \mathbb{F} . Addition and multiplication by scalars are defined formally either as the standard addition and multiplication of functions, or by adding (and multiplying by scalars) the corresponding coefficients. The two ways define the same operations.

¹ $\mathbb{F}[x]$ is an algebra over \mathbb{F} , i.e., a vector space with an additional structure, multiplication. See A.5.2.

- d.** The set $C_{\mathbb{R}}([0, 1])$ of all continuous real-valued functions f on $[0, 1]$, and the set $C([0, 1])$ of all continuous complex-valued functions f on $[0, 1]$, with the standard operations of addition and of multiplication of functions by a scalars.

$C_{\mathbb{R}}([0, 1])$ is a real vector space. $C([0, 1])$ is naturally a complex vector space, but becomes a real vector space if we limit the allowable scalars to real numbers only.

- e.** The set $C^{\infty}([-1, 1])$ of all infinitely differentiable real-valued functions f on $[-1, 1]$, with the standard operations on functions.
- f.** The set \mathcal{T}_N of 2π -periodic trigonometric polynomials of degree $\leq N$: the functions admitting a representation as a sum of the form $\sum_{|n| \leq N} a_n e^{inx}$. Standard operations on functions.
- g.** The set of functions f which satisfy the differential equation

$$3f'''(x) - \sin x f''(x) + 2f(x) = 0.$$

Standard operations.

1.1.2 ISOMORPHISM. The expression “virtually identical” in the comparison, in Example **b.** above, of $\mathcal{M}(n, m; \mathbb{F})$ with \mathbb{F}^{mn} , is not a proper mathematical term. The proper term here is *isomorphic*.

DEFINITION: A map $\varphi: \mathcal{V}_1 \mapsto \mathcal{V}_2$ is called *linear* if, for all scalars a, b and vectors $v_1, v_2 \in \mathcal{V}_1$

$$(1.1.1) \quad \varphi(av_1 + bv_2) = a\varphi(v_1) + b\varphi(v_2).$$

Two vector spaces \mathcal{V}_1 and \mathcal{V}_2 over the same field are *isomorphic* if there exist a bijective² linear map $\varphi: \mathcal{V}_1 \mapsto \mathcal{V}_2$.

²That is φ maps \mathcal{V}_1 onto \mathcal{V}_2 and the map is 1 – 1 (and linear); see Appendix A.2.

1.1.3 SUBSPACES. A (vector) *subspace* of a vector space \mathcal{V} is a subset which is closed under the operations of addition and multiplication by scalars defined in \mathcal{V} .

In other words, $\mathcal{W} \subset \mathcal{V}$ is a subspace if $a_1w_1 + a_2w_2 \in \mathcal{W}$ for all scalars a_j and vectors $w_j \in \mathcal{W}$.

EXAMPLES:

a. Solution-set of a system of homogeneous linear equations.

Here $\mathcal{V} = \mathbb{F}^n$. Given the scalars a_{ij} , $1 \leq i \leq k$, $1 \leq j \leq n$ we consider the solution-set of the system of k homogeneous linear equations

$$(1.1.2) \quad \sum_{j=1}^n a_{ij}x_j = 0, \quad i = 1, \dots, k.$$

This is the set of all n -tuples $(x_1, \dots, x_n) \in \mathbb{F}^n$ for which all k equations are satisfied. If both (x_1, \dots, x_n) and (y_1, \dots, y_n) are solutions of (1.1.2), and a and b are scalars, then for each i ,

$$\sum_{j=1}^n a_{ij}(ax_j + by_j) = a \sum_{j=1}^n a_{ij}x_j + b \sum_{j=1}^n a_{ij}y_j = 0.$$

It follows that the solution-set of (1.1.2) is a subspace of \mathbb{F}^n .

b. In the space $C^\infty(\mathbb{R})$ of all infinitely differentiable real-valued functions f on \mathbb{R} with the standard operations, the set of functions f that satisfy the differential equation

$$f'''(x) - 5f''(x) + 2f'(x) - f(x) = 0.$$

Again, we can include, if we want, complex valued functions and allow, if we want, complex scalars.

c. Subspaces of $\mathcal{M}(n)$:

The set of *diagonal matrices*—the $n \times n$ matrices with zero entries off the diagonal ($a_{ij} = 0$ for $i \neq j$).

The set of (*lower*) *triangular matrices*—the $n \times n$ matrices with zero entries above the diagonal ($a_{ij} = 0$ for $i < j$).

Similarly set of upper triangular matrices, ($a_{ij} = 0$ for $i > j$).

d. Intersection of subspaces: If \mathcal{W}_j are subspaces of a space \mathcal{V} , then $\cap \mathcal{W}_j$ is a subspace of \mathcal{V} .

e. The sum³ of subspaces: $\sum \mathcal{W}_j$ is defined by

$$\sum \mathcal{W}_j = \{ \sum v_j : v_j \in \mathcal{W}_j \}.$$

f. The span of a subset: The span of a subset $E \subset \mathcal{V}$, denoted $\text{span}[E]$, is the set $\{ \sum a_j e_j : a_j \in \mathbb{F}, e_j \in E \}$ of all the finite linear combinations of elements of E . $\text{span}[E]$ is a subspace; clearly the smallest subspace of \mathcal{V} that contains E .

1.1.4 DIRECT SUMS. If $\mathcal{V}_1, \dots, \mathcal{V}_k$ are vector spaces over \mathbb{F} , the (*formal*) *direct sum* $\bigoplus_1^k \mathcal{V}_j = \mathcal{V}_1 \oplus \dots \oplus \mathcal{V}_k$ is the set $\{ (v_1, \dots, v_k) : v_j \in \mathcal{V}_j \}$ in which we define addition:

$$(v_1, \dots, v_k) + (u_1, \dots, u_k) = (v_1 + u_1, \dots, v_k + u_k),$$

and multiplication by scalars: $a(v_1, \dots, v_k) = (av_1, \dots, av_k)$.

DEFINITION: The subspaces \mathcal{W}_j , $j = 1, \dots, k$ of a vector space \mathcal{V} are *independent* if $\sum v_j = 0$ with $v_j \in \mathcal{W}_j$ implies that $v_j = 0$ for all j .

Proposition. *If \mathcal{W}_j are subspaces of \mathcal{V} , then the map Φ of $\mathcal{W}_1 \oplus \dots \oplus \mathcal{W}_k$ into $\mathcal{W}_1 + \dots + \mathcal{W}_k$, defined by*

$$\Phi : (v_1, \dots, v_k) \mapsto v_1 + \dots + v_k,$$

is an isomorphism if, and only if, the subspaces are independent.

³Don't confuse the *sum of subspaces* with the *union of subspaces* which is seldom a subspace, see exercise **I.1.5** below.

PROOF: Φ is clearly linear and surjective. To prove it injective we need to check that every vector in the range has a unique preimage, that is, to show that

$$(1.1.3) \quad v'_j, v''_j \in \mathcal{W}_j \quad \text{and} \quad v''_1 + \cdots + v''_k = v'_1 + \cdots + v'_k$$

implies that $v''_j = v'_j$ for every j . Subtracting and writing $v_j = v''_j - v'_j$, (1.1.3) is equivalent to: $\sum v_j = 0$ with $v_j \in \mathcal{W}_j$, which implies that $v_j = 0$ for all j . \blacktriangleleft

Notice that Φ is the “natural” map of the formal direct sum onto the sum of subspaces of a given space.

In view of the proposition we refer to the sum $\sum \mathcal{W}_j$ of *independent* subspaces of a vector space as *direct sum* and write $\bigoplus \mathcal{W}_j$ instead of $\sum \mathcal{W}_j$.

If $\mathcal{V} = \mathcal{U} \oplus \mathcal{W}$, we refer to either \mathcal{U} or \mathcal{W} as a *complement* of the other in \mathcal{V} .

1.1.5 QUOTIENT SPACES. A subspace \mathcal{W} of a vector space \mathcal{V} defines an equivalence relation⁴ in \mathcal{V} :

$$(1.1.4) \quad x \equiv y \pmod{\mathcal{W}} \quad \text{if} \quad x - y \in \mathcal{W}.$$

In order to establish that this is indeed an equivalence relation we need to check that it is

a. *reflexive* (clear, since $x - x = 0 \in \mathcal{W}$),

b. *symmetric* (clear, since if $x - y \in \mathcal{W}$, then $y - x = -(x - y) \in \mathcal{W}$),

and

c. *transitive*, (if $x - y \in \mathcal{W}$ and $y - z \in \mathcal{W}$, then $x - z = (x - y) + (y - z) \in \mathcal{W}$).

The equivalence relation partitions \mathcal{V} into cosets or “translates” of \mathcal{W} , that is into sets of the form $x + \mathcal{W} = \{v : v = x + w, w \in \mathcal{W}\}$.

So far we used only the group structure and not the fact that addition in \mathcal{V} is commutative, nor the fact that we can multiply by scalars. This information will be used now.

We define the *quotient space* \mathcal{V}/\mathcal{W} to be the space whose elements are the equivalence classes mod \mathcal{W} in \mathcal{V} , and whose vector space structure, addition and multiplication by scalars, is given by:

⁴See Appendix A.1

if $\tilde{x} = x + \mathcal{W}$ and $\tilde{y} = y + \mathcal{W}$ are cosets, and $a \in \mathbb{F}$, then

$$(1.1.5) \quad \tilde{x} + \tilde{y} = x + y + \mathcal{W} = \widetilde{x + y} \quad \text{and} \quad a\tilde{x} = \widetilde{ax}.$$

The definition needs justification. We defined the sum of two cosets by taking one element of each, adding them and taking the coset containing the sum as the sum of the cosets. We need to show that the result is well defined, i.e., that it does not depend on the choice of the representatives in the cosets. In other words, we need to verify that if $x \equiv x_1 \pmod{\mathcal{W}}$ and $y \equiv y_1 \pmod{\mathcal{W}}$, then $x + y \equiv x_1 + y_1 \pmod{\mathcal{W}}$. But, $x = x_1 + w$, $y = y_1 + w'$ with $w, w' \in \mathcal{W}$ implies that $x + y = x_1 + w + y_1 + w' = x_1 + y_1 + w + w'$, and, since $w + w' \in \mathcal{W}$ we have $x + y \equiv x_1 + y_1 \pmod{\mathcal{W}}$.

Notice that the “switch” $w + y_1 = y_1 + w$ is justified by the commutativity of the addition in \mathcal{V} .

The definition of $a\tilde{x}$ is justified similarly: assuming $x \equiv x_1 \pmod{\mathcal{W}}$ then $ax - ax_1 = a(x - x_1) \in \mathcal{W}$, (since \mathcal{W} is a *subspace*, closed under multiplication by scalars) and $ax \equiv ax_1 \pmod{\mathcal{W}}$.

1.1.6 TENSOR PRODUCTS. Given vector spaces \mathcal{V} and \mathcal{U} over \mathbb{F} , the set of all the (formal) sums $\sum a_j v_j \otimes u_j$, where $a_j \in \mathbb{F}$, $v_j \in \mathcal{V}$ and $u_j \in \mathcal{U}$; with (formal) addition and multiplication by elements of \mathbb{F} , is a vector space over \mathbb{F} .

The tensor product $\mathcal{V} \otimes \mathcal{U}$ is, by definition, the quotient of this space by the subspace spanned by the elements of the form

$$(1.1.6) \quad \begin{aligned} \text{a.} & \quad (v_1 + v_2) \otimes u - (v_1 \otimes u + v_2 \otimes u), \\ \text{b.} & \quad v \otimes (u_1 + u_2) - (v \otimes u_1 + v \otimes u_2), \\ \text{c.} & \quad a(v \otimes u) - (av) \otimes u, \quad (av) \otimes u - v \otimes (au), \end{aligned}$$

for all $v, v_j \in \mathcal{V}$, $u, u_j \in \mathcal{U}$ and $a \in \mathbb{F}$.

In other words, $\mathcal{V} \otimes \mathcal{U}$ is the space of formal sums $\sum a_j v_j \otimes u_j$ modulo the the equivalence relation generated by:

$$(1.1.7) \quad \begin{aligned} \text{a.} & \quad (v_1 + v_2) \otimes u \equiv v_1 \otimes u + v_2 \otimes u, \\ \text{b.} & \quad v \otimes (u_1 + u_2) \equiv v \otimes u_1 + v \otimes u_2, \\ \text{c.} & \quad a(v \otimes u) \equiv (av) \otimes u \equiv v \otimes (au). \end{aligned}$$

Example. If $\mathcal{V} = \mathbb{F}[x]$ and $\mathcal{U} = \mathbb{F}[y]$, then $p(x) \otimes q(y)$ can be identified with the product $p(x)q(y)$ and $\mathcal{V} \otimes \mathcal{U}$ with $\mathbb{F}[x, y]$.

EXERCISES FOR SECTION 1.1

I.1.1. Verify that \mathbb{R} is a vector space over \mathbb{Q} , and that \mathbb{C} is a vector space over either \mathbb{Q} or \mathbb{R} .

I.1.2. Verify that the intersection of subspaces is a subspace.

I.1.3. Verify that the sum of subspaces is a subspace.

I.1.4. Prove that $\mathcal{M}(n, m; \mathbb{F})$ and \mathbb{F}^{mn} are isomorphic.

I.1.5. Let \mathcal{U} and \mathcal{W} be proper subspaces of a vector space \mathcal{V} , neither of them contains the other. Show that $\mathcal{U} \cup \mathcal{W}$ is *not* a subspace.

Hint: Take $u \in \mathcal{U} \setminus \mathcal{W}$, $w \in \mathcal{W} \setminus \mathcal{U}$ and consider $u + w$.

***I.1.6.** If \mathbb{F} is finite, $n > 1$, then \mathbb{F}^n is a union of a finite number of lines. Assuming that \mathbb{F} is infinite, show that the union of a finite number of subspaces of \mathcal{V} , none of which contains all others, is not a subspace.

Hint: Let \mathcal{V}_j , $j = 1, \dots, k$ be the subspaces in question. Show that there is no loss in generality in assuming that their union spans \mathcal{V} . Now you need to show that $\bigcup \mathcal{V}_j$ is not all of \mathcal{V} . Show that there is no loss of generality in assuming that \mathcal{V}_1 is not contained in the union of the others. Take $v_1 \in \mathcal{V}_1 \setminus \bigcup_{j \neq 1} \mathcal{V}_j$, and $w \notin \mathcal{V}_1$; show that $av_1 + w \in \bigcup \mathcal{V}_j$, $a \in \mathbb{F}$, for no more than k values of a .

I.1.7. Let $p > 1$ be a positive integer. Recall that two integers, m, n are *congruent* (mod p), written $n \equiv m \pmod{p}$, if $n - m$ is divisible by p . This is an equivalence relation (see Appendix A.1). For $m \in \mathbb{Z}$, denote by \tilde{m} the coset (equivalence class) of m , that is the set of all integers n such that $n \equiv m \pmod{p}$.

- a. Every integer is congruent (mod p) to one of the numbers $[0, 1, \dots, p - 1]$. In other words, there is a 1 - 1 correspondence between \mathbb{Z}_p , the set of cosets (mod p), and the integers $[0, 1, \dots, p - 1]$.
- b. As in subsection 1.1.5 above, we define the *quotient ring* $\mathbb{Z}_p = \mathbb{Z}/(p)$ (both notations are common) as the space whose elements are the cosets (mod p) in \mathbb{Z} , and define addition and multiplication by: $\tilde{m} + \tilde{n} = \widetilde{(m + n)}$ and $\tilde{m} \cdot \tilde{n} = \widetilde{m \cdot n}$. Prove that the addition and multiplication so defined are associative, commutative and satisfy the distributive law.

c. Prove that \mathbb{Z}_p , endowed with these operations, is a field if, and only if, p is prime.

Hint: You may use the following fact: if p is a prime, and both n and m are *not* divisible by p then nm is not divisible by p . Show that this implies that if $\tilde{n} \neq 0$ in \mathbb{Z}_p , then $\{\tilde{n}\tilde{m} : \tilde{m} \in \mathbb{Z}_p\}$ covers all of \mathbb{Z}_p .

1.2 LINEAR DEPENDENCE, BASES, AND DIMENSION

Let \mathcal{V} be a vector space. A *linear combination of vectors* v_1, \dots, v_k is a sum of the form $v = \sum a_j v_j$ with scalar coefficients a_j .

A linear combination is *non-trivial* if at least one of the coefficients is not zero.

1.2.1 Recall that *The span* of a set $A \subset \mathcal{V}$, denoted $\text{span}[A]$, is the set of all vectors v that can be written as linear combinations of elements in A .

DEFINITION: A set $A \subset \mathcal{V}$ is a *spanning set* if $\text{span}[A] = \mathcal{V}$.

1.2.2 DEFINITION: A set $A \subset \mathcal{V}$ is *linearly independent* if for every sequence $\{v_1, \dots, v_l\}$ of distinct vectors in A , the only vanishing linear combination of the v_j 's is trivial; that is, if $\sum a_j v_j = 0$ then $a_j = 0$ for all j .

If the set A is finite, we enumerate its elements as v_1, \dots, v_m and write the elements in its span as $\sum a_j v_j$. By definition, independence of A means that the representation of $v = 0$ is unique. Notice, however, that this implies that the representation of *every* vector in $\text{span}[A]$ is unique, since $\sum_1^l a_j v_j = \sum_1^l b_j v_j$ implies $\sum_1^l (a_j - b_j) v_j = 0$ so that $a_j = b_j$ for all j .

1.2.3 A *minimal spanning set* is a spanning set such that no proper subset thereof is spanning.

A *maximal independent set* is an independent set such that no set that contains it properly is independent.

Lemma.

a. A minimal spanning set is independent.

b. A maximal independent set is spanning.

PROOF: **a.** Let A be a minimal spanning set. If $\sum a_j v_j = 0$, with distinct $v_j \in A$, and for some k , $a_k \neq 0$, then $v_k = -a_k^{-1} \sum_{j \neq k} a_j v_j$. This permits the substitution of v_k in any linear combination by the combination of the other

v_j 's, and shows that v_k is redundant: the span of $\{v_j : j \neq k\}$ is the same as the original span, contradicting the minimality assumption.

b. If B is independent and $u \notin \text{span}[B]$, then the union $\{u\} \cup B$ is independent: assume otherwise, then there exists $\{v_1, \dots, v_l\} \subset B$ and coefficients d and c_j , not all zero, such that $du + \sum c_j v_j = 0$. Assuming $d \neq 0$ implies $u = -d^{-1} \sum c_j v_j$ and u would be in $\text{span}[v_1, \dots, v_l] \subset \text{span}[B]$, contradicting the assumption $u \notin \text{span}[B]$; so $d = 0$. But now $\sum c_j v_j = 0$ with some non-vanishing coefficients, contradicting the assumption that B is independent.

It follows that if B is maximal independent, then $u \in \text{span}[B]$ for every $u \in \mathcal{V}$, and B is spanning. ◀

DEFINITION: A *basis* for \mathcal{V} is an independent spanning set in \mathcal{V} . Thus, $\{v_1, \dots, v_n\}$ is a basis for \mathcal{V} if, and only if, every $v \in \mathcal{V}$ has a unique representation as a linear combination of $\{v_1, \dots, v_n\}$, that is a representation (or expansion) of the form $v = \sum a_j v_j$. By the lemma, a minimal spanning set is a basis, and a maximal independent set is a basis.

A *finite dimensional vector space* is a vector space that has a finite basis. (See also Definition 1.2.4.)

Theorem. *If \mathcal{V} is finite dimensional then:*

- a. Every spanning set can be trimmed to a basis.*
- b. Every independent set can be expanded to a basis.*

PROOF: **a.** Let $\{v_j\}_{j=1}^N$ be a spanning set for \mathcal{V} . Call *inessential* a vector v_l that is linearly dependent on $\{v_j\}_{j=1}^{l-1}$, and *essential* otherwise. Observe that an inessential v_l is linearly dependent on the *essential* vectors preceding it.

Remove the inessential vectors. Since every v_j is either essential or linearly dependent on the preceding essential vectors, the essential vectors span \mathcal{V} and are independent, hence form a basis.

b. Let $\{u_j\}_{j=1}^k$ be independent, and let $\{e_j\}_{j=1}^n$ be a basis for \mathcal{V} . Write $w_j = u_j$ for $j = 1, \dots, k$, and $w_{k+j} = e_j$ for $j = 1, \dots, n$. The sequence $\{w_j\}$ contains the basis $\{e_j\}$ and is therefore spanning. Now remove, as in part **a.** the inessential vectors to obtain a basis, and observe that the first k vectors, namely $\{u_j\}_{j=1}^k$ are all essential, and hence form part of the basis. ◀

EXAMPLES:

- a.** In \mathbb{F}^n we write e_j for the vector whose j 'th entry is equal to 1 and all the other entries are zero. $\{e_1, \dots, e_n\}$ is a basis for \mathbb{F}^n , and the unique representation of $v = \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix}$ in terms of this basis is $v = \sum a_j e_j$. We refer to $\{e_1, \dots, e_n\}$ as *the standard basis* for \mathbb{F}^n .
- b.** *The standard basis* for $\mathcal{M}(n, m)$: let e_{ij} denote the $n \times m$ matrix whose ij 'th entry is 1 and all the other zero. $\{e_{ij}\}$ is a basis for $\mathcal{M}(n, m)$, and $\begin{bmatrix} a_{11} & \dots & a_{1m} \\ a_{21} & \dots & a_{2m} \\ \vdots & \dots & \vdots \\ a_{n1} & \dots & a_{nm} \end{bmatrix} = \sum a_{ij} e_{ij}$ is the expansion.
- c.** The space $\mathbb{F}[x]$ is not finite dimensional. The infinite sequence $\{x^n\}_{n=0}^\infty$ is linearly independent, in fact a basis, and, as we see in the following subsection, it cannot have a finite basis.

1.2.4 STEINITZ' LEMMA AND THE DEFINITION OF DIMENSION.

Lemma (Steinitz). *Assume $\text{span}[v_1, \dots, v_n] = \mathcal{V}$ and $\{u_1, \dots, u_m\}$ linearly independent in \mathcal{V} . Claim: the vectors v_j can be (re)ordered so that, for every $k = 1, \dots, m$, the sequence $\{u_1, \dots, u_k, v_{k+1}, \dots, v_n\}$ spans \mathcal{V} .*

In particular, $m \leq n$.

PROOF: Write $u_1 = \sum a_j v_j$, possible since $\text{span}[v_1, \dots, v_n] = \mathcal{V}$. Reorder the v_j 's, if necessary, to guarantee that $a_1 \neq 0$.

Now $v_1 = a_1^{-1}(u_1 - \sum_{j=2}^n a_j v_j)$, which means that $\text{span}[u_1, v_2, \dots, v_n]$ contains every v_j and hence is equal to \mathcal{V} .

Continue recursively: assume that, having reordered the v_j 's if necessary, we have $\{u_1, \dots, u_k, v_{k+1}, \dots, v_n\}$ spans \mathcal{V} .

Observe that unless $k = m$, we have $k < n$ (since u_{k+1} is not in the span of $\{u_1, \dots, u_k\}$ at least one additional v is needed). If $k = m$ we are done. If $k < m$ we write $u_{k+1} = \sum_{j=1}^k a_j u_j + \sum_{j=k+1}^n b_j v_j$, and since $\{u_1, \dots, u_m\}$ is linearly independent, at least one of the coefficients b_j is not zero. Reordering the remaining v_j 's if necessary, we may assume that $b_{k+1} \neq 0$ and obtain,

as before, that $v_{k+1} \in \text{span}[u_1, \dots, u_{k+1}, v_{k+2}, \dots, v_n]$, and, once again, the span is \mathcal{V} . Repeating the step (a total of) m times proves the claim of the lemma. ◀

Theorem. *If $\{v_1, \dots, v_n\}$ and $\{u_1, \dots, u_m\}$ are both bases, then $m = n$.*

PROOF: Since $\{v_1, \dots, v_n\}$ is spanning and $\{u_1, \dots, u_m\}$ independent we have $m \leq n$. Reversing the roles we have $n \leq m$. ◀

Steinitz' lemma is a refinement of part **b.** of theorem 1.2.3: in a finite dimensional vector space, every independent set can be expanded to a basis by adding, if necessary, elements from any given spanning set. The additional information here, that any spanning set has at least as many elements as any independent set, that is the basis for the current theorem, is what enables the definition of dimension.

DEFINITION: A vector space \mathcal{V} is *finite dimensional* if it has a finite basis. *The dimension*, $\dim \mathcal{V}$ is the number of elements in any basis for \mathcal{V} . (Well defined since all bases have the same cardinality.)

As you are asked to check in Exercise **I.2.9** below, a subspace \mathcal{W} of a finite dimensional space \mathcal{V} is finite dimensional and, unless $\mathcal{W} = \mathcal{V}$, the dimension $\dim \mathcal{W}$ of \mathcal{W} is strictly lower than $\dim \mathcal{V}$.

The *codimension* of a subspace \mathcal{W} in \mathcal{V} is, by definition, $\dim \mathcal{V} - \dim \mathcal{W}$.

1.2.5 The following observation is sometimes useful.

Proposition. *Let \mathcal{U} and \mathcal{W} be subspaces of an n dimensional space \mathcal{V} , and assume that $\dim \mathcal{U} + \dim \mathcal{W} > n$. Then $\mathcal{U} \cap \mathcal{W} \neq \{0\}$.*

PROOF: Let $\{u_j\}_{j=1}^l$ be a basis for \mathcal{U} and $\{w_j\}_{j=1}^m$ be a basis for \mathcal{W} . Since $l + m > n$ the set $\{u_j\}_{j=1}^l \cup \{w_j\}_{j=1}^m$ is linearly dependent, i.e., there exist a nontrivial vanishing linear combination $\sum c_j u_j + \sum d_j w_j = 0$. If all the coefficients c_j were zero, we would have a vanishing nontrivial combination of the basis elements $\{w_j\}_{j=1}^m$, which is ruled out. Similarly not all the d_j 's vanish. We now have the nontrivial $\sum c_j u_j = -\sum d_j w_j$ in $\mathcal{U} \cap \mathcal{W}$. ◀

EXERCISES FOR SECTION 1.2

I.2.1. The set $\{v_j : 1 \leq j \leq k\}$ is linearly dependent if, and only if, $v_1 = 0$ or there exists $l \in [2, k]$ such that v_l is a linear combination of vectors in $\{v_j : 1 \leq j \leq l-1\}$.

I.2.2. Let \mathcal{V} be a vector space, $\mathcal{W} \subset \mathcal{V}$ a subspace. Let $v, u \in \mathcal{V} \setminus \mathcal{W}$, and assume that $u \in \text{span}[\mathcal{W}, v]$. Prove that $v \in \text{span}[\mathcal{W}, u]$.

I.2.3. What is the dimension of \mathbb{C}^5 considered as a vector space over \mathbb{R} ?

I.2.4. Is \mathbb{R} finite dimensional over \mathbb{Q} ?

I.2.5. Is \mathbb{C} finite dimensional over \mathbb{R} ?

I.2.6. Check that for every $A \subset \mathcal{V}$, $\text{span}[A]$ is a subspace of \mathcal{V} , and is the smallest subspace containing A .

I.2.7. Let \mathcal{U}, \mathcal{W} be subspaces of a vector space \mathcal{V} , and assume $\mathcal{U} \cap \mathcal{W} = \{0\}$. Assume that $\{u_1, \dots, u_k\} \subset \mathcal{U}$ and $\{w_1, \dots, w_l\} \subset \mathcal{W}$ are (each) linearly independent. Prove that $\{u_1, \dots, u_k\} \cup \{w_1, \dots, w_l\}$ is linearly independent.

I.2.8. Prove that the subspaces $\mathcal{W}_j \subset \mathcal{V}, j = 1, \dots, N$ are independent (see Definition 1.1.4) if, and only if, $\mathcal{W}_j \cap \sum_{l \neq j} \mathcal{W}_l = \{0\}$ for all j .

I.2.9. Let \mathcal{V} be finite dimensional. Prove that every subspace $\mathcal{W} \subset \mathcal{V}$ is finite dimensional, and that $\dim \mathcal{W} \leq \dim \mathcal{V}$ with equality only if $\mathcal{W} = \mathcal{V}$.

I.2.10. If \mathcal{V} is finite dimensional, every subspace $\mathcal{W} \subset \mathcal{V}$ is a direct summand.

***I.2.11.** Assume that \mathcal{V} is n -dimensional vector space over an infinite \mathbb{F} . Let $\{\mathcal{W}_j\}$ be a finite collection of distinct m -dimensional subspaces.

a. Prove that no \mathcal{W}_j is contained in the union of the others.

b. Prove that there is a subspace $\mathcal{U} \subset \mathcal{V}$ which is a complement of every \mathcal{W}_j .

Hint: See exercise **I.1.6**.

I.2.12. Let \mathcal{V} and \mathcal{W} be finite dimensional subspaces of a vector space. Prove that $\mathcal{V} + \mathcal{W}$ and $\mathcal{V} \cap \mathcal{W}$ are finite dimensional and that

$$(1.2.1) \quad \dim(\mathcal{V} \cap \mathcal{W}) + \dim(\mathcal{V} + \mathcal{W}) = \dim \mathcal{V} + \dim \mathcal{W}.$$

I.2.13. If $\mathcal{W}_j, j = 1, \dots, k$, are finite dimensional subspaces of a vector space \mathcal{V} then $\sum \mathcal{W}_j$ is finite dimensional and $\dim \sum \mathcal{W}_j \leq \sum \dim \mathcal{W}_j$, with equality if, and only if, the subspaces \mathcal{W}_j are independent.

I.2.14. Let \mathcal{V} be an n -dimensional vector space, and let $\mathcal{V}_1 \subset \mathcal{V}$ be a subspace of dimension m .

a. Prove that $\mathcal{V}/\mathcal{V}_1$ —the quotient space—is finite dimensional.

b. Let $\{v_1, \dots, v_m\}$ be a basis for \mathcal{V}_1 and let $\{\tilde{w}_1, \dots, \tilde{w}_k\}$ be a basis for $\mathcal{V}/\mathcal{V}_1$. For $j \in [1, k]$, let w_j be an element of the coset \tilde{w}_j .

Prove: $\{v_1, \dots, v_m\} \cup \{w_1, \dots, w_k\}$ is a basis for \mathcal{V} . Hence $k + m = n$.

I.2.15. Let \mathcal{V} be a real vector space. Let $r_l = (a_{l,1}, \dots, a_{l,p}) \in \mathbb{R}^p$, $1 \leq l \leq s$ be linearly independent. Let $v_1, \dots, v_p \in \mathcal{V}$ be linearly independent. Prove that the vectors $u_l = \sum_1^p a_{l,j} v_j$, $l = 1, \dots, s$, are linearly independent in \mathcal{V} .

I.2.16. Let \mathcal{V} and \mathcal{U} be finite dimensional spaces over \mathbb{F} . Prove that the tensor product $\mathcal{V} \otimes \mathcal{U}$ is finite dimensional. Specifically, show that if $\{e_j\}_{j=1}^n$ and $\{f_k\}_{k=1}^m$ are bases for \mathcal{V} and \mathcal{U} , then $\{e_j \otimes f_k\}$, $1 \leq j \leq n$, $1 \leq k \leq m$, is a basis for $\mathcal{V} \otimes \mathcal{U}$, so that $\dim \mathcal{V} \otimes \mathcal{U} = \dim \mathcal{V} \dim \mathcal{U}$.

***I.2.17.** Assume that any three of the five \mathbb{R}^3 vectors $v_j = (x_j, y_j, z_j)$, $j = 1, \dots, 5$, are linearly independent. Prove that the vectors

$$w_j = (x_j^2, y_j^2, z_j^2, x_j y_j, x_j z_j, y_j z_j)$$

are linearly independent in \mathbb{R}^6 .

Hint: Find non-zero (a, b, c) such that $ax_j + by_j + cz_j = 0$ for $j = 1, 2$. Find non-zero (d, e, f) such that $dx_j + ey_j + fz_j = 0$ for $j = 3, 4$. Observe (and use) the fact

$$(ax_5 + by_5 + cz_5)(dx_5 + ey_5 + fz_5) \neq 0$$

1.3 SYSTEMS OF LINEAR EQUATIONS.

How do we find out if a set $\{v_j\}$, $j = 1, \dots, m$ of vectors in \mathbb{F}_c^n is linearly dependent? How do we find out if a vector u belongs to $\text{span}[v_1, \dots, v_m]$?

Given the vectors $v_j = \begin{bmatrix} a_{1j} \\ \vdots \\ a_{nj} \end{bmatrix}$, $j = 1, \dots, m$, and $u = \begin{bmatrix} c_1 \\ \vdots \\ c_n \end{bmatrix}$, we express the conditions $\sum x_j v_j = 0$ for the first question, and $\sum x_j v_j = u$ for the second, in terms of the coordinates.

For the first we obtain the *system of homogeneous linear equations*:

$$(1.3.1) \quad \begin{array}{r} a_{11}x_1 + \dots + a_{1m}x_m = 0 \\ a_{21}x_1 + \dots + a_{2m}x_m = 0 \\ \vdots \\ a_{n1}x_1 + \dots + a_{nm}x_m = 0 \end{array}$$

or,

$$(1.3.2) \quad \sum_{j=1}^m a_{ij}x_j = 0, \quad i = 1, \dots, n.$$

For the second question we obtain the *non-homogeneous system*:

$$(1.3.3) \quad \sum_{j=1}^m a_{ij}x_j = c_i, \quad i = 1, \dots, n.$$

We need to determine if the solution-set of (1.3.2), namely the set of all m -tuples $(x_1, \dots, x_m) \in \mathbb{F}^m$ for which all n equations hold, is *trivial* or not, i.e., if there are solutions other than $(0, \dots, 0)$. For (1.3.3) we need to know if the solution-set is empty or not. In both cases we would like to identify the solution set as completely and as explicitly as possible.

1.3.1 Conversely, given the system (1.3.2) we can rewrite it as

$$(1.3.4) \quad x_1 \begin{bmatrix} a_{11} \\ \vdots \\ a_{n1} \end{bmatrix} + \dots + x_m \begin{bmatrix} a_{1m} \\ \vdots \\ a_{nm} \end{bmatrix} = 0$$

Our first result depends only on dimension. The m vectors in (1.3.4) are elements of the n -dimensional space \mathbb{F}_c^n . If $m > n$, any m vectors in \mathbb{F}_c^n are dependent, and since we have a nontrivial solution if, and only if, these columns are dependent, the system has nontrivial solution. This proves the following theorem.

Theorem. *A system of n homogeneous linear equations in $m > n$ unknowns has nontrivial solutions.*

Similarly, rewriting (1.3.3) in the form

$$(1.3.5) \quad x_1 \begin{bmatrix} a_{11} \\ \vdots \\ a_{n1} \end{bmatrix} + \cdots + x_m \begin{bmatrix} a_{1m} \\ \vdots \\ a_{nm} \end{bmatrix} = \begin{bmatrix} c_1 \\ \vdots \\ c_n \end{bmatrix},$$

it is clear that the system given by (1.3.3) has a solution if, and only if, the

$$\text{column } \begin{bmatrix} c_1 \\ \vdots \\ c_n \end{bmatrix} \text{ is in the span of columns } \begin{bmatrix} a_{1j} \\ \vdots \\ a_{nj} \end{bmatrix}, j \in [1, m].$$

1.3.2 The classical approach to solving systems of linear equations is the *Gaussian elimination*— an algorithm for replacing the given system by an *equivalent* system that can be solved easily. We need some terminology:

DEFINITION: The systems

$$(1.3.6) \quad \begin{aligned} (\mathfrak{A}) \quad & \sum_{j=1}^m a_{ij}x_j = c_i, \quad i = 1, \dots, k. \\ (\mathfrak{B}) \quad & \sum_{j=1}^m b_{ij}x_j = d_i, \quad i = 1, \dots, l. \end{aligned}$$

are *equivalent* if they have the same solution-set (in \mathbb{F}^m).

The matrices

$$A = \begin{bmatrix} a_{11} & \cdots & a_{1m} \\ a_{21} & \cdots & a_{2m} \\ \vdots & \cdots & \vdots \\ a_{k1} & \cdots & a_{km} \end{bmatrix} \quad \text{and} \quad A_{aug} = \begin{bmatrix} a_{11} & \cdots & a_{1m} & c_1 \\ a_{21} & \cdots & a_{2m} & c_2 \\ \vdots & \cdots & \vdots & \vdots \\ a_{k1} & \cdots & a_{km} & c_k \end{bmatrix}$$

are called *the matrix* and the *augmented matrix* of the system (\mathfrak{A}) . The augmented matrix is obtained from the matrix by adding, as additional column, the column of the *values*, that is, the right-hand side of the respective equations. The augmented matrix contains all the information of the system (\mathfrak{A}) . Any $k \times (m + 1)$ matrix is the augmented matrix of a system of linear equations in m unknowns.

1.3.3 ROW EQUIVALENCE OF MATRICES.

DEFINITION: The matrices

$$(1.3.7) \quad \begin{bmatrix} a_{11} & \cdots & a_{1m} \\ a_{21} & \cdots & a_{2m} \\ \vdots & \cdots & \vdots \\ a_{k1} & \cdots & a_{km} \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} b_{11} & \cdots & b_{1m} \\ b_{21} & \cdots & b_{2m} \\ \vdots & \cdots & \vdots \\ b_{l1} & \cdots & b_{lm} \end{bmatrix}$$

are *row equivalent* if their rows span the same subspace of \mathbb{F}_R^m ; equivalently: if each row of either matrix is a linear combination of the rows of the other.

Proposition. *Two systems of linear equations in m unknowns*

$$(\mathfrak{A}) \quad \sum_{j=1}^m a_{ij}x_j = c_i, \quad i = 1, \dots, k.$$

$$(\mathfrak{B}) \quad \sum_{j=1}^m b_{ij}x_j = d_i, \quad i = 1, \dots, l.$$

are equivalent if their respective augmented matrices are row equivalent.

PROOF: Assume that the augmented matrices are row equivalent.

If (x_1, \dots, x_m) is a solution for system (\mathfrak{A}) and

$$(b_{i1}, \dots, b_{im}, d_i) = \sum \alpha_{i,k}(a_{k1}, \dots, a_{km}, c_k)$$

then

$$\sum_{j=1}^m b_{ij}x_j = \sum_{k,j} \alpha_{i,k}a_{kj}x_j = \sum_k \alpha_{i,k}c_k = d_i$$

and (x_1, \dots, x_m) is a solution for system (\mathfrak{B}) . ◀

DEFINITION: The *row rank* of a matrix $A \in \mathcal{M}(k, m)$ is the dimension of the span of its rows in \mathbb{F}^m .

Row equivalent matrices clearly have the same rank.

1.3.4 REDUCTION TO row echelon FORM. The classical method of solving systems of linear equations, homogeneous or not, is the Gaussian elimination. It is an algorithm to replace the system at hand by an equivalent system that is easier to solve.

DEFINITION: A matrix $A = \begin{bmatrix} a_{11} & \cdots & a_{1m} \\ a_{21} & \cdots & a_{2m} \\ \vdots & \cdots & \vdots \\ a_{k1} & \cdots & a_{km} \end{bmatrix}$ is in *row echelon form* if the following conditions are satisfied

ref-1 The first q rows of A are linearly independent in \mathbb{F}^m , the remaining $k - q$ rows are zero.

ref-2 There are integers $1 \leq l_1 < l_2 < \cdots < l_q \leq m$ such that for $j \leq q$, the first nonzero entry in the j 'th row is 1, occurring in the l_j 'th column.

ref-3 The entry 1 in row j is the only nonzero entry in the l_j column.

One can rephrase the last three conditions as: The l_j 'th columns (the "main" columns) are the first q elements of the standard basis of \mathbb{F}_c^k , and every other column is a linear combination of the "main" columns that precede it.

Theorem. *Every matrix is row equivalent to a matrix in row-echelon form.*

PROOF: If $A = 0$ there's nothing to prove. Assuming $A \neq 0$, we describe an algorithm to reduce A to row-echelon form. The operations performed on the matrix are:

- Reordering (i.e., permuting) the rows,
- Multiplying a row by a non-zero constant,
- Adding a multiple of one row to another.

These operations do not change the span of the rows so that the equivalence class of the matrix is maintained. (We shall return later, in Exercise **II.3.10**, to express these operations as matrix multiplications.)

Let l_1 be the index of the first column that is not zero.

Reorder the rows so that $a_{1,l_1} \neq 0$, and multiply the first row by a_{1,l_1}^{-1} .

Subtract from the j 'th row, $j \neq 1$, the first row multiplied by a_{j,l_1} .

Now all the columns before l_1 are zero and column l_1 has 1 in the first row, and zero elsewhere.

Denote its row rank of A by q . If $q = 1$ all the entries below the first row are now zero and we are done. Otherwise let l_2 be the index of the first column that has a nonzero entry in a row beyond the first. Notice that $l_2 > l_1$. Keep the first row in its place, reorder the remaining rows so that $a_{2,l_2} \neq 0$, and multiply the second row⁵ by a_{2,l_2}^{-1} .

Subtract from the j 'th row, $j \neq 2$, the second row multiplied by a_{j,l_2} .

Repeat the sequence a total of q times. The first q rows, $\mathbf{r}_1, \dots, \mathbf{r}_q$, are (now) independent: a combination $\sum c_j \mathbf{r}_j$ has entry c_j in the l_j 'th place, and can be zero only if $c_j = 0$ for all j .

If there is a nonzero entry beyond the current q 'th row, necessarily beyond the l_q 'th column, we could continue and get a row independent of the first q , contradicting the definition of q . Thus, after q steps, all the rows beyond the q 'th are zero. ◀

Observe that the scalars used in the process belong to the smallest field that contains all the coefficients of A .

1.3.5 If A and A_{aug} are the matrix and the augmented matrix of a system (\mathfrak{A}) and we apply the algorithm of the previous subsection to both, we observe that since the augmented matrix has the additional column on the right hand side, the first q (the row rank of A) steps in the algorithm for either A or A_{aug} are identical. Having done q repetitions, A is reduced to row-echelon form, while A_{aug} may or may not be. If the row rank of A_{aug} is q , then the algorithm for A_{aug} ends as well; otherwise we have $l_{q+1} = m + 1$, and the row-echelon form for the augmented matrix is the same as that of A but with an added row and an added “main” column, both having 0 for all but the last entries, and 1 for the last entry. In the latter case, the system corresponding to the row-reduced augmented matrix has as its last equation $0 = 1$ and the system has no solutions.

On the other hand, if the row rank of the augmented matrix is the same as that of A , the row-echelon form of the augmented matrix is an augmentation of

⁵We keep referring to the entries of the successively modified matrix as a_{ij} .

the row-echelon form of A . In this case we can assign arbitrary values to the variables x_i , $i \neq l_j$, $j = 1, \dots, q$, move the corresponding terms to the right hand side and, writing C_j for the sum, we obtain

$$(1.3.8) \quad x_{l_j} = C_j, \quad j = 1, \dots, q.$$

Theorem. *A necessary and sufficient condition for the system (2) to have solutions is that the row rank of the augmented matrix be equal to that of the matrix of the system.*

The discussion preceding the statement of the theorem not only proves the theorem but offers a concrete way to solve the system. The unknowns are now split into two groups, q “main” ones and $m - q$ “secondary”. We have “ $m - q$ degrees of freedom”: the $m - q$ secondary unknowns become free parameters that can be assigned arbitrary values, and these values determine the “main” unknowns uniquely.

Remark: Notice that the split into “main” and “secondary” unknowns depends on the specific definition of “row-echelon form”; counting the columns in a different order may result in a different split, though the number q of “main” variables would be the same—the row rank of A .

Corollary. *A linear system of n equations in n unknowns with matrix A has solutions for all augmented matrices if, and only if, the only solution of the corresponding homogeneous system is the trivial solution.*

PROOF: The condition on the homogeneous system amounts to “the rows of A are independent”, and no added columns can increase the row rank. ◀

1.3.6 DEFINITION: The *column rank* of a matrix $A \in \mathcal{M}(k, m)$ is the dimension of the span of its columns in \mathbb{F}_c^k .

Linear relations between columns of A are solutions of the homogeneous system given by A . If B is row-equivalent to A , the columns of A and B have the same set of linear relations, (see Proposition 1.3.3). In particular, if B is in row-echelon form and $\{l_j\}_{j=1}^q$ are the indices of the “main” columns in B , then the l_j ’th columns in A , $j = 1, \dots, q$, are independent, and every other column is a linear combination of these.

It follows that *the column rank of A is equal to its row rank*. We shall refer to the common value simply as *the rank of A* .

EXERCISES FOR SECTION 1.3

I.3.1. Identify the matrix $A \in \mathcal{M}(n)$ of row rank n that is in row echelon form.

I.3.2. A system of linear equations with rational coefficients, that has a solution in \mathbb{C} , has a solution in \mathbb{Q} . Equivalently, vectors in \mathbb{Q}^n that are linearly dependent over \mathbb{C} , are rationally dependent.

Hint: The last sentence of Subsection 1.3.4.

I.3.3. A system of linear equations with rational coefficients, has the same number of “degrees of freedom” over \mathbb{Q} as it does over \mathbb{C} .

I.3.4. An *affine subspace* of a vector space is a translate of a subspace, that is a set of the form $v_0 + \mathcal{V}_0 = \{v_0 + v : v \in \mathcal{V}_0\}$, where v_0 is a fixed vector and $\mathcal{V}_0 \subset \mathcal{V}$ is a subspace. (Thus a *line* in \mathcal{V} is a translate of a one-dimensional subspace.)

Prove that a set $A \subset \mathcal{V}$ is an affine subspace if, and only if, $\sum a_j u_j \in A$ for all choices of $u_1, \dots, u_k \in A$, and scalars $a_j, j = 1, \dots, k$ such that $\sum a_j = 1$.

I.3.5. If $A \subset \mathcal{V}$ is an affine subspace and $u_0 \in A$, then $A - u_0 = \{u - u_0 : u \in A\}$ is a subspace of \mathcal{V} . Moreover, the subspace $A - u_0$, the “corresponding subspace” does not depend on the choice of u_0 .

I.3.6. The solution set of a system of k linear equations in m unknowns is an affine subspace of \mathbb{F}^m . The solution set of the corresponding homogeneous system is the “corresponding subspace”.

I.3.7. Consider the matrix $A = \begin{bmatrix} a_{11} & \dots & a_{1m} \\ a_{21} & \dots & a_{2m} \\ \vdots & \dots & \vdots \\ a_{k1} & \dots & a_{km} \end{bmatrix}$ and its columns $v_j = \begin{bmatrix} a_{1j} \\ \vdots \\ a_{kj} \end{bmatrix}$.

Prove that a column v_i end up as a “main column” in the *row echelon* form of A if, and only if, it is linearly independent of the columns $v_j, j < i$.

I.3.8. (continuation) Denote by $B = \begin{bmatrix} b_{11} & \dots & b_{1m} \\ b_{21} & \dots & b_{2m} \\ \vdots & \dots & \vdots \\ b_{k1} & \dots & b_{km} \end{bmatrix}$ the matrix in row echelon

form obtained from A by the algorithm described above. Let $l_1 < l_2, \dots$ be the indices

of the main columns in B and i the index of another column. Prove

$$(1.3.9) \quad v_i = \sum_{l_j < i} b_{ji} v_{l_j}.$$

I.3.9. What is the row echelon form of the 7×6 matrix A , if its columns C_j , $j = 1, \dots, 6$ satisfy the following conditions:

- a. $C_1 \neq 0$;
- b. $C_2 = 3C_1$;
- c. C_3 is not a (scalar) multiple of C_1 ;
- d. $C_4 = C_1 + 2C_2 + 3C_3$;
- e. $C_5 = 6C_3$;
- f. C_6 is not in the span of C_2 and C_3 .

I.3.10. Given polynomials $P_1 = \sum_0^n a_j x^j$, $P_2 = \sum_0^m b_j x^j$, $S = \sum_0^l s_j x^j$ of degrees n , m , and $l < n + m$ respectively, we want to find polynomials $q_1 = \sum_0^{m-1} c_j x^j$, and $q_2 = \sum_0^{n-1} d_j x^j$, such that

$$(1.3.10) \quad P_1 q_1 + P_2 q_2 = S.$$

Reduce the polynomial equation (1.3.10) to a system of linear equations, the unknown being the coefficients c_0, \dots, c_{m-1} of q_1 , and d_0, \dots, d_{n-1} of q_2 .

The associated homogeneous system corresponds to the case $S = 0$. Show that it has a nontrivial solutions if, and only if, P_1 and P_2 have a nontrivial common factor. (You may assume the unique factorization theorem, A.6.3.)

Chapter II

Linear operators and matrices

2.1 LINEAR OPERATORS (MAPS, TRANSFORMATIONS)

2.1.1 Let \mathcal{V} and \mathcal{W} be vector spaces over the same field.

DEFINITION: A map $T: \mathcal{V} \mapsto \mathcal{W}$ is *linear* if for all vectors $v_j \in \mathcal{V}$ and scalars a_j ,

$$(2.1.1) \quad T(a_1v_1 + a_2v_2) = a_1Tv_1 + a_2Tv_2.$$

This was discussed briefly in 1.1.2. Linear maps are also called linear operators, linear transformations, homomorphisms, etc. The adjective “linear” is sometimes assumed implicitly. The term we use most of the time is *operator*.

EXAMPLES:

a. If $\{v_1, \dots, v_n\}$ is a basis for \mathcal{V} and $\{w_1, \dots, w_n\} \subset \mathcal{W}$ is arbitrary, then the map $v_j \mapsto w_j, j = 1, \dots, n$ extends (uniquely) to a linear map T from \mathcal{V} to \mathcal{W} defined by

$$(2.1.2) \quad T: \sum a_jv_j \mapsto \sum a_jw_j.$$

Every linear operator from \mathcal{V} into \mathcal{W} is obtained this way.

b. Let \mathcal{V} be the space of all continuous, 2π -periodic functions on the line. For every x_0 define T_{x_0} , the *translation by* x_0 :

$$T_{x_0}: f(x) \mapsto f_{x_0}(x) = f(x - x_0).$$

c. The *transpose*.

$$(2.1.3) \quad A = \begin{bmatrix} a_{11} & \cdots & a_{1m} \\ a_{21} & \cdots & a_{2m} \\ \vdots & \cdots & \vdots \\ a_{n1} & \cdots & a_{nm} \end{bmatrix} \mapsto A^{\mathbb{T}} = \begin{bmatrix} a_{11} & \cdots & a_{n1} \\ a_{12} & \cdots & a_{n2} \\ \vdots & \cdots & \vdots \\ a_{1m} & \cdots & a_{nm} \end{bmatrix}$$

which maps $\mathcal{M}(n, m; \mathbb{F})$ onto $\mathcal{M}(m, n; \mathbb{F})$.

d. Differentiation on $\mathbb{F}[x]$:

$$(2.1.4) \quad D: \sum_0^n a_j x^j \mapsto \sum_1^n j a_j x^{j-1}.$$

There is no limiting process involved and the definition is valid for arbitrary field \mathbb{F} .

e. Differentiation on \mathcal{T}_N :

$$(2.1.5) \quad D: \sum_{-N}^N a_n e^{inx} \mapsto \sum_{-N}^N in a_n e^{inx}.$$

There is no limiting process involved.

f. Differentiation on $C^\infty[0, 1]$, the complex vector space of infinitely differentiable complex-valued functions on $[0, 1]$:

$$(2.1.6) \quad D: f \mapsto f' = \frac{df}{dx}.$$

g. If $\mathcal{V} = \mathcal{W} \oplus \mathcal{U}$ every $v \in \mathcal{V}$ has a unique representation $v = w + u$ with $w \in \mathcal{W}$, $u \in \mathcal{U}$. The map $\pi_1: v \mapsto w$ is the identity on \mathcal{W} and maps \mathcal{U} to $\{0\}$. It is called the *projection of \mathcal{V} on \mathcal{W} along \mathcal{U}* . The operator π_1 is linear since, if $v = w + u$ and $v_1 = w_1 + u_1$, then $av + bv_1 = (aw + bw_1) + (au + bu_1)$, and $\pi_1(av + bv_1) = a\pi_1 v + b\pi_1 v_1$.

Similarly, $\pi_2: v \mapsto u$ is called the *projection of \mathcal{V} on \mathcal{U} along \mathcal{W}* . π_1 and π_2 are referred to as the *projections corresponding to the direct sum decomposition*.

2.1.2 We denote the space of all linear maps from \mathcal{V} into \mathcal{W} by $\mathcal{L}(\mathcal{V}, \mathcal{W})$. Another common notation is $HOM(\mathcal{V}, \mathcal{W})$. The two most important cases in what follows are: $\mathcal{W} = \mathcal{V}$, and $\mathcal{W} = \mathbb{F}$, the field of scalars.

When $\mathcal{W} = \mathcal{V}$ we write $\mathcal{L}(\mathcal{V})$ instead of $\mathcal{L}(\mathcal{V}, \mathcal{V})$.

When \mathcal{W} is the underlying field, we refer to the linear maps as *linear functionals* or *linear forms* on \mathcal{V} . Instead of $\mathcal{L}(\mathcal{V}, \mathbb{F})$ we write \mathcal{V}^* , and refer to it as *the dual space* of \mathcal{V} .

2.1.3 If $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$ is bijective, it is invertible, and the inverse map T^{-1} is linear from \mathcal{W} onto \mathcal{V} . This is seen as follows: by (2.1.1),

$$(2.1.7) \quad \begin{aligned} T^{-1}(a_1Tv_1 + a_2Tv_2) &= T^{-1}(T(a_1v_1 + a_2v_2)) = a_1v_1 + a_2v_2 \\ &= a_1T^{-1}(Tv_1) + a_2T^{-1}(Tv_2), \end{aligned}$$

and, as T is surjective, Tv_j are arbitrary vectors in \mathcal{W} .

Recall (see 1.1.2) that an *isomorphism* of vector spaces, \mathcal{V} and \mathcal{W} is a bijective linear map $T: \mathcal{V} \mapsto \mathcal{W}$. An isomorphism of a space onto itself is called an *automorphism*.

\mathcal{V} and \mathcal{W} are *isomorphic* if there is an isomorphism of the one onto the other. The relation is clearly reflexive and, by the previous paragraph, symmetric. Since the concatenation (see 2.2.1) of isomorphisms is an isomorphism, the relation is also transitive and so is an equivalence relation. The image of a basis under an isomorphism is a basis, see exercise **II.1.2**; it follows that the dimension is an isomorphism invariant.

If \mathcal{V} is a finite dimensional vector space over \mathbb{F} , every basis $\mathbf{v} = \{v_1, \dots, v_n\}$ of \mathcal{V} defines an isomorphism $\mathbf{C}_{\mathbf{v}}$ of \mathcal{V} onto \mathbb{F}^n by:

$$(2.1.8) \quad \mathbf{C}_{\mathbf{v}}: v = \sum a_j v_j \mapsto \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} = \sum a_j e_j.$$

$\mathbf{C}_{\mathbf{v}} v$ is the *coordinate vector* of v relative to the basis \mathbf{v} . Notice that this is a special case of example **a.** above: we map the basis elements v_j on the corresponding elements e_j of the standard basis, and extend by linearity.

If \mathcal{V} and \mathcal{W} are both n -dimensional, with bases $\mathbf{v} = \{v_1, \dots, v_n\}$, and $\mathbf{w} = \{w_1, \dots, w_n\}$ respectively, the map $T: \sum a_j v_j \mapsto \sum a_j w_j$ is an isomor-

phism. This shows that the dimension is a *complete invariant*: finite dimensional vector spaces over \mathbb{F} are isomorphic if, and only if, they have the same dimension.

2.1.4 The sum of linear maps $T, S \in \mathcal{L}(\mathcal{V}, \mathcal{W})$, and the multiple of a linear map by a scalar are defined by: for every $v \in \mathcal{V}$,

$$(2.1.9) \quad (T + S)v = Tv + Sv, \quad (aT)v = a(Tv).$$

Observe that $(T + S)$ and aT , as defined, are linear maps from \mathcal{V} to \mathcal{W} , i.e., elements of $\mathcal{L}(\mathcal{V}, \mathcal{W})$.

Proposition. *Let \mathcal{V} and \mathcal{W} be vector spaces over \mathbb{F} . Then, with the addition and multiplication by a scalar defined by (2.1.9), $\mathcal{L}(\mathcal{V}, \mathcal{W})$ is a vector space defined over \mathbb{F} . If both \mathcal{V} and \mathcal{W} are finite dimensional, then so is $\mathcal{L}(\mathcal{V}, \mathcal{W})$, and $\dim \mathcal{L}(\mathcal{V}, \mathcal{W}) = \dim \mathcal{V} \dim \mathcal{W}$.*

PROOF: The proof that $\mathcal{L}(\mathcal{V}, \mathcal{W})$ is a vector space over \mathbb{F} is straightforward checking, left to the reader.

The statement about the dimension is exercise **II.1.3** below. ◀

EXERCISES FOR SECTION 2.1

II.1.1. Show that if A is linearly dependent in \mathcal{V} and $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$, then TA is linearly dependent in \mathcal{W} .

II.1.2. Prove that an injective map $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$ is an isomorphism if, and only if, it maps some basis of \mathcal{V} onto a basis of \mathcal{W} , and this is the case if, and only if, it maps every basis of \mathcal{V} onto a basis of \mathcal{W} .

II.1.3. Let \mathcal{V} and \mathcal{W} be finite dimensional with bases $\mathbf{v} = \{v_1, \dots, v_n\}$ and $\mathbf{w} = \{w_1, \dots, w_m\}$ respectively. Let $\varphi_{ij} \in \mathcal{L}(\mathcal{V}, \mathcal{W})$ be defined by $\varphi_{ij}v_i = w_j$ and $\varphi_{ij}v_k = 0$ for $k \neq i$. Prove that $\{\varphi_{ij} : 1 \leq i \leq n, 1 \leq j \leq m\}$ is a basis for $\mathcal{L}(\mathcal{V}, \mathcal{W})$.

2.2 OPERATOR MULTIPLICATION

2.2.1 For $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$ and $S \in \mathcal{L}(\mathcal{W}, \mathcal{U})$ we define $ST \in \mathcal{L}(\mathcal{V}, \mathcal{U})$ by *concatenation*, that is: $(ST)v = S(Tv)$. ST is a linear operator since

$$(2.2.1) \quad ST(a_1v_1 + a_2v_2) = S(a_1Tv_1 + a_2Tv_2) = a_1STv_1 + a_2STv_2.$$

In particular, if $\mathcal{V} = \mathcal{W} = \mathcal{U}$, we have T , S , and TS all in $\mathcal{L}(\mathcal{V})$.

Proposition. *With the product ST defined above, $\mathcal{L}(\mathcal{V})$ is an algebra over \mathbb{F} .*

PROOF: The claim is that the product is associative and, with the addition defined by (2.1.9) above, distributive. This is straightforward checking, left to the reader. ◀

The algebra $\mathcal{L}(\mathcal{V})$ is *not commutative* unless $\dim \mathcal{V} = 1$, in which case it is simply the underlying field.

The set of automorphisms, i.e., invertible elements in $\mathcal{L}(\mathcal{V})$ is a *group* under multiplication, denoted $\mathbf{GL}(\mathcal{V})$.

2.2.2 Given an operator $T \in \mathcal{L}(\mathcal{V})$ the powers T^j of T are well defined for all $j \geq 1$, and we define $T^0 = I$. Since we can take linear combinations of the powers of T we have $P(T)$ well defined for all polynomials $P \in \mathbb{F}[x]$.

We denote

$$(2.2.2) \quad \mathcal{P}(T) = \{P(T) : P \in \mathbb{F}[x]\}.$$

$\mathcal{P}(T)$ will be the main tool in understanding the way in which T acts on \mathcal{V} .

EXERCISES FOR SECTION 2.2

II.2.1. Prove that $\mathcal{P}(T)$ is a commutative subalgebra of $\mathcal{L}(\mathcal{V})$.

II.2.2. For $T \in \mathcal{L}(\mathcal{V})$ denote $\text{comm}[T] = \{S : S \in \mathcal{L}(\mathcal{V}), ST = TS\}$, the set of operators that commute with T . Prove that $\text{comm}[T]$ is a subalgebra of $\mathcal{L}(\mathcal{V})$.

II.2.3. Verify that $\mathbf{GL}(\mathcal{V})$ is in fact a group.

2.3 MATRIX MULTIPLICATION.

2.3.1 We define the product of a $1 \times n$ matrix (row) $\mathbf{r} = (a_1, \dots, a_n)$ and an $n \times 1$ matrix (column) $\mathbf{c} = \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix}$, to be the scalar given by

$$(2.3.1) \quad \mathbf{r} \cdot \mathbf{c} = \sum a_j b_j,$$

Given $A \in \mathcal{M}(l, m)$ and $B \in \mathcal{M}(m, n)$, we define *the product* AB as the $l \times n$ matrix C whose entries c_{ij} are given by

$$(2.3.2) \quad c_{ij} = \mathbf{r}_i(A) \cdot \mathbf{c}_j(B) = \sum_k a_{ik} b_{kj}$$

($\mathbf{r}_i(A)$ denotes the i 'th row in A , and $\mathbf{c}_j(B)$ denotes the j 'th column in B).

Notice that the product is defined only when the number of columns in A (the length of the row) is the same as the number of rows in B , (the height of the column).

The product is associative: given $A \in \mathcal{M}(l, m)$, $B \in \mathcal{M}(m, n)$, and $C \in \mathcal{M}(n, p)$, then $AB \in \mathcal{M}(l, n)$ and $(AB)C \in \mathcal{M}(l, p)$ is well defined. Similarly, $A(BC)$ is well defined and one checks that $A(BC) = (AB)C$ by verifying that the r, s entry in either is $\sum_{i,j} a_{rj} b_{ji} c_{is}$.

The product is distributive: for $A_j \in \mathcal{M}(l, m)$, $B_j \in \mathcal{M}(m, n)$,

$$(2.3.3) \quad (A_1 + A_2)(B_1 + B_2) = A_1 B_1 + A_1 B_2 + A_2 B_1 + A_2 B_2,$$

and commutes with multiplication by scalars: $A(aB) = aAB$.

Proposition. *The map $(A, B) \mapsto AB$, of $\mathcal{M}(l, m) \times \mathcal{M}(m, n)$ to $\mathcal{M}(l, n)$, is linear in B for every fixed A , and in A for every fixed B .*

PROOF: The statement just summarizes the properties of the multiplication discussed above. ◀

2.3.2 Write the $n \times m$ matrix $(a_{ij})_{\substack{1 \leq i \leq n \\ 1 \leq j \leq m}}$ as a “single column of rows”,

$$\begin{bmatrix} a_{11} & \cdots & a_{1m} \\ a_{21} & \cdots & a_{2m} \\ \vdots & \cdots & \vdots \\ a_{n1} & \cdots & a_{nm} \end{bmatrix} = \begin{bmatrix} (a_{11} \ \cdots \ a_{1m}) \\ (a_{21} \ \cdots \ a_{2m}) \\ \vdots \\ (a_{n1} \ \cdots \ a_{nm}) \end{bmatrix} = \begin{bmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \\ \vdots \\ \mathbf{r}_n \end{bmatrix}$$

where $\mathbf{r}_i = (a_{i,1} \ \cdots \ a_{i,m}) \in \mathbb{F}_r^m$. Notice that if $(x_1, \dots, x_n) \in \mathbb{F}_r^n$, then

$$(2.3.4) \quad (x_1, \dots, x_n) \begin{bmatrix} a_{11} & \cdots & a_{1m} \\ a_{21} & \cdots & a_{2m} \\ \vdots & \cdots & \vdots \\ a_{n1} & \cdots & a_{nm} \end{bmatrix} = (x_1, \dots, x_n) \begin{bmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \\ \vdots \\ \mathbf{r}_n \end{bmatrix} = \sum_{i=1}^n x_i \mathbf{r}_i.$$

Similarly, writing the matrix as a “single row of columns”,

$$\begin{bmatrix} a_{11} & \cdots & a_{1m} \\ a_{21} & \cdots & a_{2m} \\ \vdots & \cdots & \vdots \\ a_{n1} & \cdots & a_{nm} \end{bmatrix} = \left(\begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix} \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix} \cdots \begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix} \right) = (\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_m)$$

we have

$$(2.3.5) \quad \begin{bmatrix} a_{11} & \cdots & a_{1m} \\ a_{21} & \cdots & a_{2m} \\ \vdots & \cdots & \vdots \\ a_{n1} & \cdots & a_{nm} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{bmatrix} = (\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_m) \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{bmatrix} = \sum_{j=1}^m y_j \mathbf{c}_j.$$

2.3.3 If $l = m = n$ matrix multiplication is a product within $\mathcal{M}(n)$.

Proposition. *With the multiplication defined above, $\mathcal{M}(n)$ is an algebra over \mathbb{F} . The matrix $I = I_n = (\delta_{j,k}) = \sum_1^n e_{ii}$ is the identity¹ element in $\mathcal{M}(n)$.*

The invertible elements in $\mathcal{M}(n)$, aka the *non-singular* matrices, form a group under multiplication, the *general linear group* $\mathbf{GL}(n, \mathbb{F})$.

Theorem. *A matrix $A \in \mathcal{M}(n)$ is invertible if, and only if its rank is n .*

¹ $\delta_{j,k}$ is the Kronecker delta, equal to 1 if $j = k$, and to 0 otherwise.

PROOF: Exercise II.3.2 below (or equation (2.3.4)) give that the row rank of BA is no bigger than the row rank of A . If $BA = I$, the row rank of A is at least the row rank of I , which is clearly n .

On the other hand, if A is row equivalent to I , then its row echelon form is I , and by Exercise II.3.10 below, reduction to row echelon form amounts to multiplication on the left by a matrix, so that A has a left inverse. This implies, see Exercise II.3.12, that A is invertible. ◀

EXERCISES FOR SECTION 2.3

II.3.1. Let \mathbf{r} be the $1 \times n$ matrix all whose entries are 1, and \mathbf{c} the $n \times 1$ matrix all whose entries are 1. Compute \mathbf{rc} and \mathbf{cr} .

II.3.2. Prove that each of the columns of the matrix AB is a linear combinations of the columns of A , and that each row of AB is a linear combination of the rows of B .

II.3.3. Prove: If A is a diagonal matrix with distinct entries on the diagonal, and if B is a matrix such that $AB = BA$, then B is diagonal.

II.3.4. Denote by $\Xi(n; i, j)$, $1 \leq i, j \leq n$, the $n \times n$ matrix $\sum_{k \neq i, j} e_{kk} + e_{ij} + e_{ji}$ (the entries ξ_{lk} are all zero except for $\xi_{ij} = \xi_{ji} = 1$, and $\xi_{kk} = 1$ if $k \neq i, j$). This is the matrix obtained from the identity by interchanging rows i and j .

Let $A \in \mathcal{M}(n, m)$ and $B \in \mathcal{M}(m, n)$. Describe $\Xi(n; i, j)A$ and $B\Xi(n; i, j)$.

II.3.5. Let σ be a permutation of $[1, \dots, n]$. Let A_σ be the $n \times n$ matrix whose entries a_{ij} are defined by

$$(2.3.6) \quad a_{ij} = \begin{cases} 1 & \text{if } i = \sigma(j) \\ 0 & \text{otherwise.} \end{cases}$$

Let $B \in \mathcal{M}(n, m)$ and $C \in \mathcal{M}(m, n)$. Describe $A_\sigma B$ and CA_σ .

II.3.6. A matrix whose entries are either zero or one, with precisely one non-zero entry in each row and in each column is called a *permutation matrix*. Show that the matrix A_σ described in the previous exercise is a permutation matrix and that every permutation matrix is equal to A_σ for some $\sigma \in \mathbf{S}_n$.

II.3.7. Show that the map $\sigma \mapsto A_\sigma$ defined above is multiplicative: $A_{\sigma\tau} = A_\sigma A_\tau$. ($\sigma\tau$ is defined by concatenation: $\sigma\tau(j) = \sigma(\tau(j))$ for all $j \in [1, n]$.)

II.3.8. Denote by e_{ij} , $1 \leq i, j \leq n$, the $n \times n$ matrix whose entries are all zero except for the ij entry which is 1. With $A \in \mathcal{M}(n, m)$ and $B \in \mathcal{M}(m, n)$. Describe $e_{ij}A$ and Be_{ij} .

II.3.9. Describe an $n \times n$ matrix $A(c, i, j)$ such that multiplying on the appropriate side, an $n \times n$ matrix B by it, has the effect of replacing the i 'th row in B by the sum of the i 'th row and c times the j 'th row. Do the same for columns.

II.3.10. Show that each of the steps in the reduction of a matrix A to its row-echelon form (see 1.3.4) can be accomplished by left multiplication of A by an appropriate matrix, so that the entire reduction to row-echelon form can be accomplished by left multiplication by an appropriate matrix. Conclude that if the row rank of $A \in \mathcal{M}(n)$ is n , then A is left-invertible.

II.3.11. Let $A \in \mathcal{M}(n)$ be non-singular and let $B = (A, I)$, the matrix obtained by “augmenting” A by the identity matrix, that is by adding to A the columns of I in their given order as columns $n + 1, \dots, 2n$. Show that the matrix obtained by reducing B to row echelon form is (I, A^{-1}) .

II.3.12. Prove that if $A \in \mathcal{M}(n, m)$ and $B \in \mathcal{M}(m, l)$ then $(AB)^T = B^T A^T$. Show that if $A \in \mathcal{M}(n)$ has a left inverse then A^T has a right inverse and if A has a right inverse then A^T has a left inverse. Use the fact that A and A^T have the same rank to show that if A has a left inverse B it also has a right inverse C and since $B = B(AC) = (BA)C = C$, we have $BA = AB = I$ and A has an inverse.

Where does the fact that we deal with finite dimensional spaces enter the proof?

II.3.13. What are the ranks and the inverses (when they exist) of the matrices

$$(2.3.7) \quad \begin{bmatrix} 0 & 2 & 1 & 0 \\ 1 & 1 & 7 & 1 \\ 2 & 2 & 2 & 2 \\ 0 & 5 & 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 2 & 2 & 1 & 1 \\ 2 & 1 & 2 & 1 & 2 \\ 0 & 5 & 0 & 9 & 1 \\ 0 & 5 & 0 & 0 & 7 \end{bmatrix}, \quad \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

II.3.14. Denote $A_n = \begin{bmatrix} 1 & n \\ 0 & 1 \end{bmatrix}$. Prove that $A_m A_n = A_{m+n}$ for all integers m, n .

2.4 MATRICES AND OPERATORS.

2.4.1 Recall that we write the elements of \mathbb{F}^n as columns. A matrix A in $\mathcal{M}(m, n)$ defines, by multiplication on the left, an operator T_A from \mathbb{F}^n to \mathbb{F}^m .

The columns of A are the images, under T_A , of the standard basis vectors of \mathbb{F}^n (see (2.3.5)).

Conversly, given $T \in \mathcal{L}(\mathbb{F}^n, \mathbb{F}^m)$, if we take $A = A_T$ to be the $m \times n$ matrix whose columns are Te_j , where $\{e_1, \dots, e_n\}$ is the standard basis in \mathbb{F}^n , we have $T_A = T$.

Finally we observe that by Proposition 2.3.1 the map $A \mapsto T_A$ is linear. This proves:

Theorem. *There is a 1-1 linear correspondence $T \leftrightarrow A_T$ between $\mathcal{L}(\mathbb{F}^n, \mathbb{F}^m)$ and $\mathcal{M}(m, n)$ such that $T \in \mathcal{L}(\mathbb{F}^n, \mathbb{F}^m)$ is obtained as a left multiplication by the $m \times n$ matrix, A_T .*

2.4.2 If $T \in \mathcal{L}(\mathbb{F}^n, \mathbb{F}^m)$ and $S \in \mathcal{L}(\mathbb{F}^m, \mathbb{F}^l)$ and $A_T \in \mathcal{M}(m, n)$, resp. $A_S \in \mathcal{M}(l, m)$ are the corresponding matrices, then

$$ST \in \mathcal{L}(\mathbb{F}^n, \mathbb{F}^l), \quad A_S A_T \in \mathcal{M}(l, n), \quad \text{and} \quad A_{ST} = A_S A_T.$$

In particular, if $n = m = l$, we obtain

Theorem. *The map $T \leftrightarrow A_T$ is an algebra isomorphism between $\mathcal{L}(\mathbb{F}^n)$ and $\mathcal{M}(n)$.*

2.4.3 The special thing about \mathbb{F}^n is that it has a “standard basis”. The correspondence $T \leftrightarrow A_T$ (or $A \leftrightarrow T_A$) uses the standard basis implicitly.

Consider now general finite dimensional vector spaces \mathcal{V} and \mathcal{W} . Let $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$ and let $\mathbf{v} = \{v_1, \dots, v_n\}$ be a basis for \mathcal{V} . As mentioned earlier, the images $\{Tv_1, \dots, Tv_n\}$ of the basis elements determine T completely. In fact, expanding any vector $v \in \mathcal{V}$ as $v = \sum c_j v_j$, we must have $Tv = \sum c_j Tv_j$.

On the other hand, given any vectors $y_j \in \mathcal{W}$, $j = 1, \dots, n$ we obtain an element $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$ by declaring that $Tv_j = y_j$ for $j = 1, \dots, n$, and (necessarily) $T(\sum a_j v_j) = \sum a_j y_j$. Thus, the choice of a basis in \mathcal{V} determines a 1-1 correspondence between the elements of $\mathcal{L}(\mathcal{V}, \mathcal{W})$ and n -tuples of vectors in \mathcal{W} .

2.4.4 If $\mathbf{w} = \{w_1, \dots, w_m\}$ is a basis for \mathcal{W} , and $Tv_j = \sum_{k=1}^m t_{k,j} w_k$, then, for any vector $v = \sum c_j v_j$, we have

$$(2.4.1) \quad Tv = \sum c_j Tv_j = \sum_j \sum_k c_j t_{k,j} w_k = \sum_k \left(\sum_j c_j t_{k,j} \right) w_k.$$

Given the bases $\{v_1, \dots, v_n\}$ and $\{w_1, \dots, w_m\}$, the full information about T is contained in the matrix

$$(2.4.2) \quad A_{T, \mathbf{v}, \mathbf{w}} = \begin{bmatrix} t_{11} & \dots & t_{1n} \\ t_{21} & \dots & t_{2n} \\ \vdots & \dots & \vdots \\ t_{m1} & \dots & t_{mn} \end{bmatrix} = (\mathbf{C}_{\mathbf{w}} T v_1, \dots, \mathbf{C}_{\mathbf{w}} T v_n).$$

The “coordinates operators”, $\mathbf{C}_{\mathbf{w}}$, assign to each vector in \mathcal{W} the column of its coordinates with respect to the basis \mathbf{w} , see (2.1.8).

When $\mathcal{W} = \mathcal{V}$ and $\mathbf{w} = \mathbf{v}$ we write $A_{T, \mathbf{v}}$ instead of $A_{T, \mathbf{v}, \mathbf{v}}$.

Given the bases \mathbf{v} and \mathbf{w} , and the matrix $A_{T, \mathbf{v}, \mathbf{w}}$, the operator T is explicitly defined by (2.4.1) or equivalently by

$$(2.4.3) \quad \mathbf{C}_{\mathbf{w}} T v = A_{T, \mathbf{v}, \mathbf{w}} \mathbf{C}_{\mathbf{v}} v.$$

Let $A \in \mathcal{M}(m, n)$, and denote by Sv the vector in \mathcal{W} whose coordinates with respect to \mathbf{w} are given by the column $A \mathbf{C}_{\mathbf{v}} v$. So defined, S is clearly a linear operator in $\mathcal{L}(\mathcal{V}, \mathcal{W})$ and $A_{S, \mathbf{v}, \mathbf{w}} = A$. This gives:

Theorem. *Given the vector spaces \mathcal{V} and \mathcal{W} with bases $\mathbf{v} = \{v_1, \dots, v_n\}$ and $\mathbf{w} = \{w_1, \dots, w_m\}$ respectively, the map $T \mapsto A_{T, \mathbf{v}, \mathbf{w}}$ is a bijection of $\mathcal{L}(\mathcal{V}, \mathcal{W})$ onto $\mathcal{M}(m, n)$.*

2.4.5 CHANGE OF BASIS. Assume now that $\mathcal{W} = \mathcal{V}$, and that \mathbf{v} and \mathbf{w} are arbitrary bases. The \mathbf{v} -coordinates of a vector v are given by $\mathbf{C}_{\mathbf{v}} v$ and the \mathbf{w} -coordinates of v by $\mathbf{C}_{\mathbf{w}} v$. If we are given the \mathbf{v} -coordinates of a vector v , say $x = \mathbf{C}_{\mathbf{v}} v$, and we need the \mathbf{w} -coordinates of v , we observe that $v = \mathbf{C}_{\mathbf{v}}^{-1} x$, and hence $\mathbf{C}_{\mathbf{w}} v = \mathbf{C}_{\mathbf{w}} \mathbf{C}_{\mathbf{v}}^{-1} x$. In other words, the operator

$$(2.4.4) \quad \mathbf{C}_{\mathbf{w}, \mathbf{v}} = \mathbf{C}_{\mathbf{w}} \mathbf{C}_{\mathbf{v}}^{-1}$$

on \mathbb{F}^n assigns to the \mathbf{v} -coordinates of a vector $v \in \mathcal{V}$ its \mathbf{w} -coordinates. The factor $\mathbf{C}_{\mathbf{v}}^{-1}$ identifies the vector from its \mathbf{v} -coordinates, and $\mathbf{C}_{\mathbf{w}}$ assigns to the identified vector its \mathbf{w} -coordinates; the space \mathcal{V} remains in the background. Notice that $\mathbf{C}_{\mathbf{v}, \mathbf{w}}^{-1} = \mathbf{C}_{\mathbf{w}, \mathbf{v}}$

Suppose that we have the matrix $A_{T,\mathbf{w}}$ of an operator $T \in \mathcal{L}(\mathcal{V})$ relative to a basis \mathbf{w} , and we need to have the matrix $A_{T,\mathbf{v}}$ of the same operator T , but relative to a basis \mathbf{v} . (Much of the work in linear algebra revolves around finding a basis relative to which the matrix of a given operator is as simple as possible—a simple matrix is one that sheds light on the structure, or properties, of the operator.) Claim:

$$(2.4.5) \quad A_{T,\mathbf{v}} = \mathbf{C}_{\mathbf{v},\mathbf{w}} A_{T,\mathbf{w}} \mathbf{C}_{\mathbf{w},\mathbf{v}},$$

$\mathbf{C}_{\mathbf{w},\mathbf{v}}$ assigns to the \mathbf{v} -coordinates of a vector $v \in \mathcal{V}$ its \mathbf{w} -coordinates; $A_{T,\mathbf{w}}$ replaces the \mathbf{w} -coordinates of v by those of Tv ; $\mathbf{C}_{\mathbf{v},\mathbf{w}}$ identifies Tv from its \mathbf{w} -coordinates, and produces its \mathbf{v} -coordinates.

2.4.6 How special are the matrices (operators) $\mathbf{C}_{\mathbf{w},\mathbf{v}}$? They are clearly non-singular, and that is a complete characterization.

Proposition. *Given a basis $\mathbf{w} = \{w_1, \dots, w_n\}$ of \mathcal{V} , the map $\mathbf{v} \mapsto \mathbf{C}_{\mathbf{w},\mathbf{v}}$ is a bijection of the set of bases \mathbf{v} of \mathcal{V} onto $\mathbf{GL}(n, \mathbb{F})$.*

PROOF: Injectivity: Since $\mathbf{C}_{\mathbf{w}}$ is non-singular, the equality $\mathbf{C}_{\mathbf{w},\mathbf{v}_1} = \mathbf{C}_{\mathbf{w},\mathbf{v}_2}$ implies $\mathbf{C}_{\mathbf{v}_1}^{-1} = \mathbf{C}_{\mathbf{v}_2}^{-1}$, and since $\mathbf{C}_{\mathbf{v}_1}^{-1}$ maps the elements of the standard basis of \mathbb{F}^n onto the corresponding elements in \mathbf{v}_1 , and $\mathbf{C}_{\mathbf{v}_2}^{-1}$ maps the same vectors onto the corresponding elements in \mathbf{v}_2 , we have $\mathbf{v}_1 = \mathbf{v}_2$.

Surjectivity: Let $S \in \mathbf{GL}(n, \mathbb{F})$ be arbitrary. We shall exhibit a base \mathbf{v} such that $S = \mathbf{C}_{\mathbf{w},\mathbf{v}}$. By definition, $\mathbf{C}_{\mathbf{w}} w_j = e_j$, (recall that $\{e_1, \dots, e_n\}$ is the standard basis for \mathbb{F}^n). Define the vectors v_j by the condition: $\mathbf{C}_{\mathbf{w}} v_j = S e_j$, that is, v_j is the vector whose \mathbf{w} -coordinates are given by the j 'th column of S . As S is non-singular the v_j 's are linearly independent, hence form a basis \mathbf{v} of \mathcal{V} .

For all j we have $v_j = \mathbf{C}_{\mathbf{v}}^{-1} e_j$ and $\mathbf{C}_{\mathbf{w},\mathbf{v}} e_j = \mathbf{C}_{\mathbf{w}} v_j = S e_j$. This proves that $S = \mathbf{C}_{\mathbf{w},\mathbf{v}}$ ◀

2.4.7 SIMILARITY. The matrices B_1 and B_2 are said to be *similar* if they represent the same operator T in terms of (possibly) different bases, that is, $B_1 = A_{T,\mathbf{v}}$ and $B_2 = A_{T,\mathbf{w}}$.

If B_1 and B_2 are similar, they are related by (2.4.5). By Proposition 2.4.6 we have

Proposition. *The Matrices B_1 and B_2 are similar if, and only if there exists $C \in \mathbf{GL}(n, \mathbb{F})$ such that*

$$(2.4.6) \quad B_1 = CB_2C^{-1}.$$

We shall see later (see exercise **V.6.3**) that if there exists such C with entries in some field extension of \mathbb{F} , then one exists in $\mathcal{M}(n, \mathbb{F})$.

2.4.8 The operators $S, T \in \mathcal{L}(\mathcal{V})$ are said to be *similar* if there is an operator $R \in \mathbf{GL}(\mathcal{V})$ such that

$$(2.4.7) \quad T = RSR^{-1}.$$

EXERCISES FOR SECTION 2.4

II.4.1. Prove that $S, T \in \mathcal{L}(\mathcal{V})$ are similar if, and only if, their matrices (relative to any basis) are similar. An equivalent condition is: for any basis \mathbf{w} there is a basis \mathbf{v} such that $A_{T, \mathbf{v}} = A_{S, \mathbf{w}}$.

II.4.2. Let $\mathbb{F}_n[x]$ be the space of polynomials $\sum_0^n a_j x^j$. Let D be the differentiation operator and $T = 2D + I$.

a. What is the matrix corresponding to T relative to the basis $\{x^j\}_{j=0}^n$?

b. Verify that, if $u_j = \sum_{l=j}^n x^l$, then $\{u_j\}_{j=0}^n$ is a basis, and find the matrix corresponding to T relative to this basis.

II.4.3. Prove that if $A \in \mathcal{M}(l, m)$, the map $T: B \mapsto AB$ is a linear operator $\mathcal{M}(m, n) \mapsto \mathcal{M}(l, n)$. In particular, if $n = 1$, $\mathcal{M}(m, 1) = \mathbb{F}_c^m$ and $\mathcal{M}(l, 1) = \mathbb{F}_c^l$ and $T \in \mathcal{L}(\mathbb{F}_c^m, \mathbb{F}_c^l)$. What is the relation between A and the matrix A_T defined in 2.4.3 (for the standard bases, and with n there replaced here by l)?

2.5 KERNEL, RANGE, NULLITY, AND RANK

2.5.1 DEFINITION: The *kernel* of an operator $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$ is the set

$$\ker(T) = \{v \in \mathcal{V} : Tv = 0\}.$$

The *range* of T is the set

$$\text{range}(T) = T\mathcal{V} = \{w \in \mathcal{W} : w = Tv \text{ for some } v \in \mathcal{V}\}.$$

The kernel is also called the *nullspace* of T .

Proposition. *Assume $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$. Then $\ker(T)$ is a subspace of \mathcal{V} , and $\text{range}(T)$ is a subspace of \mathcal{W} .*

PROOF: If $v_1, v_2 \in \ker(T)$ then $T(a_1v_1 + a_2v_2) = a_1Tv_1 + a_2Tv_2 = 0$.
If $v_j = Tu_j$ then $a_1v_1 + a_2v_2 = T(a_1u_1 + a_2u_2)$. ◀

If \mathcal{V} is finite dimensional and $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$ then both $\ker(T)$ and $\text{range}(T)$ are finite dimensional; the first since it is a subspace of a finite dimensional space, the second as the image of one, (since, if $\{v_1, \dots, v_n\}$ is a basis for \mathcal{V} , $\{Tv_1, \dots, Tv_n\}$ spans $\text{range}(T)$).

We define the *rank* of T , denoted $\rho(T)$, as the dimension of $\text{range}(T)$. We define the *nullity* of T , denoted $\nu(T)$, as the dimension of $\ker(T)$.

Theorem (Rank and nullity). *Assume $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$, \mathcal{V} finite dimensional.*

$$(2.5.1) \quad \rho(T) + \nu(T) = \dim \mathcal{V}.$$

PROOF: Let $\{v_1, \dots, v_l\}$ be a basis for $\ker(T)$, $l = \nu(T)$, and extend it to a basis of \mathcal{V} by adding $\{u_1, \dots, u_k\}$. By 1.2.4 we have $l+k = \dim \mathcal{V}$. The theorem follows if we show that $k = \rho(T)$. We do it by showing that $\{Tu_1, \dots, Tu_k\}$ is a basis for $\text{range}(T)$.

Write any $v \in \mathcal{V}$ as $\sum_{i=1}^l a_i v_i + \sum_{i=1}^k b_i u_i$. Since $Tv_i = 0$, we have $Tv = \sum_{i=1}^k b_i Tu_i$, which shows that $\{Tu_1, \dots, Tu_k\}$ spans $\text{range}(T)$.

We claim that $\{Tu_1, \dots, Tu_k\}$ is also independent. To show this, assume that $\sum_{j=1}^k c_j Tu_j = 0$, then $T(\sum_{j=1}^k c_j u_j) = 0$, that is $\sum_{j=1}^k c_j u_j \in \ker(T)$. Since $\{v_1, \dots, v_l\}$ is a basis for $\ker(T)$, we have $\sum_{j=1}^k c_j u_j = \sum_{j=1}^l d_j v_j$ for appropriate constants d_j . But $\{v_1, \dots, v_l\} \cup \{u_1, \dots, u_k\}$ is independent, and we obtain $c_j = 0$ for all j . ◀

The proof gives more than is claimed in the theorem. It shows that T can be “factored” as a product of two maps. The first is the quotient map $\mathcal{V} \mapsto \mathcal{V}/\ker(T)$; vectors that are congruent modulo $\ker(T)$ have the same image under T . The second, $\mathcal{V}/\ker(T) \mapsto TV$ is an isomorphism. (This is the *Homomorphism Theorem* of groups in our context.)

2.5.2 The identity operator, defined by $Iv = v$, is an identity element in the algebra $\mathcal{L}(\mathcal{V})$. The invertible elements in $\mathcal{L}(\mathcal{V})$ are the *automorphisms* of \mathcal{V} , that is, the *bijective* linear maps. In the context of finite dimensional spaces, either *injectivity* (i.e. being 1-1) or *surjectivity* (onto) implies the other:

Theorem. *Let \mathcal{V} be a finite dimensional vector space, $T \in \mathcal{L}\mathcal{V}$. Then*

$$(2.5.2) \quad \ker(T) = \{0\} \iff \text{range}(T) = \mathcal{V},$$

and either condition is equivalent to: “ T is invertible”, aka “nonsingular”.

PROOF: $\ker(T) = \{0\}$ is equivalent to $\nu(T) = 0$, and $\text{range}(T) = \mathcal{V}$ is equivalent to $\rho(T) = \dim \mathcal{V}$. Now apply (2.5.1). ◀

2.5.3 As another illustration of how the “rank and nullity” theorem can be used, consider the following statement (which can be seen directly as a consequence of exercise **I.2.12**)

Theorem. *Let $\mathcal{V} = \mathcal{V}_1 \oplus \mathcal{V}_2$ be finite dimensional, $\dim \mathcal{V}_1 = k$. Let $\mathcal{W} \subset \mathcal{V}$ be a subspace of dimension $l > k$. Then $\dim \mathcal{W} \cap \mathcal{V}_2 \geq l - k$.*

PROOF: Denote by π_1 the restriction to \mathcal{W} of the projection of \mathcal{V} on \mathcal{V}_1 along \mathcal{V}_2 . Since the rank of π_1 is clearly $\leq k$, the nullity is $\geq l - k$. In other words, the kernel of this map, namely $\mathcal{W} \cap \mathcal{V}_2$, has dimension $\geq l - k$. ◀

EXERCISES FOR SECTION 2.5

II.5.1. Assume $T, S \in \mathcal{L}(\mathcal{V})$. Prove that $\nu(ST) \leq \nu(S) + \nu(T)$.

II.5.2. Give an example of two 2×2 matrices A and B such that $\rho(AB) = 1$ and $\rho(BA) = 0$.

II.5.3. Given vector spaces \mathcal{V} and \mathcal{W} over the same field. Let $\{v_j\}_{j=1}^n \subset \mathcal{V}$ and $\{w_j\}_{j=1}^n \subset \mathcal{W}$. Prove that there exists a linear map $T: \text{span}[v_1, \dots, v_n] \mapsto \mathcal{W}$ such that $Tv_j = w_j$, $j = 1, \dots, n$ if, and only if, the following implication holds:

$$\text{If } a_j, j = 1, \dots, n \text{ are scalars, and } \sum_1^n a_j v_j = 0, \quad \text{then } \sum_1^n a_j w_j = 0.$$

Can the definition of T be extended to the entire \mathcal{V} ?

II.5.4. What is the relationship of the previous exercise to Theorem 1.3.5?

II.5.5. The operators $T, S \in \mathcal{L}(\mathcal{V})$ are called “*equivalent*” if there exist invertible $A, B \in \mathcal{L}(\mathcal{V})$ such that

$$S = ATB \quad (\text{so that } T = A^{-1}SB^{-1}).$$

Prove that if \mathcal{V} is finite dimensional then T, S are “*equivalent*” if, and only if

$$\rho(S) = \rho(T).$$

II.5.6. Give an example of two operators on \mathbb{F}^3 that are equivalent but not similar.

II.5.7. Assume $T, S \in \mathcal{L}(\mathcal{V})$. Prove that the following statements are equivalent:

- a.** $\ker(S) \subset \ker(T)$,
- b.** There exists $R \in \mathcal{L}(\mathcal{V})$ such that $T = RS$.

Hint: For the implication **a.** \implies **b.**: Choose a basis $\{v_1, \dots, v_s\}$ for $\ker(S)$. Expand it to a basis for $\ker(T)$ by adding $\{u_1, \dots, u_{t-s}\}$, and expand further to a basis for \mathcal{V} by adding the vectors $\{w_1, \dots, w_{n-t}\}$.

The sequence $\{Su_1, \dots, Su_{t-s}\} \cup \{Sw_1, \dots, Sw_{n-t}\}$ is independent, so that R can be defined arbitrarily on it (and extended by linearity to an operator on the entire space). Define $R(Su_j) = 0$, $R(Sw_j) = Tw_j$.

The other implication is obvious.

II.5.8. Assume $T, S \in \mathcal{L}(\mathcal{V})$. Prove that the following statements are equivalent:

- a.** $\text{range}(S) \subset \text{range}(T)$,
- b.** There exists $R \in \mathcal{L}(\mathcal{V})$ such that $S = TR$.

Hint: Again, **b.** \implies **a.** is obvious.

For **a.** \implies **b.** Take a basis $\{v_1, \dots, v_n\}$ for \mathcal{V} . Let $u_j, j = 1, \dots, n$ be such that $Tu_j = Sv_j$, (use assumption **a.**). Define $Rv_j = u_j$ (and extend by linearity).

II.5.9. Find bases for the null space, $\ker(A)$, and for the range, $\text{range}(A)$, of the matrix (acting on rows in \mathbb{R}^5)

$$\begin{bmatrix} 1 & 0 & 0 & 5 & 9 \\ 0 & 1 & 0 & -3 & 2 \\ 0 & 0 & 1 & 2 & 1 \\ 3 & 2 & 1 & 11 & 32 \\ 1 & 2 & 0 & -1 & 13 \end{bmatrix}.$$

II.5.10. Let $T \in \mathcal{L}(V)$, $l \in \mathbb{N}$. Prove:

- a. $\ker(T^l) \subseteq \ker(T^{l+1})$; equality if, and only if $\text{range}(T^l) \cap \ker(T) = \{0\}$.
- b. $\text{range}(T^{l+1}) \subseteq \text{range}(T^l)$; equality if, and only if, $\ker(T^{l+1}) = \ker(T^l)$.
- c. If $\ker(T^{l+1}) = \ker(T^l)$, then $\ker(T^{l+k+1}) = \ker(T^{l+k})$ for all positive integers k .

II.5.11. An operator T is *idempotent* if $T^2 = T$. Prove that an idempotent operator is a projection on $\text{range}(T)$ along $\ker(T)$.

*2.6 NORMED FINITE DIMENSIONAL LINEAR SPACES

2.6.1 A norm on a real or complex vector space \mathcal{V} is a nonnegative function $v \mapsto \|v\|$ that satisfies the conditions

- a. Positivity: $\|0\| = 0$ and if $v \neq 0$ then $\|v\| > 0$.
- b. Homogeneity: $\|av\| = |a|\|v\|$ for scalars a and vectors v .
- c. The triangle inequality: $\|v + u\| \leq \|v\| + \|u\|$.

These properties guarantee that $\rho(v, u) = \|v - u\|$ is a metric on the space, and with a metric one can use tools and notions from point-set topology such as limits, continuity, convergence, infinite series, etc.

A vector space endowed with a norm is a *normed vector space*.

2.6.2 If \mathcal{V} and \mathcal{W} are isomorphic real or complex n -dimensional spaces and S is an isomorphism of \mathcal{V} onto \mathcal{W} , then a norm $\|\cdot\|^*$ on \mathcal{W} can be transported to \mathcal{V} by defining $\|v\| = \|Sv\|^*$. This implies that all possible norms on a real n -dimensional space are copies of norms on \mathbb{R}^n , and all norms on a complex n -dimensional space are copies of norms on \mathbb{C}^n .

A finite dimensional \mathcal{V} can be endowed with many different norms; yet, all these norms are *equivalent* in the following sense:

DEFINITION: The norms $\|\cdot\|_1$ and $\|\cdot\|_2$ are equivalent, written: $\|\cdot\|_1 \sim \|\cdot\|_2$ if there is a positive constant C such that for all $v \in \mathcal{V}$

$$C^{-1}\|v\|_1 \leq \|v\|_2 \leq C\|v\|_1$$

The metrics ρ_1, ρ_2 , defined by equivalent norms, are equivalent: for $v, u \in \mathcal{V}$

$$C^{-1}\rho_1(v, u) \leq \rho_2(v, u) \leq C\rho_1(v, u).$$

which means that they define the same topology—the familiar topology of \mathbb{R}^n or \mathbb{C}^n .

2.6.3 If \mathcal{V} and \mathcal{W} are normed vector spaces we define a norm on $\mathcal{L}(\mathcal{V}, \mathcal{W})$ by writing, for $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$,

$$(2.6.1) \quad \|T\| = \max_{\|v\|=1} \|Tv\| = \max_{v \neq 0} \frac{\|Tv\|}{\|v\|}.$$

Equivalently,

$$(2.6.2) \quad \|T\| = \inf\{C : \|Tv\| \leq C\|v\| \text{ for all } v \in \mathcal{H}\}.$$

To check that (2.6.1) defines a norm we observe that properties **a.** and **b.** are obvious, and that **c.** follows from²

$$\|(T + S)v\| \leq \|Tv\| + \|Sv\| \leq \|T\|\|v\| + \|S\|\|v\| \leq (\|T\| + \|S\|)\|v\|.$$

$\mathcal{L}(\mathcal{V})$ is an algebra and we observe that the norm defined by (2.6.1) on $\mathcal{L}(\mathcal{V})$ is *submultiplicative*: we have $\|STv\| \leq \|S\|\|Tv\| \leq \|S\|\|T\|\|v\|$, where $S, T \in \mathcal{L}(\mathcal{V})$ and $v \in \mathcal{V}$, which means

$$(2.6.3) \quad \|ST\| \leq \|S\|\|T\|.$$

EXERCISES FOR SECTION 2.6

II.6.1. Let \mathcal{V} be n -dimensional real or complex vector space, $\mathbf{v} = \{v_1, \dots, v_n\}$ a basis for \mathcal{V} . Write $\|\sum a_j v_j\|_{\mathbf{v},1} = \sum |a_j|$, and $\|\sum a_j v_j\|_{\mathbf{v},\infty} = \max |a_j|$.

Prove:

a. $\|\cdot\|_{\mathbf{v},1}$ and $\|\cdot\|_{\mathbf{v},\infty}$ are norms on \mathcal{V} , and

$$(2.6.4) \quad \|\cdot\|_{\mathbf{v},\infty} \leq \|\cdot\|_{\mathbf{v},1} \leq n\|\cdot\|_{\mathbf{v},\infty}$$

²Notice that the norms appearing in the inequalities are the ones defined on \mathcal{W} , $\mathcal{L}(\mathcal{V}, \mathcal{W})$, and \mathcal{V} , respectively.

b. If $\|\cdot\|$ is any norm on \mathcal{V} then, for all $v \in \mathcal{V}$,

$$(2.6.5) \quad \|v\|_{\mathcal{V},1} \max \|v_j\| \geq \|v\|.$$

II.6.2. Let $\|\cdot\|_j, j = 1, 2$, be norms on \mathcal{V} , and ρ_j the induced metrics. Let $\{v_n\}_{n=0}^{\infty}$ be a sequence in \mathcal{V} and assume that $\rho_1(v_n, v_0) \rightarrow 0$. Prove $\rho_2(v_n, v_0) \rightarrow 0$.

II.6.3. Let $\{v_n\}_{n=0}^{\infty}$ be bounded in \mathcal{V} . Prove that $\sum_0^{\infty} v_n z^n$ converges for every z such that $|z| < 1$.

Hint: Prove that the partial sums form a Cauchy sequence in the metric defined by the norm.

II.6.4. Let \mathcal{V} be n -dimensional real or complex normed vector space. The *unit ball* in \mathcal{V} is the set

$$B_1 = \{v \in \mathcal{V} : \|v\| \leq 1\}.$$

Prove that B_1 is

convex: If $v, u \in B_1, 0 \leq a \leq 1$, then $av + (1-a)u \in B_1$.

Bounded: For every $v \in \mathcal{V}$, there exist a (positive) constant λ such that $cv \notin B$ for $|c| > \lambda$.

Symmetric, centered at 0: If $v \in B$ and $|a| \leq 1$ then $av \in B$.

II.6.5. Let \mathcal{V} be n -dimensional real or complex vector space, and let B be a bounded symmetric convex set centered at 0. Define

$$\|u\| = \inf\{a > 0 : a^{-1}u \in B\}.$$

Prove that this defines a norm on \mathcal{V} , and the unit ball for this norm is the given B

II.6.6. Describe a norm $\|\cdot\|_0$ on \mathbb{R}^3 such that the standard unit vectors have norm 1 while $\|(1, 1, 1)\|_0 < \frac{1}{100}$.

II.6.7. Let \mathcal{V} be a normed linear space and $T \in \mathcal{L}(\mathcal{V})$. Prove that the set of vectors $v \in \mathcal{V}$ whose T -orbit, $\{T^n v\}$, is bounded is a subspace of \mathcal{V} .

Chapter III

Duality of vector spaces

3.1 LINEAR FUNCTIONALS

Let \mathcal{V} be a finite dimensional vector space with basis $\{v_1, \dots, v_n\}$. Every element $v \in \mathcal{V}$ can be written, in exactly one way, as

$$(3.1.1) \quad v = \sum_1^n a_j(v)v_j,$$

the notation $a_j(v)$ comes to emphasize the dependence of the coefficients on the vector v .

Let $v = \sum_1^n a_j(v)v_j$, and $u = \sum_1^n a_j(u)v_j$. If $c, d \in \mathbb{F}$, then

$$cv + du = \sum_1^n (ca_j(v) + da_j(u))v_j$$

so that

$$a_j(cv + du) = ca_j(v) + da_j(u).$$

In other words, $a_j(v)$ are linear functionals on \mathcal{V} .

A standard notation for the image of a vector v under a linear functional v^* is (v, v^*) . Accordingly we denote the linear functionals corresponding to $a_j(v)$ by v_j^* and write

$$(3.1.2) \quad a_j(v) = (v, v_j^*) \quad \text{so that} \quad v = \sum_1^n (v, v_j^*)v_j.$$

Proposition. *The linear functionals v_j^* , $j = 1, \dots, n$ form a basis for the dual space \mathcal{V}^* .*

PROOF: Let $u^* \in \mathcal{V}^*$. Write $b_j(u^*) = (v_j, u^*)$, then for any $v \in \mathcal{V}$,

$$(v, u^*) = \left(\sum_j (v, v_j^*) v_j, u^* \right) = \sum_j (v, v_j^*) b_j(u^*) = (v, \sum_j b_j(u^*) v_j^*),$$

and $u^* = \sum b_j(u^*) v_j^*$. It follows that $\{v_1^*, \dots, v_n^*\}$ spans \mathcal{V}^* . On the other hand, $\{v_1^*, \dots, v_n^*\}$ is independent since $\sum c_j v_j^* = 0$ implies $(v_k, \sum c_j v_j^*) = c_k = 0$ for all k . ◀

Corollary. $\dim \mathcal{V}^* = \dim \mathcal{V}$.

The basis $\{v_j^*\}_1^n, j = 1, \dots, n$ is called *the dual basis* of $\{v_1, \dots, v_n\}$. It is characterized by the condition

$$(3.1.3) \quad (v_j, v_k^*) = \delta_{j,k},$$

$\delta_{j,k}$ is the Kronecker delta, it takes the value 1 if $j = k$, and 0 otherwise.

3.1.1 The way we add linear functionals or multiply them by scalars guarantees that the form (expression) (v, v^*) , $v \in \mathcal{V}$ and $v^* \in \mathcal{V}^*$, is *bilinear*, that is linear in v for every fixed v^* , and linear in v^* for any fixed v . Thus every $v \in \mathcal{V}$ defines a linear functional on \mathcal{V}^* .

If $\{v_1, \dots, v_n\}$ is a basis for \mathcal{V} , and $\{v_1^*, \dots, v_n^*\}$ the dual basis in \mathcal{V}^* , then (3.1.3) identifies $\{v_1, \dots, v_n\}$ as the dual basis of $\{v_1^*, \dots, v_n^*\}$. The roles of \mathcal{V} and \mathcal{V}^* are perfectly symmetric and what we have is *two spaces in duality*, the duality between them defined by the bilinear form (v, v^*) . (3.1.2) works in both directions, thus if $\{v_1, \dots, v_n\}$ and $\{v_1^*, \dots, v_n^*\}$ are dual bases, then for all $v \in \mathcal{V}$ and $v^* \in \mathcal{V}^*$,

$$(3.1.4) \quad v = \sum_1^n (v, v_j^*) v_j, \quad v^* = \sum_1^n (v_j, v^*) v_j^*.$$

The dual of \mathbb{F}_c^n (i.e., \mathbb{F}^n written as columns) can be identified with \mathbb{F}_r^n (i.e., \mathbb{F}^n written as rows) and the pairing (v, v^*) as the matrix product $v^* v$ of the row v^* by the column v , (exercise **III.1.4** below). The dual of the standard basis of \mathbb{F}_c^n is the standard basis \mathbb{F}_r^n .

3.1.2 ANNIHILATOR. Given a set $A \subset \mathcal{V}$, the set of all the linear functionals $v^* \in \mathcal{V}^*$ that vanish identically on A is called *the annihilator* of A and denoted A^\perp . Clearly, A^\perp is a subspace of \mathcal{V}^* .

Functionals that annihilate A vanish on $\text{span}[A]$ as well, and functionals that annihilate $\text{span}[A]$ clearly vanish on A ; hence $A^\perp = (\text{span}[A])^\perp$.

Proposition. *Let $\mathcal{V}_1 \subset \mathcal{V}$ be a subspace, then $\dim \mathcal{V}_1 + \dim \mathcal{V}_1^\perp = \dim \mathcal{V}$.*

PROOF: Let $\{v_1, \dots, v_m\}$ be a basis for \mathcal{V}_1 , and let $\{v_{m+1}, \dots, v_n\}$ complete it to a basis for \mathcal{V} . Let $\{v_1^*, \dots, v_n^*\}$ be the dual basis.

We claim, that $\{v_{m+1}^*, \dots, v_n^*\}$ is a basis for \mathcal{V}_1^\perp ; hence $\dim \mathcal{V}_1^\perp = n - m$ proving the proposition.

By (3.1.3) we have $\{v_{m+1}^*, \dots, v_n^*\} \subset \mathcal{V}_1^\perp$, and we know these vectors to be independent. We only need to prove that they span \mathcal{V}_1^\perp .

Let $w^* \in \mathcal{V}_1^\perp$. Write $w^* = \sum_{j=1}^n a_j v_j^*$, and observe that $a_j = (v_j, w^*)$. Now $w^* \in \mathcal{V}_1^\perp$ implies $a_j = 0$ for $1 \leq j \leq m$, so that $w^* = \sum_{m+1}^n a_j v_j^*$. ◀

Theorem. *Let $A \subset \mathcal{V}$, $v \in \mathcal{V}$ and assume that $(v, u^*) = 0$ for every $u^* \in A^\perp$. Then $v \in \text{span}[A]$.*

Equivalent statement: If $v \notin \text{span}[A]$ then there exists $u^ \in A^\perp$ such that $(v, u^*) \neq 0$.*

PROOF: If $v \notin \text{span}[A]$, then $\dim \text{span}[A, v] = \dim \text{span}[A] + 1$, hence $\dim \text{span}[A, v]^\perp = \dim \text{span}[A]^\perp - 1$. It follows that $\text{span}[A]^\perp \supsetneq \text{span}[A, v]^\perp$, and since functionals in A^\perp which annihilate v annihilate $\text{span}[A, v]$, there exist functionals in A^\perp that do not annihilate v . ◀

3.1.3 Let \mathcal{V} be a finite dimensional vector space and $\mathcal{V}_1 \subset \mathcal{V}$ a subspace. Restricting the domain of a linear functional in \mathcal{V}^* to \mathcal{V}_1 defines a linear functional on \mathcal{V}_1 .

The functionals whose restriction to \mathcal{V}_1 is zero are, by definition, the elements of \mathcal{V}_1^\perp . The restrictions of v^* and u^* to \mathcal{V}_1 are equal if, and only if, $v^* - u^* \in \mathcal{V}_1^\perp$. This, combined with exercise **III.1.2** below, gives a natural identification of \mathcal{V}_1^* with the quotient space $\mathcal{V}^*/\mathcal{V}_1^\perp$.

EXERCISES FOR SECTION 3.1

III.1.1. Given a linearly independent $\{v_1, \dots, v_k\} \subset \mathcal{V}$ and scalars $\{a_j\}_{j=1}^k$. Prove that there exists $v^* \in \mathcal{V}^*$ such that $(v_j, v^*) = a_j$ for $1 \leq j \leq k$.

III.1.2. If \mathcal{V}_1 is a subspace of a finite dimensional space \mathcal{V} then every linear functional on \mathcal{V}_1 is the restriction to \mathcal{V}_1 of a linear functional on \mathcal{V} .

III.1.3. Let \mathcal{V} be a finite dimensional vector space, $\mathcal{V}_1 \subset \mathcal{V}$ a subspace. Let $\{u_k^*\}_{k=1}^r \subset \mathcal{V}^*$ be linearly independent mod \mathcal{V}_1^\perp (i.e., if $\sum c_k u_k^* \in \mathcal{V}_1^\perp$, then $c_k = 0$, $k = 1, \dots, r$). Let $\{v_j^*\}_{j=1}^s \subset \mathcal{V}_1^\perp$, be independent. Prove that $\{u_k^*\} \cup \{v_j^*\}$ is linearly independent in \mathcal{V}^* .

III.1.4. Show that every linear functional on $\mathbb{F}_\mathbf{r}^n$ is given by some $(a_1, \dots, a_n) \in \mathbb{F}_\mathbf{r}^n$ as

$$\begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \mapsto (a_1, \dots, a_n) \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} = \sum a_j x_j$$

III.1.5. Let \mathcal{V} and \mathcal{W} be finite dimensional vector spaces.

a. Prove that for every $v \in \mathcal{V}$ and $w^* \in \mathcal{W}^*$ the map

$$\varphi_{v, w^*} : T \mapsto (Tv, w^*)$$

is a linear functional on $\mathcal{L}(\mathcal{V}, \mathcal{W})$.

b. Prove that the map $v \otimes w^* \mapsto \varphi_{v, w^*}$ is an isomorphism of $\mathcal{V} \otimes \mathcal{W}^*$ onto the dual space of $\mathcal{L}(\mathcal{V}, \mathcal{W})$.

III.1.6. Let \mathcal{V} be a complex vector space, $\{v_j^*\}_{j=1}^s \subset \mathcal{V}^*$, and $w^* \in \mathcal{V}^*$ such that for all $v \in \mathcal{V}$,

$$|\langle v, w^* \rangle| \leq \max_{j=1}^s |\langle v, v_j^* \rangle|.$$

Prove that $w^* \in \text{span}[\{v_j^*\}_{j=1}^s]$.

III.1.7. Linear functionals on $\mathbb{R}_N[x]$:

1. Show that for every $x \in \mathbb{R}$ the map φ_x defined by $(P, \varphi_x) = P(x)$ is a linear functional on $\mathbb{R}_N[x]$.
2. If $\{x_1, \dots, x_m\}$ are distinct and $m \leq N + 1$, then φ_{x_j} are linearly independent.
3. For every $x \in \mathbb{R}$ and $l \in \mathbb{N}$, $l \leq N$, the map $\varphi_x^{(l)}$ defined by $(P, \varphi_x^{(l)}) = P^{(l)}(x)$ is a (non-trivial) linear functional on $\mathbb{R}_N[x]$.

III.1.8. Let $x_j \in \mathbb{R}$, $l_j \in \mathbb{N}$, and assume that the pairs (x_j, l_j) , $j = 1, \dots, N + 1$, are distinct. Denote by $\#(m)$ the number of such pairs with $l_j > m$.

a. Prove that a necessary condition for the functionals $\varphi_{x_j}^{(l_j)}$ to be independent on $\mathbb{R}_N[x]$ is:

$$(3.1.5) \quad \text{for every } m \leq N, \quad \#(m) \leq N - m.$$

b. Check that φ_1 , φ_{-1} , and $\varphi_0^{(1)}$ are linearly dependent in the dual of $\mathbb{R}_2[x]$, hence (3.1.5) is *not* sufficient. Are φ_1 , φ_{-1} , and $\varphi_0^{(1)}$ linearly dependent in the dual of $\mathbb{R}_3[x]$?

3.2 THE ADJOINT

3.2.1 The concatenation w^*T of $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$, and $w^* \in \mathcal{W}^*$, is a linear map from \mathcal{V} to the underlying field, i.e. a linear functional v^* on \mathcal{V} .

With T fixed, the mapping $w^* \mapsto w^*T$ is a linear operator $T^* \in \mathcal{L}(\mathcal{W}^*, \mathcal{V}^*)$. It is called the *adjoint* of T .

The basic relationship between T , T^* , and the bilinear forms (v, v^*) and (w, w^*) is: For all $v \in \mathcal{V}$ and $w^* \in \mathcal{W}^*$,

$$(3.2.1) \quad (Tv, w^*) = (v, T^*w^*).$$

Notice that the left-hand side is the bilinear form on (W, W^*) , while the right-hand side in (V, V^*) .

3.2.2 Proposition.

$$(3.2.2) \quad \rho(T^*) = \rho(T).$$

PROOF: Let $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$, assume $\rho(T) = r$, and let $\{v_1, \dots, v_n\}$ be a basis for \mathcal{V} such that $\{v_{r+1}, \dots, v_n\}$ is a basis for $\ker(T)$. We have seen (see the proof of theorem 2.5.1) that $\{Tv_1, \dots, Tv_r\}$ is a basis for $T\mathcal{V} = \text{range}(T)$.

Denote $w_j = Tv_j$, $j = 1, \dots, r$. Add the vectors w_j , $j = r + 1, \dots, m$ so that $\{w_1, \dots, w_m\}$ be a basis for \mathcal{W} . Let $\{w_1^*, \dots, w_m^*\}$ be the dual basis.

Fix $k > r$; for every $j \leq r$ we have $(v_j, T^*w_k^*) = (w_j, w_k^*) = 0$ which means $T^*w_k^* = 0$. Thus $T^*\mathcal{W}^*$ is spanned by $\{T^*w_j^*\}_{j=1}^r$.

For $1 \leq i, j \leq r$, $(v_i, T^*w_j^*) = (w_i, w_j^*) = \delta_{i,j}$, which implies that $\{T^*w_j^*\}_{j=1}^r$ is linearly independent in \mathcal{V}^* .

Thus, $\{T^*w_1^*, \dots, T^*w_r^*\}$ is a basis for $T^*\mathcal{W}^*$, and $\rho(T^*) = \rho(T)$. ◀

3.2.3 We have seen in 3.1.1 that if $\mathcal{V} = \mathbb{F}_{\mathbb{C}}^n$, $\mathcal{W} = \mathbb{F}_{\mathbb{C}}^m$, both with standard bases, then $\mathcal{V}^* = \mathbb{F}_{\mathbb{R}}^n$, $\mathcal{W}^* = \mathbb{F}_{\mathbb{R}}^m$, and the standard basis of $\mathbb{F}_{\mathbb{R}}^m$ is the dual basis of the standard basis of $\mathbb{F}_{\mathbb{C}}^n$.

If $A = A_T = \begin{bmatrix} t_{11} & \dots & t_{1n} \\ \vdots & \dots & \vdots \\ t_{m1} & \dots & t_{mn} \end{bmatrix}$, is the matrix of T with respect to the stan-

dard bases, then the operator T is given as left multiplication by A on $\mathbb{F}_{\mathbb{C}}^n$ and the bilinear form (Tv, w) , for $w \in \mathbb{F}_{\mathbb{R}}^m$ and $v \in \mathbb{F}_{\mathbb{C}}^n$, is just the matrix product

$$(3.2.3) \quad w(Av) = (wA)v.$$

It follows that $T^*w = wA_T$, that is, the action of T^* on the row vectors in $\mathbb{F}_{\mathbb{R}}^m$ is obtained as multiplication on the right by the *same matrix* $A = A_T$.

If we want¹ to have the matrix of T^* relative to the standard bases in $\mathbb{F}_{\mathbb{C}}^n$ and $\mathbb{F}_{\mathbb{R}}^m$, acting on columns by left multiplication, all we need to do is transpose wA and obtain

$$T^*w^{\mathbb{T}} = A^{\mathbb{T}}w^{\mathbb{T}}.$$

3.2.4 Proposition. *Let $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$. Then*

$$(3.2.4) \quad \text{range}(T)^{\perp} = \ker(T^*) \quad \text{and} \quad \text{range}(T^*)^{\perp} = \ker(T).$$

PROOF: $w^* \in \text{range}(T)^{\perp}$ is equivalent to $(Tv, w^*) = (v, T^*w^*) = 0$ for all $v \in \mathcal{V}$, and $(v, T^*w^*) = 0$ for all $v \in \mathcal{V}$ is equivalent to $T^*w^* = 0$.

The condition $v \in \text{range}(T^*)^{\perp}$ is equivalent to $(v, T^*w^*) = 0$ for all $w^* \in \mathcal{W}^*$, and $Tv = 0$ is equivalent to $(Tv, w^*) = (v, T^*w^*) = 0$ i.e. $v \in \text{range}(T)^{\perp}$. ◀

EXERCISES FOR SECTION 3.2

¹This will be the case when there is a natural way to identify the vector space with its dual, for instance when we work with *inner product spaces*. If the “identification” is sesquilinear, as is the case when $\mathbb{F} = \mathbb{C}$ the matrix for the adjoint is the complex conjugate of $A^{\mathbb{T}}$, see Chapter VI.

III.2.1. If $\mathcal{V} = \mathcal{W} \oplus \mathcal{U}$ and S is the projection of \mathcal{V} on \mathcal{W} along \mathcal{U} (see 2.1.1.g), what is the adjoint S^* ?

III.2.2. Let $A \in \mathcal{M}(m, n; \mathbb{R})$. Prove

$$\rho(A^{\text{Tr}}A) = \rho(A)$$

III.2.3. Prove that, in the notation of 3.2.2, $\{w_j^*\}_{j=r+1, \dots, m}$ is a basis for $\ker(T^*)$.

III.2.4. A vector $v \in \mathcal{V}$ is an *eigenvector* for $T \in \mathcal{L}(\mathcal{V})$ if $Tv = \lambda v$ with $\lambda \in \mathbb{F}$; λ is the corresponding *eigenvalue*.

Let $v \in \mathcal{V}$ be an eigenvector of T with eigenvalue λ , and $w \in \mathcal{V}^*$ an eigenvector of the adjoint T^* with eigenvalue $\lambda^* \neq \lambda$. Prove that $(v, w^*) = 0$.

Chapter IV

Determinants

4.1 PERMUTATIONS

A *permutation* of a set is a bijective, that is 1-1, map of the set onto itself. The set of permutations of the set $[1, \dots, n]$ is denoted \mathbf{S}_n . It is a group under concatenation—given $\sigma, \tau \in \mathbf{S}_n$ define $\tau\sigma$ by $(\tau\sigma)(j) = \tau(\sigma(j))$ for all j . The identity element of \mathbf{S}_n is the trivial permutation e defined by $e(j) = j$ for all j .

\mathbf{S}_n with this operation is called the *symmetric group on* $[1, \dots, n]$.

4.1.1 If $\sigma \in \mathbf{S}_n$ and $a \in [1, \dots, n]$ the set $\{\sigma^k(a)\}$, is called the σ -orbit of a . If $\sigma a = a$ the orbit is *trivial*, i.e., reduced to a single point (which is left unmoved by σ). A permutation σ is called a *cycle*, and denoted (a_1, \dots, a_l) , if $\{a_j\}_{j=1}^l$ is its unique nontrivial orbit, $a_{j+1} = \sigma(a_j)$ for $1 \leq j < l$, and $a_1 = \sigma a_l$. The *length of the cycle*, l , is the period of a_1 under σ , that is, the first positive integer such that $\sigma^l(a_1) = a_1$. Observe that σ is determined by the cyclic order of the entries, thus $(a_1, \dots, a_l) = (a_l, a_1, \dots, a_{l-1})$.

Given $\sigma \in \mathbf{S}_n$, the σ -orbits form a partition of $[1, \dots, n]$, the corresponding cycles commute, and their product is σ .

Cycles of length 2 are called *transpositions*.

Lemma. *Every permutation $\sigma \in \mathbf{S}_n$ is a product of transpositions.*

PROOF: Since every $\sigma \in \mathbf{S}_n$ is a product of cycles, it suffices to show that every cycle is a product of transpositions.

Observe that

$$(a_1, \dots, a_l) = (a_l, a_1, a_2, \dots, a_{l-1}) = (a_1, a_2)(a_2, a_3) \cdots (a_{l-1}, a_l)$$

(a_l trades places with a_{l-1} , then with a_{l-2} , etc., until it settles in place of a_1 ; every other a_j moves once, to the original place of a_{j+1}). Thus, every cycle of length l is a product of $l - 1$ transpositions. ◀

Another useful observation concerns conjugation in \mathbf{S}_n . If $\sigma, \tau \in \mathbf{S}_n$, and $\tau(i) = j$ then $\tau\sigma^{-1}$ maps $\sigma(i)$ to j and $\sigma\tau\sigma^{-1}$ maps $\sigma(i)$ to $\sigma(j)$. This means that the cycles of $\sigma\tau\sigma^{-1}$ are obtained from the cycles of τ by replacing the entries there by their σ images.

In particular, *all cycles of a given length are conjugate in \mathbf{S}_n .*

4.1.2 THE SIGN OF A PERMUTATION. There are several equivalent ways to define *the sign* of a permutation $\sigma \in \mathbf{S}_n$. The sign, denoted $\mathbf{sgn}[\sigma]$, is to take the values ± 1 , assign the value -1 to each transposition, and be multiplicative: $\mathbf{sgn}[\sigma\tau] = \mathbf{sgn}[\sigma]\mathbf{sgn}[\tau]$, in other words, be a homomorphism of \mathbf{S}_n onto the multiplicative group $\{1, -1\}$.

All these requirements imply that if σ can be written as a product of k transpositions, then $\mathbf{sgn}[\sigma] = (-1)^k$. But in order to use this as the *definition* of \mathbf{sgn} one needs to prove that the numbers of factors in all the representations of any $\sigma \in \mathbf{S}_n$ as products of transpositions have the same parity. Also, finding the value of $\mathbf{sgn}[\sigma]$ this way requires a concrete representation of σ as a product of transpositions.

We introduce \mathbf{sgn} in a different way:

DEFINITION: A set J of pairs $\{(k, l)\}$ is *appropriate for \mathbf{S}_n* if it contains exactly one of (j, i) , (i, j) for every pair i, j , $1 \leq i < j \leq n$.

The simplest example is $J = \{(i, j) : 1 \leq i < j \leq n\}$. A more general example of an appropriate set is: for $\tau \in \mathbf{S}_n$,

$$(4.1.1) \quad J_\tau = \{(\tau(i), \tau(j)) : 1 \leq i < j \leq n\}.$$

If J is appropriate for \mathbf{S}_n , and $\sigma \in \mathbf{S}_n$, then¹

$$(4.1.2) \quad \prod_{i < j} \operatorname{sgn}(\sigma(j) - \sigma(i)) = \prod_{(i, j) \in J} \operatorname{sgn}(\sigma(j) - \sigma(i)) \operatorname{sgn}(j - i)$$

¹The sign of integers has the usual meaning.

since reversing a pair (i, j) changes both $\text{sgn}(\sigma(j) - \sigma(i))$ and $\text{sgn}(j - i)$, and does not affect their product.

We define the sign of a permutation σ by

$$(4.1.3) \quad \mathbf{sgn}[\sigma] = \prod_{i < j} \text{sgn}(\sigma(j) - \sigma(i))$$

Proposition. *The map $\mathbf{sgn} : \sigma \mapsto \mathbf{sgn}[\sigma]$ is a homomorphism of \mathbf{S}_n onto the multiplicative group $\{1, -1\}$. The sign of any transposition is -1 .*

PROOF: The multiplicativity is shown as follows:

$$\begin{aligned} \mathbf{sgn}[\sigma\tau] &= \prod_{i < j} \text{sgn}(\sigma\tau(j) - \sigma\tau(i)) \\ &= \prod_{i < j} \text{sgn}(\sigma\tau(j) - \sigma\tau(i)) \text{sgn}(\tau(j) - \tau(i)) \prod_{i < j} \text{sgn}(\tau(j) - \tau(i)) \\ &= \mathbf{sgn}[\sigma] \mathbf{sgn}[\tau]. \end{aligned}$$

Since the sign of the identity permutation is $+1$, the multiplicativity implies that conjugate permutations have the same sign. In particular all transpositions have the same sign. The computation for $(1, 2)$ is particularly simple:

$\text{sgn}(j - 1) = \text{sgn}(j - 2) = 1$ for all $j > 2$, while $\text{sgn}(1 - 2) = -1$ and the sign of all transpositions is -1 . ◀

EXERCISES FOR SECTION 4.1

IV.1.1. Let σ be a cycle of length k ; prove that $\mathbf{sgn}[\sigma] = (-1)^{(k-1)}$.

IV.1.2. Let $\sigma \in \mathbf{S}_n$ and assume that it has s orbits (including the trivial orbits, i.e., fixed points). Prove that $\mathbf{sgn}[\sigma] = (-1)^{n-s}$.

IV.1.3. Let $\sigma_j \in \mathbf{S}_n$, $j = 1, 2$ be cycles with different orbits, Prove that the two commute if, and only if, their (nontrivial) orbits are disjoint.

4.2 MULTILINEAR MAPS

Let \mathcal{V}_j , $j = 1, \dots, k$, and \mathcal{W} be vector spaces over a field \mathbb{F} . A map

$$(4.2.1) \quad \psi : \mathcal{V}_1 \times \mathcal{V}_2 \cdots \times \mathcal{V}_k \mapsto \mathcal{W}$$

is *multilinear*, or k -linear, (*bilinear*—if $k = 2$) if $\psi(v_1, \dots, v_k)$ is linear in each entry v_j when the other entries are held fixed.

When all the \mathcal{V}_j 's are equal to some fixed \mathcal{V} we say that ψ is k -linear on \mathcal{V} . If \mathcal{W} is the underlying field \mathbb{F} , we refer to ψ as a k -linear form or just k -form.

EXAMPLES:

- a. Multiplication in an algebra, e.g., $(S, T) \mapsto ST$ in $\mathcal{L}\mathcal{V}$ or $(A, B) \mapsto AB$ in $\mathcal{M}(n)$.
- b. $\psi(v, v^*) = (v, v^*)$, the value of a linear functional $v^* \in \mathcal{V}^*$ on a vector $v \in \mathcal{V}$, is a bilinear form on $\mathcal{V} \times \mathcal{V}^*$.
- c. Given k linear functionals $v_j^* \in \mathcal{V}^*$, the product $\psi(v_1, \dots, v_k) = \prod (v_j, v_j^*)$ of is a k -form on \mathcal{V} .
- d. Let $\mathcal{V}_1 = \mathbb{F}[x]$ and $\mathcal{V}_2 = \mathbb{F}[y]$ the map $(p(x), q(y)) \mapsto p(x)q(y)$ is a bilinear map from $\mathbb{F}[x] \times \mathbb{F}[y]$ onto the space $\mathbb{F}[x, y]$ of polynomials in the two variables.

★ 4.2.1 The definition of the tensor product $\mathcal{V}_1 \otimes \mathcal{V}_2$, see 1.1.6, guarantees that the map

$$(4.2.2) \quad \Psi(v, u) = v \otimes u.$$

of $\mathcal{V}_1 \times \mathcal{V}_2$ into $\mathcal{V}_1 \otimes \mathcal{V}_2$ is bilinear. It is special in that every bilinear map from $(\mathcal{V}_1, \mathcal{V}_2)$ “factors through it”:

Theorem. *Let φ be a bilinear map from $(\mathcal{V}_1, \mathcal{V}_2)$ into \mathcal{W} . Then there is a linear map $\Phi: \mathcal{V}_1 \otimes \mathcal{V}_2 \rightarrow \mathcal{W}$ such that $\varphi = \Phi\Psi$.*

The proof consists in checking that, for $v_j \in \mathcal{V}_1$ and $u_j \in \mathcal{V}_2$,

$$\sum v_j \otimes u_j = 0 \implies \sum \varphi(v_j, u_j) = 0$$

so that writing $\Phi(v \otimes u) = \varphi(v, u)$ defines Φ unambiguously, and checking that so defined, Φ is linear. We leave the checking to the reader.

★**4.2.2** Let \mathcal{V} and \mathcal{W} be finite dimensional vector spaces. Given $v^* \in \mathcal{V}^*$, $w \in \mathcal{W}$, and $v \in \mathcal{V}$, the map $v \mapsto (v, v^*)w$ is clearly a linear map from \mathcal{V} to \mathcal{W} (a linear functional on \mathcal{V} times a fixed vector in \mathcal{W}) and we denote it (temporarily) by $\underline{v^* \otimes w}$.

Theorem. *The map $\Phi : v^* \otimes w \mapsto \underline{v^* \otimes w} \in \mathcal{L}(\mathcal{V}, \mathcal{W})$ extends by linearity to an isomorphism of $\mathcal{V}^* \otimes \mathcal{W}$ onto $\mathcal{L}(\mathcal{V}, \mathcal{W})$.*

PROOF: As in ★4.2.1 we verify that all the representations of zero in the tensor product are mapped to 0, so that we do have a linear extension.

Let $T \in \mathcal{L}(\mathcal{V}, \mathcal{W})$, $\mathbf{v} = \{v_j\}$ a basis for \mathcal{V} , and $\mathbf{v}^* = \{v_j^*\}$ the dual basis. Then, for $v \in \mathcal{V}$,

$$(4.2.3) \quad Tv = T\left(\sum (v, v_j^*)v_j\right) = \sum (v, v_j^*)Tv_j = \left(\sum \underline{v_j^* \otimes Tv_j}\right)v,$$

so that $T = \sum \underline{v_j^* \otimes Tv_j}$. This shows that Φ is surjective and, since the two spaces have the same dimension, a linear map of one *onto* the other is an isomorphism. ◀

When there is no room for confusion we omit the underlining and write the operator as $v^* \otimes w$ instead of $\underline{v^* \otimes w}$.

EXERCISES FOR SECTION 4.2

IV.2.1. Assume $\varphi(v, u)$ bilinear on $\mathcal{V}_1 \times \mathcal{V}_2$. Prove that the map $T: u \mapsto \varphi_u(v)$ is a linear map from \mathcal{V}_2 into (the dual space) \mathcal{V}_1^* . Similarly, $S: v \mapsto {}_v\varphi(u)$ is linear from \mathcal{V}_1 to \mathcal{V}_2^* .

IV.2.2. Let \mathcal{V}_1 and \mathcal{V}_2 be finite dimensional, with bases $\{v_1, \dots, v_m\}$ and $\{u_1, \dots, u_n\}$ respectively. Show that every bilinear form φ on $(\mathcal{V}_1, \mathcal{V}_2)$ is given by an $m \times n$ matrix (a_{jk}) such that if $v = \sum_1^m x_j v_j$ and $u = \sum_1^n y_k u_k$ then

$$(4.2.4) \quad \varphi(v, u) = \sum a_{jk} x_j y_k = (x_1, \dots, x_m) \begin{bmatrix} a_{11} & \dots & a_{1n} \\ \vdots & \dots & \vdots \\ a_{m1} & \dots & a_{mn} \end{bmatrix} \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}$$

IV.2.3. What is the relation between the matrix in **IV.2.2** and the maps S and T defined in **IV.2.1**?

IV.2.4. Let \mathcal{V}_1 and \mathcal{V}_2 be finite dimensional, with bases $\{v_1, \dots, v_m\}$ and $\{u_1, \dots, u_n\}$ respectively, and let $\{v_1^*, \dots, v_m^*\}$ be the dual basis of $\{v_1, \dots, v_m\}$. Let $T \in \mathcal{L}(\mathcal{V}_1, \mathcal{V}_2)$ and let

$$A_T = \begin{bmatrix} a_{11} & \dots & a_{1m} \\ \vdots & \dots & \vdots \\ a_{n1} & \dots & a_{nm} \end{bmatrix}$$

be its matrix relative to the given bases. Prove

$$(4.2.5) \quad T = \sum a_{ij}(v_j^* \otimes u_i).$$

4.2.3 If Ψ and Φ are k -linear maps of $\mathcal{V}_1 \times \mathcal{V}_2 \cdots \times \mathcal{V}_k$ into \mathcal{W} and $a, b \in \mathbb{F}$ then $a\Psi + b\Phi$ is k -linear. Thus, the k -linear maps of $\mathcal{V}_1 \times \mathcal{V}_2 \cdots \times \mathcal{V}_k$ into \mathcal{W} form a vector space which we denote by $\mathcal{ML}(\{\mathcal{V}_j\}_{j=1}^k, \mathcal{W})$.

When all the \mathcal{V}_j are the same space \mathcal{V} , the notation is: $\mathcal{ML}(\mathcal{V}^{\oplus k}, \mathcal{W})$.

The reference to \mathcal{W} is omitted when $\mathcal{W} = \mathbb{F}$.

4.2.4 Example **b.** above identifies enough k -linear forms

4.3 ALTERNATING N -FORMS

4.3.1 DEFINITION: An n -linear form $\varphi(v_1, \dots, v_n)$ on \mathcal{V} is *alternating* if $\varphi(v_1, \dots, v_n) = 0$ whenever one of the entry vectors is repeated, i.e., if $v_k = v_l$ for some $k \neq l$.

If φ is *alternating*, and $k \neq l$ then

$$(4.3.1) \quad \begin{aligned} \varphi(\dots, v_k, \dots, v_l, \dots) &= \varphi(\dots, v_k, \dots, v_l + v_k, \dots) \\ &= \varphi(\dots, -v_l, \dots, v_l + v_k, \dots) = \varphi(\dots, -v_l, \dots, v_k, \dots) \\ &= -\varphi(\dots, v_l, \dots, v_k, \dots), \end{aligned}$$

which proves that a transposition (k, l) on the entries of φ changes its sign. It follows that for any permutation $\sigma \in \mathbf{S}_n$

$$(4.3.2) \quad \varphi(v_{\sigma(1)}, \dots, v_{\sigma(n)}) = \mathbf{sgn}[\sigma]\varphi(v_1, \dots, v_n).$$

Condition (4.3.2) explains the term *alternating* and when the characteristic of \mathbb{F} is $\neq 2$, can be taken as the definition.

If φ is alternating, and if one of the entry vectors is a linear combination of the others, we use the linearity of φ in that entry and write $\varphi(v_1, \dots, v_n)$ as a linear combination of φ evaluated on several n -tuples each of which has a repeated entry. Thus, if $\{v_1, \dots, v_n\}$ is linearly dependent, $\varphi(v_1, \dots, v_n) = 0$. It follows that if $\dim \mathcal{V} < n$, there are no nontrivial alternating n -forms on \mathcal{V} .

Theorem. *Assume $\dim \mathcal{V} = n$. The space of alternating n -forms on \mathcal{V} is one dimensional: there exists one and, up to scalar multiplication, unique non-trivial alternating n -form D on \mathcal{V} . $D(v_1, \dots, v_n) \neq 0$ if, and only if, $\{v_1, \dots, v_n\}$ is a basis.*

PROOF: We show first that if φ is an alternating n -form, it is completely determined by its value on any given basis of \mathcal{V} . This will show that any two alternating n -forms are proportional, and the proof will also make it clear how to define a non-trivial alternating n -form.

If $\{v_1, \dots, v_n\}$ is a basis for \mathcal{V} and φ an alternating n -form on \mathcal{V} , then $\varphi(v_{j_1}, \dots, v_{j_n}) = 0$ unless $\{j_1, \dots, j_n\}$ is a permutation, say σ , of $\{1, \dots, n\}$, and then $\varphi(v_{\sigma(1)}, \dots, v_{\sigma(n)}) = \mathbf{sgn}[\sigma]\varphi(v_1, \dots, v_n)$.

If $\{u_1, \dots, u_n\}$ is an arbitrary n -tuple, we express each u_i in terms of the basis $\{v_1, \dots, v_n\}$:

$$(4.3.3) \quad u_j = \sum_{i=1}^n a_{i,j} v_i, \quad j = 1, \dots, n$$

and the multilinearity implies

$$(4.3.4) \quad \begin{aligned} \varphi(u_1, \dots, u_n) &= \sum a_{1,j_1} \cdots a_{n,j_n} \varphi(v_{j_1}, \dots, v_{j_n}) \\ &= \left(\sum_{\sigma \in \mathbf{S}_n} \mathbf{sgn}[\sigma] a_{1,\sigma(1)} \cdots a_{n,\sigma(n)} \right) \varphi(v_1, \dots, v_n). \end{aligned}$$

This shows that $\varphi(v_1, \dots, v_n)$ determines $\varphi(u_1, \dots, u_n)$ for all n -tuples, and all alternating n -forms are proportional. This also shows that unless φ is trivial, $\varphi(v_1, \dots, v_n) \neq 0$ for every independent (i.e., basis) $\{v_1, \dots, v_n\}$.

For the existence we fix a basis $\{v_1, \dots, v_n\}$ and set $D(v_1, \dots, v_n) = 1$. Write $D(v_{\sigma(1)}, \dots, v_{\sigma(n)}) = \mathbf{sgn}[\sigma]$ (for $\sigma \in \mathbf{S}_n$) and $D(v_{j_1}, \dots, v_{j_n}) = 0$ if there is a repeated entry.

For arbitrary n -tuple $\{u_1, \dots, u_n\}$ define $D(u_1, \dots, u_n)$ by (4.3.4), that is

$$(4.3.5) \quad D(u_1, \dots, u_n) = \sum_{\sigma \in \mathfrak{S}_n} \mathbf{sgn} [\sigma] a_{1, \sigma(1)} \cdots a_{n, \sigma(n)}.$$

The fact that D is n -linear is clear: it is defined by multilinear expansion. To check that it is alternating take $\tau \in \mathfrak{S}_n$ and write

$$(4.3.6) \quad \begin{aligned} D(u_{\tau(1)}, \dots, u_{\tau(n)}) &= \sum_{\sigma \in \mathfrak{S}_n} \mathbf{sgn} [\sigma] a_{\tau(1), \sigma(1)} \cdots a_{\tau(n), \sigma(n)} \\ &= \sum_{\sigma \in \mathfrak{S}_n} \mathbf{sgn} [\sigma] a_{1, \tau^{-1}\sigma(1)} \cdots a_{n, \tau^{-1}\sigma(n)} = \mathbf{sgn} [\tau] D(u_1, \dots, u_n) \end{aligned}$$

since $\mathbf{sgn} [\tau^{-1}\sigma] = \mathbf{sgn} [\tau] \mathbf{sgn} [\sigma]$. ◀

Observe that if $\{u_1, \dots, u_n\}$ is given by (4.3.3) then $\{Tu_1, \dots, Tu_n\}$ is given by

$$(4.3.7) \quad Tu_j = \sum_{i=1}^n a_{i,j} Tv_i, \quad j = 1, \dots, n$$

and (4.3.4) implies

$$(4.3.8) \quad D(Tu_1, \dots, Tu_n) = \frac{D(u_1, \dots, u_n)}{D(v_1, \dots, v_n)} D(Tv_1, \dots, Tv_n)$$

4.4 DETERMINANT OF AN OPERATOR

4.4.1 DEFINITION: The *determinant* $\det T$ of an operator $T \in \mathcal{L}(\mathcal{V})$ is

$$(4.4.1) \quad \det T = \frac{D(Tv_1, \dots, Tv_n)}{D(v_1, \dots, v_n)}$$

where $\{v_1, \dots, v_n\}$ is an arbitrary basis of \mathcal{V} and D is a non-trivial alternating n -form. The independence of $\det T$ from the choice of the basis is guaranteed by (4.3.8).

Proposition. $\det T = 0$ if, and only if, T is singular, (i.e., $\ker(T) \neq \{0\}$).

PROOF: T is singular if, and only if, it maps a basis onto a linearly dependent set. $D(Tv_1, \dots, Tv_n) = 0$ if, and only if, $\{Tv_1, \dots, Tv_n\}$ is linearly dependent. ◀

4.4.2 Proposition. *If $T, S \in \mathcal{L}(\mathcal{V})$ then*

$$(4.4.2) \quad \det TS = \det T \det S.$$

PROOF: If either S or T is singular both sides of (4.4.4) are zero. If $\det S \neq 0$, $\{Sv_j\}$ is a basis, and by (4.4.1),

$$\det TS = \frac{D(TSv_1, \dots, TSv_n)}{D(Sv_1, \dots, Sv_n)} \cdot \frac{D(Sv_1, \dots, Sv_n)}{D(v_1, \dots, v_n)} = \det T \det S. \quad \blacktriangleleft$$

★ **4.4.3 ORIENTATION.** When \mathcal{V} is a real vector space, a non-trivial alternating n -form D determines an equivalence relation among bases. The bases $\{v_j\}$ and $\{u_j\}$ are declared equivalent if $D(v_1, \dots, v_n)$ and $D(u_1, \dots, u_n)$ have the same sign. Using $-D$ instead of D reverses the signs of all the readings, but maintains the equivalence. An *orientation* on \mathcal{V} is a choice which of the two equivalence classes to call *positive*.

4.4.4 A subspace $\mathcal{W} \subset \mathcal{V}$ is T -invariant, ($T \in \mathcal{L}(\mathcal{V})$), if $Tw \in \mathcal{W}$ whenever $w \in \mathcal{W}$. The *restriction* $T_{\mathcal{W}}$, defined by $w \mapsto Tw$ for $w \in \mathcal{W}$, is clearly a linear operator on \mathcal{W} .

T induces also an operator $T_{\mathcal{V}/\mathcal{W}}$ on the quotient space \mathcal{V}/\mathcal{W} , see 5.1.5.

Proposition. *If $\mathcal{W} \subset \mathcal{V}$ is T -invariant, then*

$$(4.4.3) \quad \det T = \det T_{\mathcal{W}} \det T_{\mathcal{V}/\mathcal{W}}.$$

PROOF: Let $\{w_j\}_1^n$ be a basis for \mathcal{V} , such that $\{w_j\}_1^k$ is a basis for \mathcal{W} . If $T_{\mathcal{W}}$ is singular then T is singular and both sides of (4.4.3) are zero.

If $T_{\mathcal{W}}$ is nonsingular, then $\mathbf{w} = \{Tw_1, \dots, Tw_k\}$ is a basis for \mathcal{W} , and $\{Tw_1, \dots, Tw_k; w_{k+1}, \dots, w_n\}$ is a basis for \mathcal{V} .

Let D be a nontrivial alternating n -form on \mathcal{V} . Then $\Phi(u_1, \dots, u_k) = D(u_1, \dots, u_k; w_{k+1}, \dots, w_n)$ is a nontrivial alternating k -form on \mathcal{W} .

The value of $D(Tw_1, \dots, Tw_k; u_{k+1}, \dots, u_n)$ is unchanged if we replace the variables u_{k+1}, \dots, u_n by ones that are congruent to them mod \mathcal{W} , and the form $\Psi(\tilde{u}_{k+1}, \dots, \tilde{u}_n) = D(Tw_1, \dots, Tw_k; u_{k+1}, \dots, u_n)$ is therefore a well defined nontrivial alternating $n - k$ -form on \mathcal{V}/\mathcal{W} .

$$\begin{aligned} \det T &= \frac{D(Tw_1, \dots, Tw_n)}{D(w_1, \dots, w_n)} = \\ &= \frac{D(Tw_1, \dots, Tw_k; w_{k+1}, \dots, w_n)}{D(w_1, \dots, w_n)} \cdot \frac{D(Tw_1, \dots, Tw_n)}{D(Tw_1, \dots, Tw_k; w_{k+1}, \dots, w_n)} = \\ &= \frac{\Phi(Tw_1, \dots, Tw_k)}{\Phi(w_1, \dots, w_k)} \cdot \frac{\Psi(\widetilde{Tw}_{k+1}, \dots, \widetilde{Tw}_n)}{\Psi(\tilde{w}_{k+1}, \dots, \tilde{w}_n)} = \det T_{\mathcal{W}} \det T_{\mathcal{V}/\mathcal{W}}. \end{aligned}$$

◀

Corollary. If $\mathcal{V} = \bigoplus \mathcal{V}_j$ and all the \mathcal{V}_j 's are T -invariant, and $T_{\mathcal{V}_j}$ denotes the restriction of T to \mathcal{V}_j , then

$$(4.4.4) \quad \det T = \prod_j \det T_{\mathcal{V}_j}.$$

4.4.5 THE CHARACTERISTIC POLYNOMIAL OF AN OPERATOR.

DEFINITIONS: The *characteristic polynomial* of an operator $T \in \mathcal{L}(\mathcal{V})$ is the polynomial $\chi_T(\lambda) = \det(T - \lambda) \in \mathbb{F}[\lambda]$.

Opening up the expression $D(Tv_1 - \lambda v_1, \dots, Tv_n - \lambda v_n)$, we see that χ_T is a polynomial of degree $n = \dim \mathcal{V}$, with leading coefficient $(-1)^n$.

By proposition 4.4.1, $\chi_T(\lambda) = 0$ if, and only if, $T - \lambda$ is singular, that is if, and only if, $\ker(T - \lambda) \neq \{0\}$. The zeroes of χ_T are called *eigenvalues* of T and the set of eigenvalues of T is called the *spectrum* of T , and denoted $\sigma(T)$.

For $\lambda \in \sigma(T)$, (the nontrivial) $\ker(T - \lambda)$ is called *the eigenspace* of λ . The non-zero vectors $v \in \ker(T - \lambda)$ (that is the vectors $v \neq 0$ such that $Tv = \lambda v$) are *the eigenvectors of T* corresponding to the eigenvalue λ .

EXERCISES FOR SECTION 4.4

IV.4.1. Prove that if T is non-singular, then $\det T^{-1} = (\det T)^{-1}$

IV.4.2. If $\mathcal{W} \subset \mathcal{V}$ is T -invariant, then $\chi_T(\lambda) = \chi_{T_{\mathcal{W}}} \chi_{T_{\mathcal{V}/\mathcal{W}}}$.

4.5 DETERMINANT OF A MATRIX

4.5.1 Let $A = \{a_{ij}\} \in \mathcal{M}(n)$. The determinant of A can be defined in several equivalent ways: the first—as the determinant of the operator that A defines on \mathbb{F}^n by matrix multiplication; another, the standard definition, is directly by the following formula, motivated by (4.3.5):

$$(4.5.1) \quad \det A = \begin{vmatrix} a_{11} & \cdots & a_{1n} \\ a_{21} & \cdots & a_{2n} \\ \vdots & \cdots & \vdots \\ a_{n1} & \cdots & a_{nn} \end{vmatrix} = \sum_{\sigma \in \mathbf{S}_n} \mathbf{sgn}[\sigma] a_{1,\sigma(1)} \cdots a_{n,\sigma(n)}.$$

The reader should check that the two ways are in fact equivalent. They each have advantages. The first definition, in particular, makes it transparent that $\det(AB) = \det A \det B$; the second is sometimes readier for computation.

4.5.2 COFACTORS, EXPANSIONS, AND INVERSES. For a fixed pair (i, j) the elements in the sum above that have a_{ij} as a factor are those for which $\sigma(i) = j$ their sum is

$$(4.5.2) \quad \sum_{\sigma \in \mathbf{S}_n, \sigma(i)=j} \mathbf{sgn}[\sigma] a_{1,\sigma(1)} \cdots a_{n,\sigma(n)} = a_{ij} A_{ij}.$$

The sum, with the factor a_{ij} removed, denoted A_{ij} in (4.5.2), is called the *cofactor* at (i, j) .

Observe that partitioning the sum in (4.5.1) according to the value $\sigma(i)$ for some fixed i gives *the expansion of the determinant along its i 'th row*:

$$(4.5.3) \quad \det A = \sum_j a_{ij} A_{ij}.$$

If we consider a “mismatched” sum: $\sum_j a_{ij} A_{kj}$ for $i \neq k$, we obtain the determinant of the matrix obtained from A by replacing the k 'th row by the i 'th. Since this matrix has two identical rows, its determinant is zero, that is

$$(4.5.4) \quad \text{for } i \neq k, \quad \sum_j a_{ij} A_{kj} = 0.$$

Finally, write $\tilde{A} = \begin{bmatrix} A_{11} & \cdots & A_{n1} \\ A_{12} & \cdots & A_{n2} \\ \vdots & \cdots & \vdots \\ A_{1n} & \cdots & A_{nn} \end{bmatrix}$ and observe that $\sum_j a_{ij}A_{kj}$ is the

ik 'th entry of the matrix $A\tilde{A}$ so that equations (4.5.3) and (4.5.4) combined are equivalent to

$$(4.5.5) \quad A\tilde{A} = \det A I.$$

Proposition. *The inverse of a non-singular matrix $A \in \mathcal{M}(n)$ is $\frac{1}{\det(A)}\tilde{A}$.*

Historically, the matrix \tilde{A} was called the *adjoint* of A , but the term *adjoint* is now used mostly in the context of duality.

4.5.3 THE CHARACTERISTIC POLYNOMIAL OF A MATRIX.

The *characteristic polynomial* of a matrix $A \in \mathcal{M}(n)$ is the polynomial $\chi_A(\lambda) = \det(A - \lambda)$.

Proposition. *If $A, B \in \mathcal{M}(n)$ are similar then they have the same characteristic polynomial. In other words, χ_A is similarity invariant.*

PROOF: Similar matrices have the same determinant: they represent the same operator using different basis and the determinant of an operator is independent of the basis. Equivalently, if $B = CAC^{-1}$, then $\det B = \det(CAC^{-1}) = \det C \det A (\det C)^{-1} = \det A$.

Also, if $B = CAC^{-1}$, then $B - \lambda = C(A - \lambda)C^{-1}$, which implies $\det(B - \lambda) = \det(A - \lambda)$. ◀

The converse is not always true—matrices (or operators) that have the same characteristic polynomials may not be similar. See exercise **IV.5.2**.

If we write $\chi_A = \sum_0^n a_j \lambda^j$, then

$$a_n = (-1)^n, \quad a_0 = \det A, \quad \text{and} \quad a_{n-1} = (-1)^{n-1} \sum_1^n a_{ii}.$$

The sum $\sum_1^n a_{ii}$, denoted *trace* A , is called the *trace* of the matrix A . Like any part of χ_A , the trace is similarity invariant.

The trace is just one coefficient of the characteristic polynomial and is not a *complete* invariant. However, we shall see later that the traces of A^j for all $1 \leq j \leq n$ determine $\chi_A(\lambda)$ completely.

EXERCISES FOR SECTION 4.5

IV.5.1. A matrix $A = \{a_{ij}\} \in \mathcal{M}(n)$ is *upper triangular* if $a_{ij} = 0$ when $i > j$. A is *lower triangular* if $a_{ij} = 0$ when $i < j$. Prove that if A is either upper or lower triangular then $\det A = \prod_{i=1}^n a_{ii}$.

IV.5.2. Let $A \neq I$ be a lower triangular matrix with all the diagonal elements equal to 1. Prove that $\chi_A = \chi_I$ (I is the identity matrix); is A similar to I ?

IV.5.3. How can the algorithm of reduction to row echelon form be used to compute determinants?

IV.5.4. Let $A \in \mathcal{M}(n)$. A defines an operator on \mathbb{F}^n , as well as on $\mathcal{M}(n)$, both by matrix multiplication. What is the relation between the values of $\det A$ as operator in the two cases?

IV.5.5. Prove the following properties of the trace:

1. If $A, B \in \mathcal{M}(n)$, then $\text{trace}(A + B) = \text{trace } A + \text{trace } B$.
2. If $A \in \mathcal{M}(m, n)$ and $B \in \mathcal{M}(n, m)$, then $\text{trace } AB = \text{trace } BA$.

IV.5.6. If $A, B \in \mathcal{M}(2)$, then $(AB - BA)^2 = -\det(AB - BA)I$.

IV.5.7. Prove that the characteristic polynomial of the $n \times n$ matrix $A = (a_{i,j})$ is equal to $\prod_{i=1}^n (a_{i,i} - \lambda)$ plus a polynomial of degree bounded by $n - 2$.

IV.5.8. Assuming $\mathbb{F} = \mathbb{C}$, prove that $\text{trace}(a_{i,j})$ is equal to the sum (including multiplicity) of the zeros of the characteristic polynomial of $(a_{i,j})$. In other words, if the characteristic polynomial of $(a_{i,j})$ is equal to $\prod_{j=1}^n (\lambda - \lambda_j)$, then $\sum \lambda_j = \sum a_{i,i}$.

IV.5.9. Let $A = (a_{i,j}) \in \mathcal{M}(n)$ and let $m > n/2$. Assume that $a_{i,j} = 0$ whenever both $i \leq m$ and $j \leq m$. Prove that $\det(A) = 0$.

IV.5.10. The *Fibonacci sequence* is the sequence $\{f_n\}$ defined inductively by: $f_1 = 1, f_2 = 1$, and $f_n = f_{n-1} + f_{n-2}$ for $n \geq 3$, so that the start of the sequence is $1, 1, 2, 3, 5, 8, 13, 21, 34, \dots$

Let $(a_{i,j})$ be an $n \times n$ matrix such that $a_{i,j} = 0$ when $|j - i| > 1$ (that is the only non-zero elements are on the diagonal, just above it, or just below it). Prove that the

number of non-zero terms in the expansion of the determinant of $(a_{i,j})$ is at most equal to f_{n+1} .

IV.5.11. *The Vandermonde determinant.* Given scalars a_j , $j = 1, \dots, n$, the Vandermonde determinant $V(a_1, \dots, a_n)$ is defined by

$$V(a_1, \dots, a_n) = \begin{vmatrix} 1 & a_1 & a_1^2 & \dots & a_1^{n-1} \\ 1 & a_2 & a_2^2 & \dots & a_2^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & a_n & a_n^2 & \dots & a_n^{n-1} \end{vmatrix}$$

Use the following steps to compute $V(a_1, \dots, a_n)$. Observe that

$$V(a_1, \dots, a_n, x) = \begin{vmatrix} 1 & a_1 & a_1^2 & \dots & a_1^n \\ 1 & a_2 & a_2^2 & \dots & a_2^n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & a_n & a_n^2 & \dots & a_n^n \\ 1 & x & x^2 & \dots & x^n \end{vmatrix}$$

is a polynomial of degree n (in x).

a. Prove that

$$V(a_1, \dots, a_n, x) = V(a_1, \dots, a_n) \prod_{j=1}^n (x - a_j)$$

b. Use induction to prove $V(a_1, \dots, a_n) = \prod_{i < j} (a_j - a_i)$.

IV.5.12. A trigonometric polynomial $P(x) = \sum_{j=1}^m a_j e^{i\alpha_j x}$ that has a zero of order m (a point x_0 such that $P^{(j)}(x_0) = 0$ for $j = 0, \dots, m-1$) is identically zero.

IV.5.13. Let $C \in \mathcal{M}(n, \mathbb{C})$ be non-singular. Let $\Re C$, resp. $\Im C$, be the matrix whose entries are the real parts, resp. the imaginary parts, of the corresponding entries in C . Prove that for all but a finite number of values of $a \in \mathbb{R}$, the matrix $\Re C + a\Im C$ is non-singular.

Hint: Show that replacing a single column in C by the corresponding column in $\Re C + a\Im C$ creates a non-singular matrix for all but one value of a . (The determinant is a non-trivial linear function of a .)

IV.5.14. Given that the matrices $B_1, B_2 \in \mathcal{M}(n; \mathbb{R})$ are similar in $\mathcal{M}(n; \mathbb{C})$, show that they are similar in $\mathcal{M}(n; \mathbb{R})$.

Invariant subspaces

The study of linear operators on a fixed vector space \mathcal{V} (as opposed to linear maps between different spaces) takes full advantage of the fact that $\mathcal{L}(\mathcal{V})$ is an algebra. Polynomials in T play an important role in the understanding of T itself. In particular they provide a way to decompose \mathcal{V} into a direct sum of T -invariant subspaces (see below) on each of which the behaviour of T is relatively simple.

Studying the behavior of T on various subspaces justifies the following definition.

DEFINITION: A *linear system*, or simply a *system*, is a pair (\mathcal{V}, T) where \mathcal{V} is a vector space and $T \in \mathcal{L}(\mathcal{V})$. When we add adjectives they apply in the appropriate place, so that a *finite dimensional system* is a system in which \mathcal{V} is finite dimensional, while an *invertible system* is one in which T is invertible.

5.1 INVARIANT SUBSPACES

5.1.1 Let (\mathcal{V}, T) be a linear system.

DEFINITION: A subspace $\mathcal{V}_1 \subset \mathcal{V}$ is T -invariant if $T\mathcal{V}_1 \subseteq \mathcal{V}_1$. If \mathcal{V}_1 is T -invariant and $v \in \mathcal{V}_1$, then $T^j v \in \mathcal{V}_1$ for all j , and in fact $P(T)v \in \mathcal{V}_1$ for every polynomial P . Thus, \mathcal{V}_1 is $P(T)$ -invariant for all $P \in \mathbb{F}[x]$.

EXAMPLES:

a. Both $\ker(T)$ and $\text{range}(T)$ are (clearly) T -invariant.

b. If $S \in \mathcal{L}(\mathcal{V})$ and $ST = TS$, then $\ker(S)$ and $\text{range}(S)$ are T -invariant since if $Sv = 0$ then $STv = TSv = 0$, and $T S \mathcal{V} = S(T\mathcal{V}) \subset S\mathcal{V}$. In particular, if P is a polynomial then $\ker(P(T))$ and $\text{range}(P(T))$ are T -invariant.

c. Given $v \in \mathcal{V}$, the set $\text{span}[T, v] = \{P(T)v : P \in \mathbb{F}[x]\}$ is clearly a subspace, clearly T -invariant, and clearly the smallest T -invariant subspace containing v .

5.1.2 Recall (see 4.4.5) that $\lambda \in \mathbb{F}$ is an *eigenvalue* of T if $\ker(T - \lambda)$ is nontrivial, i.e., if there exists vectors $v \neq 0$ such that $Tv = \lambda v$ (called *eigenvectors* “associated with”, or “corresponding to” λ). Eigenvectors provide the simplest—namely, one dimensional— T -invariant subspaces.

The *spectrum* $\sigma(T)$ is the set of all the eigenvalues of T . It is (see 4.4.5) the set of zeros of the characteristic polynomial $\chi_T(\lambda) = \det(T - \lambda)$. If the underlying field \mathbb{F} is algebraically closed every non-constant polynomial has zeros in \mathbb{F} and every $T \in \mathcal{L}(\mathcal{V})$ has non-empty spectrum.

Proposition (Spectral Mapping theorem). *Let $T \in \mathcal{L}(\mathcal{V})$, $\lambda \in \sigma(T)$, and $P \in \mathbb{F}[x]$. Then*

a. $P(\lambda) \in \sigma(P(T))$.

b. For all $k \in \mathbb{N}$,

$$(5.1.1) \quad \ker((P(T) - P(\lambda))^k) \supset \ker((T - \lambda)^k).$$

c. If \mathbb{F} is algebraically closed, then $\sigma(P(T)) = P(\sigma(T))$.

PROOF: a. $(P(x) - P(\lambda))$ is divisible by $x - \lambda$: $(P(x) - P(\lambda)) = Q(x)(x - \lambda)$, and $(P(T) - P(\lambda)) = Q(T)(T - \lambda)$ is not invertible.

b. $(P(x) - P(\lambda)) = Q(x)(x - \lambda)$ implies: $(P(x) - P(\lambda))^k = Q^k(x - \lambda)^k$, and $(P(T) - P(\lambda))^k = Q^k(T)(T - \lambda)^k$. If $v \in \ker((T - \lambda)^k)$, i.e., $(T - \lambda)^k v = 0$, then $(P(T) - P(\lambda))^k v = Q^k(T)(T - \lambda)^k v = 0$.

c. If \mathbb{F} is algebraically closed and $\mu \in \mathbb{F}$, denote by $c_j(\mu)$ the roots of $P(x) - \mu$, and by m_j their multiplicities, so that

$$P(x) - \mu = \prod (x - c_j(\mu))^{m_j}, \quad \text{and} \quad P(T) - \mu = \prod (T - c_j(\mu))^{m_j}.$$

Unless $c_j(\mu) \in \sigma(T)$ for some j , all the factors are invertible, and so is their product. \blacktriangleleft

Remark: If \mathbb{F} is not algebraically closed, $\sigma(P(T))$ may be strictly bigger than $P(\sigma(T))$. For example, if $\mathbb{F} = \mathbb{R}$, T is a rotation by $\pi/2$ on \mathbb{R}^2 , and $P(x) = x^2$, then $\sigma(T) = \emptyset$ while $\sigma(T^2) = \{-1\}$.

5.1.3 T -invariant subspaces are $P(T)$ -invariant for all polynomials P . Notice, however, that a subspace \mathcal{W} can be T^2 -invariant, and not be T -invariant. Example: $\mathcal{V} = \mathbb{R}^2$ and T maps (x, y) to (y, x) . $T^2 = I$, the identity, so that everything is T^2 -invariant. But only the diagonal $\{(x, x) : x \in \mathbb{R}\}$ is T -invariant.

Assume that $T, S \in \mathcal{L}(\mathcal{V})$ commute.

a. T commutes with $P(S)$ for every polynomial P ; consequently (see 5.1.1 **b.**) $\ker(P(S))$ and $\text{range}(P(S))$ are T -invariant. In particular, for every $\lambda \in \mathbb{F}$, $\ker(S - \lambda)$ is T -invariant.

b. If \mathcal{W} is a S -invariant subspace, then $T\mathcal{W}$ is S -invariant. This follows from:

$$ST\mathcal{W} = TSW \subset T\mathcal{W}.$$

There is no claim that \mathcal{W} is T -invariant¹. Thus, kernels offer “a special situation.”

c. If v is an eigenvector for S with eigenvalue λ , it is contained in $\ker(S - \lambda)$ which is T invariant. If $\ker(S - \lambda)$ is one dimensional, then v is an eigenvector for T .

5.1.4 Theorem. Let $\mathcal{W} \subset \mathcal{V}$, and $T \in \mathcal{L}(\mathcal{V})$. The following statements are equivalent:

- a.** \mathcal{W} is T -invariant;
- b.** \mathcal{W}^\perp is T^* -invariant.

PROOF: For all $w \in \mathcal{W}$ and $u^* \in \mathcal{W}^\perp$ we have

$$(Tw, u^*) = (w, T^*u^*).$$

Statement **a.** is equivalent to the left-hand side being identically zero; statement **b.** to the vanishing of the right-hand side. ◀

¹An obvious example is $S = I$, which commutes with every operator T , and for which all subspaces are invariant.

5.1.5 If $\mathcal{W} \subset \mathcal{V}$ is a T -invariant subspace, we define *the restriction* $T_{\mathcal{W}}$ of T to \mathcal{W} by $T_{\mathcal{W}}v = Tv$ for $v \in \mathcal{W}$. The operator $T_{\mathcal{W}}$ is clearly linear on \mathcal{W} , and every $T_{\mathcal{W}}$ -invariant subspace $\mathcal{W}_1 \subset \mathcal{W}$ is T -invariant.

Similarly, if \mathcal{W} is T -invariant, T induces a linear operator $T_{\mathcal{V}/\mathcal{W}}$ on the quotient \mathcal{V}/\mathcal{W} as follows:

$$(5.1.2) \quad T_{\mathcal{V}/\mathcal{W}}(v + \mathcal{W}) = Tv + \mathcal{W}.$$

$v + \mathcal{W}$ is the coset of \mathcal{W} containing v and, we justify the definition by showing that it is independent of the choice of the representative: if $v_1 - v \in \mathcal{W}$ then, by the T -invariance of \mathcal{W} , $Tv_1 - Tv = T(v_1 - v) \in \mathcal{W}$.

The reader should check that $T_{\mathcal{V}/\mathcal{W}}$ is in fact linear.

5.1.6 The fact that when \mathbb{F} algebraically closed, every operator $T \in \mathcal{L}(\mathcal{V})$ has eigenvectors, applies equally to (\mathcal{V}^*, T^*) .

If \mathcal{V} is n -dimensional and $u^* \in \mathcal{V}^*$ is an eigenvector for T^* , then $\mathcal{V}_{n-1} = [u^*]^\perp = \{v \in \mathcal{V} : (v, u^*) = 0\}$ is T invariant and $\dim \mathcal{V}_{n-1} = n - 1$.

Repeating the argument in \mathcal{V}_{n-1} we find a T -invariant $\mathcal{V}_{n-2} \subset \mathcal{V}_{n-1}$ of dimension $n - 2$, and repeating the argument a total of $n - 1$ times we obtain:

Theorem. *Assume that \mathbb{F} is algebraically closed, and let \mathcal{V} be a finite dimensional vector space over \mathbb{F} . For any $T \in \mathcal{L}(\mathcal{V})$, there exist a ladder² $\{\mathcal{V}_j\}$, $j = 0, \dots, n$, of T -invariant subspaces of $\mathcal{V}_n = \mathcal{V}$, such that*

$$(5.1.3) \quad \mathcal{V}_0 = \{0\}, \quad \mathcal{V}_n = \mathcal{V}; \quad \mathcal{V}_{j-1} \subset \mathcal{V}_j, \quad \text{and} \quad \dim \mathcal{V}_j = j.$$

Corollary. *If \mathbb{F} is algebraically closed, then every matrix $A \in \mathcal{M}(n; \mathbb{F})$ is similar to an upper triangular matrix.*

PROOF: Apply the theorem to the operator T of left multiplication by A on \mathbb{F}_c^n . Choose v_j in $\mathcal{V}_j \setminus \mathcal{V}_{j-1}$, $j = 1, \dots, n$, then $\{v_1, \dots, v_n\}$ is a basis for \mathcal{V} and the matrix B corresponding to T in this basis is (upper) triangular.

The matrices A and B represent the same operator relative to two bases, hence are similar. ◀

²Also called a *complete flag*.

Observe that if the underlying field is \mathbb{R} , which is not algebraically closed, and T is a rotation by $\pi/2$ on \mathbb{R}^2 , T admits *no* invariant subspaces.

EXERCISES FOR SECTION 5.1

- V.1.1.** Let \mathcal{W} be T -invariant, P a polynomial. Prove that $P(T)_{\mathcal{W}} = P(T_{\mathcal{W}})$.
- V.1.2.** Let \mathcal{W} be T -invariant, P a polynomial. Prove that $P(T)_{\mathcal{V}/\mathcal{W}} = P(T_{\mathcal{V}/\mathcal{W}})$.
- V.1.3.** Let \mathcal{W} be T -invariant. Prove that $\ker(T_{\mathcal{W}}) = \ker(T) \cap \mathcal{W}$.
- V.1.4.** Prove that every upper triangular matrix is similar to a lower triangular one (and vice versa).
- V.1.5.** If $\mathcal{V}_1 \subset \mathcal{V}$ is a subspace, then the set $\{S : S \in \mathcal{L}(\mathcal{V}), S\mathcal{V}_1 \subset \mathcal{V}_1\}$ is a subalgebra of $\mathcal{L}(\mathcal{V})$.
- V.1.6.** Show that if S and T commute and v is an eigenvector for S , it need not be an eigenvector for T (so that the assumption in the final remark of 5.1.3 that $\ker(S - \lambda)$ is one dimensional is crucial).
- V.1.7.** Prove theorem 5.1.6 without using duality.
Hint: Start with an eigenvector u_1 of T . Set $\mathcal{U}_1 = \text{span}[u_1]$; Let $\tilde{u}_2 \in \mathcal{V}/\mathcal{U}_1$ be an eigenvector of $T_{\mathcal{V}/\mathcal{U}_1}$, $u_2 \in \mathcal{V}$ a representative of \tilde{u}_2 , and $\mathcal{U}_2 = \text{span}[u_1, u_2]$. Verify that \mathcal{U}_2 is T -invariant. Let $\tilde{u}_3 \in \mathcal{V}/\mathcal{U}_2$ be an eigenvector of $T_{\mathcal{V}/\mathcal{U}_2}$, etc.

5.2 THE MINIMAL POLYNOMIAL

5.2.1 THE MINIMAL POLYNOMIAL FOR (T, v) .

Given $v \in \mathcal{V}$, let m be the first positive integer k such that $\{T^j v\}_0^k$ is linearly dependent or, equivalently, that $T^m v$ is a linear combination of $\{T^j v\}_0^{m-1}$, say³

$$(5.2.1) \quad T^m v = - \sum_0^{m-1} a_j T^j v.$$

The coefficients a_j are uniquely determined—this since, by assumption, $\{T^j v\}_0^{m-1}$ is independent. For $k > 0$ we have $T^{m+k} v = \sum_0^{m-1} a_j T^{j+k} v$, and induction on k establishes that $T^{m+k} v \in \text{span}[v, \dots, T^{m-1} v]$. It follows that $\{T^j v\}_0^{m-1}$ is a basis for $\text{span}[T, v]$.

³The minus sign is there to give the common notation: $\min P_{T,v}(x) = x^m + \sum_0^{m-1} a_j x^j$

DEFINITION: The polynomial $\min P_{T,v}(x) = x^m + \sum_0^{m-1} a_j x^j$, with a_j defined by (5.2.1) is called the *minimal polynomial* for (T, v) .

Theorem. $\min P_{T,v}(x)$ is the monic polynomial of lowest degree that satisfies $P(T)v = 0$.

The set $\mathfrak{N}_{T,v} = \{P \in \mathbb{F}[x] : P(T)v = 0\}$ is an ideal in $\mathbb{F}[x]$. The theorem identifies $\min P_{T,v}$ as its generator.

Observe that $P(T)v = 0$ is equivalent to “ $P(T)u = 0$ for all $u \in \text{span}[T, v]$ ”.

5.2.2 CYCLIC VECTORS. A vector $v \in \mathcal{V}$ is *cyclic* for the system (\mathcal{V}, T) if it is not contained in any T -invariant proper subspace, that is, if $\text{span}[T, v] = \mathcal{V}$. Not every linear system admits cyclic vectors⁴; a system that does is called a *cyclic system*.

If v is a cyclic vector for (\mathcal{V}, T) and $\min P_{T,v}(x) = x^n + \sum_0^{n-1} a_j x^j$, then the matrix of T with respect to the basis $\mathbf{v} = \{v, Tv, \dots, T^{n-1}v\}$ has the form

$$(5.2.2) \quad A_{T,\mathbf{v}} = \begin{bmatrix} 0 & 0 & \dots & 0 & -a_0 \\ 1 & 0 & \dots & 0 & -a_1 \\ 0 & 1 & \dots & 0 & -a_2 \\ \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & \dots & 1 & -a_{n-1} \end{bmatrix}$$

We normalize v so that $D(v, Tv, \dots, T^{n-1}v) = 1$ and compute the characteristic polynomial (see 4.4.5) of T , using the basis $\mathbf{v} = \{v, Tv, \dots, T^{n-1}v\}$:

$$(5.2.3) \quad \chi_T(\lambda) = \det(T - \lambda) = D(Tv - \lambda v, \dots, T^n v - \lambda T^{n-1}v).$$

Replace $T^n v = -\sum_0^{n-1} a_k T^k v$, and observe that the only nonzero summand in the expansion of $D(Tv - \lambda v, \dots, T^j v - \lambda T^{j-1}v, \dots, T^{n-1}v - \lambda T^{n-2}v, T^k v)$ is obtained by taking $-\lambda T^j v$ for $j \leq k$ and $T^j v$ for $j > k$ so that

$$D(Tv - \lambda v, \dots, T^{n-1}v - \lambda T^{n-2}v, T^k v) = (-\lambda)^k (-1)^{n-k-1} = (-1)^{n-1} \lambda^k.$$

⁴Consider $T = I$

Adding these, with the weights $-a_k$ for $k < n-1$ and $-\lambda - a_{n-1}$ for $k = n-1$, we obtain

$$(5.2.4) \quad \chi_T(\lambda) = (-1)^n \min P_{T,v}(\lambda).$$

In particular, (5.2.4) implies that if T has a cyclic vector, then $\chi_T(T) = 0$. This is a special case, and a step in the proof, of the following theorem.

Theorem (Hamilton-Cayley). $\chi_T(T) = 0$.

PROOF: We show that χ_T is a multiple of $\min P_{T,v}$ for every $u \in \mathcal{V}$. This implies $\chi_T(T)v = 0$ for all $u \in \mathcal{V}$, i.e., $\chi_T(T) = 0$.

Let $u \in \mathcal{V}$, denote $\mathcal{U} = \text{span}[T, u]$ and $\min P_{T,u} = \lambda^m + \sum_0^{m-1} a_j \lambda^j$. The vectors $u, Tu, \dots, T^{m-1}u$ form a basis for \mathcal{U} . Complete $\{T^j u\}_0^{m-1}$ to a basis for \mathcal{V} by adding w_1, \dots, w_{n-m} . Let A_T be the matrix of T with respect to this basis. The top left $m \times m$ submatrix of A_T is the matrix of $T_{\mathcal{U}}$, and the $(n-m) \times m$ rectangle below it has only zero entries. It follows that $\chi_T = \chi_{T_{\mathcal{U}}} Q$, where Q is the characteristic polynomial of the $(n-m) \times (n-m)$ lower right submatrix of A , and since $\chi_{T_{\mathcal{U}}} = (-1)^m \min P_{T,u}$ (by (5.2.4) applied to $T_{\mathcal{U}}$) the proof is complete. ◀

An alternate way to word the proof, and to prove an additional claim along the way, is to proceed by induction on the dimension of the space \mathcal{V} .

ALTERNATE PROOF: If $n = 1$ the claim is obvious.

Assume the statement valid for all systems of dimension smaller than n .

Let $u \in \mathcal{V}$, $u \neq 0$, and $\mathcal{U} = \text{span}[T, u]$. If $\mathcal{U} = \mathcal{V}$ the claims are a consequence of (5.2.4) as explained above. Otherwise, \mathcal{U} and \mathcal{V}/\mathcal{U} have both dimension smaller than n and, by Proposition 4.4.3 applied to $T - \lambda$, (exercise IV.4.2) we have $\chi_T = \chi_{T_{\mathcal{U}}} \chi_{T_{\mathcal{V}/\mathcal{U}}}$. By the induction hypothesis, $\chi_{T_{\mathcal{V}/\mathcal{U}}}(T_{\mathcal{V}/\mathcal{U}}) = 0$, which means that $\chi_{T_{\mathcal{V}/\mathcal{U}}}(T)$ maps \mathcal{V} into \mathcal{U} , and since $\chi_{T_{\mathcal{U}}}(T)$ maps \mathcal{U} to 0 we have $\chi_T(T) = 0$. ◀

The additional claim is:

Proposition. Every prime factor of χ_T is a factor of $\min P_{T,u}$ for some $u \in \mathcal{V}$.

PROOF: We return to the proof by induction, and add the statement of the proposition to the induction hypothesis. Each prime factor of χ_T is either a factor of $\chi_{T_{\mathcal{U}}}$ or of $\chi_{T_{\mathcal{V}/\mathcal{U}}}$ and, by the strengthened induction hypothesis, is either a factor of $\min P_{T,u}$ or of $\min P_{T_{\mathcal{V}/\mathcal{U}},\tilde{v}}$ for some $\tilde{v} = v + \mathcal{U} \in \mathcal{V}/\mathcal{U}$. In the latter case, observe that $\min P_{T,v}(T)v = 0$. Reducing mod \mathcal{U} gives $\min P_{T,v}(T_{\mathcal{V}/\mathcal{U}})\tilde{v} = 0$, which implies that $\min P_{T_{\mathcal{V}/\mathcal{U}},\tilde{v}}$ divides $\min P_{T,v}$. ◀

5.2.3 Going back to the matrix defined in (5.2.2), let $P(x) = x^n + \sum b_j x^j$ be an arbitrary monic polynomial, the matrix

$$(5.2.5) \quad \begin{bmatrix} 0 & 0 & \dots & 0 & -b_0 \\ 1 & 0 & \dots & 0 & -b_1 \\ 0 & 1 & \dots & 0 & -b_2 \\ \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & \dots & 1 & -b_{n-1} \end{bmatrix}.$$

is called *the companion matrix* of the polynomial P .

If $\{u_0, \dots, u_{n-1}\}$ is a basis for \mathcal{V} , and $S \in \mathcal{L}(\mathcal{V})$ is defined by: $Su_j = u_{j+1}$ for $j < n-1$, and $Su_{n-1} = -\sum_0^{n-2} b_j u_j$. Then u_0 is cyclic for (\mathcal{V}, S) , the matrix (5.2.5) is the matrix $A_{S,u}$ of S with respect to the basis $\mathbf{u} = \{u_0, \dots, u_{n-1}\}$, and $\min P_{S,u_0} = P$.

Thus, every monic polynomial of degree n is $\min P_{S,u}$, the minimal polynomial of some cyclic vector u in an n -dimensional system (\mathcal{V}, S) .

5.2.4 THE MINIMAL POLYNOMIAL.

Let $T \in \mathcal{L}(\mathcal{V})$. The set $\mathfrak{N}_T = \{P : P \in \mathbb{F}[x], P(T) = 0\}$ is an ideal in $\mathbb{F}[x]$. The monic generator⁵ for \mathfrak{N}_T , is called the *minimal polynomial* of T and denoted $\min P_T$. To put it simply: $\min P_T$ is the monic polynomial P of least degree such that $P(T) = 0$.

Since the dimension of $\mathcal{L}(\mathcal{V})$ is n^2 , any $n^2 + 1$ powers of T are linearly dependent. This proves that \mathfrak{N}_T is non-trivial and that the degree of $\min P_T$ is at most n^2 . By the Hamilton-Cayley Theorem, $\chi_T \in \mathfrak{N}_T$ which means that $\min P_T$ divides χ_T and its degree is therefore no bigger than n .

⁵See A.6.1

The condition $P(T) = 0$ is equivalent to “ $P(T)v = 0$ for all $v \in \mathcal{V}$ ”, and the condition “ $P(T)v = 0$ ” is equivalent to $\min P_{T,v}$ divides $\min P_T$. A moment’s reflection gives:

Proposition. *$\min P_T$ is the least common multiple of $\min P_{T,v}$ for all $v \in \mathcal{V}$.*

Invoking proposition 5.2.2 we obtain

Corollary. *Every prime factor of χ_T is a factor of $\min P_T$.*

We shall see later (exercise V.3.7) that there are always vectors v such that $\min P_T$ is equal to $\min P_{T,v}$.

5.2.5 The minimal polynomial gives much information on T and on polynomials in T .

Lemma. *Let P_1 be a polynomial. Then $P_1(T)$ is invertible if, and only if, P_1 is relatively prime to $\min P_T$.*

PROOF: Denote $P = \gcd(P_1, \min P_T)$. By Theorem A.6.2, there exist polynomials q, q_1 such that $q_1 P_1 + q \min P_T = P$. Substituting T for x we have $q_1(T)P_1(T) = P(T)$.

If $P = 1$, $P_1(T)$ is invertible and $q_1(T)$ is its inverse.

If $P \neq 1$ we write $\min P_T = PQ$, so that $P(T)Q(T) = \min P_T(T) = 0$ and hence $\ker(P(T)) \supset \text{range}(Q(T))$. The minimality of $\min P_T$ guarantees that $Q(T) \neq 0$ so that $\text{range}(Q(T)) \neq \{0\}$, and since P is a factor of P_1 , $\ker(P_1(T)) \supset \ker(P(T)) \neq \{0\}$ and $P_1(T)$ is not invertible. ◀

Comments:

a. If $P_1(x) = x$, the lemma says that T itself is invertible if, and only if, $\min P_T(0) \neq 0$. The proof for this case reads: if $\min P_T = xQ(x)$, and T is invertible, then $Q(T) = 0$, contradicting the minimality. On the other hand if $\min P_T(0) = a \neq 0$, write $R(x) = a^{-1}x^{-1}(a - \min P_T)$, and observe that $TR(T) = I - a^{-1} \min P_T(T) = I$.

b. If $\min P_T$ is $P(x)$, then the minimal polynomial for $T + \lambda$ is $P(x - \lambda)$. It follows that $T - \lambda$ is invertible unless $x - \lambda$ divides $\min P_T$, that is unless $\min P_T(\lambda) = 0$.

EXERCISES FOR SECTION 5.2

V.2.1. Let $T \in \mathcal{L}(\mathcal{V})$ and $v \in \mathcal{V}$. Prove that if $u \in \text{span}[T, v]$, then $\min P_{T,u}$ divides $\min P_{T,v}$.

V.2.2. Let \mathcal{U} be a T -invariant subspace of \mathcal{V} and $T_{\mathcal{V}/\mathcal{U}}$ the operator induced on \mathcal{V}/\mathcal{U} . Let $v \in \mathcal{V}$, and let \tilde{v} be its image in \mathcal{V}/\mathcal{U} . Prove that $\min P_{T_{\mathcal{V}/\mathcal{U}}, \tilde{v}}$ divides $\min P_{T,v}$.

V.2.3. If (\mathcal{V}, T) is cyclic (has a cyclic vector), then every S that commutes with T is a polynomial in T . (In other words, $\mathcal{P}(T)$ is a maximal commutative subalgebra of $\mathcal{L}(\mathcal{V})$.)

Hint: If v is cyclic, and $Sv = P(T)v$ for some polynomial P , then $S = P(T)$.

V.2.4. (\mathcal{V}, T) is cyclic if, and only if, $\deg \min P_T = \dim \mathcal{V}$

V.2.5. If $\min P_T$ is irreducible then $\min P_{T,v} = \min P_T$ for every $v \neq 0$ in \mathcal{V} .

V.2.6. Let $P_1, P_2 \in \mathbb{F}[x]$. Prove: $\ker(P_1(T)) \cap \ker(P_2(T)) = \ker(\gcd(P_1, P_2))$.

V.2.7. (Schur's lemma) A system $\{\mathcal{W}, \mathcal{S}\}$, $\mathcal{S} \subset \mathcal{L}(\mathcal{W})$, is *minimal* if no nontrivial subspace of \mathcal{W} is invariant under every $S \in \mathcal{S}$.

Assume $\{\mathcal{W}, \mathcal{S}\}$ minimal, and $T \in \mathcal{L}(\mathcal{W})$.

a. If T commute with every $S \in \mathcal{S}$, so does $P(T)$ for every polynomial P .

b. If T commute with every $S \in \mathcal{S}$, then $\ker(T)$ is either $\{0\}$ or \mathcal{W} . That means that T is either invertible or identically zero.

c. With T as above, the minimal polynomial $\min P_T$ is irreducible.

d. If T commute with every $S \in \mathcal{S}$, and the underlying field is \mathbb{C} , then $T = \lambda I$.

Hint: The minimal polynomial of T must be irreducible, hence linear.

V.2.8. Assume T invertible and $\deg \min P_T = m$. Prove that

$$\min P_{T^{-1}}(x) = cx^m \min P_T(x^{-1}),$$

where $c = \min P_T(0)^{-1}$.

V.2.9. Let $T \in \mathcal{L}(\mathcal{V})$. Prove that $\min P_T$ vanishes at every zero of χ_T .

Hint: If $Tv = \lambda v$ then $T^k v = \lambda^k v$ and $P(T)v = P(\lambda)v$ for any polynomial.

V.2.10. What is the characteristic, resp. minimal, polynomial of the 7×7 matrix $(a_{i,j})$ defined by

$$a_{i,j} = \begin{cases} 1 & \text{if } 3 \leq j = i + 1 \leq 7, \\ 0 & \text{otherwise.} \end{cases}$$

V.2.11. Assume that A is a non-singular matrix and let $\varphi(x) = x^k + \sum_0^{k-1} a_j x^j$ be its minimal polynomial. Prove that $a_0 \neq 0$ and explain how knowing φ gives an efficient way to compute the inverse A^{-1} .

5.3 REDUCING.

5.3.1 Let (\mathcal{V}, T) be a linear system. A subspace $\mathcal{V}_1 \subset \mathcal{V}$ *reduces* T if it is T -invariant and has a T -invariant complement, that is, a T -invariant subspace \mathcal{V}_2 such that $\mathcal{V} = \mathcal{V}_1 \oplus \mathcal{V}_2$.

A system (\mathcal{V}, T) that admits no reducing subspaces is *irreducible*. We say also that T is *irreducible* on \mathcal{V} . An invariant subspace is *irreducible* if T restricted to it is irreducible.

Theorem. *Every system (\mathcal{V}, T) is completely decomposable, that is, can be decomposed into a direct sum of irreducible systems.*

PROOF: Use induction on $n = \dim \mathcal{V}$. If $n = 1$ the system is trivially irreducible. Assume the validity of the statement for $n < N$ and let (\mathcal{V}, T) be of dimension N . If (\mathcal{V}, T) is irreducible the decomposition is trivial. If (\mathcal{V}, T) is reducible, let $\mathcal{V} = \mathcal{V}_1 \oplus \mathcal{V}_2$ be a non-trivial decomposition with T -invariant \mathcal{V}_j . Then $\dim \mathcal{V}_j < N$, hence each system $(\mathcal{V}_j, T_{\mathcal{V}_j})$ is completely decomposable, $\mathcal{V}_j = \bigoplus_k \mathcal{V}_{j,k}$ with every $\mathcal{V}_{j,k}$ T -invariant, and $\mathcal{V} = \bigoplus_{j,k} \mathcal{V}_{j,k}$. ◀

Our interest in reducing subspaces is that operators can be analyzed separately on each direct summand (of a direct sum of invariant subspaces).

The effect of a direct sum decomposition into T -invariant subspaces on the matrix representing T (relative to an appropriate basis) can be seen as follows:

Assume $\mathcal{V} = \mathcal{V}_1 \oplus \mathcal{V}_2$, with T -invariant \mathcal{V}_j , and $\{v_1, \dots, v_n\}$ is a basis for \mathcal{V} such that the first k elements are a basis for \mathcal{V}_1 while the the last $l = n - k$ elements are a basis for \mathcal{V}_2 .

The entries $a_{i,j}$ of the matrix A_T of T relative to this basis are zero unless both i and j are $\leq k$, or both are $> k$. A_T consists of two square blocks centered on the diagonal. The first is the $k \times k$ matrix of T restricted to \mathcal{V}_1 (relative to the basis $\{v_1, \dots, v_k\}$), and the second is the $l \times l$ matrix of T restricted to \mathcal{V}_2 (relative to $\{v_{k+1}, \dots, v_n\}$).

Similarly, if $\mathcal{V} = \bigoplus_{j=1}^s \mathcal{V}_j$ is a decomposition with T -invariant components, and we take as basis for \mathcal{V} the union of s successive blocks—the bases of \mathcal{V}_j , then the matrix A_T relative to this basis is the *diagonal sum*⁶ of square matrices, A_j , i.e., consists of s square matrices A_1, \dots, A_s along the diagonal (and zero everywhere else). For each j , A_j is the matrix representing the action of T on \mathcal{V}_j relative to the chosen basis.

5.3.2 The *rank and nullity theorem* (see Chapter II, 2.5) gives an immediate characterization of operators whose kernels are reducing.

Proposition. *Assume \mathcal{V} finite dimensional and $T \in \mathcal{L}(\mathcal{V})$. $\ker(T)$ reduces T if, and only if, $\ker(T) \cap \text{range}(T) = \{0\}$.*

PROOF: Assume $\ker(T) \cap \text{range}(T) = \{0\}$. Then the sum $\ker(T) + \text{range}(T)$ is a direct sum and, since

$$\dim(\ker(T) \oplus \text{range}(T)) = \dim \ker(T) + \dim \text{range}(T) = \dim \mathcal{V},$$

we have $\mathcal{V} = \ker(T) \oplus \text{range}(T)$. Both $\ker(T)$ and $\text{range}(T)$ are T -invariant and the direct sum decomposition proves that they are reducing.

The opposite implication is proved in Proposition 5.3.3 below. ◀

Corollary. *$\ker(T)$ and $\text{range}(T)$ reduce T if, and only if, $\ker(T^2) = \ker(T)$.*

PROOF: For any $T \in \mathcal{L}(\mathcal{V})$ we have $\ker(T^2) \supseteq \ker(T)$ and the inclusion is proper if, and only if, there exist vectors v such that $Tv \neq 0$ but $T^2v = 0$, which amounts to $Tv \in \ker(T)$. ◀

5.3.3 Given that $\mathcal{V}_1 \subset \mathcal{V}$ reduces T —there exists a T -invariant \mathcal{V}_2 such that $\mathcal{V} = \mathcal{V}_1 \oplus \mathcal{V}_2$ —how uniquely determined is \mathcal{V}_2 . Considering the somewhat extreme example $T = I$, the condition of T -invariance is satisfied trivially and we realize that \mathcal{V}_2 is far from being unique. There are, however, cases in which the “complementary invariant subspace”, if there is one, is uniquely determined. We propose to show now that this is the case for the T -invariant subspaces $\ker(T)$ and $\text{range}(T)$.

⁶Also called the *direct sum* of A_j , $j = 1, \dots, s$

Proposition. *Let \mathcal{V} be finite dimensional and $T \in \mathcal{L}(\mathcal{V})$.*

- a. If $\mathcal{V}_2 \subset \mathcal{V}$ is T -invariant and $\mathcal{V} = \ker(T) \oplus \mathcal{V}_2$, then $\mathcal{V}_2 = \text{range}(T)$.*
- b. If $\mathcal{V}_1 \subset \mathcal{V}$ is T -invariant and $\mathcal{V} = \mathcal{V}_1 \oplus \text{range}(T)$, then $\mathcal{V}_1 = \ker(T)$.*

PROOF: *a.* As $\dim \ker(T) + \dim \mathcal{V}_2 = \dim \mathcal{V} = \dim \ker(T) + \dim \text{range}(T)$, we have $\dim \mathcal{V}_2 = \dim \text{range}(T)$. Also, since $\ker(T) \cap \mathcal{V}_2 = \{0\}$, T is 1-1 on \mathcal{V}_2 and $\dim T\mathcal{V}_2 = \dim \text{range}(T)$. Now, $T\mathcal{V}_2 = \text{range}(T) \subset \mathcal{V}_2$ and, since they have the same dimension, $\mathcal{V}_2 = \text{range}(T)$.

b. $T\mathcal{V}_1 \subset \mathcal{V}_1 \cap \text{range}(T) = \{0\}$, and hence $\mathcal{V}_1 \subset \ker(T)$. Since \mathcal{V}_1 has the same dimension as $\ker(T)$ they are equal. ◀

5.3.4 THE CANONICAL PRIME-POWER DECOMPOSITION.

Lemma. *If P_1 and P_2 are relatively prime, then*

$$(5.3.1) \quad \ker(P_1(T)) \cap \ker(P_2(T)) = \{0\}.$$

If also $P_1(T)P_2(T) = 0$ then $\mathcal{V} = \ker(P_1(T)) \oplus \ker(P_2(T))$, and the corresponding projections are polynomials in T .

PROOF: Given that P_1 and P_2 are relatively prime there exist, By Appendix A.6.1, polynomials q_1, q_2 such that $q_1P_1 + q_2P_2 = 1$. Substituting T for the variable we have

$$(5.3.2) \quad q_1(T)P_1(T) + q_2(T)P_2(T) = I.$$

If $v \in \ker(P_1(T)) \cap \ker(P_2(T))$, that is, $P_1(T)v = P_2(T)v = 0$, then $v = q_1(T)P_1(T)v + q_2(T)P_2(T)v = 0$. This proves (5.3.1) which implies, in particular, that $\dim \ker(P_1(T)) + \dim \ker(P_2(T)) \leq n$.

If $P_1(T)P_2(T) = 0$, then the range of either $P_j(T)$ is contained in the kernel of the other. By the *Rank and Nullity* theorem

$$(5.3.3) \quad \begin{aligned} n &= \dim \ker(P_1(T)) + \dim \text{range}(P_1(T)) \\ &\leq \dim \ker(P_1(T)) + \dim \ker(P_2(T)) \leq n. \end{aligned}$$

It follows that $\dim \ker(P_1(T)) + \dim \ker(P_2(T)) = n$, which implies that $\ker(P_1(T)) \oplus \ker(P_2(T))$ is all of \mathcal{V} . ◀

Observe that the proof shows that

$$(5.3.4) \quad \text{range}(P_1(T)) = \ker(P_2(T)) \quad \text{and} \quad \text{range}(P_2(T)) = \ker(P_1(T)).$$

Equation (5.3.2) implies that $\varphi_2(T) = q_1(T)P_1(T) = I - q_2(T)P_2(T)$ is the identity on $\ker(P_2(T))$ and zero on $\ker(P_1(T))$, that is, φ_2 is the projection onto $\ker(P_2(T))$ along $\ker(P_1(T))$.

Similarly, $\varphi_1(T) = q_2(T)P_2(T)$ is the projection onto $\ker(P_1(T))$ along $\ker(P_2(T))$.

Corollary. *For every factorization $\min P_T = \prod_{j=1}^l P_j$ into pairwise relatively prime factors, we have a direct sum decomposition of \mathcal{V}*

$$(5.3.5) \quad \mathcal{V} = \bigoplus_{j=1}^l \ker(P_j(T)).$$

PROOF: Use induction on the number of factors. ◀

For the *prime-power factorization* $\min P_T = \prod \Phi_j^{m_j}$, where the Φ_j 's are distinct prime (irreducible) polynomials in $\mathbb{F}[x]$, and m_j their respective multiplicities, we obtain the *canonical prime-power decomposition* of (\mathcal{V}, T) :

$$(5.3.6) \quad \mathcal{V} = \bigoplus_{j=1}^k \ker(\Phi_j^{m_j}(T)).$$

The subspaces $\ker(\Phi_j^{m_j}(T))$ are called the *primary components* of (\mathcal{V}, T)

Comments: By the Cayley-Hamilton theorem and corollary 5.2.4, the prime-power factors of χ_T are those of $\min P_T$, with at least the same multiplicities, that is:

$$(5.3.7) \quad \chi_T = \prod \Phi_j^{s_j}, \quad \text{with } s_j \geq m_j.$$

The minimal polynomial of T restricted to $\ker(\Phi_j^{m_j}(T))$ is $\Phi_j^{m_j}$ and its characteristic polynomial is $\Phi_j^{s_j}$. The dimension of $\ker(\Phi_j^{m_j}(T))$ is $s_j \deg(\Phi_j)$.

5.3.5 When the underlying field \mathbb{F} is algebraically closed, and in particular when $\mathbb{F} = \mathbb{C}$, every irreducible polynomial in $\mathbb{F}[x]$ is linear and every polynomial is a product of linear factors, see Appendix A.6.5.

Recall that the *spectrum* of T is the set $\sigma(T) = \{\lambda_j\}$ of zeros of χ_T or, equivalently, of $\min P_T$. The prime-power factorization of $\min P_T$ (for systems over an algebraically closed field) has the form $\min P_T = \prod_{\lambda \in \sigma(T)} (x - \lambda)^{m(\lambda)}$ where $m(\lambda)$ is the multiplicity of λ in $\min P_T$.

The space $\mathcal{V}_\lambda = \ker((T - \lambda)^{m(\lambda)})$ is called the *generalized eigenspace*, or, *nilspace* of λ . The canonical decomposition of (\mathcal{V}, T) is given by:

$$(5.3.8) \quad \mathcal{V} = \bigoplus_{\lambda \in \sigma(T)} \mathcal{V}_\lambda.$$

5.3.6 The projections $\varphi_j(T)$ corresponding to the canonical prime-power decomposition are given by $\varphi_j(T) = q_j(T) \prod_{i \neq j} \Phi_i^{m_i}(T)$, where the polynomials q_i are given by the representations (see Corollary A.6.2)

$$q_j \prod_{i \neq j} \Phi_i^{m_i} + q_j^* \Phi_j^{m_j} = 1.$$

An immediate consequence of the fact that these are all polynomials in T is that they all commute, and commute with T .

If $\mathcal{W} \subset \mathcal{V}$ is T -invariant then the subspaces $\varphi_j(T)\mathcal{W} = \mathcal{W} \cap \ker(\Phi_j^{m_j}(T))$, are T -invariant and we have a decomposition

$$(5.3.9) \quad \mathcal{W} = \bigoplus_{j=1}^k \varphi_j(T)\mathcal{W}$$

Proposition. *The T -invariant subspace \mathcal{W} is reducing if, and only if, $\varphi_j(T)\mathcal{W}$ is a reducing subspace of $\ker(\Phi_j^{m_j}(T))$ for every j .*

PROOF: If \mathcal{W} is reducing and \mathcal{U} is a T -invariant complement, then

$$\ker(\Phi_j^{m_j}(T)) = \varphi_j(T)\mathcal{V} = \varphi_j(T)\mathcal{W} \oplus \varphi_j(T)\mathcal{U},$$

and both components are T -invariant.

Conversely, if \mathcal{U}_j is T -invariant and $\ker(\Phi_j^{m_j}(T)) = \varphi_j(T)\mathcal{W} \oplus \mathcal{U}_j$, then $\mathcal{U} = \bigoplus \mathcal{U}_j$ is an invariant complement to \mathcal{W} . ◀

5.3.7 Recall (see 5.3.1) that if $\mathcal{V} = \bigoplus_{j=1}^s \mathcal{V}_j$ is a direct sum decomposition into T invariant subspaces, and if we take for a basis on \mathcal{V} the union of bases of the summands \mathcal{V}_j , then the matrix of T with respect to this basis is the *diagonal sum* of the matrices of the restrictions of T to the components \mathcal{V}_j . By that we mean

$$(5.3.10) \quad A_T = \begin{bmatrix} A_1 & 0 & \dots & 0 & 0 \\ 0 & A_2 & 0 & \dots & 0 \\ 0 & 0 & A_3 & 0 & \dots \\ \vdots & \vdots & & \ddots & \vdots \\ 0 & 0 & 0 & 0 & A_s \end{bmatrix}$$

where A_j is the matrix of $T|_{\mathcal{V}_j}$ (the restriction of T to the component \mathcal{V}_j in the decomposition.)

EXERCISES FOR SECTION 5.3

V.3.1. Let $T \in \mathcal{L}(\mathcal{V})$, $k > 0$ and integer. Prove that $\ker(T^k)$ reduces T if, and only if $\ker(T^{k+1}) = \ker(T^k)$.

Hint: Both $\ker(T^k)$ and $\text{range}(T^k)$ are T -invariant.

V.3.2. Let $T \in \mathcal{L}(\mathcal{V})$, and $\mathcal{V} = \mathcal{U} \oplus \mathcal{W}$ with both summands T -invariant. Let π be the projection onto \mathcal{U} along \mathcal{W} . Prove that π commutes with T .

V.3.3. Prove that if (\mathcal{V}, T) is irreducible, then its minimal polynomial is “prime power” that is, $\text{minP}_T = \Phi^m$ with Φ irreducible and $m \geq 1$.

V.3.4. If $\mathcal{V}_j = \ker(\Phi_j^{m_j}(T))$ is a primary component of (\mathcal{V}, T) , the minimal polynomial of $T|_{\mathcal{V}_j}$ is $\Phi_j^{m_j}$.

V.3.5. Show that if $\text{minP}_T = \Phi^m$, with Φ irreducible, then there exist vectors $v \in \mathcal{V}$ such that $\text{minP}_{T,v} = \text{minP}_T$.

V.3.6. Let $v_1, v_2 \in \mathcal{V}$ and assume that minP_{T,v_1} and minP_{T,v_2} are relatively prime. Prove that $\text{minP}_{T,v_1+v_2} = \text{minP}_{T,v_1} \text{minP}_{T,v_2}$.

Hint: Write $P_j = \text{minP}_{T,v_j}$, $Q = \text{minP}_{T,v_1+v_2}$, and let q_j be polynomials such that $q_1 P_1 + q_2 P_2 = 1$. Then $Q q_2 P_2(T)(v_1 + v_2) = Q(T)(v_1) = 0$, and so $P_1 \mid Q$. Similarly $P_2 \mid Q$, hence $P_1 P_2 \mid Q$. Also, $P_1 P_2(T)(v_1 + v_2) = 0$, and $Q \mid P_1 P_2$.

V.3.7. Show that there *always* exist vectors $v \in \mathcal{V}$ such that $\text{minP}_{T,v} = \text{minP}_T$.

Hint: use the prime-power decomposition and the previous exercise.

5.4 SEMISIMPLE SYSTEMS.

5.4.1 DEFINITION: The system (\mathcal{V}, T) is *semisimple* if every T -invariant subspace of \mathcal{V} is reducing.

Theorem. *The system (\mathcal{V}, T) is semisimple if, and only if, $\min P_T$ is square free (that is, the multiplicities m_j of the factors in the canonical factorization $\min P_T = \prod \Phi_j^{m_j}$ are all 1).*

PROOF: Proposition 5.3.6 reduces the general case to that in which $\min P_T$ is Φ^m with Φ irreducible.

a. When $m > 1$. $\Phi(T)$ is not invertible and hence the invariant subspace $\ker(\Phi(T))$ is non-trivial nor is it all of \mathcal{V} . $\ker(\Phi(T)^2)$ is strictly bigger than $\ker(\Phi(T))$ and, by corollary 5.3.2, $\ker(\Phi(T))$ is not $\Phi(T)$ -reducing, and hence not T -reducing.

b. When $m = 1$. Observe first that $\min P_{T,v} = \Phi$ for every non-zero $v \in \mathcal{V}$. This since $\min P_{T,v}$ divides Φ and Φ is prime. It follows that the dimension of $\text{span}[T, v]$ is equal to the degree d of Φ , and hence: *every non-trivial T -invariant subspace has dimension $\geq d$.*

Let $\mathcal{W} \subset \mathcal{V}$ be a proper T -invariant subspace, and $v_1 \notin \mathcal{W}$. The subspace $\text{span}[T, v_1] \cap \mathcal{W}$ is T -invariant and is properly contained in $\text{span}[T, v_1]$, so that its dimension is smaller than d , hence $\text{span}[T, v_1] \cap \mathcal{W} = \{0\}$. It follows that $\mathcal{W}_1 = \text{span}[T, (\mathcal{W}, v_1)] = \mathcal{W} \oplus \text{span}[T, v_1]$.

If $\mathcal{W}_1 \neq \mathcal{V}$, let $v_2 \in \mathcal{V} \setminus \mathcal{W}_1$ and define $\mathcal{W}_2 = \text{span}[T, (\mathcal{W}, v_1, v_2)]$. The argument above shows that $\mathcal{W}_2 = \mathcal{W} \oplus \text{span}[T, v_1] \oplus \text{span}[T, v_2]$. This can be repeated until, for the appropriate⁷ k , we have

$$(5.4.1) \quad \mathcal{V} = \mathcal{W} \oplus \bigoplus_{j=1}^k \text{span}[T, v_j]$$

and $\bigoplus_{j=1}^k \text{span}[T, v_j]$ is clearly T -invariant. ◀

Remark: Notice that if we take $\mathcal{W} = \{0\}$, the decomposition (5.4.1) expresses (\mathcal{V}, T) as a direct sum of cyclic subsystems.

⁷The dimension of \mathcal{W}_{i+1} is $\dim \mathcal{W}_i + d$, so that $kd = \dim \mathcal{V} - \dim \mathcal{W}$.

5.4.2 If (\mathcal{V}, T) is *semisimple* and \mathbb{F} is algebraically closed, and in particular if $\mathbb{F} = \mathbb{C}$, all irreducible polynomials in $\mathbb{F}[x]$ are linear. If $\Phi_j(T) = T - \lambda_j$ with $\lambda_j \in \mathbb{F}$, then the *canonical prime-power decomposition* has the form

$$(5.4.2) \quad \mathcal{V} = \bigoplus \ker(T - \lambda_j),$$

and, for each j , the restriction of T to $\ker(T - \lambda_j)$ is just multiplication by λ_j .

5.4.3 If \mathbb{F} is not algebraically closed and $\min P_T = \Phi$ is irreducible, but non-linear, we have much the same phenomenon, but in somewhat hidden form.

Lemma. *Let $T \in \mathcal{L}(\mathcal{V})$, and assume that $\min P_T$ is irreducible in $\mathbb{F}[x]$. Then $\mathcal{P}(T) = \{P(T) : P \in \mathbb{F}[x]\}$ is a field.*

PROOF: If $P \in \mathbb{F}[x]$ and $P(T) \neq 0$, then $\gcd(P, \Phi) = 1$ and hence $P(T)$ is invertible. Thus, every non-zero element in $\mathcal{P}(T)$ is invertible and $\mathcal{P}(T)$ is a field. ◀

★ **5.4.4** \mathcal{V} can now be considered as a vector space over the extended field $\mathcal{P}(T)$ by considering the action of $P(T)$ on v as a *multiplication of v by the “scalar”* $P(T) \in \mathcal{P}(T)$. This defines a system $(\mathcal{V}_{\mathcal{P}(T)}, T)$. A subspace of $(\mathcal{V}_{\mathcal{P}(T)})$ is precisely a T -invariant subspace of \mathcal{V} .

The subspace $\text{span}[T, v]$, in \mathcal{V} (over \mathbb{F}) becomes “the line through v in $(\mathcal{V}_{\mathcal{P}(T)})$ ”, i.e. the set of all multiples of v by scalars from $\mathcal{P}(T)$; the statement “Every subspace of a finite-dimensional vector space (here \mathcal{V} over $\mathcal{P}(T)$), has a basis.” translates here to: “Every T -invariant subspace of \mathcal{V} is a direct sum of cyclic subspaces, that is subspaces of the form $\text{span}[T, v]$.”

EXERCISES FOR SECTION 5.4

V.4.1. a. If T is diagonalizable then (\mathcal{V}, T) is semisimple.

b. If \mathbb{F} is algebraically closed and (\mathcal{V}, T) is semisimple, then T is diagonalizable

V.4.2. An algebra $B \subset \mathcal{L}(\mathcal{V})$ is *semisimple* if every $T \in B$ is semisimple. Prove that if B is commutative and semisimple, then $\dim B \leq \dim \mathcal{V}$.

V.4.3. Let $\mathcal{V} = \mathcal{V}_1 \oplus \mathcal{V}_0$ and let $B \subset \mathcal{L}(\mathcal{V})$ be the set of all the operators S such that $S\mathcal{V}_1 \subset \mathcal{V}_0$, and $S\mathcal{V}_0 = \{0\}$. Prove that B is a commutative subalgebra of $\mathcal{L}(\mathcal{V})$ and that $\dim B = \dim \mathcal{V}_0 \dim \mathcal{V}_1$. When is B semisimple?

V.4.4. Let B be the subset of $\mathcal{M}(2; \mathbb{R})$ of the matrices of the form $\begin{bmatrix} a & b \\ -b & a \end{bmatrix}$. Prove that B is an algebra over \mathbb{R} , which is in fact a field isomorphic to \mathbb{C} .

V.4.5. Let \mathcal{V} be an n -dimensional real vector space, and $T \in \mathcal{L}(\mathcal{V})$ an operator with non-linear irreducible minimal polynomial. Prove that n is even and explain: (\mathcal{V}, T) is “isomorphic” to $(\mathbb{C}^{n/2}, sI)$ (s a complex number, I the identity on $\mathbb{C}^{n/2}$).

5.5 NILPOTENT OPERATORS

The canonical prime-power decomposition reduces every system to a direct sum of systems whose minimal polynomial is a power of an irreducible polynomial Φ . If \mathbb{F} is algebraically closed, and in particular if $\mathbb{F} = \mathbb{C}$, the irreducible polynomials are linear, $\Phi(x) = (x - \lambda)$ for some scalar λ . We consider here the case of linear Φ , and discuss the general case in section $\star 5.6$.

If $\min P_T = (x - \lambda)^m$, then $\min P_{(T - \lambda)} = x^m$ and the structure of $S = T - \lambda$ clarifies the structure of T . We therefore focus on the case $\lambda = 0$.

5.5.1 DEFINITION: An operator $T \in \mathcal{L}(\mathcal{V})$ is *nilpotent* if for some positive integer k , $T^k = 0$. The *height* of (\mathcal{V}, T) , denoted $\text{height}[(\mathcal{V}, T)]$, is the smallest positive k for which $T^k = 0$.

If $T^k = 0$, $\min P_T$ divides x^k , hence it is a power of x . In other words, T is *nilpotent of height k* if $\min P_T(x) = x^k$.

For every $v \in \mathcal{V}$, $\min P_{T,v}(x) = x^l$ for an appropriate l . The height, $\text{height}[v]$, of a vector v (under the action of T) is the degree of $\min P_{T,v}$, that is smallest integer l such that $T^l v = 0$. It is the height of $T|_{\mathcal{W}}$, where $\mathcal{W} = \text{span}[T, v]$, the span of v under T . Since for $v \neq 0$, $\text{height}[Tv] = \text{height}[v] - 1$, elements of maximal height are not in $\text{range}(T)$.

EXAMPLE: \mathcal{V} is the space of all (algebraic) polynomials of degree bounded by m , (so that $\{x^j\}_{j=0}^m$ is a basis for \mathcal{V}), T the differentiation operator:

$$(5.5.1) \quad T\left(\sum_0^m a_j x^j\right) = \sum_1^m j a_j x^{j-1} = \sum_0^{m-1} (j+1) a_{j+1} x^j.$$

The vector $w = x^m$ has height $m + 1$, and $\{T^j w\}_{j=0}^m$ is a basis for \mathcal{V} (so that w is a cyclic vector). If we take $v_j = \frac{x^{m-j}}{(m-j)!}$ as basis elements, the operator takes the form of *the standard shift of height $m + 1$* .

DEFINITION: A k -shift is a k -dimensional system $\{\mathcal{V}, T\}$ with T nilpotent of height k . A *standard shift* is a k -shift for some k , that is, a cyclic nilpotent system.

If $\{\mathcal{V}, T\}$ is a k -shift, $v_0 \in \mathcal{V}$ and $\text{height}[v_0] = k$, then $\{T^j v_0\}_{j=0}^{k-1}$ is a basis for \mathcal{V} , and the action of T is to map each basis element, except for the last, to the next one, and map the last basis element to 0. The matrix of T with respect to this basis is

$$(5.5.2) \quad A_{T, \mathbf{v}} = \begin{bmatrix} 0 & 0 & \dots & 0 & 0 \\ 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & \dots & 1 & 0 \end{bmatrix}$$

Shifts are the building blocks that nilpotent systems are made of.

5.5.2 Theorem (Cyclic decomposition for nilpotent operators). *Let (\mathcal{V}, T) be a finite dimensional nilpotent system of height k . Then $\mathcal{V} = \bigoplus \mathcal{V}_j$, where \mathcal{V}_j are T -invariant, and $(\mathcal{V}_j, T|_{\mathcal{V}_j})$ is a standard shift.*

Moreover, if we arrange the direct summands so that $k_j = \text{height}[(\mathcal{V}_j, T)]$ is monotone non-increasing, then $\{k_j\}$ is uniquely determined.

PROOF: We use induction on $k = \text{height}[(\mathcal{V}, T)]$.

a. If $k = 1$, then $T = 0$ and any decomposition $\mathcal{V} = \bigoplus \mathcal{V}_j$ into one dimensional subspaces will do.

b. Assume the statement valid for systems of height less than k and let (\mathcal{V}, T) be a (finite dimensional) nilpotent system of height k .

Write $\mathcal{W}_{\text{in}} = \ker(T) \cap T\mathcal{V}$, and let $\mathcal{W}_{\text{out}} \subset \ker(T)$ be a complementary subspace, i.e., $\ker(T) = \mathcal{W}_{\text{in}} \oplus \mathcal{W}_{\text{out}}$.

$(T\mathcal{V}, T)$ is nilpotent of height $k - 1$ and, by the induction hypothesis, admits a decomposition $T\mathcal{V} = \bigoplus_{j=1}^m \tilde{\mathcal{V}}_j$ into standard shifts. Denote $l_j =$

height $[(\tilde{\mathcal{V}}_j, T)]$. Let \tilde{v}_j be of height l_j in $\tilde{\mathcal{V}}_j$ (so that $\tilde{\mathcal{V}}_j = \text{span}[T, \tilde{v}_j]$), and observe that $\{T^{l_j-1}\tilde{v}_j\}$ is a basis for \mathcal{W}_{in} .

Let v_j be such that $\tilde{v}_j = Tv_j$, write $\mathcal{V}_j = \text{span}[T, v_j]$, and let $\mathcal{W}_{\text{out}} = \bigoplus_{i \leq l} \mathcal{W}_i$ be a direct sum decomposition into one dimensional subspaces. The claim now is

$$(5.5.3) \quad \mathcal{V} = \bigoplus \mathcal{V}_j \oplus \bigoplus \mathcal{W}_i.$$

To prove (5.5.3) we need to show that the spaces $\{\mathcal{V}_j, \mathcal{W}_i\}$, $i = 1, \dots, l$, $j = 1, \dots, m$, are independent and span \mathcal{V} .

Independence: Assume there is a non-trivial relation $\sum u_j + \sum w_i = 0$ with $u_j \in \mathcal{V}_j$ and $w_i \in \mathcal{W}_i$. Let $h = \max \text{height}[u_j]$.

If $h > 1$, then $\sum T^{h-1}u_j = T^{h-1}(\sum u_j + \sum w_i) = 0$ and we obtain a non-trivial relation between the $\tilde{\mathcal{V}}_j$'s. A contradiction.

If $h = 1$ we obtain a non-trivial relation between elements of a basis of $\ker(T)$. Again a contradiction.

Spanning: Denote $\mathcal{U} = \text{span}[\{\mathcal{W}_i, \mathcal{V}_j\}]$, $i = 1, \dots, l$, $j = 1, \dots, m$. $T\mathcal{U}$ contains every \tilde{v}_j , and hence $T\mathcal{U} = T\mathcal{V}$. It follows that $\mathcal{U} \supset \mathcal{W}_{\text{in}}$ and since it contains (by its definition) \mathcal{W}_{out} , we have $\mathcal{U} \supset \ker(T)$.

For arbitrary $v \in \mathcal{V}$, let $\hat{v} \in \mathcal{U}$ be such that $Tv = T\hat{v}$. Then $v - \hat{v} \in \ker(T) \subset \mathcal{U}$ so that $v \in \mathcal{U}$, and $\mathcal{U} = \mathcal{V}$.

Finally, if we denote by $n(h)$ the number of summands \mathcal{V}_j in (5.5.3) of dimension (i.e., height) h , then $n(k) = \dim T^{k-1}\mathcal{V}$ while, for $l = 0, \dots, k-2$, we have

$$(5.5.4) \quad \dim T^l\mathcal{V} = \sum_{h=l+1}^k (h-l)n(h),$$

which determines $\{n(h)\}$ completely. ◀

Corollary. *An irreducible nilpotent system is a standard shift.*

EXERCISES FOR SECTION 5.5

V.5.1. Assume $\mathbb{F} = \mathbb{R}$ and $\min P_T = \Phi(x) = x^2 + 1$. Prove that $a + bT \mapsto a + bi$ is a (field) isomorphism of \mathbb{F}_{Φ} onto \mathbb{C} .

What is \mathbb{F}_Φ if $\min P_T = \Phi(x) = x^2 + 3$?

V.5.2. Assume $\min P_T = \Phi^m$ with irreducible Φ . Can you explain (justify) the statement: (\mathcal{V}, T) is “essentially” a standard m -shift over \mathbb{F}_Φ .

★ 5.6 THE CYCLIC DECOMPOSITION

We now show that the canonical prime-power decomposition can be refined to a *cyclic decomposition*.

DEFINITION: A *cyclic decomposition* of a system (\mathcal{V}, T) is a direct sum decomposition of the system into *irreducible cyclic subspaces*, that is, irreducible subspaces of the form $\text{span}[T, v]$.

The summands in the canonical prime-power decomposition have the form $\ker(\Phi^m(T))$ with an irreducible polynomial Φ . We show here that such systems (whose minimal polynomial is Φ^m , with irreducible Φ) admit a cyclic decomposition.

In the previous section we proved the special case⁸ in which $\Phi(x) = x$. If we use the point of view proposed in subsection ★5.4.4, the general case is nothing more than the nilpotent case over the field $\mathcal{P}(T)$ and nothing more need be proved.

For the reader not used to switching underlying fields we repeat the proof of the nilpotent case in the present context.

5.6.1 We assume now that $\min P_T = \Phi^m$ with irreducible Φ of degree d . For every $v \in \mathcal{V}$, $\min P_{T,v} = \Phi^{k(v)}$, $1 \leq k \leq m$, and $\max_v k(v) = m$; we refer to $k(v)$ as the Φ -height, or simply height, of v .

Theorem. *There exist vectors $v_j \in \mathcal{V}$ such that $\mathcal{V} = \bigoplus \text{span}[T, v_j]$. Moreover, the set of the Φ -heights of the v_j 's is uniquely determined.*

PROOF: We use induction on the Φ -height m .

a. $m = 1$. See 5.4.

b. Assume that $\min P_T = \Phi^m$, and the theorem valid for heights lower than m . Write $\mathcal{W}_{\text{in}} = \ker(\Phi(T)) \cap \Phi(T)\mathcal{V}$ and let $\mathcal{W}_{\text{out}} \subset \ker(\Phi(T))$ be a complemen-

⁸Notice that when $\Phi(x) = x$, a cyclic space is what we called a *standard shift*.

tary T -invariant subspace, i.e., $\ker(\Phi(T)) = \mathcal{W}_{\text{in}} \oplus \mathcal{W}_{\text{out}}$. Such complementary T -invariant subspace of $\ker(\Phi(T))$ exists since the system $(\ker(\Phi(T)), T)$ is semisimple, see 5.4.

$(\Phi(T)\mathcal{V}, T)$ is of height $m - 1$ and, by the induction hypothesis, admits a decomposition $\Phi(T)\mathcal{V} = \bigoplus_{j=1}^m \tilde{\mathcal{V}}_j$ into cyclic subspaces, $\tilde{\mathcal{V}}_j = \text{span}[T, \tilde{v}_j]$. Let v_j be such that $\tilde{v}_j = \Phi(T)v_j$.

Write $\mathcal{V}_j = \text{span}[T, v_j]$, and let $\mathcal{W}_{\text{out}} = \bigoplus_{i \leq l} \mathcal{W}_i$ be a direct sum decomposition into cyclic subspaces. The claim now is

$$(5.6.1) \quad \mathcal{V} = \bigoplus \mathcal{V}_j \oplus \bigoplus \mathcal{W}_i.$$

To prove (5.6.1) we need to show that the spaces $\{\mathcal{V}_j, \mathcal{W}_i\}$, $i = 1, \dots, l$, $j = 1, \dots, m$, are independent, and that they span \mathcal{V} .

Independence: Assume there is a non-trivial relation $\sum u_j + \sum w_i = 0$ with $u_j \in \mathcal{V}_j$ and $w_i \in \mathcal{W}_i$. Let $h = \max \Phi\text{-height}[u_j]$.

If $h > 1$, then $\sum \Phi(T)^{h-1}u_j = \Phi(T)^{h-1}(\sum u_j + \sum w_i) = 0$ and we obtain a non-trivial relation between the $\tilde{\mathcal{V}}_j$'s. A contradiction.

If $h = 1$ we obtain a non-trivial relation between elements of a basis of $\ker(\Phi)(T)$. Again a contradiction.

Spanning: Denote $\mathcal{U} = \text{span}[\{\mathcal{W}_i, \mathcal{V}_j\}]$, $i = 1, \dots, l$, $j = 1, \dots, m$. Notice first that $\mathcal{U} \supset \ker(T)$.

$\Phi(T)\mathcal{U}$ contains every \tilde{v}_j , and hence $T\mathcal{U} = T\mathcal{V}$. For $v \in \mathcal{V}$, let $\tilde{v} \in \mathcal{U}$ be such that $Tv = T\tilde{v}$. Then $v - \tilde{v} \in \ker(T) \subset \mathcal{U}$ so that $v \in \mathcal{U}$, and $\mathcal{U} = \mathcal{V}$.

Finally, just as in the previous subsection, denote by $n(h)$ the number of v_j 's of Φ -height h in the decomposition. Then $dn(m) = \dim \Phi(T)^{m-1}\mathcal{V}$ and, for $l = 0, \dots, m - 2$, we have

$$(5.6.2) \quad \dim \Phi(T)^l \mathcal{V} = d \sum_{h=l+1}^k (h-l)n(h),$$

which determines $\{n(h)\}$ completely. ◀

5.6.2 THE GENERAL CASE.

We now refine the canonical prime-power decomposition (5.3.6) by applying Theorem 5.6.1 to each of the summands:

Theorem (General cyclic decomposition). *Let (\mathcal{V}, T) be a linear system over a field \mathbb{F} . Let $\min P_T = \prod \Phi_j^{m_j}$ be the prime-power decomposition of its minimal polynomial. Then (\mathcal{V}, T) admits a cyclic decomposition*

$$\mathcal{V} = \bigoplus \mathcal{V}_k.$$

For each k , the minimal polynomial of T on \mathcal{V}_k is $\Phi_{j(k)}^{l(k)}$ for some $l(k) \leq m_{j(k)}$, and $m_{j(k)} = \max l(k)$.

The polynomials $\Phi_{j(k)}^{l(k)}$ are called the *elementary divisors* of T .

Remark: We defined *cyclic decomposition* as one in which the summands are *irreducible*. The requirement of irreducibility is satisfied automatically if the minimal polynomial is a “prime-power”, i.e., has the form Φ^m with irreducible Φ . If one omits this requirement and the minimal polynomial has several relatively prime factors, we no longer have uniqueness of the decomposition since the direct sum of cyclic subspaces with relatively prime minimal polynomials is itself cyclic.

EXERCISES FOR SECTION 5.6

V.6.1. Assume $\min P_{T,v} = \Phi^m$ with irreducible Φ . Let $u \in \text{span}[T, v]$, and assume $\Phi\text{-height}[u] = m$. Prove that $\text{span}[T, u] = \text{span}[T, v]$.

V.6.2. Give an example of two operators, T and S in $\mathcal{L}(\mathbb{C}^5)$, such that $\min P_T = \min P_S$ and $\chi_T = \chi_S$, and yet S and T are not similar.

V.6.3. Assume \mathbb{F} is a subfield of \mathbb{F}_1 . Let $B_1, B_2 \in \mathcal{M}(n, \mathbb{F})$ and assume that they are \mathbb{F}_1 -similar, i.e., $B_2 = C^{-1}B_1C$ for some invertible $C \in \mathcal{M}(n, \mathbb{F}_1)$. Prove that they are \mathbb{F} -similar.

5.7 THE JORDAN CANONICAL FORM

5.7.1 BASES AND CORRESPONDING MATRICES. Let (\mathcal{V}, T) be cyclic, that is $\mathcal{V} = \text{span}[T, v]$, and $\min P_T = \min P_{T,v} = \Phi^m$, with Φ irreducible of degree d . The cyclic decomposition provides several natural bases:

i. The (ordered) set $\{T^j v\}_{j=0}^{d(m-1)}$ is a basis, and the matrix of T with respect to this basis is the companion matrix of Φ^m .

ii. Another natural, and in some ways more useful, basis in this context is

$$(5.7.1) \quad \{T^k v\}_{k=0}^{d-1} \cup \{\Phi(T)T^k v\}_{k=0}^{d-1} \cup \dots \cup \{\Phi(T)^{m-1}T^k v\}_{k=0}^{d-1}$$

And the matrix A_{Φ^m} of T relative to this ordered basis consists of m copies of the companion matrix of Φ arranged on the diagonal, with 1's in the unused positions in the sub-diagonal.

If A_{Φ} is the companion matrix of Φ then the matrix A_{Φ^4} is

$$(5.7.2) \quad A_{\Phi^4} = \begin{pmatrix} \boxed{A_{\Phi}} & & & & & \\ & 1 & & & & \\ & & \boxed{A_{\Phi}} & & & \\ & & & 1 & & \\ & & & & \boxed{A_{\Phi}} & \\ & & & & & 1 \\ & & & & & & \boxed{A_{\Phi}} \end{pmatrix}$$

5.7.2 Consider the special case of linear Φ , which is the rule when the underlying field \mathbb{F} is algebraically closed, and in particular when $\mathbb{F} = \mathbb{C}$.

If $\Phi(x) = x - \lambda$ for some $\lambda \in \mathbb{F}$, then its companion matrix is 1×1 with λ its only entry.

Since now $d = 1$ the basis (5.7.1) is now simply $\{(T - \lambda)^j v\}_{j=0}^{m-1}$ and the matrix $A_{(x-\lambda)^m}$ in this case is the $m \times m$ matrix that has all its diagonal entries equal to λ , all the entries just below the diagonal (assuming $m > 1$) are equal to 1, and all the other entries are 0.

$$(5.7.3) \quad A_{(x-\lambda)^m} = \begin{bmatrix} \lambda & 0 & 0 & \dots & 0 & 0 \\ 1 & \lambda & 0 & \dots & 0 & 0 \\ 0 & 1 & \lambda & \dots & 0 & 0 \\ \vdots & \vdots & & & \vdots & \vdots \\ 0 & 0 & \dots & 1 & \lambda & 0 \\ 0 & 0 & \dots & 0 & 1 & \lambda \end{bmatrix}$$

5.7.3 THE JORDAN CANONICAL FORM. Consider a system (\mathcal{V}, T) such that all the irreducible factors of $\min P_T$ are linear, (in particular, an arbitrary system (\mathcal{V}, T) over \mathbb{C}). The prime-power factorization of $\min P_T$ is now ⁹

$$\min P_T = \prod_{\lambda \in \sigma(T)} (x - \lambda)^{m(\lambda)}$$

where $m(\lambda)$ is the multiplicity of λ in $\min P_T$.

The space $\mathcal{V}_\lambda = \ker((T - \lambda)^{m(\lambda)})$ is called the *generalized eigenspace* or *nilspace* of λ , see 5.3.4. The canonical decomposition of (\mathcal{V}, T) is given by:

$$(5.7.4) \quad \mathcal{V} = \bigoplus_{\lambda \in \sigma(T)} \mathcal{V}_\lambda.$$

For $\lambda \in \sigma(T)$, the restriction of $T - \lambda$ to \mathcal{V}_λ is nilpotent of height $m(\lambda)$. We apply to \mathcal{V}_λ the cyclic decomposition

$$\mathcal{V}_\lambda = \bigoplus \text{span}[T, v_j].$$

and take as basis in $\text{span}[T, v_j]$ the set $\{(T - \lambda)^s v_j\}_{s=0}^{h(v_j)-1}$, where $h(v_j)$ is the $(T - \lambda)$ -height of v_j .

The matrix of the restriction of T to each $\text{span}[T, v_j]$ has the form (5.7.3), the matrix $A_{T, \mathcal{V}_\lambda}$ of T on \mathcal{V}_λ is the *diagonal sum* of these, and the matrix of T on \mathcal{V} is the diagonal sum of $A_{T, \mathcal{V}_\lambda}$ for all $\lambda \in \sigma(T)$.

5.7.4 THE CANONICAL FORM FOR REAL VECTOR SPACES. When (\mathcal{V}, T) is defined over \mathbb{R} , the irreducible factors Φ of $\min P_T$ are either linear or quadratic, i.e., have the form

$$\Phi(x) = x - \lambda, \quad \text{or} \quad \Phi(x) = x^2 + 2bx + c \quad \text{with} \quad b^2 - c < 0.$$

The companion matrix in the quadratic case is

$$(5.7.5) \quad \begin{bmatrix} 0 & -c \\ 1 & -2b \end{bmatrix}.$$

⁹Recall that the *spectrum* of T is the set $\sigma(T) = \{\lambda_j\}$ of the eigenvalues of T , that is, the set of zeros of $\min P_T$.

(Over \mathbb{C} we have $x^2 + 2bx + c = (x - \lambda)(x - \bar{\lambda})$ with $\lambda = -b + \sqrt{b^2 - c}$, and the matrix similar to the diagonal matrix with λ and $\bar{\lambda}$ on the diagonal.)

EXERCISES FOR SECTION 5.7

V.7.1. Assume that v_1, \dots, v_k are eigenvectors of T with the associated eigenvalues $\lambda_1, \dots, \lambda_k$ all distinct. Prove that v_1, \dots, v_k are linearly independent.

V.7.2. Show that if we allow complex coefficients, the matrix (5.7.5) is similar to $\begin{bmatrix} \lambda & 0 \\ 0 & \bar{\lambda} \end{bmatrix}$ with $\lambda = -b + \sqrt{b^2 - c}$.

V.7.3. T is given by the matrix $A_T = \begin{bmatrix} 0 & 0 & 2 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$ acting on \mathbb{F}^3 .

a. What is the basic decomposition when $\mathbb{F} = \mathbb{C}$, when $\mathbb{F} = \mathbb{R}$, and when $\mathbb{F} = \mathbb{Q}$?

b. Prove that when $\mathbb{F} = \mathbb{Q}$ every non-zero vector is cyclic. Hence, every non-zero rational vector is cyclic when $\mathbb{F} = \mathbb{R}$ or \mathbb{C} .

c. What happens to the basic decomposition under the action of an operator S that commutes with T ?

d. Describe the set of matrices $A \in \mathcal{M}(3; \mathbb{F})$ that commute with A_T where $\mathbb{F} = \mathbb{C}$, \mathbb{R} , resp. \mathbb{Q} .

V.7.4. Prove that the matrix $\begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$ is not similar to a triangular matrix if the underlying field is \mathbb{R} , and is diagonalizable over \mathbb{C} . Why doesn't this contradict exercise **V.6.3**?

V.7.5. Let $A \in \mathcal{M}(n; \mathbb{C})$ such that $\{A^j : j \in \mathbb{N}\}$ is bounded (all the entries are uniformly bounded). Prove that all the eigenvalues of A are of absolute value not bigger than 1. Moreover, if $\lambda \in \sigma(A)$ and $|\lambda| = 1$, there are no ones under λ in the Jordan canonical form of A .

V.7.6. Let $A \in \mathcal{M}(n; \mathbb{C})$ such that $\{A^j : j \in \mathbb{Z}\}$ is bounded. Prove that A is diagonalizable, and all its eigenvalues have absolute value 1.

V.7.7. Let $T \in \mathcal{L}(\mathcal{V})$. Write $\chi_T = \prod \Phi_j^{m_j}$ with Φ_j irreducible, but not necessarily distinct, and m_j are the corresponding heights in the cyclic decomposition of the system.

Find a basis of the form (5.7.1) for each of the components. and describe the matrix of T relative to this basis.

V.7.8. Let A be the $m \times m$ matrix $A_{(x-\lambda)^m}$ defined in (5.7.3). Compute A^n for all $n > 1$.

5.8 FUNCTIONS OF AN OPERATOR

5.8.1 THEORETICAL. If $P = \sum a_j x^j$ is a polynomial with coefficients in \mathbb{F} , we defined $P(T)$ by

$$P(T) = \sum a_j T^j.$$

Is there a natural extension of the definition to a larger class of functions?

The map $P \mapsto P(T)$ is a homomorphism of $\mathbb{F}[x]$ onto a subalgebra of $\mathcal{L}(\mathcal{V})$. We can often extend the homomorphism to a bigger function space, but in most cases the range stays the same. The advantage will be in having a better match with the natural notation arising in applications.

Assume that the underlying field is either \mathbb{R} or \mathbb{C} .

Write $\min P_T(z) = \prod_{\lambda \in \sigma(T)} (z - \lambda)^{m(\lambda)}$ and observe that a necessary and sufficient condition for a polynomial Q to be divisible by $\min P_T$ is that Q be divisible by $(z - \lambda)^{m(\lambda)}$ for every $\lambda \in \sigma(T)$, that is, have a zero of order at least $m(\lambda)$ at λ . It follows that $P_1(T) = P_2(T)$ if, and only if, the Taylor expansion of the two polynomials are the same up to, and including, the term of order $m(\lambda) - 1$ at every $\lambda \in \sigma(T)$.

In particular, if $m(\lambda) = 1$ for all $\lambda \in \sigma(T)$ (i.e., if (\mathcal{V}, T) is semisimple) the condition $P_1(\lambda) = P_2(\lambda)$ for all $\lambda \in \sigma(T)$ is equivalent to $P_1(T) = P_2(T)$.

If f is an arbitrary numerical function defined on $\sigma(T)$, the only consistent way to define $f(T)$ is by setting $f(T) = P(T)$ where P is any polynomial that takes the same values as f at each point of $\sigma(T)$. This defines a homomorphism of the space of all numerical functions on $\sigma(T)$ onto the (the same old) subalgebra generated by T in $\mathcal{L}(\mathcal{V})$.

In the general (not necessarily semisimple) case, f needs to be defined and sufficiently differentiable¹⁰ in a neighborhood of every $\lambda \in \sigma(T)$, and we define $f(T) = P(T)$ where P is a polynomial whose Taylor expansion is the

¹⁰That is, differentiable at least $m(\lambda) - 1$ times.

same as that of f up to, and including, the term of order $m(\lambda) - 1$ at every $\lambda \in \sigma(T)$.

5.8.2 MORE PRACTICAL. The discussion in the previous subsection can only be put to use in practice if one has the complete spectral information about T —its minimal polynomial, its zeros including their multiplicities given explicitly.

One can often define $F(T)$ without explicit knowledge of this information if F holomorphic in a sufficiently large set, and always if F is an entire function, that is a function that admits a power series representation in the entire complex plane. This is done formally just as it was for polynomials, namely, for $F(z) = \sum_0^\infty a_n z^n$, we write $F(T) = \sum_0^\infty a_n T^n$. To verify that the definition makes sense we check the convergence of the series. Since $\mathcal{L}(\mathcal{V})$ is finite dimensional so that all the norms on it are equivalent, we can use a submultiplicative, “operator norm”, as defined by (2.6.1). This keeps the estimates a little cleaner since $\|T^n\| \leq \|T\|^n$, and if the radius of convergence of the series is bigger than $\|T\|$, the convergence of $\sum_0^\infty a_n T^n$ is assured.

Two simple examples:

- a.** Assume the norm used is submultiplicative, and $\|T\| < 1$, then $(I - T)$ is invertible and $(I - T)^{-1} = \sum_{n=0}^\infty T^n$.
- b.** Define $e^{aT} = \sum \frac{T^n}{n!}$. The series clearly convergent for every $T \in \mathcal{L}(\mathcal{V})$ and number a . As a function of the parameter a it has the usual properties of the exponential function. We can consider it as a function of T and check if $e^{(T+S)} = e^T e^S$. We find that the answer is yes if S and T commute, but *no* in general.

EXERCISES FOR SECTION 5.8

- V.8.1.** Prove that $e^{aT} e^{bT} = e^{(a+b)T}$.
 - V.8.2.** Prove that if S and T commute, then $e^{(T+S)} = e^T e^S$.
 - V.8.3.** Give an example of operators S and T such that $e^{(T+S)} \neq e^T e^S$.
- Hint:* Try for $e^S e^T \neq e^T e^S$.

Operators on inner-product spaces

6.1 INNER-PRODUCT SPACES

Inner-product spaces, are real or complex vector spaces endowed with an additional structure, called *inner-product*. The inner-product permits the introduction of a fair amount of geometry. Finite dimensional real inner-product spaces are often called *Euclidean spaces*. Complex inner-product spaces are also called *Unitary spaces*.

6.1.1 DEFINITION:

a. An *inner-product* on a *real* vector space \mathcal{V} is a symmetric, real-valued, positive definite bilinear form on \mathcal{V} . That is, a form satisfying

1. $\langle u, v \rangle = \langle v, u \rangle$
2. $\langle u, v \rangle$ is bilinear.
3. $\langle u, u \rangle \geq 0$, with $\langle u, u \rangle = 0$ if, and only if, $u = 0$.

b. An *inner-product* on a *complex* vector space \mathcal{V} is a Hermitian¹, complex-valued, positive definite, sesquilinear form on \mathcal{V} . That is, a form satisfying

1. $\langle u, v \rangle = \overline{\langle v, u \rangle}$
2. $\langle u, v \rangle$ is sesquilinear, that is, linear in u and skew linear in v :
 $\langle \lambda u, v \rangle = \lambda \langle u, v \rangle$ and $\langle u, \lambda v \rangle = \bar{\lambda} \langle u, v \rangle$.
3. $\langle u, u \rangle \geq 0$, with $\langle u, u \rangle = 0$ if and only if $u = 0$.

¹A complex-valued form φ is *Hermitian* if $\varphi(u, v) = \overline{\varphi(v, u)}$.

Notice that the sesquilinearity follows from the Hermitian symmetry, condition 1., combined with the assumption of linearity in the first entry.

EXAMPLES:

- a.** The classical Euclidean n -space \mathcal{E}^n is \mathbb{R}^n in which $\langle \mathbf{a}, \mathbf{b} \rangle = \sum a_j b_j$ where $\mathbf{a} = (a_1, \dots, a_n)$ and $\mathbf{b} = (b_1, \dots, b_n)$.
- b.** The space $C_{\mathbb{R}}([0, 1])$ of all continuous real-valued functions on $[0, 1]$. The inner-product is defined by $\langle f, g \rangle = \int f(x)g(x)dx$.
- c.** In \mathbb{C}^n for $\mathbf{a} = (a_1, \dots, a_n)$ and $\mathbf{b} = (b_1, \dots, b_n)$ we set $\langle \mathbf{a}, \mathbf{b} \rangle = \sum a_j \bar{b}_j$ which can be written as matrix multiplication: $\langle \mathbf{a}, \mathbf{b} \rangle = \mathbf{a} \bar{\mathbf{b}}^T$. If we consider the vector as columns, $\mathbf{a} = \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix}$ and $\mathbf{b} = \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix}$ then $\langle \mathbf{a}, \mathbf{b} \rangle = \bar{\mathbf{b}}^T \mathbf{a}$.
- d.** The space $C([0, 1])$ of all continuous complex-valued functions on $[0, 1]$. The inner-product is defined by $\langle f, g \rangle = \int_0^1 f(x) \overline{g(x)} dx$.

We shall reserve the notation \mathcal{H} for inner-product vector spaces.

6.1.2 Given an inner-product space \mathcal{H} we define a norm on it by:

$$(6.1.1) \quad \|v\| = \sqrt{\langle v, v \rangle}.$$

Lemma (Cauchy–Schwarz).

$$(6.1.2) \quad |\langle u, v \rangle| \leq \|u\| \|v\|.$$

PROOF: If v is a scalar multiple of u we have equality. If v, u are not proportional, then for $\lambda \in \mathbb{R}$,

$$0 < \langle u + \lambda v, u + \lambda v \rangle = \|u\|^2 + 2\lambda \Re \langle u, v \rangle + \lambda^2 \|v\|^2.$$

A quadratic polynomial with real coefficients and no real roots has negative discriminant, here $(\Re \langle u, v \rangle)^2 - \|u\|^2 \|v\|^2 < 0$.

For every τ with $|\tau| = 1$ we have $|\Re \langle \tau u, v \rangle| \leq \|u\| \|v\|$; take τ such that $\langle \tau u, v \rangle = |\langle u, v \rangle|$. ◀

The norm has the following properties:

- a.** Positivity: If $v \neq 0$ then $\|v\| > 0$; $\|0\| = 0$.
- b.** Homogeneity: $\|av\| = |a|\|v\|$ for scalars a and vectors v .
- c.** The triangle inequality: $\|v + u\| \leq \|v\| + \|u\|$.
- d.** The parallelogram law: $\|v + u\|^2 + \|v - u\|^2 = 2(\|v\|^2 + \|u\|^2)$.

Properties **a.** and **b.** are obvious. Property **c.** is equivalent to

$$\|v\|^2 + \|u\|^2 + 2\Re\langle v, u \rangle \leq \|v\|^2 + \|u\|^2 + 2\|v\|\|u\|,$$

which reduces to (6.1.2). The parallelogram law is obtained by “opening brackets” in the inner-products that correspond to the various $\| \|^2$.

The first three properties are common to all norms, whether defined by an inner-product or not. They imply that the norm can be viewed as length, and $\rho(u, v) = \|u - v\|$ has the properties of a *metric*.

The parallelogram law, on the other hand, is specific to, and in fact characteristic of, the norms defined by an inner-product.

A norm defined by an inner-product determines the inner-product, see exercises **VI.1.11** and **VI.1.12**.

6.1.3 ORTHOGONALITY. Let \mathcal{H} be an inner-product space. The vectors v, u in \mathcal{H} are (mutually) *orthogonal*, denoted $v \perp u$, if $\langle v, u \rangle = 0$. Observe that, since $\langle u, v \rangle = \overline{\langle v, u \rangle}$, the relation is symmetric: $u \perp v \iff v \perp u$.

The vector v is orthogonal to a set $A \subset \mathcal{H}$, denoted $v \perp A$, if it is orthogonal to every vector in A . If $v \perp A$, $u \perp A$, and $w \in A$ is arbitrary, then $\langle av + bu, w \rangle = a\langle v, w \rangle + b\langle u, w \rangle = 0$. It follows that for any set $A \subset \mathcal{H}$, the set² $A^\perp = \{v : v \perp A\}$ is a subspace of \mathcal{H} .

Similarly, if we assume that $v \perp A$, $w_1 \in A$, and $w_2 \in A$, we obtain $\langle v, aw_1 + bw_2 \rangle = \bar{a}\langle v, w_1 \rangle + \bar{b}\langle v, w_2 \rangle = 0$ so that $v \perp (\text{span}[A])$. In other words: $A^\perp = (\text{span}[A])^\perp$.

²This notation is consistent with 3.1.2, see 6.2.1 below.

A vector v is *normal* if $\|v\| = 1$. A sequence $\{v_1, \dots, v_m\}$ is *orthonormal* if

$$(6.1.3) \quad \langle v_i, v_j \rangle = \delta_{i,j} \quad (\text{i.e., } 1 \text{ if } i = j, \text{ and } 0 \text{ if } i \neq j);$$

that is, if the vectors v_j are normal and pairwise orthogonal.

Lemma. Let $\{u_1, \dots, u_m\}$ be orthonormal, $v, w \in \mathcal{H}$ arbitrary.

a. $\{u_1, \dots, u_m\}$ is linearly independent.

b. The vector $v_1 = v - \sum_1^m \langle v, u_j \rangle u_j$ is orthogonal to $\text{span}[u_1, \dots, u_m]$.

c. If $\{u_1, \dots, u_m\}$ is an orthonormal basis, then

$$(6.1.4) \quad v = \sum_1^m \langle v, u_j \rangle u_j.$$

d. Parseval's identity. If $\{u_1, \dots, u_m\}$ is an orthonormal basis for \mathcal{H} , then

$$(6.1.5) \quad \langle v, w \rangle = \sum_1^m \langle v, u_j \rangle \overline{\langle w, u_j \rangle}.$$

e. Bessel's inequality and identity. If $\{u_j\}$ is orthonormal then

$$(6.1.6) \quad \sum |\langle v, u_j \rangle|^2 \leq \|v\|^2.$$

If $\{u_1, \dots, u_m\}$ is an orthonormal basis for \mathcal{H} , then $\|v\|^2 = \sum_1^m |\langle v, u_j \rangle|^2$.

PROOF:

a. If $\sum a_j u_j = 0$ then $a_k = \langle \sum a_j u_j, u_k \rangle = 0$ for all $k \in [1, m]$.

b. $\langle v_1, u_k \rangle = \langle v, u_k \rangle - \langle v, u_k \rangle = 0$ for all $k \in [1, m]$; (skew-)linearity extends the orthogonality to linear combinations, that is to the span of $\{u_1, \dots, u_m\}$.

c. If the span is the entire \mathcal{H} , v_1 is orthogonal to itself, and so $v_1 = 0$.

$$\begin{aligned} \text{d.} \quad \langle v, w \rangle &= \langle \sum \langle v, u_j \rangle u_j, \sum \langle w, u_l \rangle u_l \rangle = \sum_{j,l} \langle v, u_j \rangle \overline{\langle w, u_l \rangle} \langle u_j, u_l \rangle \\ &= \sum_j \langle v, u_j \rangle \overline{\langle w, u_j \rangle} \end{aligned}$$

e. This is clearly weaker than (6.1.5). ◀

6.1.4 Proposition (Gram-Schmidt). *Let $\{v_1, \dots, v_m\}$ be independent. There exists an orthonormal $\{u_1, \dots, u_m\}$ such that for all $k \in [1, m]$,*

$$(6.1.7) \quad \text{span}[u_1, \dots, u_k] = \text{span}[v_1, \dots, v_k].$$

PROOF: (By induction on m). The independence of $\{v_1, \dots, v_m\}$ implies that $v_1 \neq 0$. Write $u_1 = v_1/\|v_1\|$. Then u_1 is normal and (6.1.7) is satisfied for $k = 1$.

Assume that $\{u_1, \dots, u_l\}$ is orthonormal and that (6.1.7) is satisfied for $k \leq l$. Since $v_{l+1} \notin \text{span}[\{v_1, \dots, v_l\}]$ the vector

$$\tilde{v}_{l+1} = v_{l+1} - \sum_{j=1}^l \langle v_{l+1}, u_j \rangle u_j$$

is non-zero and we set $u_{l+1} = \tilde{v}_{l+1}/\|\tilde{v}_{l+1}\|$. ◀

One immediate corollary is: every finite dimensional \mathcal{H} has an orthonormal basis. Another is that every orthonormal sequence $\{u_j\}_1^k$ can be completed to an orthonormal basis. For this we observe that $\{u_j\}_1^k$ is independent, complete it to a basis, apply the Gram-Schmidt process and notice that it does not change the vectors u_j , $1 \leq j \leq k$.

If $\mathcal{W} \subset \mathcal{H}$ is a subspace, $\{v_j\}_1^n$ is a basis for \mathcal{H} such that $\{v_j\}_1^m$ is a basis for \mathcal{W} , then the basis $\{u_j\}_1^n$ obtained by the Gram-Schmidt process splits into two: $\{u_j\}_1^m \cup \{u_j\}_{m+1}^n$, where $\{u_j\}_1^m$ is an o.n. basis³ for \mathcal{W} and $\{u_j\}_{m+1}^n$ is one for \mathcal{W}^\perp . This gives a direct sum (in fact, orthogonal) decomposition $\mathcal{H} = \mathcal{W} \oplus \mathcal{W}^\perp$.

The map

$$(6.1.8) \quad \pi_{\mathcal{W}}: v \mapsto \sum_1^m \langle v, u_j \rangle u_j$$

is called the *orthogonal projection onto \mathcal{W}* . It depends only on \mathcal{W} and not on the particular basis we started from. In fact, if $v = v_1 + v_2 = u_1 + u_2$ with v_1 and u_1 in \mathcal{W} , and both v_2 and u_2 in \mathcal{W}^\perp , we have

$$v_1 - u_1 = u_2 - v_2 \in \mathcal{W} \cap \mathcal{W}^\perp$$

³“o.n.” is short for “orthonormal”

which means $v_1 - u_1 = u_2 - v_2 = 0$.

6.1.5 The definition of the distance $\rho(v_1, v_2)$ ($= \|v_1 - v_2\|$) between two vectors, extends to that of the *distance between a point* ($v \in \mathcal{H}$) and a set ($E \subset \mathcal{H}$) by setting $\rho(v, E) = \inf_{u \in E} \rho(v, u)$. The *distance between two sets*, $E_j \subset \mathcal{H}$ $j = 1, 2$, is defined by

$$(6.1.9) \quad \rho(E_1, E_2) = \inf\{\|v_1 - v_2\| : v_j \in E_j\}.$$

Proposition. Let $\mathcal{W} \subset \mathcal{H}$ be a subspace, and $v \in \mathcal{H}$. Then

$$\rho(v, \mathcal{W}) = \|v - \pi_{\mathcal{W}}v\|.$$

In other words, $\pi_{\mathcal{W}}v$ is the vector in \mathcal{W} closest to v .

The proof is left as an exercise (VI.1.5) below).

EXERCISES FOR SECTION 6.1

VI.1.1. Let \mathcal{V} be a finite dimensional real or complex space, and $\{v_1, \dots, v_n\}$ a basis. Explain: “declaring $\{v_1, \dots, v_n\}$ to be orthonormal defines an inner-product on \mathcal{V} ”.

VI.1.2. Prove that if \mathcal{H} is a complex inner-product space and $T \in \mathcal{L}(\mathcal{H})$, there exists an orthonormal basis for \mathcal{H} such that the matrix of T with respect to this basis is triangular.

Hint: See corollary 5.1.6.

VI.1.3. **a.** Let \mathcal{H} be a real inner-product space. The vectors v, u are mutually orthogonal if, and only if, $\|v + u\|^2 = \|v\|^2 + \|u\|^2$.

b. If \mathcal{H} is a complex inner-product space, $v, u \in \mathcal{H}$, then $\|v + u\|^2 = \|v\|^2 + \|u\|^2$ is necessary, but not sufficient, for $v \perp u$.

Hint: Connect to the condition “ $\langle u, v \rangle$ purely imaginary”.

c. If \mathcal{H} is a complex inner-product space, and $v, u \in \mathcal{H}$, the condition: For all $a, b \in \mathbb{C}$, $\|av + bu\|^2 = |a|^2\|v\|^2 + |b|^2\|u\|^2$ is necessary and sufficient for $v \perp u$.

d. Let \mathcal{V} and \mathcal{U} be subspaces of \mathcal{H} . Prove that $\mathcal{V} \perp \mathcal{U}$ if, and only if, for $v \in \mathcal{V}$ and $u \in \mathcal{U}$, $\|v + u\|^2 = \|v\|^2 + \|u\|^2$.

e. The set $\{v_1, \dots, v_m\}$ is orthonormal if, and only if $\|\sum a_j v_j\|^2 = \sum |a_j|^2$ for all choices of scalars $a_j, j = 1, \dots, m$. (Here \mathcal{H} is either real or complex.)

VI.1.4. Show that the map $\pi_{\mathcal{W}}$ defined in (6.1.8) is an idempotent linear operator⁴ and is independent of the particular basis used in its definition.

VI.1.5. Prove proposition 6.1.5.

VI.1.6. Let $E_j = v_j + \mathcal{W}_j$ be affine subspaces in \mathcal{H} . What is $\rho(E_1, E_2)$?

VI.1.7. Show that the sequence $\{u_1, \dots, u_m\}$ obtained by the Gram-Schmidt procedure is essentially unique: each u_j is unique up to multiplication by a number of modulus 1.

Hint: If $\{v_1, \dots, v_m\}$ is independent, $\mathcal{W}_k = \text{span}[\{v_1, \dots, v_k\}]$, $k = 0, \dots, m-1$, then u_j is $c\pi_{\mathcal{W}_{j-1}^\perp} v_j$, with $|c| = \|\pi_{\mathcal{W}_{j-1}^\perp} v_j\|^{-1}$.

VI.1.8. Over \mathbb{C} : Every matrix is unitarily equivalent to a triangular matrix.

VI.1.9. Let $A \in \mathcal{M}(n, \mathbb{C})$ and assume that its rows w_j , considered as vectors in \mathbb{C}^n are pairwise orthogonal. Prove that $A\overline{A^T}$ is a diagonal matrix, and conclude that $|\det A| = \prod \|w_j\|$.

VI.1.10. Let $\{v_1, \dots, v_n\} \subset \mathbb{C}^n$ be the rows of the matrix A . Prove *Hadamard's inequality*:

$$(6.1.10) \quad |\det A| \leq \prod \|v_j\|$$

Hint: Write $\mathcal{W}_k = \text{span}[\{v_1, \dots, v_k\}]$, $k = 0, \dots, n-1$, $w_j = \pi_{\mathcal{W}_{j-1}^\perp} v_j$, and apply the previous problem.

VI.1.11. Prove that in a real inner-product space, the inner-product is determined by the norm: (*polarization formula over \mathbb{R}*)

$$(6.1.11) \quad \langle u, v \rangle = \frac{1}{4} (\|u+v\|^2 - \|u-v\|^2)$$

VI.1.12. Prove: In a complex inner-product space, the inner-product is determined by the norm, in fact, (*polarization formula over \mathbb{C}*)

$$(6.1.12) \quad \langle u, v \rangle = \frac{1}{4} (\|u+v\|^2 - \|u-v\|^2 + i\|u+iv\|^2 - i\|u-iv\|^2).$$

⁴An operator T is *idempotent* if $T^2 = T$.

VI.1.13. Show that the polarization formula (6.1.12) does not depend on positivity, to wit, define the *Hermitian quadratic form* associated with a sesquilinear Hermitian form ψ (on a vector space over \mathbb{C} or a subfield thereof) by:

$$(6.1.13) \quad Q(v) = \psi(v, v).$$

Prove

$$(6.1.14) \quad \psi(u, v) = \frac{1}{4}(Q(u+v) - Q(u-v) + iQ(u+iv) - iQ(u-iv)).$$

VI.1.14. A bilinear form φ on a vector space \mathcal{V} over a field of characteristic $\neq 2$, can be expressed uniquely as a sum of a symmetric and an alternating form: $\varphi = \varphi_{sym} + \varphi_{alt}$ where $2\varphi_{sym}(v, u) = \varphi(v, u) + \varphi(u, v)$ and $2\varphi_{alt}(v, u) = \varphi(v, u) - \varphi(u, v)$.

The *quadratic form* associated with φ is, by definition $q(v) = \varphi(v, v)$. Show that q determines φ_{sym} , in fact

$$(6.1.15) \quad \varphi_{sym}(v, u) = \frac{1}{2}(q(v+u) - q(v) - q(u)).$$

6.2 DUALITY AND THE ADJOINT.

6.2.1 \mathcal{H} AS ITS OWN DUAL. The inner-product defined in \mathcal{H} associates with every vector $u \in \mathcal{H}$ the linear functional $\varphi_u: v \mapsto \langle v, u \rangle$. In fact every linear functional is obtained this way:

Theorem. Let φ be a linear functional on a finite dimensional inner-product space \mathcal{H} . Then there exist a unique $u \in \mathcal{H}$ such that $\varphi = \varphi_u$, that is,

$$(6.2.1) \quad \varphi(v) = \langle v, u \rangle$$

for all $v \in \mathcal{H}$.

PROOF: Let $\{w_j\}$ be an orthonormal basis in \mathcal{H} , and let $u = \sum \overline{\varphi(w_j)}w_j$. For every $v \in \mathcal{H}$ we have $v = \sum \langle v, w_j \rangle w_j$, and by Parseval's identity, 6.1.3,

$$(6.2.2) \quad \varphi(v) = \sum \langle v, w_j \rangle \varphi(w_j) = \langle v, u \rangle. \quad \blacktriangleleft$$

In particular, an orthonormal basis in \mathcal{H} is its own dual basis.

6.2.2 THE ADJOINT OF AN OPERATOR. Once we identify \mathcal{H} with its dual space, the adjoint of an operator $T \in \mathcal{L}(\mathcal{H})$ is again an operator on \mathcal{H} . We repeat the argument of 3.2 in the current context. Given $u \in \mathcal{H}$, the mapping $v \mapsto \langle Tv, u \rangle$ is a linear functional and therefore equal to $v \mapsto \langle v, w \rangle$ for some $w \in \mathcal{H}$. We write $T^*u = w$ and check that $u \mapsto w$ is linear. In other words T^* is a linear operator on \mathcal{H} , characterized by

$$(6.2.3) \quad \langle Tv, u \rangle = \langle v, T^*u \rangle.$$

Lemma. For $T \in \mathcal{L}(\mathcal{H})$, $(T^*)^* = T$.

PROOF: $\langle v, (T^*)^*u \rangle = \langle T^*v, u \rangle = \overline{\langle u, T^*v \rangle} = \overline{\langle Tu, v \rangle} = \langle v, Tu \rangle$. ◀

Proposition 3.2.4 reads in the present context as

Proposition. For $T \in \mathcal{L}(\mathcal{H})$, $\text{range}(T) = (\ker(T^*))^\perp$.

PROOF: $\langle Tx, y \rangle = \langle x, T^*y \rangle$ so that $y \perp \text{range}(T)$ if, and only if $y \in \ker(T^*)$. ◀

6.2.3 THE ADJOINT OF A MATRIX.

DEFINITION: The *adjoint of a matrix* $A \in \mathcal{M}(n, \mathbb{C})$ is the matrix $A^* = \overline{A}^T$. A is *self-adjoint*, aka *Hermitian*, if $A = A^*$, that is, if $a_{ij} = \overline{a_{ji}}$ for all i, j .

If $A = A_{T, \mathbf{v}}$ is the matrix of an operator T relative to an orthonormal basis \mathbf{v} , see 2.4.3, and $A_{T^*, \mathbf{v}}$ is the matrix of T^* relative to the same basis, then, writing the inner-product as matrix multiplication:

$$(6.2.4) \quad \langle Tv, u \rangle = \overline{u}^T Av = \overline{(\overline{A}^T u)}^T v, \quad \text{and} \quad \langle v, T^*u \rangle = (\overline{A^*u})^T v,$$

we obtain $A_{T^*, \mathbf{v}} = (A_{T, \mathbf{v}})^*$. *The matrix of the adjoint is the adjoint of the matrix.*

In particular, T is self-adjoint if, and only if, $A_{T, \mathbf{v}}$, for some (every) orthonormal basis \mathbf{v} , is self-adjoint.

EXERCISES FOR SECTION 6.2

VI.2.1. Prove that if $T, S \in \mathcal{L}(\mathcal{H})$, then $(ST)^* = T^*S^*$.

VI.2.2. Prove that if $T \in \mathcal{L}(\mathcal{H})$, then $\ker(T^*T) = \ker(T)$.

VI.2.3. Prove that χ_{T^*} is the complex conjugate of χ_T .

VI.2.4. If $Tv = \lambda v$, $T^*u = \mu u$, and $\mu \neq \bar{\lambda}$, then $\langle v, u \rangle = 0$.

VI.2.5. Rewrite the proof of Theorem 6.2.1 along these lines: If $\ker(\varphi) = \mathcal{H}$ then $\varphi = 0$ and $u^* = 0$. Otherwise, $\dim \ker(\varphi) = \dim \mathcal{H} - 1$ and $(\ker(\varphi))^\perp \neq \emptyset$. Take any non-zero $\tilde{u} \in (\ker(\varphi))^\perp$ and set $u^* = c\tilde{u}$ where the constant c is the one that guarantees $\langle \tilde{u}, c\tilde{u} \rangle = \varphi(\tilde{u})$, that is $\bar{c} = \|\tilde{u}\|^{-2}\varphi(\tilde{u})$.

6.3 UNITARY AND ORTHOGONAL OPERATORS

We have mentioned that the norm in \mathcal{H} defines a *metric*, the distance between the vectors v and u given by $\rho(v, u) = \|v - u\|$. Mappings that preserve a metric are called *isometries* (of the given metric). Operators $U \in \mathcal{L}(\mathcal{H})$ which are isometries, that is such that $\|Uv\| = \|v\|$ for all $v \in \mathcal{H}$ are called *unitary operators* when \mathcal{H} is complex, and *orthogonal* when \mathcal{H} is real. The operator U is unitary if

$$\|Uv\|^2 = \langle Uv, Uv \rangle = \langle v, U^*Uv \rangle = \langle v, v \rangle$$

which is equivalent to $U^*U = I$ or $U^* = U^{-1}$.

Proposition. *Let \mathcal{H} be an inner-product space, $T \in \mathcal{L}(\mathcal{H})$. The following statements are equivalent:*

- a. T is unitary;
- b. T maps some orthonormal basis onto an orthonormal basis;
- c. T maps every orthonormal basis onto an orthonormal basis.

The columns of the matrix of a unitary operator U relative to an orthonormal basis $\{v_j\}$, are the coefficient vectors of Uv_j and, by Parseval's identity 6.1.3, are orthonormal in \mathbb{C}^n (resp. \mathbb{R}^n). Such matrices (with orthonormal columns) are called *unitary* when the underlying field is \mathbb{C} , and *orthogonal* when the field is \mathbb{R} .

The set $U(n) \subset \mathcal{M}(n, \mathbb{C})$ of unitary $n \times n$ matrices is a group under matrix multiplication. It is called the *unitary group*.

The set $O(n) \subset \mathcal{M}(n, \mathbb{R})$ of orthogonal $n \times n$ matrices is a group under matrix multiplication. It is called the *orthogonal group*.

DEFINITION: The matrices $A, B \in \mathcal{M}(n)$ are *unitarily equivalent* if there exists $U \in U(n)$ such that $A = U^{-1}BU$.

The matrices $A, B \in \mathcal{M}(n)$ are *orthogonally equivalent* if there exists $C \in O(n)$ such that $A = C^{-1}BC$.

The added condition here, compared to similarity, is that the conjugating matrix U , resp. O , be unitary, resp. orthogonal, and not just invertible.

EXERCISES FOR SECTION 6.3

VI.3.1. Prove that the set of *rows* of a unitary matrix is orthonormal.

VI.3.2. Prove that the spectrum of a unitary operator is contained in the unit circle $\{z : |z| = 1\}$.

VI.3.3. An operator T whose spectrum is contained in the unit circle is similar to a unitary operator if, and only if, it is semisimple.

VI.3.4. An operator T whose spectrum is contained in the unit circle is unitary if, and only if, it is semisimple and eigenvectors corresponding to distinct eigenvalues are mutually orthogonal.

VI.3.5. Let $T \in \mathcal{L}(\mathcal{H})$ be invertible and assume that $\|T^j\|$ is uniformly bounded for $j \in \mathbb{Z}$. Prove that T is similar to a unitary operator.

6.4 SELF-ADJOINT OPERATORS

6.4.1 **DEFINITION:** An operator $T \in \mathcal{L}(\mathcal{H})$ is *self-adjoint* if it coincides with its adjoint: $T = T^*$, (that is, if $\langle Tu, v \rangle = \langle u, Tv \rangle$ for every $u, v \in \mathcal{H}$).

For every $T \in \mathcal{L}(\mathcal{H})$, the operators $\Re T = \frac{1}{2}(T + T^*)$, $\Im T = \frac{1}{2i}(T - T^*)$, T^*T , and TT^* are all self-adjoint.

Proposition. Assume that T is self-adjoint on \mathcal{H} .

- a. $\sigma(T) \subset \mathbb{R}$.
- b. If $\mathcal{W} \subset \mathcal{H}$ is T -invariant then so is \mathcal{W}^\perp (the orthogonal complement of \mathcal{W}). In particular, every T -invariant subspace is reducing, so that T is semisimple.
- c. If $\mathcal{W} \subset \mathcal{H}$ is T -invariant then $T|_{\mathcal{W}}$, the restriction of T to \mathcal{W} , is self-adjoint.

PROOF:

a. If $\lambda \in \sigma(T)$ and v is a corresponding eigenvector, then

$$\lambda\|v\|^2 = \langle Tv, v \rangle = \langle v, Tv \rangle = \bar{\lambda}\|v\|^2, \text{ so that } \lambda = \bar{\lambda}.$$

b. If $v \in \mathcal{W}^\perp$ then, for any $w \in \mathcal{W}$, $\langle Tv, w \rangle = \langle v, Tw \rangle = 0$ (since $Tw \in \mathcal{W}$), so that $Tv \in \mathcal{W}^\perp$.

c. The condition $\langle Tw_1, w_2 \rangle = \langle w_1, Tw_2 \rangle$ is valid when $w_j \in \mathcal{W}$ since it holds for all vectors in \mathcal{H} .

◀

6.4.2 Part **b.** of the proposition implies that for self-adjoint operators T the generalized eigenspaces \mathcal{H}_λ , $\lambda \in \sigma(T)$, are not *generalized*, they are simply kernels: $\mathcal{H}_\lambda = \ker(T - \lambda)$. The Canonical Decomposition Theorem reads in this context:

Proposition. *Assume T self-adjoint. Then $\mathcal{H} = \bigoplus_{\lambda \in \sigma(T)} \ker(T - \lambda)$.*

6.4.3 The final improvement we bring to the Canonical Decomposition Theorem for self-adjoint operators is the fact that the eigenspaces corresponding to distinct eigenvalues are mutually orthogonal: if T is self-adjoint, $Tv_1 = \lambda_1 v_1$, $Tv_2 = \lambda_2 v_2$, and $\lambda_1 \neq \lambda_2$, then⁵,

$$\lambda_1 \langle v_1, v_2 \rangle = \langle Tv_1, v_2 \rangle = \langle v_1, Tv_2 \rangle = \bar{\lambda}_2 \langle v_1, v_2 \rangle = \lambda_2 \langle v_1, v_2 \rangle,$$

so that $\langle v_1, v_2 \rangle = 0$.

Theorem (The spectral theorem for self-adjoint operators). *Let \mathcal{H} be an inner-product space and T a self-adjoint operator on \mathcal{H} . Then $\mathcal{H} = \bigoplus_{\lambda \in \sigma(T)} \mathcal{H}_\lambda$ where $T_{\mathcal{H}_\lambda}$, the restriction of T to \mathcal{H}_λ , is multiplication by λ , and $\mathcal{H}_{\lambda_1} \perp \mathcal{H}_{\lambda_2}$ when $\lambda_1 \neq \lambda_2$.*

An equivalent formulation of the theorem is:

⁵Remember that $\lambda_2 \in \mathbb{R}$.

Theorem (Variant). *Let \mathcal{H} be an inner-product space and T a self-adjoint operator on \mathcal{H} . Then \mathcal{H} has an orthonormal basis all whose elements are eigenvectors for T .*

Denote by π_λ the orthogonal projection on \mathcal{H}_λ . The theorem states:

$$(6.4.1) \quad I = \sum_{\lambda \in \sigma(T)} \pi_\lambda, \quad \text{and} \quad T = \sum_{\lambda \in \sigma(T)} \lambda \pi_\lambda.$$

The decomposition $\mathcal{H} = \bigoplus_{\lambda \in \sigma(T)} \mathcal{H}_\lambda$ is often referred to as *the spectral decomposition induced by T on \mathcal{H}* . The representation of T as $\sum_{\lambda \in \sigma(T)} \lambda \pi_\lambda$ is its *spectral decomposition*.

6.4.4 If $\{u_1, \dots, u_n\}$ is an orthonormal basis whose elements are eigenvectors for T , say $Tu_j = \lambda_j u_j$, then

$$(6.4.2) \quad Tv = \sum \lambda_j \langle v, u_j \rangle u_j$$

for all $v \in \mathcal{H}$. Consequently, writing $a_j = \langle v, u_j \rangle$ and $v = \sum a_j u_j$,

$$(6.4.3) \quad \langle Tv, v \rangle = \sum \lambda_j |a_j|^2 \quad \text{and} \quad \|Tv\|^2 = \sum |\lambda_j|^2 |a_j|^2.$$

Observations. *Assume T self-adjoint.*

a. $\|T\| = \max_{\lambda \in \sigma(T)} |\lambda|.$

b. *If $\|T\| \leq 1$. Then there exists a unitary operator U that commutes with T , such that $T = \frac{1}{2}(U + U^*)$.*

PROOF: **a.** If λ_m is an eigenvalue with maximal absolute value in $\sigma(T)$, then $\|T\| \geq \|Tu_m\| = \max_{\lambda \in \sigma(T)} |\lambda|$. Conversely, by (6.4.3),

$$\|Tv\|^2 = \sum |\lambda_j|^2 |\langle v, u_j \rangle|^2 \leq \max |\lambda_j|^2 \sum |\langle v, u_j \rangle|^2 = \max |\lambda_j|^2 \|v\|^2.$$

b. By part **a.**, $\sigma(T) \subset [-1, 1]$. For $\lambda_j \in \sigma(T)$ write $\zeta_j = \lambda_j + i\sqrt{1 - \lambda_j^2}$, so that $\lambda_j = \Re \zeta_j$ and $|\zeta_j| = 1$. Define: $Uv = \sum \zeta_j \langle v, u_j \rangle u_j$. ◀

6.4.5 Theorem (Spectral theorem for Hermitian/symmetric matrices). *Every Hermitian matrix in $\mathcal{M}(n, \mathbb{C})$ is unitarily equivalent to a diagonal matrix. Every symmetric matrix in $\mathcal{M}(n, \mathbb{R})$ is orthogonally equivalent to a diagonal matrix.*

PROOF: A Hermitian matrix $A \in \mathcal{M}(n, \mathbb{C})$ is self-adjoint (i.e., the operator on \mathbb{C}^n of multiplication by A is self-adjoint). If the underlying field is \mathbb{R} the condition is being symmetric. In either case, theorem 6.4.3 guarantees that the standard \mathbb{C}^n , resp. \mathbb{R}^n , has an orthonormal basis $\{v_j\}$ all whose elements are eigenvectors for the operator of multiplication by A .

The matrix C of transition from the standard basis to $\{v_j\}$ is unitary, resp. orthogonal, and $CAC^{-1} = CAC^*$ is diagonal. ◀

6.4.6 COMMUTING SELF-ADJOINT OPERATORS.

Let T is self-adjoint, $\mathcal{H} = \bigoplus_{\lambda \in \sigma(T)} \mathcal{H}_\lambda$. If S commutes with T , then S maps each \mathcal{H}_λ into itself. Since the subspaces \mathcal{H}_λ are mutually orthogonal, if S is self-adjoint then so is its restriction to every \mathcal{H}_λ , and we can apply Theorem 6.4.3 to each one of these restrictions and obtain, in each, an orthonormal basis made up of eigenvectors of S . Since every vector in \mathcal{H}_λ is an eigenvector for T we obtained an orthonormal basis each of whose elements is an eigenvector both for T and for S . We now have the decomposition

$$\mathcal{H} = \bigoplus_{\lambda \in \sigma(T), \mu \in \sigma(S)} \mathcal{H}_{\lambda, \mu},$$

where $\mathcal{H}_{\lambda, \mu} = \ker(T - \lambda) \cap \ker(S - \mu)$.

By induction on the number of operators we obtain the following theorem.

Theorem. *Let \mathcal{H} be a finite dimensional inner-product space, and $\{T_j\}$ commuting self-adjoint operators on \mathcal{H} . Then there exists an orthonormal basis $\{u_k\}$ in \mathcal{H} such that each u_k is an eigenvector of every T_j .*

EXERCISES FOR SECTION 6.4

VI.4.1. Let $T \in \mathcal{L}(\mathcal{H})$ be self-adjoint, let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ be its eigenvalues and $\{u_j\}$ the corresponding orthonormal eigenvectors. Prove the “minmax principle”:

$$(6.4.4) \quad \lambda_l = \min_{\dim \mathcal{W}=l} \max_{v \in \mathcal{W}, \|v\|=1} \langle Tv, v \rangle.$$

Hint: Every l -dimensional subspace intersects $\text{span}\{\{u_j\}_{j=l}^n\}$, see 1.2.5.

VI.4.2. Let $\mathcal{W} \subset \mathcal{H}$ be a subspace, and $\pi_{\mathcal{W}}$ the orthogonal projection onto \mathcal{W} . Prove that if T is self-adjoint on \mathcal{H} , then $\pi_{\mathcal{W}}T$ is self-adjoint on \mathcal{W} .

VI.4.3. Use exercise **VI.2.2** to prove that a self-adjoint operator T on \mathcal{H} is semisimple (Lemma 6.4.1, part **b**).

6.5 NORMAL OPERATORS.

DEFINITION: An operator $T \in \mathcal{L}(\mathcal{H})$ is *normal* if it commutes with its adjoint: $TT^* = T^*T$.

Self-adjoint operators are clearly normal. Unitary operators are normal since for unitary U we have $U^*U = UU^* = I$.

If T is normal then $S = TT^* = T^*T$ is self-adjoint.

6.5.1 THE SPECTRAL THEOREM FOR NORMAL OPERATORS. For every operator $T \in \mathcal{L}(\mathcal{H})$, the operators

$$T_1 = \Re T = \frac{1}{2}(T + T^*) \quad \text{and} \quad T_2 = \Im T = \frac{1}{2i}(T - T^*)$$

are both self-adjoint, and $T = (T_1 + iT_2)$. T is normal if, and only if, T_1 and T_2 commute.

Theorem. Let $T \in \mathcal{L}(\mathcal{H})$ be normal. Then there is an orthonormal basis $\{u_k\}$ of \mathcal{H} such that every u_k is an eigenvector for T .

PROOF: As above, write $T_1 = T + T^*$, $T_2 = -i(T - T^*)$. Since T_1 and T_2 are commuting self-adjoint operators, Theorem 6.4.6 guarantees the existence of an orthonormal basis $\{u_k\} \subset \mathcal{H}$ such that each u_k is an eigenvector of both T_1 and T_2 . If $T_j = \sum_k t_{j,k} \pi_{u_k}$, $j = 1, 2$, then

$$(6.5.1) \quad T = \sum_k (t_{1,k} + it_{2,k}) \pi_{u_k},$$

and the vectors u_k are eigenvectors of T with eigenvalues $(t_{1,k} + it_{2,k})$. ◀

6.5.2 A subalgebra $\mathcal{A} \subset \mathcal{L}(\mathcal{H})$ is *self-adjoint* if $S \in \mathcal{A}$ implies that $S^* \in \mathcal{A}$.

Theorem. *Let $\mathcal{A} \subset \mathcal{L}(\mathcal{H})$ be a self-adjoint commutative subalgebra. Then there is an orthonormal basis $\{u_k\}$ of \mathcal{H} such that every u_k is a common eigenvector of every $T \in \mathcal{A}$.*

PROOF: The elements of \mathcal{A} are normal and \mathcal{A} is spanned by the self-adjoint elements it contains. Apply Theorem 6.4.6. ◀

EXERCISES FOR SECTION 6.5

VI.5.1. If S is normal (or just semisimple), a necessary and sufficient condition for an operator Q to commute with S is that all the eigenspaces of S be Q -invariant.

VI.5.2. If S is normal and Q commutes with S it commutes also with S^* .

VI.5.3. If $T \in \mathcal{L}(\mathcal{H})$ and $\{T^n\}_{n \in \mathbb{Z}}$ is bounded, then T is similar to a unitary operator. ($T = S^{-1}US$)

VI.5.4. Prove without using the spectral theorems:

- For any $Q \in \mathcal{L}(\mathcal{H})$, $\ker(Q^*Q) = \ker(Q)$.
- If S is normal, then $\ker(S) = \ker(S^*)$.
- If T is self-adjoint, then $\ker(T) = \ker(T^2)$.
- If S is normal, then $\ker(S) = \ker(S^2)$.
- Normal operators are semisimple.

VI.5.5. Prove without using the spectral theorems: If S is normal. then

- For all $v \in \mathcal{H}$, $\|S^*v\| = \|Sv\|$.
- If $Sv = \lambda v$ then $S^*v = \bar{\lambda}v$.

VI.5.6. If S is normal then S and S^* have the same eigenvectors with the corresponding eigenvalues complex conjugate. In particular, $\sigma(S^*) = \overline{\sigma(S)}$. If $T_1 = \Re S = \frac{S+S^*}{2}$ and $T_2 = \Im S = \frac{S-S^*}{2i}$, then if $Sv = \lambda v$, we have $T_1v = \Re \lambda v$, and $T_2v = \Im \lambda v$.

VI.5.7. Prove that the dimension of any commutative self-adjoint subalgebra of $\mathcal{L}(\mathcal{H})$ is bounded by $\dim \mathcal{H}$, and every such algebra is contained in a commutative self-adjoint subalgebra of $\mathcal{L}(\mathcal{H})$ of dimension $\dim \mathcal{H}$.

6.6 POSITIVE OPERATORS.

6.6.1 A self-adjoint operator S is *nonnegative*, written $S \geq 0$, if

$$(6.6.1) \quad \langle Sv, v \rangle \geq 0$$

for every $v \in \mathcal{H}$. S is *positive*, written $S > 0$, if, in addition, $\langle Sv, v \rangle = 0$ only for $v = 0$.

6.6.2 Lemma. A self-adjoint operator S is nonnegative, resp. positive, if, and only if, $\sigma(S) \subset [0, \infty)$, resp $\sigma(S) \subset (0, \infty)$.

PROOF: Use the spectral decomposition $S = \sum_{\lambda \in \sigma(T)} \lambda \pi_\lambda$.

We have $\langle Sv, v \rangle = \sum \lambda_j \|\pi_j v\|^2$, which, clearly, is nonnegative for all $v \in \mathcal{H}$ if, and only if, $\lambda \geq 0$ for all $\lambda \in \sigma(S)$. If $\sigma(S) \subset (0, \infty)$ and $\|v\|^2 = \sum \|\pi_\lambda v\|^2 > 0$ then $\langle Sv, v \rangle > 0$. If $0 \in \sigma(S)$ take $v \in \ker(S)$, then $\langle Sv, v \rangle = 0$ and S is not positive. ◀

6.6.3 PARTIAL ORDERS ON THE SET OF SELF-ADJOINT OPERATORS.

Let T and S be self-adjoint operators. The notions of positivity and nonnegativity define partial orders, “ $>$ ” and “ \geq ” on the set of self-adjoint operators on \mathcal{H} . We write $T > S$ if $T - S > 0$, and $T \geq S$ if $T - S \geq 0$.

Proposition. Let T and S be self-adjoint operators on \mathcal{H} , and assume $T \geq S$. Let $\sigma(T) = \{\lambda_j\}$ and $\sigma(S) = \{\mu_j\}$, both arranged in nondecreasing order. Then $\lambda_j \geq \mu_j$ for $j = 1, \dots, n$.

PROOF: Use the minmax principle, exercise VI.4.1:

$$\lambda_j = \min_{\dim \mathcal{W}=j} \max_{v \in \mathcal{W}, \|v\|=1} \langle Tv, v \rangle \geq \min_{\dim \mathcal{W}=j} \max_{v \in \mathcal{W}, \|v\|=1} \langle Sv, v \rangle = \mu_j$$

◀

Remark: The condition “ $\lambda_j \geq \mu_j$ for $j = 1, \dots, n$ ” is necessary but, even if T and S commute, *not sufficient* for $T \geq S$ (unless $n = 1$). As example consider : $\{v_1, \dots, v_n\}$ is an orthonormal basis, T defined by: $Tv_j = 2jv_j$; and S defined by: $Sv_1 = 3v_1, Sv_j = v_j$ for $j > 1$. The eigenvalues of $T - S$ are $\nu_j = 2j - 1$ for $j > 1$, but $\nu_1 = -1$.

6.7 POLAR DECOMPOSITION

6.7.1 Theorem. *A positive operator S on \mathcal{H} has a unique positive square root.*

PROOF: Write $S = \sum_{\lambda \in \sigma(T)} \lambda \pi_\lambda$ as above, and $\sqrt{S} = \sum \sqrt{\lambda} \pi_\lambda$, where we take the positive square roots of the (positive) λ 's. Then $(\sqrt{S})^2 = S$.

If T is positive and $T^2 = S$ then T and S commute so that T preserves all the eigenspaces \mathcal{H}_λ of S . On each \mathcal{H}_λ we have $S = \lambda I$, (the identity operator on \mathcal{H}_λ) so that $T = \sqrt{\lambda} J$, with positive square root, J positive, and $J^2 = I$. The eigenvalues of J are ± 1 , and the positivity of J implies that they are all 1, so $J = I$ and $T = \sqrt{S}$. ◀

A nonnegative operator S has square roots: write $\mathcal{H} = \mathcal{H}_{null} \oplus \mathcal{H}_{pos}$ where $\mathcal{H}_{null} = \ker(S)$ and $\mathcal{H}_{pos} = \bigoplus_{\lambda \in \sigma(S) \setminus \{0\}} \mathcal{H}_\lambda$. The restriction S_{pos} of S to \mathcal{H}_{pos} is positive and, by the theorem, has a unique square root $\sqrt{S_{pos}}$.

The restriction of S to \mathcal{H}_{null} is zero, and any operator T such that $T^2 = 0$ can serve as a square root of S on \mathcal{H}_{null} . This is the source of ambiguity in the definition of the square root. Setting \sqrt{S} to be the operator that keeps both \mathcal{H}_{null} and \mathcal{H}_{pos} invariant, is zero on \mathcal{H}_{null} , and is $\sqrt{S_{pos}}$ on \mathcal{H}_{pos} , is now a uniquely defined nonnegative operator whose square is S . We'll denote it, as for positive S , by \sqrt{S} or by $S^{\frac{1}{2}}$.

6.7.2 Lemma. *Let $\mathcal{H}_j \subset \mathcal{H}$, $j = 1, 2$, be isomorphic subspaces. Let U_1 be an isometry $\mathcal{H}_1 \mapsto \mathcal{H}_2$. Then there are unitary operators U on \mathcal{H} that extend U_1 .*

PROOF: Define U on \mathcal{H}_1^\perp as an arbitrary isometry onto \mathcal{H}_2^\perp (which has the same dimension) and extend by linearity. ◀

6.7.3 Lemma. *Let $A, B \in \mathcal{L}(\mathcal{H})$, and assume that $\|Av\| = \|Bv\|$ for all $v \in \mathcal{H}$. Then there exists a unitary operator U such that $B = UA$.*

PROOF: Clearly $\ker(A) = \ker(B)$. Let $\{u_1, \dots, u_n\}$ be an orthonormal basis of \mathcal{H} such that $\{u_1, \dots, u_m\}$ is a basis for $\ker(A) = \ker(B)$. The subspace $\text{range}(A)$ is spanned by $\{Au_j\}_{j=m+1}^n$ and $\text{range}(B)$ is spanned by $\{Bu_j\}_{j=m+1}^n$. The map $U_1: Au_j \mapsto Bu_j$ extends by linearity to an isometry of $\text{range}(A)$ onto $\text{range}(B)$. Now apply Lemma 6.7.2, and remember that, on the range of A , $U = U_1$. ◀

Remark: The operator U is unique if, and only if, A (or B) is invertible.

6.7.4 We observed, 6.4.1, that for any $T \in \mathcal{L}(\mathcal{H})$, the operators $S_1 = T^*T$ and $S_2 = TT^*$ are self adjoint. Notice that unless T is normal, $S_1 \neq S_2$.

For any $v \in \mathcal{H}$

$$\langle S_1 v, v \rangle = \langle Tv, Tv \rangle = \|Tv\|^2 \quad \text{and} \quad \langle S_2 v, v \rangle = \|T^*v\|^2,$$

so that both S_1 and S_2 are nonnegative, and both are positive if T is nonsingular.

Let $T \in \mathcal{L}(\mathcal{H})$. The operators $S_1 = T^*T$ and $S_2 = TT^*$ are nonnegative and hence have nonnegative square roots. Observe that

$$\begin{aligned} \|Tv\|^2 &= \langle Tv, Tv \rangle = \langle T^*Tv, v \rangle = \langle S_1 v, v \rangle = \\ &= \langle \sqrt{S_1}v, \sqrt{S_1}v \rangle = \|\sqrt{S_1}v\|^2. \end{aligned}$$

By Lemma 6.7.3, with $A = \sqrt{S_1}$ and $B = T$ there exist unitary operators U such that $T = U\sqrt{S_1}$. This proves

Theorem (Polar decomposition⁶). *Every operator $T \in \mathcal{L}(\mathcal{H})$ admits a representation*

$$(6.7.1) \quad T = UR,$$

where U is unitary and $R = \sqrt{T^*T}$ nonnegative.

Remark: Starting with T^* and taking adjoints at the end, one obtains also a representation of the form $T = R_1U_1$, with unitary U_1 and $R_1 = \sqrt{TT^*}$ nonnegative. Typically $R_1 \neq R$, as shown by following example. Let T be the map on \mathbb{C}^2 defined by $Tv_1 = v_2$, and $Tv_2 = 0$. Then R is the orthogonal projection onto the line of the scalar multiples of v_1 , R_1 is the orthogonal projection onto the multiples of v_2 , and $U = U_1$ maps each v_j on the other.

We shall use the notation $|T| = \sqrt{T^*T}$.

⁶Not to be confused with the *polarisation formula*,

6.7.5 With T , $|T|$, and U as above, let $\{\mu_1, \dots, \mu_n\}$ denote the eigenvalues of (the self-adjoint) $|T| = \sqrt{T^*T}$, let $\{u_1, \dots, u_n\}$ be the corresponding orthonormal eigenvectors, and denote $v_j = U^{-1}u_j$. Then $\{v_1, \dots, v_n\}$ is orthonormal, $|T|v = \sum \mu_j \langle v, u_j \rangle u_j$, and

$$(6.7.2) \quad Tv = \sum \mu_j \langle v, u_j \rangle v_j.$$

This is sometimes written⁷ as

$$(6.7.3) \quad T = \sum \mu_j u_j \otimes v_j.$$

EXERCISES FOR SECTION 6.7

VI.7.1. Let $\{w_1, \dots, w_n\}$ be an orthonormal basis for \mathcal{H} and let T be the (weighted) shift operator on $\{w_1, \dots, w_n\}$, defined by $Tw_j = (n-j)w_{j+1}$ for $j < n$, and $Tw_n = 0$. Describe U and R in (6.7.1), as well as R_1 and U_1 above.

VI.7.2. An operator T is *bounded below by c* , written $T \geq c$, on a subspace $\mathcal{V} \subset \mathcal{H}$ if $\|Tv\| \geq c\|v\|$ for every $v \in \mathcal{V}$. Assume that $\{u_1, \dots, u_n\}$ and $\{v_1, \dots, v_n\}$ are orthonormal sequences, $\mu_j > 0$, $\mu_{j+1} \leq \mu_j$, and $T = \sum \mu_j u_j \otimes v_j$. Show that

$$\mu_j = \max\{c: \text{there exists a } j\text{-dimensional subspace on which } T \geq c.\}$$

⁷See $\star 4.2.2$.

Chapter VII

Additional topics

Unless stated explicitly otherwise, the underlying field of the vector spaces discussed in this chapter is either \mathbb{R} or \mathbb{C} .

7.1 QUADRATIC FORMS

7.1.1 A quadratic form on an n -dimensional inner-product space \mathcal{H} is a function of the form $Q(v) = \langle Tv, v \rangle$ with $T \in \mathcal{L}(\mathcal{H})$.

A basis $\mathbf{v} = \{v_1, \dots, v_n\}$ transforms Q into a function $Q_{\mathbf{v}}$ of n variables on the underlying field, \mathbb{R} or \mathbb{C} as the case may be. We use the notation appropriate¹ for \mathbb{C} .

Write $v = \sum_1^n x_j v_j$ and $a_{i,j} = \langle Tv_i, v_j \rangle$; then $\langle Tv, v \rangle = \sum_{i,j} a_{i,j} x_i \bar{x}_j$ and

$$(7.1.1) \quad Q_{\mathbf{v}}(x_1, \dots, x_n) = \sum_{i,j} a_{i,j} x_i \bar{x}_j$$

expresses Q in terms of the variables $\{x_j\}$, (i.e., the \mathbf{v} -coordinates of v).

We denote the matrix of coefficients $(a_{i,j})$ by $A_{\mathbf{v}}$, write the coordinates as a column vector, $\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$, and observe that

$$(7.1.2) \quad Q_{\mathbf{v}}(x_1, \dots, x_n) = \langle A\mathbf{x}, \mathbf{x} \rangle = \bar{\mathbf{x}}^T A_{\mathbf{v}} \mathbf{x}$$

transfers the action to \mathbb{F}^n .

¹If the underlying field is \mathbb{R} the complex conjugation can be simply ignored

7.1.2 When the underlying field is \mathbb{R} the quadratic form Q is real-valued. It does not determine the entries $a_{i,j}$ uniquely. Since $x_j x_i = x_i x_j$, the value of Q depends on $a_{i,j} + a_{j,i}$ and not on each of the summands separately. We may therefore assume, without modifying Q , that $a_{i,j} = a_{j,i}$, thereby making the matrix $A_{\mathbf{v}} = (a_{i,j})$ symmetric.

For real-valued quadratic forms over \mathbb{C} the following lemma guarantees that the matrix of coefficients is Hermitian.

Lemma. *A quadratic form $\bar{\mathbf{x}}^{\text{Tr}} A_{\mathbf{v}} \mathbf{x}$ on \mathbb{C}^n is real-valued if, and only if, the matrix of coefficients $A_{\mathbf{v}}$ is Hermitian² i.e., $a_{i,j} = \overline{a_{j,i}}$.*

PROOF: If $a_{i,j} = \overline{a_{j,i}}$ for all i, j , then $\sum_{i,j} a_{i,j} x_i \bar{x}_j$ is its own complex conjugate.

Conversely, assume that $\sum_{i,j} a_{i,j} x_i \bar{x}_j$ is real-valued for all $x_1, \dots, x_n \in \mathbb{C}$. Taking $x_j = 0$ for $j \neq k$, and $x_k = 1$, we obtain $a_{k,k} \in \mathbb{R}$. Taking $x_k = x_l = 1$ and $x_j = 0$ for $j \neq k, l$, we obtain $a_{k,l} + a_{l,k} \in \mathbb{R}$, i.e. $\Im a_{k,l} = -\Im a_{l,k}$. For $x_k = i$, $x_l = 1$ we obtain $i(a_{k,l} - a_{l,k}) \in \mathbb{R}$, i.e., $\Re a_{k,l} = \Re a_{l,k}$ and combining the two we have $a_{k,l} = \overline{a_{l,k}}$. \blacktriangleleft

7.1.3 If we replace the basis \mathbf{v} by another, say \mathbf{w} , the coefficients undergo a linear change of variables. There exists a matrix $C \in \mathcal{M}(n)$, that trans-

forms by left multiplication the \mathbf{w} -coordinates $\mathbf{y} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}$ of a vector into its

\mathbf{v} -coordinates: $\mathbf{x} = C\mathbf{y}$. Now

$$(7.1.3) \quad Q_{\mathbf{v}}(x_1, \dots, x_n) = \bar{\mathbf{x}}^{\text{Tr}} A_{\mathbf{v}} \mathbf{x} = \bar{\mathbf{y}}^{\text{Tr}} \overline{C}^{\text{Tr}} A_{\mathbf{v}} C \mathbf{y}$$

and the matrix representing Q in terms of the variables y_j , is³

$$(7.1.4) \quad A_{\mathbf{w}} = \overline{C}^{\text{Tr}} A_{\mathbf{v}} C = C^* A_{\mathbf{v}} C.$$

Notice that the form now is $C^ A C$, rather than $C^{-1} A C$ (which defines similarity). The two notions agree if C is unitary, since then $C^* = C^{-1}$, and the matrix of coefficients for the variables $\{y_j\}$ is $C^{-1} A C$.*

²Equivalently, if the operator T is self-adjoint.

³The adjoint of a matrix is introduced in 6.2.3.

7.1.4 The fact that the matrix of coefficients of a real-valued quadratic form Q is self-adjoint makes it possible to simplify Q by a (unitary) change of variables that reduces it to a linear combination of squares. If the given matrix is A , we invoke the spectral theorem, 6.4.5, to obtain a unitary matrix U , such that $U^*AU = U^{-1}AU$ is a diagonal matrix whose diagonal consists of the complete collection, including multiplicity, of the eigenvalues $\{\lambda_j\}$ of A . In other words, if $\mathbf{x} = U\mathbf{y}$, then

$$(7.1.5) \quad Q(x_1, \dots, x_n) = \sum \lambda_j |y_j|^2.$$

There are other matrices C which diagonalize Q , and the coefficients in the diagonal representation $Q(y_1, \dots, y_n) = \sum b_j |y_j|^2$ depend on the one used. What does not depend on the particular choice of C is the number n_+ of positive coefficients, the number n_0 of zeros and the number n_- of negative coefficients. This is known as *The law of inertia*.

DEFINITION: A quadratic form $Q(v)$ on a (real or complex) vector space \mathcal{V} is *positive*, resp. *negative* if $Q(v) > 0$, resp. $Q(v) < 0$, for all $v \neq 0$ in \mathcal{V} .

On an inner-product space $Q(v) = \langle Av, v \rangle$ with a self-adjoint operator A , and our current definition is consistent with the definition in 6.6.1: the operator A is positive if so is $Q(v) = \langle Av, v \rangle$.

Denote

$$(7.1.6) \quad \begin{aligned} n_+ &= \max_{\mathcal{V}_1} \dim \mathcal{V}_1 & : & \quad Q \text{ is positive on } \mathcal{V}_1. \\ n_- &= \max_{\mathcal{V}_1} \dim \mathcal{V}_1 & : & \quad Q \text{ is negative on } \mathcal{V}_1. \end{aligned}$$

and, $n_0 = n - n_+ - n_-$.

Proposition. *Let \mathbf{v} be a basis in terms of which $Q(y_1, \dots, y_n) = \sum b_j |y_j|^2$, and arrange the coordinates so that $b_j > 0$ for $j \leq m$ and $b_j \leq 0$ for $j > m$. Then $m = n_+$.*

PROOF: Denote $\mathcal{V}_+ = \text{span}[v_1, \dots, v_m]$, and $\mathcal{V}_{\leq 0} = \text{span}[v_{m+1}, \dots, v_n]$ the complementary subspace.

$Q(y_1, \dots, y_n)$ is clearly positive on \mathcal{V}_+ , so that $m \leq n_+$. On the other hand, by Theorem 2.5.3, every subspace \mathcal{W} of dimension $> m$ has elements $v \in \mathcal{V}_{\leq 0}$, and for such v we clearly have $Q(v) \leq 0$. ◀

The proposition applied to $-Q$ shows that n_- equals the number of negative b_j 's. This proves

Theorem (Law of inertia). *Let Q be a real-valued quadratic form. Then in any representation $Q(y_1, \dots, y_n) = \sum b_j |y_j|^2$, the number of positive coefficients is n_+ , the number of negative coefficients is n_- , and the number of zeros is n_0 .*

EXERCISES FOR SECTION 7.1

VII.1.1. Prove that if $\langle Av, v \rangle = \langle Bv, v \rangle$ for all $v \in \mathbb{R}^n$, with $A, B \in \mathcal{M}(n, \mathbb{R})$, and both symmetric, then $A = B$.

7.2 POSITIVE MATRICES

A matrix $A \in \mathcal{M}(m, \mathbb{C})$ is *positive*⁴ if all the entries are positive. A is *nonnegative* if all the entries are nonnegative.

Similarly, a vector $v \in \mathbb{C}^m$ is *positive*, resp. *nonnegative*, if all its entries are positive, resp. non-negative.

With A_j denoting either matrices or vectors, $A_1 \geq A_2$, $A_1 \gneq A_2$, and $A_1 > A_2$ will mean respectively that $A_1 - A_2$ is nonnegative, nonnegative but not zero, positive.

7.2.1 We write $\|A\|_{sp} = \max\{|\tau| : \tau \in \sigma(A)\}$, and refer to it as *the spectral norm* of A .

DEFINITION: An eigenvalue λ of a matrix A is called *dominant* if

- a.** λ is simple (that is $\ker((A - \lambda)^2) = \ker(A - \lambda)$ is one dimensional), and
- b.** every other eigenvalue μ of A satisfies $|\mu| < |\lambda|$.

Notice that **b.** implies that $|\lambda| = \|A\|_{sp}$.

Theorem (Perron). *Let $A = (a_{i,j})$ be a positive matrix. Then it has a positive dominant eigenvalue and a positive corresponding eigenvector. Moreover, there is no other nonnegative eigenvector for A .*

⁴Not to be confused with positivity of *the operator* of multiplication by A .

PROOF: Let $p(A)$ be the set of all positive numbers μ such that there exist nonnegative vectors $v \neq 0$ such that

$$(7.2.1) \quad Av \geq \mu v.$$

Clearly $\min_i a_{i,i} \in p(A)$; also $\mu \leq m \max_{i,j} a_{i,j}$ for all $\mu \in p(A)$. Hence $p(A)$ is non-empty and bounded.

Write $\lambda = \sup_{\mu \in p(A)} \mu$. We propose to show that $\lambda \in p(A)$, and is the dominant eigenvalue for A .

Let $\mu_n \in p(A)$ be such that $\mu_n \rightarrow \lambda$, and $v_n = (v_n(1), \dots, v_n(m)) \succeq 0$ such that $Av_n \geq \mu_n v_n$. We normalize v_n by the condition $\sum_j v_n(j) = 1$, and since now $0 \leq v_n(j) \leq 1$ for all n and j , we can choose a (sub)sequence n_k such that $v_{n_k}(j)$ converges for each $1 \leq j \leq m$. Denote the limit by $v_*(j)$ and let $v_* = (v_*(1), \dots, v_*(m))$. Since all the entries of Av_{n_k} converge to the corresponding entries in Av_* we have $\sum_j v_*(j) = 1$, and

$$(7.2.2) \quad Av_* \geq \lambda v_*.$$

Claim: the inequality (7.2.2) is in fact an equality, so that λ is an eigenvalue and v_* a corresponding eigenvector.

If one of the entries in λv_* , say $\lambda v_*(l)$, were smaller than the l 'th entry in Av_* , we could replace v_* by $v_{**} = v_* + \varepsilon e_l$ (where e_l is the unit vector that has 1 as its l 'th entry and zero everywhere else) with $\varepsilon > 0$ small enough to have

$$Av_*(l) \geq \lambda v_{**}(l).$$

Since Ae_l is (strictly) positive, we would have $Av_{**} > Av_* \geq \lambda v_{**}$, and for $\delta > 0$ sufficiently small we would have

$$Av_{**} \geq (\lambda + \delta)v_{**}$$

contradicting the definition of λ .

Since Av is positive for any $v \succeq 0$, a nonnegative vector which is an eigenvector of A with positive eigenvalue, is positive. In particular, $v_* > 0$.

Claim: λ is a simple eigenvalue.

a. If $Au = \lambda u$ for some vector u , then $A\Re u = \lambda\Re u$ and $A\Im u = \lambda\Im u$. So it would be enough to show that u is a constant multiple of v_* under the

assumption that u has real entries. There exists a constant $c \neq 0$ such that $v_* + cu$ has nonnegative entries and at least one vanishing entry. Since $v_* + cu$ is an eigenvector for λ , the previous remark shows that $v_* + cu = 0$ and u is a multiple of v_* .

b. We need to show that $\ker((A - \lambda)^2) = \ker((A - \lambda))$. Assume the contrary, and let $u \in \ker((A - \lambda)^2) \setminus \ker((A - \lambda))$, so that

$$(7.2.3) \quad Au = \lambda u + cv_*$$

with $c \neq 0$. Splitting (7.2.3) into its real and imaginary parts we have:

$$(7.2.4) \quad A\Re u = \lambda\Re u + \Re cv_* \quad A\Im u = \lambda\Im u + \Im cv_*$$

Either $c_1 = \Re c \neq 0$ or $c_2 = \Im c \neq 0$ (or both). This shows that there is no loss of generality in assuming that u and c in (7.2.3) are real valued.

Replace u , if necessary, by $u_1 = -u$ to obtain $Au_1 = \lambda u_1 + c_1 v_*$ with $c_1 > 0$. Let $a > 0$ be such that $u_1 + av_* > 0$, and observe that

$$A(u_1 + av_*) = \lambda(u_1 + av_*) + c_1 v_*$$

so that $A(u_1 + av_*) > \lambda(u_1 + av_*)$ contradicting the maximality of λ .

c. Claim: Every eigenvalue $\mu \neq \lambda$ of A satisfies $|\mu| < \lambda$.

Let μ be an eigenvalue of A , and let $w \neq 0$ be a corresponding eigenvector: $Aw = \mu w$. Denote $|w| = (|w(1)|, \dots, |w(m)|)$.

The positivity of A implies $A|w| \geq |Aw|$ and,

$$(7.2.5) \quad A|w| \geq |Aw| \geq |\mu||w|$$

so that $|\mu| \in p(A)$, i.e., $|\mu| \leq \lambda$. If $|\mu| = \lambda$ we must have equality in (7.2.5) and $|w| = cv_*$. Equality in (7.2.5) can only happen if $A|w| = |Aw|$ which means that all the entries in w have the same argument, i.e. $w = e^{i\vartheta}|w|$, in other words, w is a constant multiple of v_* , and $\mu = \lambda$.

Finally, let $\mu \neq \lambda$ be an eigenvalue of A and w a corresponding eigenvector. The adjoint $A^* = \overline{A}^{Tr}$ is a positive matrix and has the same dominant eigenvalue λ . If v^* is the positive eigenvector corresponding to λ then $\langle w, v^* \rangle = 0$ (see exercise VI.2.4) and since v^* is strictly positive, w must have both positive and negative entries. ◀

EXERCISES FOR SECTION 7.2

VII.2.1. What part of the conclusion of Perron's theorem remains valid if the assumption is replaced by “ A is similar to a positive matrix”?

7.3 NONNEGATIVE MATRICES

Nonnegative matrices exhibit a variety of modes of behavior. Consider the following $n \times n$ matrices

- a.** The identity matrix. 1 is the only eigenvalue, multiplicity n .
- b.** The nilpotent matrix having ones below the diagonal, zeros elsewhere. The spectrum is $\{0\}$.
- c.** The matrix A_σ of a permutation $\sigma \in S_n$. The spectrum depends on the decomposition of σ into cycles. If σ is a unique cycle then the spectrum of A_σ is the set of roots of unity of order n . The eigenvalue 1 has $(1, \dots, 1)$ as a unique eigenvector. If the decomposition of σ consists of k cycles of lengths l_j , $j = 1, \dots, k$, then the spectrum of A_σ is the union of the sets of roots of unity of order l_j . The eigenvalue 1 now has multiplicity k .

7.3.1 Let $\mathbb{1}$ denote the matrix all of whose entries are 1. If $A \geq 0$ then $A + \frac{1}{m}\mathbb{1} > 0$ and has, by Perron's theorem, a dominant eigenvalue λ_m and a corresponding positive eigenvector v_m which we normalize by the condition $\sum_{j=1}^n v_m(j) = 1$.

λ_m is monotone non increasing as $m \rightarrow \infty$ and converges to a limit $\lambda \geq 0$ which clearly dominates the spectrum of A . λ can well be zero, as can be seen from example **b.** above. For a sequence $\{m_i\}$ the vectors v_{m_i} converge to a (normalized) nonnegative vector v_* which, by continuity, is an eigenvector for λ .

Thus, a nonnegative matrix has $\lambda = \|A\|_{sp}$ as an eigenvalue with nonnegative eigenvector v_* , however

1. λ may be zero,
2. λ may have high multiplicity,

3. λ may not have *positive* eigenvectors.
4. There may be other eigenvalues of modulus $\|A\|_{sp}$.

The first three problems disappear, and the last explained for *transitive nonnegative matrices*. See below.

7.3.2 DEFINITIONS. Assume $A \geq 0$. We use the following terminology:

A *connects* the index j to i (connects (j, i) for short) *directly* if $a_{i,j} \neq 0$. Since $Ae_j = \sum a_{i,j}e_i$, A connects (j, i) if e_i appears (with nonzero coefficient) in the expansion of Ae_j .

A *connects* j to i (connects (j, i) for short) if there is a *connecting chain* for (j, i) , that is, a sequence $\{s_l\}_{l=0}^k$ such that $j = s_0$, $i = s_k$ and $\prod_{l=1}^k a_{s_l, s_{l-1}} \neq 0$. The existence of connecting chain for (j, i) is equivalent to: e_i appears (with nonzero coefficient) in the expansion of $A^k e_j$.

An index j is *A-recurrent* if A connects it to itself—there is a connecting chain for (j, j) . The lengths k of connecting chains for (j, j) are called *return times* for j . Since connecting chains for (j, j) can be concatenated, the set of return times for a recurrent index is an additive semigroup of \mathbb{N} .

A is *transitive*⁵ if it connects every pair (j, i) .

Lemma. *If A is a nonnegative transitive matrix, every index is A-recurrent. In particular, A is not nilpotent.*

PROOF: Left as an exercise. ◀

Corollary. *If A is a nonnegative transitive matrix then $\lambda = \|A\|_{sp} > 0$.*

7.3.3 We write $i \leq_A j$ if A connects (i, j) . This defines a partial order and induces an equivalence relation in the set of A -recurrent indices. (The non-recurrent indices are not equivalent to themselves, nor to anybody else.)

We can reorder the indices in a way that gives each equivalence class a consecutive bloc, and is compatible with the partial order, i.e., such that for non-equivalent indices, $i \leq_A j$ implies $i \leq j$. This ordering is not unique: equivalent indices can be ordered arbitrarily within their equivalence class; pairs of

⁵Also called *irreducible*, or *ergodic*.

equivalence classes may be \leq_A comparable or not comparable in which case each may precede the other; non-recurrent indices may be placed consistently in more than one place. Yet, such order gives the matrix A a “quasi-super-triangular form”: if we denote the coefficients of the “reorganized” A again by $a_{i,j}$, then $a_{i,j} = 0$ for i greater than the end of the bloc containing j . That means that now A has square transitive matrices centered on the diagonal—the squares $J_l \times J_l$ corresponding to the equivalence classes, while the entries on the rest of diagonal, at the non-recurrent indices, as well as in the rest of the sub-diagonal, are all zeros. This reduces much of the study of the general A to that of transitive matrices.

7.3.4 We focus now on transitive matrices.

A nonnegative matrix A is transitive if, and only if, $B = \sum_{j=1}^n A^j$ is positive. Since, by 7.3.1, $\lambda = \|A\|_{sp}$ is an eigenvalue for A , it follows that $\beta = \sum_1^n \lambda^j$ is an eigenvalue for B , having the same eigenvalue v_* .

Either by observing that $\beta = \|B\|_{sp}$, or by invoking the part in Perron’s theorem stating that (up to constant multiples) there is only one nonnegative eigenvector for B (and it is in fact positive), we see that β is the dominant eigenvalue for B and v_* is positive.

Lemma. *Assume A transitive, $v \geq 0$, $\mu > 0$, $Av \not\geq \mu v$. Then there exists a positive vector $u \geq v$ such that $Au > \mu u$.*

PROOF: As in the proof of Perron’s theorem: let l be such that $Av(l) > \mu v_l$, let $0 < \varepsilon_1 < Av(l) - \mu v_l$ and $v_1 = v + \varepsilon_1 e_l$. Then $Av \geq \mu v_1$, hence

$$Av_1 = Av + \varepsilon_1 A e_l \geq \mu v_1 + \varepsilon_1 A e_l,$$

and Av_1 is strictly bigger than μv_1 at l and at all the entries on which $A e_l$ is positive, that is the i ’s such that $a_{i,l} > 0$. Now define $v_2 = v_1 + \varepsilon_2 A e_l$ with $\varepsilon_2 > 0$ sufficiently small so that $Av_2 \geq \mu v_2$ with strict inequality for l and the indices on which $A e_l + A^2 e_l$ is positive. Continue in the same manner with v_3 , and $Av_3 \geq \mu v_3$ with strict inequality on the support of $(I + A + A^2 + A^3) e_l$ etc. The transitivity of A guarantees that after $k \leq n$ such steps we obtain $u = v_k > 0$ such that $Au > \mu u$. ◀

The lemma implies in particular that if, for some $\mu > 0$, there exists a vector $v \geq 0$ such that $Av \geq \mu v$, then $\mu < \|A\|_{sp}$. This since the condition $Au > \mu u$ implies⁶ $(A + \frac{1}{m} \mathbb{I})u > (1+a)\mu u$ for $a > 0$ sufficiently small, and all m . In turn this implies $\lambda_m > (1+a)\mu$ for all m , and hence $\lambda \geq (1+a)\mu$.

In what follows we simplify the notation somewhat by normalizing (multiplying by a positive constant) the nonnegative transitive matrices under discussion so that $\|A\|_{sp} = 1$.

Corollary. *Assume $\|A\|_{sp} = 1$. If $\mu = e^{i\varphi}$ is an eigenvalue of A and u_μ a corresponding⁷ eigenvector, then $|u_\mu| = v_*$.*

PROOF: $A|u_\mu| \geq |Au_\mu| = |\mu u_\mu| = |u_\mu|$.

If $A|u_\mu| \neq |u_\mu|$ the lemma would contradict the assumption $\|A\|_{sp} = 1$. \blacktriangleleft

7.3.5 For $v \in \mathbb{C}^n$, $|v| > 0$ we write $\arg v = (\arg v_1, \dots, \arg v_n)$, and⁸ $e^{i \arg v} = (e^{i \arg v_1}, \dots, e^{i \arg v_n})$.

The key observation is: if $Au_\mu = \mu u_\mu$, then $A|u_\mu| = |Au_\mu|$ which means that every entry in Au_μ is a linear combination of entries of u_μ *having the same argument*, that is on which $\arg u_\mu$ is constant. The set $[1, \dots, n]$ is partitioned into the level sets I_j on which $\arg u_\mu = \vartheta_j$, and A maps \mathbf{e}_l for every $l \in I_j$, and hence $\text{span}[\{\mathbf{e}_l\}_{k \in I_j}]$, into $\text{span}[\{\mathbf{e}_k\}_{k \in I_s}]$ where $\vartheta_s = \vartheta_j + \varphi$.

Let $\nu = e^{i\psi}$ be another eigenvalue of A , with eigenvector $u_\nu = e^{i \arg u_\nu} v_*$, and let J_k be the level sets on which $\arg u_\nu = \gamma_k$. A maps \mathbf{e}_l for every $l \in J_k$, into $\text{span}[\{\mathbf{e}_m\}_{m \in J_s}]$ where $\gamma_s = \gamma_k + \psi$.

It follows that for $l \in I_j \cap J_k$, $A\mathbf{e}_l \in \text{span}[\{\mathbf{e}_k\}_{k \in I_s}] \cap \text{span}[\{\mathbf{e}_m\}_{m \in J_t}]$ where $\vartheta_s = \vartheta_j + \varphi$ and $\gamma_t = \gamma_k + \psi$. If we write $u_{\mu\nu} = e^{i(\arg u_\mu + \arg u_\nu)} v_*$, then $\arg Ae^{i(\vartheta_j + \gamma_k)} \mathbf{e}_l = \arg u_\mu + \arg u_\nu + \varphi + \psi$, which means: $Au_{\mu\nu} = \mu\nu u_{\mu\nu}$.

This proves that the product $\mu\nu = e^{i(\varphi + \psi)}$ of eigenvalues of A is an eigenvalue, and $\sigma(A) \cap \{z : |z| = 1\}$ is a subgroup of the multiplicative unit circle; i.e., the group of roots of unity of order m for an appropriate m .

⁶See 7.3.1 for the notation.

⁷Normalized: $\sum_j |u_\mu(j)| = 1$.

⁸The notation considers \mathbb{C}^n as an algebra of functions on the space $[1, \dots, n]$.

The group $\sigma(A) \cap \{z : |z| = 1\}$, (or $\{e^{it} : \|A\|_{sp} e^{it} \in \sigma(A)\}$ if A is not normalized), is called the *period group* of A and its order m is the *periodicity* of A .

We call the partition of $[1, \dots, n]$ into the level sets I_j of $\arg v_\mu$, where μ is a generator of the period group of A , the *basic partition*.

The subspaces $\mathcal{V}_j = \text{span}\{\mathbf{e}_l : l \in I_j\}$ are A^m -invariant and the restriction of A^m to \mathcal{V}_j is transitive with the dominant eigenvalue 1, and $v_{*,j} = \sum_{l \in I_j} v_*(l) \mathbf{e}_l$ the corresponding eigenvector.

The restriction of A^m to \mathcal{V}_j has $|I_j| - 1$ eigenvalues of modulus < 1 . Summing for $1 \leq j \leq m$ and invoking the Spectral Mapping Theorem, 5.1.2, we see that A has $n - m$ eigenvalues of modulus < 1 . This proves that the eigenvalues in the period group are simple and have no generalized eigenvectors.

Theorem (Frobenius). *Let A be a transitive nonnegative $n \times n$ matrix. Then $\lambda = \|A\|_{sp}$ is a simple eigenvalue of A and has a positive eigenvector v_* . The set $\{e^{it} : \lambda e^{it} \in \sigma(A)\}$ is a subgroup of the unit circle.*

7.3.6 DEFINITION: A matrix $A \geq 0$ is *strongly transitive* if A^m is transitive for all $m \in [1, \dots, n]$.

Theorem. *If A is strongly transitive, then $\|A\|_{sp}$ is a dominant eigenvalue for A , and has a positive corresponding eigenvector.*

PROOF: The periodicity of A has to be 1. ◀

EXERCISES FOR SECTION 7.3

VII.3.1. A nonnegative matrix A is nilpotent if, and only if, no index is A -recurrent.

VII.3.2. Prove that a nonnegative matrix A is transitive if, and only if, $B = \sum_{l=1}^n A^l$ is positive.

Hint: Check that A connects (i, j) if, and only if, $\sum_{l=1}^n A^l$ connects j to i directly.

VII.3.3. Prove that the conclusion Perron's theorem holds under the weaker assumption: "the matrix A is nonnegative and has a full row of positive entries".

VII.3.4. Prove that if the elements I_j of the basic partition are not equal in size, then $\ker(A)$ is nontrivial.

Hint: Show that $\dim \ker(A) \geq \max |I_j| - \min |I_j|$.

VII.3.5. Describe the matrix of a transitive A if the basis elements are reordered so that the elements of the basic partition are blocs of consecutive integers in $[1, \dots, n]$,

VII.3.6. Prove that if $A \geq 0$ is transitive, then so is A^* .

VII.3.7. Prove that if $A \geq 0$ is transitive, $\lambda = \|A\|_{sp}$, and v^* is the positive eigenvector of A^* , normalized by the condition $\langle v_*, v^* \rangle = 1$ then for all $v \in \mathbb{C}^n$,

$$(7.3.1) \quad \lim_{N \rightarrow \infty} \frac{1}{N} \sum_1^N \lambda^{-j} A^j v = \langle v, v^* \rangle v_*.$$

VII.3.8. Let σ be a permutation of $[1, \dots, n]$. Let A_σ be the $n \times n$ matrix whose entries a_{ij} are defined by

$$(7.3.2) \quad a_{ij} = \begin{cases} 1 & \text{if } i = \sigma(j) \\ 0 & \text{otherwise.} \end{cases}$$

What is the spectrum of A_σ , and what are the corresponding eigenvectors.

VII.3.9. Let $1 < k < n$, and let $\sigma \in S_n$, be the permutation consisting of the two cycles $(1, \dots, k)$ and $(k+1, \dots, n)$, and A_σ as defined above. (So that the corresponding operator on \mathbb{C}^n maps the basis vector e_i onto $e_{\sigma(i)}$.)

a. Describe the positive eigenvectors of A . What are the corresponding eigenvalues?

b. Let $0 < a, b < 1$. Denote by $A_{a,b}$ the matrix obtained from A by replacing the k 'th and the n 'th columns of A by $(c_{i,k})$ and $(c_{i,n})$, resp., where $c_{1,k} = 1 - a$, $c_{k+1,k} = a$ and all other entries zero; $c_{1,n} = b$, $c_{k+1,n} = 1 - b$ and all other entries zero.

Show that 1 is a simple eigenvalue of $A_{a,b}$ and find a positive corresponding eigenvector. Show also that for other eigenvalues there are no nonnegative eigenvectors.

7.4 STOCHASTIC MATRICES.

7.4.1 A *stochastic matrix* is a nonnegative matrix $A = (a_{i,j})$ such that the sum of the entries in each column⁹ is 1:

$$(7.4.1) \quad \sum_i a_{i,j} = 1.$$

⁹The action of the matrix is (left) multiplication of column vectors. The columns of the matrix are the images of the standard basis in \mathbb{R}^n or \mathbb{C}^n

A *probability vector* is a nonnegative vector $\pi = (p_i) \in \mathbb{R}^n$ such that $\sum_i p_i = 1$. Observe that if A is a stochastic matrix and π a probability vector, then $A\pi$ is a probability vector.

In applications, one considers a set of possible outcomes of an “experiment” at a given time. The outcomes are often referred to as *states*, and a probability vector assigns probabilities to the various states. The word probability is taken here in a broad sense—if one is studying the distribution of various populations, the “probability” of a given population is simply its proportion in the total population.

A (stationary) *n-state Markov chain* is a sequence $\{v_j\}_{j \geq 0}$ of probability vectors in \mathbb{R}^n , such that

$$(7.4.2) \quad v_j = Av_{j-1} = A^j v_0,$$

where A is an $n \times n$ stochastic matrix.

The matrix A is *the transition matrix*, and the vector v_0 is referred to as the *initial probability vector*. The parameter j is often referred to as *time*.

7.4.2 POSITIVE TRANSITION MATRIX. When the transition matrix A is positive, we get a clear description of the evolution of the Markov chain from Perron’s theorem 7.2.1.

Condition (7.4.1) is equivalent to $u^*A = u^*$, where u^* is the row vector $(1, \dots, 1)$. This means that the dominant eigenvalue for A^* is 1, hence the dominant eigenvalue for A is 1. If v_* is the corresponding (positive) eigenvector, normalized so as to be a probability vector, then $Av_* = v_*$ and hence $A^j v_* = v_*$ for all j .

If w is another eigenvector (or generalized eigenvector), it is orthogonal to u^* , that is: $\sum_1^n w(j) = 0$. Also, $\sum |A^l w(j)|$ is exponentially small (as a function of l).

If v_0 is any probability vector, we write $v_0 = cv_* + w$ with w in the span of the eigenspaces of the non dominant eigenvalues. By the remark above $c = \sum v_0(j) = 1$. Then $A^l v_0 = v_* + A^l w$ and, since $A^l w \rightarrow 0$ as $l \rightarrow \infty$, we have $A^l v_0 \rightarrow v_*$.

Finding the vector v_* amounts to solving a homogeneous system of n equations (knowing a-priori that the solution set is one dimensional). The observa-

tion $v_* = \lim A^l v_0$, with v_0 an arbitrary probability vector, may be a fast way to obtain a good approximation of v_* .

7.4.3 TRANSITIVE TRANSITION MATRIX. Denote v_μ the eigenvectors of A corresponding to eigenvalues μ of absolute value 1, normalized so that $v_1 = v_*$ is a probability vector, and $|v_\mu| = v_*$. If the periodicity of A is m , then, for every probability vector v_0 , the sequence $A^j v_0$ is equal to an m -periodic sequence (periodic sequence of of period m) plus a sequence that tends to zero exponentially fast.

Observe that for an eigenvalue $\mu \neq 1$ of absolute value 1, $\sum_1^m \mu^l = 0$. It follows that if v_0 is a probability vector, then

$$(7.4.3) \quad \frac{1}{m} \sum_{l=k+1}^{k+m} A^l v_0 \rightarrow v_*$$

exponential fast (as a function of k).

7.4.4 REVERSIBLE MARKOV CHAINS. One way of obtaining a transition matrix is from a nonnegative symmetric matrix $(p_{i,j})$ by writing $W_j = \sum_i p_{i,j}$ and, assuming $W_j > 0$ for all i , $a_{i,j} = \frac{p_{i,j}}{W_i}$. Then $A = (a_{i,j})$ is stochastic since $\sum_i a_{i,j} = 1$ for all j .

We can identify the “stable distribution”—the A -invariant vector—by thinking in terms of “population movement”. Assume that at a given time we have population of size b_j in state j and in the next unit of time $a_{i,j}$ proportion of this population shifts to state i . The absolute size of the j to i shift is $a_{i,j} b_j$ so that the new distribution is given by $A\mathbf{b}$, where \mathbf{b} is the column vector with entries b_j . This description applies to any stochastic matrix, and the stable distribution is given by \mathbf{b} which is invariant under A , $A\mathbf{b} = \mathbf{b}$.

The easiest way to find \mathbf{b} in this case is to go back to the matrix $(p_{i,j})$ and the weights W_j . The vector \mathbf{w} with entries W_j is A -invariant in a very strong sense. Not only is $A\mathbf{w} = \mathbf{w}$, but the exchange of mass between any two states is even:

- the mass going from i to j is: $W_i a_{j,i} = p_{j,i}$,
- the mass going from j to i is: $W_j a_{i,j} = p_{i,j}$,
- the two are equal since $p_{i,j} = p_{j,i}$.

EXERCISES FOR SECTION 7.4

VII.4.1. Let σ be a permutation of $[1, \dots, n]$. Let A_σ be the $n \times n$ matrix whose entries a_{ij} are defined by

$$(7.4.4) \quad a_{ij} = \begin{cases} 1 & \text{if } i = \sigma(j) \\ 0 & \text{otherwise.} \end{cases}$$

What is the spectrum of A_σ , and what are the corresponding eigenvectors.

VII.4.2. Let $1 < k < n$, and let $\sigma \in S_n$, be the permutation consisting of the two cycles $(1, \dots, k)$ and $(k+1, \dots, n)$, and A_σ as defined above. (So that the corresponding operator on \mathbb{C}^n maps the basis vector e_i onto $e_{\sigma(i)}$.)

a. Describe the positive eigenvectors of A . What are the corresponding eigenvalues?

b. Let $0 < a, b < 1$. Denote by $A_{a,b}$ the matrix obtained from A by replacing the k 'th and the n 'th columns of A by $(c_{i,k})$ and $(c_{i,n})$, resp., where $c_{1,k} = 1 - a$, $c_{k+1,k} = a$ and all other entries zero; $c_{1,n} = b$, $c_{k+1,n} = 1 - b$ and all other entries zero.

Show that 1 is a simple eigenvalue of $A_{a,b}$ and find a positive corresponding eigenvector. Show also that for other eigenvalues there are no nonnegative eigenvectors.

7.5 REPRESENTATION OF FINITE GROUPS

A representation of a group G in a vector space \mathcal{V} is a homomorphism $\sigma : g \mapsto \mathbf{g}$ of G into the group $\mathbf{GL}(\mathcal{V})$ of invertible elements in $\mathcal{L}\mathcal{V}$.

Throughout this section G will denote a finite group.

A representation of G in \mathcal{V} turns \mathcal{V} into a G -module, or a G -space. That means that in addition to the vector space operations there is an action of G on \mathcal{V} by linear maps: for every $g \in G$ and $v \in \mathcal{V}$ the element $\mathbf{g}v \in \mathcal{V}$ is well defined and,

$$\mathbf{g}(av_1 + bv_2) = \mathbf{a}gv_1 + \mathbf{b}gv_2 \quad \text{while} \quad (\mathbf{g}_1\mathbf{g}_2)v = \mathbf{g}_1(\mathbf{g}_2v).$$

The data (σ, \mathcal{V}) , i.e., \mathcal{V} as a G -space, is called a *representation of G in \mathcal{V}* . The representation is *faithful* if σ is injective.

Typically, σ is assumed known and is omitted from the notation. We shall use the terms G -space, G -module, and representation as synonyms.

We shall deal mainly in the case in which the underlying field is \mathbb{C} , or \mathbb{R} , and the space has an inner-product structure. The inner-product is assumed for convenience only: it identifies the space with its dual, and makes $\mathcal{L}\mathcal{V}$ self-adjoint. An inner product can always be introduced (e.g., by declaring a given basis to be orthonormal).

7.5.1 THE DUAL REPRESENTATION. If σ is a representation of G in \mathcal{V} we obtain a representation σ^* of G in \mathcal{V}^* by setting $\sigma^*(g) = (\sigma(g^{-1}))^*$ (the adjoint of the inverse of the action of G on \mathcal{V}). Since both $\mathfrak{g} \mapsto \mathfrak{g}^{-1}$ and $\mathfrak{g} \mapsto \mathfrak{g}^*$ reverse the order of factors in a product, their combination as used above preserves the order, and we have

$$\sigma^*(g_1 g_2) = \sigma^*(g_1) \sigma^*(g_2)$$

so that σ^* is in fact a homomorphism.

When \mathcal{V} is endowed with an inner product, and is thereby identified with its dual, and if σ is *unitary*, then $\sigma^* = \sigma$.

7.5.2 Let \mathcal{V}_j be G -spaces. We extend the actions of G to $\mathcal{V}_1 \oplus \mathcal{V}_2$ and $\mathcal{V}_1 \otimes \mathcal{V}_2$ by declaring¹⁰

$$(7.5.1) \quad \mathfrak{g}(v_1 \oplus v_2) = \mathfrak{g}v_1 \oplus \mathfrak{g}v_2 \quad \text{and} \quad \mathfrak{g}(v_1 \otimes v_2) = \mathfrak{g}v_1 \otimes \mathfrak{g}v_2$$

$\mathcal{L}(\mathcal{V}_1, \mathcal{V}_2) = \mathcal{V}_2 \otimes \mathcal{V}_1^*$ and as such it is a G -space.

7.5.3 G -MAPS. Let \mathcal{H}_j be G -spaces, $j = 1, 2$. A map $S : \mathcal{H}_1 \mapsto \mathcal{H}_2$ is a G -map if it commutes with the action of G . This means: for every $g \in G$, $S\mathfrak{g} = \mathfrak{g}S$. The domains of the various actions is more explicit in the diagram

$$\begin{array}{ccc} \mathcal{H}_1 & \xrightarrow{S} & \mathcal{H}_2 \\ \mathfrak{g} \downarrow & & \downarrow \mathfrak{g} \\ \mathcal{H}_1 & \xrightarrow{S} & \mathcal{H}_2 \end{array}$$

and the requirement is that it commute.

¹⁰Observe that the symbol \mathfrak{g} signifies, in (7.5.1) and elsewhere, different operators, acting on different spaces.

The prefix G - can be attached to all words describing linear maps, thus, a G -isomorphism is an isomorphism which is a G -map, etc.

If \mathcal{V}_j , $j = 1, 2$, are G -spaces, we denote by $\mathcal{L}_G(\mathcal{V}_1, \mathcal{V}_2)$ the space of linear G -maps of \mathcal{V}_1 into \mathcal{V}_2 .

7.5.4 Lemma. *Let $S : \mathcal{H}_1 \mapsto \mathcal{H}_2$ be a G -homomorphism. Then $\ker(S)$ is a subrepresentation, i.e., G -subspace, of \mathcal{H}_1 , and $\text{range}(S)$ is a subrepresentation of \mathcal{H}_2 .*

DEFINITION: Two representations \mathcal{H}_j of G are *equivalent* if there is a G -isomorphism $S : \mathcal{H}_1 \mapsto \mathcal{H}_2$, that is, if they are isomorphic as G -spaces.

7.5.5 AVERAGING, I. For a finite subgroup $\mathcal{G} \subset \mathbf{GL}(\mathcal{H})$ we write

$$(7.5.2) \quad I_{\mathcal{G}} = \{v \in \mathcal{H} : \mathbf{g}v = v \text{ for all } \mathbf{g} \in \mathcal{G}\}.$$

In words: $I_{\mathcal{G}}$ is the space of all the vectors in \mathcal{H} which are invariant under every \mathbf{g} in \mathcal{G} .

Theorem. *The operator*

$$(7.5.3) \quad \pi_{\mathcal{G}} = \frac{1}{|\mathcal{G}|} \sum_{\mathbf{g} \in \mathcal{G}} \mathbf{g}$$

is a projection onto $I_{\mathcal{G}}$.

PROOF: $\pi_{\mathcal{G}}$ is clearly the identity on $I_{\mathcal{G}}$. All we need to do is show that $\text{range}(\pi_{\mathcal{G}}) = I_{\mathcal{G}}$, and for that observe that if $v = \frac{1}{|\mathcal{G}|} \sum_{\mathbf{g} \in \mathcal{G}} \mathbf{g}u$, then

$$\mathbf{g}_1 v = \frac{1}{|\mathcal{G}|} \sum_{\mathbf{g} \in \mathcal{G}} \mathbf{g}_1 \mathbf{g} u$$

and since $\{\mathbf{g}_1 \mathbf{g} : \mathbf{g} \in \mathcal{G}\} = \mathcal{G}$, we have $\mathbf{g}_1 v = v$. ◀

7.5.6 AVERAGING, II. The operator $Q = \sum_{\mathbf{g} \in \mathcal{G}} \mathbf{g}^* \mathbf{g}$ is positive, selfadjoint, and can be used to define a new inner product

$$(7.5.4) \quad \langle v, u \rangle_Q = \langle Qv, u \rangle = \sum_{\mathbf{g} \in \mathcal{G}} \langle \mathbf{g}v, \mathbf{g}u \rangle$$

and the corresponding norm

$$\|v\|_Q^2 = \sum_{\mathbf{g} \in \mathcal{G}} \langle \mathbf{g}v, \mathbf{g}v \rangle = \sum_{\mathbf{g} \in \mathcal{G}} \|\mathbf{g}v\|^2.$$

Since $\{\mathbf{g} : \mathbf{g} \in \mathcal{G}\} = \{\mathbf{g}\mathbf{h} : \mathbf{g} \in \mathcal{G}\}$, we have

$$(7.5.5) \quad \langle \mathbf{h}v, \mathbf{h}u \rangle_Q = \sum_{\mathbf{g} \in \mathcal{G}} \langle \mathbf{g}\mathbf{h}v, \mathbf{g}\mathbf{h}u \rangle = \langle Qv, u \rangle,$$

and $\|\mathbf{h}v\|_Q = \|v\|_Q$. Thus, \mathcal{G} is a subgroup of the “unitary group” corresponding to $\langle \cdot, \cdot \rangle_Q$.

Denote by \mathcal{H}_Q the inner product space obtained by replacing the given inner-product by $\langle \cdot, \cdot \rangle_Q$. Let $\{u_1, \dots, u_n\}$ be an orthonormal basis of \mathcal{H} , and $\{v_1, \dots, v_n\}$ be an orthonormal basis of \mathcal{H}_Q . Define $S \in \mathbf{GL}(\mathcal{H})$ by imposing $Su_j = v_j$. Now, S is an isometry from \mathcal{H} onto \mathcal{H}_Q , \mathbf{g} unitary on \mathcal{H}_Q (for any $\mathbf{g} \in \mathcal{G}$), and S^{-1} an isometry from \mathcal{H}_Q back to \mathcal{H} ; hence $S^{-1}\mathbf{g}S \in U(n)$. In other words, S conjugates \mathcal{G} to a subgroup of the unitary group $U(\mathcal{H})$. This proves the following theorem

Theorem. *Every finite subgroup of $\mathbf{GL}(\mathcal{H})$ is conjugate to a subgroup of the unitary group $U(\mathcal{H})$.*

7.5.7 DEFINITION: A *unitary representation* of a group G in an inner-product space \mathcal{H} is a representation such that \mathbf{g} is unitary for all $g \in G$.

The following is an immediate corollary of Theorem 7.5.6

Theorem. *Every finite dimensional representation of a finite group is equivalent to a unitary representation.*

7.5.8 Let G be a finite group and \mathcal{H} a finite dimensional G -space (a finite dimensional representation of G).

A subspace $\mathcal{U} \subset \mathcal{H}$ is *G -invariant* if it is invariant under all the maps \mathbf{g} , $g \in G$. If $\mathcal{U} \subset \mathcal{H}$ is G -invariant, restricting the maps \mathbf{g} , $g \in G$, to \mathcal{U} defines \mathcal{U} as a representation of G and we refer to \mathcal{U} as a *subrepresentation* of \mathcal{H} .

A subspace \mathcal{U} is *G -reducing* if it is G -invariant and has a G -invariant complement, i.e., $\mathcal{H} = \mathcal{U} \oplus \mathcal{V}$ with both summands G -invariant.

Lemma. *Every G -invariant subspace is reducing.*

PROOF: Endow the space with the inner product given by (7.5.4) (which makes the representation unitary) and observe that if \mathcal{U} is a nontrivial G -invariant subspace, then so is its orthogonal complement, and we have a direct sum decomposition $\mathcal{H} = \mathcal{U} \oplus \mathcal{V}$ with both summands G -invariant. ◀

We say that (the representation) \mathcal{H} is *irreducible* if there is no non-trivial G -invariant subspace of \mathcal{H} and (*completely*) *reducible* otherwise. In the terminology of **V.2.7**, \mathcal{H} is irreducible if $(\mathcal{H}, \mathcal{G})$ is minimal.

Thus, if \mathcal{H} is reducible, there is a (non-trivial) direct sum decomposition $\mathcal{H} = \mathcal{U} \oplus \mathcal{V}$ with both summands G -invariant. We say, in this case, that σ is the sum of the representations \mathcal{U} and \mathcal{V} . If either representation is reducible we can write it as a sum of representations corresponding to a further direct sum decomposition of the space (\mathcal{U} or \mathcal{V}) into G invariant subspaces. After no more than $\dim \mathcal{H}$ such steps we obtain \mathcal{H} as a sum of irreducible representations. This proves the following theorem:

Theorem. *Every finite dimensional representation \mathcal{H} of a finite group G is a sum of irreducible representations. That is*

$$(7.5.6) \quad \mathcal{H} = \bigoplus \mathcal{U}_j$$

Uniqueness of the decomposition into irreducibles

Lemma. *Let \mathcal{V} and \mathcal{U} be irreducible subrepresentations of \mathcal{H} . Then, either $\mathcal{W} = \mathcal{U} \cap \mathcal{V} = \{0\}$, or $\mathcal{U} = \mathcal{V}$.*

PROOF: \mathcal{W} is clearly G -invariant. ◀

7.5.9 THE REGULAR REPRESENTATION. Let G be a finite group. Denote by $\ell^2(G)$ the vector space of all complex valued functions on G , and define the inner product, for $\varphi, \psi \in \ell^2(G)$, by

$$\langle \varphi, \psi \rangle = \sum_{x \in G} \varphi(x) \overline{\psi(x)}.$$

For $g \in G$, the *left translation by g* is the operator $\tau(g)$ on $\ell^2(G)$ defined by

$$(\tau(g)\varphi)(x) = \varphi(g^{-1}x).$$

Clearly $\tau(g)$ is linear and, in fact, unitary. Moreover,

$$(\tau(g_1g_2)\varphi)(x) = \varphi((g_1g_2)^{-1}x) = \varphi(g_2^{-1}(g_1^{-1}x)) = (\tau(g_1)\tau(g_2)\varphi)(x)$$

so that $\tau(g_1g_2) = \tau(g_1)\tau(g_2)$ and τ is a unitary representation of G . It is called the *regular representation of G* .

If $H \subset G$ is a subgroup we denote by $\ell^2(G/H)$ the subspace of $\ell^2(G)$ of the functions that are constant on left cosets of H .

Since multiplication on the left by arbitrary $g \in G$ maps left H -cosets onto left H -cosets, $\ell^2(G/H)$ is $\tau(g)$ invariant, and unless G is simple, that is—has no nontrivial subgroups, τ is reducible.

If H is not a maximal subgroup, that is, there exists a proper subgroup H_1 that contains H properly, then left cosets of H_1 split into left cosets of H so that $\ell^2(G/H_1) \subset \ell^2(G/H)$ and $\tau|_{\ell^2(G/H)}$ is reducible. This proves the following:

Lemma. *If the regular representation of G is irreducible, then G is simple.*

The converse is false! A cyclic group of order p , with prime p , is simple. Yet its regular representation is reducible. In fact,

Proposition. *Every representation of a finite abelian group is a direct sum of one-dimensional representations.*

PROOF: Exercise VII.5.2 ◀

7.5.10 Let \mathcal{W} be a G space and let $\langle \cdot, \cdot \rangle$ be an inner-product in \mathcal{W} . Fix a non-zero vector $u \in \mathcal{W}$ and, for $v \in \mathcal{W}$ and $g \in G$, define

$$(7.5.7) \quad f_v(g) = \langle \mathbf{g}^{-1}v, u \rangle$$

The map $S: v \mapsto f_v$ is a linear map from \mathcal{W} into $\ell^2(G)$. If \mathcal{W} is irreducible and $v \neq 0$, the set $\{\mathbf{g}v: g \in G\}$ spans \mathcal{W} which implies that $f_v \neq 0$, i.e., S is injective.

Observe that for $\gamma \in G$,

$$(7.5.8) \quad \tau(\gamma)f_v(g) = f_v(\gamma^{-1}g) = \langle \mathbf{g}^{-1}\gamma v, u \rangle = f_{\gamma v}(g),$$

so that the space $S\mathcal{W} = \mathcal{W}_S \subset \ell^2(G)$ is a reducing subspace of the regular representation of $\ell^2(G)$ and S maps σ onto the (restriction of the) regular representation τ (to) on \mathcal{W}_S .

This proves in particular

Proposition. *Every irreducible representation of G is equivalent to a subrepresentation of the regular representation.*

Corollary. *There are only a finite number of distinct irreducible representations of a finite group G .*

EXERCISES FOR SECTION 7.5

VII.5.1. If G is finite abelian group and σ a representation of G in \mathcal{H} , then the linear span of $\{\sigma(g) : g \in G\}$ is a selfadjoint commutative subalgebra of $\mathcal{L}\mathcal{H}$.

VII.5.2. Prove that every representation of a finite abelian group is a direct sum of one-dimensional representations.

Hint: 6.5.2

VII.5.3. Consider the representation of \mathbb{Z} in \mathbb{R}^2 defined by $\sigma(n) = \begin{pmatrix} 1 & n \\ 0 & 1 \end{pmatrix}$. Check the properties shown above for representations of finite groups that fail for σ .

Appendix

A.1 EQUIVALENCE RELATIONS — PARTITIONS.

A.1.1 EQUIVALENCE RELATIONS. A *binary relation* on a set X is a subset $R \subset X \times X$. We write xRy when $(x, y) \in R$.

EXAMPLES:

- a.* Equality: $R = \{(x, x) : x \in X\}$, xRy means $x = y$.
- b.* Order in \mathbb{Z} : $R = \{(x, y) : x < y\}$.

DEFINITION: An *equivalence relation* in a set X , is a binary relation (denoted here $x \equiv y$) that is

reflexive: for all $x \in X$, $x \equiv x$;

symmetric: for all $x, y \in X$, if $x \equiv y$, then $y \equiv x$;

and **transitive:** for all $x, y, z \in X$, if $x \equiv y$ and $y \equiv z$, then $x \equiv z$.

EXAMPLES:

- a.* Of the two binary relations above, *equality* is an equivalence relation, *order* is not.
- b.* Congruence modulo an integer. Here $X = \mathbb{Z}$, the set of integers. Fix an integer k . x is congruent to y modulo k and write $x \equiv y \pmod{k}$ if $x - y$ is an integer multiple of k .
- c.* For $X = \{(m, n) : m, n \in \mathbb{Z}, n \neq 0\}$, define $(m, n) \equiv (m_1, n_1)$ by the condition $mn_1 = m_1n$. This will be familiar if we write the pairs as $\frac{m}{n}$ instead of (m, n) and observe that the condition $mn_1 = m_1n$ is the one defining the equality of the rational fractions $\frac{m}{n}$ and $\frac{m_1}{n_1}$.

A.1.2 PARTITIONS.

DEFINITION: A *partition* of X is a collection \mathcal{P} of (pairwise) disjoint subsets $P_\alpha \subset X$ whose union is X .

A partition \mathcal{P} defines an equivalence relation: by definition, $x \equiv y$ if, and only if, x and y belong to the same element of the partition.

Conversely, given an equivalence relation on X , we define the *equivalence class* of $x \in X$ as the set $\mathcal{E}_x = \{y \in X : x \equiv y\}$. The defining properties of equivalence can be rephrased as: **a.** $x \in \mathcal{E}_x$, **b.** If $y \in \mathcal{E}_x$, then $x \in \mathcal{E}_y$, and **c.** If $y \in \mathcal{E}_x$, and $z \in \mathcal{E}_y$, then $z \in \mathcal{E}_x$. These conditions guarantee that different equivalence classes are disjoint and the collection of all the equivalence classes is a partition of X (which defines the given equivalence relation).

EXERCISES FOR SECTION A.1

A.1.1. Write $R_1 \subset \mathbb{R} \times \mathbb{R} = \{(x, y) : |x - y| < 1\}$ and $x \sim_1 y$ when $(x, y) \in R_1$. Is this an equivalence relation, and if not—what fails?

A.1.2. Identify the equivalence classes for congruence mod k .

A.2 MAPS

The terms used to describe properties of maps vary by author, by time, by subject matter, etc. We shall use the following:

A map $\varphi: X \mapsto Y$ is *injective* if $x_1 \neq x_2 \implies \varphi(x_1) \neq \varphi(x_2)$. Equivalent terminology: φ is *one-to-one* (or 1–1), or φ is a *monomorphism*.

A map $\varphi: X \mapsto Y$ is *surjective* if $\varphi(X) = \{\varphi(x) : x \in X\} = Y$. Equivalent terminology: φ is *onto*, or φ is an *epimorphism*.

A map $\varphi: X \mapsto Y$ is *bijective* if it is both injective and surjective: for every $y \in Y$ there is precisely one $x \in X$ such that $y = \varphi(x)$. Bijective maps are *invertible*—the inverse map defined by: $\varphi^{-1}(y) = x$ if $y = \varphi(x)$.

Maps that preserve some structure are called morphisms, often with a prefix providing additional information. Besides the *mono-* and *epi-* mentioned above, we use systematically *homomorphism*, *isomorphism*, etc.

A *permutation* of a set is a bijective map of the set onto itself.

A.3 GROUPS

A.3.1 DEFINITION: A *group* is a pair $(G, *)$, where G is a set and $*$ is a binary operation $(x, y) \mapsto x * y$, defined for all pairs $(x, y) \in G \times G$, taking values in G , and satisfying the following conditions:

- G-1 The operation is associative: For $x, y, z \in G$, $(x * y) * z = x * (y * z)$.
- G-2 There exists a unique element $e \in G$ called the *identity element* or the *unit* of G , such that $e * x = x * e = x$ for all $x \in G$.
- G-3 For every $x \in G$ there exists a unique element x^{-1} , called *the inverse* of x , such that $x^{-1} * x = x * x^{-1} = e$.

A group $(G, *)$ is *Abelian*, or *commutative* if $x * y = y * x$ for all x and y . The group operation in a commutative group is often written and referred to as addition, in which case the identity element is written as 0, and the inverse of x as $-x$.

When the group operation is written as multiplication, the operation symbol $*$ is sometimes written as a dot (i.e., $x \cdot y$ rather than $x * y$) and is often omitted altogether. We also simplify the notation by referring to the group, when the binary operation is “assumed known”, as G , rather than $(G, *)$.

EXAMPLES:

- a. $(\mathbb{Z}, +)$, the integers with standard addition.
- b. $(\mathbb{R} \setminus \{0\}, \cdot)$, the non-zero real numbers, standard multiplication.
- c. S_n , the *symmetric group on* $[1, \dots, n]$. Here n is a positive integer, the elements of S_n are all the permutations σ of the set $[1, \dots, n]$, and the operation is concatenation: for $\sigma, \tau \in S_n$ and $1 \leq j \leq n$ we set $(\tau\sigma)(j) = \tau(\sigma(j))$.

More generally, if X is a set, the collection $S(X)$ of permutations, i.e., invertible self-maps of X , is a group under concatenation. (Thus $S_n = S([1, \dots, n])$).

The first two examples are commutative; the third, if $n > 2$, is not.

A.3.2 Let G_i , $i = 1, 2$, be groups.

DEFINITION: A map $\varphi: G_1 \mapsto G_2$ is a *homomorphism* if

$$(A.3.1) \quad \varphi(xy) = \varphi(x)\varphi(y)$$

Notice that the multiplication on the left-hand side is in G_1 , while that on the right-hand side is in G_2 .

The definition of homomorphism is quite broad; we don't assume the mapping to be *injective* (1-1), nor *surjective* (onto). We use the proper adjectives explicitly whenever relevant: *monomorphism* for injective homomorphism and *epimorphism* for one that is surjective.

An *isomorphism* is a homomorphism which is *bijective*, that is both injective and surjective. Bijective maps are invertible, and the inverse of an isomorphism is an isomorphism. For the proof we only have to show that φ^{-1} is multiplicative (as in (A.3.1)), that is that for $g, h \in G_2$, $\varphi^{-1}(gh) = \varphi^{-1}(g)\varphi^{-1}(h)$. But, if $g = \varphi(x)$ and $h = \varphi(y)$, this is equivalent to $gh = \varphi(xy)$, which is the multiplicativity of φ .

If $\varphi: G_1 \mapsto G_2$ and $\psi: G_2 \mapsto G_3$ are both isomorphisms, then $\psi\varphi: G_1 \mapsto G_3$ is an isomorphism as well..

We say that two groups G and G_1 are *isomorphic* if there is an isomorphism of one onto the other. The discussion above makes it clear that this is an equivalence relation.

A.3.3 INNER AUTOMORPHISMS AND CONJUGACY CLASSES. An isomorphism of a group onto itself is called *an automorphism*. A special class of automorphisms, the *inner automorphisms*, are the *conjugations by elements* $y \in G$:

$$(A.3.2) \quad A_y x = y^{-1}xy$$

One checks easily (left as exercise) that for all $y \in G$, the map A_y is in fact an automorphism.

An important fact is that conjugacy, defined by $x \sim z$ if $z = A_y x = y^{-1}xy$ for some $y \in G$, is an equivalence relation. To check that every x is conjugate

to itself take $y = e$, the identity. If $z = A_y x$, then $x = A_{y^{-1}} z$, proving the symmetry. Finally, if $z = y^{-1} x y$ and $u = w^{-1} z w$, then

$$u = w^{-1} z w = w^{-1} y^{-1} x y w = (y w)^{-1} x (y w),$$

which proves the transitivity.

The equivalence classes defined on G by conjugation are called *conjugacy classes*.

A.3.4 SUBGROUPS AND COSETS.

DEFINITION: A subgroup of a group G is a subset $H \subset G$ such that

SG-1 H is closed under multiplication, that is, if $h_1, h_2 \in H$ then $h_1 h_2 \in H$.

SG-2 $e \in H$.

SG-3 If $h \in H$, then $h^{-1} \in H$

EXAMPLES:

- a. $\{e\}$, the subset whose only term is the identity element
- b. In \mathbb{Z} , the set $q\mathbb{Z}$ of all the integral multiples of some integer q . This is a special case of the following example.
- c. For any $x \in G$, the set $\{x^k\}_{k \in \mathbb{Z}}$ is *the subgroup generated by x* . The element x is of order m , if the group it generates is a cyclic group of order m . (That is if m is the smallest positive integer for which $x^m = e$). x has infinite order if $\{x^n\}$ is infinite, in which case $n \mapsto x^n$ is an isomorphism of \mathbb{Z} onto the group generated by x .
- d. If $\varphi: G \mapsto G_1$ is a homomorphism and e_1 denotes the identity in G_1 , then $\{g \in G: \varphi g = e_1\}$ is a subgroup of G (*the kernel of φ*).
- e. The subset of S_n of all the permutations that leave some (fixed) $l \in [1, \dots, n]$ in its place, that is, $\{\sigma \in S_n: \sigma(l) = l\}$.

Let $H \subset G$ a subgroup. For $x \in G$ write $xH = \{xz: z \in H\}$. Sets of the form xH are called *left cosets of H* .

Lemma. For any $x, y \in G$ the cosets xH and yH are either identical or disjoint. In other words, the collection of distinct xH is a partition of G .

PROOF: We check that the binary relation defined by “ $x \in yH$ ” which is usually denoted by $x \equiv y \pmod{H}$, is an equivalence relation. The cosets xH are the elements of the associated partition.

a. Reflexive: $x \in xH$, since $x = xe$ and $e \in H$.

b. Symmetric: If $y \in xH$ then $x \in yH$. $y \in xH$ means that there exists $z \in H$, such that $y = xz$. But then $yz^{-1} = x$, and since $z^{-1} \in H$, $x \in yH$.

c. Transitive: If $w \in yH$ and $y \in xH$, then $w \in xH$. For appropriate $z_1, z_2 \in H$, $y = xz_1$ and $w = yz_2 = xz_1z_2$, and $z_1z_2 \in H$. ◀

EXERCISES FOR SECTION A.3

A.3.1. Check that, for any group G and every $y \in G$, the map $A_yx = y^{-1}xy$ is an automorphism of G .

A.3.2. Let G be a finite group of order m . Let $H \subset G$ be a subgroup. Prove that the order of H divides m .

★A.4 GROUP ACTIONS

A.4.1 ACTIONS. DEFINITION: An *action of G on X* is a homomorphism φ of G into $S(X)$, the group of invertible self-maps (permutations) of X .

The action defines a map $(g, x) \mapsto \varphi(g)x$. The notation $\varphi(g)x$ often replaced, when φ is “understood”, by the simpler gx , and the assumption that φ is a homomorphism is equivalent to the conditions:

ga1. $ex = x$ for all $x \in X$, (e is the identity element of G).

ga2. $(g_1g_2)x = g_1(g_2x)$ for all $g_j \in G, x \in X$.

EXAMPLES:

a. G acts on itself ($X = G$) by left multiplication: $(x, y) \mapsto xy$.

b. G acts on itself ($X = G$) by right multiplication (by the inverse): $(x, y) \mapsto yx^{-1}$. (Remember that $(ab)^{-1} = b^{-1}a^{-1}$)

- c. G acts on itself by conjugation: $(x, y) \mapsto \varphi(x)y$ where $\varphi(x)y = xyx^{-1}$.
- d. S_n acts as mappings on $\{1, \dots, n\}$.

A.4.2 ORBITS. The orbit of an element $x \in X$ under the action of a group G is the set $\text{Orb}(x) = \{gx : g \in G\}$.

The orbits of a G action form a partition of X . This means that any two orbits, $\text{Orb}(x_1)$ and $\text{Orb}(x_2)$ are either identical (as sets) or disjoint. In fact, if $x \in \text{Orb}(y)$, then $x = g_0y$ and then $y = g_0^{-1}x$, and $gy = gg_0^{-1}x$. Since the set $\{gg_0^{-1} : g \in G\}$ is exactly G , we have $\text{Orb}(y) = \text{Orb}(x)$. If $\tilde{x} \in \text{Orb}(x_1) \cap \text{Orb}(x_2)$ then $\text{Orb}(\tilde{x}) = \text{Orb}(x_1) = \text{Orb}(x_2)$. The corresponding equivalence relation is: $x \equiv y$ when $\text{Orb}(x) = \text{Orb}(y)$.

EXAMPLES:

- a. A subgroup $H \subset G$ acts on G by right multiplication: $(h, g) \mapsto gh$. The orbit of $g \in G$ under this action is the (left) coset gH .
- b. S_n acts on $[1, \dots, n]$, $(\sigma, j) \mapsto \sigma(j)$. Since the action is transitive, there is a unique orbit— $[1, \dots, n]$.
- c. If $\sigma \in S_n$, the group $\langle \sigma \rangle$ (generated by σ) is the subgroup $\{\sigma^k\}$ of all the powers of σ . Orbit of elements $a \in [1, \dots, n]$ under the action of $\langle \sigma \rangle$, i.e. the set $\{\sigma^k(a)\}$, are called *cycles* of σ and are written (a_1, \dots, a_l) , where $a_{j+1} = \sigma(a_j)$, and l , the period of a_1 under σ , is the first positive integer such that $\sigma^l(a_1) = a_1$.

Notice that cycles are “enriched orbits”, that is orbits with some additional structure, here the cyclic order inherited from \mathbb{Z} . This cyclic order defines σ uniquely on the orbit, and is identified with the permutation that agrees with σ on the elements that appear in it, and leaves every other element in its place. For example, $(1, 2, 5)$ is the permutation that maps 1 to 2, maps 2 to 5, and 5 to 1, leaving every other element unchanged. Notice that n , the cardinality of the complete set on which S_n acts, does not enter the notation and is in fact irrelevant (provided that all the entries in the cycle are bounded by it; here $n \geq 5$). Thus, breaking $[1, \dots, n]$ into σ -orbits amounts to writing σ as a product of disjoint cycles.

A.4.3 CONJUGATION. Two actions of a group G , $\varphi_1: G \times X_1 \mapsto X_1$, and $\varphi_2: G \times X_2 \mapsto X_2$ are *conjugate to each other* if there is an invertible map $\Psi: X_1 \mapsto X_2$ such that for all $x \in G$ and $y \in X_1$,

$$(A.4.1) \quad \varphi_2(x)\Psi y = \Psi(\varphi_1(x)y) \quad \text{or, equivalently,} \quad \varphi_2 = \Psi\varphi_1\Psi^{-1}.$$

This is often stated as: *the following diagrams commute*

$$\begin{array}{ccc} X_1 & \xrightarrow{\varphi_1} & X_1 \\ \downarrow \Psi & & \downarrow \Psi \\ X_2 & \xrightarrow{\varphi_2} & X_2 \end{array} \quad \text{or} \quad \begin{array}{ccc} X_1 & \xrightarrow{\varphi_1} & X_1 \\ \Psi^{-1} \uparrow & & \downarrow \Psi \\ X_2 & \xrightarrow{\varphi_2} & X_2 \end{array}$$

meaning that the concatenation of maps associated with arrows along a path depends only on the starting and the end point, and not on the path chosen.

A.5 FIELDS, RINGS, AND ALGEBRAS

A.5.1 FIELDS.

DEFINITION: A (*commutative*) *field*, $(\mathbb{F}, +, \cdot)$ is a set \mathbb{F} endowed with two binary operations, *addition*: $(a, b) \mapsto a + b$, and *multiplication*: $(a, b) \mapsto a \cdot b$ (we often write ab instead of $a \cdot b$) such that:

F-1 $(\mathbb{F}, +)$ is a commutative group, its identity (*zero*) is denoted by 0.

F-2 $(\mathbb{F} \setminus \{0\}, \cdot)$ is a commutative group, whose identity is denoted 1, and $a \cdot 0 = 0 \cdot a = 0$ for all $a \in \mathbb{F}$.

F-3 Addition and multiplication are related by the *distributive law*:

$$a(b + c) = ab + ac.$$

EXAMPLES:

a. \mathbb{Q} , the field of rational numbers.

b. \mathbb{R} , the field of real numbers.

c. \mathbb{C} , the field of complex numbers.

d. \mathbb{Z}_2 denotes the field consisting of the two elements 0, 1, with addition and multiplication defined mod 2 (so that $1 + 1 = 0$).

Similarly, if p is a prime, the set \mathbb{Z}_p of residue classes mod p , with addition and multiplication mod p , is a field. (See exercise **I.5.2**.)

A.5.2 RINGS.

DEFINITION: A *ring* is a triplet $(R, +, \cdot)$, R is a set, $+$ and \cdot binary operations on R called addition, resp. multiplication, such that $(R, +)$ is a commutative group, the multiplication is associative (but not necessarily commutative), and the addition and multiplication are related by the *distributive laws*:

$$a(b + c) = ab + ac, \quad \text{and} \quad (b + c)a = ba + ca.$$

A *subring* R_1 of a ring R is a subset of R that is a ring under the operations induced by the ring operations, i.e., addition and multiplication, in R .

\mathbb{Z} is an example of a commutative ring with a multiplicative identity; $2\mathbb{Z}$, (the even integers), is a subring. $2\mathbb{Z}$ is an example of a commutative ring without a multiplicative identity.

A.5.3 ALGEBRAS.

DEFINITION: An *Algebra* over a field \mathbb{F} is a ring \mathcal{A} and a multiplication of elements of \mathcal{A} by scalars (elements of \mathbb{F}), that is, a map $\mathbb{F} \times \mathcal{A} \mapsto \mathcal{A}$ such that if we denote the image of (a, u) by au we have, for $a, b \in \mathbb{F}$ and $u, v \in \mathcal{A}$,

$$\begin{aligned} \text{identity:} & \quad 1u = u; \\ \text{associativity:} & \quad a(bu) = (ab)u, \quad a(uv) = (au)v; \\ \text{distributivity:} & \quad (a + b)u = au + bu, \quad \text{and} \quad a(u + v) = au + av. \end{aligned}$$

A *subalgebra* $\mathcal{A}_1 \subset \mathcal{A}$ is a subring of \mathcal{A} that is also closed under multiplication by scalars.

EXAMPLES:

a. $\mathbb{F}[x]$ – The algebra of polynomials in one variable x with coefficients from \mathbb{F} , and the standard addition, multiplication, and multiplication by scalars. It is an algebra over \mathbb{F} .

- b.** $\mathbb{C}[x, y]$ – The (algebra of) polynomials in two variables x, y with complex coefficients, and the standard operations. $\mathbb{C}[x, y]$ is “complex algebra”, that is an algebra over \mathbb{C} .

Notice that by restricting the scalar field to, say, \mathbb{R} , a complex algebra can be viewed as a “real algebra” i.e., and algebra over \mathbb{R} . The underlying field is part of the definition of an algebra. The “complex” and the “real” $\mathbb{C}[x, y]$ are *different algebras*.

- c.** $\mathcal{M}(n)$, the $n \times n$ matrices with matrix multiplication as product.

DEFINITION: A *left (resp. right) ideal* in a ring R is a subring I that is closed under multiplication on the left (resp. right) by elements of R : for $a \in R$ and $h \in I$ we have $ah \in I$ (resp. $ha \in I$). A *two-sided ideal* is a subring that is both a left ideal and a right ideal.

A left (resp. right, resp. two-sided) ideal in an algebra \mathcal{A} is a subalgebra of \mathcal{A} that is closed under left (resp. right, resp. either left or right) multiplication by elements of \mathcal{A} .

If the ring (resp. algebra) is commutative the adjectives “left”, “right” are irrelevant.

Assume that R has an identity element. For $g \in R$, the set $I_g = \{ag : a \in R\}$ is a left ideal in R , and is clearly the smallest (left) ideal that contains g .

Ideals of the form I_g are called *principal left ideals*, and g a *generator* of I_g . One defines *principal right ideals* similarly.

A.5.4 \mathbb{Z} AS A RING. Notice that since multiplication by an integer can be accomplished by repeated addition, the ring \mathbb{Z} has the (uncommon) property that every subgroup in it is in fact an ideal.

Another special property is: \mathbb{Z} is a principal ideal domain—every nontrivial¹ ideal $I \subset \mathbb{Z}$ is principal, that is, has the form $m\mathbb{Z}$ for some positive integer m .

In fact if m is the smallest positive element of I and $n \in I$, $n > 0$, we can “divide with remainder” $n = qm + r$ with q, r integers, and $0 \leq r < m$. Since both n and qm are in I so is r . Since m is the smallest positive element in I ,

¹Not reduced to $\{0\}$.

$r = 0$ and $n = qm$. Thus, all the positive elements of I are divisible by m (and so are their negatives).

If $m_j \in \mathbb{Z}$, $j = 1, 2$, the set $I_{m_1, m_2} = \{n_1 m_1 + n_2 m_2 : n_1, n_2 \in \mathbb{Z}\}$ is an ideal in \mathbb{Z} , and hence has the form $g\mathbb{Z}$. As g divides every element in I_{m_1, m_2} , it divides both m_1 and m_2 ; as $g = n_1 m_1 + n_2 m_2$ for appropriate n_j , every common divisor of m_1 and m_2 divides g . It follows that g is their *greatest common divisor*, $g = \gcd(m_1, m_2)$. We summarize:

Proposition. *If m_1 and m_2 are integers, then for appropriate integers n_1, n_2 ,*

$$\gcd(m_1, m_2) = n_1 m_1 + n_2 m_2.$$

EXERCISES FOR SECTION A.5

A.5.1. Let R be a ring with identity, $B \subset R$ a set. Prove that *the ideal generated by B* , that is the smallest ideal that contains B , is: $I = \{\sum a_j b_j : a_j \in R, b_j \in B\}$.

A.5.2. Verify that \mathbb{Z}_p is a field.

Hint: If p is a prime and $0 < m < p$ then $\gcd(m, p) = 1$.

A.5.3. Prove that the set of invertible elements in a ring with and identity is a multiplicative group.

A.5.4. Show that the set of polynomials $\{P : P = \sum_{j \geq 2} a_j x^j\}$ is an ideal in $\mathbb{F}[x]$, and that $\{P : P = \sum_{j \leq 7} a_j x^j\}$ is an additive subgroup but not an ideal.

A.6 POLYNOMIALS

Let \mathbb{F} be a field and $\mathbb{F}[x]$ the algebra of polynomials $P = \sum_0^n a_j x^j$ in the variable x with coefficients from \mathbb{F} . The *degree* of P , $\deg(P)$, is the highest power of x appearing in P with non-zero coefficient. If $\deg(P) = n$, then $a_n x^n$ is called *the leading term* of P , and a_n *the leading coefficient*. A polynomial is called *monic* if its leading coefficient is 1.

A.6.1 DIVISION WITH REMAINDER. By definition, an ideal in a ring is *principal* if it consists of all the multiples of one of its elements, called a *generator* of the ideal. The ring $\mathbb{F}[x]$ shares with \mathbb{Z} the property of being a *principal ideal domain*—every ideal is principal. The proof for $\mathbb{F}[x]$ is virtually the same as the one we had for \mathbb{Z} , and is again based on *division with remainder*.

Theorem. Let $P, F \in \mathbb{F}[x]$. There exist polynomials $Q, R \in \mathbb{F}[x]$ such that $\deg(R) < \deg(F)$, and

$$(A.6.1) \quad P = QF + R.$$

PROOF: Write $P = \sum_{j=0}^n a_j x^j$ and $F = \sum_{j=0}^m b_j x^j$ with $a_n \neq 0$ and $b_m \neq 0$, so that $\deg(P) = n$, $\deg(F) = m$.

If $n < m$ there is nothing to prove: $P = 0 \cdot F + P$.

If $n \geq m$, we write $q_{n-m} = a_n/b_m$, and $P_1 = P - q_{n-m}x^{n-m}F$, so that $P = q_{n-m}x^{n-m}F + P_1$ with $n_1 = \deg(P_1) < n$.

If $n_1 < m$ we are done. If $n_1 \geq m$, write the leading term of P_1 as $a_{1,n_1}x^{n_1}$, and set $q_{n_1-m} = a_{1,n_1}/b_m$, and $P_2 = P_1 - q_{n_1-m}x^{n_1-m}F$. Now $\deg(P_2) < \deg(P_1)$ and $P = (q_{n-m}x^{n-m} + q_{n_1-m}x^{n_1-m})F + P_2$.

Repeating the procedure a total of k times, $k \leq n - m + 1$, we obtain $P = QF + P_k$ with $\deg(P_k) < m$, and the statement follows with $R = P_k$. ◀

Corollary. Let $I \subset \mathbb{F}[x]$ be an ideal, and let P_0 be an element of minimal degree in I . Then P_0 is a generator for I .

PROOF: If $P \in I$, write $P = QP_0 + R$, with $\deg(R) < \deg(P_0)$. Since $R = P - QP_0 \in I$, and 0 is the only element of I whose degree is smaller than $\deg(P_0)$, $P = QP_0$. ◀

The generator P_0 is unique up to multiplication by a scalar. If P_1 is another generator, each of the two divides the other and since the degree has to be the same the quotients are scalars. It follows that if we normalize P_0 by requiring that it be *monic*, that is with leading coefficient 1, it is unique and we refer to it as *the* generator.

A.6.2 Given polynomials $P_j, j = 1, \dots, l$ any ideal that contains them all must contain all the polynomials $P = \sum q_j P_j$ with arbitrary polynomial coefficients q_j . On the other hand the set of all these sums is clearly an ideal in $\mathbb{F}[x]$. It follows that the ideal generated by $\{P_j\}$ is equal to the set of polynomials of the form $P = \sum q_j P_j$ with polynomial coefficients q_j .

The generator G of this ideal divides every one of the P_j 's, and, since G can be expressed as $\sum q_j P_j$, every common factor of all the P_j 's divides G . In

other words, $G = \gcd\{P_1, \dots, P_l\}$, the greatest common divisor of $\{P_j\}$. This implies

Theorem. *Given polynomials P_j , $j = 1, \dots, l$ there exist polynomials q_j such that $\gcd\{P_1, \dots, P_l\} = \sum q_j P_j$.*

In particular:

Corollary. *If P_1 and P_2 are relatively prime, there exist polynomials q_1, q_2 such that $P_1 q_1 + P_2 q_2 = 1$.*

A.6.3 FACTORIZATION. A polynomial P in $\mathbb{F}[x]$ is *irreducible* or *prime* if it has no proper factors, that is, if every factor of P is either scalar multiple of P or a scalar.

Lemma. *If $\gcd(P, P_1) = 1$ and $P \mid P_1 P_2$, then $P \mid P_2$.*

PROOF: There exist q, q_1 such that $qP + q_1 P_1 = 1$. Then the left-hand side of $qPP_2 + q_1 P_1 P_2 = P_2$ is divisible by P , and hence so is P_2 . ◀

Theorem (Prime power factorization). *Every $P \in \mathbb{F}[x]$ admits a factorization $P = \prod \Phi_j^{m_j}$, where each factor Φ_j is irreducible in $\mathbb{F}[x]$, and they are all distinct.*

The factorization is unique up to the order in which the factors are enumerated, and up to multiplication by non-zero scalars.

A.6.4 THE FUNDAMENTAL THEOREM OF ALGEBRA. A field \mathbb{F} is *algebraically closed* if it has the property that every $P \in \mathbb{F}[x]$ has roots in \mathbb{F} , that is elements $\lambda \in \mathbb{F}$ such that $P(\lambda) = 0$. The so-called *fundamental theorem of algebra* states that \mathbb{C} is algebraically closed.

Theorem. *Given a non-constant polynomial P with complex coefficients, there exist complex numbers λ such that $P(\lambda) = 0$.*

A.6.5 We now observe that $P(\lambda) = 0$ is equivalent to the statement that $(z - \lambda)$ divides P . By Theorem A.6.1, $P(z) = (z - \lambda)Q(z) + R$ with $\deg R$ smaller than $\deg(z - \lambda) = 1$, so that R is a constant. Evaluating $P(z) = (z - \lambda)Q(z) + R$ at $z = \lambda$ shows that $R = P(\lambda)$, hence the claimed equivalence. It follows that a non-constant polynomial $P \in \mathbb{C}[z]$ is prime if and only if it is linear.

Theorem. *Let $P \in \mathbb{C}[z]$ be a polynomial of degree n . There exist complex numbers $\lambda_1, \dots, \lambda_n$, (not necessarily distinct), and $a \neq 0$ (the leading coefficient of P), such that*

$$(A.6.2) \quad P(z) = a \prod_{j=1}^n (z - \lambda_j).$$

The theorem and its proof apply verbatim to polynomials over any algebraically closed field.

A.6.6 FACTORIZATION IN $\mathbb{R}[x]$. The factorization (A.6.2) applies, of course, to polynomials with real coefficients, but the roots need not be real. The basic example is $P(x) = x^2 + 1$ with the roots $\pm i$.

We observe that for polynomials P whose coefficients are all real, we have $P(\bar{\lambda}) = \overline{P(\lambda)}$, which means in particular that if λ is a root of P then so is $\bar{\lambda}$.

A second observation is that

$$(A.6.3) \quad (x - \lambda)(x - \bar{\lambda}) = x^2 - 2x\Re\lambda + |\lambda|^2$$

has real coefficients.

Combining these observations with (A.6.2) we obtain that the prime factors in $\mathbb{R}[x]$ are the linear polynomials and the quadratic of the form (A.6.3) where $\lambda \notin \mathbb{R}$.

Theorem. *Let $P \in \mathbb{R}[x]$ be a polynomial of degree n . P admits a factorization*

$$(A.6.4) \quad P(z) = a \prod (x - \lambda_j) \prod Q_j(x),$$

where a is the leading coefficient, $\{\lambda_j\}$ is the set of real zeros of P and Q_j are irreducible quadratic polynomials of the form (A.6.3) corresponding to (pairs of conjugate) non-real roots of P .

Either product may be empty, in which case it is interpreted as 1.

As mentioned above, the factors appearing in (A.6.4) need not be distinct—the same factor may be repeated several times. We can rewrite the product as

$$(A.6.5) \quad P(z) = a \prod (x - \lambda_j)^{l_j} \prod Q_j^{k_j}(x),$$

with λ_j and Q_j now distinct, and the exponents l_j resp. k_j their multiplicities. The factors $(x - \lambda_j)^{l_j}$ and $Q_j^{k_j}(x)$ appearing in (A.6.5) are pairwise relatively prime.

Index

- Adjoint
 - matrix of, 103
 - of an operator, 103
- Affine subspace, 21
- Algebra, 145
- Alternating n -form, 56
- Annihilator, 45
- Basic partition, 125
- Basis, 10
 - dual, 44
 - standard, 11
- Bilinear
 - map, 54
- Bilinear form, 44, 47
- Canonical
 - prime-power decomposition, 78
- Cauchy–Schwarz, 96
- Characteristic polynomial
 - of a matrix, 62
 - of an operator, 60
- Codimension, 12
- Complement, 6
- Coset, 141
- Cycle, 51
- Cyclic
 - decomposition, 86
 - system, 70
 - vector, 70
- Decomposition
 - cyclic, 86
- Determinant
 - of a matrix, 61
 - of an operator, 58
- Diagonal sum, 76, 80
- Dimension, 11, 12
- Direct sum
 - formal, 5
 - of subspaces, 6
- Eigenspace, 60, 66
 - generalized, 79, 90
- Eigenvalue, 49, 60, 66
- Eigenvector, 49, 60, 66
- Elementary divisor, 88
- Equivalence relation, 137
- Euclidean space, 95
- Factorization
 - in $\mathbb{R}[x]$, 150
 - prime-power, 78, 79, 149
- Field, 1, 144
- Flag, 68
- Frobenius, 125
- Gaussian elimination, 16, 18
- Group, 1, 139
- Hadamard’s inequality, 101
- Hamilton-Cayley, 71
- Hermitian
 - form, 95
 - quadratic form, 102

- Ideal, 146
- Idempotent, 101
- Independent
 - subspaces, 5
 - vectors, 9
- Inertia, law of, 118
- Inner-product, 95
- Irreducible
 - polynomial, 149
 - system, 75
- Isomorphism, 3
- Jordan canonical form, 88, 90
- Kernel, 35
- Ladder, 68
- Linear
 - system, 65
- Linear equations
 - homogeneous, 14
 - non-homogeneous, 14
- Markov chain, 127
 - reversible, 128
- Matrix
 - orthogonal, 104
 - unitary, 104
 - augmented, 16
 - companion, 72
 - diagonal, 5
 - Hermitian, 103
 - nonnegative, 121
 - permutation, 30
 - positive, 118
 - self-adjoint, 103
 - stochastic, 126
 - strongly transitive, 125
 - transitive, 122
 - triangular, 5, 63, 68
- Minimal
 - system, 74
- Minimal polynomial, 72
 - for (T, v) , 70
- Minmax principle, 108
- Monic polynomial, 147
- Multilinear
 - form, 54
 - map, 53
- Nilpotent, 83
- Nilspace, 79, 90
- Nonsingular, 37
- Norm, 39
- Normal
 - operator, 109
- Nullity, 35
- Nullspace, 35
- Operator
 - nonnegative, 111
 - normal, 109
 - orthogonal, 104
 - positive, 111
 - self-adjoint, 105
 - unitary, 104
- Orientation, 59
- Orthogonal
 - operator, 104
 - projection, 99
 - vectors, 97
- Orthogonal equivalence, 105
- Orthonormal, 98
- Period group, 125
- Permutation, 51, 138
- Perron, 118
- Polarization, 101

- Primary components, 78
- Probability vector, 127
- Projection
 - along a subspace, 24
 - orthogonal, 99
- Quadratic form
 - positive, 117
- Quadratic forms, 115
- Quotient space, 6
- Range, 35
- Rank
 - column, 20
 - of a matrix, 21
 - of an operator, 35
 - row, 17
- Reducing subspace, 75
- Regular representation, 133
- Ring, 145
- Row echelon form, 18
- Row equivalence, 17
- Schur's lemma, 74
- Self-adjoint
 - algebra, 110
 - matrix, 108
 - operator, 105
- Semisimple, 81
- k -shift, 84
- Similar, 35
- Similarity, 34
- Solution-set, 4
- Span, 5, 9
- Spectral mapping theorem, 66
- Spectral norm, 118
- Spectral Theorems, 106–110
- Spectrum, 60, 66, 79
- Steinitz' lemma, 11
- Symmetric group, 51
- Tensor product, 7
- Trace, 62
- Transition matrix, 127
- Transposition, 51
- Unitary
 - operator, 104
 - space, 95
- Unitary equivalence, 105
- Vandermonde, 64
- Vector space, 1
 - complex, 1
 - real, 1

Symbols

\mathbb{C} , 1
 \mathbb{Q} , 1
 \mathbb{R} , 1
 χ_T , 60
 \mathbf{C}_v , 25
 $\mathbf{C}_{w,v}$, 33
 $\dim \mathcal{V}$, 11
 $\{e_1, \dots, e_n\}$, 10
 \mathbb{F}^n , 2
 $\mathbb{F}[x]$, 2
 $\mathbf{GL}(\mathcal{H})$, 131
 $\mathbf{GL}(\mathcal{V})$, 27
 $\text{height}[v]$, 83
 $\mathcal{M}(n; \mathbb{F})$, 2
 $\mathcal{M}(n, m; \mathbb{F})$, 2
 minP_T , 72
 $\text{minP}_{T,v}$, 69
 $O(n)$, 104
 $\mathcal{P}(T)$, 27, 82
 \mathbf{S}_n , 51
 $\text{span}[E]$, 5
 $\text{span}[T, v]$, 66
 $\| \cdot \|_{sp}$, 118
 $T_{\mathcal{W}}$, 68
 $U(n)$, 104