LTCC notes on Lie groups and their representations

Michael Singer
University College London
Jan–Feb 2020

Contents

CHAPTER I

Background, notation, definitions

1. Definitions and examples of Lie groups

I assume that you are all comfortable with the definition of smooth manifolds, smooth maps between manifolds, etc.

> The default option for all manifolds in this course will be smooth, boundaryless and *finite-dimensional.*

DEFINITION I.1.1. A *Lie group* G is a smooth (finite-dimensional) manifold, which is also a group in such a way that the group operations of multiplication and inversion

$$\mu : G \times G \to G,\ \mu(g_1, g_2) = g_1 g_2,\ \ \iota : G \to G,\ \ \iota(g) = g^{-1}$$

are smooth maps.

If G and H are Lie groups then a map $f : G \to H$ is called a *homomorphism of Lie groups* if it is a homomorphism that is also smooth.

EXAMPLE I.1.2. The additive group $\mathbb{R}^n$ is a Lie group. The torus $\mathbb{T}^n = \mathbb{R}^n/\mathbb{Z}^n$ is a Lie group. The smooth structure is inherited from $\mathbb{R}^n$ as a quotient space, and so is the group operation. The lattice $\mathbb{Z}^n$ is also a (0-dimensional) Lie group, and the inclusion $\mathbb{Z}^n \hookrightarrow \mathbb{R}^n$ and quotient map $\mathbb{R}^n \longrightarrow \mathbb{T}^n$ are both homomorphisms of Lie groups.

We allow ourselves the flexibility of writing the multiplication law additively for abelian groups, in accordance with usual practice.

2. The general linear groups

DEFINITION I.2.1. The group of invertible (linear) endomorphisms of $\mathbb{R}^n$ is denoted $\mathrm{GL}_n(\mathbb{R})$ (or often $\mathrm{GL}(n, \mathbb{R})$) and is called the *(real) general linear group.* Similarly, the group of complex linear invertible endomorphisms of $\mathbb{C}^n$ is denoted $\mathrm{GL}_n(\mathbb{C})$ and is called the *(complex) general linear group.*

PROPOSITION I.2.2. $\mathrm{GL}_n(\mathbb{R})$ *and* $\mathrm{GL}_n(\mathbb{C})$ *are Lie groups, respectively of real dimension* n^2 *and* $2n^2$*.*

PROOF. We have $\mathrm{GL}_n(\mathbb{R}) \subset \mathrm{End}(\mathbb{R}^n) = \mathbb{R}^{n^2}$, consisting of those matrices A with $\det A \neq 0$. Now the determinant map is smooth (as it is polynomial in the coefficients of A). In particular det is continuous and so $\mathrm{GL}_n(\mathbb{R})$ is an open subset, hence an open submanifold, of $\mathrm{End}(\mathbb{R}^n)$. The group operation $(A, B) \mapsto AB$ is matrix multiplication. This is a smooth map, since it is again polynomial in the coefficients of A and B.

Inversion $A \mapsto A^{-1}$ is also smooth because it is a rational function of the coefficients of A (determinant of minors of A divided by $\det A$).

The proof for $\mathrm{GL}_n(\mathbb{C})$ is precisely analogous. The dimensions are obvious: in each case the dimension is the same as the dimension of the space of endomorphisms of the vector space. □

REMARK I.2.3. If V is a real n-dimensional vector space, then $\mathrm{Aut}(V)$ is the group of invertible endomorphisms of V. It is an abstract version of $\mathrm{GL}_n(\mathbb{R})$, in the sense that if an isomorphism $V \to \mathbb{R}^n$ is chosen, then $\mathrm{Aut}(V)$ is identified with $\mathrm{GL}_n(\mathbb{R})$. The same remark applies if V is a complex n-dimensional vector space.

Remark that $\mathrm{GL}_n(\mathbb{R})$ is not connected: the sign of $\det A$ distinguishes the two connected components. (See Exercise I.2.2 below.)

Since $\mathrm{GL}_n(\mathbb{R})$ is a smooth manifold, it has a tangent space $T_A \mathrm{GL}_n(\mathbb{R})$ at every point $A \in \mathrm{GL}_n(\mathbb{R})$. It is an open subset of the vector space $\mathrm{End}(\mathbb{R}^n)$, so the tangent space can itself be identified with this vector space

$$T_A \mathrm{GL}_n(\mathbb{R}) = \mathrm{End}(\mathbb{R}^n) = \mathbb{R}^{n^2}. \tag{I.2.1}$$

It is often convenient to identify T_A with T_1, the tangent space at the identity element of GL_n by left- or right-multiplication. These are just the maps

$$T_1 \longrightarrow T_A, \quad X \mapsto AX \text{ or } X \mapsto XA. \tag{I.2.2}$$

2.1. The special linear, orthogonal, and special orthogonal groups.

DEFINITION I.2.4. The set

$$\{A \in \mathrm{GL}_n(\mathbb{R}) : \det A = 1\}$$

is denoted $\mathrm{SL}_n(\mathbb{R})$ and is called the special linear group. The set

$$\{A \in \mathrm{GL}_n(\mathbb{R}) : A^t A = 1\}$$

is denoted O_n and is called the orthogonal group. The intersection is denoted

$$\mathrm{SO}_n = \mathrm{SL}_n(\mathbb{R}) \cap O_n$$

and is called the special orthogonal group.

PROPOSITION I.2.5. *As the notation anticipates,* $\mathrm{SL}_n(\mathbb{R})$, O_n *and* SO_n *are all groups, in fact Lie groups.*

The dimensions of these groups are

$$\dim \mathrm{SL}_n(\mathbb{R}) = n^2 - 1, \dim O_n = \frac{n(n-1)}{2}, \dim \mathrm{SO}_n = \frac{n(n-1)}{2} - 1. \tag{I.2.3}$$

PROOF. That SL_n is a subgroup of GL_n follows from the fact that det is a homomorphism, $\det(AB) = \det A \det B$. If $A^t A = 1$ and $B^t B = 1$ then

$$(AB)^t(AB) = B^t A^t AB = B^t(A^t A)B = B^t B = 1$$

so O_n is a group. Hence SO_n is also a group.

That these groups are in fact Lie groups follows from a general theorem that a closed subgroup of a Lie group is again a Lie group (see later). Let us show directly that O_n

is a closed submanifold of $\mathrm{GL}_n(\mathbb{R})$. Let $\Phi(A) = A^tA - 1$. Then $\Phi(A)$ is symmetric for any A, and hence lies in a real vector space of dimension $N = n(n+1)/2$. Thus

$$\Phi : \mathrm{GL}_n(\mathbb{R}) \longrightarrow \mathbb{R}^N$$

is a smooth map. Its zero-set will be a closed submanifold of $\mathrm{GL}_n(\mathbb{R})$ if the derivative $\mathrm{d}\Phi$ is surjective at every point $A \in O_n$. Now the tangent space $T_A\,\mathrm{GL}_n(\mathbb{R})$ is just the $n \times n$ matrices $\mathrm{End}(\mathbb{R}^n)$, since GL_n is an open subset of $\mathrm{End}(\mathbb{R}^n)$. Differentiating Φ, we obtain the map

$$\mathrm{d}\Phi_A(a) = \text{linear part of } a \mapsto (a^t + A^t)(A + a) - 1, \text{ i.e. } a \mapsto a^tA + A^ta.$$

We want to show this is onto the space of all symmetric matrices. If s is a symmetric $n \times n$ matrix, set

$$a = \frac{1}{2}As.$$

Then

$$A^ta = \frac{1}{2}A^tAs = \frac{1}{2}s \text{ and so } \mathrm{d}\Phi_A(a) = s.$$

Thus $\mathrm{d}\Phi_A$ is indeed onto the space of symmetric $n \times n$ matrics, and hence O_n is a closed submanifold of $\mathrm{GL}_n(\mathbb{R})$. The group operations are automatically smooth as they are restrictions from GL_n. We have shown that O_n is a Lie group.

We omit the proofs of the remaining parts of the Proposition. □

PROPOSITION I.2.6. *The orthogonal group O_n is compact and has two connected components.*

PROOF. We have already seen that $O_n \subset \mathrm{GL}_n(\mathbb{R})$ is *closed* so for compactness, we just need to show it is also bounded. Now if $A = (a_{ij})$, then

$$(A^tA)_{jk} = \sum_j a_{ij}a_{ik}$$

and so

$$\operatorname{tr} A^tA = \sum a_{ij}^2 = n$$

Thus O_n is contained in a sphere of radius $\sqrt{n}$ inside $\mathbb{R}^{n^2}$, so certainly bounded. Thus O_n, being a closed and bounded subspace of $\mathbb{R}^{n^2}$, is compact.

If $A \in O_n$, then

$$\det A^tA = (\det A)^2 = 1$$

so $\det A = \pm 1$. det is continuous, so O_n has at least two connected components, distinguished by the sign of $\det A$. We now show that SO_n is a connected set. For this we use the following fact, proved below.

> Any $A \in O_n$ is a product of at most n reflections.

Recall that the reflection in the hyperplane $\{e \cdot x = 0\}$, where e is a unit vector in $\mathbb{R}^n$, is the linear map

$$R(x) = x - 2(x \cdot e)e. \tag{I.2.4}$$

This acts as the identity on the mirror plane $\{e \cdot x = 0\}$ and reverses vectors normal to it. (By default, 'reflection' will mean a map of the form (I.2.4), with a codimension-1 mirror plane.) Obviously, $R^2 = 1$ and it is easy to check that $R \in O_n$.

To show that SO_n is connected, it is sufficient to show that if $A \in \mathrm{SO}_n$, then A can be connected by a path to 1. Such A must be the product of an *even* number of reflections

$$A = R_1 R_2 \cdots R_{2k-1} R_{2k}.$$

since the determinant of any reflection is -1. Such a product of an even number of reflections can be connected to the identity by pairing off the mirror planes and deforming them to each other. More precisely, pick any pair of reflections, say R_1 and R_2, with normals e_1 and e_2. It is clear that we can find a path in the plane of e_1 and e_2 which rotates e_2 onto e_1, thus deforming R_2 continuously to R_1. Thus A is connected by a path to the product

$$A' = R_3 R_4 \cdots R_{2k-1} R_{2k}$$

of $2k-2$ reflections. Continuing to cancel reflections in pairs like this, we eventually see that A can be connected by a path to 1. (Another way to think of this argument is that $R_1 R_2$ is a rotation in the plane spanned by e_1 and e_2, fixing the $(n-2)$-plane othogonal to the span of these vectors. Any such rotation can be deformed to the identity, leaving us with A' as before.) Hence SO_n is connected.

Since SO_n is mapped diffeomorphically to the set $\{A \in O_n : \det A = -1\}$ by any given reflection, it follows that the latter set is also connected. $\square$

We now prove the boxed statement above. For this suppose that $(e_1, \dots, e_n)$ and $(f_1, \dots, f_n)$ are two orthonormal bases of $\mathbb{R}^n$. Let k be the largest integer such that

$$e_j = f_j \text{ for } j = 1, \dots, k.$$

We find a reflection R such that

$$Re_j = f_j \text{ for } j = 1, \dots, k+1.$$

Since $e_{k+1} \neq f_{k+1}$,

$$e := \frac{e_{k+1} - f_{k+1}}{|e_{k+1} - f_{k+1}|}$$

is well-defined, and the reflection R with mirror plane $\{e \cdot x = 0\}$ switches e_{k+1} and f_{k+1},

$$Re_{k+1} = f_{k+1}, \quad Rf_{k+1} = e_{k+1}.$$

Since the e- and f-bases are both orthonormal, $e_1, \dots, e_k$ both lie in the mirror plane, and so are fixed by R. If we set Re_j then $Re_j = e_j = f_j$ for $j = 1, \dots, k$ and $Re_{k+1} = f_{k+1}$ by construction. Applying this argument inductively we see that for any pair of orthonormal bases (e_j) and (f_j) there are $\leqslant n$ reflections $R_1, \dots, R_m$ such

$$(R_1 \dots R_m) e_j = f_j \text{ for all } j.$$

Now given $A \in O_n$, pick the standard orthonormal basis (e_j) and set $f_j = Ae_j$. Applying the above construction to these two bases, we have $A = R_1 \dots R_m$ as claimed.

Exercise I.2.1. Verify the statements about the dimensions of the groups in Proposition I.2.5.

Exercise I.2.2. Show that $\mathrm{GL}_n(\mathbb{R})$ has two connected components. [Hint: recall that if $A \in \mathrm{GL}_n(\mathbb{R})$, then there exist unique P and R, such that $A = PR$, where P is a diagonal matrix with positive entries and $R \in O_n$. Deduce from this that $\mathrm{GL}_n(\mathbb{R})$ is homotopy equivalent to O_n. For $A = PR$, start by diagonalizing A^tA.]

2.2. Special linear, unitary and special unitary groups. The following section parallels the previous one, with $\mathbb{R}$ replaced by $\mathbb{C}$.

Definition I.2.7. The set

$$\{A \in \mathrm{GL}_n(\mathbb{C}) : \det A = 1\}$$

is denoted $\mathrm{SL}_n(\mathbb{C})$ and is called the (complex) special linear group. The set

$$\{A \in \mathrm{GL}_n(\mathbb{C}) : A^*A = 1\}$$

is denoted U_n and is called the unitary group (Here A^* is the conjugate transpose of A.) The intersection is denoted

$$\mathrm{SU}_n = \mathrm{SL}_n(\mathbb{C}) \cap U_n$$

and is called the special unitary group.

Proposition I.2.8. *As the notation anticipates,* $\mathrm{SL}_n(\mathbb{C})$*,* U_n *and* SU_n *are all groups, in fact Lie groups.*

The dimensions of these groups are

$$\dim_{\mathbb{R}} \mathrm{SL}_n(\mathbb{C}) = 2n^2 - 2, \ \ \dim_{\mathbb{R}} U_n = n^2, \ \ \dim_{\mathbb{R}} \mathrm{SU}_n = n^2 - 1. \tag{I.2.5}$$

Proof. The proof that these sets are groups is the same as in the real case. We defer the proof that they are Lie groups, since it will follow from a general fact, to be proved below. On the other hand, one can prove that U_n is a Lie group can be proved in exactly the same way as we proved that O_n is a Lie group in Proposition I.2.5. The dimensions are also more easily computed once we have defined the Lie algebra of a Lie group. □

Proposition I.2.9. *The orthogonal group* U_n *is compact and connected.*

Proof. The set U_n is closed in $\mathrm{GL}_n(\mathbb{C})$ and, as in the real case, it is bounded because $A^*A = 1$ implies

$$\sum_{i,j} |a_{ij}|^2 = n$$

where the a_{ij} are the (complex) coefficients of A. Hence U_n is contained in a sphere in $\mathbb{C}^{n^2}$ and is thus compact.

To prove connectedness, consider the eigenvalue decomposition of A. First of all, we recall the elementary fact that A *is* diagonalizable. This follows because if $x \neq 0$ is any eigenvector then A preserves the decomposition $\mathbb{C}^n = \mathbb{C}x \oplus x^\perp$. This fact leads to a proof by induction that $\mathbb{C}^n$ has a decomposition into the eigenspaces of A. That A preserves $x^\perp$ is a one-line computation: if $Ax = \lambda x$ and $y \in x^\perp$, then

$$x^*Ay = \overline{y^*A^*x} = \overline{y^*A^{-1}x} = \overline{\lambda^{-1}y^*x} = \lambda x^*y = 0,$$

where we've used the elementary fact that every eigenvalue λ lies on the unit circle. Hence, given A, there exists a unitary matrix Q (whose columns are an orthonormal basis of eigenvectors of A) such that

$$QAQ^{-1} = \mathrm{diag}(\lambda_1, \ldots, \lambda_n), \quad |\lambda_j| = 1.$$

Any such diagonal matrix can be smoothly joined to the identity, showing that A itself can be smoothly joined to $1 \in \mathrm{SU}_n$. Thus U_n is connected. □

EXERCISE I.2.3. Show that if $A \in U_n$ then there exists a unitary matrix Q and a diagonal matrix Λ such that $A = Q^{-1}\Lambda Q$. Show further that $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_n)$, where $|\lambda_j| = 1$.

EXERCISE I.2.4. Verify the statements about the dimensions in Proposition I.2.8.

3. Quaternions and the symplectic group

Recall that the (skew) field $\mathbb{H}$ of quaternions may be defined by introducing three 'imaginary units' i, j, k satisfying

$$i^2 = j^2 = k^2 = ijk = -1 \tag{I.3.1}$$

and setting

$$\mathbb{H} = \{x_0 + x_1 i + x_2 j + x_3 k : (x_0, x_1, x_2, x_3) \in \mathbb{R}^4\}. \tag{I.3.2}$$

Then $\mathbb{H}$ is regarded as a four-dimensional real vector space, and quaternions $q, q' \in \mathbb{H}$ are multiplied in the obvious way using Hamilton's identities (I.3.1). It follows from these identities that

$$ij = k, ji = -k$$

so the multiplication in $\mathbb{H}$ is not commutative. This apart, addition and multiplication in $\mathbb{H}$ satisfy all the usual axioms for a field. The existence of multiplicative inverses follows from the properties of quaternionic conjugation. The *conjugate* of the quaternion

$$q = x_0 + x_1 i + x_2 j + x_3 k,$$

is defined to be

$$\bar{q} = x_0 - x_1 i - x_2 j - x_3 k.$$

Then

$$q\bar{q} = \bar{q}q = x_0^2 + x_1^2 + x_2^2 + x_3^2 =: |q|^2$$

It follows that if $q \neq 0$, then $\bar{q}/|q|^2$ is the multiplicative inverse of q, a precise analogue of familiar formula for the multiplicative inverse of a non-zero complex number.

When considering quaternionic vector spaces, one has to make a decision about whether scalar multiplication should act on the left or the right, and different authors do different things. We shall think of $\mathbb{H}^n$ as column vectors of quaternions, and $\mathrm{GL}_n(\mathbb{H})$ as invertible matrices with entries in the quaternions, acting in the usual way on $\mathbb{H}^n$ by matrix multiplication, $x \mapsto Ax$, where $x \in HH^n$. For this to be $\mathbb{H}$-linear, we need to make the scalars act by multiplication on the right, $x \mapsto x\bar{q}$.

It is often convenient to think of a quaternionic vector space as a complex vector space 'with extra structure' in such a way that the action of $\mathrm{GL}_n(\mathbb{H})$ is *complex-linear*. In order to do this, identify $\mathbb{C}^2$ with $\mathbb{H}$ by the map

$$(z_1, z_2) \leftrightarrow z_1 + jz_2. \tag{I.3.3}$$

If $a = a_1 + ja_2$ is another quaternion, then

$$(a_1 + ja_2)(z_1 + jz_2) = (a_1 z_1 - \bar{a}_2 z_2) + j(a_2 z_1 + \bar{a}_1 z_2)$$

so this induces a complex-linear map in (z_1, z_2). Explicitly, we have

$$\begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \longmapsto \begin{bmatrix} a_1 & -\bar{a}_2 \\ a_2 & \bar{a}_1 \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \tag{I.3.4}$$

Thus we see explicitly $\mathrm{GL}_1(\mathbb{H})$ as a subgroup of $\mathrm{GL}_2(\mathbb{C})$.

More generally, if we identify $\mathbb{H}^n = \mathbb{C}^n + j\mathbb{C}^n$, then $\mathrm{GL}_n(\mathbb{H})$ is identified with a subgroup of $\mathrm{GL}_{2n}(\mathbb{C})$. To understand this subgroup, note that a complex-linear map of $\mathbb{C}^{2n}$ is $\mathbb{H}$-linear if it commutes with multiplication *on the right* by j. If $q = x + jy$, where $x, y \in \mathbb{C}^n$, then

$$qj = xj + jyj = -\bar{y} + j\bar{x}. \tag{I.3.5}$$

Denoting this map by σ, we see that σ is $\mathbb{C}$-antilinear and $\sigma^2 = -1$.

Then

PROPOSITION I.3.1. *When $\mathbb{H}^n$ is identified with $\mathbb{C}^{2n}$ as above ($\mathbb{H}^n = \mathbb{C}^n + j\mathbb{C}^n$),*

$$\mathrm{GL}_n(\mathbb{H}) = \{A \in \mathrm{GL}_{2n}(\mathbb{C}) : \sigma A = A\sigma\}. \tag{I.3.6}$$

DEFINITION I.3.2. The (compact) symplectic group Sp_n is the intersection

$$\mathrm{Sp}_n = U_{2n} \cap \mathrm{GL}_n(\mathbb{H})$$

where $\mathbb{H}^n = \mathbb{C}^n + j\mathbb{C}^n$ as above. Thus, with σ defined as in (I.3.5)

$$\mathrm{Sp}_n = \{P \in U_{2n} : P\sigma = \sigma P\}, \tag{I.3.7}$$

More explicitly, Sp_n consists of the unitary $2n \times 2n$ complex matrices of block form

$$P = \begin{bmatrix} A & -\bar{B} \\ B & \bar{A} \end{bmatrix}. \tag{I.3.8}$$

4. Abstract version

Let V be a complex vector space of complex dimension n. Then we have $\mathrm{Aut}(V)$, the group of complex-linear invertible endomorphisms of V. Any complex linear isomorphism $V \simeq \mathbb{C}^n$ (i.e. choice of basis) gives an identification of $\mathrm{Aut}(V)$ with $\mathrm{GL}_n(\mathbb{C})$. If h is a hermitian inner product on V, then $\mathrm{Aut}(V, h)$ is the subgroup of $\mathrm{Aut}(V)$ of h-isometries:

$$h(Av, Av') = h(v, v') \text{ for all } v, v' \in V.$$

If we choose an h-orthonormal basis of V, we obtain an isomorphism $\mathrm{Aut}(V, h) \simeq U_n$.

In order to define the abstract version of $\mathrm{SL}_n(\mathbb{C})$, recall that $A : V \to V$ defines a map $\Lambda^n A : \Lambda^n V \to \Lambda^n V$. Since $\Lambda^n V$ is a complex 1-dimensional vector space, $\Lambda^n A$ is given by multiplication by a complex number. Requiring this to be 1 defines the subgroup of 'special' automorphisms inside $\mathrm{Aut}(V)$.

Following Adams, observe that we can deal with both real and quaternionic versions of these groups by the introduction of a so-called *structure map*:

DEFINITION I.4.1. A *structure map* σ on a hermitian vector space (V, h) is an antilinear map $\sigma : V \to V$ which satisfies $\sigma^2 = 1$ (real) or $\sigma^2 = -1$ (quaternionic) and is an h-isometry.

We obtain subgroups of

$$\mathrm{Aut}_\sigma(V) \subset \mathrm{Aut}(V), \quad \mathrm{Aut}_\sigma(V, h) \subset \mathrm{Aut}(V, h)$$

where the subscript indicates that we restrict to elements which commute with σ. Then it is easy to check that if σ is real, $\mathrm{Aut}_\sigma(V)$ is an abstract version of $\mathrm{GL}_n(\mathbb{R})$ and $\mathrm{Aut}_\sigma(V, h)$ is an abstract version of O_n. Similarly if σ is quaternionic, then $\dim_{\mathbb{C}} V = 2m$ must be even and $\mathrm{Aut}_\sigma(V, h)$ is an abstract version of Sp_m.

This point of view allows some arguments to be streamlined: we prove results in the complex setting, and then get corresponding results in the real and quaternionic settings by adding a structure map.

5. Actions and representations

If G is a group and X is a set, then an *action* of G on X is a map

$$(g, x) \mapsto gx \tag{I.5.1}$$

which satisifies

$$(g_1 g_2)x = g_1(g_2 x), \quad ex = x \tag{I.5.2}$$

(where e is the identity element of e). If G is a Lie group and X is a smooth manifold, then a *Lie group action* of G on X is an action which is also smooth as a map $G \times X \to X$.

Definition I.5.1. A representation of a group G on a vector space V is a *linear* action of G on V. (In other words, $x \mapsto gx$ is linear for every $g \in G$. If G is a Lie group, then a Lie group representation of G is a representation that is also smooth.

Equivalently, a (Lie) group representation of G on V is a (Lie) group homomorphism $\rho : G \longrightarrow \mathrm{GL}(V)$.

Definition I.5.2. A representation is irreducible if and only if the only subspaces which are preserved by G are 0 and V. The representation ρ is *unitary* if V is a hermitian vector space and $\rho(g) \in U(V)$ for every $g \in G$.

If $\rho : G \to \mathrm{Aut}(V)$ is a representation, it is convenient to call V a G-space, particularly if there is no possible confusion about how G might be acting on V. Another handy piece of notation is to write $g|V$ for $\rho(g)$ to mean the action of the group element g on V. If $\sigma : G \to \mathrm{Aut}(W)$ is another representation, then $T : V \to W$ is a homomorphism of representations (or simply a G-map) if

$$T \circ \rho = \sigma \circ T. \tag{I.5.3}$$

If this equation holds we also say that T intertwines the representations V and W. Using the notation just introduced, we may write the intertwining relation as

$$T \circ (g|V) = (g|W) \circ T \text{ for all } g \in G. \tag{I.5.4}$$

The class of G-spaces and G-maps forms a *category*. Moreover, the standard operations of (multi-)linear algebra also work in the category of G-spaces. For example if V and W are G-spaces then so are V^*, $V \otimes W$ and $\mathrm{Hom}(V, W)$. In particular, the action of G on $\mathrm{Hom}(V, W)$ is given by

$$g : T \mapsto gTg^{-1} \text{ where the RHS is more formally written } \sigma(g)T\rho(g)^{-1}. \tag{I.5.5}$$

Then

The space $\mathrm{Hom}_G(V, W) \subset \mathrm{Hom}(V, W)$ of G-maps $V \to W$ is precisely the same as the G-invariant part of $\mathrm{Hom}(V, W)$.

DEFINITION I.5.3. More generally, if V is any G-space, we write V_G for the set of G-invariant elements.

5.1. Examples.

EXAMPLE I.5.4. The basic representations of $\mathbb{R}$ on $\mathbb{C}$ are

$$\rho_a(t) = \exp(at), \tag{I.5.6}$$

where $a \in \mathbb{C}$ is any number. If we think of $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, then ρ_a gives a representation of $\mathbb{T}$ iff and only if $\rho_a(n) = 1$ for every $n \in \mathbb{Z}$. Thus

$$\rho_a \text{ defines a representation of } \mathbb{T} \Leftrightarrow a \in 2\pi i\mathbb{Z}. \tag{I.5.7}$$

The integer $a/2\pi i$ is called the weight of the representation.

EXAMPLE I.5.5. Let Σ_n be the symmetric group on $\{1, \ldots, n\}$. Then Σ_n acts on $\mathbb{C}^n$ by the obvious permutation action,

$$\sigma(x_1, \ldots, x_n) = (x_{\sigma^{-1}(1)}, \ldots, x_{\sigma^{-1}(n)}).$$

The diagonal subspace spanned by $(1, \ldots, 1)$ is clearly fixed by Σ_n. Its orthogonal complement is a representation of dimension $(n - 1)$.

6. Complete reducibility and Schur's Lemma

Let G be a finite group and let V be a (finite-dimensional) complex representation of G.

We give two fundamental results. The first is Schur's Lemma:

THEOREM I.6.1. *Let V and W be irreducible representations. Then for the space of G-maps from V to W we have*

$$\mathrm{Hom}_G(V, W) = \begin{cases} \mathbb{C} \text{ if } V \text{ are isomorphic as } G\text{-spaces}; \\ 0, \text{ otherwise.} \end{cases} \tag{I.6.1}$$

PROOF. Suppose $T : V \to W$ is a G-map. Then obviously $\ker(T)$ and $\mathrm{Im}(T)$ are G-spaces. So if V is irreducible, $\ker(T) = 0$ or $\ker(T) = V$. The second case only happens if $T = 0$. In the opposite case, T is injective and its image is a subrepresentation of W. If W is irreducible this must be the whole of W and V and W are isomorphic G-spaces. This shows the second part of the theorem. To prove the first part, we may suppose that $W = V$ and suppose that $T \in \mathrm{Hom}_G(V, V)$, $T \neq 0$. Now T has a non-zero eigenvalue $\lambda \in \mathbb{C}$, and in particular $\ker(T - \lambda 1_V) \neq 0$. This kernel is a G-space and non-zero, so by irreducibility, it is V. Hence $T = \lambda 1_V$. This completes the proof. □

THEOREM I.6.2. *There is a unique decomposition*

$$V = \bigoplus m_i V_i \tag{I.6.2}$$

of V into irreducible representation of G.

Here we have written m_iV_i for the direct sum of m_i copies of V_i. The (non-negative) integer m_i is called the multiplicity of V_i in V.

PROOF. It is enough to show that if $W \subset V$ is a proper G-subspace of V, then there is a complementary subspace which is also preserved by the action of G. The easiest way to do this is to pick a G-invariant (hermitian) inner product on V. Recall that G-invariant means

$$\langle gx, gy\rangle = \langle x, y\rangle \text{ for all } g \in G, x, y \in V.$$

If now W is preserved by G, then for all $x \in W$, $y \in W^{\perp}$,

$$\langle x, gy\rangle = \langle g^{-1}x, y\rangle = 0 \text{ since } x \in W \Rightarrow g^{-1}x \in W.$$

If $0 < W < V$ then $0 < W^{\perp} < V$, so we have decomposed V into complementary subrepresentations, and after a finite number of steps we obtain (I.6.2).

To prove uniqueness, observe that the same argument as in Schur's Lemma shows that if V_i is irreducible and V is decomposed as in (I.6.2), then

$$\dim \operatorname{Hom}_G(V_i, V) = m_i. \tag{I.6.3}$$

This shows that the multiplicity m_i of V_i in V is uniquely determined by V. □

7. The regular representation and characters

Let G be a finite group acting on a finite set X. Then we have a representation of G on $L^2(X) = \mathbb{C}^{|X|}$. This is given by

$$(g \cdot f)(x) = f(g^{-1}x). \tag{I.7.1}$$

This is, confusingly, an action. For

$$(g_1g_2 \cdot f)(x) = f((g_1g_2)^{-1}x) = f(g_2^{-1}g_1^{-1}x) = (g_2 \cdot f)(g_1^{-1}x) = (g_1 \cdot (g_2 \cdot f))(x).$$

(In fact a the action of Σ_n on $\mathbb{C}^n$ is a special case, taking $X = \{1, \dots, n\}$.)

DEFINITION I.7.1. The regular representation of G is the representation given by the above construction applied to $X = G$ with the standard action $(g, x) \mapsto gx$, (left-translation) so

$$(g \cdot f)(x) = f(g^{-1}x), g, x \in G.$$

The main result here is that $L^2(G)$ contains all irreducible representations.

THEOREM I.7.2. *Let V_i denote the irreducible representations of G. Then the decomposition of $L^2(G)$ into irreducibles is*

$$L^2(G) = \bigoplus_i d_iV_i \textit{ where } d_i = \dim V_i. \tag{I.7.2}$$

In particular, $L^2(G)$ contains all irreducible representations of G.

The usual proof uses characters.

DEFINITION I.7.3. If V is any representation of G then the *character* χ_V of V is given by the formula $\chi_V(g) = \operatorname{tr}(g|V)$.

The significance of characters in 'detecting' irreducible summands in a representation comes from the following circle of ideas. As before, let V_G be the set of elements of V fixed by G. Then it is easy to see that

$$\pi = \frac{1}{|G|}\sum_g (g|V) \tag{I.7.3}$$

is a G-map $V \to V_G$, in fact a projection ($\pi^2 = \pi$) onto this subspace. For any projection, the trace is equal to the dimension of the image. Applying this to (I.7.3),

$$\operatorname{tr}\pi = \dim V_G = \frac{1}{|G|}\sum_g \chi_V(g). \tag{I.7.4}$$

Thus the dimension of the G-invariant subspace of V is the average, taken over G, of the character χ_V.

The other property we need is that if W and V are unitary representations, then

$$\chi_{\mathrm{Hom}(V,W)}(g) = \overline{\chi_V}(g)\chi_W(g). \tag{I.7.5}$$

Define an inner product on characters,

$$(\chi_V, \chi_W) = \frac{1}{|G|}\sum_g \overline{\chi_V}(g)\chi_W(g). \tag{I.7.6}$$

Then we have an orthonormality property

Proposition I.7.4. *If V and W are irreducible unitary representations of G,*

$$(\chi_V, \chi_W) = \begin{cases} 1 \text{ if } V \simeq W; \\ 0 \text{ otherwise.} \end{cases} \tag{I.7.7}$$

Proof. By formula (I.7.5) the LHS is the average of the character $\chi_{\mathrm{Hom}(V,W)}$ over the group. By the interpretation (I.7.4) of this average, (I.7.7) is equal to $\dim \mathrm{Hom}_G(V, W)$. Since V and W are irreducible, the result now follows from Schur's Lemma. □

We can now embark on the proof of Theorem I.7.2. Let $V = L^2(G)$. We know that this has decomposition into irreducibles, and after the above Proposition, if V_i occurs with multiplicity m_i in V, we have

$$m_i = (\chi_i, \chi_V). \tag{I.7.8}$$

On the other hand, we can compute explicitly $\chi_V(g)$. A little thought shows that for $V = L^2(G)$,

$$\chi_V(g) = \begin{cases} |G| \text{ if } g = e; \\ 0 \text{ otherwise.} \end{cases} \tag{I.7.9}$$

Hence

$$m_i = (\chi_i, \chi_V) = \chi_i(e) = \dim V_i = d_i.$$

Modulo a few facts about characters which we did not prove, this proves Theorem I.7.2.

This result has some important consequences.

- The number of irreducible representations of a finite group G is finite. This is because they all occur in $L^2(G)$ and this is itself finite-dimensional.

- With d_i as above,
$$|G| = \sum_i d_i^2.$$
An interpretation of this result is the following. For any fixed elements $x, y \in V_i$, we have the 'matrix element' $\langle x, \rho_i(g)y \rangle$, which is a function on the group. Then these matrix elements, as x and y run over a basis of V_i, give d_i^2 linearly independent functions on the group. Repeating this for all i gives a basis of $L^2(G)$.
- The number of irreducible representations is equal to the number of conjugacy classes in G. To see this, recall that a *class function* is a function $\varphi : G \to \mathbb{C}$ which is constant on the elements of each conjugacy class, $\varphi(gxg^{-1}) = \varphi(x)$. Clearly the dimension of the space of class functions is equal to the number N of conjugacy classes. Now any character is a class function (because it's a trace) and the characters χ_i of the irreducible representations are orthonormal, hence linearly independent in the space of class functions. So $N \geqslant$ the number of irreducible representations. To prove equality, suppose if possible that $\varphi \neq 0$ is a class function orthogonal to all χ_i. Following Fulton and Harris, Proposition 2.30, for any representation V, form
$$\Pi_V = \sum_g \varphi(g)\rho_V(g) \tag{I.7.10}$$
This is a G-map, $V \to V$, so if $V = V_i$ is irreducible, it must be a multiple of the identity by Schur's Lemma again. One computes the multiple by taking the trace of φ: the answer is a fixed (positive) multiple of the inner product $\overline{(\varphi, \chi_{V^*})}$ and this is zero by assumption. Thus for *any* representation V, we have $\Pi_V = 0$. Now if we take the regular representation $L^2(G)$, The elements $\rho(g)$ are all linearly independent, and it follows that $\varphi(g) = 0$ for all g. Thus the set of characters χ_i of the irreducible representations of G form a basis for the space of class functions, and so the number of irreducible representations is equal to N.