Proof. Let $\{\mathbf{u}_1, \cdots, \mathbf{u}_m\}$ be a basis for U and extend this to a basis $\{\mathbf{u}_1, \cdots, \mathbf{u}_m, \mathbf{u}_m, \mathbf{u}_m\}$ $\mathbf{v}_{m+1}, \cdots, \mathbf{v}_n$ for V. We want to show that $\{\mathbf{v}_{m+1} + U, \cdots, \mathbf{v}_n + U\}$ is a basis for V/U.

It is easy to see that this spans V/U. If $\mathbf{v} + U \in V/U$, then we can write

$$\mathbf{v} = \sum \lambda_i \mathbf{u}_i + \sum \mu_i \mathbf{v}_i.$$

Then

$$\mathbf{v} + U = \sum \mu_i(\mathbf{v}_i + U) + \sum \lambda_i(\mathbf{u}_i + U) = \sum \mu_i(\mathbf{v}_i + U).$$

So done.

To show that they are linearly independent, suppose that

$$\sum \lambda_i (\mathbf{v}_i + U) = \mathbf{0} + U = U.$$

Then this requires

$$\sum \lambda_i \mathbf{v}_i \in U.$$

Then we can write this as a linear combination of the \mathbf{u}_i 's. So

$$\sum \lambda_i \mathbf{v}_i = \sum \mu_j \mathbf{u}_j$$

for some μ_j . Since $\{\mathbf{u}_1, \cdots, \mathbf{u}_m, \mathbf{v}_{n+1}, \cdots, \mathbf{v}_n\}$ is a basis for V, we must have lesale.co.1 $\lambda_i = \mu_j = 0$ for all i, j. So $\{\mathbf{v}_i + U\}$ is linearly independent.

1.3Direct sums

s in order to confuse students. We are going to define direct sums in

Definition ((Internation nic sum). Suppose V vector space over \mathbb{F} and $U, W \subseteq V$ are Vternal) direct sum of U andWe say that 🖊 is the U + W =

(ii)
$$U \cap W = 0$$
.

We write $V = U \oplus W$.

Equivalently, this requires that every $\mathbf{v} \in V$ can be written uniquely as $\mathbf{u} + \mathbf{w}$ with $\mathbf{u} \in U, \mathbf{w} \in W$. We say that U and W are complementary subspaces of V.

You will show in the example sheets that given any subspace $U \subseteq V, U$ must have a complementary subspace in V.

Example. Let $V = \mathbb{R}^2$, and $U = \langle \begin{pmatrix} 0 \\ 1 \end{pmatrix} \rangle$. Then $\langle \begin{pmatrix} 1 \\ 1 \end{pmatrix} \rangle$ and $\langle \begin{pmatrix} 1 \\ 0 \end{pmatrix} \rangle$ are both complementary subspaces to U in V.

Definition ((External) direct sum). If U, W are vector spaces over \mathbb{F} , the (external) direct sum is

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

with addition and scalar multiplication componentwise:

$$(\mathbf{u}_1, \mathbf{w}_1) + (\mathbf{u}_2, \mathbf{w}_2) = (\mathbf{u}_1 + \mathbf{u}_2, \mathbf{w}_1 + \mathbf{w}_2), \quad \lambda(\mathbf{u}, \mathbf{w}) = (\lambda \mathbf{u}, \lambda \mathbf{w}).$$

- (ii) $AE_{Ii}^n(\lambda)$ is obtained by adding $\lambda \times$ column *i* to column *j*.
- (iii) $AT_i^n(\lambda)$ is obtained from A by rescaling the *i*th column by λ .

Multiplying on the left instead of the right would result in the same operations performed on the rows instead of the columns.

Proposition. If $A \in \operatorname{Mat}_{n,m}(\mathbb{F})$, then there exists invertible matrices $P \in \operatorname{GL}_m(\mathbb{F}), Q \in \operatorname{GL}_n(\mathbb{F})$ so that

$$Q^{-1}AP = \begin{pmatrix} I_r & 0\\ 0 & 0 \end{pmatrix}$$

for some $0 \le r \le \min(m, n)$.

We are going to start with A, and then apply these operations to get it into this form.

Proof. We claim that there are elementary matrices E_1^m, \dots, E_a^m and F_1^n, \dots, F_b^n (these E are not necessarily the shears, but any elementary matrix) such that

$$E_1^m \cdots E_a^m A F_1^n \cdots F_b^n = \begin{pmatrix} I_r & 0\\ 0 & 0 \end{pmatrix}$$

This suffices since the $E_i^m \in \operatorname{GL}_M(\mathbb{F})$ and $F_j^n \in \operatorname{GL}_n(\mathbb{F})$. Moreover, to prove the claim, it suffices to find a sequence of elementary row and to man operations reducing A to this form.

If A = 0, then done. If not, there is the i, there is that $A_{ij} \neq 0$. By swapping row 1 and row i; and then column 1 and column j, we can a some $A_{11} \neq 0$. By rescaling row 1 by $\frac{1}{A_{11}}$ we can further assume $A_{11} = 1$.

Now we can add $-A_{1j}$ times column 1 to robumn j for each $j \neq 1$, and then add $-A_1$ threes row 1 to row $j \neq 0$. Then we now have

5	(1)	0		0)
4	0			
A =			B	
	$\setminus 0$)

Now B is smaller than A. So by induction on the size of A, we can reduce B to a matrix of the required form, so done. \Box

It is an exercise to show that the row and column operations do not change the row rank or column rank, and deduce that they are equal. *Proof.* We just have to compute

$$\psi_L(\mathbf{e}_i)(\mathbf{f}_j) = A_{ij} = \left(\sum A_{i\ell}\eta_\ell\right)(\mathbf{f}_j).$$

So we get

$$\psi_L(\mathbf{e}_i) = \sum A_{\ell i}^T \eta_\ell.$$

 $\psi_R(\mathbf{f}_i)(\mathbf{e}_i) = A_{ij}.$

So A^T represents ψ_L . We also have

So

$$\psi_R(\mathbf{f}_j) = \sum A_{kj} arepsilon_k.$$

Definition (Left and right kernel). The kernel of ψ_L is *left kernel* of ψ , while the kernel of ψ_R is the *right kernel* of ψ .

Then by definition, **v** is in the left kernel if $\psi(\mathbf{v}, \mathbf{w}) = 0$ for all $\mathbf{w} \in W$. More generally, if $T \subseteq V$, then we write

$$T^{\perp} = \{ \mathbf{w} \in W : \psi(\mathbf{t}, \mathbf{w}) = 0 \text{ for all } \mathbf{t} \in T \}.$$

Similarly, if $U \subseteq W$, then we write

$$^{\perp}U = \{ \mathbf{v} \in V : \psi(\mathbf{v}, \mathbf{u}) = 0 \text{ for all } \mathbf{u} \in V \}$$

sale.co.uk In particular, $V^{\perp} = \ker \psi_R$ and ${}^{\perp}W =$

If we have a non-trivial left come sense, some elements in V (or W) and we don't

Definition D egenerate bilinear for is non-degenerate if the left and 187 is *degenerate* otherwise. Ŵ

Definition (Rank of bilinear form). If $\psi: V \to W$ is a bilinear form \mathbb{F} on a finite-dimensional vector space V, then the rank of V is the rank of any matrix representing ϕ . This is well-defined since $r(P^T A Q) = r(A)$ if P and Q are invertible.

Alternatively, it is the rank of ψ_L (or ψ_R).

Lemma. Let V and W be finite-dimensional vector spaces over \mathbb{F} with bases $(\mathbf{e}_1, \cdots, \mathbf{e}_n)$ and $(\mathbf{f}_1, \cdots, \mathbf{f}_m)$ be their basis respectively.

Let $\psi: V \times W \to \mathbb{F}$ be a bilinear form represented by A with respect to these bases. Then ϕ is non-degenerate if and only if A is (square and) invertible. In particular, V and W have the same dimension.

We can understand this as saying if there are too many things in V (or W), then some of them are bound to be useless.

Proof. Since ψ_R and ψ_L are represented by A and A^T (in some order), they both have trivial kernel if and only if $n(A) = n(A^T) = 0$. So we need $r(A) = \dim V$ and $r(A^T) = \dim W$. So we need $\dim V = \dim W$ and A have full rank, i.e. the corresponding linear map is bijective. So done.

Lemma. If A is an upper triangular matrix, i.e.

$$A = \begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ 0 & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_{nn} \end{pmatrix}$$

Then

$$\det A = \prod_{i=1}^{n} a_{ii}.$$

Proof. We have

$$\det A = \sum_{\sigma \in S_n} \varepsilon(\sigma) \prod_{i=1}^n A_{i\sigma(i)}$$

But $A_{i\sigma(i)} = 0$ whenever $i > \sigma(i)$. So

$$\prod_{i=1}^{n} A_{i\sigma(i)} = 0$$

if there is some $i \in \{1, \dots, n\}$ such that $i > \sigma(i)$.

However, the only permutation in which $i \le \sigma(i)$. However, the only permutation in which $i \le \sigma(i)$ for all i is the identity. So the only thing that contributes in the sum is $\sigma = \text{id}$. So $\det A = \prod_{i=1}^{n} (A_i)$. the only thing that contributes in the sum is $\sigma = id$. So

olume. How can we define To motivate ve need n lo ion oi It should be that the "volume" cannot be for the second se volume r space? It termined. aying the volume is \mathbf{C} modules unless we provide the units, e.g. cm³. 1" So we have an axiomatic definition for what it means for something to denote a "volume".

Definition (Volume form). A volume form on \mathbb{F}^n is a function $d: \mathbb{F}^n \times \cdots \times \mathbb{F}^n \to$ \mathbb{F} that is

(i) Multilinear, i.e. for all i and all $\mathbf{v}_1, \cdots, \mathbf{v}_{i-1}, \mathbf{v}_{i+1}, \cdots, \mathbf{v}_n \in \mathbb{F}^n$, we have

 $\mathbf{d}(\mathbf{v}_1,\cdots,\mathbf{v}_{i-1},\cdot,\mathbf{v}_{i+1},\cdots,\mathbf{v}_n)\in(\mathbb{F}^n)^*.$

(ii) Alternating, i.e. if $\mathbf{v}_i = \mathbf{v}_j$ for some $i \neq j$, then

$$d(\mathbf{v}_1,\cdots,\mathbf{v}_n)=0.$$

We should think of $d(\mathbf{v}_1, \cdots, \mathbf{v}_n)$ as the *n*-dimensional volume of the parallelopiped spanned by $\mathbf{v}_1, \cdots, \mathbf{v}_n$.

We can view $A \in \operatorname{Mat}_n(\mathbb{F})$ as *n*-many vectors in \mathbb{F}^n by considering its columns $A = (A^{(1)} A^{(2)} \cdots A^{(n)})$, with $A^{(i)} \in \mathbb{F}^n$. Then we have

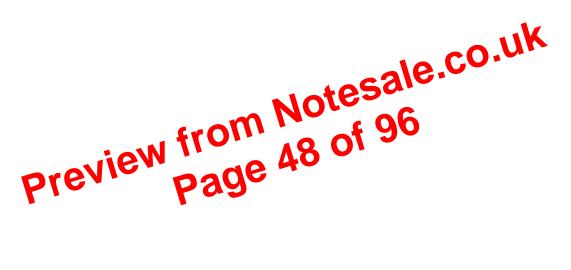
Lemma. det A is a volume form.

 $\sigma = \sigma_1 \sigma_2$, where σ_1 is a permutation of $\{1, \dots, k\}$ and fixes the remaining things, while σ_2 fixes $\{1, \dots, k\}$, and permutes the remaining. Then

$$\det X = \sum_{\sigma=\sigma_1\sigma_2} \varepsilon(\sigma_1\sigma_2) \prod_{i=1}^k X_{i\sigma_1(i)} \prod_{j=1}^\ell X_{k+j\sigma_2(k+j)}$$
$$= \left(\sum_{\sigma_1 \in S_k} \varepsilon(\sigma_1) \prod_{i=1}^k A_{i\sigma_1(i)}\right) \left(\sum_{\sigma_2 \in S_\ell} \varepsilon(\sigma_2) \prod_{j=1}^\ell B_{j\sigma_2(j)}\right)$$
$$= (\det A)(\det B)$$

Corollary.

$$\det \begin{pmatrix} A_1 & & \text{stuff} \\ & A_2 & & \\ & & \ddots & \\ 0 & & & A_n \end{pmatrix} = \prod_{i=1}^n \det A_i$$



Our initial strategy is to identify basis-independent invariants for endomorphisms. For example, we will show that the rank, trace, determinant and characteristic polynomial are all such invariants.

Recall that the trace of a matrix $A \in Mat_n(\mathbb{F})$ is the sum of the diagonal elements:

Definition (Trace). The *trace* of a matrix of $A \in Mat_n(\mathbb{F})$ is defined by

$$\operatorname{tr} A = \sum_{i=1}^{n} A_{ii}.$$

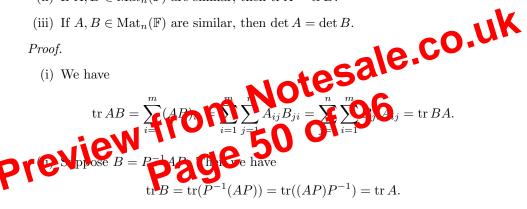
We want to show that the trace is an invariant. In fact, we will show a stronger statement (as well as the corresponding statement for determinants):

Lemma.

(i) If $A \in Mat_{m,n}(\mathbb{F})$ and $B \in Mat_{n,m}(\mathbb{F})$, then

$$\operatorname{tr} AB = \operatorname{tr} BA.$$

(ii) If $A, B \in Mat_n(\mathbb{F})$ are similar, then tr A = tr B.



(iii) We have

$$\det(P^{-1}AP) = \det P^{-1} \det A \det P = (\det P)^{-1} \det A \det P = \det A. \square$$

This allows us to define the trace and determinant of an *endomorphism*.

Definition (Trace and determinant of endomorphism). Let $\alpha \in \text{End}(V)$, and A be a matrix representing α under any basis. Then the *trace* of α is tr $\alpha = \text{tr } A$, and the *determinant* is det $\alpha = \text{det } A$.

The lemma tells us that the determinant and trace are well-defined. We can also define the determinant without reference to a basis, by defining more general volume forms and define the determinant as a scaling factor.

The trace is slightly more tricky to define without basis, but in IB Analysis II example sheet 4, you will find that it is the directional derivative of the determinant at the origin.

To talk about the characteristic polynomial, we need to know what eigenvalues are.

We can use the last lemma and induction to show that any non-zero $f \in \mathbb{F}[t]$ can be written as

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

where $\lambda_1, \dots, \lambda_k$ are all distinct, $a_i > 1$, and g is a polynomial with no roots in $\mathbb{F}.$

Hence we obtain the following:

Lemma. A non-zero polynomial $f \in \mathbb{F}[t]$ has at most deg f roots, counted with multiplicity.

Corollary. Let $f, g \in \mathbb{F}[t]$ have degree < n. If there are $\lambda_1, \dots, \lambda_n$ distinct such that $f(\lambda_i) = g(\lambda_i)$ for all *i*, then f = g.

Proof. Given the lemma, consider f - g. This has degree less than n, and $(f-g)(\lambda_i) = 0$ for $i = 1, \dots, n$. Since it has at least $n \ge \deg(f-g)$ roots, we must have f - g = 0. So f = g.

Corollary. If \mathbb{F} is infinite, then f and g are equal if and only if they agree on all points.

More importantly, we have the following:

Theorem (The fundamental theorem of algebra). Every non-constant by yao mial over \mathbb{C} has a root in \mathbb{C} . We will not prove this. We say \mathbb{C} is an *algebraically*

We say \mathbb{C} is an *algebraically closed* \bigcirc has precisely nIt thus follows that ever \mathcal{T} by homial over \mathbb{C} of degree roots, counted with multiplicity, since if we write $f(t) = g(t) \prod (t - \lambda_i)^{a_i}$ and g has not easy then g is constant. So the number of roots is $\sum a_i = \deg f$, a super-write with multiplicity the number of roots is $\sum a_i = \deg f$,

Previous that we consider the set of the se quadratic polynomials with no real roots (since complex roots of real polynomials come in complex conjugate pairs).

6.2.2 Minimal polynomial

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

$$f(A) = \sum_{i=0}^{m} a_i A^i, \quad f(\alpha) = \sum_{i=0}^{m} a_i \alpha$$

where $A^0 = I$ and $\alpha^0 = \iota$.

Theorem (Diagonalizability theorem). Suppose $\alpha \in \text{End}(V)$. Then α is diagonalizable if and only if there exists non-zero $p(t) \in \mathbb{F}[t]$ such that $p(\alpha) = 0$, and p(t) can be factored as a product of *distinct* linear factors.

Proof. Suppose α is diagonalizable. Let $\lambda_1, \dots, \lambda_k$ be the distinct eigenvalues of α . We have Ŀ

$$V = \bigoplus_{i=1}^{n} E(\lambda_i).$$

We can now re-state our diagonalizability theorem.

Theorem (Diagonalizability theorem 2.0). Let $\alpha \in \text{End}(V)$. Then α is diagonalizable if and only if $M_{\alpha}(t)$ is a product of its distinct linear factors.

Proof. (\Leftarrow) This follows directly from the previous diagonalizability theorem.

 (\Rightarrow) Suppose α is diagonalizable. Then there is some $p \in \mathbb{F}[t]$ non-zero such that $p(\alpha) = 0$ and p is a product of distinct linear factors. Since M_{α} divides p, M_{α} also has distinct linear factors. \Box

Theorem. Let $\alpha, \beta \in \text{End}(V)$ be both diagonalizable. Then α and β are simultaneously diagonalizable (i.e. there exists a basis with respect to which both are diagonal) if and only if $\alpha\beta = \beta\alpha$.

This is important in quantum mechanics. This means that if two operators do not commute, then they do not have a common eigenbasis. Hence we have the uncertainty principle.

Proof. (\Rightarrow) If there exists a basis ($\mathbf{e}_1, \dots, \mathbf{e}_n$) for V such that α and β are represented by A and B respectively, with both diagonal, then by direct computation, AB = BA. But AB represents $\alpha\beta$ and BA represents $\beta\alpha$. So $\alpha\beta = \beta\alpha$.

(\Leftarrow) Suppose $\alpha\beta = \beta\alpha$. The idea is to consider each eigenspace of α individually, and then diagonalize β in each of the eigenspaces. Since α is diagonalizable, we can write

$$V = \bigoplus_{i=1}^{k} E_{\alpha}(\lambda_i),$$

 $=\beta(\lambda_i \mathbf{v}) = \lambda_i \beta(\mathbf{v}).$

where λ_i are the eigenvalues of V. We will E_i for $E_{\alpha}(\lambda_i)$. We want to show that β sends E_i to its df $\dot{\mu} = 0$ ($E_i = E_i$. Let $\mathbf{v} \in E_i$. Then ψ want $\beta(\mathbf{v})$ to be in E_i . This is true since

So $\beta(\mathbf{v})$ is an eigenvector of α with eigenvalue λ_i . Now we can view $\beta|_{E_i} \in \text{End}(E_i)$. Note that

$$M_{\beta}(\beta|_{E_i}) = M_{\beta}(\beta)|_{E_i} = 0.$$

Since $M_{\beta}(t)$ is a product of its distinct linear factors, it follows that $\beta|_{E_i}$ is diagonalizable. So we can choose a basis B_i of eigenvectors for $\beta|_{E_i}$. We can do this for all *i*.

Then since V is a direct sum of the E_i 's, we know that $B = \bigcup_{i=1}^k B_i$ is a basis for V consisting of eigenvectors for both α and β . So done.

6.3 The Cayley-Hamilton theorem

We will first state the theorem, and then prove it later.

Recall that $\chi_{\alpha}(t) = \det(t\iota - \alpha)$ for $\alpha \in \operatorname{End}(V)$. Our main theorem of the section (as you might have guessed from the title) is

Theorem (Cayley-Hamilton theorem). Let V be a finite-dimensional vector space and $\alpha \in \text{End}(V)$. Then $\chi_{\alpha}(\alpha) = 0$, i.e. $M_{\alpha}(t) \mid \chi_{\alpha}(t)$. In particular, deg $M_{\alpha} \leq n$.

Then we get the result

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

We would like to just throw in t = A, and get the desired result, but in all these derivations, t is assumed to be a real number, and, $tI_n - A$ is the matrix

1	$t - a_{11}$	a_{12}	• • •	a_{1n}
	a_{21}	$t - a_{22}$		a_{2n}
	÷	÷	·	:
(a_{n1}	a_{n2}		$t - a_{nn}$

It doesn't make sense to put our A in there.

However, what we *can* do is to note that since this is true for all values of t, the coefficients on both sides must be equal. Equating coefficients in t^k , we have

$$-AB_{0} = a_{0}I_{n}$$

$$B_{0} - AB_{1} = a_{1}I_{n}$$

$$\vdots$$

$$B_{n-2} - AB_{n-1} = a_{n-1}I_{n}$$

$$AB_{n-1} - 0 = I_{n}$$
We now multiply each row a suitable power of A to cheater the second state of the second

Summing this up then gives $\chi_{\alpha}(A) = 0$.

This proof suggests that we *really* ought to be able to just substitute in $t = \alpha$ and be done. In fact, we can do this, after we develop sufficient machinery. This will be done in the IB Groups, Rings and Modules course.

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

- (i) λ is an eigenvalue of α .
- (ii) λ is a root of $\chi_{\alpha}(t)$.
- (iii) λ is a root of $M_{\alpha}(t)$.

Proof.

- (i) \Leftrightarrow (ii): λ is an eigenvalue of α if and only if $(\alpha \lambda \iota)(\mathbf{v}) = 0$ has a non-trivial root, iff det $(\alpha \lambda \iota) = 0$.
- (iii) \Rightarrow (ii): This follows from Cayley-Hamilton theorem since $M_{\alpha} \mid \chi_{\alpha}$.

with $\lambda, \mu \in \mathbb{C}$ distinct. We see that M_A determines the Jordan normal form of A, but χ_A does not.

Every 3×3 matrix in Jordan normal form is one of the six types. Here λ_1, λ_2 and λ_3 are distinct complex numbers.

Jordan normal form	χ_A	M_A		
$\begin{pmatrix} \lambda_1 & 0 & 0\\ 0 & \lambda_2 & 0\\ 0 & 0 & \lambda_3 \end{pmatrix}$	$(t-\lambda_1)(t-\lambda_2)(t-\lambda_3)$	$(t-\lambda_1)(t-\lambda_2)(t-\lambda_3)$		
$\begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_1 & 0 \\ 0 & 0 & \lambda_2 \end{pmatrix}$	$(t-\lambda_1)^2(t-\lambda_2)$	$(t - \lambda_1)(t - \lambda_2)$		
$\begin{pmatrix} \lambda_1 & 1 & 0\\ 0 & \lambda_1 & 0\\ 0 & 0 & \lambda_2 \end{pmatrix}$	$(t-\lambda_1)^2(t-\lambda_2)$	$(t-\lambda_1)^2(t-\lambda_2)$		
$\begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_1 & 0 \\ 0 & 0 & \lambda_1 \end{pmatrix}$	$(t-\lambda_1)^3$	$(t - \lambda_1)$		
$\begin{pmatrix} \lambda_1 & 1 & 0 \\ 0 & \lambda_1 & 0 \\ 0 & 0 & \lambda_1 \end{pmatrix}$	$(t-\lambda_1)^3$	$(t - \lambda_1)^2 CO$		
$\begin{pmatrix} \lambda_1 & 1 & 0 \\ 0 & \lambda_1 & 1 \\ 0 & 0 & \lambda_1 \end{pmatrix}$	(t-Note	$\frac{53}{(t-\lambda_1)^3}$		

Notice that χ_A and M_{44} together determine the locan normal form of a 3×3 complexing χ_A in the second case, since if we are given $M_{42} = 1 - \lambda_1 (t - \lambda_{42})$ we now doe of the roots is double, but not which one.

In general, though, even χ_A and M_A together does not suffice.

We now want to understand the Jordan normal blocks better. Recall the definition ()

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

If $(\mathbf{e}_1, \cdots, \mathbf{e}_n)$ is the standard basis for \mathbb{C}^n , we have

$$J_n(0)(\mathbf{e}_1) = 0, \quad J_n(0)(\mathbf{e}_i) = \mathbf{e}_{i-1} \text{ for } 2 \le i \le n.$$

Thus we know

$$J_n(0)^k(\mathbf{e}_i) = \begin{cases} 0 & i \le k \\ \mathbf{e}_{i-k} & k < i \le n \end{cases}$$

In other words, for k < n, we have

$$(J_n(\lambda) - \lambda I)^k = J_n(0)^k = \begin{pmatrix} 0 & I_{n-k} \\ 0 & 0 \end{pmatrix}.$$

Proof. We work blockwise for

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

We have previously computed

$$n((J_m(\lambda) - \lambda I_m)^r) = \begin{cases} r & r \le m \\ m & r > m \end{cases}$$

Hence we know

$$n((J_m(\lambda) - \lambda I_m)^r) - n((J_m(\lambda) - \lambda I_m)^{r-1}) = \begin{cases} 1 & r \le m \\ 0 & \text{otherwise.} \end{cases}$$

It is also easy to see that for $\mu \neq \lambda$,

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

Adding up for each block, for $r \ge 1$, we have

$$n((\alpha - \lambda \iota)^r) - n((\alpha - \lambda \iota)^{r-1}) =$$
 number of Jordan blocks I_n

Evit $n \ge r$. take an Contractional we take an additional We can interpret this result as follo s. i 🕐 power of $J_m(\lambda) - \lambda I_m$ we get to . So we kill off one more column in the matrix, and the noisity prease by one. This happens until $(J_n) = \lambda I_m)^r = 0$, in which case increasing the power no longer affects in matrix. So when work a the difference in nullity, we are counting the number of blocks that are affected by the increase in power, which is the number of blocks of size at least r.

We have now proved uniqueness, but existence is not yet clear. To show this, we will reduce it to the case where there is exactly one eigenvalue. This reduction is easy if the matrix is diagonalizable, because we can decompose the matrix into each eigenspace and then work in the corresponding eigenspace. In general, we need to work with "generalized eigenspaces".

Theorem (Generalized eigenspace decomposition). Let V be a finite-dimensional vector space \mathbb{C} such that $\alpha \in \text{End}(V)$. Suppose that

$$M_{\alpha}(t) = \prod_{i=1}^{k} (t - \lambda_i)^{c_i},$$

with $\lambda_1, \dots, \lambda_k \in \mathbb{C}$ distinct. Then

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

where $V_i = \ker((\alpha - \lambda_i \iota)^{c_i})$ is the generalized eigenspace.

j > r

We see that the diagonal matrix we get is not unique. We can re-scale our basis by any constant, and get an equivalent expression.

Theorem. Let ϕ be a symmetric bilinear form over a complex vector space V. Then there exists a basis $(\mathbf{v}_1, \cdots, \mathbf{v}_m)$ for V such that ϕ is represented by

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

with respect to this basis, where $r = r(\phi)$.

Proof. We've already shown that there exists a basis $(\mathbf{e}_1, \cdots, \mathbf{e}_n)$ such that $\phi(\mathbf{e}_i,\mathbf{e}_j) = \lambda_i \delta_{ij}$ for some λ_{ij} . By reordering the \mathbf{e}_i , we can assume that $\lambda_1, \cdots, \lambda_r \neq 0 \text{ and } \lambda_{r+1}, \cdots, \lambda_n = 0.$

For each $1 \le i \le r$, there exists some μ_i such that $\mu_i^2 = \lambda_i$. For $r+1 \le r \le n$, we let $\mu_i = 1$ (or anything non-zero). We define

$$\mathbf{v}_i = \frac{\mathbf{e}_i}{\mu_i}.$$

Then

$$\phi(\mathbf{v}_i, \mathbf{v}_j) = \frac{1}{\mu_i \mu_j} \phi(\mathbf{e}_i, \mathbf{e}_j) = \begin{cases} 0 & i \neq j \text{ or } i = \\ 1 & i = j < r. \end{cases}$$

So done.

e.co.u Note that it follows that for the correspond form a, we have $\operatorname{at}_n(\mathbb{C})$ is congruent to a unique matrix of

Now this theorem is a bit too strong, and we are going to fix that next lecture, by talking about Hermitian forms and sesquilinear forms. Before that, we do the equivalent result for real vector spaces.

Theorem. Let ϕ be a symmetric bilinear form of a finite-dimensional vector space over \mathbb{R} . Then there exists a basis $(\mathbf{v}_1, \cdots, \mathbf{v}_n)$ for V such that ϕ is represented

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

with $p + q = r(\phi)$, $p, q \ge 0$. Equivalently, the corresponding quadratic forms is given by

$$q\left(\sum_{i=1}^{n} a_i \mathbf{v}_i\right) = \sum_{i=1}^{p} a_i^2 - \sum_{j=p+1}^{p+q} a_j^2.$$

Definition (Signature). The signature of a bilinear form ϕ is the number p - q, where p and q are as above.

Of course, we can recover p and q from the signature and the rank of ϕ .

Corollary. Every real symmetric matrix is congruent to precisely one matrix of the form \sim

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

7.2Hermitian form

The above result was nice for real vector spaces. However, if ϕ is a bilinear form on a \mathbb{C} -vector space V, then $\phi(i\mathbf{v}, i\mathbf{v}) = -\phi(\mathbf{v}, \mathbf{v})$. So there can be no good notion of positive definiteness for complex bilinear forms. To make them work for complex vector spaces, we need to modify the definition slightly to obtain Hermitian forms.

Definition (Sesquilinear form). Let V, W be complex vector spaces. Then a sesquilinear form is a function $\phi: V \times W \to \mathbb{C}$ such that

(i) $\phi(\lambda \mathbf{v}_1 + \mu \mathbf{v}_2, \mathbf{w}) = \bar{\lambda}\phi(\mathbf{v}_1, \mathbf{w}) + \bar{\mu}\phi(\mathbf{v}_2, \mathbf{w}).$

(ii)
$$\phi(\mathbf{v}, \lambda \mathbf{w}_1 + \mu \mathbf{w}_2) = \lambda \phi(\mathbf{v}, \mathbf{w}_1) + \mu \phi(\mathbf{v} \mathbf{w}_2).$$

for all $\mathbf{v}, \mathbf{v}_1, \mathbf{v}_2 \in V$, $\mathbf{w}, \mathbf{w}_1, \mathbf{w}_2 \in W$ and $\lambda, \mu \in \mathbb{C}$. Note that some people have an opposite of the solution of the first argument.

Note that some people have an opposite d function, where we have linearity in the first argument and conjugate linearity in the second These are called second in since "sesqui" more "one and a half", and this is linear in the sound argument and "tarf linearies in the first.

Afternatively, to define a set pline form, we can define a new complex renor space \bar{V} struct possible caking the same abelian group (i.e. the same underlying set and addition), but with the scalar multiplication $\mathbb{C} \times \bar{V} \to \bar{V}$ defined as

$$(\lambda, \mathbf{v}) \mapsto \lambda \mathbf{v}$$

Then a sesquilinear form on $V \times W$ is a bilinear form on $\overline{V} \times W$. Alternatively, this is a linear map $W \to \overline{V}^*$.

Definition (Representation of sesquilinear form). Let V, W be finite-dimensional complex vector spaces with basis $(\mathbf{v}_1, \cdots, \mathbf{v}_n)$ and $(\mathbf{w}_1, \cdots, \mathbf{w}_m)$ respectively, and $\phi: V \times W \to \mathbb{C}$ be a sesquilinear form. Then the matrix representing ϕ with respect to these bases is

$$A_{ij} = \phi(\mathbf{v}_i, \mathbf{w}_j).$$

for $1 \leq i \leq n, 1 \leq j \leq m$.

As usual, this determines the whole sesquilinear form. This follows from the analogous fact for the bilinear form on $\overline{V} \times W \to \mathbb{C}$. Let $\mathbf{v} = \sum \lambda_i \mathbf{v}_i$ and $W = \sum \mu_j w_j$. Then we have

$$\phi(\mathbf{v}, \mathbf{w}) = \sum_{i,j} \overline{\lambda}_i \mu_j \phi(\mathbf{v}_i, \mathbf{w}_j) = \lambda^{\dagger} A \mu.$$

8.2 Gram-Schmidt orthogonalization

As mentioned, we want to make sure every vector space has an orthonormal basis, and we can extend any orthonormal set to an orthonormal basis, at least in the case of finite-dimensional vector spaces. The idea is to start with an arbitrary basis, which we know exists, and produce an orthonormal basis out of it. The way to do this is the Gram-Schmidt process.

Theorem (Gram-Schmidt process). Let V be an inner product space and $\mathbf{e}_1, \mathbf{e}_2, \cdots$ a linearly independent set. Then we can construct an orthonormal set $\mathbf{v}_1, \mathbf{v}_2, \cdots$ with the property that

$$\langle \mathbf{v}_1, \cdots, \mathbf{v}_k \rangle = \langle \mathbf{e}_1, \cdots, \mathbf{e}_k \rangle$$

for every k.

Note that we are not requiring the set to be finite. We are just requiring it to be countable.

Proof. We construct it iteratively, and prove this by induction on k. The base case k = 0 is contentless.

Suppose we have already found $\mathbf{v}_1, \cdots, \mathbf{v}_k$ that satisfies the properties. We e.co.u define

$$\mathbf{u}_{k+1} = \mathbf{e}_{k+1} - \sum_{i=1}^{k} (\mathbf{v}_i, \mathbf{e}_{i+1})$$

We want to prove that this is orthogon for $i \leq k$. We have

is non-zero. Note that We want to argue that

$$\langle \mathbf{v}_1, \cdots, \mathbf{v}_k, \mathbf{u}_{k+1} \rangle = \langle \mathbf{v}_1, \cdots, \mathbf{v}_k, \mathbf{e}_{k+1} \rangle$$

since we can recover \mathbf{e}_{k+1} from $\mathbf{v}_1, \cdots, \mathbf{v}_k$ and \mathbf{u}_{k+1} by construction. We also know

$$\langle \mathbf{v}_1, \cdots, \mathbf{v}_k, \mathbf{e}_{k+1} \rangle = \langle \mathbf{e}_1, \cdots, \mathbf{e}_k, \mathbf{e}_{k+1} \rangle$$

by assumption. We know $\langle \mathbf{e}_1, \cdots, \mathbf{e}_k, \mathbf{e}_{k+1} \rangle$ has dimension k+1 since the \mathbf{e}_i are linearly independent. So we must have \mathbf{u}_{k+1} non-zero, or else $\langle \mathbf{v}_1, \cdots, \mathbf{v}_k \rangle$ will be a set of size k spanning a space of dimension k + 1, which is clearly nonsense.

Therefore, we can define

$$\mathbf{v}_{k+1} = \frac{\mathbf{u}_{k+1}}{\|\mathbf{u}_{k+1}\|}.$$

Then $\mathbf{v}_1, \cdots, \mathbf{v}_{k+1}$ is orthonormal and $\langle \mathbf{v}_1, \cdots, \mathbf{v}_{k+1} \rangle = \langle \mathbf{e}_1, \cdots, \mathbf{e}_{k+1} \rangle$ as required.

Corollary. If V is a finite-dimensional inner product space, then any orthonormal set can be extended to an orthonormal basis.

(ii) This is just Pythagoras' theorem. Note that if \mathbf{x} and \mathbf{y} are orthogonal, then

$$\begin{aligned} \|\mathbf{x} + \mathbf{y}\|^2 &= (\mathbf{x} + \mathbf{y}, \mathbf{x} + \mathbf{y}) \\ &= (\mathbf{x}, \mathbf{x}) + (\mathbf{x}, \mathbf{y}) + (\mathbf{y}, \mathbf{x}) + (\mathbf{y}, \mathbf{y}) \\ &= \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2. \end{aligned}$$

We apply this to our projection. For any $\mathbf{w} \in W$, we have

$$\|\mathbf{v} - \mathbf{w}\|^2 = \|\mathbf{v} - \pi(\mathbf{v})\|^2 + \|\pi(\mathbf{v}) - \mathbf{w}\|^2 \ge \|\mathbf{v} - \pi(\mathbf{v})\|^2$$

with equality if and only if $\|\pi(\mathbf{v}) - \mathbf{w}\| = 0$, i.e. $\pi(\mathbf{v}) = \mathbf{w}$.

8.3 Adjoints, orthogonal and unitary maps

Adjoints

Lemma. Let V and W be finite-dimensional inner product spaces and $\alpha: V \to W$ is a linear map. Then there exists a unique linear map $\alpha^*: W \to V$ such that

$$(\alpha \mathbf{v}, \mathbf{w}) = (\mathbf{v}, \alpha^* \mathbf{w}) \tag{(*)}$$

for all $\mathbf{v} \in V$, $\mathbf{w} \in W$.

Proof. There are two parts. We have to prove existence and uniqueness. We'll first prove it concretely using matrices, and then provide a concepturity as not what this means.

Let $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ and $(\mathbf{w}_1, \dots, \mathbf{w}_m)$ be orthogonal classis for V and W. Suppose α is represented by A. To show uniqueness, suppose $\alpha^* : W \to V$ satisfies $(\mathbf{v}, \alpha^* \mathbf{w}) = (\mathbf{v}, \alpha^* \mathbf{w})$ for all $\mathbf{v} \in V$, $\mathbf{w} \in W$, therefore $(\mathbf{v}_i, \alpha^* (\mathbf{w}_j)) = (\alpha(\mathbf{v}_i, \mathbf{w}_j))$ $(\mathbf{v}_i, \alpha^* (\mathbf{w}_j)) = (\alpha(\mathbf{v}_i, \mathbf{w}_j))$ $= \left(\sum_k A_{ki} \mathbf{w}_k, \mathbf{w}_j\right)$ $= \sum_k \bar{A}_{ki} (\mathbf{w}_k, \mathbf{w}_j) = \bar{A}_{ji}.$

So we get

$$\alpha^*(\mathbf{w}_j) = \sum_i (\mathbf{v}_i, \alpha^*(\mathbf{w}_j)) \mathbf{v}_i = \sum_i \bar{A}_{ji} \mathbf{v}_i$$

Hence α^* must be represented by A^{\dagger} . So α^* is unique.

To show existence, all we have to do is to show A^{\dagger} indeed works. Now let α^* be represented by A^{\dagger} . We can compute the two sides of (*) for arbitrary \mathbf{v}, \mathbf{w} . We have

$$\left(\alpha\left(\sum\lambda_{i}\mathbf{v}_{i}\right),\sum\mu_{j}\mathbf{w}_{j}\right)=\sum_{i,j}\bar{\lambda}_{i}\mu_{j}(\alpha(\mathbf{v}_{i}),\mathbf{w}_{j})$$
$$=\sum_{i,j}\bar{\lambda}_{i}\mu_{j}\left(\sum_{k}A_{ki}\mathbf{w}_{k},\mathbf{w}_{j}\right)$$
$$=\sum_{i,j}\bar{\lambda}_{i}\bar{A}_{ji}\mu_{j}.$$

Corollary. If $A, B \in \operatorname{Mat}_n(\mathbb{R})$ are symmetric and A is positive definitive (i.e. $\mathbf{v}^T A \mathbf{v} > 0$ for all $\mathbf{v} \in \mathbb{R}^n \setminus \{0\}$). Then there exists an invertible matrix Q such that $Q^T A Q$ and $Q^T B Q$ are both diagonal.

We can deduce similar results for complex finite-dimensional vector spaces, with the same proofs. In particular,

Proposition.

(i) If $A \in \operatorname{Mat}_n(\mathbb{C})$ is Hermitian, then there exists a unitary matrix $U \in \operatorname{Mat}_n(\mathbb{C})$ such that

$$U^{-1}AU = U^{\dagger}AU$$

is diagonal.

- (ii) If ψ is a Hermitian form on a finite-dimensional complex inner product space V, then there is an orthonormal basis for V diagonalizing ψ .
- (iii) If ϕ, ψ are Hermitian forms on a finite-dimensional complex vector space and ϕ is positive definite, then there exists a basis for which ϕ and ψ are diagonalized.
- (iv) Let $A, B \in \operatorname{Mat}_n(\mathbb{C})$ be Hermitian, and A positive definitive (i.e. $\mathbf{v}^{\dagger}A\mathbf{v} > 0$ for $\mathbf{v} \in V \setminus \{0\}$). Then there exists some invertible Q such that $Q^{\dagger}AQ$ and $Q^{\dagger}BQ$ are diagonal.

That's all for self-adjoint matrices. How about under an uncer-

Theorem. Let V be a finite-dimension provided provided by the vector space and $\alpha \in U(V)$ be unitary. Then V has an ext or translation of α eigenvectors.

Proof. By the final amendal theorem of algebra, there exists
$$\mathbf{v} \in V \setminus \{0\}$$
 and $\lambda \in \mathbb{C}$ is to that $\alpha \mathbf{v} = \lambda \mathbf{v}$. Now consider $W = \langle \mathbf{v} \rangle^{\perp}$. Then $V = W \perp \langle \mathbf{v} \rangle$.

We want to show α restricts to a (unitary) endomorphism of W. Let $\mathbf{w} \in W$. We need to show $\alpha(\mathbf{w})$ is orthogonal to \mathbf{v} . We have

$$(\alpha \mathbf{w}, \mathbf{v}) = (\mathbf{w}, \alpha^{-1}\mathbf{v}) = (\mathbf{w}, \lambda^{-1}\mathbf{v}) = 0.$$

So $\alpha(\mathbf{w}) \in W$ and $\alpha|_W \in \text{End}(W)$. Also, $\alpha|_W$ is unitary since α is. So by induction on dim V, W has an orthonormal basis of α eigenvectors. If we add $\mathbf{v}/||\mathbf{v}||$ to this basis, we get an orthonormal basis of V itself comprised of α eigenvectors.

This theorem and the analogous one for self-adjoint endomorphisms have a common generalization, at least for complex inner product spaces. The key fact that leads to the existence of an orthonormal basis of eigenvectors is that α and α^* commute. This is clearly a necessary condition, since if α is diagonalizable, then α^* is diagonal in the same basis (since it is just the transpose (and conjugate)), and hence they commute. It turns out this is also a sufficient condition, as you will show in example sheet 4.