

Part IB — Linear Algebra

Based on lectures by S. J. Wadsley

Notes taken by Dexter Chua

Michaelmas 2015

These notes are not endorsed by the lecturers, and I have modified them (often significantly) after lectures. They are nowhere near accurate representations of what was actually lectured, and in particular, all errors are almost surely mine.

Definition of a vector space (over \mathbb{R} or \mathbb{C}), subspaces, the space spanned by a subset. Linear independence, bases, dimension. Direct sums and complementary subspaces. [3]

Linear maps, isomorphisms. Relation between rank and nullity. The space of linear maps from U to V , representation by matrices. Change of basis. Row rank and column rank. [4]

Determinant and trace of a square matrix. Determinant of a product of two matrices and of the inverse matrix. Determinant of an endomorphism. The adjugate matrix. [3]

Eigenvalues and eigenvectors. Diagonal and triangular forms. Characteristic and minimal polynomials. Cayley-Hamilton Theorem over \mathbb{C} . Algebraic and geometric multiplicity of eigenvalues. Statement and illustration of Jordan normal form. [4]

Dual of a finite-dimensional vector space, dual bases and maps. Matrix representation, rank and determinant of dual map. [2]

Bilinear forms. Matrix representation, change of basis. Symmetric forms and their link with quadratic forms. Diagonalisation of quadratic forms. Law of inertia, classification by rank and signature. Complex Hermitian forms. [4]

Inner product spaces, orthonormal sets, orthogonal projection, $V = W \oplus W^\perp$. Gram-Schmidt orthogonalisation. Adjoints. Diagonalisation of Hermitian matrices. Orthogonality of eigenvectors and properties of eigenvalues. [4]

Contents

0	Introduction	3
1	Vector spaces	4
1.1	Definitions and examples	4
1.2	Linear independence, bases and the Steinitz exchange lemma . .	6
1.3	Direct sums	13
2	Linear maps	15
2.1	Definitions and examples	15
2.2	Linear maps and matrices	19
2.3	The first isomorphism theorem and the rank-nullity theorem . .	21
2.4	Change of basis	24
2.5	Elementary matrix operations	27
3	Duality	29
3.1	Dual space	29
3.2	Dual maps	32
4	Bilinear forms I	37
5	Determinants of matrices	41
6	Endomorphisms	49
6.1	Invariants	49
6.2	The minimal polynomial	53
6.2.1	Aside on polynomials	53
6.2.2	Minimal polynomial	54
6.3	The Cayley-Hamilton theorem	57
6.4	Multiplicities of eigenvalues and Jordan normal form	63
7	Bilinear forms II	72
7.1	Symmetric bilinear forms and quadratic forms	72
7.2	Hermitian form	80
8	Inner product spaces	83
8.1	Definitions and basic properties	83
8.2	Gram-Schmidt orthogonalization	86
8.3	Adjoint, orthogonal and unitary maps	89
8.4	Spectral theory	92

0 Introduction

In IA Vectors and Matrices, we have learnt about vectors (and matrices) in a rather applied way. A vector was just treated as a “list of numbers” representing a point in space. We used these to represent lines, conics, planes and many other geometrical notions. A matrix is treated as a “physical operation” on vectors that stretches and rotates them. For example, we studied the properties of rotations, reflections and shears of space. We also used matrices to express and solve systems of linear equations. We mostly took a practical approach in the course.

In IB Linear Algebra, we are going to study vectors in an abstract setting. Instead of treating vectors as “lists of numbers”, we view them as things we can add and scalar-multiply. We will write down axioms governing how these operations should behave, just like how we wrote down the axioms of group theory. Instead of studying matrices as an array of numbers, we instead look at linear maps between vector spaces abstractly.

In the course, we will, of course, prove that this abstract treatment of linear algebra is just “the same as” our previous study of “vectors as a list of numbers”. Indeed, in certain cases, results are much more easily proved by working with matrices (as an array of numbers) instead of abstract linear maps, and we don’t shy away from doing so. However, most of the time, looking at these abstractly will provide a much better fundamental understanding of how things work.

1 Vector spaces

1.1 Definitions and examples

Notation. We will use \mathbb{F} to denote an arbitrary field, usually \mathbb{R} or \mathbb{C} .

Intuitively, a vector space V over a field \mathbb{F} (or an \mathbb{F} -vector space) is a space with two operations:

- We can add two vectors $\mathbf{v}_1, \mathbf{v}_2 \in V$ to obtain $\mathbf{v}_1 + \mathbf{v}_2 \in V$.
- We can multiply a scalar $\lambda \in \mathbb{F}$ with a vector $\mathbf{v} \in V$ to obtain $\lambda\mathbf{v} \in V$.

Of course, these two operations must satisfy certain axioms before we can call it a vector space. However, before going into these details, we first look at a few examples of vector spaces.

Example.

- (i) $\mathbb{R}^n = \{\text{column vectors of length } n \text{ with coefficients in } \mathbb{R}\}$ with the usual addition and scalar multiplication is a vector space.

An $m \times n$ matrix A with coefficients in \mathbb{R} can be viewed as a linear map from \mathbb{R}^n to \mathbb{R}^m via $\mathbf{v} \mapsto A\mathbf{v}$.

This is a motivational example for vector spaces. When confused about definitions, we can often think what the definition means in terms of \mathbb{R}^n and matrices to get some intuition.

- (ii) Let X be a set and define $\mathbb{R}^X = \{f : X \rightarrow \mathbb{R}\}$ with addition $(f + g)(x) = f(x) + g(x)$ and scalar multiplication $(\lambda f)(x) = \lambda f(x)$. This is a vector space.

More generally, if V is a vector space, X is a set, we can define $V^X = \{f : X \rightarrow V\}$ with addition and scalar multiplication as above.

- (iii) Let $[a, b] \subseteq \mathbb{R}$ be a closed interval, then

$$C([a, b], \mathbb{R}) = \{f \in \mathbb{R}^{[a, b]} : f \text{ is continuous}\}$$

is a vector space, with operations as above. We also have

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

- (iv) The set of $m \times n$ matrices with coefficients in \mathbb{R} is a vector space, using componentwise addition and scalar multiplication, is a vector space.

Of course, we cannot take a random set, define some random operations called addition and scalar multiplication, and call it a vector space. These operations have to behave sensibly.

Definition (Vector space). An \mathbb{F} -vector space is an (additive) abelian group V together with a function $\mathbb{F} \times V \rightarrow V$, written $(\lambda, \mathbf{v}) \mapsto \lambda\mathbf{v}$, such that

$$(i) \quad \lambda(\mu\mathbf{v}) = \lambda\mu\mathbf{v} \text{ for all } \lambda, \mu \in \mathbb{F}, \mathbf{v} \in V \quad (\text{associativity})$$

$$(ii) \quad \lambda(\mathbf{u} + \mathbf{v}) = \lambda\mathbf{u} + \lambda\mathbf{v} \text{ for all } \lambda \in \mathbb{F}, \mathbf{u}, \mathbf{v} \in V \quad (\text{distributivity in } V)$$

- (iii) $(\lambda + \mu)\mathbf{v} = \lambda\mathbf{v} + \mu\mathbf{v}$ for all $\lambda, \mu \in \mathbb{F}$, $\mathbf{v} \in V$ (distributivity in \mathbb{F})
- (iv) $1\mathbf{v} = \mathbf{v}$ for all $\mathbf{v} \in V$ (identity)

We always write $\mathbf{0}$ for the additive identity in V , and call this the identity. By abuse of notation, we also write 0 for the trivial vector space $\{\mathbf{0}\}$.

In a general vector space, there is no notion of “coordinates”, length, angle or distance. For example, it would be difficult to assign these quantities to the vector space of real-valued continuous functions in $[a, b]$.

From the axioms, there are a few results we can immediately prove.

Proposition. In any vector space V , $0\mathbf{v} = \mathbf{0}$ for all $\mathbf{v} \in V$, and $(-1)\mathbf{v} = -\mathbf{v}$, where $-\mathbf{v}$ is the additive inverse of \mathbf{v} .

Proof is left as an exercise.

In mathematics, whenever we define “something”, we would also like to define a “sub-something”. In the case of vector spaces, this is a subspace.

Definition (Subspace). If V is an \mathbb{F} -vector space, then $U \subseteq V$ is an (\mathbb{F} -linear) *subspace* if

- (i) $\mathbf{u}, \mathbf{v} \in U$ implies $\mathbf{u} + \mathbf{v} \in U$.
- (ii) $\mathbf{u} \in U, \lambda \in \mathbb{F}$ implies $\lambda\mathbf{u} \in U$.
- (iii) $\mathbf{0} \in U$.

These conditions can be expressed more concisely as “ U is non-empty and if $\lambda, \mu \in \mathbb{F}, \mathbf{u}, \mathbf{v} \in U$, then $\lambda\mathbf{u} + \mu\mathbf{v} \in U$ ”.

Alternatively, U is a subspace of V if it is itself a vector space, inheriting the operations from V .

We sometimes write $U \leq V$ if U is a subspace of V .

Example.

- (i) $\{(x_1, x_2, x_3) \in \mathbb{R}^3 : x_1 + x_2 + x_3 = t\}$ is a subspace of \mathbb{R}^3 iff $t = 0$.
- (ii) Let X be a set. We define the *support* of f in \mathbb{F}^X to be $\text{supp}(f) = \{x \in X : f(x) \neq 0\}$. Then the set of functions with finite support forms a vector subspace. This is since $\text{supp}(f + g) \subseteq \text{supp}(f) \cup \text{supp}(g)$, $\text{supp}(\lambda f) = \text{supp}(f)$ (for $\lambda \neq 0$) and $\text{supp}(0) = \emptyset$.

If we have two subspaces U and V , there are several things we can do with them. For example, we can take the intersection $U \cap V$. We will shortly show that this will be a subspace. However, taking the union will in general not produce a vector space. Instead, we need the sum:

Definition (Sum of subspaces). Suppose U, W are subspaces of an \mathbb{F} vector space V . The *sum* of U and V is

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

Proposition. Let U, W be subspaces of V . Then $U + W$ and $U \cap W$ are subspaces.

Proof. Let $\mathbf{u}_i + \mathbf{w}_i \in U + W$, $\lambda, \mu \in \mathbb{F}$. Then

$$\lambda(\mathbf{u}_1 + \mathbf{w}_1) + \mu(\mathbf{u}_2 + \mathbf{w}_2) = (\lambda\mathbf{u}_1 + \mu\mathbf{u}_2) + (\lambda\mathbf{w}_1 + \mu\mathbf{w}_2) \in U + W.$$

Similarly, if $\mathbf{v}_i \in U \cap W$, then $\lambda\mathbf{v}_1 + \mu\mathbf{v}_2 \in U$ and $\lambda\mathbf{v}_1 + \mu\mathbf{v}_2 \in W$. So $\lambda\mathbf{v}_1 + \mu\mathbf{v}_2 \in U \cap W$.

Both $U \cap W$ and $U + W$ contain $\mathbf{0}$, and are non-empty. So done. \square

In addition to sub-somethings, we often have quotient-somethings as well.

Definition (Quotient spaces). Let V be a vector space, and $U \subseteq V$ a subspace. Then the quotient group V/U can be made into a vector space called the *quotient space*, where scalar multiplication is given by $(\lambda, \mathbf{v} + U) = (\lambda\mathbf{v}) + U$.

This is well defined since if $\mathbf{v} + U = \mathbf{w} + U \in V/U$, then $\mathbf{v} - \mathbf{w} \in U$. Hence for $\lambda \in \mathbb{F}$, we have $\lambda\mathbf{v} - \lambda\mathbf{w} \in U$. So $\lambda\mathbf{v} + U = \lambda\mathbf{w} + U$.

1.2 Linear independence, bases and the Steinitz exchange lemma

Recall that in \mathbb{R}^n , we had the “standard basis” made of vectors of the form $\mathbf{e}_i = (0, \dots, 0, 1, 0, \dots, 0)$, with 1 in the i th component and 0 otherwise. We call this a *basis* because everything in \mathbb{R}^n can be (uniquely) written as a sum of (scalar multiples of) these basis elements. In other words, the whole \mathbb{R}^n is generated by taking sums and multiples of the basis elements.

We would like to capture this idea in general vector spaces. The most important result in this section is to prove that for any vector space V , any two basis must contain the same number of elements. This means we can define the “dimension” of a vector space as the number of elements in the basis.

While this result sounds rather trivial, it is a very important result. We will in fact prove a slightly stronger statement than what was stated above, and this ensures that the dimension of a vector space is well-behaved. For example, the subspace of a vector space has a smaller dimension than the larger space (at least when the dimensions are finite).

This is not the case when we study modules in IB Groups, Rings and Modules, which are generalizations of vector spaces. Not all modules have basis, which makes it difficult to define the dimension. Even for those that have basis, the behaviour of the “dimension” is complicated when, say, we take submodules. The existence and well-behavedness of basis and dimension is what makes linear algebra different from modules.

Definition (Span). Let V be a vector space over \mathbb{F} and $S \subseteq V$. The *span* of S is defined as

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

This is the smallest subspace of V containing S .

Note that the sums must be finite. We will not play with infinite sums, since the notion of convergence is not even well defined in a general vector space.

Example.

(i) Let $V = \mathbb{R}^3$ 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

$$\langle S \rangle = \left\{ \begin{pmatrix} a \\ b \\ b \end{pmatrix} : a, b \in \mathbb{R} \right\}.$$

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

(ii) Let X be a set, $x \in X$. Define the function $\delta x : X \rightarrow \mathbb{F}$ by

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

Then $\langle \delta x : x \in X \rangle$ is the set of all functions with finite support.

Definition (Spanning set). Let V be a vector space over \mathbb{F} and $S \subseteq V$. S spans V if $\langle S \rangle = V$.

Definition (Linear independence). Let V be a vector space over \mathbb{F} and $S \subseteq V$. Then S is *linearly independent* if whenever

$$\sum_{i=1}^n \lambda_i \mathbf{s}_i = \mathbf{0} \text{ with } \lambda_i \in \mathbb{F}, \mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_n \in S \text{ distinct,}$$

we must have $\lambda_i = 0$ for all i .

If S is not linearly independent, we say it is *linearly dependent*.

Definition (Basis). Let V be a vector space over \mathbb{F} and $S \subseteq V$. Then S is a *basis* for V if S is linearly independent and spans V .

Definition (Finite dimensional). A vector space is *finite dimensional* if there is a finite basis.

Ideally, we would want to define the *dimension* as the number of vectors in the basis. However, we must first show that this is well-defined. It is certainly plausible that a vector space has a basis of size 7 as well as a basis of size 3. We must show that this can never happen, which is something we'll do soon.

We will first have an example:

Example. Again, let $V = \mathbb{R}^3$ 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

$$1 \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + 2 \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} + (-1) \begin{pmatrix} 1 \\ 2 \\ 2 \end{pmatrix} = \mathbf{0}.$$

S also does not span V since $\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \notin \langle S \rangle$.

Note that no linearly independent set can contain $\mathbf{0}$, as $1 \cdot \mathbf{0} = \mathbf{0}$. We also have $\langle \emptyset \rangle = \{\mathbf{0}\}$ and \emptyset is a basis for this space.

There is an alternative way in which we can define linear independence.

Lemma. $S \subseteq V$ is linearly dependent if and only if there are distinct $\mathbf{s}_0, \dots, \mathbf{s}_n \in S$ and $\lambda_1, \dots, \lambda_n \in \mathbb{F}$ such that

$$\sum_{i=1}^n \lambda_i \mathbf{s}_i = \mathbf{s}_0.$$

Proof. If S is linearly dependent, then there are some $\lambda_1, \dots, \lambda_n \in \mathbb{F}$ not all zero and $\mathbf{s}_1, \dots, \mathbf{s}_n \in S$ such that $\sum \lambda_i \mathbf{s}_i = \mathbf{0}$. Wlog, let $\lambda_1 \neq 0$. Then

$$\mathbf{s}_1 = \sum_{i=2}^n -\frac{\lambda_i}{\lambda_1} \mathbf{s}_i.$$

Conversely, if $\mathbf{s}_0 = \sum_{i=1}^n \lambda_i \mathbf{s}_i$, then

$$(-1)\mathbf{s}_0 + \sum_{i=1}^n \lambda_i \mathbf{s}_i = \mathbf{0}.$$

So S is linearly dependent. □

This in turn gives an alternative characterization of what it means to be a basis:

Proposition. If $S = \{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ is a subset of V over \mathbb{F} , then it is a basis if and only if every $\mathbf{v} \in V$ can be written uniquely as a finite linear combination of elements in S , i.e. as

$$\mathbf{v} = \sum_{i=1}^n \lambda_i \mathbf{e}_i.$$

Proof. We can view this as a combination of two statements: it can be spanned in at least one way, and it can be spanned in at most one way. We will see that the first part corresponds to S spanning V , and the second part corresponds to S being linearly independent.

In fact, S spanning V is defined exactly to mean that every item $\mathbf{v} \in V$ can be written as a finite linear combination in at least one way.

Now suppose that S is linearly independent, and we have

$$\mathbf{v} = \sum_{i=1}^n \lambda_i \mathbf{e}_i = \sum_{i=1}^n \mu_i \mathbf{e}_i.$$

Then we have

$$\mathbf{0} = \mathbf{v} - \mathbf{v} = \sum_{i=1}^n (\lambda_i - \mu_i) \mathbf{e}_i.$$

Linear independence implies that $\lambda_i - \mu_i = 0$ for all i . Hence $\lambda_i = \mu_i$. So \mathbf{v} can be expressed in a unique way.

On the other hand, if S is not linearly independent, then we have

$$\mathbf{0} = \sum_{i=1}^n \lambda_i \mathbf{e}_i$$

where $\lambda_i \neq 0$ for some i . But we also know that

$$\mathbf{0} = \sum_{i=1}^n 0 \cdot \mathbf{e}_i.$$

So there are two ways to write $\mathbf{0}$ as a linear combination. So done. \square

Now we come to the key theorem:

Theorem (Steinitz exchange lemma). Let V be a vector space over \mathbb{F} , and $S = \{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ a finite linearly independent subset of V , and T a spanning subset of V . Then there is some $T' \subseteq T$ of order n such that $(T \setminus T') \cup S$ still spans V . In particular, $|T| \geq n$.

What does this actually say? This says if T is spanning and S is independent, there is a way of grabbing $|S|$ many elements away from T and replace them with S , and the result will still be spanning.

In some sense, the final remark is the most important part. It tells us that we cannot have a independent set larger than a spanning set, and most of our corollaries later will only use this remark.

This is sometimes stated in the following alternative way for $|T| < \infty$.

Corollary. Let $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ be a linearly independent subset of V , and suppose $\{\mathbf{f}_1, \dots, \mathbf{f}_m\}$ spans V . Then there is a re-ordering of the $\{\mathbf{f}_i\}$ such that $\{\mathbf{e}_1, \dots, \mathbf{e}_n, \mathbf{f}_{n+1}, \dots, \mathbf{f}_m\}$ spans V .

The proof is going to be slightly technical and notationally daunting. So it helps to give a brief overview of what we are going to do in words first. The idea is to do the replacement one by one. The first one is easy. Start with \mathbf{e}_1 . Since T is spanning, we can write

$$\mathbf{e}_1 = \sum \lambda_i \mathbf{t}_i$$

for some $\mathbf{t}_i \in T, \lambda_i \in \mathbb{F}$ non-zero. We then replace with \mathbf{t}_1 with \mathbf{e}_1 . The result is still spanning, since the above formula allows us to write \mathbf{t}_1 in terms of \mathbf{e}_1 and the other \mathbf{t}_i .

We continue inductively. For the r th element, we again write

$$\mathbf{e}_r = \sum \lambda_i \mathbf{t}_i.$$

We would like to just pick a random \mathbf{t}_i and replace it with \mathbf{e}_r . However, we cannot do this arbitrarily, since the lemma wants us to replace something *in* T with with \mathbf{e}_r . After all that replacement procedure before, some of the \mathbf{t}_i might have actually come from S .

This is where the linear independence of S kicks in. While some of the \mathbf{t}_i might be from S , we cannot possibly have all *all* of them being from S , or else this violates the linear independence of S . Hence there is something genuinely from T , and we can safely replace it with \mathbf{e}_r .

We now write this argument properly and formally.

Proof. Suppose that we have already found $T'_r \subseteq T$ of order $0 \leq r < n$ such that

$$T_r = (T \setminus T'_r) \cup \{\mathbf{e}_1, \dots, \mathbf{e}_r\}$$

spans V .

(Note that the case $r = 0$ is trivial, since we can take $T'_r = \emptyset$, and the case $r = n$ is the theorem which we want to achieve.)

Suppose we have these. Since T_r spans V , we can write

$$\mathbf{e}_{r+1} = \sum_{i=1}^k \lambda_i \mathbf{t}_i, \quad \lambda_i \in \mathbb{F}, \mathbf{t}_i \in T_r.$$

We know that the \mathbf{e}_i are linearly independent, so not all \mathbf{t}_i 's are \mathbf{e}_i 's. So there is some j such that $\mathbf{t}_j \in (T \setminus T'_r)$. We can write this as

$$\mathbf{t}_j = \frac{1}{\lambda_j} \mathbf{e}_{r+1} + \sum_{i \neq j} -\frac{\lambda_i}{\lambda_j} \mathbf{t}_i.$$

We let $T'_{r+1} = T'_r \cup \{\mathbf{t}_j\}$ of order $r + 1$, and

$$T_{r+1} = (T \setminus T'_{r+1}) \cup \{\mathbf{e}_1, \dots, \mathbf{e}_{r+1}\} = (T_r \setminus \{\mathbf{t}_j\}) \cup \{\mathbf{e}_{r+1}\}$$

Since \mathbf{t}_j is in the span of $T_r \cup \{\mathbf{e}_{r+1}\}$, we have $\mathbf{t}_j \in \langle T_{r+1} \rangle$. So

$$V \supseteq \langle T_{r+1} \rangle \supseteq \langle T_r \rangle = V.$$

So $\langle T_{r+1} \rangle = V$.

Hence we can inductively find T_n . □

From this lemma, we can immediately deduce a lot of important corollaries.

Corollary. Suppose V is a vector space over \mathbb{F} with a basis of order n . Then

- (i) Every basis of V has order n .
- (ii) Any linearly independent set of order n is a basis.
- (iii) Every spanning set of order n is a basis.
- (iv) Every finite spanning set contains a basis.
- (v) Every linearly independent subset of V can be extended to basis.

Proof. Let $S = \{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ be the basis for V .

- (i) Suppose T is another basis. Since S is independent and T is spanning, $|T| \geq |S|$.

The other direction is less trivial, since T might be infinite, and Steinitz does not immediately apply. Instead, we argue as follows: since T is linearly independent, every finite subset of T is independent. Also, S is spanning. So every finite subset of T has order at most $|S|$. So $|T| \leq |S|$. So $|T| = |S|$.

- (ii) Suppose now that T is a linearly independent subset of order n , but $\langle T \rangle \neq V$. Then there is some $\mathbf{v} \in V \setminus \langle T \rangle$. We now show that $T \cup \{\mathbf{v}\}$ is independent. Indeed, if

$$\lambda_0 \mathbf{v} + \sum_{i=1}^m \lambda_i \mathbf{t}_i = \mathbf{0}$$

with $\lambda_i \in \mathbb{F}$, $\mathbf{t}_1, \dots, \mathbf{t}_m \in T$ distinct, then

$$\lambda_0 \mathbf{v} = \sum_{i=1}^m (-\lambda_i) \mathbf{t}_i.$$

Then $\lambda_0 \mathbf{v} \in \langle T \rangle$. So $\lambda_0 = 0$. As T is linearly independent, we have $\lambda_0 = \dots = \lambda_m = 0$. So $T \cup \{\mathbf{v}\}$ is a linearly independent subset of size $> n$. This is a contradiction since S is a spanning set of size n .

- (iii) Let T be a spanning set of order n . If T were linearly dependent, then there is some $\mathbf{t}_0, \dots, \mathbf{t}_m \in T$ distinct and $\lambda_1, \dots, \lambda_m \in \mathbb{F}$ such that

$$\mathbf{t}_0 = \sum \lambda_i \mathbf{t}_i.$$

So $\mathbf{t}_0 \in \langle T \setminus \{\mathbf{t}_0\} \rangle$, i.e. $\langle T \setminus \{\mathbf{t}_0\} \rangle = V$. So $T \setminus \{\mathbf{t}_0\}$ is a spanning set of order $n - 1$, which is a contradiction.

- (iv) Suppose T is any finite spanning set. Let $T' \subseteq T$ be a spanning set of least possible size. This exists because T is finite. If $|T'|$ has size n , then done by (iii). Otherwise by the Steinitz exchange lemma, it has size $|T'| > n$. So T' must be linearly dependent because S is spanning. So there is some $\mathbf{t}_0, \dots, \mathbf{t}_m \in T$ distinct and $\lambda_1, \dots, \lambda_m \in \mathbb{F}$ such that $\mathbf{t}_0 = \sum \lambda_i \mathbf{t}_i$. Then $T' \setminus \{\mathbf{t}_0\}$ is a smaller spanning set. Contradiction.

- (v) Suppose T is a linearly independent set. Since S spans, there is some $S' \subseteq S$ of order $|T|$ such that $(S \setminus S') \cup T$ spans V by the Steinitz exchange lemma. So by (ii), $(S \setminus S') \cup T$ is a basis of V containing T . \square

Note that the last part is where we actually use the full result of Steinitz. Finally, we can use this to define the dimension.

Definition (Dimension). If V is a vector space over \mathbb{F} with finite basis S , then the *dimension* of V , written

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

By the corollary, $\dim V$ does not depend on the choice of S . However, it does depend on \mathbb{F} . For example, $\dim_{\mathbb{C}} \mathbb{C} = 1$ (since $\{1\}$ is a basis), but $\dim_{\mathbb{R}} \mathbb{C} = 2$ (since $\{1, i\}$ is a basis).

After defining the dimension, we can prove a few things about dimensions.

Lemma. If V is a finite dimensional vector space over \mathbb{F} , $U \subseteq V$ is a proper subspace, then U is finite dimensional and $\dim U < \dim V$.

Proof. Every linearly independent subset of V has size at most $\dim V$. So let $S \subseteq U$ be a linearly independent subset of largest size. We want to show that S spans U and $|S| < \dim V$.

If $\mathbf{v} \in V \setminus \langle S \rangle$, then $S \cup \{\mathbf{v}\}$ is linearly independent. So $\mathbf{v} \notin U$ by maximality of S . This means that $\langle S \rangle = U$.

Since $U \neq V$, there is some $\mathbf{v} \in V \setminus U = V \setminus \langle S \rangle$. So $S \cup \{\mathbf{v}\}$ is a linearly independent subset of order $|S| + 1$. So $|S| + 1 \leq \dim V$. In particular, $\dim U = |S| < \dim V$. \square

Proposition. If U, W are subspaces of a finite dimensional vector space V , then

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

The proof is not hard, as long as we manage to pick the right basis to do the proof. This is our slogan:

When you choose a basis, always choose the right basis.

We need a basis for all four of them, and we want to compare the bases. So we want to pick bases that are compatible.

Proof. Let $R = \{\mathbf{v}_1, \dots, \mathbf{v}_r\}$ be a basis for $U \cap W$. This is a linearly independent subset of U . So we can extend it to be a basis of U by

$$S = \{\mathbf{v}_1, \dots, \mathbf{v}_r, \mathbf{u}_{r+1}, \dots, \mathbf{u}_s\}.$$

Similarly, for W , we can obtain a basis

$$T = \{\mathbf{v}_1, \dots, \mathbf{v}_r, \mathbf{w}_{r+1}, \dots, \mathbf{w}_t\}.$$

We want to show that $\dim(U + W) = s + t - r$. It is sufficient to prove that $S \cup T$ is a basis for $U + W$.

We first show spanning. Suppose $\mathbf{u} + \mathbf{w} \in U + W$, $\mathbf{u} \in U, \mathbf{w} \in W$. Then $\mathbf{u} \in \langle S \rangle$ and $\mathbf{w} \in \langle T \rangle$. So $\mathbf{u} + \mathbf{w} \in \langle S \cup T \rangle$. So $U + W = \langle S \cup T \rangle$.

To show linear independence, suppose we have a linear relation

$$\sum_{i=1}^r \lambda_i \mathbf{v}_i + \sum_{j=r+1}^s \mu_j \mathbf{u}_j + \sum_{k=r+1}^t \nu_k \mathbf{w}_k = \mathbf{0}.$$

So

$$\sum \lambda_i \mathbf{v}_i + \sum \mu_j \mathbf{u}_j = - \sum \nu_k \mathbf{w}_k.$$

Since the left hand side is something in U , and the right hand side is something in W , they both lie in $U \cap W$.

Since S is a basis of U , there is only one way of writing the left hand vector as a sum of \mathbf{v}_i and \mathbf{u}_j . However, since R is a basis of $U \cap W$, we can write the left hand vector just as a sum of \mathbf{v}_i 's. So we must have $\mu_j = 0$ for all j . Then we have

$$\sum \lambda_i \mathbf{v}_i + \sum \nu_k \mathbf{w}_k = \mathbf{0}.$$

Finally, since T is linearly independent, $\lambda_i = \nu_k = 0$ for all i, k . So $S \cup T$ is linearly independent. \square

Proposition. If V is a finite dimensional vector space over \mathbb{F} and $U \cup V$ is a subspace, then

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

We can view this as a linear algebra version of Lagrange's theorem. Combined with the first isomorphism theorem for vector spaces, this gives the rank-nullity theorem.

Proof. Let $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$ be a basis for U and extend this to a basis $\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{v}_{m+1}, \dots, \mathbf{v}_n\}$ for V . We want to show that $\{\mathbf{v}_{m+1} + U, \dots, \mathbf{v}_n + U\}$ 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 $\{\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{v}_{m+1}, \dots, \mathbf{v}_n\}$ is a basis for V , we must have $\lambda_i = \mu_j = 0$ for all i, j . So $\{\mathbf{v}_i + U\}$ is linearly independent. \square

1.3 Direct sums

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

Definition ((Internal) direct sum). Suppose V is a vector space over \mathbb{F} and $U, W \subseteq V$ are subspaces. We say that V is the (*internal*) *direct sum* of U and W if

$$(i) \quad U + W = V$$

$$(ii) \quad U \cap W = \mathbf{0}.$$

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:

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

The difference between these two definitions is that the first is decomposing V into smaller spaces, while the second is building a bigger space based on two spaces.

Note, however, that the external direct sum $U \oplus W$ is the internal direct sum of U and W viewed as subspaces of $U \oplus W$, i.e. as the internal direct sum of $\{(\mathbf{u}, \mathbf{0}) : \mathbf{u} \in U\}$ and $\{(\mathbf{0}, \mathbf{v}) : \mathbf{v} \in W\}$. So these two are indeed compatible notions, and this is why we give them the same name and notation.

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

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

if every $\mathbf{v} \in V$ can be written uniquely as $\mathbf{v} = \sum \mathbf{u}_i$ with $\mathbf{u}_i \in U_i$.

This can be extended to an infinite sum with the same definition, just noting that the sum $\mathbf{v} = \sum \mathbf{u}_i$ has to be finite.

For more details, see example sheet 1 Q. 10, where we prove in particular that $\dim V = \sum \dim U_i$.

Definition ((Multiple) (external) direct sum). If U_1, \dots, U_n are vector spaces over \mathbb{F} , the external direct sum is

$$U_1 \oplus \dots \oplus U_n = \bigoplus_{i=1}^n U_i = \{(\mathbf{u}_1, \dots, \mathbf{u}_n) : \mathbf{u}_i \in U_i\},$$

with pointwise operations.

This can be made into an infinite sum if we require that all but finitely many of the \mathbf{u}_i have to be zero.

2 Linear maps

In mathematics, apart from studying objects, we would like to study functions between objects as well. In particular, we would like to study functions that respect the structure of the objects. With vector spaces, the kinds of functions we are interested in are *linear maps*.

2.1 Definitions and examples

Definition (Linear map). Let U, V be vector spaces over \mathbb{F} . Then $\alpha : U \rightarrow V$ is a *linear map* if

- (i) $\alpha(\mathbf{u}_1 + \mathbf{u}_2) = \alpha(\mathbf{u}_1) + \alpha(\mathbf{u}_2)$ for all $\mathbf{u}_i \in U$.
- (ii) $\alpha(\lambda\mathbf{u}) = \lambda\alpha(\mathbf{u})$ for all $\lambda \in \mathbb{F}, \mathbf{u} \in U$.

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

There are a few things we should take note of:

- If we are lazy, we can combine the two requirements to the single requirement that

$$\alpha(\lambda\mathbf{u}_1 + \mu\mathbf{u}_2) = \lambda\alpha(\mathbf{u}_1) + \mu\alpha(\mathbf{u}_2).$$

- It is easy to see that if α is linear, then it is a group homomorphism (if we view vector spaces as groups). In particular, $\alpha(\mathbf{0}) = \mathbf{0}$.
- If we want to stress the field \mathbb{F} , we say that α is \mathbb{F} -linear. For example, complex conjugation is a map $\mathbb{C} \rightarrow \mathbb{C}$ that is \mathbb{R} -linear but not \mathbb{C} -linear.

Example.

- (i) Let A be an $n \times m$ matrix with coefficients in \mathbb{F} . We will write $A \in M_{n,m}(\mathbb{F})$. Then $\alpha : \mathbb{F}^m \rightarrow \mathbb{F}^n$ defined by $\mathbf{v} \mapsto A\mathbf{v}$ is linear.

Recall matrix multiplication is defined by: if A_{ij} is the ij th coefficient of A , then the i th coefficient of $A\mathbf{v}$ is $A_{ij}v_j$. So we have

$$\begin{aligned} \alpha(\lambda\mathbf{u} + \mu\mathbf{v})_i &= \sum_{j=1}^m A_{ij}(\lambda u_j + \mu v_j) \\ &= \lambda \sum_{j=1}^m A_{ij}u_j + \mu \sum_{j=1}^m A_{ij}v_j \\ &= \lambda\alpha(\mathbf{u})_i + \mu\alpha(\mathbf{v})_i. \end{aligned}$$

So α is linear.

- (ii) Let X be a set and $g \in \mathbb{F}^X$. Then we define $m_g : \mathbb{F}^X \rightarrow \mathbb{F}^X$ by $m_g(f)(x) = g(x)f(x)$. Then m_g is linear. For example, $f(x) \mapsto 2x^2f(x)$ is linear.
- (iii) Integration $I : (C([a, b]), \mathbb{R}) \rightarrow (C([a, b]), \mathbb{R})$ defined by $f \mapsto \int_a^x f(t) dt$ is linear.
- (iv) Differentiation $D : (C^\infty([a, b]), \mathbb{R}) \rightarrow (C^\infty([a, b]), \mathbb{R})$ by $f \mapsto f'$ is linear.

(v) If $\alpha, \beta \in \mathcal{L}(U, V)$, then $\alpha + \beta$ defined by $(\alpha + \beta)(\mathbf{u}) = \alpha(\mathbf{u}) + \beta(\mathbf{u})$ is linear. Also, if $\lambda \in \mathbb{F}$, then $\lambda\alpha$ defined by $(\lambda\alpha)(\mathbf{u}) = \lambda(\alpha(\mathbf{u}))$ is also linear.

In this way, $\mathcal{L}(U, V)$ is also a vector space over \mathbb{F} .

(vi) Composition of linear maps is linear. Using this, we can show that many things are linear, like differentiating twice, or adding and then multiplying linear maps.

Just like everything else, we want to define isomorphisms.

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

If there exists an isomorphism $U \rightarrow V$, we say U and V are *isomorphic*, and write $U \cong V$.

Lemma. If U and V are vector spaces over \mathbb{F} and $\alpha : U \rightarrow V$, then α is an isomorphism iff α is a bijective linear map.

Proof. If α is an isomorphism, then it is clearly bijective since it has an inverse function.

Suppose α is a linear bijection. Then as a function, it has an inverse $\beta : V \rightarrow U$. We want to show that this is linear. Let $\mathbf{v}_1, \mathbf{v}_2 \in V$, $\lambda, \mu \in \mathbb{F}$. We have

$$\alpha\beta(\lambda\mathbf{v}_1 + \mu\mathbf{v}_2) = \lambda\mathbf{v}_1 + \mu\mathbf{v}_2 = \lambda\alpha\beta(\mathbf{v}_1) + \mu\alpha\beta(\mathbf{v}_2) = \alpha(\lambda\beta(\mathbf{v}_1) + \mu\beta(\mathbf{v}_2)).$$

Since α is injective, we have

$$\beta(\lambda\mathbf{v}_1 + \mu\mathbf{v}_2) = \lambda\beta(\mathbf{v}_1) + \mu\beta(\mathbf{v}_2).$$

So β is linear. □

Definition (Image and kernel). Let $\alpha : U \rightarrow V$ be a linear map. Then the *image* of α is

$$\text{im } \alpha = \{\alpha(\mathbf{u}) : \mathbf{u} \in U\}.$$

The *kernel* of α is

$$\ker \alpha = \{\mathbf{u} : \alpha(\mathbf{u}) = \mathbf{0}\}.$$

It is easy to show that these are subspaces of V and U respectively.

Example.

(i) Let $A \in M_{m,n}(\mathbb{F})$ and $\alpha : \mathbb{F}^n \rightarrow \mathbb{F}^m$ be the linear map $\mathbf{v} \mapsto A\mathbf{v}$. Then the system of linear equations

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

has a solution iff $(b_1, \dots, b_n) \in \text{im } \alpha$.

The kernel of α contains all solutions to $\sum_j A_{ij}x_j = 0$.

(ii) Let $\beta : C^\infty(\mathbb{R}, \mathbb{R}) \rightarrow C^\infty(\mathbb{R}, \mathbb{R})$ that sends

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

for some $p, q \in C^\infty(\mathbb{R}, \mathbb{R})$.

Then if $y(t) \in \text{im } \beta$, then there is a solution (in $C^\infty(\mathbb{R}, \mathbb{R})$) to the differential equation

$$f''(t) + p(t)f'(t) + q(t)f(t) = y(t).$$

Similarly, $\ker \beta$ contains the solutions to the homogeneous equation

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

If two vector spaces are isomorphic, then it is not too surprising that they have the same dimension, since isomorphic spaces are “the same”. Indeed this is what we are going to show.

Proposition. Let $\alpha : U \rightarrow V$ be an \mathbb{F} -linear map. Then

- (i) If α is injective and $S \subseteq U$ is linearly independent, then $\alpha(S)$ is linearly independent in V .
- (ii) If α is surjective and $S \subseteq U$ spans U , then $\alpha(S)$ spans V .
- (iii) If α is an isomorphism and $S \subseteq U$ is a basis, then $\alpha(S)$ is a basis for V .

Here (iii) immediately shows that two isomorphic spaces have the same dimension.

Proof.

- (i) We prove the contrapositive. Suppose that α is injective and $\alpha(S)$ is linearly dependent. So there are $\mathbf{s}_0, \dots, \mathbf{s}_n \in S$ distinct and $\lambda_1, \dots, \lambda_n \in \mathbb{F}$ not all zero such that

$$\alpha(\mathbf{s}_0) = \sum_{i=1}^n \lambda_i \alpha(\mathbf{s}_i) = \alpha \left(\sum_{i=1}^n \lambda_i \mathbf{s}_i \right).$$

Since α is injective, we must have

$$\mathbf{s}_0 = \sum_{i=1}^n \lambda_i \mathbf{s}_i.$$

This is a non-trivial relation of the \mathbf{s}_i in U . So S is linearly dependent.

- (ii) Suppose α is surjective and S spans U . Pick $\mathbf{v} \in V$. Then there is some $\mathbf{u} \in U$ such that $\alpha(\mathbf{u}) = \mathbf{v}$. Since S spans U , there is some $\mathbf{s}_1, \dots, \mathbf{s}_n \in S$ and $\lambda_1, \dots, \lambda_n \in \mathbb{F}$ such that

$$\mathbf{u} = \sum_{i=1}^n \lambda_i \mathbf{s}_i.$$

Then

$$\mathbf{v} = \alpha(\mathbf{u}) = \sum_{i=1}^n \lambda_i \alpha(\mathbf{s}_i).$$

So $\alpha(S)$ spans V .

(iii) Follows immediately from (i) and (ii). \square

Corollary. If U and V are finite-dimensional vector spaces over \mathbb{F} and $\alpha : U \rightarrow V$ is an isomorphism, then $\dim U = \dim V$.

Note that we restrict it to finite-dimensional spaces since we've only shown that dimensions are well-defined for finite dimensional spaces. Otherwise, the proof works just fine for infinite dimensional spaces.

Proof. Let S be a basis for U . Then $\alpha(S)$ is a basis for V . Since α is injective, $|S| = |\alpha(S)|$. So done. \square

How about the other way round? If two vector spaces have the same dimension, are they necessarily isomorphic? The answer is yes, at least for finite-dimensional ones.

However, we will not just prove that they are isomorphic. We will show that they are isomorphic in *many ways*.

Proposition. Suppose V is a \mathbb{F} -vector space of dimension $n < \infty$. Then writing $\mathbf{e}_1, \dots, \mathbf{e}_n$ for the standard basis of \mathbb{F}^n , there is a bijection

$$\Phi : \{\text{isomorphisms } \mathbb{F}^n \rightarrow V\} \rightarrow \{(\text{ordered}) \text{ basis } (\mathbf{v}_1, \dots, \mathbf{v}_n) \text{ for } V\},$$

defined by

$$\alpha \mapsto (\alpha(\mathbf{e}_1), \dots, \alpha(\mathbf{e}_n)).$$

Proof. We first make sure this is indeed a function — if α is an isomorphism, then from our previous proposition, we know that it sends a basis to a basis. So $(\alpha(\mathbf{e}_1), \dots, \alpha(\mathbf{e}_n))$ is indeed a basis for V .

We now have to prove surjectivity and injectivity.

Suppose $\alpha, \beta : \mathbb{F}^n \rightarrow V$ are isomorphism such that $\Phi(\alpha) = \Phi(\beta)$. In other words, $\alpha(\mathbf{e}_i) = \beta(\mathbf{e}_i)$ for all i . We want to show that $\alpha = \beta$. We have

$$\alpha \left(\begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \right) = \alpha \left(\sum_{i=1}^n x_i \mathbf{e}_i \right) = \sum x_i \alpha(\mathbf{e}_i) = \sum x_i \beta(\mathbf{e}_i) = \beta \left(\begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \right).$$

Hence $\alpha = \beta$.

Next, suppose that $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ is an ordered basis for V . Then define

$$\alpha \left(\begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \right) = \sum x_i \mathbf{v}_i.$$

It is easy to check that this is well-defined and linear. We also know that α is injective since $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ is linearly independent. So if $\sum x_i \mathbf{v}_i = \sum y_i \mathbf{v}_i$, then $x_i = y_i$. Also, α is surjective since $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ spans V . So α is an isomorphism, and by construction $\Phi(\alpha) = (\mathbf{v}_1, \dots, \mathbf{v}_n)$. \square

2.2 Linear maps and matrices

Recall that our first example of linear maps is matrices acting on \mathbb{F}^n . We will show that in fact, *all* linear maps come from matrices. Since we know that all vector spaces are isomorphic to \mathbb{F}^n , this means we can represent arbitrary linear maps on vector spaces by matrices.

This is a useful result, since it is sometimes easier to argue about matrices than linear maps.

Proposition. Suppose U, V are vector spaces over \mathbb{F} and $S = \{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ is a basis for U . Then every function $f : S \rightarrow V$ extends uniquely to a linear map $U \rightarrow V$.

The slogan is “to define a linear map, it suffices to define its values on a basis”.

Proof. For uniqueness, first suppose $\alpha, \beta : U \rightarrow V$ are linear and extend $f : S \rightarrow V$. We have sort-of proved this already just now.

If $\mathbf{u} \in U$, we can write $\mathbf{u} = \sum_{i=1}^n u_i \mathbf{e}_i$ with $u_i \in \mathbb{F}$ since S spans. Then

$$\alpha(\mathbf{u}) = \alpha\left(\sum u_i \mathbf{e}_i\right) = \sum u_i \alpha(\mathbf{e}_i) = \sum u_i f(\mathbf{e}_i).$$

Similarly,

$$\beta(\mathbf{u}) = \sum u_i f(\mathbf{e}_i).$$

So $\alpha(\mathbf{u}) = \beta(\mathbf{u})$ for every \mathbf{u} . So $\alpha = \beta$.

For existence, if $\mathbf{u} \in U$, we can write $\mathbf{u} = \sum u_i \mathbf{e}_i$ in a unique way. So defining

$$\alpha(\mathbf{u}) = \sum u_i f(\mathbf{e}_i)$$

is unambiguous. To show linearity, let $\lambda, \mu \in \mathbb{F}$, $\mathbf{u}, \mathbf{v} \in U$. Then

$$\begin{aligned} \alpha(\lambda\mathbf{u} + \mu\mathbf{v}) &= \alpha\left(\sum(\lambda u_i + \mu v_i)\mathbf{e}_i\right) \\ &= \sum(\lambda u_i + \mu v_i)f(\mathbf{e}_i) \\ &= \lambda\left(\sum u_i f(\mathbf{e}_i)\right) + \mu\left(\sum v_i f(\mathbf{e}_i)\right) \\ &= \lambda\alpha(\mathbf{u}) + \mu\alpha(\mathbf{v}). \end{aligned}$$

Moreover, α does extend f . □

Corollary. If U and V are finite-dimensional vector spaces over \mathbb{F} with bases $(\mathbf{e}_1, \dots, \mathbf{e}_m)$ and $(\mathbf{f}_1, \dots, \mathbf{f}_n)$ respectively, then there is a bijection

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

sending A to the unique linear map α such that $\alpha(\mathbf{e}_i) = \sum a_{ji} \mathbf{f}_j$.

We can interpret this as follows: the i th column of A tells us how to write $\alpha(\mathbf{e}_i)$ in terms of the \mathbf{f}_j .

We can also draw a fancy diagram to display this result. Given a basis $\mathbf{e}_1, \dots, \mathbf{e}_m$, by our bijection, we get an isomorphism $s(\mathbf{e}_i) : U \rightarrow \mathbb{F}^m$. Similarly, we get an isomorphism $s(\mathbf{f}_i) : V \rightarrow \mathbb{F}^n$.

Since a matrix is a linear map $A : \mathbb{F}^m \rightarrow \mathbb{F}^n$, given a matrix A , we can produce a linear map $\alpha : U \rightarrow V$ via the following composition

$$U \xrightarrow{s(\mathbf{e}_i)} \mathbb{F}^m \xrightarrow{A} \mathbb{F}^n \xrightarrow{s(\mathbf{f}_i)^{-1}} V.$$

We can put this into a square:

$$\begin{array}{ccc} \mathbb{F}^m & \xrightarrow{A} & \mathbb{F}^n \\ s(\mathbf{e}_i) \uparrow & & \uparrow s(\mathbf{f}_i) \\ U & \xrightarrow{\alpha} & V \end{array}$$

Then the corollary tells us that every A gives rise to an α , and every α corresponds to an A that fit into this diagram.

Proof. If α is a linear map $U \rightarrow V$, then for each $1 \leq i \leq m$, we can write $\alpha(\mathbf{e}_i)$ uniquely as

$$\alpha(\mathbf{e}_i) = \sum_{j=1}^n a_{ji} \mathbf{f}_j$$

for some $a_{ji} \in \mathbb{F}$. This gives a matrix $A = (a_{ij})$. The previous proposition tells us that every matrix A arises in this way, and α is determined by A . \square

Definition (Matrix representation). We call the matrix corresponding to a linear map $\alpha \in \mathcal{L}(U, V)$ under the corollary the *matrix representing* α with respect to the bases $(\mathbf{e}_1, \dots, \mathbf{e}_m)$ and $(\mathbf{f}_1, \dots, \mathbf{f}_n)$.

It is an exercise to show that the bijection $\text{Mat}_{n,m}(\mathbb{F}) \rightarrow \mathcal{L}(U, V)$ is an isomorphism of the vector spaces and deduce that $\dim \mathcal{L}(U, V) = (\dim U)(\dim V)$.

Proposition. Suppose U, V, W are finite-dimensional vector spaces over \mathbb{F} with bases $R = (\mathbf{u}_1, \dots, \mathbf{u}_r)$, $S = (\mathbf{v}_1, \dots, \mathbf{v}_s)$ and $T = (\mathbf{w}_1, \dots, \mathbf{w}_t)$ respectively.

If $\alpha : U \rightarrow V$ and $\beta : V \rightarrow W$ are linear maps represented by A and B respectively (with respect to R, S and T), then $\beta\alpha$ is linear and represented by BA with respect to R and T .

$$\begin{array}{ccccc} \mathbb{F}^r & \xrightarrow{A} & \mathbb{F}^s & \xrightarrow{B} & \mathbb{F}^t \\ s(R) \uparrow & & \uparrow s(S) & & \uparrow s(T) \\ U & \xrightarrow{\alpha} & V & \xrightarrow{\beta} & W \end{array}$$

Proof. Verifying $\beta\alpha$ is linear is straightforward. Next we write $\beta\alpha(\mathbf{u}_i)$ as a linear

combination of $\mathbf{w}_1, \dots, \mathbf{w}_t$:

$$\begin{aligned}
 \beta\alpha(\mathbf{u}_i) &= \beta\left(\sum_k A_{ki}\mathbf{v}_k\right) \\
 &= \sum_k A_{ki}\beta(\mathbf{v}_k) \\
 &= \sum_k A_{ki}\sum_j B_{jk}\mathbf{w}_j \\
 &= \sum_j \left(\sum_k B_{jk}A_{ki}\right)\mathbf{w}_j \\
 &= \sum_j (BA)_{ji}\mathbf{w}_j
 \end{aligned}$$

□

2.3 The first isomorphism theorem and the rank-nullity theorem

The main theorem of this section is the *rank-nullity theorem*, which relates the dimensions of the kernel and image of a linear map. This is in fact an easy corollary of a stronger result, known as the *first isomorphism theorem*, which directly relates the kernel and image themselves. This first isomorphism is an exact analogy of that for groups, and should not be unfamiliar. We will also provide another proof that does not involve quotients.

Theorem (First isomorphism theorem). Let $\alpha : U \rightarrow V$ be a linear map. Then $\ker \alpha$ and $\operatorname{im} \alpha$ are subspaces of U and V respectively. Moreover, α induces an isomorphism

$$\begin{aligned}
 \bar{\alpha} : U / \ker \alpha &\rightarrow \operatorname{im} \alpha \\
 (\mathbf{u} + \ker \alpha) &\mapsto \alpha(\mathbf{u})
 \end{aligned}$$

Note that if we view a vector space as an abelian group, then this is exactly the first isomorphism theorem of groups.

Proof. We know that $\mathbf{0} \in \ker \alpha$ and $\mathbf{0} \in \operatorname{im} \alpha$.

Suppose $\mathbf{u}_1, \mathbf{u}_2 \in \ker \alpha$ and $\lambda_1, \lambda_2 \in \mathbb{F}$. Then

$$\alpha(\lambda_1\mathbf{u}_1 + \lambda_2\mathbf{u}_2) = \lambda_1\alpha(\mathbf{u}_1) + \lambda_2\alpha(\mathbf{u}_2) = \mathbf{0}.$$

So $\lambda_1\mathbf{u}_1 + \lambda_2\mathbf{u}_2 \in \ker \alpha$. So $\ker \alpha$ is a subspace.

Similarly, if $\alpha(\mathbf{u}_1), \alpha(\mathbf{u}_2) \in \operatorname{im} \alpha$, then $\lambda\alpha(\mathbf{u}_1) + \lambda_2\alpha(\mathbf{u}_2) = \alpha(\lambda_1\mathbf{u}_1 + \lambda_2\mathbf{u}_2) \in \operatorname{im} \alpha$. So $\operatorname{im} \alpha$ is a subspace.

Now by the first isomorphism theorem of groups, $\bar{\alpha}$ is a well-defined isomorphism of groups. So it remains to show that $\bar{\alpha}$ is a linear map. Indeed, we have

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

So $\bar{\alpha}$ is a linear map. □

Definition (Rank and nullity). If $\alpha : U \rightarrow V$ is a linear map between finite-dimensional vector spaces over \mathbb{F} (in fact we just need U to be finite-dimensional), the *rank* of α is the number $r(\alpha) = \dim \operatorname{im} \alpha$. The *nullity* of α is the number $n(\alpha) = \dim \ker \alpha$.

Corollary (Rank-nullity theorem). If $\alpha : U \rightarrow V$ is a linear map and U is finite-dimensional, then

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

Proof. By the first isomorphism theorem, we know that $U/\ker \alpha \cong \operatorname{im} \alpha$. So we have

$$\dim \operatorname{im} \alpha = \dim(U/\ker \alpha) = \dim U - \dim \ker \alpha.$$

So the result follows. \square

We can also prove this result without the first isomorphism theorem, and say a bit more in the meantime.

Proposition. If $\alpha : U \rightarrow V$ is a linear map between finite-dimensional vector spaces over \mathbb{F} , then there are bases $(\mathbf{e}_1, \dots, \mathbf{e}_m)$ for U and $(\mathbf{f}_1, \dots, \mathbf{f}_n)$ for V such that α is represented by the matrix

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

where $r = r(\alpha)$ and I_r is the $r \times r$ identity matrix.

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

Proof. Let $\mathbf{e}_{k+1}, \dots, \mathbf{e}_m$ be a basis for the kernel of α . Then we can extend this to a basis of the $(\mathbf{e}_1, \dots, \mathbf{e}_m)$.

Let $\mathbf{f}_i = \alpha(\mathbf{e}_i)$ for $1 \leq i \leq k$. We now show that $(\mathbf{f}_1, \dots, \mathbf{f}_k)$ is a basis for $\operatorname{im} \alpha$ (and thus $k = r$). We first show that it spans. Suppose $\mathbf{v} \in \operatorname{im} \alpha$. Then we have

$$\mathbf{v} = \alpha \left(\sum_{i=1}^m \lambda_i \mathbf{e}_i \right)$$

for some $\lambda_i \in \mathbb{F}$. By linearity, we can write this as

$$\mathbf{v} = \sum_{i=1}^m \lambda_i \alpha(\mathbf{e}_i) = \sum_{i=1}^k \lambda_i \mathbf{f}_i + \mathbf{0}.$$

So $\mathbf{v} \in \langle \mathbf{f}_1, \dots, \mathbf{f}_k \rangle$.

To show linear dependence, suppose that

$$\sum_{i=1}^k \mu_i \mathbf{f}_i = \mathbf{0}.$$

So we have

$$\alpha \left(\sum_{i=1}^k \mu_i \mathbf{e}_i \right) = \mathbf{0}.$$

So $\sum_{i=1}^k \mu_i \mathbf{e}_i \in \ker \alpha$. Since $(\mathbf{e}_{k+1}, \dots, \mathbf{e}_m)$ is a basis for $\ker \alpha$, we can write

$$\sum_{i=1}^k \mu_i \mathbf{e}_i = \sum_{i=k+1}^m \mu_i \mathbf{e}_i$$

for some μ_i ($i = k+1, \dots, m$). Since $(\mathbf{e}_1, \dots, \mathbf{e}_m)$ is a basis, we must have $\mu_i = 0$ for all i . So they are linearly independent.

Now we extend $(\mathbf{f}_1, \dots, \mathbf{f}_r)$ to a basis for V , and

$$\alpha(\mathbf{e}_i) = \begin{cases} \mathbf{f}_i & 1 \leq i \leq k \\ 0 & k+1 \leq i \leq m \end{cases}. \quad \square$$

Example. Let

$$W = \{x \in \mathbb{R}^5 : x_1 + x_2 + x_3 = 0 = x_3 - x_4 - x_5\}.$$

What is $\dim W$? Well, it clearly is 3, but how can we prove it?

We can consider the map $\alpha : \mathbb{R}^5 \rightarrow \mathbb{R}^2$ given by

$$\begin{pmatrix} x_1 \\ \vdots \\ x_5 \end{pmatrix} \mapsto \begin{pmatrix} x_1 + x_2 + x_5 \\ x_3 - x_4 - x_5 \end{pmatrix}$$

Then $\ker \alpha = W$. So $\dim W = 5 - r(\alpha)$. We know that $\alpha(1, 0, 0, 0, 0) = (1, 0)$ and $\alpha(0, 0, 1, 0, 0) = (0, 1)$. So $r(\alpha) = \dim \operatorname{im} \alpha = 2$. So $\dim W = 3$.

More generally, the rank-nullity theorem gives that m linear equations in n have a space of solutions of dimension at least $n - m$.

Example. Suppose that U and W are subspaces of V , all of which are finite-dimensional vector spaces of \mathbb{F} . We let

$$\begin{aligned} \alpha : U \oplus W &\rightarrow V \\ (\mathbf{u}, \mathbf{w}) &\mapsto \mathbf{u} + \mathbf{w}, \end{aligned}$$

where the \oplus is the *external* direct sum. Then $\operatorname{im} \alpha = U + W$ and

$$\ker \alpha = \{(\mathbf{u}, -\mathbf{u}) : \mathbf{u} \in U \cap W\} \cong \dim(U \cap W).$$

Then we have

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

This is a result we've previously obtained through fiddling with basis and horrible stuff.

Corollary. Suppose $\alpha : U \rightarrow V$ is a linear map between vector spaces over \mathbb{F} both of dimension $n < \infty$. Then the following are equivalent

- (i) α is injective;
- (ii) α is surjective;

(iii) α is an isomorphism.

Proof. It is clear that, (iii) implies (i) and (ii), and (i) and (ii) together implies (iii). So it suffices to show that (i) and (ii) are equivalent.

Note that α is injective iff $n(\alpha) = 0$, and α is surjective iff $r(\alpha) = \dim V = n$. By the rank-nullity theorem, $n(\alpha) + r(\alpha) = n$. So the result follows immediately. \square

Lemma. Let $A \in M_{n,n}(\mathbb{F}) = M_n(\mathbb{F})$ be a square matrix. The following are equivalent

(i) There exists $B \in M_n(\mathbb{F})$ such that $BA = I_n$.

(ii) There exists $C \in M_n(\mathbb{F})$ such that $AC = I_n$.

If these hold, then $B = C$. We call A *invertible* or *non-singular*, and write $A^{-1} = B = C$.

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

We note that (i) is equivalent to saying that there exists β such that $\beta\alpha = \iota$. This is true iff α is injective, which is true iff α is an isomorphism, which is true iff α has an inverse α^{-1} .

Similarly, (ii) is equivalent to saying that there exists γ such that $\alpha\gamma = \iota$. This is true iff α is injective, which is true iff α is isomorphism, which is true iff α has an inverse α^{-1} .

So these are the same things, and we have $\beta = \alpha^{-1} = \gamma$. \square

2.4 Change of basis

Suppose we have a linear map $\alpha : U \rightarrow V$. Given a basis $\{\mathbf{e}_i\}$ for U , and a basis $\{\mathbf{f}_i\}$ for V , we can obtain a matrix A .

$$\begin{array}{ccc} U & \xrightarrow{\alpha} & V \\ \uparrow (\mathbf{e}_i) & & \uparrow (\mathbf{f}_i) \\ \mathbb{F}^m & \xrightarrow{A} & \mathbb{F}^n \end{array}$$

We now want to consider what happens when we have two different basis $\{\mathbf{u}_i\}$ and $\{\mathbf{e}_i\}$ of U . These will then give rise to two different maps from \mathbb{F}^m to our space U , and the two basis can be related by a change-of-basis map P . We can put them in the following diagram:

$$\begin{array}{ccc} U & \xrightarrow{\iota_U} & U \\ \uparrow (\mathbf{u}_i) & & \uparrow (\mathbf{e}_i) \\ \mathbb{F}^m & \xrightarrow{P} & \mathbb{F}^m \end{array}$$

where ι_U is the identity map. If we perform a change of basis for both U and V , we can stitch the diagrams together as

$$\begin{array}{ccccccc}
 U & \xrightarrow{\iota_U} & U & \xrightarrow{\alpha} & V & \xleftarrow{\iota_V} & V \\
 (\mathbf{u}_i) \uparrow & & (\mathbf{e}_i) \uparrow & & (\mathbf{f}_i) \uparrow & & (\mathbf{v}_i) \uparrow \\
 \mathbb{F}^m & \xrightarrow{P} & \mathbb{F}^m & \xrightarrow{A} & \mathbb{F}^n & \xleftarrow{Q} & \mathbb{F}^n \\
 & & & \searrow & & \nearrow & \\
 & & & & B & &
 \end{array}$$

Then if we want a matrix representing the map $U \rightarrow V$ with respect to bases (\mathbf{u}_i) and (\mathbf{v}_i) , we can write it as the composition

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

We can write this as a theorem:

Theorem. Suppose that $(\mathbf{e}_1, \dots, \mathbf{e}_m)$ and $(\mathbf{u}_1, \dots, \mathbf{u}_m)$ are basis for a finite-dimensional vector space U over \mathbb{F} , and $(\mathbf{f}_1, \dots, \mathbf{f}_n)$ and $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ are basis of a finite-dimensional vector space V over \mathbb{F} .

Let $\alpha : U \rightarrow V$ be a linear map represented by a matrix A with respect to (\mathbf{e}_i) and (\mathbf{f}_i) and by B with respect to (\mathbf{u}_i) and (\mathbf{v}_i) . Then

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

where P and Q are given by

$$\mathbf{u}_i = \sum_{k=1}^m P_{ki} \mathbf{e}_k, \quad \mathbf{v}_i = \sum_{k=1}^n Q_{ki} \mathbf{f}_k.$$

Note that one can view P as the matrix representing the identity map i_U from U with basis (\mathbf{u}_i) to U with basis (\mathbf{e}_i) , and similarly for Q . So both are invertible.

Proof. On the one hand, we have

$$\alpha(\mathbf{u}_i) = \sum_{j=1}^n B_{ji} \mathbf{v}_j = \sum_j \sum_{\ell} B_{ji} Q_{\ell j} \mathbf{f}_{\ell} = \sum_{\ell} [QB]_{\ell i} \mathbf{f}_{\ell}.$$

On the other hand, we can write

$$\alpha(\mathbf{u}_i) = \alpha \left(\sum_{k=1}^m P_{ki} \mathbf{e}_k \right) = \sum_{k=1}^m P_{ki} \sum_{\ell} A_{\ell k} \mathbf{f}_{\ell} = \sum_{\ell} [AP]_{\ell i} \mathbf{f}_{\ell}.$$

Since the \mathbf{f}_{ℓ} are linearly independent, we conclude that

$$QB = AP.$$

Since Q is invertible, we get $B = Q^{-1}AP$. \square

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

Since $\text{GL}_K(\mathbb{F}) = \{A \in \text{Mat}_k(\mathbb{F}) : A \text{ is invertible}\}$ is a group, for each $k \geq 1$, this is indeed an equivalence relation. The equivalence classes are orbits under the action of $\text{GL}_m(\mathbb{F}) \times \text{GL}_n(\mathbb{F})$, given by

$$\begin{aligned} \text{GL}_m(\mathbb{F}) \times \text{GL}_n(\mathbb{F}) \times \text{Mat}_{n,m}(\mathbb{F}) &\rightarrow \text{Mat}(\mathbb{F}) \\ (P, Q, A) &\mapsto QAP^{-1}. \end{aligned}$$

Two matrices are equivalent if and only if they represent the same linear map with respect to different basis.

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

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

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

This is just a rephrasing of the proposition we had last time. But this tells us there are $\min(m, n) + 1$ orbits of the action above parametrized by r .

Definition (Column and row rank). If $A \in \text{Mat}_{n,m}(\mathbb{F})$, then

- The *column rank* of A , written $r(A)$, is the dimension of the subspace of \mathbb{F}^n spanned by the columns of A .
- The *row rank* of A , written $r(A)$, is the dimension of the subspace of \mathbb{F}^m spanned by the rows of A . Alternatively, it is the column rank of A^T .

There is no a priori reason why these should be equal to each other. However, it turns out they are always equal.

Note that if $\alpha : \mathbb{F}^m \rightarrow \mathbb{F}^n$ is the linear map represented by A (with respect to the standard basis), then $r(A) = r(\alpha)$, i.e. the column rank is the rank. Moreover, since the rank of a map is independent of the basis, equivalent matrices have the same column rank.

Theorem. If $A \in \text{Mat}_{n,m}(\mathbb{F})$, then $r(A) = r(A^T)$, i.e. the row rank is equivalent to the column rank.

Proof. We know that there are some invertible P, Q such that

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

where $r = r(A)$. We can transpose this whole equation to obtain

$$(Q^{-1}AP)^T = P^T A^T (Q^T)^{-1} = \begin{pmatrix} I_r & 0 \\ 0 & 0 \end{pmatrix}$$

So $r(A^T) = r$. □

(ii) $AE_{I_j}^n(\lambda)$ 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 $A \in \text{Mat}_{n,m}(\mathbb{F})$, then there exists invertible matrices $P \in \text{GL}_m(\mathbb{F})$, $Q \in \text{GL}_n(\mathbb{F})$ so that

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

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

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

If $A = 0$, then done. If not, there is some i, j such that $A_{ij} \neq 0$. By swapping row 1 and row i ; and then column 1 and column j , we can assume $A_{11} \neq 0$. 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 column j for each $j \neq 1$, and then add $-A_{i1}$ times row 1 to row $i \neq 1$. Then we now have

$$A = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & & & \\ \vdots & & B & \\ 0 & & & \end{pmatrix}$$

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. \square

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.

3 Duality

Duality is a principle we will find throughout mathematics. For example, in IB Optimisation, we considered the dual problems of linear programs. Here we will look for the dual of vector spaces. In general, we try to look at our question in a “mirror” and hope that the mirror problem is easier to solve than the original mirror.

At first, the definition of the dual might seem a bit arbitrary and weird. We will try to motivate it using what we will call *annihilators*, but they are much more useful than just for these. Despite their usefulness, though, they can be confusing to work with at times, since the dual space of a vector space V will be constructed by considering linear maps on V , and when we work with maps on dual spaces, things explode.

3.1 Dual space

To specify a subspace of \mathbb{F}^n , we can write down linear equations that its elements satisfy. For example, if we have the subspace $U = \left\langle \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix} \right\rangle \subseteq \mathbb{F}^3$, we can specify

this by saying $\begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \in U$ if and only if

$$\begin{aligned} x_1 - x_3 &= 0 \\ 2x_1 - x_2 &= 0. \end{aligned}$$

However, characterizing a space in terms of equations involves picking some particular equations out of the many possibilities. In general, we do not like making arbitrary choices. Hence the solution is to consider *all* possible such equations. We will show that these form a subspace in some space.

We can interpret these equations in terms of linear maps $\mathbb{F}^n \rightarrow \mathbb{F}$. For example $x_1 - x_3 = 0$ if and only if $\begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \in \ker \theta$, where $\theta : \mathbb{F}^3 \rightarrow \mathbb{F}$ is defined

$$\text{by } \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \mapsto x_1 - x_3.$$

This works well with the vector space operations. If $\theta_1, \theta_2 : \mathbb{F}^n \rightarrow \mathbb{F}$ vanish on some subspace of \mathbb{F}^n , and $\lambda, \mu \in \mathbb{F}$, then $\lambda\theta_1 + \mu\theta_2$ also vanishes on the subspace. So the set of all maps $\mathbb{F}^n \rightarrow \mathbb{F}$ that vanishes on U forms a vector space.

To formalize this notion, we introduce dual spaces.

Definition (Dual space). Let V be a vector space over \mathbb{F} . The *dual* of V is defined as

$$V^* = \mathcal{L}(V, \mathbb{F}) = \{\theta : V \rightarrow \mathbb{F} : \theta \text{ linear}\}.$$

Elements of V^* are called *linear functionals* or *linear forms*.

By convention, we use Roman letters for elements in V , and Greek letters for elements in V^* .

Example.

- If $V = \mathbb{R}^3$ and $\theta : V \rightarrow \mathbb{R}$ that sends $\begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \mapsto x_1 - x_3$, then $\theta \in V^*$.
- Let $V = \mathbb{F}^X$. Then for any fixed x , $\theta : V \rightarrow \mathbb{F}$ defined by $f \mapsto f(x)$ is in V^* .
- Let $V = C([0, 1], \mathbb{R})$. Then $f \mapsto \int_0^1 f(t) dt \in V^*$.
- The trace $\text{tr} : M_n(\mathbb{F}) \rightarrow \mathbb{F}$ defined by $A \mapsto \sum_{i=1}^n A_{ii}$ is in $M_n(\mathbb{F})^*$.

It turns out it is rather easy to specify how the dual space looks like, at least in the case where V is finite dimensional.

Lemma. If V is a finite-dimensional vector space over \mathbb{F} with basis $(\mathbf{e}_1, \dots, \mathbf{e}_n)$, then there is a basis $(\varepsilon_1, \dots, \varepsilon_n)$ for V^* (called the *dual basis* to $(\mathbf{e}_1, \dots, \mathbf{e}_n)$) such that

$$\varepsilon_i(\mathbf{e}_j) = \delta_{ij}.$$

Proof. Since linear maps are characterized by their values on a basis, there exists unique choices for $\varepsilon_1, \dots, \varepsilon_n \in V^*$. Now we show that $(\varepsilon_1, \dots, \varepsilon_n)$ is a basis.

Suppose $\theta \in V^*$. We show that we can write it uniquely as a combination of $\varepsilon_1, \dots, \varepsilon_n$. We have $\theta = \sum_{i=1}^n \lambda_i \varepsilon_i$ if and only if $\theta(\mathbf{e}_j) = \sum_{i=1}^n \lambda_i \varepsilon_i(\mathbf{e}_j)$ (for all j) if and only if $\lambda_j = \theta(\mathbf{e}_j)$. So we have uniqueness and existence. \square

Corollary. If V is finite dimensional, then $\dim V = \dim V^*$.

When V is not finite dimensional, this need not be true. However, we know that the dimension of V^* is at least as big as that of V , since the above gives a set of $\dim V$ many independent vectors in V^* . In fact for any infinite dimensional vector space, $\dim V^*$ is strictly larger than $\dim V$, if we manage to define dimensions for infinite-dimensional vector spaces.

It helps to come up with a more concrete example of how dual spaces look like. Consider the vector space \mathbb{F}^n , where we treat each element as a column vector (with respect to the standard basis). Then we can regard elements of V^* as just row vectors $(a_1, \dots, a_n) = \sum_{j=1}^n a_j \varepsilon_j$ with respect to the dual basis. We have

$$\left(\sum a_j \varepsilon_j \right) \left(\sum_{x_i} \mathbf{e}_i \right) = \sum_{i,j} a_j x_i \delta_{ij} = \sum_{i=1}^n a_i x_i = (a_1 \quad \dots \quad a_n) \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}.$$

This is exactly what we want.

Now what happens when we change basis? How will the dual basis change?

Proposition. Let V be a finite-dimensional vector space over \mathbb{F} with bases $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ and $(\mathbf{f}_1, \dots, \mathbf{f}_n)$, and that P is the change of basis matrix so that

$$\mathbf{f}_i = \sum_{k=1}^n P_{ki} \mathbf{e}_k.$$

Let $(\varepsilon_1, \dots, \varepsilon_n)$ and (η_1, \dots, η_n) be the corresponding dual bases so that

$$\varepsilon_i(\mathbf{e}_j) = \delta_{ij} = \eta_i(\mathbf{f}_j).$$

Then the change of basis matrix from $(\varepsilon_1, \dots, \varepsilon_n)$ to (η_1, \dots, η_n) is $(P^{-1})^T$, i.e.

$$\varepsilon_i = \sum_{\ell=1}^n P_{\ell i}^T \eta_\ell.$$

Proof. For convenience, write $Q = P^{-1}$ so that

$$\mathbf{e}_j = \sum_{k=1}^n Q_{kj} \mathbf{f}_k.$$

So we can compute

$$\begin{aligned} \left(\sum_{\ell=1}^{\infty} P_{i\ell} \eta_\ell \right) (\mathbf{e}_j) &= \left(\sum_{\ell=1}^{\infty} P_{i\ell} \eta_\ell \right) \left(\sum_{k=1}^n Q_{kj} \mathbf{f}_k \right) \\ &= \sum_{k,\ell} P_{i\ell} \delta_{\ell k} Q_{kj} \\ &= \sum_{k,\ell} P_{i\ell} Q_{\ell j} \\ &= [PQ]_{ij} \\ &= \delta_{ij}. \end{aligned}$$

$$\text{So } \varepsilon_i = \sum_{\ell=1}^n P_{\ell i}^T \eta_\ell. \quad \square$$

Now we'll return to our original motivation, and think how we can define subspaces of V^* in terms of subspaces of V , and vice versa.

Definition (Annihilator). Let $U \subseteq V$. Then the *annihilator* of U is

$$U^0 = \{\theta \in V^* : \theta(\mathbf{u}) = 0, \forall \mathbf{u} \in U\}.$$

If $W \subseteq V^*$, then the *annihilator* of W is

$$W^0 = \{\mathbf{v} \in V : \theta(\mathbf{v}) = 0, \forall \theta \in W\}.$$

One might object that W^0 should be a subset of V^{**} and not V . We will later show that there is a canonical isomorphism between V^{**} and V , and this will all make sense.

Example. Consider \mathbb{R}^3 with standard basis $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$; $(\mathbb{R}^3)^*$ with dual basis $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$. If $U = \langle \mathbf{e}_1 + 2\mathbf{e}_2 + \mathbf{e}_3 \rangle$ and $W = \langle \varepsilon_1 - \varepsilon_3, 2\varepsilon_1 - \varepsilon_2 \rangle$, then $U^0 = W$ and $W^0 = U$.

We see that the dimension of U and U^0 add up to three, which is the dimension of \mathbb{R}^3 . This is typical.

Proposition. Let V be a vector space over \mathbb{F} and U a subspace. Then

$$\dim U + \dim U^0 = \dim V.$$

We are going to prove this in many ways.

Proof. Let $(\mathbf{e}_1, \dots, \mathbf{e}_k)$ be a basis for U and extend to $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ a basis for V . Consider the dual basis for V^* , say $(\varepsilon_1, \dots, \varepsilon_n)$. Then we will show that

$$U^0 = \langle \varepsilon_{k+1}, \dots, \varepsilon_n \rangle.$$

So $\dim U^0 = n - k$ as required. This is easy to prove — if $j > k$, then $\varepsilon_j(\mathbf{e}_i) = 0$ for all $i \leq k$. So $\varepsilon_{k+1}, \dots, \varepsilon_n \in U^0$. On the other hand, suppose $\theta \in U^0$. Then we can write

$$\theta = \sum_{j=1}^n \lambda_j \varepsilon_j.$$

But then $0 = \theta(\mathbf{e}_i) = \lambda_i$ for $i \leq k$. So done. \square

Proof. Consider the restriction map $V^* \rightarrow U^*$, given by $\theta \mapsto \theta|_U$. This is obviously linear. Since every linear map $U \rightarrow \mathbb{F}$ can be extended to $V \rightarrow \mathbb{F}$, this is a surjection. Moreover, the kernel is U^0 . So by rank-nullity theorem,

$$\dim V^* = \dim U^0 + \dim U^*.$$

Since $\dim V^* = \dim V$ and $\dim U^* = \dim U$, we're done. \square

Proof. We can show that $U^0 \simeq (V/U)^*$, and then deduce the result. Details are left as an exercise. \square

3.2 Dual maps

Since linear algebra is the study of vector spaces and linear maps between them, after dualizing vector spaces, we should be able to dualize linear maps as well. If we have a map $\alpha : V \rightarrow W$, then after dualizing, the map will go *the other direction*, i.e. $\alpha^* : W^* \rightarrow V^*$. This is a characteristic common to most dualization processes in mathematics.

Definition (Dual map). Let V, W be vector spaces over \mathbb{F} and $\alpha : V \rightarrow W \in \mathcal{L}(V, W)$. The *dual map* to α , written $\alpha^* : W^* \rightarrow V^*$ is given by $\theta \mapsto \theta \circ \alpha$. Since the composite of linear maps is linear, $\alpha^*(\theta) \in V^*$. So this is a genuine map.

Proposition. Let $\alpha \in \mathcal{L}(V, W)$ be a linear map. Then $\alpha^* \in \mathcal{L}(W^*, V^*)$ is a linear map.

This is *not* the same as what we remarked at the end of the definition of the dual map. What we remarked was that given any θ , $\alpha^*(\theta)$ is a linear map. What we want to show here is that α^* itself as a map $W^* \rightarrow V^*$ is linear.

Proof. Let $\lambda, \mu \in \mathbb{F}$ and $\theta_1, \theta_2 \in W^*$. We want to show

$$\alpha^*(\lambda\theta_1 + \mu\theta_2) = \lambda\alpha^*(\theta_1) + \mu\alpha^*(\theta_2).$$

To show this, we show that for every $\mathbf{v} \in V$, the left and right give the same result. We have

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

So $\alpha^* \in \mathcal{L}(W^*, V^*)$. \square

What happens to the matrices when we take the dual map? The answer is that we get the transpose.

Proposition. Let V, W be finite-dimensional vector spaces over \mathbb{F} and $\alpha : V \rightarrow W$ be a linear map. Let $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ be a basis for V and $(\mathbf{f}_1, \dots, \mathbf{f}_m)$ be a basis for W ; $(\varepsilon_1, \dots, \varepsilon_n)$ and (η_1, \dots, η_m) the corresponding dual bases.

Suppose α is represented by A with respect to (\mathbf{e}_i) and (\mathbf{f}_i) for V and W . Then α^* is represented by A^T with respect to the corresponding dual bases.

Proof. We are given that

$$\alpha(\mathbf{e}_i) = \sum_{k=1}^m A_{ki} \mathbf{f}_k.$$

We must compute $\alpha^*(\eta_i)$. To do so, we evaluate it at \mathbf{e}_j . We have

$$\alpha^*(\eta_i)(\mathbf{e}_j) = \eta_i(\alpha(\mathbf{e}_j)) = \eta_i\left(\sum_{k=1}^m A_{kj} \mathbf{f}_k\right) = \sum_{k=1}^m A_{kj} \delta_{ik} = A_{ij}.$$

We can also write this as

$$\alpha^*(\eta_i)(\mathbf{e}_j) = \sum_{k=1}^n A_{ik} \varepsilon_k(\mathbf{e}_j).$$

Since this is true for all j , we have

$$\alpha^*(\eta_i) \sum_{k=1}^n A_{ik} \varepsilon_k = \sum_{k=1}^n A_{ki}^T \varepsilon_k.$$

So done. □

Note that if $\alpha : U \rightarrow V$ and $\beta : V \rightarrow W$, $\theta \in W^*$, then

$$(\beta\alpha)^*(\theta) = \theta\beta\alpha = \alpha^*(\theta\beta) = \alpha^*(\beta^*(\theta)).$$

So we have $(\beta\alpha)^* = \alpha^*\beta^*$. This is obviously true for the finite-dimensional case, since that's how transposes of matrices work.

Similarly, if $\alpha, \beta : U \rightarrow V$, then $(\lambda\alpha + \mu\beta)^* = \lambda\alpha^* + \mu\beta^*$.

What happens when we change basis? If $B = Q^{-1}AP$ for some invertible P and Q , then

$$B^T = (Q^{-1}AP)^T = P^T A^T (Q^{-1})^T = ((P^{-1})^T)^{-1} A^T (Q^{-1})^T.$$

So in the dual space, we conjugate by the dual of the change-of-basis matrices.

As we said, we can use dualization to translate problems about a vector space to its dual. The following lemma gives us some good tools to do so:

Lemma. Let $\alpha \in \mathcal{L}(V, W)$ with V, W finite dimensional vector spaces over \mathbb{F} . Then

(i) $\ker \alpha^* = (\operatorname{im} \alpha)^0$.

(ii) $r(\alpha) = r(\alpha^*)$ (which is another proof that row rank is equal to column rank).

(iii) $\text{im } \alpha^* = (\ker \alpha)^0$.

At first sight, (i) and (iii) look quite similar. However, (i) is almost trivial to prove, but (iii) is rather hard.

Proof.

(i) If $\theta \in W^*$, then

$$\begin{aligned} \theta \in \ker \alpha^* &\Leftrightarrow \alpha^*(\theta) = 0 \\ &\Leftrightarrow (\forall \mathbf{v} \in V) \theta\alpha(\mathbf{v}) = 0 \\ &\Leftrightarrow (\forall \mathbf{w} \in \text{im } \alpha) \theta(\mathbf{w}) = 0 \\ &\Leftrightarrow \theta \in (\text{im } \alpha)^0. \end{aligned}$$

(ii) As $\text{im } \alpha \leq W$, we've seen that

$$\dim \text{im } \alpha + \dim(\text{im } \alpha)^0 = \dim W.$$

Using (i), we see

$$n(\alpha^*) = \dim(\text{im } \alpha)^0.$$

So

$$r(\alpha) + n(\alpha^*) = \dim W = \dim W^*.$$

By the rank-nullity theorem, we have $r(\alpha) = r(\alpha^*)$.

(iii) The proof in (i) doesn't quite work here. We can only show that one includes the other. To draw the conclusion, we will show that the two spaces have the dimensions, and hence must be equal.

Let $\theta \in \text{im } \alpha^*$. Then $\theta = \phi\alpha$ for some $\phi \in W^*$. If $\mathbf{v} \in \ker \alpha$, then

$$\theta(\mathbf{v}) = \phi(\alpha(\mathbf{v})) = \phi(\mathbf{0}) = \mathbf{0}.$$

So $\text{im } \alpha^* \subseteq (\ker \alpha)^0$.

But we know

$$\dim(\ker \alpha)^0 + \dim \ker \alpha = \dim V,$$

So we have

$$\dim(\ker \alpha)^0 = \dim V - n(\alpha) = r(\alpha) = r(\alpha^*) = \dim \text{im } \alpha^*.$$

Hence we must have $\text{im } \alpha^* = (\ker \alpha)^0$. □

Not only do we want to get from V to V^* , we want to get back from V^* to V . We can take the dual of V^* to get a V^{**} . We already know that V^{**} is isomorphic to V , since V^* is isomorphic to V already. However, the isomorphism between V^* and V are not "natural". To define such an isomorphism, we needed to pick a basis for V and consider a dual basis. If we picked a different basis, we would get a different isomorphism. There is no natural, canonical, uniquely-defined isomorphism between V and V^* .

However, this is not the case when we want to construct an isomorphism $V \rightarrow V^{**}$. The construction of this isomorphism is obvious once we think hard what V^{**} actually means. Unwrapping the definition, we know $V^{**} = \mathcal{L}(V^*, \mathbb{F})$.

Our isomorphism has to produce something in V^{**} given any $\mathbf{v} \in V$. This is equivalent to saying given any $\mathbf{v} \in V$ and a function $\theta \in V^*$, produce something in \mathbb{F} .

This is easy, by definition $\theta \in V^*$ is just a linear map $V \rightarrow \mathbb{F}$. So given \mathbf{v} and θ , we just return $\theta(\mathbf{v})$. We now just have to show that this is linear and is bijective.

Lemma. Let V be a vector space over \mathbb{F} . Then there is a linear map $\text{ev} : V \rightarrow (V^*)^*$ given by

$$\text{ev}(\mathbf{v})(\theta) = \theta(\mathbf{v}).$$

We call this the *evaluation* map.

We call this a “canonical” map since this does not require picking a particular basis of the vector spaces. It is in some sense a “natural” map.

Proof. We first show that $\text{ev}(\mathbf{v}) \in V^{**}$ for all $\mathbf{v} \in V$, i.e. $\text{ev}(\mathbf{v})$ is linear for any \mathbf{v} . For any $\lambda, \mu \in \mathbb{F}$, $\theta_1, \theta_2 \in V^*$, then for $\mathbf{v} \in V$, we have

$$\begin{aligned} \text{ev}(\mathbf{v})(\lambda\theta_1 + \mu\theta_2) &= (\lambda\theta_1 + \mu\theta_2)(\mathbf{v}) \\ &= \lambda\theta_1(\mathbf{v}) + \mu\theta_2(\mathbf{v}) \\ &= \lambda\text{ev}(\mathbf{v})(\theta_1) + \mu\text{ev}(\mathbf{v})(\theta_2). \end{aligned}$$

So done. Now we show that ev itself is linear. Let $\lambda, \mu \in \mathbb{F}$, $\mathbf{v}_1, \mathbf{v}_2 \in V$. We want to show

$$\text{ev}(\lambda\mathbf{v}_1 + \mu\mathbf{v}_2) = \lambda\text{ev}(\mathbf{v}_1) + \mu\text{ev}(\mathbf{v}_2).$$

To show these are equal, pick $\theta \in V^*$. Then

$$\begin{aligned} \text{ev}(\lambda\mathbf{v}_1 + \mu\mathbf{v}_2)(\theta) &= \theta(\lambda\mathbf{v}_1 + \mu\mathbf{v}_2) \\ &= \lambda\theta(\mathbf{v}_1) + \mu\theta(\mathbf{v}_2) \\ &= \lambda\text{ev}(\mathbf{v}_1)(\theta) + \mu\text{ev}(\mathbf{v}_2)(\theta) \\ &= (\lambda\text{ev}(\mathbf{v}_1) + \mu\text{ev}(\mathbf{v}_2))(\theta). \end{aligned}$$

So done. □

In the special case where V is finite-dimensional, this is an isomorphism.

Lemma. If V is finite-dimensional, then $\text{ev} : V \rightarrow V^{**}$ is an isomorphism.

This is very false for infinite dimensional spaces. In fact, this is true *only* for finite-dimensional vector spaces (assuming the axiom of choice), and some (weird) people use this as the definition of finite-dimensional vector spaces.

Proof. We first show it is injective. Suppose $\text{ev}(\mathbf{v}) = \mathbf{0}$ for some $\mathbf{v} \in V$. Then $\theta(\mathbf{v}) = \text{ev}(\mathbf{v})(\theta) = 0$ for all $\theta \in V^*$. So $\dim\langle \mathbf{v} \rangle^0 = \dim V^* = \dim V$. So $\dim\langle \mathbf{v} \rangle = 0$. So $\mathbf{v} = \mathbf{0}$. So ev is injective. Since V and V^{**} have the same dimension, this is also surjective. So done. □

From now on, we will just pretend that V and V^{**} are the same thing, at least when V is finite dimensional.

Note that this lemma does not just say that V is isomorphic to V^{**} (we already know that since they have the same dimension). This says there is a completely canonical way to choose the isomorphism.

In general, if V is infinite dimensional, then ev is injective, but not surjective. So we can think of V as a subspace of V^{**} in a canonical way.

Lemma. Let V, W be finite-dimensional vector spaces over \mathbb{F} after identifying $(V$ and $V^{**})$ and $(W$ and $W^{**})$ by the evaluation map. Then we have

- (i) If $U \leq V$, then $U^{00} = U$.
- (ii) If $\alpha \in \mathcal{L}(V, W)$, then $\alpha^{**} = \alpha$.

Proof.

- (i) Let $\mathbf{u} \in U$. Then $\mathbf{u}(\theta) = \theta(\mathbf{u}) = 0$ for all $\theta \in U^0$. So \mathbf{u} annihilates everything in U^0 . So $\mathbf{u} \in U^{00}$. So $U \subseteq U^{00}$. We also know that

$$\dim U = \dim V - \dim U^0 = \dim V - (\dim V - \dim U^{00}) = \dim U^{00}.$$

So we must have $U = U^{00}$.

- (ii) The proof of this is basically — the transpose of the transpose is the original matrix. The only work we have to do is to show that the dual of the dual basis is the original basis.

Let $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ be a basis for V and $(\mathbf{f}_1, \dots, \mathbf{f}_m)$ be a basis for W , and let $(\varepsilon_1, \dots, \varepsilon_n)$ and (η_1, \dots, η_m) be the corresponding dual basis. We know that

$$\mathbf{e}_i(\varepsilon_j) = \delta_{ij} = \varepsilon_j(\mathbf{e}_i), \quad \mathbf{f}_i(\eta_j) = \delta_{ij} = \eta_j(\mathbf{f}_i).$$

So $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ is dual to $(\varepsilon_1, \dots, \varepsilon_n)$, and similarly for \mathbf{f} and η .

If α is represented by A , then α^* is represented by A^T . So α^{**} is represented by $(A^T)^T = A$. So done. \square

Proposition. Let V be a finite-dimensional vector space \mathbb{F} and U_1, U_2 are subspaces of V . Then we have

- (i) $(U_1 + U_2)^0 = U_1^0 \cap U_2^0$
- (ii) $(U_1 \cap U_2)^0 = U_1^0 + U_2^0$

Proof.

- (i) Suppose $\theta \in V^*$. Then

$$\begin{aligned} \theta \in (U_1 + U_2)^0 &\Leftrightarrow \theta(\mathbf{u}_1 + \mathbf{u}_2) = 0 \text{ for all } \mathbf{u}_i \in U_i \\ &\Leftrightarrow \theta(\mathbf{u}) = 0 \text{ for all } \mathbf{u} \in U_1 \cup U_2 \\ &\Leftrightarrow \theta \in U_1^0 \cap U_2^0. \end{aligned}$$

- (ii) We have

$$(U_1 \cap U_2)^0 = ((U_1^0)^0 \cap (U_2^0)^0)^0 = (U_1^0 + U_2^0)^{00} = U_1^0 + U_2^0.$$

So done. \square

4 Bilinear forms I

So far, we have been looking at linear things only. This can get quite boring. For a change, we look at *bilinear* maps instead. In this chapter, we will look at bilinear forms in general. It turns out there isn't much we can say about them, and hence this chapter is rather short. Later, in Chapter 7, we will study some special kinds of bilinear forms which are more interesting.

Definition (Bilinear form). Let V, W be vector spaces over \mathbb{F} . Then a function $\phi : V \times W \rightarrow \mathbb{F}$ is a *bilinear form* if it is linear in each variable, i.e. for each $\mathbf{v} \in V$, $\phi(\mathbf{v}, \cdot) : W \rightarrow \mathbb{F}$ is linear; for each $\mathbf{w} \in W$, $\phi(\cdot, \mathbf{w}) : V \rightarrow \mathbb{F}$ is linear.

Example. The map defined by

$$\begin{aligned} V \times V^* &\rightarrow \mathbb{F} \\ (\mathbf{v}, \theta) &\mapsto \theta(\mathbf{v}) = \text{ev}(\mathbf{v})(\theta) \end{aligned}$$

is a bilinear form.

Example. Let $V = W = \mathbb{F}^n$. Then the function $(\mathbf{v}, \mathbf{w}) = \sum_{i=1}^n v_i w_i$ is bilinear.

Example. If $V = W = C([0, 1], \mathbb{R})$, then

$$(f, g) \mapsto \int_0^a fg \, dt$$

is a bilinear form.

Example. Let $A \in \text{Mat}_{m,n}(\mathbb{F})$. Then

$$\begin{aligned} \phi : \mathbb{F}^m \times \mathbb{F}^n &\rightarrow \mathbb{F} \\ (\mathbf{v}, \mathbf{w}) &\mapsto \mathbf{v}^T A \mathbf{w} \end{aligned}$$

is bilinear. Note that the (real) dot product is the special case of this, where $n = m$ and $A = I$.

In fact, this is the most general form of bilinear forms on finite-dimensional vector spaces.

Definition (Matrix representing bilinear form). Let $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ be a basis for V and $(\mathbf{f}_1, \dots, \mathbf{f}_m)$ be a basis for W , and $\psi : V \times W \rightarrow \mathbb{F}$. Then the *matrix A representing ψ* with respect to the basis is defined to be

$$A_{ij} = \psi(\mathbf{e}_i, \mathbf{f}_j).$$

Note that if $\mathbf{v} = \sum \lambda_i \mathbf{e}_i$ and $\mathbf{w} = \sum \mu_j \mathbf{f}_j$, then by linearity, we get

$$\begin{aligned} \psi(\mathbf{v}, \mathbf{w}) &= \psi\left(\sum \lambda_i \mathbf{e}_i, \mathbf{w}\right) \\ &= \sum_i \lambda_i \psi(\mathbf{e}_i, \mathbf{w}) \\ &= \sum_i \lambda_i \psi\left(\mathbf{e}_i, \sum \mu_j \mathbf{f}_j\right) \\ &= \sum_{i,j} \lambda_i \mu_j \psi(\mathbf{e}_i, \mathbf{f}_j) \\ &= \lambda^T A \mu. \end{aligned}$$

So ψ is determined by A .

We have identified linear maps with matrices, and we have identified bilinear maps with matrices. However, you shouldn't think linear maps are bilinear maps. They are, obviously, two different things. In fact, the matrices representing matrices and bilinear forms transform differently when we change basis.

Proposition. Suppose $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ and $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ are basis for V such that

$$\mathbf{v}_i = \sum P_{ki} \mathbf{e}_k \text{ for all } i = 1, \dots, n;$$

and $(\mathbf{f}_1, \dots, \mathbf{f}_m)$ and $(\mathbf{w}_1, \dots, \mathbf{w}_m)$ are bases for W such that

$$\mathbf{w}_i = \sum Q_{\ell j} \mathbf{f}_\ell \text{ for all } j = 1, \dots, m.$$

Let $\psi : V \times W \rightarrow \mathbb{F}$ be a bilinear form represented by A with respect to $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ and $(\mathbf{f}_1, \dots, \mathbf{f}_m)$, and by B with respect to the bases $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ and $(\mathbf{w}_1, \dots, \mathbf{w}_m)$. Then

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

The difference with the transformation laws of matrices is this time we are taking *transposes*, not *inverses*.

Proof. We have

$$\begin{aligned} B_{ij} &= \phi(\mathbf{v}_i, \mathbf{w}_j) \\ &= \phi\left(\sum P_{ki} \mathbf{e}_k, \sum Q_{\ell j} \mathbf{f}_\ell\right) \\ &= \sum P_{ki} Q_{\ell j} \phi(\mathbf{e}_k, \mathbf{f}_\ell) \\ &= \sum_{k, \ell} P_{ki}^T A_{k\ell} Q_{\ell j} \\ &= (P^T A Q)_{ij}. \end{aligned} \quad \square$$

Note that while the transformation laws for bilinear forms and linear maps are different, we still get that two matrices are representing the same bilinear form with respect to different bases if and only if they are equivalent, since if $B = P^{-1} A Q$, then $B = ((P^{-1})^T)^T A Q$.

If we are given a bilinear form $\psi : V \times W \rightarrow \mathbb{F}$, we immediately get two linear maps:

$$\psi_L : V \rightarrow W^*, \quad \psi_R : W \rightarrow V^*,$$

defined by $\psi_L(\mathbf{v}) = \psi(\mathbf{v}, \cdot)$ and $\psi_R(\mathbf{w}) = \psi(\cdot, \mathbf{w})$.

For example, if $\psi : V \times V^* \rightarrow \mathbb{F}$, is defined by $(\mathbf{v}, \theta) \mapsto \theta(\mathbf{v})$, then $\psi_L : V \rightarrow V^{**}$ is the evaluation map. On the other hand, $\psi_R : V^* \rightarrow V^*$ is the identity map.

Lemma. Let $(\varepsilon_1, \dots, \varepsilon_n)$ be a basis for V^* dual to $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ of V ; (η_1, \dots, η_m) be a basis for W^* dual to $(\mathbf{f}_1, \dots, \mathbf{f}_m)$ of W .

If A represents ψ with respect to $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ and $(\mathbf{f}_1, \dots, \mathbf{f}_m)$, then A also represents ψ_R with respect to $(\mathbf{f}_1, \dots, \mathbf{f}_m)$ and $(\varepsilon_1, \dots, \varepsilon_n)$; and A^T represents ψ_L with respect to $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ and (η_1, \dots, η_m) .

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

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

So A^T represents ψ_L .

We also have

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

So

$$\psi_R(\mathbf{f}_j) = \sum A_{kj} \varepsilon_k. \quad \square$$

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, \mathbf{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

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

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

If we have a non-trivial left (or right) kernel, then in some sense, some elements in V (or W) are “useless”, and we don’t like these.

Definition (Non-degenerate bilinear form). ψ is *non-degenerate* if the left and right kernels are both trivial. We say ψ is *degenerate* otherwise.

Definition (Rank of bilinear form). If $\psi : V \rightarrow W$ 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 $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ and $(\mathbf{f}_1, \dots, \mathbf{f}_m)$ be their basis respectively.

Let $\psi : V \times W \rightarrow \mathbb{F}$ 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. \square

Example. The map

$$\begin{aligned} \mathbb{F}^2 \times \mathbb{F}^2 &\rightarrow \mathbb{F} \\ \begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix} &\mapsto ad - bc \end{aligned}$$

is a bilinear form. This, obviously, corresponds to the determinant of a 2-by-2 matrix. We have $\psi(\mathbf{v}, \mathbf{w}) = -\psi(\mathbf{w}, \mathbf{v})$ for all $\mathbf{v}, \mathbf{w} \in \mathbb{F}^2$.

5 Determinants of matrices

We probably all know what the determinant is. Here we are going to give a slightly more abstract definition, and spend quite a lot of time trying to motivate this definition.

Recall that S_n is the group of permutations of $\{1, \dots, n\}$, and there is a unique group homomorphism $\varepsilon : S_n \rightarrow \{\pm 1\}$ such that $\varepsilon(\sigma) = 1$ if σ can be written as a product of an even number of transpositions; $\varepsilon(\sigma) = -1$ if σ can be written as an odd number of transpositions. It is proved in IA Groups that this is well-defined.

Definition (Determinant). Let $A \in \text{Mat}_{n,n}(\mathbb{F})$. Its *determinant* is

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

This is a big scary definition. Hence, we will spend the first half of the chapter trying to understand what this really means, and how it behaves. We will eventually prove a formula that is useful for computing the determinant, which is probably how you were first exposed to the determinant.

Example. If $n = 2$, then $S_2 = \{\text{id}, (1\ 2)\}$. So

$$\det A = A_{11}A_{22} - A_{12}A_{21}.$$

When $n = 3$, then S_3 has 6 elements, and

$$\begin{aligned} \det A &= A_{11}A_{22}A_{33} + A_{12}A_{23}A_{31} + A_{13}A_{21}A_{32} \\ &\quad - A_{11}A_{23}A_{32} - A_{22}A_{31}A_{13} - A_{33}A_{12}A_{21}. \end{aligned}$$

We will first prove a few easy and useful lemmas about the determinant.

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

Proof.

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

Since $\varepsilon(\tau) = \varepsilon(\tau^{-1})$, we get

$$\begin{aligned} &= \sum_{\tau \in S_n} \varepsilon(\tau) \prod_{j=1}^n A_{j\tau(j)} \\ &= \det A. \end{aligned}$$

□

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 $i \leq \sigma(i)$ for all i is the identity. So the only thing that contributes in the sum is $\sigma = \text{id}$. So

$$\det A = \prod_{i=1}^n A_{ii}. \quad \square$$

To motivate this definition, we need a notion of volume. How can we define *volume* on a vector space? It should be clear that the “volume” cannot be uniquely determined, since it depends on what units we are using. For example, saying the volume is “1” is meaningless unless we provide the units, e.g. cm^3 . 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 $d : \mathbb{F}^n \times \cdots \times \mathbb{F}^n \rightarrow \mathbb{F}$ that is

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

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

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

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

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

We can view $A \in \text{Mat}_n(\mathbb{F})$ 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.

Proof. To see that \det is multilinear, it is sufficient to show that each

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

is multilinear for all $\sigma \in S_n$, since linear combinations of multilinear forms are multilinear. But each such product contains precisely one entry from each column, and so is multilinear.

To show it is alternating, suppose now there are some k, ℓ distinct such that $A^{(k)} = A^{(\ell)}$. We let τ be the transposition $(k \ell)$. By Lagrange's theorem, we can write

$$S_n = A_n \amalg \tau A_n,$$

where $A_n = \ker \varepsilon$ and \amalg is the disjoint union. We also know that

$$\sum_{\sigma \in A_n} \prod_{i=1}^n A_{i\sigma(i)} = \sum_{\sigma \in A_n} \prod_{i=1}^n A_{i,\tau\sigma(i)},$$

since if $\sigma(i)$ is not k or ℓ , then τ does nothing; if $\sigma(i)$ is k or ℓ , then τ just swaps them around, but $A^{(k)} = A^{(\ell)}$. So we get

$$\sum_{\sigma \in A_n} \prod_{i=1}^n A_{i\sigma(i)} = \sum_{\sigma' \in \tau A_n} \prod_{i=1}^n A_{i\sigma'(i)},$$

But we know that

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

So done. □

We have shown that determinants are volume forms, but is this the only volume form? Well obviously not, since $2 \det A$ is also a valid volume form. However, in some sense, all volume forms are “derived” from the determinant. Before we show that, we need the following

Lemma. Let d be a volume form on \mathbb{F}^n . Then swapping two entries changes the sign, i.e.

$$d(\mathbf{v}_1, \dots, \mathbf{v}_i, \dots, \mathbf{v}_j, \dots, \mathbf{v}_n) = -d(\mathbf{v}_1, \dots, \mathbf{v}_j, \dots, \mathbf{v}_i, \dots, \mathbf{v}_n).$$

Proof. By linearity, we have

$$\begin{aligned} 0 &= d(\mathbf{v}_1, \dots, \mathbf{v}_i + \mathbf{v}_j, \dots, \mathbf{v}_i + \mathbf{v}_j, \dots, \mathbf{v}_n) \\ &= d(\mathbf{v}_1, \dots, \mathbf{v}_i, \dots, \mathbf{v}_i, \dots, \mathbf{v}_n) \\ &\quad + d(\mathbf{v}_1, \dots, \mathbf{v}_i, \dots, \mathbf{v}_j, \dots, \mathbf{v}_n) \\ &\quad + d(\mathbf{v}_1, \dots, \mathbf{v}_j, \dots, \mathbf{v}_i, \dots, \mathbf{v}_n) \\ &\quad + d(\mathbf{v}_1, \dots, \mathbf{v}_j, \dots, \mathbf{v}_j, \dots, \mathbf{v}_n) \\ &= d(\mathbf{v}_1, \dots, \mathbf{v}_i, \dots, \mathbf{v}_j, \dots, \mathbf{v}_n) \\ &\quad + d(\mathbf{v}_1, \dots, \mathbf{v}_j, \dots, \mathbf{v}_i, \dots, \mathbf{v}_n). \end{aligned}$$

So done. □

Corollary. If $\sigma \in S_n$, then

$$d(\mathbf{v}_{\sigma(1)}, \dots, \mathbf{v}_{\sigma(n)}) = \varepsilon(\sigma)d(\mathbf{v}_1, \dots, \mathbf{v}_n)$$

for any $\mathbf{v}_i \in \mathbb{F}^n$.

Theorem. Let d be any volume form on \mathbb{F}^n , and let $A = (A^{(1)} \ \dots \ A^{(n)}) \in \text{Mat}_n(\mathbb{F})$. Then

$$d(A^{(1)}, \dots, A^{(n)}) = (\det A)d(\mathbf{e}_1, \dots, \mathbf{e}_n),$$

where $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ is the standard basis.

Proof. We can compute

$$\begin{aligned} d(A^{(1)}, \dots, A^{(n)}) &= d\left(\sum_{i=1}^n A_{i1}\mathbf{e}_i, A^{(2)}, \dots, A^{(n)}\right) \\ &= \sum_{i=1}^n A_{i1}d(\mathbf{e}_i, A^{(2)}, \dots, A^{(n)}) \\ &= \sum_{i,j=1}^n A_{i1}A_{j2}d(\mathbf{e}_i, \mathbf{e}_j, A^{(3)}, \dots, A^{(n)}) \\ &= \sum_{i_1, \dots, i_n} d(\mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_n}) \prod_{j=1}^n A_{i_j j}. \end{aligned}$$

We know that lots of these are zero, since if $i_k = i_j$ for some k, j , then the term is zero. So we are just summing over distinct tuples, i.e. when there is some σ such that $i_j = \sigma(j)$. So we get

$$d(A^{(1)}, \dots, A^{(n)}) = \sum_{\sigma \in S_n} d(\mathbf{e}_{\sigma(1)}, \dots, \mathbf{e}_{\sigma(n)}) \prod_{j=1}^n A_{\sigma(j)j}.$$

However, by our corollary up there, this is just

$$d(A^{(1)}, \dots, A^{(n)}) = \sum_{\sigma \in S_n} \varepsilon(\sigma)d(\mathbf{e}_1, \dots, \mathbf{e}_n) \prod_{j=1}^n A_{\sigma(j)j} = (\det A)d(\mathbf{e}_1, \dots, \mathbf{e}_n).$$

So done. □

We can rewrite the formula as

$$d(A\mathbf{e}_1, \dots, A\mathbf{e}_n) = (\det A)d(\mathbf{e}_1, \dots, \mathbf{e}_n).$$

It is not hard to see that the same proof gives for any $\mathbf{v}_1, \dots, \mathbf{v}_n$, we have

$$d(A\mathbf{v}_1, \dots, A\mathbf{v}_n) = (\det A)d(\mathbf{v}_1, \dots, \mathbf{v}_n).$$

So we know that $\det A$ is the volume rescaling factor of an arbitrary parallelepiped, and this is true for *any* volume form d .

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

Proof. Let d be a non-zero volume form on \mathbb{F}^n (e.g. the “determinant”). Then we can compute

$$d(AB\mathbf{e}_1, \dots, AB\mathbf{e}_n) = (\det AB)d(\mathbf{e}_1, \dots, \mathbf{e}_n),$$

but we also have

$$d(AB\mathbf{e}_1, \dots, AB\mathbf{e}_n) = (\det A)d(B\mathbf{e}_1, \dots, B\mathbf{e}_n) = (\det A)(\det B)d(\mathbf{e}_1, \dots, \mathbf{e}_n).$$

Since d is non-zero, we must have $\det AB = \det A \det B$. □

Corollary. If $A \in \text{Mat}_n(\mathbb{F})$ is invertible, then $\det A \neq 0$. In fact, when A is invertible, then $\det(A^{-1}) = (\det A)^{-1}$.

Proof. We have

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

So done. □

Definition (Singular matrices). A matrix A is *singular* if $\det A = 0$. Otherwise, it is *non-singular*.

We have just shown that if $\det A = 0$, then A is not invertible. Is the converse true? If $\det A \neq 0$, then can we conclude that A is invertible? The answer is yes. We are now going to prove it in an abstract and clean way. We will later prove this fact again by constructing an explicit formula for the inverse, which involves dividing by the determinant. So if the determinant is non-zero, then we know an inverse exists.

Theorem. Let $A \in \text{Mat}_n(\mathbb{F})$. Then the following are equivalent:

- (i) A is invertible.
- (ii) $\det A \neq 0$.
- (iii) $r(A) = n$.

Proof. We have proved that (i) \Rightarrow (ii) above, and the rank-nullity theorem implies (iii) \Rightarrow (i). We will prove (ii) \Rightarrow (iii). In fact we will show the contrapositive. Suppose $r(A) < n$. By rank-nullity theorem, $n(A) > 0$. So there is some

$\mathbf{x} = \begin{pmatrix} \lambda_1 \\ \vdots \\ \lambda_n \end{pmatrix}$ such that $A\mathbf{x} = \mathbf{0}$. Suppose $\lambda_k \neq 0$. We define B as follows:

$$B = \begin{pmatrix} 1 & & & \lambda_1 & & & & & & & \\ & \ddots & & \vdots & & & & & & & \\ & & 1 & \lambda_{k-1} & & & & & & & \\ & & & \lambda_k & & & & & & & \\ & & & \lambda_{k+1} & 1 & & & & & & \\ & & & \vdots & & \ddots & & & & & \\ & & & \lambda_n & & & & & & & 1 \end{pmatrix}$$

So AB has the k th column identically zero. So $\det(AB) = 0$. So it is sufficient to prove that $\det(B) \neq 0$. But $\det B = \lambda_k \neq 0$. So done. □

We are now going to come up with an alternative formula for the determinant (which is probably the one you are familiar with). To do so, we introduce the following notation:

Notation. Write \hat{A}_{ij} for the matrix obtained from A by deleting the i th row and j th column.

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

(i) We can expand $\det A$ along the j th column by

$$\det A = \sum_{i=1}^n (-1)^{i+j} A_{ij} \det \hat{A}_{ij}.$$

(ii) We can expand $\det A$ along the i th row by

$$\det A = \sum_{j=1}^n (-1)^{i+j} A_{ij} \det \hat{A}_{ij}.$$

We could prove this directly from the definition, but that is messy and scary, so let's use volume forms instead.

Proof. Since $\det A = \det A^T$, (i) and (ii) are equivalent. So it suffices to prove just one of them. We have

$$\det A = d(A^{(1)}, \dots, A^{(n)}),$$

where d is the volume form induced by the determinant. Then we can write as

$$\begin{aligned} \det A &= d\left(A^{(1)}, \dots, \sum_{i=1}^n A_{ij} \mathbf{e}_i, \dots, A^{(n)}\right) \\ &= \sum_{i=1}^n A_{ij} d(A^{(1)}, \dots, \mathbf{e}_i, \dots, A^{(n)}) \end{aligned}$$

The volume form on the right is the determinant of a matrix with the j th column replaced with \mathbf{e}_i . We can move our columns around so that our matrix becomes

$$B = \begin{pmatrix} \hat{A}_{ij} & 0 \\ \text{stuff} & 1 \end{pmatrix}$$

We get that $\det B = \det \hat{A}^{ij}$, since the only permutations that give a non-zero sum are those that send n to n . In the row and column swapping, we have made $n - j$ column transpositions and $n - i$ row transpositions. So we have

$$\begin{aligned} \det A &= \sum_{i=1}^n A_{ij} (-1)^{n-j} (-1)^{n-i} \det B \\ &= \sum_{i=1}^n A_{ij} (-1)^{i+j} \det \hat{A}_{ij}. \end{aligned} \quad \square$$

This is not only useful for computing determinants, but also computing inverses.

Definition (Adjugate matrix). Let $A \in \text{Mat}_n(\mathbb{F})$. The *adjugate matrix* of A , written $\text{adj } A$, is the $n \times n$ matrix such that $(\text{adj } A)_{ij} = (-1)^{i+j} \det \hat{A}_{ji}$.

The relevance is the following result:

Theorem. If $A \in \text{Mat}_n(\mathbb{F})$, then $A(\text{adj } A) = (\det A)I_n = (\text{adj } A)A$. In particular, if $\det A \neq 0$, then

$$A^{-1} = \frac{1}{\det A} \text{adj } A.$$

Note that this is *not* an efficient way to compute the inverse.

Proof. We compute

$$[(\text{adj } A)A]_{jk} = \sum_{i=1}^n (\text{adj } A)_{ji} A_{ik} = \sum_{i=1}^n (-1)^{i+j} \det \hat{A}_{ij} A_{ik}. \quad (*)$$

So if $j = k$, then $[(\text{adj } A)A]_{jk} = \det A$ by the lemma.

Otherwise, if $j \neq k$, consider the matrix B obtained from A by replacing the j th column by the k th column. Then the right hand side of $(*)$ is just $\det B$ by the lemma. But we know that if two columns are the same, the determinant is zero. So the right hand side of $(*)$ is zero. So

$$[(\text{adj } A)A]_{jk} = \det A \delta_{jk}$$

The calculation for $[A \text{adj } A] = (\det A)I_n$ can be done in a similar manner, or by considering $(A \text{adj } A)^T = (\text{adj } A)^T A^T = (\text{adj } (A^T))A^T = (\det A)I_n$. \square

Note that the coefficients of $(\text{adj } A)$ are just given by polynomials in the entries of A , and so is the determinant. So if A is invertible, then its inverse is given by a rational function (i.e. ratio of two polynomials) in the entries of A .

This is very useful theoretically, but not computationally, since the polynomials are very large. There are better ways computationally, such as Gaussian elimination.

We'll end with a useful tricks to compute the determinant.

Lemma. Let A, B be square matrices. Then for any C , we have

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

Proof. Suppose $A \in \text{Mat}_k(\mathbb{F})$, and $B \in \text{Mat}_\ell(\mathbb{F})$, so $C \in \text{Mat}_{k,\ell}(\mathbb{F})$. Let

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

Then by definition, we have

$$\det X = \sum_{\sigma \in S_{k+\ell}} \varepsilon(\sigma) \prod_{i=1}^{k+\ell} X_{i\sigma(i)}.$$

If $j \leq k$ and $i > k$, then $X_{ij} = 0$. We only want to sum over permutations σ such that $\sigma(i) > k$ if $i > k$. So we are permuting the last j things among themselves, and hence the first k things among themselves. So we can decompose this into $\sigma = \sigma_1\sigma_2$, 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

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

□

Corollary.

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

6 Endomorphisms

Endomorphisms are linear maps from a vector space V to itself. One might wonder — why would we want to study these linear maps in particular, when we can just work with arbitrary linear maps from any space to any other space?

When we work with arbitrary linear maps, we are free to choose any basis for the domain, and any basis for the co-domain, since it doesn't make sense to require they have the “same” basis. Then we proved that by choosing the right bases, we can put matrices into a nice form with only 1's in the diagonal.

However, when working with endomorphisms, we can require ourselves to use the same basis for the domain and co-domain, and there is much more we can say. One major objective is to classify all matrices up to similarity, where two matrices are similar if they represent the same endomorphism under different bases.

6.1 Invariants

Definition. If V is a (finite-dimensional) vector space over \mathbb{F} . An *endomorphism* of V is a linear map $\alpha : V \rightarrow V$. We write $\text{End}(V)$ for the \mathbb{F} -vector space of all such linear maps, and I for the identity map $V \rightarrow V$.

When we think about matrices representing an endomorphism of V , we'll use the same basis for the domain and the range. We are going to study some properties of these endomorphisms that are not dependent on the basis we pick, known as *invariants*.

Lemma. Suppose $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ and $(\mathbf{f}_1, \dots, \mathbf{f}_n)$ are bases for V and $\alpha \in \text{End}(V)$. If A represents α with respect to $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ and B represents α with respect to $(\mathbf{f}_1, \dots, \mathbf{f}_n)$, then

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

where P is given by

$$\mathbf{f}_i = \sum_{j=1}^n P_{ji} \mathbf{e}_j.$$

Proof. This is merely a special case of an earlier more general result for arbitrary maps and spaces. \square

Definition (Similar matrices). We say matrices A and B are *similar* or *conjugate* if there is some P invertible such that $B = P^{-1}AP$.

Recall that $\text{GL}_n(\mathbb{F})$, the group of invertible $n \times n$ matrices. $\text{GL}_n(\mathbb{F})$ acts on $\text{Mat}_n(\mathbb{F})$ by conjugation:

$$(P, A) \mapsto PAP^{-1}.$$

We are conjugating it this way so that the associativity axiom holds (otherwise we get a *right* action instead of a *left* action). Then A and B are similar iff they are in the same orbit. Since orbits always partition the set, this is an equivalence relation.

Our main goal is to classify the orbits, i.e. find a “nice” representative for each orbit.

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 $A \in \text{Mat}_n(\mathbb{F})$ is the sum of the diagonal elements:

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

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

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 $A \in \text{Mat}_{m,n}(\mathbb{F})$ and $B \in \text{Mat}_{n,m}(\mathbb{F})$, then

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

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

(iii) If $A, B \in \text{Mat}_n(\mathbb{F})$ are similar, then $\det A = \det B$.

Proof.

(i) We have

$$\text{tr } AB = \sum_{i=1}^m (AB)_{ii} = \sum_{i=1}^m \sum_{j=1}^n A_{ij} B_{ji} = \sum_{j=1}^n \sum_{i=1}^m B_{ji} A_{ij} = \text{tr } BA.$$

(ii) Suppose $B = P^{-1}AP$. Then we have

$$\text{tr } B = \text{tr}(P^{-1}(AP)) = \text{tr}((AP)P^{-1}) = \text{tr } A.$$

(iii) We have

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

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 $\text{tr } \alpha = \text{tr } A$, and the *determinant* is $\det \alpha = \det A$.

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.

Definition (Eigenvalue and eigenvector). Let $\alpha \in \text{End}(V)$. Then $\lambda \in \mathbb{F}$ is an *eigenvalue* (or E-value) if there is some $\mathbf{v} \in V \setminus \{0\}$ such that $\alpha\mathbf{v} = \lambda\mathbf{v}$.

\mathbf{v} is an *eigenvector* if $\alpha(\mathbf{v}) = \lambda\mathbf{v}$ for some $\lambda \in \mathbb{F}$.

When $\lambda \in \mathbb{F}$, the λ -*eigenspace*, written $E_\alpha(\lambda)$ or $E(\lambda)$ is the subspace of V containing all the λ -eigenvectors, i.e.

$$E_\alpha(\lambda) = \ker(\lambda\iota - \alpha).$$

where ι is the identity function.

Definition (Characteristic polynomial). The *characteristic polynomial* of α is defined by

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

You might be used to the definition $\chi_\alpha(t) = \det(\alpha - t\iota)$ instead. These two definitions are obviously equivalent up to a factor of -1 , but this definition has an advantage that $\chi_\alpha(t)$ is always monic, i.e. the leading coefficient is 1. However, when doing computations in reality, we often use $\det(\alpha - t\iota)$ instead, since it is easier to negate $t\iota$ than α .

We know that λ is an eigenvalue of α iff $n(\alpha - \lambda\iota) > 0$ iff $r(\alpha - \lambda\iota) < \dim V$ iff $\chi_\alpha(\lambda) = \det(\lambda\iota - \alpha) = 0$. So the eigenvalues are precisely the roots of the characteristic polynomial.

If $A \in \text{Mat}_n(\mathbb{F})$, we can define $\chi_A(t) = \det(tI - A)$.

Lemma. If A and B are similar, then they have the same characteristic polynomial.

Proof.

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

Lemma. Let $\alpha \in \text{End}(V)$ and $\lambda_1, \dots, \lambda_k$ distinct eigenvalues of α . Then

$$E(\lambda_1) + \dots + E(\lambda_k) = \bigoplus_{i=1}^k E(\lambda_i)$$

is a direct sum.

Proof. Suppose

$$\sum_{i=1}^k \mathbf{x}_i = \sum_{i=1}^k \mathbf{y}_i,$$

with $\mathbf{x}_i, \mathbf{y}_i \in E(\lambda_i)$. We want to show that they are equal. We are going to find some clever map that tells us what \mathbf{x}_i and \mathbf{y}_i are. Consider $\beta_j \in \text{End}(V)$ defined by

$$\beta_j = \prod_{r \neq j} (\alpha - \lambda_r \iota).$$

Then

$$\begin{aligned} \beta_j \left(\sum_{i=1}^k \mathbf{x}_i \right) &= \sum_{i=1}^k \prod_{r \neq j} (\alpha - \lambda_r \iota)(\mathbf{x}_i) \\ &= \sum_{i=1}^k \prod_{r \neq j} (\lambda_i - \lambda_r)(\mathbf{x}_i). \end{aligned}$$

Each summand is zero, unless $i \neq j$. So this is equal to

$$\beta_j \left(\sum_{i=1}^k \mathbf{x}_i \right) = \prod_{r \neq j} (\lambda_j - \lambda_r) (\mathbf{x}_j).$$

Similarly, we obtain

$$\beta_j \left(\sum_{i=1}^k \mathbf{y}_i \right) = \prod_{r \neq j} (\lambda_j - \lambda_r) (\mathbf{y}_j).$$

Since we know that $\sum \mathbf{x}_i = \sum \mathbf{y}_i$, we must have

$$\prod_{r \neq j} (\lambda_j - \lambda_r) \mathbf{x}_j = \prod_{r \neq j} (\lambda_j - \lambda_r) \mathbf{y}_j.$$

Since we know that $\prod_{r \neq j} (\lambda_r - \lambda_j) \neq 0$, we must have $\mathbf{x}_i = \mathbf{y}_i$ for all i .

So each expression for $\sum \mathbf{x}_i$ is unique. \square

The proof shows that any set of non-zero eigenvectors with distinct eigenvalues is linearly independent.

Definition (Diagonalizable). We say $\alpha \in \text{End}(V)$ is diagonalizable if there is some basis for V such that α is represented by a diagonal matrix, i.e. all terms not on the diagonal are zero.

These are in some sense the nice matrices we like to work with.

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

- (i) α is diagonalizable.
- (ii) V has a basis of eigenvectors for α .
- (iii) $V = \bigoplus_{i=1}^k E_i$.
- (iv) $\dim V = \sum_{i=1}^k \dim E_i$.

Proof.

- (i) \Leftrightarrow (ii): Suppose $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ is a basis for V . Then

$$\alpha(\mathbf{e}_i) = A_{ji} \mathbf{e}_j,$$

where A represents α . Then A is diagonal iff each \mathbf{e}_i is an eigenvector. So done

- (ii) \Leftrightarrow (iii): It is clear that (ii) is true iff $\sum E_i = V$, but we know that this must be a direct sum. So done.
- (iii) \Leftrightarrow (iv): This follows from example sheet 1 Q10, which says that $V = \bigoplus_{i=1}^k E_i$ iff the bases for E_i are disjoint and their union is a basis of V . \square

6.2 The minimal polynomial

6.2.1 Aside on polynomials

Before we talk about minimal polynomials, we first talk about polynomials in general.

Definition (Polynomial). A *polynomial* over \mathbb{F} is an object of the form

$$f(t) = a_m t^m + a_{m-1} t^{m-1} + \cdots + a_1 t + a_0,$$

with $m \geq 0, a_0, \dots, a_m \in \mathbb{F}$.

We write $\mathbb{F}[t]$ for the set of polynomials over \mathbb{F} .

Note that we don't identify a polynomial f with the corresponding function it represents. For example, if $\mathbb{F} = \mathbb{Z}/p\mathbb{Z}$, then t^p and t are different polynomials, even though they define the same function (by Fermat's little theorem/Lagrange's theorem). Two polynomials are equal if and only if they have the same coefficients.

However, we will later see that if \mathbb{F} is \mathbb{R} or \mathbb{C} , then polynomials are equal if and only if they represent the same function, and this distinction is not as important.

Definition (Degree). Let $f \in \mathbb{F}[t]$. Then the *degree* of f , written $\deg f$ is the largest n such that $a_n \neq 0$. In particular, $\deg 0 = -\infty$.

Notice that $\deg fg = \deg f + \deg g$ and $\deg f + g \leq \max\{\deg f, \deg g\}$.

Lemma (Polynomial division). If $f, g \in \mathbb{F}[t]$ (and $g \neq 0$), then there exists $q, r \in \mathbb{F}[t]$ with $\deg r < \deg g$ such that

$$f = qg + r.$$

Proof is omitted.

Lemma. If $\lambda \in \mathbb{F}$ is a root of f , i.e. $f(\lambda) = 0$, then there is some g such that

$$f(t) = (t - \lambda)g(t).$$

Proof. By polynomial division, let

$$f(t) = (t - \lambda)g(t) + r(t)$$

for some $g(t), r(t) \in \mathbb{F}[t]$ with $\deg r < \deg(t - \lambda) = 1$. So r has to be constant, i.e. $r(t) = a_0$ for some $a_0 \in \mathbb{F}$. Now evaluate this at λ . So

$$0 = f(\lambda) = (\lambda - \lambda)g(\lambda) + r(\lambda) = a_0.$$

So $a_0 = 0$. So $r = 0$. So done. \square

Definition (Multiplicity of a root). Let $f \in \mathbb{F}[t]$ and λ a root of f . We say λ has *multiplicity* k if $(t - \lambda)^k$ is a factor of f but $(t - \lambda)^{k+1}$ is not, i.e.

$$f(t) = (t - \lambda)^k g(t)$$

for some $g(t) \in \mathbb{F}[t]$ with $g(\lambda) \neq 0$.

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 $n \geq \deg(f - g)$ roots, we must have $f - g = 0$. So $f = g$. \square

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 polynomial over \mathbb{C} has a root in \mathbb{C} .

We will not prove this.

We say \mathbb{C} is an *algebraically closed field*.

It thus follows that every polynomial over \mathbb{C} of degree $n > 0$ has precisely n roots, counted with multiplicity, since if we write $f(t) = g(t) \prod (t - \lambda_i)^{a_i}$ and g has no roots, then g is constant. So the number of roots is $\sum a_i = \deg f$, counted with multiplicity.

It also follows that every polynomial in \mathbb{R} factors into linear polynomials and quadratic polynomials with no real roots (since complex roots of real polynomials come in complex conjugate pairs).

6.2.2 Minimal polynomial

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

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

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

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

So each $\mathbf{v} \in V$ can be written (uniquely) as

$$\mathbf{v} = \sum_{i=1}^k \mathbf{v}_i \text{ with } \alpha(\mathbf{v}_i) = \lambda_i \mathbf{v}_i.$$

Now let

$$p(t) = \prod_{i=1}^k (t - \lambda_i).$$

Then for any \mathbf{v} , we get

$$p(\alpha)(\mathbf{v}) = \sum_{i=1}^k p(\alpha)(\mathbf{v}_i) = \sum_{i=1}^k p(\lambda_i) \mathbf{v}_i = \mathbf{0}.$$

So $p(\alpha) = 0$. By construction, p has distinct linear factors.

Conversely, suppose we have our polynomial

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

with $\lambda_1, \dots, \lambda_k \in \mathbb{F}$ distinct, and $p(\alpha) = 0$ (we can wlog assume p is monic, i.e. the leading coefficient is 1). We will show that

$$V = \sum_{i=1}^k E_{\alpha}(\lambda_i).$$

In other words, we want to show that for all $\mathbf{v} \in V$, there is some $\mathbf{v}_i \in E_{\alpha}(\lambda_i)$ for $i = 1, \dots, k$ such that $\mathbf{v} = \sum \mathbf{v}_i$.

To find these \mathbf{v}_i out, we let

$$q_j(t) = \prod_{i \neq j} \frac{t - \lambda_i}{\lambda_j - \lambda_i}.$$

This is a polynomial of degree $k - 1$, and $q_j(\lambda_i) = \delta_{ij}$.

Now consider

$$q(t) = \sum_{i=1}^k q_i(t).$$

We still have $\deg q \leq k - 1$, but $q(\lambda_i) = 1$ for any i . Since q and 1 agree on k points, we must have $q = 1$.

Let $\pi_j : V \rightarrow V$ be given by $\pi_j = q_j(\alpha)$. Then the above says that

$$\sum_{j=1}^k \pi_j = \iota.$$

Hence given $\mathbf{v} \in V$, we know that $\mathbf{v} = \sum \pi_j \mathbf{v}$.

We now check that $\pi_j \mathbf{v} \in E_{\alpha}(\lambda_j)$. This is true since

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

So

$$\alpha \mathbf{v}_j = \lambda_j \mathbf{v}_j.$$

So done. \square

In the above proof, if $\mathbf{v} \in E_\alpha(\lambda_i)$, then $\pi_j(\mathbf{v}) = \delta_{ij}\mathbf{v}$. So π_i is a projection onto the $E_\alpha(\lambda_i)$.

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

The monic requirement is just for things to look nice, since we can always divide by the leading coefficient of a polynomial to get a monic version.

Note that if A represents α , then for all $p \in \mathbb{F}[t]$, $p(A)$ represents $p(\alpha)$. Thus $p(\alpha)$ is zero iff $p(A) = 0$. So the minimal polynomial of α is the minimal polynomial of A if we define M_A analogously.

There are two things we want to know — whether the minimal polynomial exists, and whether it is unique.

Existence is always guaranteed in finite-dimensional cases. If $\dim V = n < \infty$, then $\dim \text{End}(V) = n^2$. So $I, \alpha, \alpha^2, \dots, \alpha^{n^2}$ are linearly dependent. So there are some $\lambda_0, \dots, \lambda_{n^2} \in \mathbb{F}$ not all zero such that

$$\sum_{i=0}^{n^2} \lambda_i \alpha^i = 0.$$

So $\deg M_\alpha \leq n^2$. So we must have a minimal polynomial.

To show that the minimal polynomial is unique, we will prove the following stronger characterization of the minimal polynomial:

Lemma. Let $\alpha \in \text{End}(V)$, and $p \in \mathbb{F}[t]$. Then $p(\alpha) = 0$ if and only if $M_\alpha(t)$ is a factor of $p(t)$. In particular, M_α is unique.

Proof. For all such p , we can write $p(t) = q(t)M_\alpha(t) + r(t)$ for some r of degree less than $\deg M_\alpha$. Then

$$p(\alpha) = q(\alpha)M_\alpha(\alpha) + r(\alpha).$$

So if $r(\alpha) = 0$ iff $p(\alpha) = 0$. But $\deg r < \deg M_\alpha$. By the minimality of M_α , we must have $r(\alpha) = 0$ iff $r = 0$. So $p(\alpha) = 0$ iff $M_\alpha(t) \mid p(t)$.

So if M_1 and M_2 are both minimal polynomials for α , then $M_1 \mid M_2$ and $M_2 \mid M_1$. So M_2 is just a scalar multiple of M_1 . But since M_1 and M_2 are monic, they must be equal. \square

Example. Let $V = \mathbb{F}^2$, and consider the following matrices:

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad B = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

Consider the polynomial $p(t) = (t-1)^2$. We can compute $p(A) = p(B) = 0$. So $M_A(t)$ and $M_B(t)$ are factors of $(t-1)^2$. There aren't many factors of $(t-1)^2$. So the minimal polynomials are either $(t-1)$ or $(t-1)^2$. Since $A - I = 0$ and $B - I \neq 0$, the minimal polynomial of A is $t-1$ and the minimal polynomial of B is $(t-1)^2$.

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. \square

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

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

where λ_i are the eigenvalues of V . We write E_i for $E_\alpha(\lambda_i)$. We want to show that β sends E_i to itself, i.e. $\beta(E_i) \subseteq E_i$. Let $\mathbf{v} \in E_i$. Then we want $\beta(\mathbf{v})$ to be in E_i . This is true since

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

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. \square

6.3 The Cayley-Hamilton theorem

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

Recall that $\chi_\alpha(t) = \det(tI - \alpha)$ for $\alpha \in \text{End}(V)$. 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, $\deg M_\alpha \leq n$.

We will not prove this yet, but just talk about it first. It is tempting to prove this by substituting $t = \alpha$ into $\det(tI - \alpha)$ and get $\det(\alpha - \alpha) = 0$, but this is meaningless, since what the statement $\chi_\alpha(t) = \det(tI - \alpha)$ tells us to do is to expand the determinant of the matrix

$$\begin{pmatrix} t - a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & t - a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & t - a_{nn} \end{pmatrix}$$

to obtain a polynomial, and we clearly cannot substitute $t = A$ in this expression. However, we can later show that we can use this idea to prove it, but just be a bit more careful.

Note also that if $\rho(t) \in \mathbb{F}[t]$ and

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

then

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

Since $\chi_A(t)$ is defined as $\prod_{i=1}^n (t - \lambda_i)$, it follows that $\chi_A(A) = 0$. So if α is diagonalizable, then the theorem is clear.

This was easy. Diagonalizable matrices are nice. The next best thing we can look at is upper-triangular matrices.

Definition (Triangulable). An endomorphism $\alpha \in \text{End}(V)$ is *triangulable* if there is a basis for V such that α is represented by an upper triangular matrix

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

We have a similar lemma telling us when matrices are triangulable.

Lemma. An endomorphism α is triangulable if and only if $\chi_\alpha(t)$ can be written as a product of linear factors, not necessarily distinct. In particular, if $\mathbb{F} = \mathbb{C}$ (or any algebraically closed field), then every endomorphism is triangulable.

Proof. Suppose that α is triangulable and represented by

$$\begin{pmatrix} \lambda_1 & * & \cdots & * \\ 0 & \lambda_2 & \cdots & * \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_n \end{pmatrix}.$$

Then

$$\chi_\alpha(t) = \det \begin{pmatrix} t - \lambda_1 & * & \cdots & * \\ 0 & t - \lambda_2 & \cdots & * \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & t - \lambda_n \end{pmatrix} = \prod_{i=1}^n (t - \lambda_i).$$

So it is a product of linear factors.

We are going to prove the converse by induction on the dimension of our space. The base case $\dim V = 1$ is trivial, since every 1×1 matrix is already upper triangular.

Suppose $\alpha \in \text{End}(V)$ and the result holds for all spaces of dimensions $< \dim V$, and χ_α is a product of linear factors. In particular, $\chi_\alpha(t)$ has a root, say $\lambda \in \mathbb{F}$.

Now let $U = E(\lambda) \neq 0$, and let W be a complementary subspace to U in V , i.e. $V = U \oplus W$. Let $\mathbf{u}_1, \dots, \mathbf{u}_r$ be a basis for U and $\mathbf{w}_{r+1}, \dots, \mathbf{w}_n$ be a basis for W so that $\mathbf{u}_1, \dots, \mathbf{u}_r, \mathbf{w}_{r+1}, \dots, \mathbf{w}_n$ is a basis for V , and α is represented by

$$\begin{pmatrix} \lambda I_r & \text{stuff} \\ 0 & B \end{pmatrix}$$

We know $\chi_\alpha(t) = (t - \lambda)^r \chi_B(t)$. So $\chi_B(t)$ is also a product of linear factors. We let $\beta : W \rightarrow W$ be the map defined by B with respect to $\mathbf{w}_{r+1}, \dots, \mathbf{w}_n$.

(Note that in general, β is not $\alpha|_W$ in general, since α does not necessarily map W to W . However, we can say that $(\alpha - \beta)(\mathbf{w}) \in U$ for all $\mathbf{w} \in W$. This can be much more elegantly expressed in terms of quotient spaces, but unfortunately that is not officially part of the course)

Since $\dim W < \dim V$, there is a basis $\mathbf{v}_{r+1}, \dots, \mathbf{v}_n$ for W such that β is represented by C , which is upper triangular.

For $j = 1, \dots, n - r$, we have

$$\alpha(\mathbf{v}_{j+r}) = \mathbf{u} + \sum_{k=1}^{n-r} C_{kj} \mathbf{v}_{k+r}$$

for some $\mathbf{u} \in U$. So α is represented by

$$\begin{pmatrix} \lambda I_r & \text{stuff} \\ 0 & C \end{pmatrix}$$

with respect to $(\mathbf{u}_1, \dots, \mathbf{u}_r, \mathbf{v}_{r+1}, \dots, \mathbf{v}_n)$, which is upper triangular. \square

Example. Consider the real rotation matrix

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

This is *not* similar to a real upper triangular matrix (if θ is not an integer multiple of π). This is since the eigenvalues are $e^{\pm i\theta}$ and are not real. On the other hand, as a complex matrix, it is triangulable, and in fact diagonalizable since the eigenvalues are distinct.

For this reason, in the rest of the section, we are mostly going to work in \mathbb{C} . We can now prove the Cayley-Hamilton theorem.

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, $\deg M_\alpha \leq n$.

Proof. In this proof, we will work over \mathbb{C} . By the lemma, we can choose a basis $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ is represented by an upper triangular matrix.

$$A = \begin{pmatrix} \lambda_1 & * & \cdots & * \\ 0 & \lambda_2 & \cdots & * \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_n \end{pmatrix}.$$

We must prove that

$$\chi_\alpha(\alpha) = \chi_A(\alpha) = \prod_{i=1}^n (\alpha - \lambda_i \iota) = 0.$$

Write $V_j = \langle \mathbf{e}_1, \dots, \mathbf{e}_j \rangle$. So we have the inclusions

$$V_0 = 0 \subseteq V_1 \subseteq \cdots \subseteq V_{n-1} \subseteq V_n = V.$$

We also know that $\dim V_j = j$. This increasing sequence is known as a *flag*.

Now note that since A is upper-triangular, we get

$$\alpha(\mathbf{e}_i) = \sum_{k=1}^i A_{ki} \mathbf{e}_k \in V_i.$$

So $\alpha(V_j) \subseteq V_j$ for all $j = 0, \dots, n$.

Moreover, we have

$$(\alpha - \lambda_j \iota)(\mathbf{e}_j) = \sum_{k=1}^{j-1} A_{kj} \mathbf{e}_k \subseteq V_{j-1}$$

for all $j = 1, \dots, n$. So every time we apply one of these things, we get to a smaller space. Hence by induction on $n - j$, we have

$$\prod_{i=j}^n (\alpha - \lambda_i \iota)(V_n) \subseteq V_{j-1}.$$

In particular, when $j = 1$, we get

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

So $\chi_\alpha(\alpha) = 0$ as required.

Note that if our field \mathbb{F} is not \mathbb{C} but just a subfield of \mathbb{C} , say \mathbb{R} , we can just pretend it is a complex matrix, do the same proof. \square

We can see this proof more “visually” as follows: for simplicity of expression, we suppose $n = 4$. In the basis where α is upper-triangular, the matrices $A - \lambda_i I$ look like this

$$\begin{aligned}
 A - \lambda_1 I &= \begin{pmatrix} 0 & * & * & * \\ 0 & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & * \end{pmatrix} & A - \lambda_2 I &= \begin{pmatrix} * & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & * \end{pmatrix} \\
 A - \lambda_3 I &= \begin{pmatrix} * & * & * & * \\ 0 & * & * & * \\ 0 & 0 & 0 & * \\ 0 & 0 & 0 & * \end{pmatrix} & A - \lambda_4 I &= \begin{pmatrix} * & * & * & * \\ 0 & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & 0 \end{pmatrix}
 \end{aligned}$$

Then we just multiply out directly (from the right):

$$\begin{aligned}
 \prod_{i=1}^4 (A - \lambda_i I) &= \begin{pmatrix} 0 & * & * & * \\ 0 & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & * \end{pmatrix} \begin{pmatrix} * & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & * \end{pmatrix} \begin{pmatrix} * & * & * & * \\ 0 & * & * & * \\ 0 & 0 & 0 & * \\ 0 & 0 & 0 & * \end{pmatrix} \begin{pmatrix} * & * & * & * \\ 0 & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & 0 \end{pmatrix} \\
 &= \begin{pmatrix} 0 & * & * & * \\ 0 & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & * \end{pmatrix} \begin{pmatrix} * & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & * \end{pmatrix} \begin{pmatrix} * & * & * & * \\ 0 & * & * & * \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \\
 &= \begin{pmatrix} 0 & * & * & * \\ 0 & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & * \end{pmatrix} \begin{pmatrix} * & * & * & * \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \\
 &= \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.
 \end{aligned}$$

This is exactly what we showed in the proof — after multiplying out the first k elements of the product (counting from the right), the image is contained in the span of the first $n - k$ basis vectors.

Proof. We’ll now prove the theorem again, which is somewhat a formalization of the “nonsense proof” where we just substitute $t = \alpha$ into $\det(\alpha - tI)$.

Let α be represented by A , and $B = tI - A$. Then

$$B \operatorname{adj} B = \det BI_n = \chi_\alpha(t)I_n.$$

But we know that $\operatorname{adj} B$ is a matrix with entries in $\mathbb{F}[t]$ of degree at most $n - 1$. So we can write

$$\operatorname{adj} B = B_{n-1}t^{n-1} + B_{n-2}t^{n-2} + \cdots + B_0,$$

with $B_i \in \operatorname{Mat}_n(\mathbb{F})$. We can also write

$$\chi_\alpha(t) = t^n + a_{n-1}t^{n-1} + \cdots + a_0.$$

Then we get the result

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

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

$$\begin{pmatrix} t - a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & t - a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & t - a_{nn} \end{pmatrix}$$

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

$$\begin{aligned} -AB_0 &= a_0I_n \\ B_0 - AB_1 &= a_1I_n \\ &\vdots \\ B_{n-2} - AB_{n-1} &= a_{n-1}I_n \\ AB_{n-1} - 0 &= I_n \end{aligned}$$

We now multiply each row a suitable power of A to obtain

$$\begin{aligned} -AB_0 &= a_0I_n \\ AB_0 - A^2B_1 &= a_1A \\ &\vdots \\ A^{n-1}B_{n-2} - A^nB_{n-1} &= a_{n-1}A^{n-1} \\ A^nB_{n-1} - 0 &= A^n. \end{aligned}$$

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 $(\alpha - \lambda I)(\mathbf{v}) = 0$ has a non-trivial root, iff $\det(\alpha - \lambda I) = 0$.
- (iii) \Rightarrow (ii): This follows from Cayley-Hamilton theorem since $M_\alpha \mid \chi_\alpha$.

- (i) \Rightarrow (iii): Let λ be an eigenvalue, and \mathbf{v} be a corresponding eigenvector. Then by definition of M_α , we have

$$M_\alpha(\alpha)(\mathbf{v}) = 0(\mathbf{v}) = 0.$$

We also know that

$$M_\alpha(\alpha)(\mathbf{v}) = M_\alpha(\lambda)\mathbf{v}.$$

Since \mathbf{v} is non-zero, we must have $M_\alpha(\lambda) = 0$.

- (iii) \Rightarrow (i): This is not necessary since it follows from the above, but we could as well do it explicitly. Suppose λ is a root of $M_\alpha(t)$. Then $M_\alpha(t) = (t - \lambda)g(t)$ for some $g \in \mathbb{F}[t]$. But $\deg g < \deg M_\alpha$. Hence by minimality of M_α , we must have $g(\alpha) \neq 0$. So there is some $\mathbf{v} \in V$ such that $g(\alpha)(\mathbf{v}) \neq 0$. Then

$$(\alpha - \lambda)g(\alpha)(\mathbf{v}) = M_\alpha(\alpha)\mathbf{v} = 0.$$

So we must have $\alpha(g(\alpha)(\mathbf{v})) = \lambda g(\alpha)(\mathbf{v})$. So $g(\alpha)(\mathbf{v}) \in E_\alpha(\lambda) \setminus \{0\}$. So (i) holds. \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 we know that the minimal polynomial is one of $(t-1)^2(t-2)$ and $(t-1)(t-2)$.

By direct and boring computations, we can find $(A - I)(A - 2I) = 0$. So we know that $M_A(t) = (t-1)(t-2)$. So A is diagonalizable.

6.4 Multiplicities of eigenvalues and Jordan normal form

We will want to put our matrices in their “Jordan normal forms”, which is a unique form for each equivalence class of similar matrices. The following properties will help determine which Jordan normal form a matrix can have.

Definition (Algebraic and geometry multiplicity). Let $\alpha \in \text{End}(V)$ and λ an eigenvalue of α . Then

- (i) The *algebraic multiplicity* of λ , written a_λ , is the multiplicity of λ as a root of $\chi_\alpha(t)$.
- (ii) The *geometric multiplicity* of λ , written g_λ , is the dimension of the corresponding eigenspace, $\dim E_\alpha(\lambda)$.
- (iii) c_λ is the multiplicity of λ as a root of the minimal polynomial $m_\alpha(t)$.

We will look at some extreme examples:

Example.

– Let

$$A = \begin{pmatrix} \lambda & 1 & \cdots & 0 \\ 0 & \lambda & \ddots & \vdots \\ \vdots & \vdots & \ddots & 1 \\ 0 & 0 & \cdots & \lambda \end{pmatrix}.$$

We will later show that $a_\lambda = n = c_\lambda$ and $g_\lambda = 1$.

– Consider $A = \lambda I$. Then $a_\lambda = g_\lambda = n$, $c_\lambda = 1$.

Lemma. If λ is an eigenvalue of α , then

(i) $1 \leq g_\lambda \leq a_\lambda$

(ii) $1 \leq c_\lambda \leq a_\lambda$.

Proof.

(i) The first inequality is easy. If λ is an eigenvalue, then $E(\lambda) \neq 0$. So $g_\lambda = \dim E(\lambda) \geq 1$. To prove the other inequality, if $\mathbf{v}_1, \dots, \mathbf{v}_g$ is a basis for $E(\lambda)$, then we can extend it to a basis for V , and then α is represented by

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

So $\chi_\alpha(t) = (t - \lambda)^g \chi_B(t)$. So $a_\lambda > g = g_\lambda$.

(ii) This is straightforward since $M_\alpha(\lambda) = 0$ implies $1 \leq c_\lambda$, and since $M_\alpha(t) \mid \chi_\alpha(t)$, we know that $c_\lambda \leq a_\lambda$. \square

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

(i) α is diagonalizable.

(ii) $g_\lambda = a_\lambda$ for all eigenvalues of α .

(iii) $c_\lambda = 1$ for all λ .

Proof.

– (i) \Leftrightarrow (ii): α is diagonalizable iff $\dim V = \sum \dim E_\alpha(\lambda_i)$. But this is equivalent to

$$\dim V = \sum g_{\lambda_i} \leq \sum a_{\lambda_i} = \deg \chi_\alpha = \dim V.$$

So we must have $\sum g_{\lambda_i} = \sum a_{\lambda_i}$. Since each g_{λ_i} is at most a_{λ_i} , they must be individually equal.

– (i) \Leftrightarrow (iii): α is diagonalizable if and only if $M_\alpha(t)$ is a product of distinct linear factors if and only if $c_\lambda = 1$ for all eigenvalues λ . \square

Definition (Jordan normal form). We say $A \in \text{Mat}_N(\mathbb{C})$ is in *Jordan normal form* if it is a block diagonal of the form

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

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
$\begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix}$	$(t - \lambda_1)(t - \lambda_2)(t - \lambda_3)$	$(t - \lambda_1)(t - \lambda_2)(t - \lambda_3)$
$\begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_1 & 0 \\ 0 & 0 & \lambda_2 \end{pmatrix}$	$(t - \lambda_1)^2(t - \lambda_2)$	$(t - \lambda_1)(t - \lambda_2)$
$\begin{pmatrix} \lambda_1 & 1 & 0 \\ 0 & \lambda_1 & 0 \\ 0 & 0 & \lambda_2 \end{pmatrix}$	$(t - \lambda_1)^2(t - \lambda_2)$	$(t - \lambda_1)^2(t - \lambda_2)$
$\begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_1 & 0 \\ 0 & 0 & \lambda_1 \end{pmatrix}$	$(t - \lambda_1)^3$	$(t - \lambda_1)$
$\begin{pmatrix} \lambda_1 & 1 & 0 \\ 0 & \lambda_1 & 0 \\ 0 & 0 & \lambda_1 \end{pmatrix}$	$(t - \lambda_1)^3$	$(t - \lambda_1)^2$
$\begin{pmatrix} \lambda_1 & 1 & 0 \\ 0 & \lambda_1 & 1 \\ 0 & 0 & \lambda_1 \end{pmatrix}$	$(t - \lambda_1)^3$	$(t - \lambda_1)^3$

Notice that χ_A and M_A together determine the Jordan normal form of a 3×3 complex matrix. We do indeed need χ_A in the second case, since if we are given $M_A = (t - \lambda_1)(t - \lambda_2)$, we know one 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 $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ is the standard basis for \mathbb{C}^n , we have

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

Thus we know

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

In other words, for $k < n$, we have

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

If $k \geq n$, then we have $(J_n(\lambda) - \lambda I)^k = 0$. Hence we can summarize this as

$$n((J_m(\lambda) - \lambda I_m)^r) = \min\{r, m\}.$$

Note that if $A = J_n(\lambda)$, then $\chi_A(t) = M_A(t) = (t - \lambda)^n$. So λ is the only eigenvalue of A . So $a_\lambda = c_\lambda = n$. We also know that $n(A - \lambda I) = n - r(A - \lambda I) = 1$. So $g_\lambda = 1$.

Recall that a general Jordan normal form is a block diagonal matrix of Jordan blocks. We have just studied individual Jordan blocks. Next, we want to look at some properties of block diagonal matrices in general. If A is the block diagonal matrix

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

then

$$\chi_A(t) = \prod_{i=1}^k \chi_{A_i}(t).$$

Moreover, if $\rho \in \mathbb{F}[t]$, then

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

Hence

$$M_A(t) = \text{lcm}(M_{A_1}(t), \dots, M_{A_k}(t)).$$

By the rank-nullity theorem, we have

$$n(\rho(A)) = \sum_{i=1}^k n(\rho(A_i)).$$

Thus if A is in Jordan normal form, we get the following:

- g_λ is the number of Jordan blocks in A with eigenvalue λ .
- a_λ is the sum of sizes of the Jordan blocks of A with eigenvalue λ .
- c_λ is the size of the largest Jordan block with eigenvalue λ .

We are now going to prove the uniqueness part of Jordan normal form theorem.

Theorem. Let $\alpha \in \text{End}(V)$, and A in Jordan normal form representing α . Then the number of Jordan blocks $J_n(\lambda)$ in A with $n \geq r$ is

$$n((\alpha - \lambda I)^r) - n((\alpha - \lambda I)^{r-1}).$$

This is clearly independent of the choice of basis. Also, given this information, we can figure out how many Jordan blocks of size exactly n by doing the right subtraction. Hence this tells us that Jordan normal forms are unique up to permutation of blocks.

Proof. We work blockwise for

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

We have previously computed

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

Hence we know

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

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 $r \geq 1$, we have

$$n((\alpha - \lambda I)^r) - n((\alpha - \lambda I)^{r-1}) = \text{number of Jordan blocks } J_n(\lambda) \text{ with } n \geq r. \quad \square$$

We can interpret this result as follows: if $r \leq m$, when we take an additional power of $J_m(\lambda) - \lambda I_m$, we get from $\begin{pmatrix} 0 & I_{m-r} \\ 0 & 0 \end{pmatrix}$ to $\begin{pmatrix} 0 & I_{m-r-1} \\ 0 & 0 \end{pmatrix}$. So we kill off one more column in the matrix, and the nullity increase by one. This happens until $(J_m(\lambda) - \lambda I_m)^r = 0$, in which case increasing the power no longer affects the matrix. So when we look at 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

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

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

This allows us to decompose V into a block diagonal matrix, and then each block will only have one eigenvalue.

Note that if $c_1 = \cdots = c_k = 1$, then we recover the diagonalizability theorem. Hence, it is not surprising that the proof of this is similar to the diagonalizability theorem. We will again prove this by constructing projection maps to each of the V_i .

Proof. Let

$$p_j(t) = \prod_{i \neq j} (t - \lambda_i)^{c_i}.$$

Then p_1, \dots, p_k have no common factors, i.e. they are coprime. Thus by Euclid's algorithm, there exists $q_1, \dots, q_k \in \mathbb{C}[t]$ such that

$$\sum p_i q_i = 1.$$

We now define the endomorphism

$$\pi_j = q_j(\alpha) p_j(\alpha)$$

for $j = 1, \dots, k$.

Then $\sum \pi_j = \iota$. Since $M_\alpha(\alpha) = 0$ and $M_\alpha(t) = (t - \lambda_j \iota)^{c_j} p_j(t)$, we get

$$(\alpha - \lambda_j \iota)^{c_j} \pi_j = 0.$$

So $\text{im } \pi_j \subseteq V_j$.

Now suppose $\mathbf{v} \in V$. Then

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

So

$$V = \sum V_j.$$

To show this is a direct sum, note that $\pi_i \pi_j = 0$, since the product contains $M_\alpha(\alpha)$ as a factor. So

$$\pi_i = \iota \pi_i = \left(\sum \pi_j \right) \pi_i = \pi_i^2.$$

So π is a projection, and $\pi_j|_{V_j} = \iota|_{V_j}$. So if $\mathbf{v} = \sum \mathbf{v}_i$, then applying π_i to both sides gives $\mathbf{v}_i = \pi_i(\mathbf{v})$. Hence there is a unique way of writing \mathbf{v} as a sum of things in V_i . So $V = \bigoplus V_j$ as claimed. \square

Note that we didn't really use the fact that the vector space is over \mathbb{C} , except to get that the minimal polynomial is a product of linear factors. In fact, for arbitrary vector spaces, if the minimal polynomial of a matrix is a product of linear factors, then it can be put into Jordan normal form. The converse is also true — if it can be put into Jordan normal form, then the minimal polynomial is a product of linear factors, since we've seen that a necessary and sufficient condition for the minimal polynomial to be a product of linear factors is for there to be a basis in which the matrix is upper triangular.

Using this theorem, by restricting α to its generalized eigenspaces, we can reduce the existence part of the Jordan normal form theorem to the case $M_\alpha(t) = (t - \lambda)^c$. Further by replacing α by $\alpha - \lambda \iota$, we can reduce this to the case where 0 is the only eigenvalue.

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

Over \mathbb{C} , α is nilpotent if and only if the only eigenvalue of α is 0. This is since α is nilpotent if the minimal polynomial is t^r for some r .

We've now reduced the problem of classifying complex endomorphisms to the classifying nilpotent endomorphisms. This is the point where we stop. For the remainder of the proof, see IB Groups, Rings and Modules. There is in fact an elementary proof of it, and we're not doing it not because it's hard, but because we don't have time.

Example. Let

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

We know we can find the Jordan normal form by just computing the minimal polynomial and characteristic polynomial. But we can do better and try to find a P such that $P^{-1}AP$ is in Jordan normal form.

We first compute the eigenvalues of A . The characteristic polynomial is

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

We now compute the eigenspaces of A . We have

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

We see this has rank 2 and hence nullity 1, and the eigenspace is the kernel

$$E_A(1) = \left\langle \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right\rangle$$

We can also compute the other eigenspace. We have

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

This has rank 2 and

$$E_A(2) = \left\langle \begin{pmatrix} 2 \\ 1 \\ 2 \end{pmatrix} \right\rangle.$$

Since

$$\dim E_A(1) + \dim E_A(2) = 2 < 3,$$

this is not diagonalizable. So the minimal polynomial must also be $M_A(t) = \chi_A(t) = (t-1)^2(t-2)$. From the classification last time, we know that A is similar to

$$\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 2 \end{pmatrix}$$

We now want to compute a basis that transforms A to this. We want a basis $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ of \mathbb{C}^3 such that

$$A\mathbf{v}_1 = \mathbf{v}_1, \quad A\mathbf{v}_2 = \mathbf{v}_1 + \mathbf{v}_2, \quad A\mathbf{v}_3 = 2\mathbf{v}_3.$$

Equivalently, we have

$$(A - I)\mathbf{v}_1 = \mathbf{0}, \quad (A - I)\mathbf{v}_2 = \mathbf{v}_1, \quad (A - 2I)\mathbf{v}_3 = \mathbf{0}.$$

There is an obvious choice for \mathbf{v}_3 , namely the eigenvector of eigenvalue 2.

To find \mathbf{v}_1 and \mathbf{v}_2 , the idea is to find some \mathbf{v}_2 such that $(A - I)\mathbf{v}_2 \neq \mathbf{0}$ but $(A - I)^2\mathbf{v}_2 = \mathbf{0}$. Then we can let $\mathbf{v}_1 = (A - I)\mathbf{v}_2$.

We can compute the kernel of $(A - I)^2$. We have

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

The kernel of this is

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

We need to pick our \mathbf{v}_2 that is in this kernel but not in the kernel of $A - I$ (which is the eigenspace E_1 we have computed above). So we have

$$\mathbf{v}_2 = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}, \quad \mathbf{v}_1 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \quad \mathbf{v}_3 = \begin{pmatrix} 2 \\ 1 \\ 2 \end{pmatrix}.$$

Hence we have

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

and

$$P^{-1}AP = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 2 \end{pmatrix}.$$

7 Bilinear forms II

In Chapter 4, we have looked at bilinear forms in general. Here, we want to look at bilinear forms on a single space, since often there is just one space we are interested in. We are also not looking into general bilinear forms on a single space, but just those that are symmetric.

7.1 Symmetric bilinear forms and quadratic forms

Definition (Symmetric bilinear form). Let V be a vector space over \mathbb{F} . A bilinear form $\phi : V \times V \rightarrow \mathbb{F}$ is *symmetric* if

$$\phi(\mathbf{v}, \mathbf{w}) = \phi(\mathbf{w}, \mathbf{v})$$

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

Example. If $S \in \text{Mat}_n(\mathbb{F})$ is a symmetric matrix, i.e. $S^T = S$, the bilinear form $\phi : \mathbb{F}^n \times \mathbb{F}^n \rightarrow \mathbb{F}$ defined by

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

is a symmetric bilinear form.

This example is typical in the following sense:

Lemma. Let V be a finite-dimensional vector space over \mathbb{F} , and $\phi : V \times V \rightarrow \mathbb{F}$ is a symmetric bilinear form. Let $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ be a basis for V , and let M be the matrix representing ϕ with respect to this basis, i.e. $M_{ij} = \phi(\mathbf{e}_i, \mathbf{e}_j)$. Then ϕ is symmetric if and only if M is symmetric.

Proof. If ϕ is symmetric, then

$$M_{ij} = \phi(\mathbf{e}_i, \mathbf{e}_j) = \phi(\mathbf{e}_j, \mathbf{e}_i) = M_{ji}.$$

So $M^T = M$. So M is symmetric.

If M is symmetric, then

$$\begin{aligned} \phi(\mathbf{x}, \mathbf{y}) &= \phi\left(\sum x_i \mathbf{e}_i, \sum y_j \mathbf{e}_j\right) \\ &= \sum_{i,j} x_i M_{ij} y_j \\ &= \sum_{i,j} y_j M_{ji} x_i \\ &= \phi(\mathbf{y}, \mathbf{x}). \end{aligned} \quad \square$$

We are going to see what happens when we change basis. As in the case of endomorphisms, we will require to change basis in the same ways on both sides.

Lemma. Let V be a finite-dimensional vector space, and $\phi : V \times V \rightarrow \mathbb{F}$ a bilinear form. Let $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ and $(\mathbf{f}_1, \dots, \mathbf{f}_n)$ be bases of V such that

$$\mathbf{f}_i = \sum_{k=1}^n P_{ki} \mathbf{e}_k.$$

If A represents ϕ with respect to (\mathbf{e}_i) and B represents ϕ with respect to (\mathbf{f}_i) , then

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

Proof. Special case of general formula proven before. \square

This motivates the following definition:

Definition (Congruent matrices). Two square matrices A, B are *congruent* if there exists some invertible P such that

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

It is easy to see that congruence is an equivalence relation. Two matrices are congruent if and only if represent the same bilinear form with respect to different bases.

Thus, to classify (symmetric) bilinear forms is the same as classifying (symmetric) matrices up to congruence.

Before we do the classification, we first look at quadratic forms, which are something derived from bilinear forms.

Definition (Quadratic form). A function $q : V \rightarrow \mathbb{F}$ is a *quadratic form* if there exists some bilinear form ϕ such that

$$q(\mathbf{v}) = \phi(\mathbf{v}, \mathbf{v})$$

for all $\mathbf{v} \in V$.

Note that quadratic forms are *not* linear maps (they are quadratic).

Example. Let $V = \mathbb{R}^2$ and ϕ be represented by A with respect to the standard basis. Then

$$q\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} x & y \end{pmatrix} \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = A_{11}x^2 + (A_{12} + A_{21})xy + A_{22}y^2.$$

Notice that if A is replaced the symmetric matrix

$$\frac{1}{2}(A + A^T),$$

then we get a different ϕ , but the same q . This is in fact true in general.

Proposition (Polarization identity). Suppose that $\text{char } \mathbb{F} \neq 2$, i.e. $1 + 1 \neq 0$ on \mathbb{F} (e.g. if \mathbb{F} is \mathbb{R} or \mathbb{C}). If $q : V \rightarrow \mathbb{F}$ is a quadratic form, then there exists a *unique* symmetric bilinear form $\phi : V \times V \rightarrow \mathbb{F}$ such that

$$q(\mathbf{v}) = \phi(\mathbf{v}, \mathbf{v}).$$

Proof. Let $\psi : V \times V \rightarrow \mathbb{F}$ be a bilinear form such that $\psi(\mathbf{v}, \mathbf{v}) = q(\mathbf{v})$. We define $\phi : V \times V \rightarrow \mathbb{F}$ by

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

for all $\mathbf{v}, \mathbf{w} \in \mathbb{F}$. This is clearly a bilinear form, and it is also clearly symmetric and satisfies the condition we wants. So we have proved the existence part.

To prove uniqueness, we want to find out the values of $\phi(\mathbf{v}, \mathbf{w})$ in terms of what q tells us. Suppose ϕ is such a symmetric bilinear form. We compute

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

So we have

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

So it is determined by q , and hence unique. \square

Theorem. Let V be a finite-dimensional vector space over \mathbb{F} , and $\phi : V \times V \rightarrow \mathbb{F}$ a symmetric bilinear form. Then there exists a basis $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ for V such that ϕ is represented by a diagonal matrix with respect to this basis.

This tells us classifying symmetric bilinear forms is easier than classifying endomorphisms, since for endomorphisms, even over \mathbb{C} , we cannot always make it diagonal, but we can for bilinear forms over arbitrary fields.

Proof. We induct over $n = \dim V$. The cases $n = 0$ and $n = 1$ are trivial, since all matrices are diagonal.

Suppose we have proven the result for all spaces of dimension less than n . First consider the case where $\phi(\mathbf{v}, \mathbf{v}) = 0$ for all $\mathbf{v} \in V$. We want to show that we must have $\phi = 0$. This follows from the polarization identity, since this ϕ induces the zero quadratic form, and we know that there is a unique bilinear form that induces the zero quadratic form. Since we know that the zero bilinear form also induces the zero quadratic form, we must have $\phi = 0$. Then ϕ will be represented by the zero matrix with respect to any basis, which is trivially diagonal.

If not, pick $\mathbf{e}_1 \in V$ such that $\phi(\mathbf{e}_1, \mathbf{e}_1) \neq 0$. Let

$$U = \ker \phi(\mathbf{e}_1, \cdot) = \{\mathbf{u} \in V : \phi(\mathbf{e}_1, \mathbf{u}) = 0\}.$$

Since $\phi(\mathbf{e}_1, \cdot) \in V^* \setminus \{0\}$, we know that $\dim U = n - 1$ by the rank-nullity theorem.

Our objective is to find other basis elements $\mathbf{e}_2, \dots, \mathbf{e}_n$ such that $\phi(\mathbf{e}_1, \mathbf{e}_j) = 0$ for all $j > 1$. For this to happen, we need to find them inside U .

Now consider $\phi|_{U \times U} : U \times U \rightarrow \mathbb{F}$, a symmetric bilinear form. By the induction hypothesis, we can find a basis $\mathbf{e}_2, \dots, \mathbf{e}_n$ for U such that $\phi|_{U \times U}$ is represented by a diagonal matrix with respect to this basis.

Now by construction, $\phi(\mathbf{e}_i, \mathbf{e}_j) = 0$ for all $1 \leq i \neq j \leq n$ and $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ is a basis for V . So we're done. \square

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

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

We want to find a basis $\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3$ for \mathbb{R}^3 such that q is of the form

$$q(a\mathbf{f}_1 + b\mathbf{f}_2 + c\mathbf{f}_3) = \lambda a^2 + \mu b^2 + \nu c^2$$

for some $\lambda, \mu, \nu \in \mathbb{R}$.

There are two ways to do this. The first way is to follow the proof we just had. We first find our symmetric bilinear form. It is the bilinear form represented by the matrix

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

We then find \mathbf{f}_1 such that $\phi(\mathbf{f}_1, \mathbf{f}_1) \neq 0$. We note that $q(\mathbf{e}_1) = 1 \neq 0$. So we pick

$$\mathbf{f}_1 = \mathbf{e}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}.$$

Then

$$\phi(\mathbf{e}_1, \mathbf{v}) = (1 \ 0 \ 0) \begin{pmatrix} 1 & 1 & 3 \\ 1 & 1 & 2 \\ 3 & 2 & 1 \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} = v_1 + v_2 + 3v_3.$$

Next we need to pick our \mathbf{f}_2 . Since it is in the kernel of $\phi(\mathbf{f}_1, \cdot)$, it must satisfy

$$\phi(\mathbf{f}_1, \mathbf{f}_2) = 0.$$

To continue our proof inductively, we also have to pick an \mathbf{f}_2 such that

$$\phi(\mathbf{f}_2, \mathbf{f}_2) \neq 0.$$

For example, we can pick

$$\mathbf{f}_2 = \begin{pmatrix} 3 \\ 0 \\ -1 \end{pmatrix}.$$

Then we have $q(\mathbf{f}_2) = -8$.

Then we have

$$\phi(\mathbf{f}_2, \mathbf{v}) = (3 \ 0 \ -1) \begin{pmatrix} 1 & 1 & 3 \\ 1 & 1 & 2 \\ 3 & 2 & 1 \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} = v_2 + 8v_3$$

Finally, we want $\phi(\mathbf{f}_1, \mathbf{f}_3) = \phi(\mathbf{f}_2, \mathbf{f}_3) = 0$. Then

$$\mathbf{f}_3 = \begin{pmatrix} 5 \\ -8 \\ 1 \end{pmatrix}$$

works. We have $q(\mathbf{f}_3) = 8$.

With these basis elements, we have

$$\begin{aligned} q(a\mathbf{f}_1 + b\mathbf{f}_2 + c\mathbf{f}_3) &= \phi(a\mathbf{f}_1 + b\mathbf{f}_2 + c\mathbf{f}_3, a\mathbf{f}_1 + b\mathbf{f}_2 + c\mathbf{f}_3) \\ &= a^2q(\mathbf{f}_1) + b^2q(\mathbf{f}_2) + c^2q(\mathbf{f}_3) \\ &= a^2 - 8b^2 + 8c^2. \end{aligned}$$

Alternatively, we can solve the problem by completing the square. We have

$$\begin{aligned}x^2 + y^2 + z^2 + 2xy + 4yz + 6xz &= (x + y + 3z)^2 - 2yz - 8z^2 \\ &= (x + y + 3z)^2 - 8\left(z + \frac{y}{8}\right)^2 + \frac{1}{8}y^2.\end{aligned}$$

We now see

$$\phi\left(\begin{pmatrix} x \\ y \\ z \end{pmatrix}, \begin{pmatrix} x' \\ y' \\ z' \end{pmatrix}\right) = (x + y + 3z)(x' + y' + 3z') - 8\left(z + \frac{y}{8}\right)\left(z' + \frac{y'}{8}\right) + \frac{1}{8}yy'.$$

Why do we know this? This is clearly a symmetric bilinear form, and this also clearly induces the q given above. By uniqueness, we know that this is the right symmetric bilinear form.

We now use this form to find our $\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3$ such that $\phi(\mathbf{f}_i, \mathbf{f}_j) = \delta_{ij}$.

To do so, we just solve the equations

$$\begin{aligned}x + y + 3z &= 1 \\ z + \frac{1}{8}y &= 0 \\ y &= 0.\end{aligned}$$

This gives our first vector as

$$\mathbf{f}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}.$$

We then solve

$$\begin{aligned}x + y + 3z &= 0 \\ z + \frac{1}{8}y &= 1 \\ y &= 0.\end{aligned}$$

So we have

$$\mathbf{f}_2 = \begin{pmatrix} -3 \\ 0 \\ 1 \end{pmatrix}.$$

Finally, we solve

$$\begin{aligned}x + y + 3z &= 0 \\ z + \frac{1}{8}y &= 0 \\ y &= 1.\end{aligned}$$

This gives

$$\mathbf{f}_3 = \begin{pmatrix} -5/8 \\ 1 \\ -1/8 \end{pmatrix}.$$

Then we can see that the result follows, and we get

$$q(a\mathbf{f}_1 + b\mathbf{f}_2 + c\mathbf{f}_3) = a^2 - 8b^2 + \frac{1}{8}c^2.$$

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 $(\mathbf{v}_1, \dots, \mathbf{v}_m)$ 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 $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ 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 $\lambda_1, \dots, \lambda_r \neq 0$ and $\lambda_{r+1}, \dots, \lambda_n = 0$.

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

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

Then

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

So done. □

Note that it follows that for the corresponding quadratic form q , we have

$$q\left(\sum_{i=1}^n a_i \mathbf{v}_i\right) = \sum_{i=1}^r a_i^2.$$

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

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

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 $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ for V such that ϕ is represented

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

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

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

Note that we have seen these things in special relativity, where the Minkowski inner product is given by the symmetric bilinear form represented by

$$\begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

in units where $c = 1$.

Proof. We've already shown that there exists a basis $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ such that $\phi(\mathbf{e}_i, \mathbf{e}_j) = \lambda_i \delta_{ij}$ for some $\lambda_1, \dots, \lambda_n \in \mathbb{R}$. By reordering, we may assume

$$\begin{cases} \lambda_i > 0 & 1 \leq i \leq p \\ \lambda_i < 0 & p+1 \leq i \leq r \\ \lambda_i = 0 & i > r \end{cases}$$

We let μ_i be defined by

$$\mu_i = \begin{cases} \sqrt{\lambda_i} & 1 \leq i \leq p \\ \sqrt{-\lambda_i} & p+1 \leq i \leq r \\ 1 & i > r \end{cases}$$

Defining

$$\mathbf{v}_i = \frac{1}{\mu_i} \mathbf{e}_i,$$

we find that ϕ is indeed represented by

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

We will later show that this form is indeed unique. Before that, we will have a few definitions, that really only make sense over \mathbb{R} .

Definition (Positive/negative (semi-)definite). Let ϕ be a symmetric bilinear form on a finite-dimensional real vector space V . We say

- (i) ϕ is *positive definite* if $\phi(\mathbf{v}, \mathbf{v}) > 0$ for all $\mathbf{v} \in V \setminus \{0\}$.
- (ii) ϕ is *positive semi-definite* if $\phi(\mathbf{v}, \mathbf{v}) \geq 0$ for all $\mathbf{v} \in V$.
- (iii) ϕ is *negative definite* if $\phi(\mathbf{v}, \mathbf{v}) < 0$ for all $\mathbf{v} \in V \setminus \{0\}$.
- (iv) ϕ is *negative semi-definite* if $\phi(\mathbf{v}, \mathbf{v}) \leq 0$ for all $\mathbf{v} \in V$.

We are going to use these notions to prove uniqueness. It is easy to see that if $p = 0$ and $q = n$, then we are negative definite; if $p = 0$ and $q \neq n$, then we are negative semi-definite etc.

Example. Let ϕ be a symmetric bilinear form on \mathbb{R}^n represented by

$$\begin{pmatrix} I_p & 0 \\ 0 & 0_{n-p} \end{pmatrix}.$$

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

If instead ϕ is represented by

$$\begin{pmatrix} -I_p & 0 \\ 0 & 0_{n-p} \end{pmatrix},$$

then ϕ is negative semi-definite. ϕ is negative definite precisely if $n = q$.

We are going to use this to prove the uniqueness part of our previous theorem.

Theorem (Sylvester's law of inertia). Let ϕ be a symmetric bilinear form on a finite-dimensional real vector space V . Then there exists unique non-negative integers p, q such that ϕ is represented by

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

with respect to some basis.

Proof. We have already proved the existence part, and we just have to prove uniqueness. To do so, we characterize p and q in a basis-independent way. We already know that $p + q = r(\phi)$ does not depend on the basis. So it suffices to show p is unique.

To see that p is unique, we show that p is the largest dimension of a subspace $P \subseteq V$ such that $\phi|_{P \times P}$ is positive definite.

First we show we can find such a P . Suppose ϕ is represented by

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

with respect to $(\mathbf{e}_1, \dots, \mathbf{e}_n)$. Then ϕ restricted to $\langle \mathbf{e}_1, \dots, \mathbf{e}_p \rangle$ is represented by I_p with respect to $\mathbf{e}_1, \dots, \mathbf{e}_p$. So ϕ restricted to this is positive definite.

Now suppose P is any subspace of V such that $\phi|_{P \times P}$ is positive definite. To show P has dimension at most p , we find a subspace complementary to P with dimension $n - p$.

Let $Q = \langle \mathbf{e}_{p+1}, \dots, \mathbf{e}_n \rangle$. Then ϕ restricted to $Q \times Q$ is represented by

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

Now if $\mathbf{v} \in P \cap Q \setminus \{0\}$, then $\phi(\mathbf{v}, \mathbf{v}) > 0$ since $\mathbf{v} \in P \setminus \{0\}$ and $\phi(\mathbf{v}, \mathbf{v}) \leq 0$ since $\mathbf{v} \in Q$, which is a contradiction. So $P \cap Q = 0$.

We have

$$\dim V \geq \dim(P + Q) = \dim P + \dim Q = \dim P + (n - p).$$

Rearranging gives

$$\dim P \leq p.$$

A similar argument shows that q is the maximal dimension of a subspace $Q \subseteq V$ such that $\phi|_{Q \times Q}$ is negative definite. \square

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

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

7.2 Hermitian 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 $\phi : V \times W \rightarrow \mathbb{C}$ such that

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

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

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 definition, where we have linearity in the first argument and conjugate linearity in the second.

These are called sesquilinear since “sesqui” means “one and a half”, and this is linear in the second argument and “half linear” in the first.

Alternatively, to define a sesquilinear form, we can define a new complex vector space \bar{V} structure on V by taking the same abelian group (i.e. the same underlying set and addition), but with the scalar multiplication $\mathbb{C} \times \bar{V} \rightarrow \bar{V}$ defined as

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

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

Definition (Representation of sesquilinear form). Let V, W be finite-dimensional complex vector spaces with basis $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ and $(\mathbf{w}_1, \dots, \mathbf{w}_m)$ respectively, and $\phi : V \times W \rightarrow \mathbb{C}$ 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 $\bar{V} \times W \rightarrow \mathbb{C}$. Let $\mathbf{v} = \sum \lambda_i \mathbf{v}_i$ and $\mathbf{w} = \sum \mu_j \mathbf{w}_j$. Then we have

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

We now want the right definition of symmetric sesquilinear form. We cannot just require $\phi(\mathbf{v}, \mathbf{w}) = \phi(\mathbf{w}, \mathbf{v})$, since ϕ is linear in the second variable and conjugate linear on the first variable. So in particular, if $\phi(\mathbf{v}, \mathbf{w}) \neq 0$, we have $\phi(i\mathbf{v}, \mathbf{w}) \neq \phi(\mathbf{v}, i\mathbf{w})$.

Definition (Hermitian sesquilinear form). A sesquilinear form on $V \times V$ is *Hermitian* if

$$\phi(\mathbf{v}, \mathbf{w}) = \overline{\phi(\mathbf{w}, \mathbf{v})}.$$

Note that if ϕ is Hermitian, then $\phi(\mathbf{v}, \mathbf{v}) = \overline{\phi(\mathbf{v}, \mathbf{v})} \in \mathbb{R}$ for any $\mathbf{v} \in V$. So it makes sense to ask if it is positive or negative. Moreover, for any complex number λ , we have

$$\phi(\lambda\mathbf{v}, \lambda\mathbf{v}) = |\lambda|^2 \phi(\mathbf{v}, \mathbf{v}).$$

So multiplying by a scalar does not change the sign. So it makes sense to talk about positive (semi-)definite and negative (semi-)definite Hermitian forms.

We will prove results analogous to what we had for real symmetric bilinear forms.

Lemma. Let $\phi : V \times V \rightarrow \mathbb{C}$ be a sesquilinear form on a finite-dimensional vector space over \mathbb{C} , and $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ a basis for V . Then ϕ is Hermitian if and only if the matrix A representing ϕ is Hermitian (i.e. $A = A^\dagger$).

Proof. If ϕ is Hermitian, then

$$A_{ij} = \phi(\mathbf{e}_i, \mathbf{e}_j) = \overline{\phi(\mathbf{e}_j, \mathbf{e}_i)} = A_{ij}^\dagger.$$

If A is Hermitian, then

$$\phi\left(\sum \lambda_i \mathbf{e}_i, \sum \mu_j \mathbf{e}_j\right) = \lambda^\dagger A \mu = \overline{\mu^\dagger A^\dagger \lambda} = \overline{\phi\left(\sum \mu_j \mathbf{e}_j, \sum \lambda_i \mathbf{e}_i\right)}.$$

So done. □

Proposition (Change of basis). Let ϕ be a Hermitian form on a finite dimensional vector space V ; $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ and $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ are bases for V such that

$$\mathbf{v}_i = \sum_{k=1}^n P_{ki} \mathbf{e}_k;$$

and A, B represent ϕ with respect to $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ and $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ respectively. Then

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

Proof. We have

$$\begin{aligned} B_{ij} &= \phi(\mathbf{v}_i, \mathbf{v}_j) \\ &= \phi\left(\sum P_{ki} \mathbf{e}_k, \sum P_{\ell j} \mathbf{e}_\ell\right) \\ &= \sum_{k, \ell=1}^n \bar{P}_{ki} P_{\ell j} A_{k\ell} \\ &= (P^\dagger A P)_{ij}. \end{aligned} \quad \square$$

Lemma (Polarization identity (again)). A Hermitian form ϕ on V is determined by the function $\psi : \mathbf{v} \mapsto \phi(\mathbf{v}, \mathbf{v})$.

The proof this time is slightly more involved.

Proof. We have the following:

$$\begin{aligned}\psi(\mathbf{x} + \mathbf{y}) &= \phi(\mathbf{x}, \mathbf{x}) + \phi(\mathbf{x}, \mathbf{y}) + \phi(\mathbf{y}, \mathbf{x}) + \phi(\mathbf{y}, \mathbf{y}) \\ -\psi(\mathbf{x} - \mathbf{y}) &= -\phi(\mathbf{x}, \mathbf{x}) + \phi(\mathbf{x}, \mathbf{y}) + \phi(\mathbf{y}, \mathbf{x}) - \phi(\mathbf{y}, \mathbf{y}) \\ i\psi(\mathbf{x} - i\mathbf{y}) &= i\phi(\mathbf{x}, \mathbf{x}) + \phi(\mathbf{x}, \mathbf{y}) - \phi(\mathbf{y}, \mathbf{x}) + i\phi(\mathbf{y}, \mathbf{y}) \\ -i\psi(\mathbf{x} + i\mathbf{y}) &= -i\phi(\mathbf{x}, \mathbf{x}) + \phi(\mathbf{x}, \mathbf{y}) - \phi(\mathbf{y}, \mathbf{x}) - i\phi(\mathbf{y}, \mathbf{y})\end{aligned}$$

So

$$\phi(\mathbf{x}, \mathbf{y}) = \frac{1}{4}(\psi(\mathbf{x} + \mathbf{y}) - \psi(\mathbf{x} - \mathbf{y}) + i\psi(\mathbf{x} - i\mathbf{y}) - i\psi(\mathbf{x} + i\mathbf{y})). \quad \square$$

Theorem (Hermitian form of Sylvester's law of inertia). Let V be a finite-dimensional complex vector space and ϕ a hermitian form on V . Then there exists unique non-negative integers p and q such that ϕ is represented by

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

with respect to some basis.

Proof. Same as for symmetric forms over \mathbb{R} . □

8 Inner product spaces

Welcome to the last chapter, where we discuss inner products. Technically, an inner product is just a special case of a positive-definite symmetric bilinear or hermitian form. However, they are usually much more interesting and useful. Many familiar notions such as orthogonality only make sense when we have an inner product.

In this chapter, we will adopt the convention that \mathbb{F} always means either \mathbb{R} or \mathbb{C} , since working with other fields doesn't make much sense here.

8.1 Definitions and basic properties

Definition (Inner product space). Let V be a vector space. An *inner product* on V is a positive-definite symmetric bilinear/hermitian form. We usually write (x, y) instead of $\phi(x, y)$.

A vector space equipped with an inner product is an *inner product space*.

We will see that if we have an inner product, then we can define lengths and distances in a sensible way.

Example.

- (i) \mathbb{R}^n or \mathbb{C}^n with the usual inner product

$$(x, y) = \sum_{i=1}^n \bar{x}_i y_i$$

forms an inner product space.

In some sense, this is the only inner product on finite-dimensional spaces, by Sylvester's law of inertia. However, we would not like to think so, and instead work with general inner products.

- (ii) Let $C([0, 1], \mathbb{F})$ be the vector space of real/complex valued continuous functions on $[0, 1]$. Then the following is an inner product:

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

- (iii) More generally, for any $w : [0, 1] \rightarrow \mathbb{R}^+$ continuous, we can define the inner product on $C([0, 1], \mathbb{F})$ as

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

If V is an inner product space, we can define a norm on V by

$$\|\mathbf{v}\| = \sqrt{(\mathbf{v}, \mathbf{v})}.$$

This is just the usual notion of norm on \mathbb{R}^n and \mathbb{C}^n . This gives the notion of length in inner product spaces. Note that $\|\mathbf{v}\| > 0$ with equality if and only if $\mathbf{v} = \mathbf{0}$.

Note also that the norm $\|\cdot\|$ determines the inner product by the polarization identity.

We want to see that this indeed satisfies the definition of a norm, as you might have seen from Analysis II. To prove this, we need to prove the Cauchy-Schwarz inequality.

Theorem (Cauchy-Schwarz inequality). Let V be an inner product space and $\mathbf{v}, \mathbf{w} \in V$. Then

$$|(\mathbf{v}, \mathbf{w})| \leq \|\mathbf{v}\| \|\mathbf{w}\|.$$

Proof. If $\mathbf{w} = 0$, then this is trivial. Otherwise, since the norm is positive definite, for any λ , we get

$$0 \leq (\mathbf{v} - \lambda\mathbf{w}, \mathbf{v} - \lambda\mathbf{w}) = (\mathbf{v}, \mathbf{v}) - \bar{\lambda}(\mathbf{w}, \mathbf{v}) - \lambda(\mathbf{v}, \mathbf{w}) + |\lambda|^2(\mathbf{w}, \mathbf{w}).$$

We now pick a clever value of λ . We let

$$\lambda = \frac{(\mathbf{w}, \mathbf{v})}{(\mathbf{w}, \mathbf{w})}.$$

Then we get

$$0 \leq (\mathbf{v}, \mathbf{v}) - \frac{|(\mathbf{w}, \mathbf{v})|^2}{(\mathbf{w}, \mathbf{w})} - \frac{|(\mathbf{w}, \mathbf{v})|^2}{(\mathbf{w}, \mathbf{w})} + \frac{|(\mathbf{w}, \mathbf{v})|^2}{(\mathbf{w}, \mathbf{w})}.$$

So we get

$$|(\mathbf{w}, \mathbf{v})|^2 \leq (\mathbf{v}, \mathbf{v})(\mathbf{w}, \mathbf{w}).$$

So done. □

With this, we can prove the triangle inequality.

Corollary (Triangle inequality). Let V be an inner product space and $\mathbf{v}, \mathbf{w} \in V$. Then

$$\|\mathbf{v} + \mathbf{w}\| \leq \|\mathbf{v}\| + \|\mathbf{w}\|.$$

Proof. We compute

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

So done. □

The next thing we do is to define orthogonality. This generalizes the notion of being “perpendicular”.

Definition (Orthogonal vectors). Let V be an inner product space. Then $\mathbf{v}, \mathbf{w} \in V$ are *orthogonal* if $(\mathbf{v}, \mathbf{w}) = 0$.

Definition (Orthonormal set). Let V be an inner product space. A set $\{\mathbf{v}_i : i \in I\}$ is an *orthonormal set* if for any $i, j \in I$, we have

$$(\mathbf{v}_i, \mathbf{v}_j) = \delta_{ij}$$

It should be clear that an orthonormal set must be linearly independent.

Definition (Orthonormal basis). Let V be an inner product space. A subset of V is an *orthonormal basis* if it is an orthonormal set and is a basis.

In an inner product space, we almost always want orthonormal basis only. If we pick a basis, we should pick an orthonormal one.

However, we do not know there is always an orthonormal basis, even in the finite-dimensional case. Also, given an orthonormal set, we would like to extend it to an orthonormal basis. This is what we will do later.

Before that, we first note that given an orthonormal basis, it is easy to find the coordinates of any vector in this basis. Suppose V is a finite-dimensional inner product space with an orthonormal basis $\mathbf{v}_1, \dots, \mathbf{v}_n$. Given

$$\mathbf{v} = \sum_{i=1}^n \lambda_i \mathbf{v}_i,$$

we have

$$(\mathbf{v}_j, \mathbf{v}) = \sum_{i=1}^n \lambda_i (\mathbf{v}_j, \mathbf{v}_i) = \lambda_j.$$

So $\mathbf{v} \in V$ can always be written as

$$\sum_{i=1}^n (\mathbf{v}_i, \mathbf{v}) \mathbf{v}_i.$$

Lemma (Parseval's identity). Let V be a finite-dimensional inner product space with an orthonormal basis $\mathbf{v}_1, \dots, \mathbf{v}_n$, and $\mathbf{v}, \mathbf{w} \in V$. Then

$$(\mathbf{v}, \mathbf{w}) = \sum_{i=1}^n \overline{(\mathbf{v}_i, \mathbf{v})} (\mathbf{v}_i, \mathbf{w}).$$

In particular,

$$\|\mathbf{v}\|^2 = \sum_{i=1}^n |(\mathbf{v}_i, \mathbf{v})|^2.$$

This is something we've seen in IB Methods, for infinite dimensional spaces. However, we will only care about finite-dimensional ones now.

Proof.

$$\begin{aligned} (\mathbf{v}, \mathbf{w}) &= \left(\sum_{i=1}^n (\mathbf{v}_i, \mathbf{v}) \mathbf{v}_i, \sum_{j=1}^n (\mathbf{v}_j, \mathbf{w}) \mathbf{v}_j \right) \\ &= \sum_{i,j=1}^n \overline{(\mathbf{v}_i, \mathbf{v})} (\mathbf{v}_j, \mathbf{w}) (\mathbf{v}_i, \mathbf{v}_j) \\ &= \sum_{i,j=1}^n \overline{(\mathbf{v}_i, \mathbf{v})} (\mathbf{v}_j, \mathbf{w}) \delta_{ij} \\ &= \sum_{i=1}^n \overline{(\mathbf{v}_i, \mathbf{v})} (\mathbf{v}_i, \mathbf{w}). \end{aligned}$$

□

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 $\mathbf{e}_1, \mathbf{e}_2, \dots$ a linearly independent set. Then we can construct an orthonormal set $\mathbf{v}_1, \mathbf{v}_2, \dots$ with the property that

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

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 $\mathbf{v}_1, \dots, \mathbf{v}_k$ that satisfies the properties. We define

$$\mathbf{u}_{k+1} = \mathbf{e}_{k+1} - \sum_{i=1}^k \langle \mathbf{v}_i, \mathbf{e}_{k+1} \rangle \mathbf{v}_i.$$

We want to prove that this is orthogonal to all the other \mathbf{v}_i 's for $i \leq k$. We have

$$\langle \mathbf{v}_j, \mathbf{u}_{k+1} \rangle = \langle \mathbf{v}_j, \mathbf{e}_{k+1} \rangle - \sum_{i=1}^k \langle \mathbf{v}_i, \mathbf{e}_{k+1} \rangle \delta_{ij} = \langle \mathbf{v}_j, \mathbf{e}_{k+1} \rangle - \langle \mathbf{v}_j, \mathbf{e}_{k+1} \rangle = 0.$$

So it is orthogonal.

We want to argue that \mathbf{u}_{k+1} is non-zero. Note that

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

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

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

by assumption. We know $\langle \mathbf{e}_1, \dots, \mathbf{e}_k, \mathbf{e}_{k+1} \rangle$ has dimension $k+1$ since the \mathbf{e}_i are linearly independent. So we must have \mathbf{u}_{k+1} non-zero, or else $\langle \mathbf{v}_1, \dots, \mathbf{v}_k \rangle$ 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 $\mathbf{v}_1, \dots, \mathbf{v}_{k+1}$ is orthonormal and $\langle \mathbf{v}_1, \dots, \mathbf{v}_{k+1} \rangle = \langle \mathbf{e}_1, \dots, \mathbf{e}_{k+1} \rangle$ as required. \square

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

Proof. Let $\mathbf{v}_1, \dots, \mathbf{v}_k$ be an orthonormal set. Since this is linearly independent, we can extend it to a basis $(\mathbf{v}_1, \dots, \mathbf{v}_k, \mathbf{x}_{k+1}, \dots, \mathbf{x}_n)$.

We now apply the Gram-Schmidt process to this basis to get an orthonormal basis of V , say $(\mathbf{u}_1, \dots, \mathbf{u}_n)$. Moreover, we can check that the process does not modify our $\mathbf{v}_1, \dots, \mathbf{v}_k$, i.e. $\mathbf{u}_i = \mathbf{v}_i$ for $1 \leq i \leq k$. So done. \square

Definition (Orthogonal internal direct sum). Let V be an inner product space and $V_1, V_2 \leq V$. Then V is the *orthogonal internal direct sum* of V_1 and V_2 if it is a direct sum and V_1 and V_2 are orthogonal. More precisely, we require

- (i) $V = V_1 + V_2$
- (ii) $V_1 \cap V_2 = 0$
- (iii) $(\mathbf{v}_1, \mathbf{v}_2) = 0$ for all $\mathbf{v}_1 \in V_1$ and $\mathbf{v}_2 \in V_2$.

Note that condition (iii) implies (ii), but we write it for the sake of explicitness.

We write $V = V_1 \perp V_2$.

Definition (Orthogonal complement). If $W \leq V$ is a subspace of an inner product space V , then the *orthogonal complement* of W in V is the subspace

$$W^\perp = \{\mathbf{v} \in V : (\mathbf{v}, \mathbf{w}) = 0, \forall \mathbf{w} \in W\}.$$

It is true that the orthogonal complement is a complement and orthogonal, i.e. V is the orthogonal direct sum of W and W^\perp .

Proposition. Let V be a finite-dimensional inner product space, and $W \leq V$. Then

$$V = W \perp W^\perp.$$

Proof. There are three things to prove, and we know (iii) implies (ii). Also, (iii) is obvious by definition of W^\perp . So it remains to prove (i), i.e. $V = W + W^\perp$.

Let $\mathbf{w}_1, \dots, \mathbf{w}_k$ be an orthonormal basis for W , and pick $\mathbf{v} \in V$. Now let

$$\mathbf{w} = \sum_{i=1}^k (\mathbf{w}_i, \mathbf{v}) \mathbf{w}_i.$$

Clearly, we have $\mathbf{w} \in W$. So we need to show $\mathbf{v} - \mathbf{w} \in W^\perp$. For each j , we can compute

$$\begin{aligned} (\mathbf{w}_j, \mathbf{v} - \mathbf{w}) &= (\mathbf{w}_j, \mathbf{v}) - \sum_{i=1}^k (\mathbf{w}_i, \mathbf{v})(\mathbf{w}_j, \mathbf{w}_i) \\ &= (\mathbf{w}_j, \mathbf{v}) - \sum_{i=1}^k (\mathbf{w}_i, \mathbf{v}) \delta_{ij} \\ &= 0. \end{aligned}$$

Hence for any λ_j , we have

$$\left(\sum \lambda_j \mathbf{w}_j, \mathbf{v} - \mathbf{w} \right) = 0.$$

So we have $\mathbf{v} - \mathbf{w} \in W^\perp$. So done. \square

Definition (Orthogonal external direct sum). Let V_1, V_2 be inner product spaces. The *orthogonal external direct sum* of V_1 and V_2 is the vector space $V_1 \oplus V_2$ with the inner product defined by

$$(\mathbf{v}_1 + \mathbf{v}_2, \mathbf{w}_1 + \mathbf{w}_2) = (\mathbf{v}_1, \mathbf{w}_1) + (\mathbf{v}_2, \mathbf{w}_2),$$

with $\mathbf{v}_1, \mathbf{w}_1 \in V_1, \mathbf{v}_2, \mathbf{w}_2 \in V_2$.

Here we write $\mathbf{v}_1 + \mathbf{v}_2 \in V_1 \oplus V_2$ instead of $(\mathbf{v}_1, \mathbf{v}_2)$ to avoid confusion.

This external direct sum is equivalent to the internal direct sum of $\{(\mathbf{v}_1, \mathbf{0}) : \mathbf{v}_1 \in V_1\}$ and $\{(\mathbf{0}, \mathbf{v}_2) : \mathbf{v}_2 \in V_2\}$.

Proposition. Let V be a finite-dimensional inner product space and $W \leq V$. Let $(\mathbf{e}_1, \dots, \mathbf{e}_k)$ be an orthonormal basis of W . Let π be the orthonormal projection of V onto W , i.e. $\pi : V \rightarrow W$ is a function that satisfies $\ker \pi = W^\perp$, $\pi|_W = \text{id}$. Then


(i) π is given by the formula

$$\pi(\mathbf{v}) = \sum_{i=1}^k (\mathbf{e}_i, \mathbf{v}) \mathbf{e}_i.$$

(ii) For all $\mathbf{v} \in V, \mathbf{w} \in W$, we have

$$\|\mathbf{v} - \pi(\mathbf{v})\| \leq \|\mathbf{v} - \mathbf{w}\|,$$

with equality if and only if $\pi(\mathbf{v}) = \mathbf{w}$. This says $\pi(\mathbf{v})$ is the point on W that is closest to \mathbf{v} .



Proof.

(i) Let $\mathbf{v} \in V$, and define

$$\mathbf{w} = \sum_{i=1}^k (\mathbf{e}_i, \mathbf{v}) \mathbf{e}_i.$$

We want to show this is $\pi(\mathbf{v})$. We need to show $\mathbf{v} - \mathbf{w} \in W^\perp$. We can compute

$$(\mathbf{e}_j, \mathbf{v} - \mathbf{w}) = (\mathbf{e}_j, \mathbf{v}) - \sum_{i=1}^k (\mathbf{e}_i, \mathbf{v}) (\mathbf{e}_j, \mathbf{e}_i) = 0.$$

So $\mathbf{v} - \mathbf{w}$ is orthogonal to every basis vector in \mathbf{w} , i.e. $\mathbf{v} - \mathbf{w} \in W^\perp$. So

$$\pi(\mathbf{v}) = \pi(\mathbf{w}) + \pi(\mathbf{v} - \mathbf{w}) = \mathbf{w}$$

as required.

(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

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

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

8.3 Adjoints, orthogonal and unitary maps

Adjoints

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

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

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 conceptual reason of what this means.

Let $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ and $(\mathbf{w}_1, \dots, \mathbf{w}_m)$ be orthonormal basis for V and W . Suppose α is represented by A .

To show uniqueness, suppose $\alpha^* : W \rightarrow V$ satisfies $(\alpha\mathbf{v}, \mathbf{w}) = (\mathbf{v}, \alpha^*\mathbf{w})$ for all $\mathbf{v} \in V$, $\mathbf{w} \in W$, then for all i, j , by definition, we know

$$\begin{aligned}(\mathbf{v}_i, \alpha^*(\mathbf{w}_j)) &= (\alpha(\mathbf{v}_i), \mathbf{w}_j) \\ &= \left(\sum_k A_{ki} \mathbf{w}_k, \mathbf{w}_j \right) \\ &= \sum_k \bar{A}_{ki} (\mathbf{w}_k, \mathbf{w}_j) = \bar{A}_{ji}.\end{aligned}$$

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

$$\begin{aligned}\left(\alpha \left(\sum \lambda_i \mathbf{v}_i \right), \sum \mu_j \mathbf{w}_j \right) &= \sum_{i,j} \bar{\lambda}_i \mu_j (\alpha(\mathbf{v}_i), \mathbf{w}_j) \\ &= \sum_{i,j} \bar{\lambda}_i \mu_j \left(\sum_k A_{ki} \mathbf{w}_k, \mathbf{w}_j \right) \\ &= \sum_{i,j} \bar{\lambda}_i \bar{A}_{ji} \mu_j.\end{aligned}$$

We can compute the other side and get

$$\begin{aligned} \left(\sum \lambda_i \mathbf{v}_i, \alpha^* \left(\sum \mu_j \mathbf{w}_j \right) \right) &= \sum_{i,j} \bar{\lambda}_i \mu_j \left(\mathbf{v}_i, \sum_k A_{kj}^\dagger \mathbf{v}_k \right) \\ &= \sum_{i,j} \bar{\lambda}_i \bar{A}_{ji} \mu_j. \end{aligned}$$

So done. □

What does this mean, conceptually? Note that the inner product V defines an isomorphism $V \rightarrow \bar{V}^*$ by $\mathbf{v} \mapsto (\cdot, \mathbf{v})$. Similarly, we have an isomorphism $W \rightarrow \bar{W}^*$. We can then put them in the following diagram:

$$\begin{array}{ccc} V & \xrightarrow{\alpha} & W \\ \downarrow \cong & & \downarrow \cong \\ \bar{V}^* & \xleftarrow{\alpha^*} & \bar{W}^* \end{array}$$

Then α^* is what fills in the dashed arrow. So α^* is in some sense the “dual” of the map α .

Definition (Adjoint). We call the map α^* the *adjoint* of α .

We have just seen that if α is represented by A with respect to some orthonormal bases, then α^* is represented by A^\dagger .

Definition (Self-adjoint). Let V be an inner product space, and $\alpha \in \text{End}(V)$. Then α is *self-adjoint* if $\alpha = \alpha^*$, i.e.

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

for all \mathbf{v}, \mathbf{w} .

Thus if $V = \mathbb{R}^n$ with the usual inner product, then $A \in \text{Mat}_n(\mathbb{R})$ is self-adjoint if and only if it is symmetric, i.e. $A = A^T$. If $V = \mathbb{C}^n$ with the usual inner product, then $A \in \text{Mat}_n(\mathbb{C})$ is self-adjoint if and only if A is Hermitian, i.e. $A = A^\dagger$.

Self-adjoint endomorphisms are important, as you may have noticed from IB Quantum Mechanics. We will later see that these have real eigenvalues with an orthonormal basis of eigenvectors.

Orthogonal maps

Another important class of endomorphisms is those that preserve lengths. We will first do this for real vector spaces, since the real and complex versions have different names.

Definition (Orthogonal endomorphism). Let V be a real inner product space. Then $\alpha \in \text{End}(V)$ is *orthogonal* if

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

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

By the polarization identity, α is orthogonal if and only if $\|\alpha(\mathbf{v})\| = \|\mathbf{v}\|$ for all $\mathbf{v} \in V$.

A real square matrix (as an endomorphism of \mathbb{R}^n with the usual inner product) is orthogonal if and only if its columns are an orthonormal set.

There is also an alternative way of characterizing these orthogonal maps.

Lemma. Let V be a finite-dimensional space and $\alpha \in \text{End}(V)$. Then α is orthogonal if and only if $\alpha^{-1} = \alpha^*$.

Proof. (\Leftarrow) Suppose $\alpha^{-1} = \alpha^*$. If $\alpha^{-1} = \alpha^*$, then

$$(\alpha\mathbf{v}, \alpha\mathbf{v}) = (\mathbf{v}, \alpha^*\alpha\mathbf{v}) = (\mathbf{v}, \alpha^{-1}\alpha\mathbf{v}) = (\mathbf{v}, \mathbf{v}).$$

(\Rightarrow) If α is orthogonal and $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ is an orthonormal basis for V , then for $1 \leq i, j \leq n$, we have

$$\delta_{ij} = (\mathbf{v}_i, \mathbf{v}_j) = (\alpha\mathbf{v}_i, \alpha\mathbf{v}_j) = (\mathbf{v}_i, \alpha^*\alpha\mathbf{v}_j).$$

So we know

$$\alpha^*\alpha(\mathbf{v}_j) = \sum_{i=1}^n (\mathbf{v}_i, \alpha^*\alpha\mathbf{v}_j) \mathbf{v}_i = \mathbf{v}_j.$$

So by linearity of $\alpha^*\alpha$, we know $\alpha^*\alpha = \text{id}_V$. So $\alpha^* = \alpha^{-1}$. \square

Corollary. $\alpha \in \text{End}(V)$ is orthogonal if and only if α is represented by an orthogonal matrix, i.e. a matrix A such that $A^T A = A A^T = I$, with respect to any orthonormal basis.

Proof. Let $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ be an orthonormal basis for V . Then suppose α is represented by A . So α^* is represented by A^T . Then $A^* = A^{-1}$ if and only if $A A^T = A^T A = I$. \square

Definition (Orthogonal group). Let V be a real inner product space. Then the *orthogonal group* of V is

$$\text{O}(V) = \{\alpha \in \text{End}(V) : \alpha \text{ is orthogonal}\}.$$

It follows from the fact that $\alpha^* = \alpha^{-1}$ that α is invertible, and it is clear from definition that $\text{O}(V)$ is closed under multiplication and inverses. So this is indeed a group.

Proposition. Let V be a finite-dimensional real inner product space and $(\mathbf{e}_1, \dots, \mathbf{e}_n)$ is an orthonormal basis of V . Then there is a bijection

$$\begin{aligned} \text{O}(V) &\rightarrow \{\text{orthonormal basis for } V\} \\ \alpha &\mapsto (\alpha(\mathbf{e}_1, \dots, \mathbf{e}_n)). \end{aligned}$$

This is analogous to our result for general vector spaces and general bases, where we replace $\text{O}(V)$ with $\text{GL}(V)$.

Proof. Same as the case for general vector spaces and general bases. \square

Unitary maps

We are going to study the complex version of orthogonal maps, known as *unitary maps*. The proofs are almost always identical to the real case, and we will not write the proofs again.

Definition (Unitary map). Let V be a finite-dimensional complex vector space. Then $\alpha \in \text{End}(V)$ is *unitary* if

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

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

By the polarization identity, α is unitary if and only if $\|\alpha(\mathbf{v})\| = \|\mathbf{v}\|$ for all $\mathbf{v} \in V$.

Lemma. Let V be a finite dimensional complex inner product space and $\alpha \in \text{End}(V)$. Then α is unitary if and only if α is invertible and $\alpha^* = \alpha^{-1}$.

Corollary. $\alpha \in \text{End}(V)$ is unitary if and only if α is represented by a unitary matrix A with respect to any orthonormal basis, i.e. $A^{-1} = A^\dagger$.

Definition (Unitary group). Let V be a finite-dimensional complex inner product space. Then the *unitary group* of V is

$$U(V) = \{\alpha \in \text{End}(V) : \alpha \text{ is unitary}\}.$$

Proposition. Let V be a finite-dimensional complex inner product space. Then there is a bijection

$$\begin{aligned} U(V) &\rightarrow \{\text{orthonormal basis of } V\} \\ \alpha &\mapsto \{\alpha(\mathbf{e}_1), \dots, \alpha(\mathbf{e}_n)\}. \end{aligned}$$

8.4 Spectral theory

We are going to classify matrices in inner product spaces. Recall that for general vector spaces, what we effectively did was to find the orbits of the conjugation action of $\text{GL}(V)$ on $\text{Mat}_n(\mathbb{F})$. If we have inner product spaces, we will want to look at the action of $O(V)$ or $U(V)$ on $\text{Mat}_n(\mathbb{F})$. In a more human language, instead of allowing arbitrary basis transformations, we only allow transforming between orthonormal basis.

We are not going to classify all endomorphisms, but just self-adjoint and orthogonal/unitary ones.

Lemma. Let V be a finite-dimensional inner product space, and $\alpha \in \text{End}(V)$ self-adjoint. Then

- (i) α has a real eigenvalue, and all eigenvalues of α are real.
- (ii) Eigenvectors of α with distinct eigenvalues are orthogonal.

Proof. We are going to do real and complex cases separately.

- (i) Suppose first V is a complex inner product space. Then by the fundamental theorem of algebra, α has an eigenvalue, say λ . We pick $\mathbf{v} \in V \setminus \{0\}$ such that $\alpha\mathbf{v} = \lambda\mathbf{v}$. Then

$$\bar{\lambda}(\mathbf{v}, \mathbf{v}) = (\lambda\mathbf{v}, \mathbf{v}) = (\alpha\mathbf{v}, \mathbf{v}) = (\mathbf{v}, \alpha\mathbf{v}) = (\mathbf{v}, \lambda\mathbf{v}) = \lambda(\mathbf{v}, \mathbf{v}).$$

Since $\mathbf{v} \neq \mathbf{0}$, we know $(\mathbf{v}, \mathbf{v}) \neq 0$. So $\lambda = \bar{\lambda}$.

For the real case, we pretend we are in the complex case. Let $\mathbf{e}_1, \dots, \mathbf{e}_n$ be an orthonormal basis for V . Then α is represented by a symmetric matrix A (with respect to this basis). Since real symmetric matrices are Hermitian viewed as complex matrices, this gives a self-adjoint endomorphism of \mathbb{C}^n . By the complex case, A has real eigenvalues only. But the eigenvalues of A are the eigenvalues of α and $M_A(t) = M_\alpha(t)$. So done.

Alternatively, we can prove this without reducing to the complex case. We know every irreducible factor of $M_\alpha(t)$ in $\mathbb{R}[t]$ must have degree 1 or 2, since the roots are either real or come in complex conjugate pairs. Suppose $f(t)$ were an irreducible factor of degree 2. Then

$$\left(\frac{m_\alpha}{f}\right)(\alpha) \neq 0$$

since it has degree less than the minimal polynomial. So there is some $\mathbf{v} \in V$ such that

$$\left(\frac{M_\alpha}{f}\right)(\alpha)(\mathbf{v}) \neq \mathbf{0}.$$

So it must be that $f(\alpha)(\mathbf{v}) = \mathbf{0}$. Let $U = \langle \mathbf{v}, \alpha(\mathbf{v}) \rangle$. Then this is an α -invariant subspace of V since f has degree 2.

Now $\alpha|_U \in \text{End}(U)$ is self-adjoint. So if $(\mathbf{e}_1, \mathbf{e}_2)$ is an orthonormal basis of U , then α is represented by a real symmetric matrix, say

$$\begin{pmatrix} a & b \\ b & a \end{pmatrix}$$

But then $\chi_{\alpha|_U}(t) = (t-a)^2 - b^2$, which has real roots, namely $a \pm b$. This is a contradiction, since $M_{\alpha|_U} = f$, but f is irreducible.

- (ii) Now suppose $\alpha\mathbf{v} = \lambda\mathbf{v}$, $\alpha\mathbf{w} = \mu\mathbf{w}$ and $\lambda \neq \mu$. We need to show $(\mathbf{v}, \mathbf{w}) = 0$. We know

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

by definition. This then gives

$$\lambda(\mathbf{v}, \mathbf{w}) = \mu(\mathbf{v}, \mathbf{w})$$

Since $\lambda \neq \mu$, we must have $(\mathbf{v}, \mathbf{w}) = 0$. □

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 previous lemma, α has a real eigenvalue, say λ . Then we can find an eigenvector $\mathbf{v} \in V \setminus \{0\}$ such that $\alpha\mathbf{v} = \lambda\mathbf{v}$.

Let $U = \langle \mathbf{v} \rangle^\perp$. Then we can write

$$V = \langle \mathbf{v} \rangle \perp U.$$

We now want to prove α sends U into U . Suppose $\mathbf{u} \in U$. Then

$$(\mathbf{v}, \alpha(\mathbf{u})) = (\alpha\mathbf{v}, \mathbf{u}) = \lambda(\mathbf{v}, \mathbf{u}) = 0.$$

So $\alpha(\mathbf{u}) \in \langle \mathbf{v} \rangle^\perp = U$. So $\alpha|_U \in \text{End}(U)$ and is self-adjoint.

By induction on $\dim V$, U has an orthonormal basis $(\mathbf{v}_2, \dots, \mathbf{v}_n)$ of α eigenvectors. Now let

$$\mathbf{v}_1 = \frac{\mathbf{v}}{\|\mathbf{v}\|}.$$

Then $(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n)$ is an orthonormal basis of eigenvectors for α . \square

Corollary. Let V be a finite-dimensional vector space and α self-adjoint. Then V is the orthogonal (internal) direct sum of its α -eigenspaces.

Corollary. Let $A \in \text{Mat}_n(\mathbb{R})$ be symmetric. Then there exists an orthogonal matrix P such that $P^T A P = P^{-1} A P$ is diagonal.

Proof. Let (\cdot, \cdot) be the standard inner product on \mathbb{R}^n . Then A is self-adjoint as an endomorphism of \mathbb{R}^n . So \mathbb{R}^n has an orthonormal basis of eigenvectors for A , say $(\mathbf{v}_1, \dots, \mathbf{v}_n)$. Taking $P = (\mathbf{v}_1 \ \mathbf{v}_2 \ \dots \ \mathbf{v}_n)$ gives the result. \square

Corollary. Let V be a finite-dimensional real inner product space and $\psi : V \times V \rightarrow \mathbb{R}$ a symmetric bilinear form. Then there exists an orthonormal basis $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ for V with respect to which ψ is represented by a diagonal matrix.

Proof. Let $(\mathbf{u}_1, \dots, \mathbf{u}_n)$ be any orthonormal basis for V . Then ψ is represented by a symmetric matrix A . Then there exists an orthogonal matrix P such that $P^T A P$ is diagonal. Now let $\mathbf{v}_i = \sum P_{ki} \mathbf{u}_k$. Then $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ is an orthonormal basis since

$$\begin{aligned} (\mathbf{v}_i, \mathbf{v}_j) &= \left(\sum P_{ki} \mathbf{u}_k, \sum P_{\ell j} \mathbf{u}_\ell \right) \\ &= \sum P_{ik}^T (\mathbf{u}_k, \mathbf{u}_\ell) P_{\ell j} \\ &= [P^T A P]_{ij} \\ &= \delta_{ij}. \end{aligned}$$

Also, ψ is represented by $P^T A P$ with respect to $(\mathbf{v}_1, \dots, \mathbf{v}_n)$. \square

Note that the diagonal values of $P^T A P$ are just the eigenvalues of A . So the signature of ψ is just the number of positive eigenvalues of A minus the number of negative eigenvalues of A .

Corollary. Let V be a finite-dimensional real vector space and ϕ, ψ symmetric bilinear forms on V such that ϕ is positive-definite. Then we can find a basis $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ for V such that both ϕ and ψ are represented by diagonal matrices with respect to this basis.

Proof. We use ϕ to define an inner product. Choose an orthonormal basis for V (equipped with ϕ) $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ with respect to which ψ is diagonal. Then ϕ is represented by I with respect to this basis, since $\psi(\mathbf{v}_i, \mathbf{v}_j) = \delta_{ij}$. So done. \square

Corollary. If $A, B \in \text{Mat}_n(\mathbb{R})$ are symmetric and A is positive definite (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 $A \in \text{Mat}_n(\mathbb{C})$ is Hermitian, then there exists a unitary matrix $U \in \text{Mat}_n(\mathbb{C})$ such that

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

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 $A, B \in \text{Mat}_n(\mathbb{C})$ be Hermitian, and A positive definite (i.e. $\mathbf{v}^\dagger A \mathbf{v} > 0$ for $\mathbf{v} \in V \setminus \{0\}$). Then there exists some invertible Q such that $Q^\dagger A Q$ and $Q^\dagger B Q$ are diagonal.

That's all for self-adjoint matrices. How about unitary matrices?

Theorem. Let V be a finite-dimensional complex vector space and $\alpha \in U(V)$ be unitary. Then V has an orthonormal basis of α eigenvectors.

Proof. By the fundamental theorem of algebra, there exists $\mathbf{v} \in V \setminus \{0\}$ and $\lambda \in \mathbb{C}$ such 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

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

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. \square

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.

However, we cannot generalize this in the real orthogonal case. For example,

$$\begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \in O(\mathbb{R}^2)$$

cannot be diagonalized (if $\theta \notin \pi\mathbb{Z}$). However, in example sheet 4, you will find a classification of $O(V)$, and you will see that the above counterexample is the worst that can happen in some sense.