

LINEAR ALGEBRA

SIMON WADSLEY

CONTENTS

1. Vector spaces	2
1.1. Definitions and examples	2
1.2. Linear independence, bases and the Steinitz exchange lemma	4
1.3. Direct sum	8
2. Linear maps	9
2.1. Definitions and examples	9
2.2. Linear maps and matrices	11
2.3. The first isomorphism theorem and the rank-nullity theorem	13
2.4. Change of basis	16
2.5. Elementary matrix operations	18
3. Duality	19
3.1. Dual spaces	19
3.2. Dual maps	21
4. Bilinear Forms (I)	23
5. Determinants of matrices	26
6. Endomorphisms	30
6.1. Invariants	30
6.2. Minimal polynomials	33
6.3. The Cayley-Hamilton Theorem	36
6.4. Multiplicities of eigenvalues and Jordan Normal Form	39
7. Bilinear forms (II)	45
7.1. Symmetric bilinear forms and quadratic forms	45
7.2. Hermitian forms	49
8. Inner product spaces	51
8.1. Definitions and basic properties	51
8.2. Gram–Schmidt orthogonalisation	52
8.3. Adjoints	54
8.4. Spectral theory	56

LECTURE 1

1. VECTOR SPACES

Linear algebra can be summarised as the study of vector spaces and linear maps between them. This is a second ‘first course’ in Linear Algebra. That is to say, we will define everything we use but will assume some familiarity with the concepts (picked up from the IA course Vectors & Matrices for example).

1.1. Definitions and examples.

Examples.

- (1) For each non-negative integer n , the set \mathbf{R}^n of column vectors of length n with real entries is a vector space (over \mathbf{R}). An $(m \times n)$ -matrix A with real entries can be viewed as a linear map $\mathbf{R}^n \rightarrow \mathbf{R}^m$ via $v \mapsto Av$. In fact, as we will see, every linear map from $\mathbf{R}^n \rightarrow \mathbf{R}^m$ is of this form. This is the motivating example and can be used for intuition throughout this course. However, it comes with a specified system of co-ordinates given by taking the various entries of the column vectors. A substantial difference between this course and Vectors & Matrices is that we will work with vector spaces without a specified set of co-ordinates. We will see a number of advantages to this approach as we go.
- (2) Let X be a set and $\mathbf{R}^X := \{f: X \rightarrow \mathbf{R}\}$ be equipped with an addition given by $(f + g)(x) := f(x) + g(x)$ and a multiplication by scalars (in \mathbf{R}) given by $(\lambda f)(x) = \lambda(f(x))$. Then \mathbf{R}^X is a vector space (over \mathbf{R}) in some contexts called the space of scalar fields on X . More generally, if V is a vector space over \mathbf{R} then $\mathbf{V}^X = \{f: X \rightarrow V\}$ is a vector space in a similar manner — a space of vector fields on X .
- (3) If $[a, b]$ is a closed interval in \mathbf{R} then $C([a, b], \mathbf{R}) := \{f \in \mathbf{R}^{[a, b]} \mid f \text{ is continuous}\}$ is an \mathbf{R} -vector space by restricting the operations on $\mathbf{R}^{[a, b]}$. Similarly

$$C^\infty([a, b], \mathbf{R}) := \{f \in C([a, b], \mathbf{R}) \mid f \text{ is infinitely differentiable}\}$$

is an \mathbf{R} -vector space.

- (4) The set of $(m \times n)$ -matrices with real entries is a vector space over \mathbf{R} .

Notation. We will use \mathbf{F} to denote an arbitrary field. However the schedules only require consideration of \mathbf{R} and \mathbf{C} in this course. If you prefer you may understand \mathbf{F} to always denote either \mathbf{R} or \mathbf{C} (and the examiners **must** take this view).

What do our examples of vector spaces above have in common? In each case we have a notion of addition of ‘vectors’ and scalar multiplication of ‘vectors’ by elements in \mathbf{R} .

Definition. An \mathbf{F} -vector space is an abelian group $(V, +)$ equipped with a function $\mathbf{F} \times V \rightarrow V$; $(\lambda, v) \mapsto \lambda v$ such that

- (i) $\lambda(\mu v) = (\lambda\mu)v$ for all $\lambda, \mu \in \mathbf{F}$ and $v \in V$;
- (ii) $\lambda(u + v) = \lambda u + \lambda v$ for all $\lambda \in \mathbf{F}$ and $u, v \in V$;
- (iii) $(\lambda + \mu)v = \lambda v + \mu v$ for all $\lambda, \mu \in \mathbf{F}$ and $v \in V$;
- (iv) $1v = v$ for all $v \in V$.

Note that this means that we can add, subtract and rescale elements in a vector space and these operations behave in the ways that we are used to. Note also that in general a vector space does not come equipped with a co-ordinate system, or

notions of length, volume or angle. We will discuss how to recover these later in the course. At that point particular properties of the field \mathbf{F} will be important.

Convention. We will always write 0 to denote the additive identity of a vector space V . By slight abuse of notation we will also write 0 to denote the vector space $\{0\}$.

Exercise.

- (1) Convince yourself that all the vector spaces mentioned thus far do indeed satisfy the axioms for a vector space.
- (2) Show that for any v in any vector space V , $0v = 0$ and $(-1)v = -v$

Definition. Suppose that V is a vector space over \mathbf{F} . A subset $U \subset V$ is an (\mathbf{F} -linear) subspace if

- (i) for all $u_1, u_2 \in U$, $u_1 + u_2 \in U$;
- (ii) for all $\lambda \in \mathbf{F}$ and $u \in U$, $\lambda u \in U$;
- (iii) $0 \in U$.

Remarks.

- (1) It is straightforward to see that $U \subset V$ is a subspace if and only if $U \neq \emptyset$ and $\lambda u_1 + \mu u_2 \in U$ for all $u_1, u_2 \in U$ and $\lambda, \mu \in \mathbf{F}$.
- (2) If U is a subspace of V then U is a vector space under the inherited operations.

Examples.

- (1) $\left\{ \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \in \mathbf{R}^3 : x_1 + x_2 + x_3 = t \right\}$ is a subspace of \mathbf{R}^3 if and only if $t = 0$.
- (2) Let X be a set. We define the *support* of a function $f: X \rightarrow \mathbf{F}$ to be

$$\text{supp } f := \{x \in X : f(x) \neq 0\}.$$

Then $\{f \in \mathbf{F}^X : |\text{supp } f| < \infty\}$ is a subspace of \mathbf{F}^X since we can compute $\text{supp } 0 = \emptyset$, $\text{supp}(f + g) \subset \text{supp } f \cup \text{supp } g$ and $\text{supp } \lambda f = \text{supp } f$ if $\lambda \neq 0$.

Definition. Suppose that U and W are subspaces of a vector space V over \mathbf{F} . Then the *sum* of U and W is the set

$$U + W := \{u + w : u \in U, w \in W\}.$$

Proposition. If U and W are subspaces of a vector space V over \mathbf{F} then $U \cap W$ and $U + W$ are also subspaces of V .

Proof. Certainly both $U \cap W$ and $U + W$ contain 0 . Suppose that $v_1, v_2 \in U \cap W$, $u_1, u_2 \in U$, $w_1, w_2 \in W$, and $\lambda, \mu \in \mathbf{F}$. Then $\lambda v_1 + \mu v_2 \in U \cap W$ and

$$\lambda(u_1 + w_1) + \mu(u_2 + w_2) = (\lambda u_1 + \mu u_2) + (\lambda w_1 + \mu w_2) \in U + W.$$

So $U \cap W$ and $U + W$ are subspaces of V . □

LECTURE 2

Definition. Suppose that V is a vector space over \mathbf{F} and U is a subspace of V . Then the *quotient space* V/U is the abelian group V/U equipped with the scalar multiplication $\mathbf{F} \times V/U \rightarrow V/U$ given by

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

for $\lambda \in \mathbf{F}$ and $v + U \in V/U$.

Proposition. V/U with the above structure is an \mathbf{F} -vector space.

Proof. First suppose $v_1 + U = v_2 + U \in V/U$. Then $(v_1 - v_2) \in U$ and so

$$\lambda v_1 - \lambda v_2 = \lambda(v_1 - v_2) \in U$$

for each $\lambda \in \mathbf{F}$ since U is a subspace. Thus $\lambda v_1 + U = \lambda v_2 + U$ and the scalar multiplication function $\mathbf{F} \times V/U \rightarrow V/U$ is well-defined.

Notice that each of the four axioms (i)-(iv) that must be verified to show that this scalar multiplication makes V/U into a vector space is now an almost immediate consequence of the fact that the same axiom must hold for the scalar multiplication on V . For example to see (i) we observe that for $v + U \in V/U$ and $\lambda, \mu \in \mathbf{F}$

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

□

1.2. Linear independence, bases and the Steinitz exchange lemma.

Definition. Let V be a vector space over \mathbf{F} and $S \subset V$ a subset of V . Then the *span* of S in V is the set of all finite \mathbf{F} -linear combinations of elements of S ,

$$\langle S \rangle := \left\{ \sum_{i=1}^n \lambda_i s_i : \lambda_i \in \mathbf{F}, s_i \in S, n \geq 0 \right\}$$

Remarks.

- (1) $\langle S \rangle$ only consists of *finite* linear combinations of elements of S .
- (2) For any subset $S \subset V$, $\langle S \rangle$ is the smallest subspace of V containing S .

Example. Suppose that V is \mathbf{R}^3 .

$$\text{If } S = \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 2 \\ 2 \end{pmatrix} \right\} \text{ then } \langle S \rangle = \left\{ \begin{pmatrix} a \\ b \\ b \end{pmatrix} : a, b \in \mathbf{R} \right\}.$$

Note also that every subset of S of order 2 has the same span as S .

Example. Let X be a set and for each $x \in X$, define $\delta_x : X \rightarrow \mathbf{F}$ by

$$\delta_x(y) = \begin{cases} 1 & \text{if } y = x \\ 0 & \text{if } y \neq x. \end{cases}$$

Then $\langle \delta_x : x \in X \rangle = \{f \in \mathbf{F}^X : |\text{supp} f| < \infty\}$.

Definition. Let V be a vector space over \mathbf{F} and $S \subset V$.

- (i) We say that S *spans* V if $V = \langle S \rangle$.
- (ii) We say that S is *linearly independent (LI)* if, whenever

$$\sum_{i=1}^n \lambda_i s_i = 0$$

with $\lambda_i \in \mathbf{F}$, and s_i distinct elements of S , it follows that $\lambda_i = 0$ for all i . If S is not linearly independent then we say that S is *linearly dependent (LD)*.

- (iii) We say that S is a *basis* for V if S spans and is linearly independent.

If V has a finite basis we say that V is *finite-dimensional*.

Note that it is not yet clear that if V is finite-dimensional then all bases must have the same size. So we cannot define the dimension of V yet even when it is known to be finite-dimensional. Fixing this is our next main goal.

Example. Suppose that V is \mathbf{R}^3 and $S = \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 2 \\ 2 \end{pmatrix} \right\}$. Then S is linearly

dependent since $1 \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + 2 \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} + (-1) \begin{pmatrix} 1 \\ 2 \\ 2 \end{pmatrix} = 0$. Moreover S does not span V since

$\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$ is not in $\langle S \rangle$. However, every subset of S of order 2 is linearly independent and forms a basis for $\langle S \rangle$.

Remark. Note that no linearly independent set can contain the zero vector since $1 \cdot 0 = 0$.

Convention. The span of the empty set $\langle \emptyset \rangle$ is the zero subspace 0 . Thus the empty set is a basis of 0 . One may consider this to not be so much a convention as the only reasonable interpretation of the definitions of span, linearly independent and basis in this case.

Lemma. A subset S of a vector space V over \mathbf{F} is linearly dependent if and only if there exist $s_0, s_1, \dots, s_n \in S$ distinct and $\lambda_1, \dots, \lambda_n \in \mathbf{F}$ such that $s_0 = \sum_{i=1}^n \lambda_i s_i$.

Proof. Suppose that S is linearly dependent so that $\sum \lambda_i s_i = 0$ for some $s_i \in S$ distinct and $\lambda_i \in \mathbf{F}$ with $\lambda_j \neq 0$ say. Then

$$s_j = \sum_{i \neq j} \frac{-\lambda_i}{\lambda_j} s_i.$$

Conversely, if $s_0 = \sum_{i=1}^n \lambda_i s_i$ then $(-1)s_0 + \sum_{i=1}^n \lambda_i s_i = 0$. □

Proposition. Let V be a vector space over \mathbf{F} . Then $S \subset V$ is a basis for V if and only if every element $v \in V$ can be written uniquely as $v = \sum_{s \in S} \lambda_s s$ with $\lambda_s \in \mathbf{F}$ and all but finitely many $\lambda_s = 0$.

Remark. Note that $\sum_{s \in S} \lambda_s s$ makes sense whenever all but finitely many λ_s are zero since we then only summing finitely many non-zero terms — if you are concerned you may define

$$\sum_{s \in S} \lambda_s s := \sum_{\substack{s \in S \\ \lambda_s \neq 0}} \lambda_s s$$

in this case. We will here and elsewhere abbreviate ‘all but finitely many’ as *almost all*.

Proof. First we observe that by definition S spans V if and only if every element v of V can be written in at least one way as $v = \sum_{s \in S} \lambda_s s$ with $\lambda_s \in \mathbf{F}$ and almost all $\lambda_s = 0$.

So it suffices to show that S is linearly independent if and only if there is at most one such expression for every $v \in V$.

Suppose that S is linearly independent and $v = \sum_{s \in S} \lambda_s s = \sum_{s \in S} \mu_s s$ with $\lambda_s, \mu_s \in \mathbf{F}$ and almost all zero. Then, $\sum_s (\lambda_s - \mu_s) s = 0$. Thus by definition of linear independence, $\lambda_s - \mu_s = 0$ for all $s \in S$ and so $\lambda_s = \mu_s$ for all s .

Conversely if S is linearly dependent then we can write

$$\sum_{s \in S} \lambda_s s = 0 = \sum_{s \in S} 0 \cdot s$$

for some $\lambda_s \in \mathbf{F}$ almost all zero but not all zero. Thus there are two ways to write 0 as an \mathbf{F} -linear combination of the s . \square

The following result is necessary for a good notion of dimension for vector spaces.

Theorem (Steinitz exchange lemma). *Let V be a vector space over \mathbf{F} . Suppose that $S = \{e_1, \dots, e_n\}$ is a linearly independent subset of V and $T \subset V$ spans V . Then there is a subset D of T of order n such that $(T \setminus D) \cup S$ spans V . In particular $|S| \leq |T|$.*

This is sometimes stated as follows (with the assumption that T is finite).

Corollary. *If $\{e_1, \dots, e_n\} \subset V$ is linearly independent and $\{f_1, \dots, f_m\}$ spans V . Then $n \leq m$ and, possibly after reordering the f_i , $\{e_1, \dots, e_n, f_{n+1}, \dots, f_m\}$ spans V .*

LECTURE 3

Corollary (Corollary of Steinitz exchange lemma). *Let V be a vector space with a basis of order n .*

- (a) *Every basis of V has order n .*
- (b) *Any n LI vectors in V form a basis for V .*
- (c) *Any n vectors in V that span V form a basis for V .*
- (d) *Any set of linearly independent vectors in V can be extended to a basis for V .*
- (e) *Any finite spanning set in V contains a basis for V .*

Proof. Suppose that $S = \{e_1, \dots, e_n\}$ is a basis for V .

(a) Suppose that T is another basis of V . Since S spans V and any finite subset of T is linearly independent $|T| \leq n$. Since T spans and S is linearly independent $|T| \geq n$. Thus $|T| = n$ as required.

(b) Suppose T is a LI subset of V of order n . If T did not span we could choose $v \in V \setminus \langle T \rangle$. Then $T \cup \{v\}$ is a LI subset of V of order $n + 1$, a contradiction.

(c) Suppose T spans V and has order n . If T were LD we could find t_0, t_1, \dots, t_m in T distinct such that $t_0 = \sum_{i=1}^m \lambda_i t_i$ for some $\lambda_i \in \mathbf{F}$. Thus $V = \langle T \rangle = \langle T \setminus \{t_0\} \rangle$ so $T \setminus \{t_0\}$ is a spanning set for V of order $n - 1$, a contradiction.

(d) Let $T = \{t_1, \dots, t_m\}$ be a linearly independent subset of V . Since S spans V we can find s_1, \dots, s_m in S such that $(S \setminus \{s_1, \dots, s_m\}) \cup T$ spans V . Since this set has order (at most) n it is a basis containing T .

(e) Suppose that T is a finite spanning set for V and let $T' \subset T$ be a subset of minimal size that still spans V . If $|T'| = n$ we're done by (c). Otherwise $|T'| > n$ and so T' is LD as S spans. Thus there are t_0, \dots, t_m in T' distinct such that $t_0 = \sum \lambda_i t_i$ for some $\lambda_i \in \mathbf{F}$. Then $V = \langle T' \rangle = \langle T' \setminus \{t_0\} \rangle$ contradicting the minimality of T' \square

Exercise. Prove (e) holds for any spanning set in a f.d. V .

We prove the theorem by replacing elements of T by elements of S one by one.

Proof of the Steinitz exchange lemma. Suppose that we've already found a subset D_r of T of order $0 \leq r < n$ such that $T_r := (T \setminus D_r) \cup \{e_1, \dots, e_r\}$ spans V (the case $r = 0$ is clear and the case $r = n$ is the result). Then we can write

$$e_{r+1} = \sum_{i=1}^k \lambda_i t_i$$

with $\lambda_i \in \mathbf{F}$ and $t_i \in T_r$. Since $\{e_1, \dots, e_{r+1}\}$ is linearly independent there must be some $1 \leq j \leq k$ such that $\lambda_j \neq 0$ and $t_j \notin \{e_1, \dots, e_r\}$. Let $D_{r+1} = D_r \cup \{t_j\}$ and

$$T_{r+1} = (T \setminus D_{r+1}) \cup \{e_1, \dots, e_{r+1}\} = (T_r \setminus \{t_j\}) \cup \{e_{r+1}\}$$

Now

$$t_j = \frac{1}{\lambda_j} e_{r+1} - \sum_{i \neq j} \frac{\lambda_i}{\lambda_j} t_i,$$

so $t_j \in \langle T_{r+1} \rangle$ and $\langle T_{r+1} \rangle = \langle T_{r+1} \cup \{t_j\} \rangle \supset \langle T_r \rangle = V$.

Now we can inductively construct $D = D_n$ with the required properties. \square

Definition. If a vector space V over \mathbf{F} is finite-dimensional with basis S , we define the *dimension* of V by

$$\dim_{\mathbf{F}} V = \dim V = |S|.$$

Remarks.

- (1) By the last corollary the dimension of a finite dimensional space V does not depend on the choice of basis S . However the dimension *does* depend on \mathbf{F} . For example \mathbf{C} has dimension 1 viewed as a vector space over \mathbf{C} (since $\{1\}$ is a basis) but dimension 2 viewed as a vector space over \mathbf{R} (since $\{1, i\}$ is a basis).
- (2) If we wanted to be more precise then we could define the dimension of an infinite-dimensional space to be the cardinality of any basis for V . We have not proven enough to see that this would be well-defined; in fact there are no problems.

Lemma. *If V is f.d. and $U \subset V$ is a subspace then U is also f.d. and*

$$\dim U \leq \dim V.$$

Proof. Let $S \subset U$ be a LI subset of U of maximal possible size. Then every finite subset of S has size at most $\dim V$ (by the Steinitz Exchange Lemma). Thus $|S| \leq \dim V$.

If $u \in U \setminus \langle S \rangle$ then $S \cup \{u\}$ is LI contradicting the maximality of S . Thus $U = \langle S \rangle$ and S is a basis for U . \square

Proposition. *If V is a finite dimensional vector space over \mathbf{F} and U is a subspace then*

$$\dim V = \dim U + \dim V/U.$$

Proof. Since dimension is defined in terms of bases, and we have no way to compute it at this stage of the course except by finding bases and counting the number of elements, we must find suitable bases. The key idea is to be careful about how we choose our bases.

Slogan When choosing bases always choose the best basis for the job.

Let $\{u_1, \dots, u_m\}$ be a basis for U and extend to a basis $\{u_1, \dots, u_m, v_{m+1}, \dots, v_n\}$ for V . It suffices to show that $S := \{v_{m+1} + U, \dots, v_n + U\}$ is a basis for V/U . Suppose that $v + U \in V/U$. Then we can write

$$v = \sum_{i=1}^m \lambda_i u_i + \sum_{j=m+1}^n \mu_j v_j.$$

Thus

$$v + U = \sum_{i=1}^m \lambda_i (u_i + U) + \sum_{j=m+1}^n \mu_j (v_j + U) = 0 + \sum_{j=m+1}^n \mu_j (v_j + U)$$

and so S spans V/U . To show that S is LI, suppose that

$$\sum_{j=m+1}^n \mu_j (v_j + U) = 0.$$

Then $\sum_{j=m+1}^n \mu_j v_j \in U$ so we can write $\sum_{j=m+1}^n \mu_j v_j = \sum_{i=1}^m \lambda_i u_i$ for some $\lambda_i \in \mathbf{F}$. Since the set $\{u_1, \dots, u_m, v_{m+1}, \dots, v_n\}$ is LI we deduce that each μ_j (and λ_i) is zero as required. \square

Corollary. *If U is a proper subspace of V then $\dim U < \dim V$.*

Proof. Since U is proper, V/U is non-zero and so the empty set does not span V/U . Thus $\dim V/U > 0$ and $\dim U = \dim V - \dim V/U < \dim V$. \square

Exercise. Prove this last corollary directly by strengthening the proof that $U \leq V$ implies $\dim U \leq \dim V$.

LECTURE 4

1.3. Direct sum. There are two related notions of direct sum of vector spaces and the distinction between them can often cause confusion to newcomers to the subject. The first is sometimes known as the *internal* direct sum and the latter as the *external* direct sum. However it is common to gloss over the difference between them.

Definition. Suppose that V is a vector space over \mathbf{F} and U and W are subspaces of V . Recall that the *sum* of U and W is defined to be

$$U + W = \{u + w : u \in U, w \in W\}.$$

We say that V is the (*internal*) *direct sum* of U and W , written $V = U \oplus W$, if $V = U + W$ and $U \cap W = 0$. Equivalently $V = U \oplus W$ if every element $v \in V$ can be written uniquely as $u + w$ with $u \in U$ and $w \in W$.

We also say that U and W are *complementary subspaces* in V .

Example. Suppose that $V = \mathbf{R}^3$ and

$$U = \left\langle \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} : x_1 + x_2 + x_3 = 0 \right\rangle, W_1 = \left\langle \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\rangle \text{ and } W_2 = \left\langle \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \right\rangle$$

then $V = U \oplus W_1 = U \oplus W_2$.

Note in particular that U does not have only one complementary subspace in V .

Definition. Given any two vector spaces U and W over \mathbf{F} the (*external*) *direct sum* $U \oplus W$ of U and W is defined to be the set of pairs

$$\{(u, w) : u \in U, w \in W\}$$

with addition given by

$$(u_1, w_1) + (u_2, w_2) = (u_1 + u_2, w_1 + w_2)$$

and scalar multiplication given by

$$\lambda(u, w) = (\lambda u, \lambda w).$$

Exercise. Show that $U \oplus W$ is a vector space over \mathbf{F} with the given operations and that it is the internal direct sum of its subspaces

$$\{(u, 0) : u \in U\} \text{ and } \{(0, w) : w \in W\}.$$

More generally we can make the following definitions.

Definition. If U_1, \dots, U_n are subspaces of V then V is the (*internal*) *direct sum* of U_1, \dots, U_n written

$$V = U_1 \oplus \dots \oplus U_n = \bigoplus_{i=1}^n U_i$$

if every element v of V can be written uniquely as $v = \sum_{i=1}^n u_i$ with $u_i \in U_i$.

Definition. If U_1, \dots, U_n are any vector spaces over \mathbf{F} their (*external*) *direct sum* is the vector space

$$\bigoplus_{i=1}^n U_i := \{(u_1, \dots, u_n) \mid u_i \in U_i\}$$

with natural coordinate-wise operations.

From now on we will drop the adjectives ‘internal’ and ‘external’ from ‘direct sum’.

2. LINEAR MAPS

2.1. Definitions and examples.

Definition. Suppose that U and V are vector spaces over a field \mathbf{F} . Then a function $\alpha: U \rightarrow V$ is a *linear map* if

- (i) $\alpha(u_1 + u_2) = \alpha(u_1) + \alpha(u_2)$ for all $u_1, u_2 \in U$;
- (ii) $\alpha(\lambda u) = \lambda \alpha(u)$ for all $u \in U$ and $\lambda \in \mathbf{F}$.

Notation. We write $\mathcal{L}(U, V)$ for the set of linear maps $U \rightarrow V$.

Remarks.

- (1) We can combine the two parts of the definition into one as: α is linear if and only if $\alpha(\lambda u_1 + \mu u_2) = \lambda \alpha(u_1) + \mu \alpha(u_2)$ for all $\lambda, \mu \in \mathbf{F}$ and $u_1, u_2 \in U$. Linear maps should be viewed as functions between vector spaces that respect their structure as vector spaces.
- (2) If α is a linear map then α is a homomorphism of the underlying abelian groups. In particular $\alpha(0) = 0$.
- (3) If we want to stress the field \mathbf{F} then we will say a map is \mathbf{F} -linear. For example, complex conjugation defines an \mathbf{R} -linear map from \mathbf{C} to \mathbf{C} but it is not \mathbf{C} -linear.

Examples.

- (1) Let A be an $n \times m$ matrix with coefficients in \mathbf{F} — write $A \in M_{n,m}(\mathbf{F})$. Then $\alpha: \mathbf{F}^m \rightarrow \mathbf{F}^n$; $\alpha(v) = Av$ is a linear map.

To see this let $\lambda, \mu \in \mathbf{F}$ and $u, v \in \mathbf{F}^m$. As usual, let A_{ij} denote the ij th entry of A and u_j , (resp. v_j) for the j th coordinate of u (resp. v). Then for $1 \leq i \leq n$,

$$(\alpha(\lambda u + \mu v))_i = \sum_{j=1}^m A_{ij}(\lambda u_j + \mu v_j) = \lambda \alpha(u)_i + \mu \alpha(v)_i$$

so $\alpha(\lambda u + \mu v) = \lambda \alpha(u) + \mu \alpha(v)$ as required.

- (2) If X is any set and $g \in \mathbf{F}^X$ then $m_g: \mathbf{F}^X \rightarrow \mathbf{F}^X$; $m_g(f)(x) := g(x)f(x)$ for $x \in X$ is linear.
- (3) For all $x \in [a, b]$, $\delta_x: C([a, b], \mathbf{R}) \rightarrow \mathbf{R}$; $f \mapsto f(x)$ is linear.
- (4) $I: C([a, b], \mathbf{R}) \rightarrow C([a, b], \mathbf{R})$; $I(f)(x) = \int_a^x f(t) dt$ is linear.
- (5) $D: C^\infty([a, b], \mathbf{R}) \rightarrow C^\infty([a, b], \mathbf{R})$; $(Df)(t) = f'(t)$ is linear.
- (6) If $\alpha, \beta: U \rightarrow V$ are linear and $\lambda \in \mathbf{F}$ then $\alpha + \beta: U \rightarrow V$ given by $(\alpha + \beta)(u) = \alpha(u) + \beta(u)$ and $\lambda\alpha: U \rightarrow V$ given by $(\lambda\alpha)(u) = \lambda(\alpha(u))$ are linear. In this way $\mathcal{L}(U, V)$ is a vector space over \mathbf{F} .

Definition. We say that a linear map $\alpha: U \rightarrow V$ is an *isomorphism* if there is a linear map $\beta: V \rightarrow U$ such that $\beta\alpha = \text{id}_U$ and $\alpha\beta = \text{id}_V$.

LECTURE 5

Lemma. Suppose that U and V are vector spaces over \mathbf{F} . A linear map $\alpha: U \rightarrow V$ is an isomorphism if and only if α is a bijection.

Proof. Certainly an isomorphism $\alpha: U \rightarrow V$ is a bijection since it has an inverse as a function between the underlying sets U and V . Suppose that $\alpha: U \rightarrow V$ is a linear bijection and let $\beta: V \rightarrow U$ be its inverse as a function. We must show that β is also linear. Let $\lambda, \mu \in \mathbf{F}$ and $v_1, v_2 \in V$. Then

$$\alpha\beta(\lambda v_1 + \mu v_2) = \lambda\alpha\beta(v_1) + \mu\alpha\beta(v_2) = \alpha(\lambda\beta(v_1) + \mu\beta(v_2)).$$

Since α is injective it follows that β is linear as required. \square

Proposition. Suppose that $\alpha: U \rightarrow V$ is an \mathbf{F} -linear map.

- (a) If α is injective and $S \subset U$ is linearly independent then $\alpha(S) \subset V$ is linearly independent.
- (b) If α is surjective and $S \subset U$ spans U then $\alpha(S)$ spans V .
- (c) If α is an isomorphism and S is a basis then $\alpha(S)$ is a basis.

Proof. (a) Suppose α is injective, $S \subset U$ and $\alpha(S)$ is linearly dependent. Then there are $s_0, \dots, s_n \in S$ distinct and $\lambda_1, \dots, \lambda_n \in \mathbf{F}$ such that

$$\alpha(s_0) = \sum \lambda_i \alpha(s_i) = \alpha\left(\sum_{i=1}^n \lambda_i s_i\right).$$

Since α is injective it follows that $s_0 = \sum_{i=1}^n \lambda_i s_i$ and S is LD.

(b) Now suppose that α is surjective, $S \subset U$ spans U and let v in V . There is $u \in U$ such that $\alpha(u) = v$ and there are $s_1, \dots, s_n \in S$ and $\lambda_1, \dots, \lambda_n \in \mathbf{F}$ such that $\sum \lambda_i s_i = u$. Then $\sum \lambda_i \alpha(s_i) = v$. Thus $\alpha(S)$ spans V .

(c) Follows immediately from (a) and (b). \square

Corollary. *If two finite dimensional vector spaces are isomorphic then they have the same dimension.*

Proof. If $\alpha: U \rightarrow V$ is an isomorphism and S is a finite basis for U then $\alpha(S)$ is a basis of V by the proposition. Since α is an injection $|S| = |\alpha(S)|$. \square

Proposition. *Suppose that V is a vector space over \mathbf{F} of dimension $n < \infty$. Writing e_1, \dots, e_n for the standard basis for \mathbf{F}^n , there is a bijection Φ between the set of isomorphisms $\mathbf{F}^n \rightarrow V$ and the set of (ordered) bases for V that sends the isomorphism $\alpha: \mathbf{F}^n \rightarrow V$ to the (ordered) basis $(\alpha(e_1), \dots, \alpha(e_n))$.*

Proof. That the map Φ is well-defined follows immediately from part (c) of the last Proposition.

If $\Phi(\alpha) = \Phi(\beta)$ then

$$\alpha \left(\begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \right) = \sum_{i=1}^n x_i \alpha(e_i) = \sum_{i=1}^n x_i \beta(e_i) = \beta \left(\begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \right)$$

for all $\begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \in \mathbf{F}^n$ so $\alpha = \beta$. Thus Φ is injective.

Suppose now that (v_1, \dots, v_n) is an ordered basis for V and define $\alpha: \mathbf{F}^n \rightarrow V$ by

$$\alpha \left(\begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \right) = \sum_{i=1}^n x_i v_i.$$

Then α is injective since v_1, \dots, v_n are LI and α is surjective since v_1, \dots, v_n span V and α is easily seen to be linear. Thus α is an isomorphism such that $\Phi(\alpha) = (v_1, \dots, v_n)$ and Φ is surjective as required. \square

Thus choosing a basis for an n -dimensional vector space V corresponds to choosing an identification of V with \mathbf{F}^n .

2.2. Linear maps and matrices.

Proposition. *Suppose that U and V are vector spaces over \mathbf{F} and $S := \{e_1, \dots, e_n\}$ is a basis for U . Then every function $f: S \rightarrow V$ extends uniquely to a linear map $\alpha: U \rightarrow V$.*

Slogan To define a linear map it suffices to specify its values on a basis.

Proof. First we prove uniqueness: suppose that $f: S \rightarrow V$ and α and β are two linear maps $U \rightarrow V$ extending f . Let $u \in U$ so that $u = \sum u_i e_i$ for some $u_i \in \mathbf{F}$. Then

$$\alpha(u) = \alpha \left(\sum_{i=1}^n u_i e_i \right) = \sum_{i=1}^n u_i \alpha(e_i).$$

Similarly, $\beta(u) = \sum_{i=1}^n u_i \beta(e_i)$. Since $\alpha(e_i) = f(e_i) = \beta(e_i)$ for each $1 \leq i \leq n$ we see that $\alpha(u) = \beta(u)$ for all $u \in U$ and so $\alpha = \beta$.

That argument also shows us how to construct a linear map α that extends f . Every $u \in U$ can be written uniquely as $u = \sum_{i=1}^n u_i e_i$ with $u_i \in \mathbf{F}$. Thus we can

define $\alpha(u) = \sum u_i f(e_i)$ without ambiguity. Certainly α extends f so it remains to show that α is linear. So we compute for $u = \sum u_i e_i$ and $v = \sum v_i e_i$,

$$\begin{aligned} \alpha(\lambda u + \mu v) &= \alpha\left(\sum_{i=1}^n (\lambda u_i + \mu v_i) e_i\right) \\ &= \sum_{i=1}^n (\lambda u_i + \mu v_i) f(e_i) \\ &= \lambda \sum_{i=1}^n u_i f(e_i) + \mu \sum_{i=1}^n v_i f(e_i) \\ &= \lambda \alpha(u) + \mu \alpha(v) \end{aligned}$$

as required. □

Remarks.

- (1) With a little care the proof of the proposition can be extended to the case U is not assumed finite dimensional.
- (2) It is not hard to see that the only subsets S of U that satisfy the conclusions of the proposition are bases: spanning is necessary for the uniqueness part and linear independence is necessary for the existence part. The proposition should be considered a key motivation for the definition of a basis.

Corollary. *If U and V are finite dimensional vector spaces over \mathbf{F} with (ordered) bases (e_1, \dots, e_m) and (f_1, \dots, f_n) respectively then there is a bijection*

$$\text{Mat}_{n,m}(\mathbf{F}) \leftrightarrow \mathcal{L}(U, V)$$

that sends a matrix A to the unique linear map α such that $\alpha(e_i) = \sum A_{ji} f_j$.

Interpretation The i th column of the matrix A tells where the i th basis vector of U goes (as a linear combination of the basis vectors of V).

Proof. If $\alpha: U \rightarrow V$ is a linear map then for each $1 \leq i \leq m$ we can write $\alpha(e_i)$ uniquely as $\alpha(e_i) = \sum a_{ji} f_j$ with $a_{ji} \in \mathbf{F}$. The proposition tells us that every matrix $A = (a_{ij})$ arises in this way from some linear map and that α is determined by A . □

Exercise. Show that the bijection given by the corollary is even an isomorphism of vector spaces. By finding a basis for $\text{Mat}_{n,m}(\mathbf{F})$, deduce that $\dim \mathcal{L}(U, V) = \dim U \dim V$. Show moreover that if U_1, \dots, U_n are vector spaces then

$$\mathcal{L}(V, \bigoplus_{i=1}^n U_i) \cong \bigoplus_{i=1}^n \mathcal{L}(V, U_i)$$

and

$$\mathcal{L}\left(\bigoplus_{i=1}^n U_i, V\right) \cong \bigoplus_{i=1}^n \mathcal{L}(U_i, V).$$

LECTURE 6

Definition. If $\alpha \in \mathcal{L}(U, V)$, (e_1, \dots, e_m) is a basis for U and (f_1, \dots, f_n) is a basis for V then we call the matrix A such that $\alpha(e_i) = \sum_{j=1}^n A_{ji} f_j$ the *matrix representing α with respect to (e_1, \dots, e_m) and (f_1, \dots, f_n)* .

Proposition. Suppose that U, V and W are finite dimensional vector spaces over \mathbf{F} with bases $R := (u_1, \dots, u_r)$, $S := (v_1, \dots, v_s)$ and $T := (w_1, \dots, w_t)$ respectively. If $\alpha: U \rightarrow V$ is a linear map represented by the matrix A with respect to R and S and $\beta: V \rightarrow W$ is a linear map represented by the matrix B with respect to S and T then $\beta\alpha$ is the linear map $U \rightarrow W$ represented by BA with respect to R and T .

Proof. Verifying that $\beta\alpha$ is linear is straightforward: suppose $x, y \in U$ and $\lambda, \mu \in \mathbf{F}$ then

$$\beta\alpha(\lambda x + \mu y) = \beta(\lambda\alpha(x) + \mu\alpha(y)) = \lambda\beta\alpha(x) + \mu\beta\alpha(y).$$

Next we compute $\beta\alpha(u_i)$ as a linear combination of w_j .

$$\beta\alpha(u_i) = \beta\left(\sum_k A_{ki} v_k\right) = \sum_k A_{ki} \beta(v_k) = \sum_{k,j} A_{ki} B_{jk} w_j = \sum_j (BA)_{ji} w_j$$

as required. □

2.3. The first isomorphism theorem and the rank-nullity theorem.

Definition. Suppose that $\alpha: U \rightarrow V$ is a linear map.

- The *image* of α , $\text{Im } \alpha := \{\alpha(u) : u \in U\}$.
- The *kernel* of α , $\ker \alpha := \{u \in U : \alpha(u) = 0\}$.

Note that α is injective if and only if $\ker \alpha = 0$ and that α is surjective if and only if $\text{Im } \alpha = V$.

Examples.

- (1) Let $A \in M_{n,m}(\mathbf{F})$ and let $\alpha: \mathbf{F}^m \rightarrow \mathbf{F}^n$ be the linear map defined by $x \mapsto Ax$. Then the system of equations

$$\sum_{j=1}^m A_{ij} x_j = b_i; \quad 1 \leq i \leq n$$

has a solution if and only if $\begin{pmatrix} b_1 \\ \vdots \\ b_n \end{pmatrix} \in \text{Im } \alpha$. The kernel of α consists of the

solutions $\begin{pmatrix} x_1 \\ \vdots \\ x_m \end{pmatrix}$ to the homogeneous equations

$$\sum_{j=1}^m A_{ij} x_j = 0; \quad 1 \leq i \leq n$$

- (2) Let $\beta: C^\infty(\mathbf{R}, \mathbf{R}) \rightarrow C^\infty(\mathbf{R}, \mathbf{R})$ be given by

$$\beta(f)(t) = f''(t) + p(t)f'(t) + q(t)f(t)$$

for some $p, q \in C^\infty(\mathbf{R}, \mathbf{R})$. A function $g \in C^\infty(\mathbf{R}, \mathbf{R})$ is in the image of β precisely if

$$f''(t) + p(t)f'(t) + q(t)f(t) = g(t)$$

has a solution in $C^\infty(\mathbf{R}, \mathbf{R})$. Moreover, $\ker \beta$ consists of the solutions to the differential equation

$$f''(t) + p(t)f'(t) + q(t)f(t) = 0$$

in $C^\infty(\mathbf{R}, \mathbf{R})$.

The following analogue of the first isomorphism theorem for groups holds for vector spaces.

Theorem (The first isomorphism theorem). *Let $\alpha: U \rightarrow V$ be a linear map between vector spaces over \mathbf{F} . Then $\ker \alpha$ is a subspace of U and $\text{Im } \alpha$ is a subspace of V . Moreover α induces an isomorphism $U/\ker \alpha \rightarrow \text{Im } \alpha$ given by*

$$\bar{\alpha}(u + \ker \alpha) = \alpha(u).$$

Proof. Since $\alpha(0) = 0$, $0 \in \ker \alpha$. Suppose that $u_1, u_2 \in \ker \alpha$ and $\lambda, \mu \in \mathbf{F}$. Then

$$\alpha(\lambda u_1 + \mu u_2) = \lambda \alpha(u_1) + \mu \alpha(u_2) = 0 + 0 = 0.$$

Thus $\ker \alpha \leq U$. Similarly $0 \in \text{Im } \alpha$ and for $u_1, u_2 \in U$,

$$\lambda \alpha(u_1) + \mu \alpha(u_2) = \alpha(\lambda u_1 + \mu u_2) \in \text{Im } \alpha$$

so $\text{Im } \alpha \leq V$ and $\bar{\alpha}$ is linear if it is well-defined.

But if $u + \ker \alpha = u' + \ker \alpha \in U/\ker \alpha$ then $u - u' \in \ker \alpha$ so $\alpha(u) = \alpha(u')$ and $\bar{\alpha}$ is well-defined. So it remains to verify that $\bar{\alpha}$ is a bijection.

If $\bar{\alpha}(u + \ker \alpha) = 0$ then $\alpha(u) = 0$ so $u \in \ker \alpha$ and $u + \ker \alpha = 0$. So $\bar{\alpha}$ is injective. That $\bar{\alpha}$ is surjective is clear. \square

Definition. Suppose that $\alpha: U \rightarrow V$ is a linear map between finite dimensional vector spaces.

- The number $n(\alpha) := \dim \ker \alpha$ is called the *nullity* of α .
- The number $r(\alpha) := \dim \text{Im } \alpha$ is called the *rank* of α .

Corollary (The rank-nullity theorem). *If $\alpha: U \rightarrow V$ is a linear map between f.d. vector spaces over \mathbf{F} then*

$$r(\alpha) + n(\alpha) = \dim U.$$

Proof. Since $U/\ker \alpha \cong \text{Im } \alpha$ they have the same dimension. But

$$\dim U = \dim(U/\ker \alpha) + \dim \ker \alpha$$

by an earlier computation so $\dim U = r(\alpha) + n(\alpha)$ as required. \square

We are about to give another more direct proof of the rank-nullity theorem. This one looks slick but we've hidden the work in the proof that if $U \leq V$ then $\dim V = \dim U + \dim V/U$.

Exercise. Deduce that $\dim V = \dim U + \dim V/U$ from the rank-nullity theorem.

Proposition. *Suppose that $\alpha: U \rightarrow V$ is a linear map between finite dimensional vector spaces then there are bases (e_1, \dots, e_n) for U and (f_1, \dots, f_m) for V such that the matrix representing α is*

$$\begin{pmatrix} I_r & 0 \\ 0 & 0 \end{pmatrix}$$

where $r = r(\alpha)$.

In particular $r(\alpha) + n(\alpha) = \dim U$.

Proof. Let (e_{k+1}, \dots, e_n) be a basis for $\ker \alpha$ (here $n(\alpha) = n - k$) and extend it to a basis (e_1, \dots, e_n) for U (we're being careful about ordering now so that we don't have to change it later). Let $f_i = \alpha(e_i)$ for $1 \leq i \leq k$.

We claim that (f_1, \dots, f_k) form a basis for $\text{Im } \alpha$ (so that $k = r(\alpha)$).

Suppose first that $\sum_{i=1}^k \lambda_i f_i = 0$ for some $\lambda_i \in \mathbf{F}$. Then $\alpha\left(\sum_{i=1}^k \lambda_i e_i\right) = 0$ and so $\sum_{i=1}^k \lambda_i e_i \in \ker \alpha$. But $\ker \alpha \cap \langle e_1, \dots, e_k \rangle = 0$ by construction and so $\sum_{i=1}^k \lambda_i e_i = 0$. Since e_1, \dots, e_k are LI, each $\lambda_i = 0$. Thus we have shown that $\{f_1, \dots, f_k\}$ is LI.

Now suppose that $v \in \text{Im } \alpha$, so that $v = \alpha(\sum_{i=1}^n \mu_i e_i)$ for some $\mu_i \in \mathbf{F}$. Since $\alpha(e_i) = 0$ for $i > k$ and $\alpha(e_i) = f_i$ for $i \leq k$, $v = \sum_{i=1}^k \mu_i f_i \in \langle f_1, \dots, f_k \rangle$. So (f_1, \dots, f_k) is a basis for $\text{Im } \alpha$ as claimed (and $k = r$).

We can extend $\{f_1, \dots, f_r\}$ to a basis $\{f_1, \dots, f_m\}$ for V .

Now

$$\alpha(e_i) = \begin{cases} f_i & 1 \leq i \leq r \\ 0 & r + 1 \leq i \leq m \end{cases}$$

so the matrix representing α with respect to our choice of basis is as in the statement. \square

The proposition says that the rank of a linear map between two finite dimensional vector spaces (together with the dimensions of the spaces) is its only basis-independent invariant (or more precisely any other invariant can be deduced from it).

LECTURE 7

The rank-nullity theorem is very useful for computing dimensions of vector spaces in terms of known dimensions of other spaces.

Examples.

(1) Let $W = \{\mathbf{x} \in \mathbf{R}^5 \mid x_1 + x_2 + x_5 = 0 \text{ and } x_3 - x_4 - x_5 = 0\}$. What is $\dim W$?

Consider $\alpha: \mathbf{R}^5 \rightarrow \mathbf{R}^2$ given by $\alpha(\mathbf{x}) = \begin{pmatrix} x_1 + x_2 + x_5 \\ x_3 - x_4 - x_5 \end{pmatrix}$. Then α is a linear map with image \mathbf{R}^2 (since

$$\alpha \left(\begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \right) = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ and } \alpha \left(\begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} \right) = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.)$$

and $\ker \alpha = W$. Thus $\dim W = n(\alpha) = 5 - r(\alpha) = 5 - 2 = 3$.

More generally, one can use the rank-nullity theorem to see that m linear equations in n unknowns have a space of solutions of dimension at least $n - m$.

- (2) Suppose that U and W are subspaces of a finite dimensional vector space V then let $\alpha: U \oplus W \rightarrow V$ be the linear map given by $\alpha((u, w)) = u + w$. Then $\ker \alpha = \{(u, -u) \mid u \in U \cap W\} \cong U \cap W$, and $\text{Im } \alpha = U + W$. Thus

$$\dim(U) + \dim(W) = \dim U \oplus W = \dim(U + W) + \dim(U \cap W).$$

Corollary (of the rank-nullity theorem). *Suppose that $\alpha: U \rightarrow V$ is a linear map between two vector spaces of dimension $n < \infty$. Then the following are equivalent:*

- (a) α is injective;
- (b) α is surjective;
- (c) α is an isomorphism.

Proof. It suffices to see that (a) is equivalent to (b) since these two together are already known to be equivalent to (c). Now α is injective if and only if $n(\alpha) = 0$. By the rank-nullity theorem $n(\alpha) = 0$ if and only if $r(\alpha) = n$ and the latter is equivalent to α being surjective. \square

This enables us to prove the following fact about matrices.

Lemma. *Let A be an $n \times n$ matrix over \mathbf{F} . The following are equivalent*

- (a) *there is a matrix B such that $BA = I_n$;*
- (b) *there is a matrix C such that $AC = I_n$.*

Moreover, if (a) and (b) hold then $B = C$ and we write $A^{-1} = B = C$; we say A is invertible.

Proof. Let $\alpha, \beta, \gamma, \iota: \mathbf{F}^n \rightarrow \mathbf{F}^n$ be the linear maps represented by A, B, C and I_n respectively (with respect to the standard basis for \mathbf{F}^n).

Note first that (a) holds if and only if there exists β such that $\beta\alpha = \iota$. This last implies that α is injective which in turn implies that α is an isomorphism by the previous result. Conversely if α is an isomorphism there does exist β such that $\beta\alpha = \iota$ by definition. Thus (a) holds if and only if α is an isomorphism.

Similarly (b) holds if and only if there exists γ such that $\alpha\gamma = \iota$. This last implies that α is surjective and so an isomorphism by the previous result. Thus (b) also holds if and only if α is an isomorphism.

If α is an isomorphism then β and γ must both be the set-theoretic inverse of α and so $B = C$ as claimed. \square

2.4. Change of basis.

Theorem. *Suppose that (e_1, \dots, e_m) and (u_1, \dots, u_m) are two bases for a vector space U over \mathbf{F} and (f_1, \dots, f_n) and (v_1, \dots, v_n) are two bases of another vector space V . Let $\alpha: U \rightarrow V$ be a linear map, A be the matrix representing α with respect to (e_1, \dots, e_m) and (f_1, \dots, f_n) and B be the matrix representing α with respect to (u_1, \dots, u_m) and (v_1, \dots, v_n) then*

$$B = Q^{-1}AP$$

where $u_i = \sum P_{ki}e_k$ for $i = 1, \dots, m$ and $v_j = \sum Q_{lj}f_l$ for $j = 1, \dots, n$.

Note that one can view P as the matrix representing the identity map from U with basis (u_1, \dots, u_m) to U with basis (e_1, \dots, e_m) and Q as the matrix representing the identity map from V with basis (v_1, \dots, v_n) to V with basis (f_1, \dots, f_n) . Thus both are invertible with inverses represented by the identity maps going in the opposite directions.

Proof. On the one hand, by definition

$$\alpha(u_i) = \sum_j B_{ji}v_j = \sum_{j,l} B_{ji}Q_{lj}f_l = \sum_l (QB)_{li}f_l.$$

On the other hand, also by definition

$$\alpha(u_i) = \alpha\left(\sum_k P_{ki}e_k\right) = \sum_{k,l} P_{ki}A_{lk}f_l = \sum_l (AP)_{li}f_l.$$

Thus $QB = AP$ as the f_l are LI. Since Q is invertible the result follows. \square

Definition. We say two matrices $A, B \in \text{Mat}_{n,m}(\mathbf{F})$ are *equivalent* if there are invertible matrices $P \in \text{Mat}_m(\mathbf{F})$ and $Q \in \text{Mat}_n(\mathbf{F})$ such that $Q^{-1}AP = B$.

Note that equivalence is an equivalence relation. It can be reinterpreted as follows: two matrices are equivalent precisely if they represent the same linear map with respect to different bases.

We saw earlier that for every linear map α between f.d. vector spaces there are bases for the domain and codomain such that α is represented by a matrix of the form

$$\begin{pmatrix} I_r & 0 \\ 0 & 0 \end{pmatrix}.$$

Moreover $r = r(\alpha)$ is independent of the choice of bases. We can now rephrase this as follows.

Corollary. If $A \in \text{Mat}_{n,m}(\mathbf{F})$ there are invertible matrices $P \in \text{Mat}_m(\mathbf{F})$ and $Q \in \text{Mat}_n(\mathbf{F})$ such that $Q^{-1}AP$ is of the form

$$\begin{pmatrix} I_r & 0 \\ 0 & 0 \end{pmatrix}.$$

Moreover r is uniquely determined by A . i.e. every equivalence class contains precisely one matrix of this form. \square

Definition. If $A \in \text{Mat}_{n,m}(\mathbf{F})$ then

- *column rank* of A , written $r(A)$ is the dimension of the subspace of \mathbf{F}^n spanned by the columns of A ;
- the *row rank* of A is the column rank of A^T .

Note that if we take α to be a linear map represented by A with respect to the standard bases of \mathbf{F}^m and \mathbf{F}^n then $r(A) = r(\alpha)$. i.e. ‘column rank=rank’. Moreover, since $r(\alpha)$ is defined in a basis-invariant way, the column rank of A is constant on equivalence classes.

Corollary (Row rank equals column rank). If $A \in \text{Mat}_{m,n}(\mathbf{F})$ then $r(A) = r(A^T)$.

Pure matrix proof of the Proposition. We claim that there are elementary matrices E_1^n, \dots, E_a^n and F_1^m, \dots, F_b^m such that $E_a^n \cdots E_1^n A F_1^m \cdots F_b^m$ is of the required form. This suffices since all the elementary matrices are invertible and products of invertible matrices are invertible.

Moreover, to prove the claim it suffices to show that there is a sequence of elementary row and column operations that reduces A to the required form.

If $A = 0$ there is nothing to do. Otherwise, we can find a pair i, j such that $A_{ij} \neq 0$. By swapping rows 1 and i and then swapping columns 1 and j we can reduce to the case that $A_{11} \neq 0$. By multiplying row 1 by $\frac{1}{A_{11}}$ we can further assume that $A_{11} = 1$.

Now, given $A_{11} = 1$ we can add $-A_{1j}$ times column 1 to column j for each $1 < j \leq m$ and then add $-A_{i1}$ times row 1 to row i for each $1 < i \leq n$ to reduce further to the case that A is of the form

$$\begin{pmatrix} 1 & 0 \\ 0 & B \end{pmatrix}.$$

Now by induction on the size of A we can find elementary row and column operations that reduces B to the required form. Applying these ‘same’ operations to A we complete the proof. \square

Note that the algorithm described in the proof can easily be implemented on a computer in order to actually compute the matrices P and Q .

Exercise. Show that elementary row and column operations do not alter $r(A)$ or $r(A^T)$. Conclude that the r in the statement of the proposition is thus equal to $r(A)$ and to $r(A^T)$.

3. DUALITY

3.1. Dual spaces. To specify a subspace of \mathbf{F}^n we can write down a set of linear equations that every vector in the space satisfies. For example if $U = \left\langle \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix} \right\rangle \subset \mathbf{F}^3$

we can see that

$$U = \left\{ \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} : 2x_1 - x_2 = 0, x_1 - x_3 = 0 \right\}.$$

These equations are determined by linear maps $\mathbf{F}^n \rightarrow \mathbf{F}$. Moreover if $\theta_1, \theta_2: \mathbf{F}^n \rightarrow \mathbf{F}$ are linear maps that vanish on U and $\lambda, \mu \in \mathbf{F}$ then $\lambda\theta_1 + \mu\theta_2$ vanishes on U . Since the 0 map vanishes on every subspace, one may study the subspace of linear maps $\mathbf{F}^n \rightarrow \mathbf{F}$ that vanish on U .

Definition. Let V be a vector space over \mathbf{F} . The *dual space* of V is the vector space

$$V^* := \mathcal{L}(V, \mathbf{F}) = \{\alpha: V \rightarrow \mathbf{F} \text{ linear}\}$$

with pointwise addition and scalar multiplication. The elements of V^* are sometimes called *linear forms* or *linear functionals* on V .

Examples.

$$(1) V = \mathbf{R}^3, \theta: V \rightarrow \mathbf{R}; \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \mapsto x_3 - x_1 \in V^*.$$

- (2) $V = \mathbf{F}^X$, $x \in X$ then $f \mapsto f(x) \in V^*$.
(3) $V = C([0, 1], \mathbf{R})$, then $V \rightarrow \mathbf{R}$; $f \mapsto \int_0^1 \sin(2n\pi t)f(t)dt \in V^*$.
(4) $\text{tr}: \text{Mat}_n(\mathbf{F}) \rightarrow \mathbf{F}$; $A \mapsto \sum_{i=1}^n A_{ii} \in \text{Mat}_n(\mathbf{F})^*$.

Lemma. Suppose that V is a f.d. vector space over \mathbf{F} with basis (e_1, \dots, e_n) . Then V^* has a basis $(\epsilon_1, \dots, \epsilon_n)$ such that $\epsilon_i(e_j) = \delta_{ij}$.

Definition. We call the basis $(\epsilon_1, \dots, \epsilon_n)$ the *dual basis* of V^* with respect to (e_1, \dots, e_n) .

Proof of Lemma. We know that to define a linear map it suffices to define it on a basis so there are unique elements $\epsilon_1, \dots, \epsilon_n$ such that $\epsilon_i(e_j) = \delta_{ij}$. We must show that they span and are LI.

Suppose that $\theta \in V^*$ is any linear map. Then let $\lambda_i = \theta(e_i) \in \mathbf{F}$. We claim that $\theta = \sum_{i=1}^n \lambda_i \epsilon_i$. It suffices to show that the two elements agree on the basis e_1, \dots, e_n of V . But $\sum_{i=1}^n \lambda_i \epsilon_i(e_j) = \lambda_j = \theta(e_j)$. So the claim is true that $\epsilon_1, \dots, \epsilon_n$ do span V^* .

Next, suppose that $\sum \mu_i \epsilon_i = 0 \in V^*$ for some $\mu_1, \dots, \mu_n \in \mathbf{F}$. Then $0 = \sum \mu_i \epsilon_i(e_j) = \mu_j$ for each $j = 1, \dots, n$. Thus $\epsilon_1, \dots, \epsilon_n$ are LI as claimed. \square

Remark. If we think of elements of V as column vectors with respect to some basis

$$\sum x_i e_i = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix},$$

then we can view elements of V^* as row vectors with respect to the dual basis

$$\sum a_i \epsilon_i = (a_1 \quad \cdots \quad a_n).$$

Then

$$\left(\sum a_i \epsilon_i \right) \left(\sum x_j e_j \right) = \sum a_i x_i = (a_1 \quad \cdots \quad a_n) \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$$

Corollary. If V is f.d. then $\dim V^* = \dim V$. \square

Definition. If $U \subset V$ then the *annihilator* of U , $U^\circ := \{\theta \in V^* \mid \theta(u) = 0 \quad \forall u \in U\} \subset V^*$.

Example. Consider \mathbf{R}^3 with standard basis (e_1, e_2, e_3) and $(\mathbf{R}^3)^*$ with dual basis $(\epsilon_1, \epsilon_2, \epsilon_3)$, $U = \langle e_1 + 2e_2 + e_3 \rangle \subset \mathbf{R}^3$ and $W = \langle \epsilon_1 - \epsilon_3, \epsilon_1 - 2\epsilon_2 \rangle \subset (\mathbf{R}^3)^*$. Then $U^\circ = W$.

Proposition. Suppose that V is f.d. over \mathbf{F} and $U \subset V$ is a subspace. Then

$$\dim U + \dim U^\circ = \dim V.$$

Proof 1. Let (e_1, \dots, e_k) be a basis for U extend to a basis (e_1, \dots, e_n) for V and consider the dual basis $(\epsilon_1, \dots, \epsilon_n)$ for V^* .

We claim that U° is spanned by $\epsilon_{k+1}, \dots, \epsilon_n$.

Certainly if $j > k$, then $\epsilon_j(e_i) = 0$ for each $1 \leq i \leq k$ and so $\epsilon_j \in U^\circ$. Suppose now that $\theta \in U^\circ$. We can write $\theta = \sum_{i=1}^n \lambda_i \epsilon_i$ with $\lambda_i \in \mathbf{F}$. Now,

$$0 = \theta(e_j) = \lambda_j \text{ for each } 1 \leq j \leq k.$$

So $\theta = \sum_{j=k+1}^n \lambda_j \epsilon_j$. Thus U° is the span of $\epsilon_{k+1}, \dots, \epsilon_n$ and

$$\dim U^\circ = n - k = \dim V - \dim U$$

as claimed. □

LECTURE 9

Proof 2. Consider the restriction map $V^* \rightarrow U^*$ given by $\theta \mapsto \theta|_U$. Since every linear map $U \rightarrow \mathbf{F}$ can be extended to a linear map $V \rightarrow \mathbf{F}$ this map is a linear surjection. Moreover its kernel is U° . Thus $\dim V^* = \dim U^* + \dim U^\circ$ by the rank-nullity theorem. The proposition follows from the statements $\dim U = \dim U^*$ and $\dim V = \dim V^*$. □

Exercise (Proof 3). Show that $(V/U)^* \cong U^\circ$ and deduce the result.

Of course all three of these proofs are really the same but presented with differing levels of sophistication.

Proposition. *Suppose that V is a f.d. vector space over \mathbf{F} with bases (e_1, \dots, e_n) and (f_1, \dots, f_n) , and that P is the change of basis matrix from (e_1, \dots, e_n) to (f_1, \dots, f_n) i.e. $f_i = \sum_{k=1}^n P_{ki} e_k$ for $1 \leq i \leq n$.*

Let $(\epsilon_1, \dots, \epsilon_n)$ and (η_1, \dots, η_n) be the corresponding dual bases so that

$$\epsilon_i(e_j) = \delta_{ij} = \eta_i(f_j) \text{ for } 1 \leq i, j \leq n.$$

Then the change of basis matrix from $(\epsilon_1, \dots, \epsilon_n)$ to (η_1, \dots, η_n) is given by $(P^{-1})^T$ i.e. $\epsilon_i = \sum_l P_{li}^T \eta_l$.

Proof. Let $Q = P^{-1}$. Then $e_j = \sum_k Q_{kj} f_k$, so we can compute

$$\left(\sum_l P_{il} \eta_l \right) (e_j) = \sum_{k,l} (P_{il} \eta_l) (Q_{kj} f_k) = \sum_{k,l} P_{il} \delta_{kl} Q_{kj} = \delta_{ij}.$$

Thus $\epsilon_i = \sum_l P_{il} \eta_l$ as claimed. □

3.2. Dual maps.

Definition. Let V and W be vector spaces over \mathbf{F} and suppose that $\alpha: V \rightarrow W$ is a linear map. The *dual map* to α is the map $\alpha^*: W^* \rightarrow V^*$ is given by $\theta \mapsto \theta\alpha$.

Note that $\theta\alpha$ is the composite of two linear maps and so is linear. Moreover, if $\lambda, \mu \in \mathbf{F}$ and $\theta_1, \theta_2 \in W^*$ and $v \in V$ then

$$\begin{aligned} \alpha^*(\lambda\theta_1 + \mu\theta_2)(v) &= (\lambda\theta_1 + \mu\theta_2)\alpha(v) \\ &= \lambda\theta_1\alpha(v) + \mu\theta_2\alpha(v) \\ &= (\lambda\alpha^*(\theta_1) + \mu\alpha^*(\theta_2))(v). \end{aligned}$$

Therefore $\alpha^*(\lambda\theta_1 + \mu\theta_2) = \lambda\alpha^*(\theta_1) + \mu\alpha^*(\theta_2)$ and α^* is linear i.e. $\alpha^* \in \mathcal{L}(W^*, V^*)$.

Lemma. *Suppose that V and W are f.d. with bases (e_1, \dots, e_n) and (f_1, \dots, f_m) respectively. Let $(\epsilon_1, \dots, \epsilon_n)$ and (η_1, \dots, η_m) be the corresponding dual bases. If $\alpha: V \rightarrow W$ is represented by A with respect to (e_1, \dots, e_n) and (f_1, \dots, f_m) then α^* is represented by A^T with respect to (η_1, \dots, η_m) and $(\epsilon_1, \dots, \epsilon_n)$.*

Proof. We're given that $\alpha(e_i) = \sum A_{ki}f_k$ and must compute $\alpha^*(\eta_i)$ in terms of $\epsilon_1, \dots, \epsilon_n$.

$$\begin{aligned}\alpha^*(\eta_i)(e_j) &= \eta_i(\alpha(e_j)) \\ &= \eta_i\left(\sum_k A_{kj}f_k\right) \\ &= \sum_k A_{kj}\delta_{ik} = A_{ij}\end{aligned}$$

Thus $\alpha^*(\eta_i)(e_j) = \sum_k A_{ik}\epsilon_k(e_j) = \sum_k A_{ki}^T\epsilon_k(e_j)$ and so $\alpha^*(\eta_i) = \sum_K A_{ki}^T\epsilon_k$ as required. \square

Remarks.

- (1) If $\alpha: U \rightarrow V$ and $\beta: V \rightarrow W$ are linear maps then $(\beta\alpha)^* = \alpha^*\beta^*$.
- (2) If $\alpha, \beta: U \rightarrow V$ then $(\alpha + \beta)^* = \alpha^* + \beta^*$.
- (3) If $B = Q^{-1}AP$ is an equality of matrices with P and Q invertible, then

$$B^T = P^T A^T (Q^{-1})^T = \left((P^{-1})^T\right)^{-1} A^T (Q^{-1})^T$$

as we should expect at this point.

Lemma. Suppose that $\alpha \in \mathcal{L}(V, W)$ with V, W f.d. over \mathbf{F} . Then

- (a) $\ker \alpha^* = (\text{Im } \alpha)^\circ$;
- (b) $r(\alpha^*) = r(\alpha)$; and
- (c) $\text{Im } \alpha^* = (\ker \alpha)^\circ$.

Proof. (a) Suppose $\theta \in W^*$. Then $\theta \in \ker \alpha^*$ if and only if $\alpha^*(\theta) = 0$ if and only if $\theta\alpha(v) = 0$ for all $v \in V$ if and only if $\theta \in (\text{Im } \alpha)^\circ$.

(b) As $\text{Im } \alpha$ is a subspace of W , we've seen that $\dim \text{Im } \alpha + \dim(\text{Im } \alpha)^\circ = \dim W$. Using part (a) we can deduce that $r(\alpha) + n(\alpha^*) = \dim W = \dim W^*$. But the rank-nullity theorem gives $r(\alpha^*) + n(\alpha^*) = \dim W^*$.

(c) Suppose that $\phi \in \text{Im } \alpha^*$. Then there is some $\theta \in W^*$ such that $\phi = \alpha^*(\theta) = \theta\alpha$. Therefore for all $v \in \ker \alpha$, $\phi(v) = \theta\alpha(v) = \theta(0) = 0$. Thus $\text{Im } \alpha^* \subset (\ker \alpha)^\circ$.

But $\dim \ker \alpha + \dim(\ker \alpha)^\circ = \dim V$. So

$$\dim(\ker \alpha)^\circ = \dim V - n(\alpha) = r(\alpha) = r(\alpha^*) = \dim \text{Im } \alpha^*.$$

and so the inclusion must be an equality. \square

Notice that we have reproven that row-rank=column rank in a more conceptually satisfying way.

Lemma. Let V be a vector space over \mathbf{F} there is a canonical linear map $\text{ev}: V \rightarrow V^{**}$ given by $\text{ev}(v)(\theta) = \theta(v)$ for each $\theta \in V^*$.

Proof. First we must show that $\text{ev}(v) \in V^{**}$ whenever $v \in V$. Suppose that $\theta_1, \theta_2 \in V^*$ and $\lambda, \mu \in \mathbf{F}$. Then

$$\text{ev}(v)(\lambda\theta_1 + \mu\theta_2) = \lambda\theta_1(v) + \mu\theta_2(v) = \lambda\text{ev}(v)(\theta_1) + \mu\text{ev}(v)(\theta_2).$$

Next, we must show ev is linear, ie $\text{ev}(\lambda v_1 + \mu v_2) = \lambda\text{ev}(v_1) + \mu\text{ev}(v_2)$ whenever $v_1, v_2 \in V$, $\lambda, \mu \in \mathbf{F}$. We can show this by evaluating both sides at each $\theta \in V^*$. Then

$$\text{ev}(\lambda v_1 + \mu v_2)(\theta) = \theta(\lambda v_1 + \mu v_2) = (\lambda\text{ev}(v_1) + \mu\text{ev}(v_2))(\theta)$$

so ev is linear. \square

LECTURE 10

Lemma. *Suppose that V is f.d. then the canonical linear map $\text{ev}: V \rightarrow V^{**}$ is an isomorphism.*

Proof. Suppose that $\text{ev}(v) = 0$. Then $\theta(v) = \text{ev}(v)(\theta) = 0$ for all $\theta \in V^*$. Thus $\langle v \rangle^\circ$ has dimension $\dim V$. It follows that $\langle v \rangle$ is a space of dimension 0 so $v = 0$. In particular we've proven that ev is injective.

To complete the proof it suffices to observe that $\dim V = \dim V^* = \dim V^{**}$ so any injective linear map $V \rightarrow V^{**}$ is an isomorphism. \square

Remark. Although the canonical linear map $\text{ev}: V \rightarrow V^{**}$ always exists it is not an isomorphism in general if V is not f.d. In particular if V has a basis ev will be injective but not surjective when V is not finite-dimensional.

Lemma. *Suppose V and W are f.d. over \mathbf{F} and $\alpha \in \mathcal{L}(V, W)$ then $\alpha^{**} \circ \text{ev} = \text{ev} \circ \alpha$.*

Proof. Note that $\alpha^{**} \circ \text{ev}$ and $\text{ev} \circ \alpha$ are both linear maps $V \rightarrow W^{**}$. Suppose $v \in V$ and $\theta \in W^*$. We must show that $\alpha^{**}(\text{ev}(v))(\theta) = \text{ev}(\alpha(v))(\theta)$.

But $\alpha^{**}(\text{ev}(v))(\theta) = (\text{ev}(v)\alpha^*)(\theta) = \text{ev}(v)(\theta\alpha) = \theta(\alpha(v)) = \text{ev}(\alpha(v))(\theta)$ as required. \square

Proposition. *Suppose V is f.d. over \mathbf{F} and U, U_1 and U_2 are subspaces of V then*

- (a) $U^{\circ\circ} = \text{ev}(U)$;
- (b) $\text{ev}(U)^\circ = \text{ev}(U^\circ)$;
- (c) $(U_1 + U_2)^\circ = U_1^\circ \cap U_2^\circ$; and
- (d) $(U_1 \cap U_2)^\circ = U_1^\circ + U_2^\circ$.

Proof. (a) Let $u \in U$. Then $\text{ev}(u)(\theta) = \theta(u) = 0$ for all $\theta \in U^\circ$. Thus $\text{ev}(u) \in U^{\circ\circ}$. ie $\text{ev}(U) \subset U^{\circ\circ}$. But

$$\dim \text{ev}(U) = \dim U = \dim V - \dim U^\circ = \dim V^* - \dim U^\circ = \dim U^{\circ\circ}$$

so we have equality of subspaces since V is f.d.

(b) by (a) $\text{ev}(U)^\circ = (U^{\circ\circ})^\circ = (U^\circ)^{\circ\circ} = \text{ev}(U^\circ)$.

(c) Suppose that $\theta \in V^*$. Then $\theta \in (U_1 + U_2)^\circ$ if and only if $\theta(u_1 + u_2) = 0$ for all $u_1 \in U_1$ and $u_2 \in U_2$ if and only if $\theta(u) = 0$ for all $u \in U_1 \cup U_2$ if and only if $\theta \in U_1^\circ \cap U_2^\circ$.

(d) by parts (a) and (c),

$$\text{ev}(U_1) \cap \text{ev}(U_2) = U_1^{\circ\circ} \cap U_2^{\circ\circ} = (U_1^\circ + U_2^\circ)^\circ.$$

Thus since ev is an isomorphism and (b) holds,

$$\text{ev}((U_1 \cap U_2)^\circ) = (\text{ev}(U_1) \cap \text{ev}(U_2))^\circ = (U_1^\circ + U_2^\circ)^{\circ\circ} = \text{ev}(U_1^\circ + U_2^\circ).$$

Again using that ev is an isomorphism we deduce the result. \square

4. BILINEAR FORMS (I)

Let U and V be vector spaces over \mathbf{F} .

Definition. $\phi: U \times V \rightarrow \mathbf{F}$ is a *bilinear form* if it is linear in both arguments; i.e. if $\psi(u, -): V \rightarrow \mathbf{F} \in V^*$ for all $u \in U$ and $\phi(-, v): U \rightarrow \mathbf{F} \in U^*$ for all $v \in V$.

Examples.

(0) The map $V \times V^* \rightarrow \mathbf{F}; (v, \theta) \mapsto \theta(v)$ is a bilinear form.

- (1) $U = V = \mathbf{R}^n$; $\psi(x, y) = \sum_{i=1}^n x_i y_i$ is a bilinear form.
- (2) Suppose that $A \in \text{Mat}_{m,n}(\mathbf{F})$ then $\phi: \mathbf{F}^m \times \mathbf{F}^n \rightarrow \mathbf{F}$; $\psi(u, v) = u^T A v$ is a bilinear form.
- (3) If $U = V = C([0, 1], \mathbf{R})$ then $\phi(f, g) = \int_0^1 f(t)g(t) dt$ is a bilinear form.
- (4) The map $\mathbf{F}^2 \times \mathbf{F}^2 \rightarrow \mathbf{F}$; $\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix} \right) \mapsto ad - bc$ is a bilinear form.

Definition. Let (e_1, \dots, e_n) be a basis for U and (f_1, \dots, f_m) be a basis of V and $\phi: U \times V \rightarrow \mathbf{F}$ a bilinear form. Then the *matrix* A representing ϕ with respect to (e_1, \dots, e_n) and (f_1, \dots, f_m) is given by $A_{ij} = \phi(e_i, f_j)$.

Remark. If $u = \sum \lambda_i e_i$ and $v = \sum \mu_j f_j$ then

$$\phi\left(\sum \lambda_i e_i, \sum \mu_j f_j\right) = \sum_{i=1}^n \lambda_i \phi\left(e_i, \sum \mu_j f_j\right) = \sum_{i=1}^n \sum_{j=1}^m \lambda_i \mu_j \phi(e_i, f_j).$$

Therefore if A is the matrix representing ϕ with respect to (e_1, \dots, e_n) and (f_1, \dots, f_m) we have

$$\phi(v, w) = (\lambda_1 \quad \dots \quad \lambda_n) A \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_m \end{pmatrix}$$

and ϕ is determined by the matrix representing it.

A bilinear form $\phi: U \times V \rightarrow \mathbf{F}$ defines linear maps $\phi_L: U \rightarrow V^*$ and $\phi_R: V \rightarrow U^*$ by the formulae

$$\phi_L(u)(v) = \phi(u, v) = \phi_R(v)(u)$$

for $u \in U$ and $v \in V$.

Example. If $\phi: V \times V^* \rightarrow \mathbf{F}$; $\phi(v, \theta) = \theta(v)$ then $\phi_L: V \rightarrow V^{**}$ is ev and $\phi_R: V^* \rightarrow V^*$ is ι_{V^*} .

LECTURE 11

Recall that given a bilinear form $\phi: U \times V \rightarrow \mathbf{F}$ there are associated linear maps $\phi_L: U \rightarrow V^*$ and $\phi_R: V \rightarrow U^*$ given by

$$\phi_L(u)(v) = \phi(u, v) = \phi_R(v)(u).$$

In fact one can easily show that $\phi_R = \phi_L^* \circ \text{ev}$ and $\phi_L = \phi_R^* \circ \text{ev}$.

Lemma. Let (e_1, \dots, e_n) be a basis of U with dual basis $(\epsilon_1, \dots, \epsilon_n)$ and (f_1, \dots, f_m) be a basis for V with dual basis (η_1, \dots, η_m) . If A represents the bilinear form $\phi: U \times V \rightarrow \mathbf{F}$ with respect to (e_1, \dots, e_n) and (f_1, \dots, f_m) then A represents ϕ_R with respect to (f_1, \dots, f_m) and $(\epsilon_1, \dots, \epsilon_n)$ and A^T represents ϕ_L with respect to (e_1, \dots, e_n) and (η_1, \dots, η_m) .

Proof. We can compute $\phi_L(e_i)(f_j) = \phi(e_i, f_j) = A_{ij}$ and so $\phi_L(e_i) = \sum_{j=1}^m A_{ji}^T \eta_j$ and $\phi_R(f_j)(e_i) = \phi(e_i, f_j) = A_{ij}$ and so $\phi_R(f_j) = \sum_{i=1}^n A_{ij} \epsilon_i$. \square

Definition. We call $\ker \phi_L$ the *left kernel* of ϕ and $\ker \phi_R$ the *right kernel* of ϕ .

Note that

$$\ker \phi_L = \{u \in U \mid \phi(u, v) = 0 \text{ for all } v \in V\}$$

and

$$\ker \phi_R = \{v \in V \mid \phi(u, v) = 0 \text{ for all } u \in U\}.$$

More generally, if $T \subset U$ we write

$$T^\perp := \{v \in V \mid \phi(t, v) = 0 \text{ for all } t \in T\}$$

and if $S \subset V$ we write

$${}^\perp S := \{u \in U \mid \phi(u, s) = 0 \text{ for all } s \in S\}.$$

Definition. We say a bilinear form $\phi: U \times V \rightarrow \mathbf{F}$ is *non-degenerate* if $\ker \phi_L = 0$ and $\ker \phi_R = 0$. Otherwise we say that ϕ is *degenerate*.

Lemma. Let U and V be f.d. vector spaces over \mathbf{F} with bases (e_1, \dots, e_n) and (f_1, \dots, f_m) and let $\phi: U \times V \rightarrow \mathbf{F}$ be a bilinear form represented by the matrix A with respect to those bases. Then ϕ is non-degenerate if and only if the matrix A is invertible. In particular, if ϕ non-degenerate then $\dim U = \dim V$.

Proof. The condition that ϕ is non-degenerate is equivalent to $\ker \phi_L = 0$ and $\ker \phi_R = 0$ which is in turn equivalent to $n(A) = 0 = n(A^T)$. This last is equivalent to $r(A) = \dim U$ and $r(A^T) = \dim V$. Since row-rank and column-rank agree we can see that this final statement is equivalent to A being invertible as required. \square

It follows that, when U and V are f.d., defining a non-degenerate bilinear form $\phi: U \times V \rightarrow \mathbf{F}$ is equivalent to defining an isomorphism $\phi_L: U \rightarrow V^*$ (or equivalently an isomorphism $\phi_R: V \rightarrow U^*$).

Proposition. Suppose that (e_1, \dots, e_n) and (u_1, \dots, u_n) are two bases of U such that $u_i = \sum_{k=1}^n P_{ki} e_k$ for $i = 1, \dots, n$ and (f_1, \dots, f_m) and (v_1, \dots, v_m) are two bases of V such that $v_i = \sum_{l=1}^m Q_{li} f_l$ for $i = 1, \dots, m$. Let $\phi: U \times V \rightarrow \mathbf{F}$ be a bilinear form represented by A with respect to (e_1, \dots, e_n) and (f_1, \dots, f_m) and by B with respect to (u_1, \dots, u_n) and (v_1, \dots, v_m) then

$$B = P^T A Q.$$

Proof.

$$\begin{aligned} B_{ij} &= \phi(u_i, v_j) \\ &= \phi\left(\sum_{k=1}^n P_{ki} e_k, \sum_{l=1}^m Q_{lj} f_l\right) \\ &= \sum_{k,l} P_{ki} Q_{lj} \phi(e_k, f_l) \\ &= (P^T A Q)_{ij} \end{aligned}$$

\square

Definition. We define the *rank* of ϕ , $r(\phi)$, to be the rank of any matrix representing ϕ . Since $r(P^T A Q) = r(A)$ for any invertible matrices P and Q we see that this is independent of any choices.

Note that $r(\phi) = r(\phi_R) = r(\phi_L)$.

5. DETERMINANTS OF MATRICES

Recall that S_n is the group of permutations of the set $\{1, \dots, n\}$. Moreover we can define a group homomorphism $\epsilon: S_n \rightarrow \{\pm 1\}$ such that $\epsilon(\sigma) = 1$ whenever σ is a product of an even number of transpositions and $\epsilon(\sigma) = -1$ whenever σ is a product of an odd number of transpositions.

Definition. If $A \in \text{Mat}_n(\mathbf{F})$ then the *determinant* of A

$$\det A := \sum_{\sigma \in S_n} \epsilon(\sigma) \left(\prod_{i=1}^n A_{i\sigma(i)} \right).$$

Example. If $n = 2$ then $\det A = A_{11}A_{22} - A_{12}A_{21}$.

Lemma. $\det A = \det A^T$.

Proof.

$$\begin{aligned} \det A^T &= \sum_{\sigma \in S_n} \epsilon(\sigma) \prod_{i=1}^n A_{\sigma(i)i} \\ &= \sum_{\sigma \in S_n} \epsilon(\sigma) \prod_{i=1}^n A_{i\sigma^{-1}(i)} \\ &= \sum_{\tau \in S_n} \epsilon(\tau^{-1}) \prod_{i=1}^n A_{i\tau(i)} \\ &= \det A \end{aligned}$$

□

Definition. A *volume form* d on \mathbf{F}^n is a function $\mathbf{F}^n \times \mathbf{F}^n \times \dots \times \mathbf{F}^n \rightarrow \mathbf{F}$; $(v_1, \dots, v_n) \mapsto d(v_1, \dots, v_n)$ such that

(i) d is *multi-linear* i.e. for each $1 \leq i \leq n$, and $v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n \in V$

$$d(v_1, \dots, v_{i-1}, -, v_{i+1}, \dots, v_n) \in (\mathbf{F}^n)^*$$

(ii) d is *alternating* i.e. whenever $v_i = v_j$ for some $i \neq j$ then $d(v_1, \dots, v_n) = 0$.

One may view a matrix $A \in \text{Mat}_n(\mathbf{F})$ as an n -tuple of elements of \mathbf{F}^n given by its columns $A = (A^{(1)} \dots A^{(n)})$ with $A^{(1)}, \dots, A^{(n)} \in \mathbf{F}^n$.

Lemma. $\det: \mathbf{F}^n \times \dots \times \mathbf{F}^n \rightarrow \mathbf{F}$; $(A^{(1)}, \dots, A^{(n)}) \mapsto \det A$ is a *volume form*.

Proof. To see that \det is multilinear it suffices to see that $\prod_{i=1}^n A_{i\sigma(i)}$ is multilinear for each $\sigma \in S_n$ since a sum of (multi)-linear functions is (multi)-linear. Since one term from each column appears in each such product this is easy to see.

Suppose now that $A^{(k)} = A^{(l)}$ for some $k \neq l$. Let τ be the transposition (kl) . Then $a_{ij} = a_{i\tau(j)}$ for every i, j in $\{1, \dots, n\}$. We can write S_n is a disjoint union of cosets $A_n \coprod \tau A_n$.

Then

$$\sum_{\sigma \in A_n} \prod a_{i\sigma(i)} = \sum_{\sigma \in A_n} \prod a_{i\tau\sigma(i)} = \sum_{\sigma \in \tau A_n} \prod a_{i\sigma(i)}$$

Thus $\det A = \text{LHS} - \text{RHS} = 0$.

□

LECTURE 12

Lemma. *Let d be a volume form. Swapping two entries changes the sign. i.e.*

$$d(v_1, \dots, v_i, \dots, v_j, \dots, v_n) = -d(v_1, \dots, v_j, \dots, v_i, \dots, v_n).$$

Proof. Consider $d(v_1, \dots, v_i + v_j, \dots, v_i + v_j, \dots, v_n) = 0$. Expanding the left-hand-side using linearity of the i th and j th coordinates we obtain

$$\begin{aligned} d(v_1, \dots, v_i, \dots, v_i, \dots, v_n) + d(v_1, \dots, v_i, \dots, v_j, \dots, v_n) + \\ d(v_1, \dots, v_j, \dots, v_i, \dots, v_n) + d(v_1, \dots, v_j, \dots, v_j, \dots, v_n) = 0. \end{aligned}$$

Since the first and last terms on the left are zero, the statement follows immediately. \square

Corollary. *If $\sigma \in S_n$ then $d(v_{\sigma(1)}, \dots, v_{\sigma(n)}) = \epsilon(\sigma)d(v_1, \dots, v_n)$.* \square

Theorem. *Let d be a volume form on \mathbf{F}^n . Let A be a matrix with i th column $A^{(i)} \in \mathbf{F}^n$. Then*

$$d(A^{(1)}, \dots, A^{(n)}) = \det A \cdot d(e_1, \dots, e_n).$$

In other words \det is the unique volume form d such that $d(e_1, \dots, e_n) = 1$.

Proof. We compute

$$\begin{aligned} d(A^{(1)}, \dots, A^{(n)}) &= d\left(\sum_{i=1}^n A_{i1}e_i, A^{(2)}, \dots, A^{(n)}\right) \\ &= \sum_i A_{i1}d(e_i, A^{(2)}, \dots, A^{(n)}) \\ &= \sum_{i,j} A_{i1}A_{j2}d(e_i, e_j, \dots, A^{(n)}) \\ &= \sum_{i_1, \dots, i_n} \left(\prod_{j=1}^n A_{i_j j}\right) d(e_{i_1}, \dots, e_{i_n}) \end{aligned}$$

But $d(e_{i_1}, \dots, e_{i_n}) = 0$ unless i_1, \dots, i_n are distinct. That is unless there is some $\sigma \in S_n$ such that $i_j = \sigma(j)$. Thus

$$d(A^{(1)}, \dots, A^{(n)}) = \sum_{\sigma \in S_n} \left(\prod_{j=1}^n A_{\sigma(j)j}\right) d(e_{\sigma(1)}, \dots, e_{\sigma(n)}).$$

But $d(e_{\sigma(1)}, \dots, e_{\sigma(n)}) = \epsilon(\sigma)d(e_1, \dots, e_n)$ so we're done. \square

Remark. We can interpret this as saying that for every matrix A ,

$$d(Ae_1, \dots, Ae_n) = \det A \cdot d(e_1, \dots, e_n).$$

Theorem. *Let $A, B \in \text{Mat}_n(\mathbf{F})$. Then $\det(AB) = \det A \det B$.*

Proof. Let d be a non-zero volume form on \mathbf{F}^n , for example \det as $\det I_n = 1$. Let $d_A: \mathbf{F}^n \times \dots \times \mathbf{F}^n \rightarrow \mathbf{F}$ be given by $d_A(v_1, \dots, v_n) = d(Av_1, \dots, Av_n)$. Then d_A is a volume form since it is multi-linear ($v_i \mapsto Av_i$ is linear) and alternating (if $v_i = v_j$ then $Av_i = Av_j$). Thus

$$d_A(Be_1, \dots, Be_n) = \det(B)d_A(e_1, \dots, e_n) = \det(B) \det(A)d(e_1, \dots, e_n)$$

by the last theorem. But we can also compute

$$d_A(Be_1, \dots, Be_n) = d(ABe_1, \dots, ABe_n) = \det(AB)d(e_1, \dots, e_n)$$

also by the last theorem.

Thus as $d(e_1, \dots, e_n) \neq 0$ we can see that $\det(AB) = \det A \det B$. \square

It follows that for any volume form d any $v_1, \dots, v_n \in \mathbf{F}^n$ and any $A \in \text{Mat}_n(\mathbf{F})$,

$$d(Av_1, \dots, Av_n) = (\det A)d(v_1, \dots, v_n)$$

since there is a matrix B such that $v_i = Be_i$ and

$$\begin{aligned} d(Av_1, \dots, Av_n) &= d(ABe_1, \dots, ABe_n) \\ &= \det(A) \det(B) d(e_1, \dots, e_n) \\ &= \det(A) d(Be_1, \dots, Be_n) \\ &= \det(A) d(v_1, \dots, v_n). \end{aligned}$$

Definition. We say A is *singular* if $\det A = 0$.

Corollary. If A is invertible then A is non-singular and $\det(A^{-1}) = \frac{1}{\det A}$.

Proof. We can compute

$$1 = \det I_n = \det(AA^{-1}) = \det A \det A^{-1}.$$

Thus $\det A \neq 0$ and $\det A^{-1} = \frac{1}{\det A}$ as required. \square

Theorem. Let $A \in \text{Mat}_n(\mathbf{F})$. The following statements are equivalent:

- (a) A is invertible;
- (b) A is non-singular;
- (c) $r(A) = n$.

Proof. We've seen that (a) implies (b) above.

Suppose that $r(A) < n$. Then by the rank-nullity theorem $n(A) > 0$ and so there is some $\lambda \in \mathbf{F}^n \setminus \{0\}$ such that $A\lambda = 0$ i.e. there is a linear relation between the columns of A ; $\sum_{i=1}^n \lambda_i A^{(i)} = 0$ for some $\lambda_i \in \mathbf{F}$ not all zero.

Suppose that $\lambda_k \neq 0$ and let B be the matrix with i th column e_i for $i \neq k$ and k th column λ . Then AB has k th column 0. Thus $\det AB = 0$. But we can compute $\det AB = \det A \det B = \lambda_k \det A$. Since $\lambda_k \neq 0$, $\det A = 0$. Thus (b) implies (c).

Finally (c) implies (a) by the rank-nullity theorem: $r(A) = n$ implies $n(A) = 0$ and the linear map corresponding to A is bijective as required. \square

Notation. Let \widehat{A}_{ij} denote the submatrix of A obtained by deleting the i th row and the j th column.

Lemma. Let $A \in \text{Mat}_n(\mathbf{F})$. Then

- (a) (expanding determinant along the j th column) $\det A = \sum_{i=1}^n (-1)^{i+j} A_{ij} \det \widehat{A}_{ij}$;
- (b) (expanding determinant along the i th row) $\det A = \sum_{j=1}^n (-1)^{i+j} A_{ij} \det \widehat{A}_{ij}$.

Proof. Since $\det A = \det A^T$ it suffices to verify (a).

Now

$$\begin{aligned}
 \det A &= \det(A^{(1)}, \dots, A^{(n)}) \\
 &= \det(A^{(1)}, \dots, \sum_i A_{ij} e_i, \dots, A^{(n)}) \\
 &= \sum_i A_{ij} \det(A^{(1)}, \dots, e_i, \dots, A^{(n)}) \\
 &= \sum_i A_{ij} (-1)^{i+j} \det B
 \end{aligned}$$

where

$$B = \begin{pmatrix} \widehat{A_{ij}} & 0 \\ * & 1 \end{pmatrix}.$$

Finally for $\sigma \in S_n$, $\prod_{i=1}^n B_{i\sigma(i)} = 0$ unless $\sigma(n) = n$ and we see easily that $\det B = \det \widehat{A_{ij}}$ as required. \square

Definition. Let $A \in \text{Mat}_n(\mathbf{F})$. The *adjugate matrix* $\text{adj } A$ is the element of $\text{Mat}_n(\mathbf{F})$ such that

$$(\text{adj } A)_{ij} = (-1)^{i+j} \det \widehat{A_{ji}}.$$

Theorem. Let $A \in \text{Mat}_n(\mathbf{F})$. Then

$$(\text{adj } A)A = A(\text{adj } A) = (\det A)I_n.$$

Thus if $\det A \neq 0$ then $A^{-1} = \frac{1}{\det A} \text{adj } A$

Proof. We compute

$$\begin{aligned}
 ((\text{adj } A)A)_{jk} &= \sum_{i=1}^n (\text{adj } A)_{ji} A_{ik} \\
 &= \sum_{i=1}^n (-1)^{j+i} \det \widehat{A_{ij}} A_{ik}
 \end{aligned}$$

The right-hand-side is $\det A$ if $k = j$. If $k \neq j$ then the right-hand-side is the determinant of the matrix obtained by replacing the j th column of A by the k th column. Since the resulting matrix has two identical columns $((\text{adj } A)A)_{jk} = 0$ in this case. Therefore $(\text{adj } A)A = (\det A)I_n$ as required.

We can now obtain $A \text{adj } A = (\det A)I_n$ either by using a similar argument using the rows or by considering the transpose of $A \text{adj } A$. The final part follows immediately. \square

Remark. Note that the entries of the adjugate matrix are all given by polynomials in the entries of A . Since the determinant is also a polynomial, it follows that the entries of the inverse of an invertible square matrix are given by a rational function (i.e. a ratio of two polynomial functions) in the entries of A . Whilst this is a very useful fact from a theoretical point of view, computationally there are better ways of computing the determinant and inverse of a matrix than using these formulae.

LECTURE 13

We'll complete this section on determinants of matrices with a couple of results about block triangular matrices.

Lemma. *Let $A \in \text{Mat}_k(\mathbf{F})$ and $B \in \text{Mat}_l(\mathbf{F})$ with $k, l \geq 1$. Then*

$$\det \begin{pmatrix} A & C \\ 0 & B \end{pmatrix} = \det(A) \det(B).$$

Proof. Define

$$X = \begin{pmatrix} A & C \\ 0 & B \end{pmatrix}$$

and $n = k + l$. Then

$$\det X = \sum_{\sigma \in S_n} \epsilon(\sigma) \left(\prod_{i=1}^n X_{i\sigma(i)} \right).$$

Since $X_{ij} = 0$ whenever $i > k$ and $j \leq k$ the terms with σ such that $\sigma(i) \leq k$ for some $i > k$ are all zero. So we may restrict the sum to those σ such that $\sigma(i) > k$ for $i > k$ i.e. those σ that restrict to a permutation of $\{1, \dots, k\}$. We may factorise these σ as $\sigma = \sigma_1 \sigma_2$ with $\sigma_1 \in S_k$ and σ_2 a permutation of $\{k+1, \dots, n\}$. Thus

$$\begin{aligned} \det X &= \sum_{\sigma_1} \sum_{\sigma_2} \epsilon(\sigma_1 \sigma_2) \left(\prod_{i=1}^k X_{i\sigma_1(i)} \right) \left(\prod_{j=1}^l X_{j+k, \sigma_2(j+k)} \right) \\ &= \left(\sum_{\sigma_1 \in S_k} \epsilon(\sigma_1) \left(\prod_{i=1}^k A_{i\sigma_1(i)} \right) \right) \left(\sum_{\sigma_2 \in S_l} \epsilon(\sigma_2) \left(\prod_{j=1}^l B_{j\sigma_2(j)} \right) \right) \\ &= \det A \det B \end{aligned}$$

□

Corollary. *If A_1, \dots, A_k are square matrices (possibly 1×1)*

$$\det \begin{pmatrix} A_1 & * & * \\ 0 & \ddots & * \\ 0 & 0 & A_k \end{pmatrix} = \prod_{i=1}^k \det A_i \quad \square$$

Warning: it is *not* true in general that if $A, B, C, D \in \text{Mat}_n(\mathbf{F})$ and M is the element of $\text{Mat}_{2n}(\mathbf{F})$ given by

$$M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$$

then $\det M = \det A \det D - \det B \det C$.

6. ENDOMORPHISMS

6.1. Invariants.

Definition. Suppose that V is a finite dimensional vector space over \mathbf{F} . An *endomorphism* of V is a linear map $\alpha: V \rightarrow V$. We'll write $\text{End}(V)$ denote the vector space of endomorphisms of V and ι to denote the identity endomorphism of V .

When considering endomorphisms as matrices we will always choose the same basis for V for both the domain and the range.

Lemma. Suppose that (e_1, \dots, e_n) and (f_1, \dots, f_n) are bases for V such that $f_i = \sum P_{ki}e_k$. Let $\alpha \in \text{End}(V)$, A be the matrix representing α with respect to (e_1, \dots, e_n) and B the matrix representing α with respect to (f_1, \dots, f_n) . Then $B = P^{-1}AP$.

Proof. This is a special case of the change of basis formula for all linear maps between f.d. vector spaces. \square

Definition. We say matrices A and B are *similar* (or *conjugate*) if $B = P^{-1}AP$ for some invertible matrix P .

Recall $GL_n(\mathbf{F})$ denotes all the invertible matrices in $\text{Mat}_n(\mathbf{F})$. Then $GL_n(\mathbf{F})$ acts on $\text{Mat}_n(\mathbf{F})$ by conjugation and two such matrices are similar precisely if they lie in the same orbit. Thus similarity is an equivalence relation.

An important problem is to classify elements of $\text{Mat}_n(\mathbf{F})$ up to similarity (ie classify $GL_n(\mathbf{F})$ -orbits). It will help us to find basis independent invariants of the corresponding endomorphisms.

Recall that the *trace* of $A \in \text{Mat}_n(\mathbf{F})$ is defined by $\text{tr } A = \sum A_{ii} \in \mathbf{F}$.

Lemma.

- (a) If $A \in \text{Mat}_{n,m}(\mathbf{F})$ and $B \in \text{Mat}_{m,n}(\mathbf{F})$ then $\text{tr } AB = \text{tr } BA$.
- (b) If A and B are similar then $\text{tr } A = \text{tr } B$.
- (c) If A and B are similar then $\det A = \det B$.

Proof. (a)

$$\begin{aligned} \text{tr } AB &= \sum_{i=1}^n \left(\sum_{j=1}^m A_{ij} B_{ji} \right) \\ &= \sum_{j=1}^m \left(\sum_{i=1}^n B_{ji} A_{ij} \right) \\ &= \text{tr } BA \end{aligned}$$

If $B = P^{-1}AP$ then,

- (b) $\text{tr } B = \text{tr}(P^{-1}A)P = \text{tr } P(P^{-1}A) = \text{tr } A$.
- (c) $\det B = \det P^{-1} \det A \det P = \frac{1}{\det P} \det A \det P = \det A$. \square

Definition. Let $\alpha \in \text{End}(V)$, (e_1, \dots, e_n) be a basis for V and A the matrix representing α with respect to (e_1, \dots, e_n) . Then the *trace* of α written $\text{tr } \alpha$ is defined to be the trace of A and the *determinant* of α written $\det \alpha$ is defined to be the determinant of A .

We've proven that the trace and determinant of α do not depend on the choice of basis (e_1, \dots, e_n) .

Remark. Since $\det A = \det A^T$ and $\text{tr } A = \text{tr } A^T$, if $\alpha \in \text{End}(V)$ then its dual $\alpha^* \in \text{End}(V^*)$ satisfies $\det \alpha^* = \det \alpha$ and $\text{tr } \alpha^* = \text{tr } \alpha$.

Definition. Let $\alpha \in \text{End}(V)$.

- (a) $\lambda \in \mathbf{F}$ is an *eigenvalue* of α if there is $v \in V \setminus \{0\}$ such that $\alpha v = \lambda v$.
- (b) $v \in V$ is an *eigenvector* for α if $\alpha(v) = \lambda v$ for some $\lambda \in \mathbf{F}$.
- (c) When $\lambda \in \mathbf{F}$, the λ -*eigenspace* of α , written $E_\alpha(\lambda)$ or simply $E(\lambda)$ is the set of λ -eigenvectors of α ; i.e. $E(\lambda) = \ker(\alpha - \lambda I)$.

(d) The *characteristic polynomial* of α is defined by

$$\chi_\alpha(t) = \det(tI - \alpha).$$

Remarks.

- (1) $\chi_\alpha(t)$ is a monic polynomial in t of degree $\dim V$.
- (2) $\lambda \in \mathbf{F}$ is an eigenvalue of α if and only if $\ker(\alpha - \lambda I) \neq 0$ if and only if λ is a root of $\chi_\alpha(t)$.
- (3) If $A \in \text{Mat}_n(\mathbf{F})$ we can define $\chi_A(t) = \det(tI_n - A)$. Then similar matrices have the same characteristic polynomials.

Lemma. *Let $\alpha \in \text{End}(V)$ and $\lambda_1, \dots, \lambda_d$ be the distinct eigenvalues of α . Then $E(\lambda_1) + \dots + E(\lambda_d)$ is a direct sum of the $E(\lambda_i)$.*

Proof. Suppose that $\sum_{i=1}^d x_i = \sum_{i=1}^d y_i$ with $x_i, y_i \in E(\lambda_i)$. Consider the linear maps

$$\beta_j := \prod_{k \neq j} (\alpha - \lambda_k I).$$

Then

$$\begin{aligned} \beta_j\left(\sum_{i=1}^d x_i\right) &= \sum_{i=1}^d \beta_j(x_i) \\ &= \sum_{i=1}^d \left(\prod_{k \neq j} (\alpha - \lambda_k I) \right) (x_i) \\ &= \sum_{i=1}^d \left(\prod_{k \neq j} (\lambda_i - \lambda_k) \right) x_i \\ &= \left(\prod_{k \neq j} (\lambda_j - \lambda_k) \right) x_j \end{aligned}$$

Similarly, $\beta_j(\sum_{i=1}^k y_i) = \prod_{r \neq j} (\lambda_j - \lambda_r) y_i$. Thus since $\prod_{k \neq j} (\lambda_j - \lambda_r) \neq 0$, $x_j = y_j$ and the expression is unique. \square

Note that the proof of this lemma shows that any set of non-zero eigenvectors with distinct eigenvalues is LI.

LECTURE 14

Definition. $\alpha \in \text{End}(V)$ is *diagonalisable* if there is a basis for V such that the corresponding matrix is diagonal.

Theorem. *Let $\alpha \in \text{End}(V)$. Let $\lambda_1, \dots, \lambda_k$ be the distinct eigenvalues of α . Write $E_i = E(\lambda_i)$. Then the following are equivalent*

- (a) α is diagonalisable;
- (b) V has a basis consisting of eigenvectors of α ;
- (c) $V = \sum_{i=1}^k E_i$;
- (d) $V = \bigoplus_{i=1}^k E_i$;
- (e) $\sum \dim E_i = \dim V$.

Proof. Suppose that (e_1, \dots, e_n) is a basis for V and A is the matrix representing α with respect to this basis. Then $\alpha(e_i) = \sum A_{ji}e_j$. Thus A is diagonal if and only if each e_i is an eigenvector for α . i.e. (a) and (b) are equivalent.

Now it follows from (b) that the E_i span V and thus $V = \sum E_i$ ie (b) implies (c). We've proven above that $\sum E_i = \bigoplus_{i=1}^k E_i$ so (c) and (d) are equivalent. Moreover (d) implies (b) since we may take a basis for each E_i and then take the union to form a basis of eigenvectors of α .

That (d) is equivalent to (e) given that $\sum E_i$ is a direct sum of the E_i is a basic fact about direct sums that follows from Example Sheet 1 Q9. \square

6.2. Minimal polynomials.

6.2.1. An aside on polynomials.

Definition. A polynomial over \mathbf{F} is something of the form

$$f(t) = a_m t^m + \dots + a_1 t + a_0$$

for some $m \geq 0$ and $a_0, \dots, a_m \in \mathbf{F}$. The largest n such that $a_n \neq 0$ is the *degree* of f written $\deg f$. Thus $\deg 0 = -\infty$.

It is straightforward to show that

$$\deg(f + g) \leq \max(\deg f, \deg g)$$

and

$$\deg fg = \deg f + \deg g.$$

Notation. We write $\mathbf{F}[t] := \{\text{polynomials with coefficients in } \mathbf{F}\}$.

Note that a polynomial over \mathbf{F} defines a function $\mathbf{F} \rightarrow \mathbf{F}$ but we don't identify the polynomial with this function. For example if $\mathbf{F} = \mathbf{Z}/p\mathbf{Z}$ then $t^p \neq t$ even though they define the same function on \mathbf{F} . If you are restricting \mathbf{F} to be just \mathbf{R} or \mathbf{C} this point is not important.

Lemma (Polynomial division). *Given $f, g \in \mathbf{F}[t]$, $g \neq 0$ there exist $q, r \in \mathbf{F}[t]$ such that $f(t) = q(t)g(t) + r(t)$ and $\deg r < \deg g$.*

Lemma. *If $\lambda \in \mathbf{F}$ is a root of a polynomial $f(t)$, i.e. $f(\lambda) = 0$, then $f(t) = (t - \lambda)g(t)$ for some $g(t) \in \mathbf{F}[t]$.*

Proof. There are $q, r \in \mathbf{F}[t]$ such that $f(t) = (t - \lambda)q(t) + r(t)$ with $\deg r < 1$. But $\deg r < 1$ means $r(t) = r_0$ some $r_0 \in \mathbf{F}$. But then $0 = f(\lambda) = (\lambda - \lambda)q(\lambda) + r_0 = r_0$. So $r_0 = 0$ and we're done. \square

Definition. If $f \in \mathbf{F}[t]$ and $\lambda \in \mathbf{F}$ is a root of f we say that λ is a *root of multiplicity k* if $(t - \lambda)^k$ is a factor of $f(t)$ but $(t - \lambda)^{k+1}$ is not a factor of f . i.e. if $f(t) = (t - \lambda)^k g(t)$ for some $g(t) \in \mathbf{F}[t]$ with $g(\lambda) \neq 0$.

We can use the last lemma and induction to show that every $f(t)$ can be written as

$$f(t) = \prod_{i=1}^r (t - \lambda_i)^{a_i} g(t)$$

with $r \geq 0$, $a_1, \dots, a_r \geq 1$, $\lambda_1, \dots, \lambda_r \in \mathbf{F}$ and $g(t) \in \mathbf{F}[t]$ with no roots in \mathbf{F} .

Lemma. *A polynomial $f \in \mathbf{F}[t]$ of degree $n \geq 0$ has at most n roots counted with multiplicity.*

Corollary. *Suppose $f, g \in \mathbf{F}[t]$ have degrees less than n and $f(\lambda_i) = g(\lambda_i)$ for $\lambda_1, \dots, \lambda_n \in \mathbf{F}$ distinct. Then $f = g$.*

Proof. Consider $f - g$ which has degree less than n but at least n roots, namely $\lambda_1, \dots, \lambda_n$. Thus $\deg(f - g) = -\infty$ and so $f = g$. \square

Theorem (Fundamental Theorem of Algebra). *Every polynomial $f \in \mathbf{C}[t]$ of degree at least 1 has a root in \mathbf{C} .*

It follows that $f \in \mathbf{C}[t]$ has precisely n roots in \mathbf{C} counted with multiplicity. It also follows that every $f \in \mathbf{R}[t]$ can be written as a product of its linear and quadratic factors.

6.2.2. Minimal polynomials.

Notation. Given $f(t) = \sum_{i=0}^m a_i t^i \in \mathbf{F}[t]$, $A \in \text{Mat}_n(\mathbf{F})$ and $\alpha \in \text{End}(V)$ we write

$$f(A) := \sum_{i=0}^m a_i A^i$$

and

$$f(\alpha) := \sum_{i=0}^m a_i \alpha^i.$$

Here $A^0 = I_n$ and $\alpha^0 = \iota$.

Theorem. *Suppose that $\alpha \in \text{End}(V)$. Then α is diagonalisable if and only if there is a non-zero polynomial $p(t) \in \mathbf{F}[t]$ that can be expressed as a product of distinct linear factors such that $p(\alpha) = 0$.*

Proof. Suppose that α is diagonalisable and $\lambda_1, \dots, \lambda_k \in \mathbf{F}$ are the distinct eigenvalues of α . Thus if v is an eigenvector for α then $\alpha(v) = \lambda_i v$ for some $i = 1, \dots, k$. Let $p(t) = \prod_{j=1}^k (t - \lambda_j)$

Since α is diagonalisable, $V = \bigoplus_{i=1}^k E(\lambda_i)$ and each $v \in V$ can be written as $v = \sum v_i$ with $v_i \in E(\lambda_i)$. Then

$$p(\alpha)(v) = \sum_{i=1}^k p(\alpha)(v_i) = \sum_{i=1}^k \prod_{j=1}^k (\lambda_i - \lambda_j) v_i = 0.$$

Thus $p(\alpha)(v) = 0$ for all $v \in V$ and so $p(\alpha) = 0 \in \text{End}(V)$.

LECTURE 15

Conversely, if $p(\alpha) = 0$ for $p(t) = \prod_{i=1}^k (t - \lambda_i)$ for $\lambda_1, \dots, \lambda_k \in \mathbf{F}$ distinct (note that without loss of generality we may assume that p is monic). We will show that $V = \sum_{i=1}^k E(\lambda_i)$.

Let

$$q_j(t) := \prod_{\substack{i=1 \\ i \neq j}}^k \frac{(t - \lambda_i)}{(\lambda_j - \lambda_i)}$$

for $j = 1, \dots, k$. Thus $q_j(\lambda_i) = \delta_{ij}$.

Now $q(t) = \sum_{j=1}^k q_j(t) \in \mathbf{F}[t]$ has degree at most $k - 1$ and $q(\lambda_i) = 1$ for each $i = 1, \dots, k$. It follows that $q(t) = 1$.

Let $\pi_j: V \rightarrow V$ be defined by $\pi_j = q_j(\alpha)$. Then $\sum \pi_j = q(\alpha) = \iota$.

Let $v \in V$. Then $v = \iota(v) = \sum \pi_j(v)$. But

$$(\alpha - \lambda_j \iota)q_j(\alpha) = \frac{1}{\prod_{i \neq j} (\lambda_j - \lambda_i)} p(\alpha)v = 0.$$

Thus $\pi_j(v) \in \ker(\alpha - \lambda_j \iota) = E(\lambda_j)$ and we're done. \square

Remark. In the above proof, if $v \in E(\lambda_i)$ then $\pi_j(v) = q_j(\lambda_i)v = \delta_{ij}v$. So π_j is a projection onto $E(\lambda_j)$ along $\oplus_{i \neq j} E(\lambda_i)$.

Definition. The *minimal polynomial* of $\alpha \in \text{End}(V)$ is the non-zero monic polynomial $m_\alpha(t)$ of least degree such that $m_\alpha(\alpha) = 0$.

Note that if $\dim V = n < \infty$ then $\dim \text{End}(V) = n^2$, so $\iota, \alpha, \alpha^2, \dots, \alpha^{n^2}$ are linearly dependent since there are $n^2 + 1$ of them. Thus there is some non-trivial linear equation $\sum_{i=0}^{n^2} a_i \alpha^i = 0$. i.e. there is a non-zero polynomial $p(t)$ of degree at most n^2 such that $p(\alpha) = 0$.

Note also that if A represents α with respect to some basis then $p(A)$ represents $p(\alpha)$ with respect to the same basis for any polynomial $p(t) \in \mathbf{F}[t]$. Thus if we define the minimal polynomial of A in a similar fashion then $m_A(t) = m_\alpha(t)$ i.e. minimal polynomial can be viewed as an invariant on square matrices that is constant on GL_n -orbits.

Lemma. Let $\alpha \in \text{End}(V)$, $p \in \mathbf{F}[t]$ then $p(\alpha) = 0$ if and only if $m_\alpha(t)$ is a factor of $p(t)$. In particular $m_\alpha(t)$ is well-defined.

Proof. We can find $q, r \in \mathbf{F}[t]$ such that $p(t) = q(t)m_\alpha(t) + r(t)$ with $\deg r < \deg m_\alpha$. Then $p(\alpha) = q(\alpha)m_\alpha(\alpha) + r(\alpha) = 0 + r(\alpha)$. Thus $p(\alpha) = 0$ if and only if $r(\alpha) = 0$. But the minimality of the degree of m_α means that $r(\alpha) = 0$ if and only if $r = 0$ i.e. if and only if m_α is a factor of p .

Now if m_1, m_2 are both minimal polynomials for α then m_1 divides m_2 and m_2 divides m_1 so as both are monic $m_2 = m_1$. \square

Example. If $V = \mathbf{F}^2$ then

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \text{ and } B = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

both satisfy the polynomial $(t-1)^2$. Thus their minimal polynomials must be either $(t-1)$ or $(t-1)^2$. One can see that $m_A(t) = t-1$ but $m_B(t) = (t-1)^2$ so minimal polynomials distinguish these two similarity classes.

Theorem (Diagonalisability Theorem). Let $\alpha \in \text{End}(V)$ then α is diagonalisable if and only if $m_\alpha(t)$ is a product of distinct linear factors.

Proof. If α is diagonalisable there is some polynomial $p(t)$ that is a product of distinct linear factors such that $p(\alpha) = 0$ then m_α divides $p(t)$ so must be a product of distinct linear factors. The converse is already proven. \square

Theorem. Let $\alpha, \beta \in \text{End}(V)$ be diagonalisable. Then α, β are simultaneously diagonalisable (i.e. there is a single basis with respect to which the matrices representing α and β are both diagonal) if and only if α and β commute.

Proof. Certainly if there is a basis (e_1, \dots, e_n) such that α and β are represented by diagonal matrices, A and B respectively, then α and β commute since A and B commute and $\alpha\beta$ is represented by AB and $\beta\alpha$ by BA .

For the converse, suppose that α and β commute. Let $\lambda_1, \dots, \lambda_k$ denote the distinct eigenvalues of α and let $E_i = E_\alpha(\lambda_i)$ for $i = 1, \dots, k$. Then as α is diagonalisable we know that $V = \bigoplus_{i=1}^k E_i$.

We claim that $\beta(E_i) \subset E_i$ for each $i = 1, \dots, k$. To see this, suppose that $v \in E_i$ for some such i . Then

$$\alpha\beta(v) = \beta\alpha(v) = \beta(\lambda_i v) = \lambda_i\beta(v)$$

and so $\beta(v) \in E_i$ as claimed. Thus we can view $\beta|_{E_i}$ as an endomorphism of E_i .

Now since β is diagonalisable, the minimal polynomial m_β of β has distinct linear factors. But $m_\beta(\beta|_{E_i}) = m_\beta(\beta)|_{E_i} = 0$. Thus $\beta|_{E_i}$ is diagonalisable for each E_i and we can find B_i a basis of E_i consisting of eigenvectors of β . Then $B = \bigcup_{i=1}^k B_i$ is a basis for V . Moreover α and β are both diagonal with respect to this basis. \square

6.3. The Cayley-Hamilton Theorem. Recall that the characteristic polynomial of an endomorphism $\alpha \in \text{End}(V)$ is defined by $\chi_\alpha(t) = \det(tI - \alpha)$.

Theorem (Cayley–Hamilton Theorem). *Suppose that V is a f.d. vector space over \mathbf{F} and $\alpha \in \text{End}(V)$. Then $\chi_\alpha(\alpha) = 0$. In particular m_α divides χ_α (and so $\deg m_\alpha \leq \dim V$).*

Remarks.

- (1) It is tempting to substitute ‘ $t = A$ ’ into $\chi_A(t) = \det(tI_n - A)$ but it is not possible to make sense of this.
- (2) If $p(t) \in \mathbf{F}[t]$ and

$$A = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \lambda_n \end{pmatrix}$$

is diagonal then

$$p(A) = \begin{pmatrix} p(\lambda_1) & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & p(\lambda_n) \end{pmatrix}.$$

So as $\chi_A(t) = \prod_{i=1}^n (t - \lambda_i)$ we see $\chi_A(A) = 0$. So Cayley–Hamilton is obvious when α is diagonalisable.

Definition. $\alpha \in \text{End}(V)$ is *triangulable* if there is a basis for V such that the corresponding matrix is upper triangular.

Lemma. *An endomorphism α is triangulable if and only if $\chi_\alpha(t)$ can be written as a product of linear factors. In particular if $\mathbf{F} = \mathbf{C}$ then every matrix is triangulable.*

Proof. Suppose that α is triangulable and is represented by

$$\begin{pmatrix} a_1 & * & * \\ 0 & \ddots & * \\ 0 & 0 & a_n \end{pmatrix}$$

with respect to some basis. Then

$$\begin{aligned}\chi_\alpha(t) &= \det \left(tI_n - \begin{pmatrix} a_1 & * & * \\ 0 & \ddots & * \\ 0 & 0 & a_n \end{pmatrix} \right) \\ &= \prod (t - a_i).\end{aligned}$$

Thus χ_α is a product of linear factors.

LECTURE 16

We'll prove the converse by induction on $n = \dim V$. If $n = 0, 1$ every matrix is upper triangular.

Suppose that $n > 1$ and the result holds for all endomorphisms of spaces of smaller dimension. By hypothesis $\chi_\alpha(t)$ has a root $\lambda \in \mathbf{F}$. Let $U = E(\lambda) \neq 0$. Then $\alpha(U) \leq U$ and so α induces $\bar{\alpha} \in \text{End}(V/U)$ given by $\bar{\alpha}(v + U) = \alpha(v) + U$. Moreover if (v_1, \dots, v_n) is a basis for V such that v_1, \dots, v_k is a basis for U then α is represented by a matrix of the form

$$\begin{pmatrix} \lambda I_k & * \\ 0 & B \end{pmatrix}$$

where B represents $\bar{\alpha}$ with respect to $(v_{k+1} + U, \dots, v_n + U)$ (Example Sheet 1 Q10). Thus

$$\chi_\alpha = \det(tI - \alpha) = (t - \lambda)^k \det(tI_{n-k} - B) = (t - \lambda)^k \chi_{\bar{\alpha}}.$$

In particular $\chi_{\bar{\alpha}}$ is a product of linear factors and by the induction hypothesis there is a basis $(f_{k+1} + U, \dots, f_n + U)$ for V/U with respect to which $\bar{\alpha}$ is represented by an upper triangular matrix C . Then with respect to $(v_1, \dots, v_k, f_{k+1}, \dots, f_n)$, α is represented by the upper triangular matrix

$$\begin{pmatrix} \lambda I_r & * \\ 0 & C \end{pmatrix}.$$

□

Remark. In general if $\alpha \in \text{End}(V)$ and $U \leq V$ such that $\alpha(U) \leq U$ then $\chi_\alpha = \chi_{\alpha|_U} \chi_{\bar{\alpha}}$ where $\bar{\alpha}: V/U \rightarrow V/U$ is given by $\bar{\alpha}(v + U) = \alpha(v) + U$.

Example. The real matrix

$$\begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}$$

is not similar to an upper triangular matrix over \mathbf{R} for $\theta \notin \pi\mathbf{Z}$ since

$$\chi_A(t) = t^2 - 2\cos \theta + 1$$

has no real roots unless $\cos \theta = \pm 1$. Of course it is similar to a diagonal matrix over \mathbf{C} .

Theorem (Cayley–Hamilton Theorem). *Suppose that V is a f.d. vector space over \mathbf{F} and $\alpha \in \text{End}(V)$. Then $\chi_\alpha(\alpha) = 0$.*

Proof of Cayley–Hamilton when $\mathbf{F} = \mathbf{C}$. Since $\mathbf{F} = \mathbf{C}$ we’ve seen that there is a basis (e_1, \dots, e_n) for V such that α is represented by an upper triangular matrix

$$A = \begin{pmatrix} \lambda_1 & * & * \\ 0 & \ddots & * \\ 0 & 0 & \lambda_n \end{pmatrix}.$$

Then we can compute $\chi_\alpha(t) = \prod_{i=1}^n (t - \lambda_i)$. Let $V_j = \langle e_1, \dots, e_j \rangle$ for $j = 0, \dots, n$ so we have

$$0 = V_0 \subset V_1 \subset \dots \subset V_{n-1} \subset V_n = V$$

with $\dim V_j = j$. Since $\alpha(e_i) = \sum_{k=1}^n A_{ki}e_k = \sum_{k=1}^i A_{ki}e_k$, we see that

$$\alpha(V_j) \subset V_j \text{ for each } j = 0, \dots, n.$$

Moreover $(\alpha - \lambda_j \iota)(e_j) = \sum_{k=1}^{j-1} A_{kj}e_k$ so

$$(\alpha - \lambda_j \iota)(V_j) \subset V_{j-1} \text{ for each } j = 1, \dots, n.$$

Thus we see inductively that $\prod_{i=j}^n (\alpha - \lambda_i \iota)(V_n) \subset V_{j-1}$. In particular

$$\prod_{i=1}^n (\alpha - \lambda_i \iota)(V) \subset V_0 = 0.$$

Thus $\chi_\alpha(\alpha) = 0$ as claimed. \square

Remark. It is straightforward to extend this to the case $\mathbf{F} = \mathbf{R}$: since $\mathbf{R} \subset \mathbf{C}$, if $A \in \text{Mat}_n(\mathbf{R})$ then we can view A as an element of $\text{Mat}_n(\mathbf{C})$ to see that $\chi_A(A) = 0$. But then if $\alpha \in \text{End}(V)$ for any vector space V over \mathbf{R} we can take A to be the matrix representing α over some basis. Then $\chi_\alpha(\alpha) = \chi_A(\alpha)$ is represented by $\chi_A(A)$ and so it zero.

Second proof of Cayley–Hamilton. Let $A \in \text{Mat}_n(\mathbf{F})$ and let $B = tI_n - A$. We can compute that $\text{adj } B$ is an $n \times n$ -matrix with entries elements of $\mathbf{F}[t]$ of degree at most $n - 1$. So we can write

$$\text{adj } B = B_{n-1}t^{n-1} + B_{n-2}t^{n-2} + \dots + B_1t + B_0$$

with each $B_i \in \text{Mat}_n(\mathbf{F})$. Now we know that $B \text{adj } B = (\det B)I_n = \chi_A(t)I_n$. ie

$$(tI_n - A)(B_{n-1}t^{n-1} + B_{n-2}t^{n-2} + \dots + B_1t + B_0) = (t^n + a_{n-1}t^{n-1} + \dots + a_0)I_n$$

where $\chi_A(t) = t^n + a_{n-1}t^{n-1} + \dots + a_0$. Comparing coefficients in t^k for $k = n, \dots, 0$ we see

$$\begin{aligned} B_{n-1} - 0 &= I_n \\ B_{n-2} - AB_{n-1} &= a_{n-1}I_n \\ B_{n-3} - AB_{n-2} &= a_{n-2}I_n \\ &\dots = \dots \\ 0 - AB_0 &= a_0I_n \end{aligned}$$

Thus

$$\begin{aligned} A^n B_{n-1} - 0 &= A^n \\ A^{n-1} B_{n-2} - A^n B_{n-1} &= a_{n-1} A^{n-1} \\ A^{n-2} B_{n-3} - A^{n-1} B_{n-2} &= a_{n-2} A^{n-2} \\ &\dots = \dots \\ 0 - AB_0 &= a_0 I_n \end{aligned}$$

Summing we get $0 = \chi_A(A)$ as required. \square

Lemma. *Let $\alpha \in \text{End}(V)$, $\lambda \in \mathbf{F}$. Then the following are equivalent*

- (a) λ is an eigenvalue of α ;
- (b) λ is a root of $\chi_\alpha(t)$;
- (c) λ is a root of $m_\alpha(t)$.

Proof. We see the equivalence of (a) and (b) in section 6.1

Suppose that λ is an eigenvalue of α . There is some $v \in V$ non-zero such that $\alpha v = \lambda v$. Then for any polynomial $p \in \mathbf{F}[t]$, $p(\alpha)v = p(\lambda)v$ so

$$0 = m_\alpha(\alpha)v = m_\alpha(\lambda)(v).$$

Since $v \neq 0$ it follows that $m_\alpha(\lambda) = 0$. Thus (a) implies (c).

Since $m_\alpha(t)$ is a factor of $\chi_\alpha(t)$ by the Cayley–Hamilton Theorem, we see that (c) implies (b).

Alternatively, we could prove (c) implies (a) directly: suppose that $m_\alpha(\lambda) = 0$. Then $m_\alpha(t) = (t - \lambda)g(t)$ for some $g \in \mathbf{F}[t]$. Since $\deg g < \deg m_\alpha$, $g(\alpha) \neq 0$. Thus there is some $v \in V$ such that $g(\alpha)(v) \neq 0$. But then $(\alpha - \lambda)(g(\alpha)(v)) = m_\alpha(\alpha)(v) = 0$. So $g(\alpha)(v) \neq 0$ is a λ -eigenvector of α and so λ is an eigenvalue of α . \square

Example. What is the minimal polynomial of

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

We can compute $\chi_A(t) = (t - 1)^2(t - 2)$. So by Cayley–Hamilton m_A is a factor of $(t - 1)^2(t - 2)$. Moreover by the lemma it must be a multiple of $(t - 1)(t - 2)$. So m_A is one of $(t - 1)(t - 2)$ and $(t - 1)^2(t - 2)$.

We can compute

$$(A - I)(A - 2I) = \begin{pmatrix} 0 & 0 & -2 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} -1 & 0 & -2 \\ 0 & -1 & 1 \\ 0 & 0 & 0 \end{pmatrix} = 0.$$

Thus $m_A(t) = (t - 1)(t - 2)$. Since this has distinct roots, A is diagonalisable.

6.4. Multiplicities of eigenvalues and Jordan Normal Form.

Definition (Multiplicity of eigenvalues). Suppose that $\alpha \in \text{End}(V)$ and λ is an eigenvalue of α :

- (a) the *algebraic multiplicity* of λ is

$$a_\lambda := \text{the multiplicity of } \lambda \text{ as a root of } \chi_\alpha(t);$$

(b) the *geometric multiplicity* of λ is

$$g_\lambda := \dim E_\alpha(\lambda);$$

(c) another useful number is

$$c_\lambda := \text{the multiplicity of } \lambda \text{ as a root of } m_\alpha(t).$$

LECTURE 17

Lemma. *Let $\alpha \in \text{End}(V)$ and $\lambda \in \mathbf{F}$ an eigenvalue of α . Then*

- (a) $1 \leq g_\lambda \leq a_\lambda$ and
 (b) $1 \leq c_\lambda \leq a_\lambda$.

Proof. (a) By definition if λ is an eigenvalue of α then $E_\alpha(\lambda) \neq 0$ so $g_\lambda \geq 1$. Moreover $\chi_\alpha(t) = (t - \lambda)^{g_\lambda} \chi_{\bar{\alpha}}$ where $\bar{\alpha} \in \text{End}(V/E(\lambda))$. So $a_\lambda \geq g_\lambda$.

(b) We've seen that if λ is an eigenvalue of α then α is a root of $m_\alpha(t)$ so $c_\lambda \geq 1$. Cayley–Hamilton says $m_\alpha(t)$ divides $\chi_\alpha(t)$ so $c_\lambda \leq a_\lambda$. \square

Examples.

$$(1) \text{ If } A = \begin{pmatrix} \lambda & 1 & & 0 \\ & \lambda & \ddots & \\ & & \ddots & 1 \\ 0 & & & \lambda \end{pmatrix} \in \text{Mat}_n(\mathbf{F}) \text{ then } g_\lambda = 1 \text{ and } a_\lambda = c_\lambda = n.$$

(2) If $A = \lambda I_n$ then $g_\lambda = a_\lambda = n$ and $c_\lambda = 1$.

Lemma. *Suppose that $\mathbf{F} = \mathbf{C}$ and $\alpha \in \text{End}(V)$. Then the following are equivalent:*

- (a) α is diagonalisable;
 (b) $a_\lambda = g_\lambda$ for all eigenvalues λ of α ;
 (c) $c_\lambda = 1$ for all eigenvalues λ of α .

Proof. To see that (a) is equivalent to (b) suppose that the distinct eigenvalues of α are $\lambda_1, \dots, \lambda_k$. Then α is diagonalisable if and only if $\dim V = \sum_{i=1}^k \dim E(\lambda_i) = \sum_{i=1}^k g_{\lambda_i}$. But $g_\lambda \leq a_\lambda$ for each eigenvalue λ and $\sum_{i=1}^k a_{\lambda_i} = \deg \chi_\alpha = \dim V$ by the Fundamental Theorem of Algebra. Thus α is diagonalisable if and only if $g_{\lambda_i} = a_{\lambda_i}$ for each $i = 1, \dots, k$.

Since by the Fundamental Theorem of Algebra for any such α , $m_\alpha(t)$ may be written as a product of linear factors, α is diagonalisable if and only if these factors are distinct. This is equivalent to $c_\lambda = 1$ for every eigenvalue λ of α . \square

Definition. We say that a matrix $A \in \text{Mat}_n(\mathbf{C})$ is in *Jordan Normal Form (JNF)* if it is a block diagonal matrix

$$A = \begin{pmatrix} J_{n_1}(\lambda_1) & 0 & 0 & 0 \\ 0 & J_{n_2}(\lambda_2) & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & J_{n_k}(\lambda_k) \end{pmatrix}$$

where $k \geq 1$, $n_1, \dots, n_k \in \mathbb{N}$ such that $\sum_{i=1}^k n_i = n$ and $\lambda_1, \dots, \lambda_k \in \mathbf{C}$ (not necessarily distinct) and $J_m(\lambda) \in \text{Mat}_m(\mathbf{C})$ has the form

$$J_m(\lambda) := \begin{pmatrix} \lambda & 1 & 0 & 0 \\ 0 & \lambda & \ddots & 0 \\ 0 & 0 & \ddots & 1 \\ 0 & 0 & 0 & \lambda \end{pmatrix}.$$

We call the $J_m(\lambda)$ *Jordan blocks*

Note $J_m(\lambda) = \lambda I_m + J_m(0)$.

Theorem (Jordan Normal Form). *Every matrix $A \in \text{Mat}_n(\mathbf{C})$ is similar to a matrix in JNF. Moreover this matrix in JNF is uniquely determined by A up to reordering the Jordan blocks.*

Remarks.

- (1) Of course, we can rephrase this as whenever α is an endomorphism of a f.d. \mathbf{C} -vector space V , there is a basis of V such that α is represented by a matrix in JNF. Moreover, this matrix is uniquely determined by α up to reordering the Jordan blocks.
- (2) Two matrices in JNF that differ only in the ordering of the blocks are similar. A corresponding basis change arises as a reordering of the basis vectors.

Examples.

- (1) Every 2×2 matrix in JNF is of the form $\begin{pmatrix} \lambda & 0 \\ 0 & \mu \end{pmatrix}$ with $\lambda \neq \mu$ or $\begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix}$ or $\begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}$. The minimal polynomials are $(t - \lambda)(t - \mu)$, $(t - \lambda)$ and $(t - \lambda)^2$ respectively. The characteristic polynomials are $(t - \lambda)(t - \mu)$, $(t - \lambda)^2$ and $(t - \lambda)^2$ respectively. Thus we see that the JNF is determined by the minimal polynomial of the matrix in this case (but not by just the characteristic polynomial).
- (2) Suppose now that λ_1, λ_2 and λ_3 are distinct complex numbers. Then every 3×3 matrix in JNF is one of six forms

$$\begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix}, \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_2 \end{pmatrix}, \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 1 \\ 0 & 0 & \lambda_2 \end{pmatrix}$$

$$\begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_1 & 0 \\ 0 & 0 & \lambda_1 \end{pmatrix}, \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_1 & 1 \\ 0 & 0 & \lambda_1 \end{pmatrix} \text{ and } \begin{pmatrix} \lambda_1 & 1 & 0 \\ 0 & \lambda_1 & 1 \\ 0 & 0 & \lambda_1 \end{pmatrix}.$$

The minimal polynomials are $(t - \lambda_1)(t - \lambda_2)(t - \lambda_3)$, $(t - \lambda_1)(t - \lambda_2)$, $(t - \lambda_1)(t - \lambda_2)^2$, $(t - \lambda_1)$, $(t - \lambda_1)^2$ and $(t - \lambda_1)^3$ respectively. The characteristic polynomials are $(t - \lambda_1)(t - \lambda_2)(t - \lambda_3)$, $(t - \lambda_1)(t - \lambda_2)^2$, $(t - \lambda_1)(t - \lambda_2)^2$, $(t - \lambda_1)^3$, $(t - \lambda_1)^3$ and $(t - \lambda_1)^3$ respectively. So in this case the minimal polynomial does not determine the JNF by itself (when the minimal polynomial is of the form $(t - \lambda_1)(t - \lambda_2)$ the JNF must be diagonal but it cannot be determined whether λ_1 or λ_2 occurs twice on the diagonal) but the minimal and characteristic polynomials together do determine the JNF. In general even these two bits of data together don't suffice to determine everything.

We recall that

$$J_n(\lambda) = \begin{pmatrix} \lambda & 1 & 0 & 0 \\ 0 & \lambda & \ddots & 0 \\ 0 & 0 & \ddots & 1 \\ 0 & 0 & 0 & \lambda \end{pmatrix}.$$

Thus if (e_1, \dots, e_n) is the standard basis for \mathbb{C}^n we can compute $(J_n(\lambda) - \lambda I_n)e_1 = 0$ and $(J_n(\lambda) - \lambda I_n)e_i = e_{i-1}$ for $1 < i \leq n$. Thus $(J_n(\lambda) - \lambda I_n)^k$ maps e_1, \dots, e_k to 0 and e_{k+j} to e_j for $1 \leq j \leq n - k$. That is

$$(J_n(\lambda) - \lambda I_n)^k = \begin{pmatrix} 0 & I_{n-k} \\ 0 & 0 \end{pmatrix} \text{ for } k < n$$

and $(J_n(\lambda) - \lambda I_n)^k = 0$ for $k \geq n$.

Thus if $A = J_n(\lambda)$ is a single Jordan block, then $\chi_A(t) = m_A(t) = (t - \lambda)^n$, so λ is the only eigenvalue of A . Moreover $\dim E(\lambda) = 1$. Thus $a_\lambda = c_\lambda = n$ and $g_\lambda = 1$.

Now if A be a block diagonal square matrix; ie

$$A = \begin{pmatrix} A_1 & 0 & 0 & 0 \\ 0 & A_2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & A_k \end{pmatrix}$$

then $\chi_A(t) = \prod_{i=1}^k \chi_{A_i}(t)$. Moreover, if $p \in \mathbf{F}[t]$ then

$$p(A) = \begin{pmatrix} p(A_1) & 0 & 0 & 0 \\ 0 & p(A_2) & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & p(A_k) \end{pmatrix}$$

so $m_A(t)$ is the lowest common multiple of $m_{A_1}(t), \dots, m_{A_k}(t)$.

We also have, by a straightforward computation, $n(p(A)) = \sum_{i=1}^k n(p(A_i))$ for any $p \in \mathbf{F}[t]$.

In general a_λ is the sum of the sizes of the blocks with eigenvalue λ which is the same as the number of λ s on the diagonal. g_λ is the number of blocks with eigenvalue λ and c_λ is the size of the largest block with eigenvalue λ .

Theorem. *If $\alpha \in \text{End}(V)$ and A in JNF represents α with respect to some basis then the number of Jordan blocks $J_n(\lambda)$ of A with eigenvalue λ and size $n \geq r \geq 1$ is given by*

$$|\{\text{Jordan blocks } J_n(\lambda) \text{ in } A \text{ with } n \geq r\}| = n((\alpha - \lambda I)^r) - n((\alpha - \lambda I)^{r-1})$$

Proof. We work blockwise with

$$A = \begin{pmatrix} J_{n_1}(\lambda_1) & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & J_{n_k}(\lambda_k) \end{pmatrix}.$$

We can compute that

$$n((J_m(\lambda) - \lambda I_m)^j) = \min(m, j)$$

and

$$n((J_m(\mu) - \lambda I_m)^j) = 0$$

when $\mu \neq \lambda$.

Adding up for each block we get for $r \geq 1$

$$\begin{aligned} n((\alpha - \lambda I)^r) - n((\alpha - \lambda I)^{r-1}) &= n((A - \lambda I)^r) - n((A - \lambda I)^{r-1}) \\ &= \sum_{\substack{i=1 \\ \lambda_i = \lambda}}^k (\min(r, n_i) - \min(r-1, n_i)) \\ &= |\{1 \leq i \leq k \mid \lambda_i = \lambda, n_i \geq r\}| \\ &= |\{\text{Jordan blocks } J_n(\lambda) \text{ in } A \text{ with } n \geq r\}| \end{aligned}$$

as required. \square

Because these nullities are basis-invariant, it follows that if it exists then the Jordan normal form representing α is unique up to reordering the blocks as claimed.

LECTURE 18

Theorem (Generalised eigenspace decomposition). *Let V be a f.d. \mathbb{C} -vector space and $\alpha \in \text{End}(V)$. Suppose that*

$$m_\alpha(t) = (t - \lambda_1)^{c_1} \cdots (t - \lambda_k)^{c_k}$$

with $\lambda_1, \dots, \lambda_k$ distinct. Then

$$V = V_1 \oplus V_2 \oplus \cdots \oplus V_k$$

where $V_j = \ker((\alpha - \lambda_j I)^{c_j})$ is an α -invariant subspace (called a generalised eigenspace).

Note that in the case $c_1 = c_2 = \cdots = c_k = 1$ we recover the diagonalisability theorem and indeed the proof is a generalisation of that one.

Sketch of proof. Let $p_j(t) = \prod_{i \neq j}^k (t - \lambda_i)^{c_i}$. Then p_1, \dots, p_k have no common factor i.e. they are coprime. Thus by Euclid's algorithm we can find $q_1, \dots, q_k \in \mathbb{C}[t]$ such that $\sum_{i=1}^k q_i p_i = 1$.

Let $\pi_j = q_j(\alpha) p_j(\alpha)$ for $j = 1, \dots, k$. Then $\sum_{j=1}^k \pi_j = I$. Since $m_\alpha(\alpha) = 0$, $(\alpha - \lambda_j I)^{c_j} \pi_j = 0$, thus $\text{Im } \pi_j \subset V_j$.

Suppose that $v \in V$ then

$$v = I(v) = \sum_{j=1}^k \pi_j(v) \in \sum V_j.$$

Thus $V = \sum V_j$.

But $\pi_i \pi_j = 0$ for $i \neq j$ and so $\pi_i = \pi_i (\sum_{j=1}^k \pi_j) = \pi_i^2$ for $1 \leq i \leq k$. Thus $\pi_j|_{V_j} = I_{V_j}$ and if $v = \sum v_j$ with $v_j \in V_j$ then $v_j = \pi_j(v)$. So $V = \bigoplus V_j$ as claimed. \square

Using the generalised eigenspace decomposition theorem we can, by considering the action of α on its generalised eigenspaces separately, reduce the proof of the existence of Jordan normal form for α to the case that it has only one eigenvalue λ . By considering $(\alpha - \lambda I)$ we can even reduce to the case that 0 is the only eigenvalue.

Definition. We say that $\alpha \in \text{End}(V)$ is *nilpotent* if there is some $k \geq 0$ such that $\alpha^k = 0$.

Note that α is nilpotent if and only if $m_\alpha(t) = t^k$ for some $1 \leq k \leq n$. When $\mathbf{F} = \mathbf{C}$ this is equivalent to 0 being the only eigenvalue for α .

Example. Let

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

Find an invertible matrix P such that $P^{-1}AP$ is in JNF.

First we compute the eigenvalues of A :

$$\chi_A(t) = \det \begin{pmatrix} t-3 & 2 & 0 \\ -1 & t & 0 \\ -1 & 0 & t-1 \end{pmatrix} = (t-1)(t(t-3)+2) = (t-1)^2(t-2).$$

Next we compute the eigenspaces

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

which has rank 2 and kernel spanned by $\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$. Thus $E_A(1) = \left\langle \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right\rangle$. Similarly

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

also has rank 1 and kernel spanned by $\begin{pmatrix} 2 \\ 1 \\ 2 \end{pmatrix}$ thus $E_A(2) = \left\langle \begin{pmatrix} 2 \\ 1 \\ 2 \end{pmatrix} \right\rangle$. Since $\dim E_A(1) + \dim E_A(2) = 2 < 3$, A is not diagonalisable. Thus

$$m_A(t) = \chi_A(t) = (t-1)^2(t-2)$$

and the JNF of A is

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

So we want to find a basis (v_1, v_2, v_3) such that $Av_1 = v_1$, $Av_2 = v_1 + v_2$ and $Av_3 = 2v_3$ or equivalently $(A-I)v_2 = v_1$, $(A-I)v_1 = 0$ and $(A-2I)v_3 = 0$. Note that under these conditions $(A-I)^2v_2 = 0$ but $(A-I)v_2 \neq 0$.

We compute

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

Thus

$$\ker(A-I)^2 = \left\langle \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right\rangle$$

Take $v_2 = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}$, $v_1 = (A - I)v_2 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$ and $v_3 = \begin{pmatrix} 2 \\ 1 \\ 2 \end{pmatrix}$. Then these are LI so form a basis for \mathbf{C}^3 and if we take P to have columns v_1, v_2, v_3 we see that $P^{-1}AP = J$ as required.

7. BILINEAR FORMS (II)

7.1. Symmetric bilinear forms and quadratic forms. We return to the study of bilinear forms but we will think particular about bilinear forms $\phi: V \times V \rightarrow \mathbf{F}$ on a single vector space. As for endomorphisms we won't allow the freedom to choose different bases on each side.

Lemma. *Suppose that V is a f.d. vector space over \mathbf{F} , $\phi: V \times V \rightarrow \mathbf{F}$ is a bilinear form and (e_1, \dots, e_n) and (f_1, \dots, f_n) are two bases of V such that $f_i = \sum P_{ki}e_k$ for $i = 1, \dots, n$. If A represents ϕ with respect to (e_1, \dots, e_n) and B represents ϕ with respect to (f_1, \dots, f_n) then*

$$B = P^T A P$$

Proof. This is a special case of a result from section 4 □

Definition. We say that square matrices A and B are *congruent* if there is an invertible matrix P such that $B = P^T A P$.

Congruence is an equivalence relation. Two matrices are congruent precisely if they represent the same bilinear form $\phi: V \times V \rightarrow \mathbf{F}$ with respect to different bases for V . Thus to classify (symmetric) bilinear forms on a f.d. vector space is to classify (symmetric) matrices up to congruence.

Definition. Let V be a vector space over \mathbf{F} . A bilinear form $\phi: V \times V \rightarrow \mathbf{F}$ is *symmetric* if $\phi(v_1, v_2) = \phi(v_2, v_1)$ for all $v_1, v_2 \in V$.

Example. Suppose $S \in \text{Mat}_n(\mathbf{F})$ is a symmetric matrix (ie $S^T = S$), then we can define a symmetric bilinear form $\phi: \mathbf{F}^n \times \mathbf{F}^n \rightarrow \mathbf{F}$ by

$$\phi(x, y) = x^T S y = \sum_{i,j=1}^n x_i S_{ij} y_j$$

In fact that example is completely typical.

Lemma. *Suppose that V is a f.d. vector space over \mathbf{F} and $\phi: V \times V \rightarrow \mathbf{F}$ is a bilinear form. Let (e_1, \dots, e_n) be a basis for V and M be the matrix representing ϕ with respect to this basis, i.e. $M_{ij} = \phi(e_i, e_j)$. Then ϕ is symmetric if and only if M is symmetric.*

Proof. If ϕ is symmetric then $M_{ij} = \phi(e_i, e_j) = \phi(e_j, e_i) = M_{ji}$ so M is symmetric. Conversely if M is symmetric, then

$$\phi(x, y) = \sum_{i,j=1}^n x_i M_{ij} y_j = \sum_{i,j=1}^n y_j M_{ji} x_i = \phi(y, x).$$

Thus ϕ is symmetric. □

It follows that if ϕ is represented by a symmetric matrix with respect to one basis then it is represented by a symmetric matrix with respect to every basis.

LECTURE 19

Definition. If $\phi: V \times V \rightarrow \mathbf{F}$ is a bilinear form then we call the map $V \rightarrow \mathbf{F}; v \mapsto \phi(v, v)$ a *quadratic form* on V .

Examples.

(1) If $V = \mathbf{R}^3$ and ϕ is represented by I with respect to the standard basis then

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \mapsto x^2 + y^2 + z^2 \text{ is the corresponding quadratic form.}$$

(2) If $V = \mathbf{R}^4$ and ϕ is represented by $\begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$ with respect to the stan-

dard basis then $\begin{pmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \end{pmatrix} \mapsto -x_0^2 + x_1^2 + x_2^2 + x_3^2$ is the corresponding quadratic form.

(3) If $V = \mathbf{R}^2$ and ϕ is represented by the matrix A with respect to the standard basis then the corresponding quadratic form is

$$\begin{pmatrix} x \\ y \end{pmatrix} \mapsto (x \ y) A \begin{pmatrix} x \\ y \end{pmatrix} = A_{11}x^2 + (A_{12} + A_{21})xy + A_{22}y^2$$

Note that if we replace A by the symmetric matrix $\frac{1}{2}(A + A^T)$ we get the same quadratic form.

Proposition (Polarisation identity). (*Suppose that $1+1 \neq 0$ in \mathbf{F} .*) If $q: V \rightarrow \mathbf{F}$ is a quadratic form then there exists a unique symmetric bilinear form $\phi: V \times V \rightarrow \mathbf{F}$ such that $q(v) = \phi(v, v)$ for all $v \in V$.

Proof. Let ψ be a bilinear form on $V \times V$ such that $\psi(v, v) = q(v)$ for all $v \in V$. Then

$$\phi(v, w) := \frac{1}{2}(\psi(v, w) + \psi(w, v))$$

is a symmetric bilinear form such that $\phi(v, v) = q(v)$ for all $v \in V$.

It remains to prove uniqueness. Suppose that ϕ is such a symmetric bilinear form. Then for $v, w \in V$,

$$\begin{aligned} q(v+w) &= \phi(v+w, v+w) \\ &= \phi(v, v) + \phi(v, w) + \phi(w, v) + \phi(w, w) \\ &= q(v) + 2\phi(v, w) + q(w). \end{aligned}$$

Thus $\phi(v, w) = \frac{1}{2}(q(v+w) - q(v) - q(w))$. □

Theorem (Diagonal form for symmetric bilinear forms). (*Suppose that $1+1 \neq 0 \in \mathbf{F}$*) If $\phi: V \times V \rightarrow \mathbf{F}$ is a symmetric bilinear form on a f.d. vector space V over \mathbf{F} , then there is a basis (e_1, \dots, e_n) for V such that ϕ is represented by a diagonal matrix.

Proof. By induction on $n = \dim V$. If $n = 0, 1$ the result is clear. Suppose that we have proven the result for all spaces of dimension strictly smaller than n .

If $\phi(v, v) = 0$ for all $v \in V$, then by the polarisation identity ϕ is identically zero and is represented by the zero matrix with respect to every basis. Otherwise, we can choose $e_1 \in V$ such that $\phi(e_1, e_1) \neq 0$. Let

$$U = \{u \in V \mid \phi(e_1, u) = 0\} = \ker \phi(e_1, -): V \rightarrow \mathbf{F}.$$

By the rank-nullity theorem, U has dimension $n - 1$. Since also $e_1 \notin U$, U is a complement to $\langle e_1 \rangle$ in V .

Consider $\phi|_{U \times U}: U \times U \rightarrow \mathbf{F}$, a symmetric bilinear form on U . By the induction hypothesis, there is a basis (e_2, \dots, e_n) for U such that $\phi|_{U \times U}$ is represented by a diagonal matrix. The basis (e_1, \dots, e_n) satisfies $\phi(e_i, e_j) = 0$ for $i \neq j$ and we're done. \square

Example. Let q be the quadratic form on \mathbf{R}^3 given by

$$q \left(\begin{pmatrix} x \\ y \\ z \end{pmatrix} \right) = x^2 + y^2 + z^2 + 2xy + 4yz + 6xz.$$

Find a basis (f_1, f_2, f_3) for \mathbf{R}^3 such that q is of the form

$$q(af_1 + bf_2 + cf_3) = \lambda a^2 + \mu b^2 + \nu c^2.$$

Method 1 Let ϕ be the bilinear form represented by the matrix

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

so that $q(v) = \phi(v, v)$ for all $v \in \mathbf{R}^3$.

Now $q(e_1) = 1 \neq 0$ so let $f_1 = e_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$. Then $\phi(f_1, v) = f_1^T A v = v_1 + v_2 + 3v_3$.

So we choose f_2 such that $\phi(f_1, f_2) = 0$ but $\phi(f_2, f_2) \neq 0$. For example

$$q \left(\begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix} \right) = 0 \text{ but } q \left(\begin{pmatrix} 3 \\ 0 \\ -1 \end{pmatrix} \right) = -8 \neq 0.$$

So we can take $f_2 = \begin{pmatrix} 3 \\ 0 \\ -1 \end{pmatrix}$. Then $\phi(f_2, v) = f_2^T A v = (0 \ 1 \ 8) v = v_2 + 8v_3$.

Now we want $\phi(f_1, f_3) = \phi(f_2, f_3) = 0$, $f_3 = (5 \ -8 \ 1)^T$ will work. Then

$$q(af_1 + bf_2 + cf_3) = a^2 + (-8)b^2 + 8c^2.$$

Method 2 Complete the square

$$\begin{aligned} x^2 + y^2 + z^2 + 2xy + 4yz + 6xz &= (x + y + 3z)^2 + (-2yz) - 8z^2 \\ &= (x + y + 3z)^2 - 8 \left(z + \frac{y}{8} \right)^2 + \frac{y^2}{8} \end{aligned}$$

Now solve $x + y + 3z = 1$, $z + \frac{y}{8} = 0$ and $y = 0$ to obtain $f_1 = (1 \ 0 \ 0)^T$, solve $x + y + 3z = 0$, $z + \frac{y}{8} = 1$ and $y = 0$ to obtain $f_2 = (-3 \ 0 \ 1)^T$ and solve $x + y + 3z = 0$, $z + \frac{y}{8} = 0$ and $y = 1$ to obtain $f_3 = (-\frac{5}{8} \ 1 \ -\frac{1}{8})^T$.

Corollary. Let ϕ be a symmetric bilinear form on a f.d \mathbf{C} -vector space V . Then there is a basis (v_1, \dots, v_n) for V such that ϕ is represented by a matrix of the form

$$\begin{pmatrix} I_r & 0 \\ 0 & 0 \end{pmatrix}$$

with $r = r(\phi)$ or equivalently such that the corresponding quadratic form q is given by $q(\sum_{i=1}^n a_i v_i) = \sum_{i=1}^r a_i^2$.

Proof. We have already shown that there is a basis (e_1, \dots, e_n) such that $\phi(e_i, e_j) = \delta_{ij} \lambda_j$ for some $\lambda_1, \dots, \lambda_n \in \mathbf{C}$. By reordering the e_i we can assume that $\lambda_i \neq 0$ for $1 \leq i \leq r$ and $\lambda_i = 0$ for $i > r$. Since we're working over \mathbf{C} for each $1 \leq i \leq r$, λ_i has a non-zero square root μ_i , say. Defining $v_i = \frac{1}{\mu_i} e_i$ for $1 \leq i \leq r$ and $v_i = e_i$ for $r + 1 \leq i \leq n$, we see that $\phi(v_i, v_j) = 0$ if $i \neq j$ or $i = j > r$ and $\phi(v_i, v_i) = 1$ if $1 \leq i \leq r$ as required. \square

Corollary. Every symmetric matrix in $\text{Mat}_n(\mathbf{C})$ is congruent to a unique matrix of the form

$$\begin{pmatrix} I_r & 0 \\ 0 & 0 \end{pmatrix}.$$

\square

LECTURE 20

Corollary. Let ϕ be a symmetric bilinear form on a f.d \mathbf{R} -vector space V . Then there is a basis (v_1, \dots, v_n) for V such that ϕ is represented by a matrix of the form

$$\begin{pmatrix} I_p & 0 & 0 \\ 0 & -I_q & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

with $p, q \geq 0$ and $p + q = r(\phi)$ or equivalently such that the corresponding quadratic form q is given by $q(\sum_{i=1}^n a_i v_i) = \sum_{i=1}^p a_i^2 - \sum_{i=p+1}^{p+q} a_i^2$.

Proof. We have already shown that there is a basis (e_1, \dots, e_n) such that $\phi(e_i, e_j) = \delta_{ij} \lambda_j$ for some $\lambda_1, \dots, \lambda_n \in \mathbf{R}$. By reordering the e_i we can assume that there is a p such that $\lambda_i > 0$ for $1 \leq i \leq p$ and $\lambda_i < 0$ for $p + 1 \leq i \leq r(\phi)$ and $\lambda_i = 0$ for $i > r(\phi)$. Since we're working over \mathbf{R} we can define $\mu_i = \sqrt{\lambda_i}$ for $1 \leq i \leq p$, $\mu_i = \sqrt{-\lambda_i}$ for $p + 1 \leq i \leq r(\phi)$ and $\mu_i = 1$ for $i > r(\phi)$. Defining $v_i = \frac{1}{\mu_i} e_i$ we see that ϕ is represented by the given matrix with respect to (v_1, \dots, v_n) . \square

Definition. A symmetric bilinear form ϕ on a real vector space V is

- (a) *positive definite* if $\phi(v, v) > 0$ for all $v \in V \setminus \{0\}$;
- (b) *positive semi-definite* if $\phi(v, v) \geq 0$ for all $v \in V$;
- (c) *negative definite* if $\phi(v, v) < 0$ for all $v \in V \setminus \{0\}$;
- (d) *negative semi-definite* if $\phi(v, v) \leq 0$ for all $v \in V$.

We say a quadratic form is *...-definite* if the corresponding bilinear form is so.

Examples. If ϕ is a symmetric bilinear form on \mathbf{R}^n represented by

$$\begin{pmatrix} I_p & 0 \\ 0 & 0 \end{pmatrix}$$

then ϕ is positive semi-definite. Moreover ϕ is positive definite if and only if $n = p$.

If instead ϕ is represented by

$$\begin{pmatrix} -I_q & 0 \\ 0 & 0 \end{pmatrix}$$

then ϕ is negative semi-definite. Moreover ϕ is negative definite if and only if $n = q$.

Theorem (Sylvester's Law of Inertia). *Let V be a finite-dimensional real vector space and let ϕ be a symmetric bilinear form on V . Then there are unique integers p, q such that V has a basis v_1, \dots, v_n with respect to which ϕ is represented by the matrix*

$$\begin{pmatrix} I_p & 0 & 0 \\ 0 & -I_q & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

Proof. We've already done the existence part. We also already know that $p + q = r(\phi)$ is unique. To see p is unique we'll prove that p is the largest dimension of a subspace P of V such that $\phi|_{P \times P}$ is positive definite.

Let v_1, \dots, v_n be some basis with respect to which ϕ is represented by

$$\begin{pmatrix} I_p & 0 & 0 \\ 0 & -I_q & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

for some choice of p . Then ϕ is positive definite on the space spanned by v_1, \dots, v_p . Thus it remains to prove that there is no larger such subspace.

Let P be any subspace of V such that $\phi|_{P \times P}$ is positive definite and let Q be the space spanned by v_{p+1}, \dots, v_n . The restriction of ϕ to $Q \times Q$ is negative semi-definite so $P \cap Q = 0$. So $\dim P + \dim Q = \dim(P + Q) \leq n$. Thus $\dim P \leq p$ as required. \square

Definition. The *signature* of the symmetric bilinear form ϕ given in the Theorem is defined to be $p - q$.

Corollary. *Every real symmetric matrix is congruent to a unique matrix of the form*

$$\begin{pmatrix} I_p & 0 & 0 \\ 0 & -I_q & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

\square

7.2. Hermitian forms. Recall that the standard inner product of \mathbb{C}^n is given by

$$\langle x, y \rangle = \sum_{i=1}^n \bar{x}_i y_i.$$

This is not a bilinear form since it is not linear in the first variable. However it is linear 'up to complex conjugation'. This motivates the following definition.

Definition. Let V and W be vector spaces over \mathbf{C} . Then a *sesquilinear form* is a function $\phi: V \times W \rightarrow \mathbf{C}$ such that

$$\begin{aligned} \phi(\lambda_1 v_1 + \lambda_2 v_2, w) &= \bar{\lambda}_1 \phi(v_1, w) + \bar{\lambda}_2 \phi(v_2, w) \text{ and} \\ \phi(v, \mu_1 w_1 + \mu_2 w_2) &= \mu_1 \phi(v, w_1) + \mu_2 \phi(v, w_2) \end{aligned}$$

for all $\lambda_1, \lambda_2, \mu_1, \mu_2 \in \mathbf{C}$, $v, v_1, v_2 \in V$ and $w, w_1, w_2 \in W$.

Definition. Let ϕ be a sesquilinear form on $V \times W$ and let V have basis (v_1, \dots, v_m) and W have basis (w_1, \dots, w_n) . The *matrix A representing ϕ* with respect to these bases is defined by $A_{ij} = \phi(v_i, w_j)$.

Suppose that $\sum \lambda_i v_i \in V$ and $\sum \mu_j w_j \in W$ then

$$\phi\left(\sum \lambda_i v_i, \sum \mu_j w_j\right) = \sum_{i=1}^m \sum_{j=1}^n \bar{\lambda}_i A_{ij} \mu_j = \bar{\lambda}^T A \mu.$$

Definition. A sesquilinear form $\phi: V \times V \rightarrow \mathbf{C}$ is said to be *Hermitian* if $\phi(x, y) = \overline{\phi(y, x)}$ for all $x, y \in V$.

Notice that if ϕ is a Hermitian form on V then $\phi(v, v) \in \mathbf{R}$ for all $v \in V$ and $\phi(\lambda v, \lambda v) = |\lambda|^2 \phi(v, v)$ for all $\lambda \in \mathbf{C}$. Thus is it meaningful to speak of positive/negative (semi)-definite Hermitian forms and we will do so.

Lemma. Let $\phi: V \times V \rightarrow \mathbf{C}$ be a sesquilinear form on a complex vector space V with basis (v_1, \dots, v_n) . Then ϕ is Hermitian if and only if the matrix A representing ϕ with respect to this basis satisfies $A = \overline{A}^T =: A^\dagger$ (we also say the matrix A is Hermitian).

Proof. If ϕ is Hermitian then

$$A_{ij} = \phi(v_i, v_j) = \overline{\phi(v_j, v_i)} = \overline{A_{ji}}.$$

Conversely if $A = A^\dagger$ then

$$\phi\left(\sum \lambda_i v_i, \sum \mu_j v_j\right) = \bar{\lambda}^T A \mu = \mu^T A^T \bar{\lambda} = \overline{\mu^T A \lambda} = \overline{\phi\left(\sum \mu_j v_j, \sum \lambda_i v_i\right)}$$

as required □

Proposition (Change of basis). Suppose that ϕ is a Hermitian form on a f.d. complex vector space V and that (e_1, \dots, e_n) and (v_1, \dots, v_n) are bases for V such that $v_i = \sum_{k=1}^n P_{ki} e_k$. Let A be the matrix representing ϕ with respect to (e_1, \dots, e_n) and B be the matrix representing ϕ with respect to (v_1, \dots, v_n) then

$$B = P^\dagger A P.$$

Proof. We compute

$$B_{ij} = \phi\left(\sum_{k=1}^n P_{ki} e_k, \sum_{l=1}^n P_{lj} e_l\right) = \sum_{k,l} (\overline{P}^T)_{ik} \phi(e_k, e_l) P_{lj} = [\overline{P}^T A P]_{ij}$$

as required. □

LECTURE 21

Lemma (Polarisation Identity). A Hermitian form ϕ on a complex vector space V is determined by the function $\psi: V \rightarrow \mathbf{R}; v \mapsto \phi(v, v)$.

Proof. It can be checked that

$$\phi(x, y) = \frac{1}{4} (\psi(x+y) - i\psi(x+iy) - \psi(x-y) + i\psi(x-iy))$$

□

Theorem (Hermitian version of Sylvester’s Law of Inertia). *Let V be a f.d. complex vector space and suppose that $\phi: V \times V \rightarrow \mathbf{C}$ is a Hermitian form on V . Then there is a basis (v_1, \dots, v_n) of V with respect to which ϕ is represented by a matrix of the form*

$$\begin{pmatrix} I_p & 0 & 0 \\ 0 & -I_q & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

Moreover p and q depend only on ϕ not on the basis.

Notice that for such a basis $\phi(\sum \lambda_i v_i, \sum \lambda_i v_i) = \sum_{i=1}^p |\lambda_i|^2 - \sum_{j=p+1}^{p+q} |\lambda_j|^2$.

Sketch of Proof. This is nearly identical to the real case. For existence: if ϕ is identically zero then any basis will do. If not, then by the Polarisation Identity there is some $v_1 \in V$ such that $\phi(v_1, v_1) \neq 0$. By replacing v_1 by $\frac{v_1}{|\phi(v_1, v_1)|^{1/2}}$ we can assume that $\phi(v_1, v_1) = \pm 1$. Define $U := \ker \phi(v_1, -): V \rightarrow \mathbf{C}$ a subspace of V of dimension $\dim V - 1$. Since $v_1 \notin U$, U is a complement to the span of v_1 in V . By induction on $\dim V$, there is a basis (v_2, \dots, v_n) of U such that $\phi|_{U \times U}$ is represented by a matrix of the required form. Now (v_1, v_2, \dots, v_n) is a basis for V that after suitable reordering works.

For uniqueness: $p + q$ is the rank of the matrix representing ϕ with respect to any basis and p arises as the dimension of a maximal positive definite subspace as in the real symmetric case. \square

8. INNER PRODUCT SPACES

From now on \mathbf{F} will always denote \mathbf{R} or \mathbf{C} .

8.1. Definitions and basic properties.

Definition. Let V be a vector space over \mathbf{F} . An *inner product* on V is a positive definite symmetric/Hermitian form ϕ on V . Usually instead of writing $\phi(x, y)$ we’ll write $\langle x, y \rangle$. A vector space equipped with an inner product $\langle -, - \rangle$ is called an *inner product space*.

Examples.

- (1) The usual scalar product on \mathbf{R}^n or \mathbf{C}^n : $\langle x, y \rangle = \sum_{i=1}^n \bar{x}_i y_i$.
- (2) Let $C([0, 1], \mathbf{F})$ be the space of continuous real/complex valued functions on $[0, 1]$ and define

$$\langle f, g \rangle = \int_0^1 \overline{f(t)} g(t) dt.$$

- (3) A weighted version of (2). Let $w: [0, 1] \rightarrow \mathbf{R}$ take only positive values and define

$$\langle f, g \rangle = \int_0^1 w(t) \overline{f(t)} g(t) dt.$$

If V is an inner product space then we can define a norm $\| \cdot \|$ on V by $\|v\| = \langle v, v \rangle^{1/2}$. Note $\|v\| \geq 0$ with equality if and only if $v = 0$. Note that the norm determines the inner product because of the polarisation identity.

Lemma (Cauchy–Schwarz inequality). *Let V be an inner product space and take $v, w \in V$. Then $|\langle v, w \rangle| \leq \|v\| \|w\|$.*

Proof. Since $\langle -, - \rangle$ is positive-definite,

$$0 \leq \langle v - \lambda w, v - \lambda w \rangle = \langle v, v \rangle - \lambda \langle v, w \rangle - \bar{\lambda} \langle w, v \rangle + |\lambda|^2 \langle w, w \rangle$$

for all $\lambda \in \mathbf{F}$. Now when $\lambda = \frac{\langle w, v \rangle}{\langle w, w \rangle}$ (the case $w = 0$ is clear) then we get

$$0 \leq \langle v, v \rangle - \frac{2|\langle v, w \rangle|^2}{\langle w, w \rangle} + \frac{|\langle v, w \rangle|^2}{\langle w, w \rangle^2} \langle w, w \rangle = \langle v, v \rangle - \frac{|\langle v, w \rangle|^2}{\langle w, w \rangle}.$$

The inequality follows by multiplying by $\langle w, w \rangle$ rearranging and taking square roots. \square

Corollary (Triangle inequality). *Let V be an inner product space and take $v, w \in V$. Then $\|v + w\| \leq \|v\| + \|w\|$.*

Proof.

$$\begin{aligned} \|v + w\|^2 &= \langle v + w, v + w \rangle \\ &= \|v\|^2 + \langle v, w \rangle + \langle w, v \rangle + \|w\|^2 \\ &\leq \|v\|^2 + 2\|v\|\|w\| + \|w\|^2 \\ &= (\|v\| + \|w\|)^2 \end{aligned}$$

Taking square roots gives the result. \square

Definition. Let V be an inner product space. Then $v, w \in V$ are said to be *orthogonal* if $\langle v, w \rangle = 0$. A set $\{v_i \mid i \in I\}$ is *orthonormal* if $\langle v_i, v_j \rangle = \delta_{ij}$ for $i, j \in I$. An *orthormal basis* (*o.n. basis*) for V is a basis for V that is orthonormal.

Suppose that V is a f.d. inner product space with o.n. basis v_1, \dots, v_n . Then given $v \in V$, we can write $v = \sum_{i=1}^n \lambda_i v_i$. But then $\langle v_j, v \rangle = \sum_{i=1}^n \lambda_i \langle v_j, v_i \rangle = \lambda_j$. Thus $v = \sum_{i=1}^n \langle v_i, v \rangle v_i$.

Lemma (Parseval's identity). *Suppose that V is a f.d. inner product space with o.n. basis (v_1, \dots, v_n) then $\langle v, w \rangle = \sum_{i=1}^n \overline{\langle v_i, v \rangle} \langle v_i, w \rangle$. In particular*

$$\|v\|^2 = \sum_{i=1}^n |\langle v_i, v \rangle|^2.$$

Proof. $\langle v, w \rangle = \langle \sum_{i=1}^n \langle v_i, v \rangle v_i, \sum_{j=1}^n \langle v_j, w \rangle v_j \rangle = \sum_{i=1}^n \overline{\langle v_i, v \rangle} \langle v_i, w \rangle$. \square

LECTURE 22

8.2. Gram–Schmidt orthogonalisation.

Theorem (Gram-Schmidt process). *Let V be an inner product space and e_1, e_2, \dots be LI vectors. Then there is a sequence v_1, v_2, \dots of orthonormal vectors such that $\langle e_1, \dots, e_k \rangle = \langle v_1, \dots, v_k \rangle$ for each $k \geq 0$.*

Proof. We proceed by induction on k . The case $k = 0$ is clear. Suppose we've found v_1, \dots, v_k . Let

$$u_{k+1} = e_{k+1} - \sum_{i=1}^k \langle v_i, e_{k+1} \rangle v_i.$$

Then for $j \leq k$,

$$\langle v_j, u_{k+1} \rangle = \langle v_j, e_{k+1} \rangle - \sum_{i=1}^k \langle v_i, e_{k+1} \rangle \langle v_j, v_i \rangle = 0.$$

Since $\langle v_1, \dots, v_k \rangle = \langle e_1, \dots, e_k \rangle$, and e_1, \dots, e_{k+1} are LI, $\{v_1, \dots, v_k, e_{k+1}\}$ are LI and so $u_{k+1} \neq 0$. Let $v_{k+1} = \frac{u_{k+1}}{\|u_{k+1}\|}$. \square

Corollary. *Let V be a f.d. inner product space. Then any orthonormal sequence v_1, \dots, v_k can be extended to an orthonormal basis.*

Proof. Let $v_1, \dots, v_k, x_{k+1}, \dots, x_n$ be any basis of V extending v_1, \dots, v_k . If we apply the Gram–Schmidt process to this basis we obtain an o.n. basis w_1, \dots, w_n . Moreover one can check that $w_i = v_i$ for $1 \leq i \leq k$. \square

Definition. Let V be an inner product space and let V_1, V_2 be subspaces of V . Then V is the *orthogonal (internal) direct sum* of V_1 and V_2 , written $V = V_1 \perp V_2$, if

- (i) $V = V_1 + V_2$;
- (ii) $V_1 \cap V_2 = 0$;
- (iii) $\langle v_1, v_2 \rangle = 0$ for all $v_1 \in V_1$ and $v_2 \in V_2$.

Note that condition (iii) implies condition (ii).

Definition. If $W \subset V$ is a subspace of an inner product space V then the *orthogonal complement* of W in V , written W^\perp , is the subspace of V

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

Proposition. *Let V be a f.d. inner product space and W a subspace of V . Then $V = W \perp W^\perp$.*

Proof. Of course if $w \in W$ and $w^\perp \in W^\perp$ then $\langle w, w^\perp \rangle = 0$. So it remains to show that $V = W + W^\perp$. Let w_1, \dots, w_k be an o.n. basis of W . For $v \in V$ and $1 \leq j \leq k$,

$$\langle w_j, v - \sum_{i=1}^k \langle w_i, v \rangle w_i \rangle = \langle w_j, v \rangle - \sum_{i=1}^k \langle w_i, v \rangle \delta_{ij} = 0.$$

Thus $\langle \sum_{j=1}^k \lambda_j w_j, v - \sum_{i=1}^k \langle w_i, v \rangle w_i \rangle = 0$ for all $\lambda_1, \dots, \lambda_k \in \mathbf{F}$ and so

$$v - \sum_{i=1}^k \langle w_i, v \rangle w_i \in W^\perp$$

which suffices. \square

Notice that unlike general vector space complements, orthogonal complements are unique.

Definition. We can also define the *orthogonal (external) direct sum* of two inner product spaces V_1 and V_2 by endowing the vector space direct sum $V_1 \oplus V_2$ with the inner product

$$\langle (v_1, v_2), (w_1, w_2) \rangle = \langle v_1, w_1 \rangle + \langle v_2, w_2 \rangle$$

for $v_1, w_1 \in V_1$ and $v_2, w_2 \in V_2$.

Definition. Suppose that $V = U \oplus W$. Then there is linear map $\pi: V \rightarrow W$ given by $\pi(u + w) = w$ for $u \in U$ and $w \in W$. We call π a *projection map* onto W along U . If $U = W^\perp$ we call π the *orthogonal projection* onto W .

Proposition. Let V be a f.d. inner product space and $W \subset V$ be a subspace with o.n. basis (e_1, \dots, e_k) . Let π be the orthogonal projection onto W . Then

- (a) $\pi(v) = \sum_{i=1}^k \langle e_i, v \rangle e_i$ for each $v \in V$;
 (b) $\|v - \pi(v)\| \leq \|v - w\|$ for all $w \in W$ with equality if and only if $\pi(v) = w$; that is $\pi(v)$ is the closest point to v in W .

Proof. (a) Put $w = \sum_{i=1}^k \langle e_i, v \rangle e_i \in W$. Then

$$\langle e_j, v - w \rangle = \langle e_j, v \rangle - \sum_{i=1}^k \langle e_i, v \rangle \langle e_j, e_i \rangle = 0 \text{ for } 1 \leq j \leq k.$$

Thus $v - w \in W^\perp$. Now $v = w + (v - w)$ so $\pi(v) = w$.

(b) If $x, y \in V$ are orthogonal then

$$\|x + y\|^2 = \langle x + y, x + y \rangle = \|x\|^2 + \langle x, y \rangle + \langle y, x \rangle + \|y\|^2 = \|x\|^2 + \|y\|^2$$

so

$$\|v - w\|^2 = \|(v - \pi(v)) + (\pi(v) - w)\|^2 = \|v - \pi(v)\|^2 + \|(\pi(v) - w)\|^2$$

and $\|v - w\|^2 \geq \|v - \pi(v)\|^2$ with equality if and only if $\|\pi(v) - w\|^2 = 0$ ie $\pi(v) = w$. \square

8.3. Adjoints.

Lemma. Suppose V and W are f.d. inner product spaces and $\alpha: V \rightarrow W$ is linear. Then there is a unique linear map $\alpha^*: W \rightarrow V$ such that $\langle \alpha(v), w \rangle = \langle v, \alpha^*(w) \rangle$ for all $v \in V$ and $w \in W$.

Proof. Let (v_1, \dots, v_m) be an o.n. basis for V and (w_1, \dots, w_n) be an o.n. basis for W and suppose that α is represented by the matrix A with respect to these bases. Then if $\alpha^*: W \rightarrow V$ satisfies $\langle \alpha(v), w \rangle = \langle v, \alpha^*(w) \rangle$ for all $v \in V$ and $w \in W$, we can compute

$$\langle v_i, \alpha^*(w_j) \rangle = \langle \alpha(v_i), w_j \rangle = \left\langle \sum_k A_{ki} w_k, w_j \right\rangle = \overline{A_{ji}}.$$

Thus $\alpha^*(w_j) = \sum_k \overline{A_{kj}} v_k$ ie α^* is represented by the matrix A^\dagger . In particular α^* is unique if it exists.

But to prove existence we can now take α^* to be the linear map represented by the matrix A^\dagger . Then

$$\begin{aligned} \left\langle \alpha \left(\sum_i \lambda_i v_i \right), \sum_j \mu_j w_j \right\rangle &= \sum_{i,j} \overline{\lambda_i} \mu_j \left\langle \sum_k A_{ki} w_k, w_j \right\rangle \\ &= \sum_{i,j} \overline{\lambda_i} \overline{A_{ji}} \mu_j \end{aligned}$$

whereas

$$\begin{aligned} \left\langle \sum_i \lambda_i v_i, \sum_j \alpha^*(\mu_j w_j) \right\rangle &= \sum_{i,j} \overline{\lambda_i} \mu_j \left\langle v_i, \sum_l \overline{A^T}_{lj} v_l \right\rangle \\ &= \sum_{i,j} \overline{\lambda_i A_{ji}} \mu_j \end{aligned}$$

Thus $\langle \alpha(v), w \rangle = \langle v, \alpha^*(w) \rangle$ for all $v \in V$ and $w \in W$ as required. \square

Definition. We call the linear map α^* characterised by the lemma the *adjoint* of α .

LECTURE 23

We've seen that if α is represented by A with respect to some o.n. bases then α^* is represented by A^\dagger with respect to the same bases.

Definition. Suppose that V is an inner product space. Then $\alpha \in \text{End}(V)$ is *self-adjoint* if $\alpha^* = \alpha$; i.e. if $\langle \alpha(v), w \rangle = \langle v, \alpha(w) \rangle$ for all $v, w \in V$.

Thus if $V = \mathbf{R}^n$ with the standard inner product then a matrix $A \in \text{Mat}_n(\mathbf{R})$ is self-adjoint if and only if $A = A^T$; ie A is symmetric. If $V = \mathbf{C}^n$ with the standard inner product then a matrix $A \in \text{Mat}_n(\mathbf{C})$ is self-adjoint if and only if $A = A^\dagger$; ie A is Hermitian.

Definition. If V is a real (resp. complex) inner product space then we say that $\alpha \in \text{End}(V)$ is *orthogonal* (resp. *unitary*) if

$$\langle \alpha(v), \alpha(w) \rangle = \langle v, w \rangle \text{ for all } v, w \in V.$$

By the polarisation identity α is orthogonal (resp. unitary) if and only if $\|\alpha(v)\| = \|v\|$ for all $v \in V$.

Lemma. Suppose that V is a f.d. real (resp. complex) inner product space. Let $\alpha \in \text{End}(V)$. Then α is orthogonal (resp. unitary) if and only if α is invertible and $\alpha^* = \alpha^{-1}$.

Proof. If $\alpha^* = \alpha^{-1}$ then $\langle v, v \rangle = \langle v, \alpha^* \alpha(v) \rangle = \langle \alpha(v), \alpha(v) \rangle$ for all $v \in V$ ie α is orthogonal.

Conversely, suppose $\langle \alpha(v), \alpha(w) \rangle = \langle v, w \rangle$ for all $v, w \in V$ and let v_1, \dots, v_n be an o.n. basis for V . Then for each $1 \leq i, j \leq n$,

$$\delta_{ij} = \langle v_i, v_j \rangle = \langle \alpha(v_i), \alpha(v_j) \rangle = \langle v_i, \alpha^* \alpha(v_j) \rangle.$$

and $\alpha^* \alpha(v_j) = v_j$ as required. \square

Note that a real (resp. complex) square matrix is orthogonal (resp. unitary) as an endomorphism of $V = \mathbf{R}^n$ (resp. \mathbf{C}^n) with the standard inner product if and only if $A^T A = I$ (resp. $A^\dagger A = I$) or equivalently if and only if the columns of A form an o.n. basis.

Corollary. With notation as in the lemma, $\alpha \in \text{End}(V)$ is orthogonal (resp. unitary) if and only if α is represented by an orthogonal (resp. unitary) matrix with respect to any orthonormal basis.

Proof. Let (v_1, \dots, v_n) be an o.n. basis then α is represented by A with respect to this basis if and only if α^* is represented by A^T (resp. A^\dagger). Thus α is orthogonal (resp. unitary) if and only if A is invertible with inverse A^T (resp. A^\dagger) i.e. A is orthogonal (resp. unitary). \square

Definition.

- If V is a f.d. real inner product space then

$$O(V) := \{\alpha \in \text{End}(V) \mid \alpha \text{ is orthogonal}\}$$

forms a group under composition called the *orthogonal group* of V .

- If V is a f.d. complex inner product space then

$$U(V) := \{\alpha \in \text{End}(V) \mid \alpha \text{ is unitary}\}$$

forms a group under composition called the *unitary group* of V .

Proposition. *Suppose that V is an inner product space with o.n. basis (e_1, \dots, e_n) .*

(a) *If $\mathbf{F} = \mathbf{R}$ there is a natural bijection*

$$O(V) \longrightarrow \{\text{o.n. bases of } V\}$$

given by

$$\alpha \mapsto (\alpha(e_1), \dots, \alpha(e_n)).$$

(b) *If $\mathbf{F} = \mathbf{C}$ there is a natural bijection*

$$U(V) \longrightarrow \{\text{o.n. bases of } V\}$$

given by

$$\alpha \mapsto (\alpha(e_1), \dots, \alpha(e_n)).$$

\square

8.4. Spectral theory.

Lemma. *Suppose that V is an inner product space and $\alpha \in \text{End}(V)$ is self-adjoint then*

- (a) *α has a real eigenvalue;*
- (b) *all eigenvalues of α are real;*
- (c) *eigenvectors of α with distinct eigenvalues are orthogonal.*

Proof. (a) and (b) Suppose first that V is a complex inner product space. By the fundamental theorem of algebra α has an eigenvalue (since the minimal polynomial has a root). Suppose that $\alpha(v) = \lambda v$ with $v \in V \setminus \{0\}$ and $\lambda \in \mathbf{C}$. Then

$$\lambda \langle v, v \rangle = \langle v, \lambda v \rangle = \langle v, \alpha(v) \rangle = \langle \alpha(v), v \rangle = \langle \lambda v, v \rangle = \bar{\lambda} \langle v, v \rangle.$$

Since $\langle v, v \rangle \neq 0$ we can deduce $\lambda \in \mathbf{R}$.

Now, suppose that V is a real inner product space. Let (v_1, \dots, v_n) be an o.n. basis. Then α is represented by a real symmetric matrix A . But A viewed as a complex matrix is also Hermitian so all its eigenvalues are real by the above. Finally, the eigenvalues of α are precisely the eigenvalues of A .

- (c) Suppose $\alpha(v) = \lambda v$ and $\alpha(w) = \mu w$ with $\lambda \neq \mu \in \mathbf{R}$. Then

$$\lambda \langle v, w \rangle = \langle \lambda v, w \rangle = \langle \alpha(v), w \rangle = \langle v, \alpha(w) \rangle = \langle v, \mu(w) \rangle = \mu \langle v, w \rangle.$$

Since $\lambda \neq \mu$ we must have $\langle v, w \rangle = 0$. \square

Theorem. *Let V be a finite dimensional inner product space and $\alpha \in \text{End}(V)$ self-adjoint. Then V has an orthonormal basis of eigenvectors of α .*

Proof. By the lemma, α has a real eigenvalue λ , say. Thus we can find $v_1 \in V \setminus \{0\}$ such that $\alpha(v_1) = \lambda v_1$. Let $U := \ker \langle v_1, - \rangle: V \rightarrow \mathbf{F}$ the orthogonal complement of the span of v_1 in V .

If $u \in U$, then

$$\langle v_1, \alpha(u) \rangle = \langle \alpha(v_1), u \rangle = \langle \lambda v_1, u \rangle = \lambda \langle v_1, u \rangle = 0.$$

Thus $\alpha(u) \in U$ and α restricts to an element of $\text{End}(U)$. Since $\langle \alpha(v), w \rangle = \langle v, \alpha(w) \rangle$ for all $v, w \in V$ also for all $v, w \in U$ ie $\alpha|_U$ is also self-adjoint. By induction on $\dim V$ we can conclude that U has an o.n. basis of eigenvectors (v_2, \dots, v_n) of $\alpha|_U$. Then $(\frac{v_1}{\|v_1\|}, v_2, \dots, v_n)$ is an o.n. basis for V consisting of eigenvectors of α . \square

LECTURE 24

Corollary. *If V is an inner product space and $\alpha \in \text{End}(V)$ is self adjoint then V is the orthogonal direct sum of the α -eigenspaces (and all eigenvalues are real).*

Corollary. *Let $A \in \text{Mat}_n(\mathbf{R})$ (resp. $\text{Mat}_n(\mathbf{C})$) be a symmetric bilinear (resp. Hermitian) matrix. Then there is an orthogonal (resp. unitary) matrix P such that $P^T A P$ (resp. $P^\dagger A P$) is diagonal with real entries.*

Proof. Let $\langle -, - \rangle$ be the standard inner product on \mathbf{F}^n . Then $A \in \text{End}(\mathbf{F}^n)$ is self-adjoint so \mathbf{F}^n has an o.n. basis (e_1, \dots, e_n) consisting of eigenvectors of A . Let P be the matrix whose columns are given by e_1, \dots, e_n . Then P is orthogonal (resp. unitary) and $P^T A P = P^{-1} A P$ (resp. $P^\dagger A P$) is diagonal with real entries. \square

Corollary. *Let V be a f.d. real (resp. complex) inner product space and $\psi: V \times V \rightarrow \mathbf{R}$ a symmetric bilinear (resp. Hermitian) form. Then there is an orthonormal basis of V such that ψ is represented by a diagonal matrix.*

Proof. Let (u_1, \dots, u_n) be any o.n. basis for V and suppose that A represents ψ with respect to this basis. Then A is symmetric (resp. Hermitian) and there is an orthogonal (resp. unitary) matrix P such that $P^T A P$ (resp. $P^\dagger A P$) is diagonal. Let $v_i = \sum_k P_{ki} u_k$. Then (v_1, \dots, v_n) is an o.n. basis and ψ is represented by $P^T A P$ (resp. $P^\dagger A P$) with respect to it. \square

Remark. Note that in the proof the diagonal entries of $P^T A P$ (resp. $P^\dagger A P$) are the eigenvalues of A . Thus it is easy to see that the signature of ψ is given by

$$\# \text{ of positive eigenvalues of } A - \# \text{ of negative eigenvalues of } A.$$

Corollary. *Let V be a f.d. real (resp. complex) vector space and let ϕ and ψ be symmetric bilinear (resp. Hermitian) forms on V . If ϕ is positive-definite there is a basis (v_1, \dots, v_n) for V with respect to which both forms are represented by a diagonal matrix.*

Proof. Use ϕ to make V into an inner product space and then use the last result to find an o.n basis with respect to which ψ is represented by a diagonal matrix. Then ϕ is represented by I_n with respect to this basis. \square

Corollary. Let $A, B \in \text{Mat}_n(\mathbf{R})$ (resp. $\text{Mat}_n(\mathbf{C})$) be symmetric (resp. Hermitian) matrices such that A is positive definite (ie $\bar{v}^T A v > 0$ for all $v \in \mathbf{F}^n \setminus \{0\}$). Then there is an invertible matrix Q such that $Q^T A Q$ and $Q^T B Q$ (resp. $Q^\dagger A Q$ and $Q^\dagger B Q$) are both diagonal.

We can also prove a diagonalisability theorem for unitary endomorphisms.

Theorem. Let V be a f.d. complex inner product space and $\alpha \in \text{End}(V)$ be unitary. Then V has an o.n. basis consisting of eigenvectors of α . Moreover all the eigenvalues have length 1.

Proof. By the fundamental theorem of algebra, α has an eigenvalue λ , say and there is $v \in V$ such that $\|v\| = 1$ and $\alpha v = \lambda v$. Then

$$1 = \langle v, v \rangle = \langle \alpha(v), \alpha(v) \rangle = \langle \lambda v, \lambda v \rangle = |\lambda|^2.$$

Thus $|\lambda| = 1$. Let $W = \ker \langle v, - \rangle: V \rightarrow \mathbf{C}$ a $\dim V - 1$ dimensional subspace. Then if $w \in W$,

$$\langle v, \alpha(w) \rangle = \langle \alpha^{-1} v, w \rangle = \langle \frac{1}{\lambda} v, w \rangle = \lambda \langle v, w \rangle = 0.$$

Thus α restricts to a unitary endomorphism of W . By induction W has an o.n. basis consisting of eigenvectors of α . By adding v to this basis of W we obtain a suitable basis of V . \square

Remarks.

- (1) This theorem and its self-adjoint version have a common generalisation in the complex case. The key point is that α and α^* commute — see Example Sheet 4 Question 9.
- (2) It is not possible to diagonalise a real orthogonal matrix in general. For example a rotation in \mathbf{R} through an angle that is not an integer multiple of π . However, one can classify orthogonal maps in a similar fashion — see Example Sheet 4 Question 14.