

Representations of finite groups

Bill Casselman
University of British Columbia
cass@math.ubc.ca

Representations of locally compact groups are increasingly important in number theory and physics, as well as other domains of mathematics, but they can be technically very complicated. Representations of finite groups are much simpler than those of locally compact groups, but they offer a model for the general theory, and for that reason are indispensable as an introduction to the subject. The challenge is to present the material in such a way as to give some idea of techniques that might generalize, rather than use tricks that definitely work only for finite groups. Pursuing this goal systematically often clears up puzzles about finite groups as well. Representations of finite groups always decompose into a direct sum of irreducible components, but this is not canonical. Relying on it amounts to burying a distinction between sub- and quotient-representations that springs to life, causing some distress, in many practical problems.

One standard introduction to the subject is [Serre:1967/1977], and I have often followed it closely. I'll sometimes skip details when they are covered well by Serre, and to some extent I intend this to be a supplement to his treatment. One difference between his exposition and mine is that I am more interested in questions about explicit realizations of representations, rather than just some characterization of an isomorphism class. I have tried hard to state and prove things without a choice of coordinates. Keeping this in mind, I have included possibly superfluous discussions of several topics in linear algebra that in my experience are not treated carefully in the literature—change of bases, duality, tensor and exterior products, and the traces of linear operators.

All representations in this essay will be on finite-dimensional complex vector spaces, although in matters of basic linear algebra other coefficient fields will be allowed. All groups will be finite. If G is a group, it possesses a unique G -invariant measure, assigning to each point of G measure $1/|G|$, hence to all of G measure 1. I'll write

$$\int_G f(g) dg = \frac{1}{|G|} \sum_{g \in G} f(g).$$



The sign at left indicates a dangerous bend on the road, or maybe a patch that's slippery when wet. This is an idea I have copied happily from Don Knuth's **T_EXBook**. But then he copied it from Bourbaki.

Contents

1. Introduction
2. Decomposition
3. Duality
4. Tensor products
5. Trace and determinant
6. Schur orthogonality
7. Characters
8. Central functions
9. Conjugacy classes
10. Induced representations
11. Normal subgroups
12. The Fourier transform on abelian groups
13. Harmonic analysis on G
14. Examples
15. References

1. Introduction

A representation (π, V) of G on a finite-dimensional complex vector space V is a homomorphism π from the group G to the group $\text{GL}(V)$ of invertible complex linear maps from V to itself. This is a very simple definition, and it gives no idea at all of why looking at such representations is such a fruitful idea. I'll try to give first a hint of what's to come.

The theory of representations is a natural extension of eigenvector expansions. If S and T are two linear operators on a vector space V that commute with each other, any eigenspace of S is taken into itself by T , and therefore decomposes into T -eigenspaces. We get therefore a decomposition of V into joint eigenspaces. It is very natural to ask, suppose we look at larger sets of commuting operators? Suppose operators in the set don't commute? There is presumably not much of a theory unless some assumptions are made about the set, and it is natural to assume they form a group.

The case of commuting operators is already interesting. For example, suppose G to be the group \mathbb{Z}/n for $n > 0$. A character of G is a homomorphism from G to \mathbb{C}^\times . Its image is necessarily contained in the group μ_n of n -th roots of unity. In fact, every character of G is completely specified by the image of 1, and the group of all characters is isomorphic to μ_n —to an n -th root of unity ζ corresponds the character $k \mapsto \zeta^k$. Suppose (π, V) to be a representation of G . Then V decomposes into eigenspaces for the operator $\pi(1)$. The eigenvalues are n -th roots of unity, so

$$V = \bigoplus V_\zeta$$

in which the sum is over μ_n , and

$$\pi(k)v = \zeta^k v$$

for v in V_ζ . We can deduce from this a formula for the component v_ζ of v in V_ζ . If

$$v = \sum_{\zeta \in \mu_n} v_\zeta$$

then for α in μ_n

$$\begin{aligned} \pi(0)v &= \sum v_\zeta \\ \alpha^{-1}\pi(1)v &= \sum (\zeta/\alpha)v_\zeta \\ &\dots \\ \alpha^{-k}\pi(k)v &= \sum (\zeta/\alpha)^k v_\zeta \\ \sum_k \alpha^{-k}\pi(k)v &= \sum_\zeta \sum_k (\zeta/\alpha)^k v_\zeta \\ &= n \cdot v_\alpha, \end{aligned}$$

since if β is an n -th root other than 1

$$S = \sum_{k=0}^{n-1} \beta^k = 0,$$

since, as one can prove easily, $(\beta - 1)S = 0$. Therefore

$$v_\zeta = \frac{1}{n} \sum_0^{n-1} \zeta^{-k} \pi(k)v.$$

This is a finite analogue of Fourier series.

For another example, let G be the symmetric group \mathfrak{S}_3 . It is generated by $\tau = (123)$ and $\sigma = (12)$, with relations

$$\tau^3 = 1, \sigma^2 = 1, \sigma\tau = \tau^2\sigma.$$

The group acts on \mathbb{C}^3 by permuting coordinates, and takes the plane $x + y + z = 0$ into itself. The generator τ has order three on this plane, which decomposes into eigenspaces with eigenvalues $e^{\pm 2\pi i/3}$. The generator σ swaps these two spaces, and it is easy to see that this representation is **irreducible**, which means here that there are no proper G -stable lines. There are two other irreducible representations of G , namely the trivial one and the character sgn taking τ to 1 and σ to -1 . If (π, V) is any representation of G , then V contains eigenspaces with respect to 1 and sgn . It is not hard to see that there exists a G -stable complement U to the sum of these two subspaces. This decomposes into a sum of eigenspaces for τ , and as we shall see eventually it turns out to be a sum of two-dimensional G -stable subspaces on which G acts as it does on $x + y + z = 0$. We can ask: How to prove this? Can we find a projection operator onto this complement? If V is the space $\mathbb{C}(G)$ of complex-valued functions on G , can we find a formula for a function f that decomposes it into components parametrized by these three irreducible representations?

The decomposition of a representation into irreducible subspaces is exactly the analogue of eigenspace decompositions.

COORDINATES AND MATRICES. If V is assigned a coordinate system, a representation becomes a homomorphism from G to $\text{GL}_n(\mathbb{C})$, since the image of each g in G is expressed in terms of a matrix. For reasons that should become clearer later on, I want to be somewhat pedantic about what this means.

It is important distinguish carefully between the groups $\text{GL}(V)$ and $\text{GL}_n(\mathbb{C})$. We can identify them only if we are given a choice of coordinates on V . The group $\text{GL}(V)$ acts naturally from the left on V , and there is a right action of $\text{GL}_n(\mathbb{C})$ on the bases of V . First of all, if (e_i) is a matrix whose columns are vectors e_i and a is a column of n coefficients (a_i) , then

$$e \cdot a = a_1 e_1 + \cdots + a_n e_n$$

is a linear combination of the e_i , hence a vector. Then if (a_i) is an $n \times n$ matrix with columns a_i , we can define

$$e \cdot (a_i) = (e \cdot a_i).$$

If (e_i) is the basis of the coordinate system and v a vector, then

$$v = \sum v_i e_i$$

where the v_i are the coordinates of v in this coordinate system. This may be expressed conveniently as a matrix equation

$$v = [e_1 \ \dots \ e_n] \begin{bmatrix} v_1 \\ \dots \\ v_n \end{bmatrix} = e \cdot v_e.$$

(The difference between left and right is not important here, since coefficients lie in a commutative field, although they become important if one is looking at vectors over arbitrary division algebras.) This is in keeping with the usual convention that vectors are represented as column matrices. What happens to coordinates if we change basis?

Let's look at an example. Suppose V has dimension 2, and a basis e_1, e_2 . Fix a new basis

$$\begin{aligned} f_1 &= 2e_2 \\ f_2 &= e_1 - e_2. \end{aligned}$$

This can be expressed by the single equation

$$[f_1 \ f_2] = [e_1 \ e_2] \begin{bmatrix} 0 & 1 \\ 2 & -1 \end{bmatrix}.$$

But we can also write

$$\begin{aligned} e_1 &= (1/2)f_1 + f_2 \\ e_2 &= (1/2)f_1. \end{aligned}$$

This means that

$$\begin{aligned} v &= x_1 e_1 + x_2 e_2 \\ &= x_1((1/2)f_1 + f_2) + x_2(1/2)f_1 \\ &= (x_1/2 + x_2/2)f_1 + x_1 f_2. \end{aligned}$$

This can be expressed in the matrix equation

$$v_f = \begin{bmatrix} 1/2 & 1/2 \\ 1 & 0 \end{bmatrix} v_e.$$

Very generally, if

$$f = eA$$

then

$$e = A^{-1}f, \quad v = ev_e = fA^{-1}v_e, \quad v_f = A^{-1}v_e.$$

Given a coordinate system, any linear transformation T is assigned a matrix M_e according to the recipe that the j -th column of M_e records the coordinates of the vector Te_j . This means that

$$Tv = e \cdot M_e \cdot v_e.$$

What happens to M_e if we change bases?

If (f_i) is another basis then

$$f = [f_1 \dots f_n] = [e_1 \dots e_n] \cdot A = e \cdot A, \quad e = f \cdot A^{-1}$$

for some invertible matrix A in $\text{GL}_n(\mathbb{C})$. I want to emphasize: the group $\text{GL}(V)$ acts naturally on vectors in V , but the group $\text{GL}_n(\mathbb{C})$ acts naturally on bases of V . The first action is naturally on the left, the second on the right. Hence we have

$$Tv = e \cdot M_e \cdot v_e = f \cdot A^{-1} M_e A \cdot v_f$$

which means:

1.1. Proposition. *Suppose e and f to be two bases of V with $eA = f$. If M_e is the matrix of a linear transformation in the coordinate associated to e , that associated to f is*

$$M_f = A^{-1} M_e A.$$

EQUIVALENT REPRESENTATIONS. For operators in $\pi(G)$ this reads

$$M_f(g) = A^{-1} M_e(g) A, \quad A M_f(g) = M_e(g) A$$

for all g in G . The ‘essential’ identity of a representation shouldn’t depend on a choice of coordinate system, so the two representations taking

$$g \mapsto M(g), \quad g \mapsto A^{-1} M(g) A$$

are to be considered equivalent. More generally, two representations (π_1, V_1) and (π_2, V_2) are to be considered **equivalent** or **isomorphic** if there exists an invertible linear map $T: V_1 \rightarrow V_2$ such that

$$\pi_1(g) = T^{-1} \pi_2(g) T, \quad T \pi_1(g) = \pi_2(g) T$$

for all g in G . *They are really the same representation, but with different names.*

More generally, the vector space $\text{Hom}_G(V_1, V_2)$ of G -**equivariant** linear maps from the space of one representation (π_1, V_1) to that of another one (π_2, V_2) is that of all linear maps

$$F: V_1 \rightarrow V_2 \text{ such that } F\pi_1(g) = \pi_2(g)F.$$

Some of the important questions one asks in this business are: *Can we classify, up to equivalence, all representations of G ? Is there a useful criterion for determining whether two given representations are the same? Can we describe the space $\text{Hom}_G(V_1, V_2)$ for two given representations?*

One class of representations arises naturally, those in which one is given a homomorphism from G to the group of permutations of a set X . If $V = \mathbb{C}(X)$ is the vector space of all complex-valued functions on X , we have a representation of G on V according to the formula

$$[L_g F](x) = F(g^{-1}(x)),$$

where $x \mapsto g(x)$ is the permutation of X associated to g . One needs the inverse in this formula in order to have

$$L_{g_1} L_{g_2} = L_{g_1 g_2},$$

as you can check. (The “ L ” here stands for “left”.) One of the most important examples is the case in which the product $G \times G$ acts on G — (g_1, g_2) takes

$$g \mapsto g_1 g g_2^{-1}.$$

From this we get two representations of G on $\mathbb{C}(G)$, the **left regular** and **right regular**:

$$\begin{aligned} L_g F(x) &= F(g^{-1}x) \\ R_g F(x) &= F(xg). \end{aligned}$$

An important special class of representations is made up of the **characters** of G , which are the one-dimensional representations, or homomorphisms from G to the multiplicative group \mathbb{C}^\times . Every group has at least one of these, taking every g to 1 (the **trivial** representation of G), and there may not be any others. (We shall see another use of the word ‘characters’ later on; the two uses are compatible, if not the same. One refers to characters of the group in the sense used here, and the other, introduced later on, to characters of a representation. For one-dimensional representations, the two uses coincide.)

If (π_1, V_1) and (π_2, V_2) are two representations of G , then we can make a new one π on the space $\text{Hom}(V_1, V_2)$ of all linear maps from V_1 to V_2 according to the formula

$$[\pi(g)F](v_1) = \pi_2(g)F(\pi_1(g^{-1})v_1).$$

That is to say, $\pi(g)F$ is to be a linear map from V_1 to V_2 , and is hence defined in terms of what it does to vectors in V_1 . I leave it as an exercise to show that this is indeed a representation—that is to say, $\pi(g_1 g_2) = \pi(g_1)\pi(g_2)$. The space $\text{Hom}_G(V_1, V_2)$ of G -equivariant linear maps from V_1 to V_2 is the subspace of G -invariant elements of $\text{Hom}(V_1, V_2)$.

Given coordinate systems on U and V , the space $\text{Hom}(U, V)$ may be identified with all matrices of size $\dim V \times \dim U$ (the columns represent vectors in V). Its dimension is therefore the product $\dim U \cdot \dim V$.

2. Decomposition

Suppose (π, V) to be a representation of G . A vector subspace $U \subseteq V$ is said to be G -**stable** or G -**invariant** if all operators $\pi(g)$ take U into itself. In that case, the restriction of the operators of $\pi(G)$ to U is said to be a **subrepresentation** of π . For example, the subspace of all vectors fixed by all $\pi(g)$ is G -stable, on which G acts trivially. If U is a G -stable subspace of V then there is also an associated representation on the quotient V/U . It is important to distinguish sub-objects from quotient objects, but with representations of finite groups this is not so apparent, because of the following theorem, which is one of the fundamental results in the subject:

2.1. Theorem. *If (π, V) is a representation of G and U is G -stable, then there exists at least one G -stable linear complement W of U .*

That is to say, both U and W are G -stable and $V = U \oplus W$. In this case π is said to be the **direct sum** of the representations on U and W .

Proof. I'll offer two of many possible proofs. The first offers us a tool that is generally useful to have at hand:

2.2. Lemma. *If (π, V) is a representation of G , there exists at least one positive definite Hermitian inner product $u \bullet v$ on V with respect to which the operators in G are unitary.*

I recall that an Hermitian inner product $u \bullet v$ is a map from $V \times V$ to \mathbb{C} such that

$$\begin{aligned} (u + v) \bullet w &= u \bullet w + v \bullet w \\ v \bullet u &= \overline{u \bullet v} \\ (cu) \bullet v &= c(u \bullet v) \\ u \bullet (cv) &= \overline{c}(u \bullet v). \end{aligned}$$

It is **positive definite** if

$$u \bullet u \geq 0 \quad \text{and} \quad u \bullet u = 0 \quad \text{if and only if} \quad u = 0.$$

An example is the inner product

$$x_1 \overline{y_1} + \cdots + x_n \overline{y_n}$$

on \mathbb{C}^n . The positive definite Hermitian inner products are an open convex cone in the space of all Hermitian inner products.

An important example is the natural Hermitian metric on $\mathbb{C}(G)$ defined by

$$f \bullet g = \int_G f(x) \overline{g(x)} dx.$$

I recall that a linear transformation T is unitary with respect to a Hermitian inner product $u \bullet v$ if $(Tu) \bullet (Tv) = u \bullet v$ for all u, v .

Proof. Start with any positive definite Hermitian product $u \diamond v$. Define the new one to be

$$u \bullet v = \int_G (\pi(g)u) \diamond (\pi(g)v) dg.$$

This is the first of many averaging formulas we shall encounter. ▣

Now for the first proof of the Theorem. Given $U \subseteq V$, choose a G -invariant Hermitian inner product on V and take W to be the subspace of V perpendicular to U . Q.E.D.

The second proof also uses the averaging formula. Given U , choose any linear complement W_\diamond to U in V , and let $P_\diamond: V \rightarrow U$ be the corresponding projection onto U , annihilating W_\diamond . Then

$$P(v) = \int_G \pi(g) P_\diamond(\pi(g^{-1})v) dg$$

is a G -equivariant projection. Let W be its kernel. ▮

A G -stable complement is not in general unique. For example, if G acts trivially on a vector space, any subspace is G -stable.

There is a useful diagrammatic way to formulate the existence of G -stable complements. An **exact sequence** of linear maps is one in which the kernel of each map coincides with the image of the preceding one. A **short exact sequence** is a sequence of maps

$$0 \longrightarrow U \xrightarrow{f} V \xrightarrow{g} W \longrightarrow 0$$

in which f is injective, g is surjective, and the kernel of g is the same as the image of f . A **splitting** of this exact sequence is a linear map φ back from W to V such that $g \circ \varphi = I$.

2.3. Corollary. *An exact sequence of representations of G always possesses a G -equivariant splitting.*

Proof. A G -stable complement W_\diamond to U in V projects isomorphically to W .

Or, one can just repeat the second proof above—start with a linear splitting, which always exists, and average it over G to make it G -equivariant. ▮

 *These results, in so far as they suggest that for finite groups there is no distinction between quotient and subrepresentations, are somewhat misleading. It is true that G -stable complements always exist, that there always exists a G -invariant positive definite Hermitian inner product, and that short exact sequences always possess a G -equivariant splitting, but the complements, the Hermitian inner product, and the splitting are not generally unique. Even when they are, the confusion of sub- and quotient representations can cause trouble.*

A representation is said to be **irreducible** if it possesses no G -stable subspaces other than itself and $\{0\}$.

2.4. Corollary. *If (π, V) is any representation of G , there exists a decomposition of V into irreducible summands*

$$V = V_1 \oplus \cdots \oplus V_m.$$

Proof. By induction on dimension. If $\dim_\pi = 1$, then π is irreducible, so we may suppose $\dim_\pi > 1$. If V itself is irreducible, we are done. Otherwise, it possesses a G -stable subspace U other than $\{0\}$ or V . Express $V = U \oplus W$, and apply induction. ▮

One can ask if such a decomposition is unique—if the subspaces in a decomposition into irreducible components are unique. The answer is negative, for trivial reasons, since for example if $G = \{1\}$ then any linear decomposition of a vector space into one-dimensional subspaces will also be an irreducible decomposition of a representation of G . However, we shall next see that each irreducible representation of G occurs with a certain multiplicity that is in fact uniquely determined. In addition, for each irreducible representation π of G there exists a unique G -stable subspace of V that (a) contains all copies of (π, V_π) in V and (b) is spanned by such copies.

2.5. Proposition. (Schur's Lemma) *Suppose π, ρ to be two irreducible representations of G . Then*

$$\dim \operatorname{Hom}_G(V_\pi, V_\rho) \cong \begin{cases} 1 & \text{if } \pi \cong \rho \\ 0 & \text{otherwise.} \end{cases}$$

In other words, if $\pi \cong \rho$ then all functions in $\operatorname{Hom}_G(V_\pi, V_\rho)$ are scalar multiples of some given isomorphism, and if π is not isomorphic to ρ all such functions are trivial. If π is actually equal to ρ —that is to say, if they act on the same space—the identity operator spans the space, but otherwise there is not necessarily one distinguished candidate.

Proof. I'll leave it as an exercise. ▮

There are several, more or less equivalent, results called ‘Schur’s Lemma’. Here is another one:

2.6. Corollary. *Suppose π to be an irreducible representation of G . Every endomorphism of V_π commuting with operators in $\pi(G)$ is a scalar multiplication.*

If π is a representation of G , let $(n\pi, V_\pi^n)$ be the direct sum of n copies of π , with G acting on V_π^n .

2.7. Corollary. *If π is irreducible and V is a direct sum of irreducible representations of G , then the number of times π occurs in this decomposition is*

$$\dim \text{Hom}_G(V_\pi, V).$$

Proof. If $V = \bigoplus V_i$ is the space of a direct sum of representations of G , then

$$\text{Hom}(V_\pi, V) \cong \bigoplus_i \text{Hom}(V_\pi, V_i), \quad \text{Hom}_G(V_\pi, V) \cong \bigoplus_i \text{Hom}_G(V_\pi, V_i),$$

Hence the Corollary follows from Schur’s Lemma. ▮

Yet another version:

2.8. Corollary. *Suppose π, ρ to be irreducible. Then for any F in $\text{Hom}(V_\pi, V_\rho)$*

$$\int_G \rho(g) \cdot F \cdot \pi(g)^{-1} dg = \begin{cases} 0 & \text{if } \pi \text{ is not isomorphic to } \rho \\ \text{some scalar} & \text{if } \pi = \rho. \end{cases}$$

Proof. Because the sum commutes with G . ▮

3. Duality

LINEAR DUALITY. For the moment, I shall work with vector spaces over an arbitrary field F . If V is any vector space F , its **linear dual** is the set \widehat{V} of all linear functions from V to F . For \widehat{v} in \widehat{V} and v in V , I’ll write the value of \widehat{v} at v as $\langle \widehat{v}, v \rangle$.

This \widehat{V} is again a vector space:

$$\begin{aligned} \langle \widehat{v}_1 + \widehat{v}_2, v \rangle &= \langle \widehat{v}_1, v \rangle + \langle \widehat{v}_2, v \rangle \\ \langle c\widehat{v}, v \rangle &= c \langle \widehat{v}, v \rangle. \end{aligned}$$

 *If we assign a coordinate system to V , then every u in V determines an element of \widehat{V} :*

$$\langle u, v \rangle = \sum u_i v_i,$$

but this is a grievous deception! In general, it is a very bad idea to identify V with \widehat{V} . The basic reason is that with duals things go backwards, as we shall now see.

If $f: U \rightarrow V$ is a linear map, then it induces a map backwards

$$\widehat{f}: \widehat{V} \rightarrow \widehat{U}$$

defined by

$$\langle \widehat{f}(\widehat{v}), u \rangle = \langle \widehat{v}, f(u) \rangle.$$

The map f is injective if and only if \widehat{f} is surjective, and f surjective if and only if \widehat{f} is injective. A better way to phrase this is to say that if

$$0 \longrightarrow U \xrightarrow{f} V \xrightarrow{g} W \longrightarrow 0$$

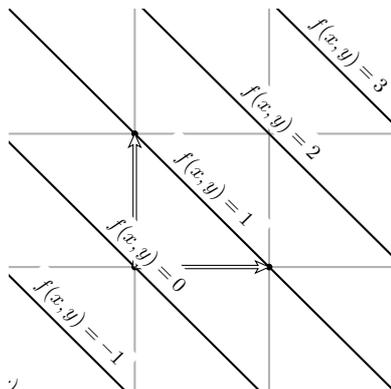
is an exact sequence of finite-dimensional vector spaces so is the dual sequence

$$0 \longrightarrow \widehat{W} \xrightarrow{\widehat{g}} \widehat{V} \xrightarrow{\widehat{f}} \widehat{U} \longrightarrow 0.$$

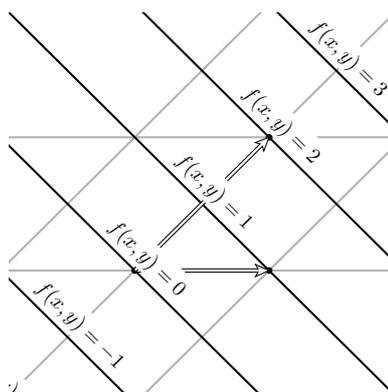
In current jargon $V \rightsquigarrow \widehat{V}$ is an **exact contravariant functor**.

The map \widehat{f} is often called the **transpose** of f . As we have seen already, the transpose of the composite $f \circ g$ is the composite $\widehat{g} \circ \widehat{f}$ of transposes, in the opposite order. The term ‘transpose’ might appear confusing, since it is also a term used for what we get when we flip a matrix around its NW-SE axis. But the two usages are essentially the same, if we write dual vectors as columns instead of rows.

Perhaps the way duality works will become clear if we consider in visual terms how things look. If V is a two-dimensional real plane. There are several different notions that often get confounded: (1) A **point** is a location. (2) A **vector** is a relative displacement. Thus if P and Q are points, then $Q - P$ is a vector. (3) A **linear function** is a function f on vectors that is linear: $f(au + bv) = af(u) + bf(v)$. If we are given a coordinate system on the n -dimensional vector space V , these all look somewhat alike, since each is specified by n coordinates. But it is important to distinguish them. They each have distinct intrinsic representations. *How do we represent points?* Well, you know, . . . by little dots. *How do we represent vectors?* By arrows. *How do we represent linear or affine functions?* By their **level sets**:



It is extremely important to realize that the three notions are all very different, although the distinctions are blurred if you have chosen a coordinate system. Points and vectors are equivalent, as long as we don't change the origin, because we can associate to every point the vector from the origin to P . Vectors and linear functions are equivalent, as long as we are working in an orthogonal coordinate system, with a fixed Euclidean metric. But vectors and linear functions change coordinates in different ways if you make a non-orthogonal change of coordinates. For example, suppose we choose a coordinate system in which the new basis is $f_1 = e_1$ and $f_2 = e_1 + e_2$. A function remains the same old function, *it assigns the same values to the same points*. But when the coordinate system changes, the *names* of points change, and the formula for a function changes—what used to be the line $x + y = 1$ in the old coordinates is now $x + 2y = 1$. In contrast, the point that used to be $(1, 1)$ is now $(0, 1)$.



The difference between vectors and linear functions is also taken into account by the standard convention that makes vectors out to be columns of coordinates, functions rows of them. So if φ is a linear function, v a vector, and e a basis, evaluation is in terms of a matrix product: then

$$\langle \varphi, v \rangle = \varphi_e \cdot v_e.$$

If we change coordinates from e to $f = e \cdot T$ then

$$\begin{aligned} v &= e \cdot v_e \\ &= f \cdot v_f \\ f &= e \cdot T, \quad f = e \cdot T^{-1} \\ v &= f \cdot T^{-1} \cdot v_e \\ v_f &= T^{-1} \cdot v_e. \end{aligned}$$

This is how the coordinates of vectors change. But

$$\begin{aligned} \langle \varphi, v \rangle &= \varphi_e \cdot v_e \\ &= \varphi_e \cdot T \cdot v_f \\ &= \varphi_f \cdot v_f, \end{aligned}$$

leading to:

3.1. Proposition. *If we change bases from e to $f = e \cdot M$, then the coordinates of vectors change according to the rule*

$$v_f = M^{-1} \cdot v_e$$

and those of linear functions according to the rule

$$\varphi_f = \varphi_e \cdot M.$$

DUAL REPRESENTATIONS. Revert to $F = \mathbb{C}$. If (π, V) is a representation of G , we want to define an associated action of G on \widehat{V} . It is determined by the condition that the canonical pairing of \widehat{V} with V is preserved by G :

$$\langle \widehat{\pi}(g)\widehat{v}, \pi(g)v \rangle = \langle \widehat{v}, v \rangle,$$

This forces the definition of $\widehat{\pi}(g)\widehat{v}$ by the formula

$$\langle \widehat{\pi}(g)\widehat{v}, v \rangle = \langle \widehat{v}, \pi(g^{-1})v \rangle.$$

In other words, $\widehat{\pi}(g) = \pi(\widehat{g^{-1}})$. Since $\widehat{f \circ g} = \widehat{g} \circ \widehat{f}$, we get what we expect— $\widehat{\pi}(g_1 g_2) = \widehat{\pi}(g_1) \widehat{\pi}(g_2)$. (Note also that this is a special case of the definition of the representation on the space $\text{Hom}(V_1, V_2)$.) For example, if χ is a character of G then $\widehat{\chi}$ is the inverse character $g \mapsto \chi(g)^{-1}$. For each g in G , the operator $\pi(g)$ has finite order, so V_π is the direct sum of eigenspaces on which g acts by a scalar $\chi(g)$. In the dual representation, $\widehat{\pi}(g)$ acts as the direct sum of the eigenspaces on which g acts by the inverses $\chi^{-1}(g)$.

Any v in V determines an element of the linear dual of \widehat{V} , taking

$$\widehat{v} \mapsto \langle \widehat{v}, v \rangle.$$

This is injective, so is an isomorphism as long as V is finite-dimensional. The dual of $\widehat{\pi}$ is then again π .

If we are given a basis (e_i) of V then there exists a unique **dual basis** (\widehat{e}_i) of \widehat{V} such that

$$\langle \widehat{e}_i, e_j \rangle = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

Assume U, V now to be two complex vector spaces. Assume given bases (e_i) of U , (f_j) of V with corresponding dual bases $(\widehat{e}_i), (\widehat{f}_j)$. If T is a linear transformation from U to V , we get from it a matrix $M_{i,j}$ according to the specification

$$T(e_j) = \sum_i M_{i,j} f_i.$$

The interpretation of the coefficients of a matrix is therefore

$$M_{i,j} = \langle \widehat{f}_i, T(e_j) \rangle.$$

CONJUGATE DUALITY. If π is a representation, its **conjugate** $\overline{\pi}$ is that in which g is taken to the matrix whose entries are the conjugates of those in $\pi(g)$. In general, the dual of π is not isomorphic to π , but Lemma 2.2 is that π is equivalent to the conjugate of its dual. For example, if χ is a non-trivial character of \mathbb{Z}/n with n a prime, then $\chi^{-1} \neq \chi$ but $\overline{\chi^{-1}} = \chi$.

4. Tensor products

This section can be best motivated by this question: *Suppose G and H are two finite groups. How do we classify the irreducible representations of $G \times H$ if we know all those of G and H ?* The answer is in terms of tensor products.

LINEAR TENSOR PRODUCTS. The tensor product of two vector spaces is an extremely useful—even indispensable—construction. In fact, tensor products can be defined in great generality, for modules over a ring. But tensor products are a bit tricky to define in great generality, and I shall confine myself to the simple case we shall require in this essay. In this section also, I allow U and V to be finite-dimensional vector spaces over an arbitrary field F .

What is the question to which tensor products are the answer? A **bilinear** map from $U \times V$ to any other vector space over F is a map

$$B: U \times V \rightarrow W$$

that is separately linear in each factor, or in other words satisfies the equations

$$\begin{aligned} B(u_1 + u_2, v) &= B(u_1, v) + B(u_2, v) \\ B(u, v_1 + v_2) &= B(u, v_1) + B(u, v_2) \\ B(cu, v) &= c B(u, v) \\ B(u, cv) &= c B(u, v). \end{aligned}$$

For example, the map from $F^n \times F^n$ to F taking

$$(x, y) \mapsto x_1y_1 + \cdots + x_ny_n$$

is bilinear. The set $\mathfrak{B}(U, V)$ of all bilinear maps from $U \times V$ to \mathbb{C} is a vector space of finite dimension. Given bases e_i of U and f_j of V there exists for each pair (i, j) a unique bilinear map $\varepsilon_{i,j}$ from $U \times V$ to \mathbb{C} which is 1 on (e_i, f_j) and 0 on all other basic pairs. Conversely, if B is any bilinear map then

$$B = \sum_{i,j} B(e_i, f_j) \varepsilon_{i,j}.$$

Therefore the $\varepsilon_{i,j}$ form a basis of the space $\mathfrak{B}(U, V)$. It has dimension $\dim_U \cdot \dim_V$.

What we are looking for is a vector space $U \otimes V$ which serves as the target of a universal bilinear map, in the sense that all bilinear maps from $U \times V$ to W factor through it. More precisely, we want an identification of

bilinear maps from $U \times V$ to W

with

linear maps from $U \otimes V$ to W .

Even more precisely, I require that there exist a *universal* bilinear map from $U \times V$ to $U \otimes V$, taking (u, v) to (say) $u \otimes v$, through which any bilinear map B from $U \times V$ to W factors—there must exist a linear map β from $U \otimes V$ to W filling in this diagram:

$$\begin{array}{ccc} U \times V & \xrightarrow{(u,v) \mapsto u \otimes v} & U \otimes V \\ & \searrow B & \swarrow \beta \\ & & W \end{array}$$

This is to be true in particular if $W = \mathbb{C}$, so we are to identify $\text{Hom}(U \otimes V, \mathbb{C})$ with $B(U, V)$. In effect, we are given the vector space dual to $U \otimes V$ and want to recover $U \otimes V$ itself. Now if X is any finite-dimensional vector space it may be defined in terms of its dual \widehat{X} , since it is the linear dual of \widehat{X} . So we define

$$U \otimes V = \text{the linear dual of the space } \mathfrak{B}(U, V) \text{ of all bilinear maps from } U \times V \text{ to } \mathbb{C}.$$

(⊗1) A pair u, v from U, V defines a unique element of $U \otimes V$:

$$u \otimes v: B \mapsto B(u, v).$$

I'll call $u \otimes v$ the **tensor** defined by the pair. The tensor product construction has these properties:

- (⊗2) The map $(u, v) \mapsto u \otimes v$ is bilinear in u and v .
- (⊗3) Every element of $U \otimes V$ is a linear combination of tensors $u \otimes v$.

Since $(e_i \otimes f_j)$ is the basis dual to $(\varepsilon_{i,j})$.

- (⊗4) If F is a bilinear map from $U \times V$ to the vector space W , there exists a unique linear map from $U \otimes V$ to W taking $u \otimes v$ to $F(u, v)$.

For $W = \mathbb{C}$, this is a matter of definition. But if we choose a basis (w_k) of W then any linear map into W is a sum of one-dimensional ones.

I define a **tensor product** of U, V to be a vector space T together with a bilinear map ι from $U \times V$ to T with the following property: *If B is a bilinear map from $U \times V$ to a vector space W , there exists a unique linear map from T to W making this diagram commutative:*

$$\begin{array}{ccc}
 U \times V & \xrightarrow{\iota} & T \\
 & \searrow B & \nearrow \beta \\
 & & W
 \end{array}$$

Thus the space $U \otimes V$ together with $(u, v) \mapsto u \otimes v$ is a tensor product. Tensor products are unique, completely characterized by the properties above, in a very strong sense:

4.1. Theorem. *Suppose U, V to be finite-dimensional vector spaces over F . If (T, ι) is a tensor product then there exists a unique map $f: U \otimes V \rightarrow T$ making this diagram commutative:*

$$\begin{array}{ccc}
 U \times V & \xrightarrow{(u,v) \mapsto u \otimes v} & U \otimes V \\
 & \searrow \iota & \nearrow f \\
 & & T
 \end{array}$$

It is an isomorphism.

Proof. Since the map ι from $U \times V$ is bilinear, property $(\otimes 4)$ assures us of a linear map f from $U \otimes V$ to T making the diagram commutative. I leave verification that it is an isomorphism as an exercise. ▣

- The map taking (\hat{v}, v) to $\langle \hat{v}, v \rangle v$ is bilinear, hence by $(\otimes 4)$ there exists a canonical linear map from $\hat{V} \otimes V$ to \mathbb{C} , which is G -invariant.
- A pair \hat{u} in \hat{U} and v in V determine a linear map from U to V , taking

$$w \mapsto \langle \hat{u}, w \rangle v$$

It is a simple map, taking every u in U to some multiple of v , thus highly non-invertible in dimensions greater than one. This is a bilinear map from $\hat{U} \times V$ to $\text{Hom}(U, V)$, thus inducing a linear map from $\hat{U} \otimes V$ to $\text{Hom}(U, V)$.

4.2. Proposition. *The canonical linear map from $\hat{U} \otimes V$ to $\text{Hom}(U, V)$ is an isomorphism.*

One can define similarly define multilinear maps and multi-tensor products

$$V_1 \otimes \cdots \otimes V_n .$$

If all $V_i = V$, we get

$$\otimes^n V .$$

The direct sum

$$\otimes^\bullet V = \bigoplus_{n \geq 0} \otimes^n V$$

is an associative algebra, called the **tensor algebra** of V .

⚠ *If U and V are arbitrary vector spaces, there still exists a tensor product satisfying properties $(\otimes 1)$ – $(\otimes 4)$, even though the construction I have used here certainly is not valid. In fact, one can construct tensor products of modules over arbitrary rings, but here too one needs a different idea to construct them.*

The construction of tensor products is functorial. Every linear map $f: U \rightarrow V$ gives rise to a unique sequence of linear maps

$$\otimes^n U \xrightarrow{\otimes^n f} \otimes^n V$$

compatible with tensor products and behaving well with respect to composition.

TENSOR PRODUCTS AND REPRESENTATIONS. Suppose π and ρ to be representations of G . Then for each g the map

$$(u, v) \mapsto \pi(g)u \otimes \rho(g)v$$

is bilinear, therefore inducing a map

$$[\pi(g) \otimes \rho(g)]: U \otimes V \rightarrow U \otimes V.$$

Thus we get a representation $\pi \otimes \rho$ of G on $U \otimes V$.

4.3. Proposition. *If π_1 and π_2 are irreducible representations of G_1 and G_2 , then so is $\pi_1 \otimes \pi_2$.*

If V is any irreducible representation of $G_1 \times G_2$, suppose π_1 to be irreducible representation of G_1 occurring in it. The natural map

$$V_1 \otimes \text{Hom}_{G_1}(V_1, V) \longrightarrow V$$

is an isomorphism. Therefore all irreducible representations of $G_1 \times G_2$ are tensor products of irreducible representations of its factors.

5. Trace and determinant

In this section, continue to let F be an arbitrary field.

• **Trace.** As long as V is a finite dimensional vector space, there exists according to Proposition 4.2 a canonical isomorphism of $\widehat{V} \otimes V$ with $\text{End}(V)$. It takes $\widehat{v} \otimes v$ to the linear map taking

$$u \mapsto \langle \widehat{v}, u \rangle v.$$

The tautologous pairing induces a canonical linear map from $\widehat{V} \otimes V$ to \mathbb{C} , hence by ‘transport of structure’ there exists also one from $\text{End}(V)$ to \mathbb{C} .

5.1. Proposition. *The canonical map from $\text{End}(V)$ to \mathbb{C} takes a matrix to the sum of its diagonal entries.*

This is called the trace of the matrix. In other words, the coordinate-free definition I have given agrees with the usual one.

Proof. Since the trace is linear, it suffices to prove this for a basis of $\text{End}(V)$. It is easy to see it is true for the map $v \mapsto \langle \widehat{e}_i, v \rangle e_j$. ▮

5.2. Corollary. *If S is invertible then*

$$\text{TR}(STS^{-1}) = \text{TR}(T).$$

Proof. It suffices to prove this for the linear map

$$T_{\widehat{v}, v}: u \mapsto \langle \widehat{v}, u \rangle v.$$

Its trace is $\langle \widehat{v}, v \rangle$. But then

$$STS^{-1}: u \mapsto \langle \widehat{v}, S^{-1}u \rangle Sv = \langle S\widehat{v}, u \rangle Sv = T_{S\widehat{v}, Sv}(u),$$

so STS^{-1} has trace $\langle S\widehat{v}, Sv \rangle = \langle \widehat{v}, v \rangle$. ▮

◊ *One curiosity is that the canonical pairing on $\widehat{V} \otimes V$ exists for spaces of infinite dimension, but the trace on $\text{End}(V)$ does not exist, except in certain circumstances. (For example, what could the trace of I be?) This is because the identification of $\widehat{V} \otimes V$ with $\text{End}(V)$ is no longer valid.*

One reason for the importance of the trace is that in many circumstances it counts fixed points of maps. Suppose X a finite set, τ a map from X to itself. Then it induces a map τ^* from $\mathbb{C}(X)$ to itself:

$$[\tau^*F](x) = F(\tau x).$$

This works like transposes: $(\sigma\tau)^* = \tau^*\sigma^*$. For example, if G acts by permutations $\pi(g)$ on X the map $g \mapsto \pi^*(g^{-1})$ is a representation of G .

5.3. Proposition. *If τ is a map from the set X to itself, the trace of τ acting on $\mathbb{C}(X)$ is the same as the number of points x of X such that $\tau(x) = x$.*

Proof. Take as basis of $\mathbb{C}(X)$ the functions ε_x with

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

For one of these, the claim is immediate. ▣

There is a useful generalization. A **complex vector bundle** over a finite set X can be defined in any one of a few equivalent ways. The simplest is to assign to each x in X a complex vector space E_x . Let E be the union of spaces E_x , and $p: E \rightarrow X$ the map that projects E_x onto x . A **section** of the bundle is a map s from X to E , taking each x to a vector $s(x)$ in E_x —i.e. such that $p(s(x)) = x$. The set $\Gamma(X, E)$ of sections of the bundle is then a complex vector space—if s and t are two sections then

$$[s + t](x) = s(x) + t(x), \quad [cs](x) = c \cdot s(x).$$

The set $\Gamma(X, E)$ is the direct sum of the spaces $E_x: s \mapsto (s(x))$.

The space $\mathbb{C}(X)$ is a ring:

$$[f + g](x) = f(x) + g(x), \quad [f \cdot g](x) = f(x)g(x).$$

The space $\Gamma(X, E)$ is a module over this ring:

$$[f \cdot s](x) = f(x)s(x).$$

The bundle E is completely determined by this module, in a strong sense. For each x in X , let \mathfrak{m}_x be the maximal ideal of $\mathbb{C}(X)$ consisting of functions vanishing at x . It is the kernel of the ring homomorphism e_x from $\mathbb{C}(X)$ to \mathbb{C} that takes f to $f(x)$. If M is any finitely generated module over $\mathbb{C}(X)$, define the space M_x to be the quotient of M by $\mathfrak{m}_x M$. The set $\{M_x\}$ is a vector bundle over X whose space of sections is M .

If the group G acts on the space X , a G -bundle over X is a vector bundle E over X together with an action of G on E compatible with its action on X —that is to say, there exists for every g a map $\lambda_g: E_x \mapsto E_{g(x)}$ and for all g, h in G we have $\lambda_{gh} = \lambda_g \lambda_h$. This gives rise to a representation of G on $\Gamma(X, E)$:

$$[L_g s](x) = \lambda_g s(g^{-1}(x)).$$

If G_x is the subgroup of G fixing x —i.e. G_x is the subgroup of g in G with $g(x) = x$ —then G_x takes E_x into itself, and we get a representation (L_x, E_x) of G_x .

5.4. Proposition. *Suppose that G acts on X and $E \mapsto X$ is a G -bundle. The trace of L_g on the space of sections of E is the sum*

$$\sum_{g(x)=x} \text{TR}(L_x(g)).$$

• **Determinant.** Determinants are also conjugation-invariant, and intimately related to traces. To explain exactly how this works I'll recall something about **exterior products**. Let I be the left ideal of $T(V) = \bigotimes^{\bullet} V$ generated by elements $v \otimes v$. It also contains all elements $u \otimes v + v \otimes u$ since

$$(u + v) \otimes (u + v) - u \otimes u - v \otimes v = u \otimes v + v \otimes u.$$

For example, if (e_i) is a basis of V , then $I \cap (V \otimes V)$ has basis

$$e_i \otimes e_i, \quad e_i \otimes e_j + e_j \otimes e_i \quad (i \neq j).$$

If the characteristic of F is not two, then we can write

$$u \otimes u = \frac{u \otimes u + u \otimes u}{2},$$

so it would have been equivalent to specify I as the ideal generated by the $u \otimes v + v \otimes u$, but this would not be satisfactory in characteristic 2.

Define the **exterior algebra** $\bigwedge^\bullet V$ of V to be $T(V)/I$. It is an algebra, since the left ideal I is also a right ideal:

$$\begin{aligned} v \otimes v \otimes u - u \otimes v \otimes v &= v \otimes v \otimes u + v \otimes u \otimes v - v \otimes u \otimes v - u \otimes v \otimes v \\ &= v \otimes (v \otimes u + u \otimes v) - (v \otimes u + u \otimes v) \otimes v. \end{aligned}$$

Define the product $u \wedge v$ to be the image of $u \otimes v$ in $\bigwedge^2 V$. Since $u \otimes v + v \otimes u$ lies in the ideal I , we have $u \wedge v = -v \wedge u$. If (e_i) is a basis of V , it is not difficult to figure out a basis for $I^n = I \cap \bigotimes^n V$. With very little work, you can then deduce:

5.5. Proposition. *If (e_i) is a basis of V , then a basis of $\bigwedge^m V$ is made up of all the*

$$e_{i_1} \wedge \dots \wedge e_{i_m}$$

with $i_1 < i_2 < \dots < i_m$.

Proof. It is an easy pleasant exercise to show that these span $\bigwedge^m V$, using the rule that if σ is a permutation of $[1, m]$ then

$$(5.6) \quad e_{\sigma(1)} \wedge \dots \wedge e_{\sigma(m)} = \text{sgn}(\sigma) e_{i_1} \wedge \dots \wedge e_{i_m}.$$

Here $\text{sgn}(\sigma)$ of a permutation σ is defined by

$$\prod_{1 \leq i < j \leq m} (x_{\sigma(i)} - x_{\sigma(j)}) = \text{sgn}(\sigma) \prod_{1 \leq i < j \leq m} (x_i - x_j).$$

It is more difficult to show they are linearly independent. Knowing a basis of I^n will do the trick. ▢

5.7. Corollary. *If V has dimension n , then $\bigwedge^n V$ has dimension one, spanned by*

$$e_{i_1} \wedge \dots \wedge e_{i_n},$$

and all $\bigwedge^m V$ with $m > n$ vanish.

To put ourselves in familiar territory, let $V = F^3$. Rather than the basis of Proposition 5.5, I choose basis

$$e_1 \wedge e_2, \quad e_3 \wedge e_1, \quad e_2 \wedge e_3$$

of $\bigwedge^2 V$. Thus if $u = \sum u_i e_i$ and $v = \sum v_i e_i$

$$\begin{aligned} u \wedge v &= (u_1 e_1 + u_2 e_2 + u_3 e_3) \wedge (v_1 e_1 + v_2 e_2 + v_3 e_3) \\ &= (u_1 v_2 - u_2 v_1) e_1 \wedge e_2 + (u_3 v_1 - u_1 v_3) e_3 \wedge e_1 + (u_2 v_3 - u_3 v_2) e_2 \wedge e_3. \end{aligned}$$

In other words, we are computing nothing other than the coefficients of cross-product.

If T is a linear transformation of V , it gives rise to maps $\otimes^m T$, which in turn induces maps $\wedge^m T$ of $\wedge^m V$. For example, suppose $V = F^2$, with basis e_1, e_2 , and T with the matrix

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}.$$

It takes

$$\begin{aligned} e_1 &\mapsto ae_1 + ce_2 \\ e_2 &\mapsto be_1 + de_2 \\ e_1 \otimes e_2 &\mapsto (ae_1 + ce_2) \otimes (be_1 + de_2) \\ e_1 \wedge e_2 &\mapsto ab(e_1 \wedge e_1) + bc(e_2 \wedge e_1) + ad(e_1 \wedge e_2) + cd(e_2 \wedge e_2) \\ &= (ad - bc)(e_1 \wedge e_2). \end{aligned}$$

More generally, if $n = \dim V$ then $\wedge^n T$ is a scalar multiplication, and

5.8. Proposition. *If T is in $\text{End}(V)$ and $n = \dim V$, then*

$$\wedge^n T = \det(T) I.$$

Proof. The proof depends on how the determinant is defined. Here, I take its definition to be

$$\sum_{\sigma} \text{sgn}(\sigma) t_{i, \sigma(i)}$$

with the sum overall permutations σ . From this, I deduce the Proposition from (5.6). ▢

This Proposition is a special case of a result that tells how to find the matrix of any of the $\wedge^m T$ in terms of determinants of square sub-matrices of the matrix of T . The point is to evaluate $f_1 \wedge \dots \wedge f_m$ where $f_i = Te_i$.

Since $\wedge^m S \circ \wedge^m T = \wedge^m (S \circ T)$, this explains naturally the product formula for determinants. This in turn shows that the determinant is certainly conjugation-invariant. The trace and determinant are related via the characteristic polynomial:

$$\det(T - xI) = \sum_{m=0}^n (-1)^{n-m} x^m \text{TR} \wedge^{n-m} T,$$

if the dimension of V is n .

 *If the field F has characteristic other than two, exterior products play the same role for alternating bilinear maps that tensor products do for multi-linear ones. It is because of problems in characteristic two that I define them in terms of relations $u \wedge u = 0$ rather than $u \wedge v = -v \wedge u$.*

6. Schur orthogonality

Suppose (π, V) to be a representation of G . Then for every v in V , \widehat{v} in \widehat{V} I define the corresponding **matrix coefficient** to be the function on G :

$$\Phi_{\widehat{v}, v}: G \rightarrow \mathbb{C}, \quad g \mapsto \langle \widehat{v}, \pi(g)v \rangle.$$

The product $G \times G$ acts on G , hence on the space $C[G]$ of \mathbb{C} -valued functions on G :

$$f(x) \mapsto f(g_1^{-1} x g_2).$$

The matrix coefficient map taking (\widehat{v}, v) to the function $\langle \widehat{v}, \pi(g)v \rangle$ is bilinear, hence induces a $G \times G$ -equivariant map from $\widehat{V} \otimes V$ to $\mathbb{C}(G)$. If π is irreducible, then $\widehat{\pi} \otimes \pi$ is also an irreducible representation of $G \times G$, and this is an embedding.

The map from $V \otimes \widehat{V}$ to $C[G]$ is $G \times G$ -equivariant:

$$\Phi_{\widehat{\pi}(g_1)(\widehat{v}), \pi(g_2)v}(x) = \langle \widehat{\pi}(g_1)\widehat{v}, \pi(x)\pi(g_2v) \rangle = \langle \widehat{v}, \pi(g_1^{-1}xg_2v) \rangle.$$

6.1. Proposition. (Schur orthogonality relations) *Suppose (π, V) , (ρ, U) to be irreducible representations of G , with u in U , v in V , \widehat{u} in \widehat{V} , \widehat{v} in \widehat{V} . Then*

$$\begin{aligned} \int_G \Phi_{\widehat{v}, v}(g)\Phi_{\widehat{u}, u}(g^{-1}) dg &= \int_G \langle \widehat{v}, \pi(x)v \rangle \langle \widehat{u}, \rho(x^{-1})u \rangle dx \\ &= 0 && \text{if } \pi \text{ is not isomorphic to } \rho \\ &= \frac{1}{\dim_\pi} \langle \widehat{v}, u \rangle \langle \widehat{u}, v \rangle && \text{if } \pi = \rho. \end{aligned}$$

Proof. Let

$$B_{\widehat{v}, v, \widehat{u}, u} = \int_G \Phi_{\widehat{v}, v}(g)\Phi_{\widehat{u}, u}(g^{-1}) dg.$$

Fix $\widehat{v} = \widehat{v}_0$, $u = u_0$ for the moment. The sum above gives a pairing of \widehat{U} and V with values in \mathbb{C} :

$$B(\widehat{u}, v) = B_{\widehat{v}_0, v, \widehat{u}, u_0} = \int_G \langle \widehat{v}_0, \pi(x)v \rangle \langle \widehat{u}, \rho(x^{-1})u_0 \rangle dx.$$

Then

$$\begin{aligned} B(\widehat{\rho}(g)\widehat{u}, \pi(g)v) &= \int_G \langle \widehat{u}, \pi(g^{-1}x)u_0 \rangle \langle \widehat{v}_0, \rho(x^{-1}g)v \rangle dx \\ &= \int_G \langle \widehat{u}, \pi(y)u_0 \rangle \langle \widehat{v}_0, \rho(y^{-1})v \rangle dy \\ &= B(\widehat{u}, v), \end{aligned}$$

where I have changed variables $y = g^{-1}x$ in the ‘integral’. The pairing is thus G -invariant, which means that it gives a G -map from \widehat{U} to the dual of V . If π and ρ are not isomorphic, this map must be 0. If π and ρ are isomorphic, we may as well assume them to be the same, and we must have, by Schur’s Lemma, $B(\widehat{u}, v) = C\langle \widehat{u}, v \rangle$ for some C . Now we fix $\widehat{u} = \widehat{u}_0$, $v = v_0$, with $\langle \widehat{u}_0, v_0 \rangle \neq 0$, but let the other two vary. This gives us

$$\int_G \langle \widehat{v}, \pi(x)v \rangle \langle \widehat{u}, \pi(x^{-1})u \rangle dx = \Gamma \cdot \langle \widehat{v}, u \rangle \langle \widehat{u}, v \rangle$$

for some constant Γ independent of all the choices of \widehat{v} etc. Now we choose a basis (e_i) for $U = V$, set

$$u = e_j, v = e_k, \widehat{u} = \widehat{e}_i, \widehat{v} = \widehat{e}_j,$$

and then sum $B_{\widehat{e}_i, e_j, \widehat{e}_j, e_k}(x)$ over j . The result is

$$\sum_j \langle \widehat{e}_i, \pi(x)e_j \rangle \langle \widehat{e}_j, \pi(x^{-1})e_k \rangle$$

which is the (i, k) -matrix entry of the product of $\pi(x)$ and $\pi(x^{-1})$ —it does not depend on x . So when we sum over x we get on the one hand $\langle \widehat{e}_i, e_k \rangle$ and on the other $\dim_\pi \cdot \Gamma \cdot \langle \widehat{e}_i, e_k \rangle$. Therefore $\Gamma = 1/\dim_\pi$. ■

I recall that the space $\mathbb{C}(G)$ possesses a $G \times G$ -invariant positive definite Hermitian inner product

$$f_1 \bullet f_2 = \int_G f_1(g)\overline{f_2(g)} dg.$$

6.2. Corollary. *The spaces of matrix coefficients of non-isomorphic representations of G are orthogonal with respect to this inner product.*

Proof. Since $\pi(g)$ has finite order, it is diagonalizable. This implies that the eigenvalues of $\pi(g^{-1})$ are the conjugates of those of $\pi(g)$, hence the trace of $\pi(g^{-1})$ is the conjugate of the trace of $\pi(g)$. ■

7. Characters

If π is a representation of G , its **character** χ_π is the function

$$g \longmapsto \text{TR } \pi(g).$$

Since TR is conjugation-invariant, this depends only on the equivalence class of π —i.e. isomorphic representations have the same character. As we shall see, the converse is also true: representations with the same character are isomorphic.

There are several properties of the character that are straightforward to verify:

$$\begin{aligned} \chi_\pi(1) &= \dim_\pi \\ \chi_\pi(ghg^{-1}) &= \chi_\pi(h) \\ \chi_{\pi \oplus \rho} &= \chi_\pi + \chi_\rho \\ \chi_{\pi \otimes \rho} &= \chi_\pi \cdot \chi_\rho \\ \chi_\pi(g^{-1}) &= \overline{\chi_{\bar{\pi}}(g)} \\ &= \overline{\chi_\pi(g)}. \end{aligned}$$

Only the last two are worth saying more about—the last is valid because g has finite order, so its trace is a sum of roots of unity ζ with $\bar{\zeta} = \zeta^{-1}$. The other does not depend on this, and remains valid whenever it makes sense.

7.1. Proposition. *If χ and ρ are characters of irreducible representations of G , then*

$$\chi_\pi \bullet \chi_\rho = \begin{cases} 1 & \text{if } \pi \cong \rho \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Assign coordinates to the spaces of the representations. What we have to evaluate is

$$\int_G \left(\sum_i \langle \hat{e}_i, \pi(g)e_i \rangle \right) \left(\sum_j \langle \hat{f}_j, \rho(g^{-1})f_j \rangle \right) dg = \sum_{i,j} \int_G \langle \hat{e}_i, \pi(g)e_i \rangle \langle \hat{f}_j, \rho(g^{-1})f_j \rangle dg$$

which by Schur's orthogonality is equal to 0 if π and ρ are not isomorphic, and

$$\frac{1}{\dim_\pi} \sum_{i,j} \langle \hat{e}_i, e_j \rangle \langle \hat{e}_j, e_i \rangle = \frac{\dim_\pi}{\dim_\pi} = 1$$

if $\pi = \rho$. ▮

7.2. Corollary. *Suppose $\pi = \bigoplus \pi_i$ to be an irreducible decomposition of π and ρ is an irreducible representation of G . Let $N_{\pi,\rho}$ be the number of i with $\pi|_{V_i} \cong \rho$. Then*

$$N_{\pi,\rho} = \chi_\pi \bullet \chi_\rho = \dim \text{Hom}_G(V_\rho, V_\pi) = \dim \text{Hom}_G(V_\pi, V_\rho).$$

When one, hence all, of these numbers are positive, I'll say that π **contains** ρ .

⚠ *The Corollary must not be misinterpreted. It is true that the dimensions of $\text{Hom}_G(V_\rho, V_\pi)$ and $\text{Hom}_G(V_\pi, V_\rho)$ are the same, but there is not in general any natural identity of the two. If we are given a G -equivariant embedding of V_ρ into V_π , there does exist a G -equivariant projection from V_π onto the image, but it depends on a choice of G -stable complement of that image, which will in general not be unique. Even if it is unique, getting from an embedding to a projection may be awkward.*

7.3. Corollary. *Two representations are isomorphic if and only if they have the same characters.*

8. Central functions

The space $Z[G]$ is that of all conjugation-invariant functions on G —i.e. f is in $Z[G]$ if and only if $f(gxg^{-1}) = f(x)$ for all x, g . These are also called **central functions**, for reasons that will appear in time. (The “ Z ” here is from German “Zentrum”.) This is not to be confused with Z_G , which is the center of G , the subgroup of elements of G commuting with all of G . For any f in $\mathbb{C}(G)$ one can define an operator

$$\pi(f)v = \int_G f(x)\pi(x)v \, dx.$$

If f is in $Z[G]$ then $\pi(f)$ commutes with G on the space of any representation of G :

$$\pi(g)\pi(f)\pi(g)^{-1} = \int_G f(x)\pi(gxg^{-1}) \, dx = \int_G f(y)\pi(y) \, dy = \pi(f).$$

Hence by Schur’s Lemma it acts by a scalar λ_π if π is irreducible. Explicitly:

8.1. Lemma. *If f is a central function in $\mathbb{C}(G)$ then*

$$\lambda_\pi(f) = \frac{1}{\dim_\pi} \cdot (f \bullet \chi_\pi).$$

8.2. Proposition. *The characters are an orthonormal basis of conjugation-invariant functions.*

Proof. We need to prove that the orthogonal complement is trivial. Suppose f to be a central function such that

$$\langle f, \chi \rangle = 0$$

for all characters χ . Then $\lambda_\pi(\overline{f}) = 0$ for all irreducible π , hence for all π , in particular the regular representation. But then set $F = \varepsilon_1$, which is 1 at 1 and 0 everywhere else.

$$0 = R_{\overline{f}} \varepsilon_1(x) = \int_G \overline{f(g)} \varepsilon_1(xg) \, dg = \frac{1}{|G|} \overline{f(x^{-1})}$$

so $f(x^{-1}) = 0$ for all x . ▣

The character χ_ρ is in $Z[G]$. What can we say about $\pi(\chi_\rho)$ if ρ is irreducible? It acts as a scalar on each π , which we have a formula for. In fact, it acts as 1 on the representation ρ , 0 on every other. This in fact answers a fundamental question. Consider V^G , the space of vectors fixed by G . Let \mathcal{P}_1 be the operator on V defined as

$$\mathcal{P}_1: v \mapsto \int_G \pi(g)v \, dg.$$

8.3. Proposition. *A vector v lies in V^G if and only if $\mathcal{P}_1 v = v$.*

Something similar occurs for every irreducible representation of G .

8.4. Proposition. *The map*

$$\mathcal{P}_\rho = \dim_\rho \int_G \chi_\rho(g^{-1})\pi(g) \, dg$$

is a projection onto the subspace of V_π decomposing into copies of V_ρ .

9. Conjugacy classes

A more obvious basis for $Z[G]$ is made up of the characteristic functions of conjugacy classes \mathcal{O}_x , the orbits of G acting by conjugation. One of these may be expanded in terms of the irreducible characters:

$$\begin{aligned}\mathcal{O}_x(g) &= \sum_{\chi} \langle \chi, \mathcal{O}_x \rangle \chi(g) \\ &= \sum_{\chi} \chi(g) \int_G \mathcal{O}_x(y) \chi(y^{-1}) dy \\ &= \frac{|\mathcal{O}_x|}{|G|} \sum_{\chi} \chi(x^{-1}) \chi(g) \\ \mathcal{O}_x &= \sum_{\chi} \left(\frac{|\mathcal{O}_x|}{|G|} \chi(x^{-1}) \right) \chi.\end{aligned}$$

Since the \mathcal{O}_x may be expressed in terms of the characters, and vice-versa::

9.1. Proposition. *The number of irreducible representations of G is the same as the number of its conjugacy classes.*

Example. Let $G = \mathfrak{S}_3$, the group of permutations of 3 elements. Here is a list of all elements, written in cycle form:

$$\begin{aligned}(1)(2)(3) \\ (12)(3), (13)(2), (23)(1) \\ (123), (321).\end{aligned}$$

The conjugacy classes of any \mathfrak{S}_n correspond to cycle counts, hence here to partitions of 3. For example

$$(12)(123)(12)^{-1} = (213) = (321)$$

Hence there are three conjugacy classes, therefore three irreducible representations. One is the trivial representation. Another is the **sign** representation, which takes a permutation to the sign of its effect on

$$(x_1 - x_2)(x_2 - x_3)(x_3 - x_1)$$

It takes g to $(-1)^n$ if n is the number of even-length cycles in g . For example $(12) \mapsto -1$, $(123) \mapsto 1$.

What's the third? If it has dimension n then $1^2 + 1^2 + n^2 = 6$, so $n = 2$. In fact you can see this representation rather directly. The group \mathfrak{S}_3 acts on \mathbb{C}^3 by permuting coordinates, and leaves fixed the diagonal line $\Delta = \{(x, x, x)\}$. It also leaves invariant the plane $x_1 + x_2 + x_3 = 0$ perpendicular to Δ . On this the cycles act by rotations, the transpositions by reflection, and the representation is irreducible.

 *The number of irreducible representations of G and the number of its conjugacy classes is the same, but there is no 'natural' way to match these two sets. One simple example is \mathbb{Z}/N , Here the irreducible representations are the characters $n \mapsto \zeta_N^n$ where $\zeta_N = e^{2\pi i/N}$, but there is no canonical way to identify \mathbb{Z}/N with $\{\zeta_N^n\}$.*

10. Induced representations

There is one simple idea that can be used to construct representations of a group G . Suppose G to contain a proper subgroup H . Given an irreducible representation σ of H , what can we say about representations of G whose restriction to H contains σ ? We shall see that in some sense there is a universal representation of G whose restriction to H contains U .

I recall that what ‘contains’ means is explained in the remark just following Corollary 7.2. The remarks there about confusing $\text{Hom}_G(V_\pi, V_\rho)$ and $\text{Hom}_G(V_\rho, V_\pi)$ will become significant here.

We can get some idea of what’s to come by looking at the simplest case, with σ the trivial representation of H . So suppose (π, V) to be an irreducible representation of G ‘containing’ the trivial representation of H . According to Corollary 7.2 this means that either (a) we are given an H -equivariant linear map $\varphi: \mathbb{C} \rightarrow V_\pi$ or (b) we are given an H -equivariant map φ from V_π to \mathbb{C} . Somewhat but not totally arbitrarily, I’ll generally choose interpretation (b), but in truth it is not easy to separate the two notions. Our hypothesis means that $\varphi(\pi(h)v) = \varphi(v)$ for all h in H . I am going to derive from this datum an embedding of π into the space $\mathbb{C}(H \setminus G)$ of functions on $H \setminus G$. This is simple, if at first sight mysterious. For every v in V define the function Φ_v :

$$\Phi_v(g) = \varphi(\pi(g)v).$$

This is a function on $H \setminus G$, since

$$\Phi_v(hg) = \varphi(\pi(hg)v) = \varphi(\pi(g)v) = \Phi_v(g).$$

It is easy to check that $\Phi_{\pi(g)v} = R_g \Phi_v$. So we have embedded the representation π into the regular representation of G on $\mathbb{C}(H \setminus G)$. Conversely, if we are given an embedding $v \mapsto \Phi_v$ of π into $\mathbb{C}(H \setminus G)$ then $v \mapsto \Phi_v \mapsto \Phi_v(1)$ is a map in $\text{Hom}_H(\pi, \mathbb{C})$. In summary, embeddings of π into $\mathbb{C}[H \setminus G]$ may be identified with H -equivariant maps from π to \mathbb{C} . Roughly speaking, $\mathbb{C}[H \setminus G]$ contains π as often as π contains \mathbb{C} , but two meanings for ‘contain’ are being used here.)

We deduce that *the irreducible representation π contains the trivial representation of H (in the sense I have interpreted this) if and only if it can be embedded into $\mathbb{C}(H \setminus G)$* . We shall see later how to decompose $\mathbb{C}(H \setminus G)$ into irreducible components. For the moment, my goal is to generalize what we have just seen to representations of H other than the trivial one.

Given (σ, U) of H , define the representation of G induced by σ to be the right regular representation of G on

$$\text{Ind}(\sigma | H, G) = \{f : G \rightarrow U \mid f(hg) = \sigma(h)f(g) \text{ for all } h \in H, g \in G\}.$$

If σ is the trivial representation of H , this is the right regular representation of G on $\mathbb{C}(H \setminus G)$.

We can picture this space by decomposing it into $|H \setminus G|$ copies of U . For each coset x let U_x be the subspace of functions in $\text{Ind}(\sigma)$ that have support on x . The space U_x is the fibre of a vector bundle over X . The fibre U_1 is the space of functions f on H such that $f(hg) = \sigma(h)f(g)$, and $f \mapsto f(1)$ identifies it with U . Similarly if g is any element of the coset x the $f \mapsto f(g)$ identifies U_x as a representation of H (but not usually isomorphic to σ a representation of H). In effect, the representation of G on a certain vector bundle associated to σ .

Remark. An alternative definition of the induced representation would be

$$\mathbb{C}[G] \otimes_{\mathbb{C}[H]} U.$$

This is the definition of Serre’s book. I prefer the one made here, for various reasons I’ll not specify.

We can now apply Proposition 5.4 to computing the character of $\text{Ind}(\sigma)$. It tells us:

10.1. Proposition. The trace of R_g on $\text{Ind}(\sigma | H, G)$ is

$$\sum_{\substack{x \in H \setminus G \\ xg=x}} \chi_\sigma(xgx^{-1}).$$

Note that $xg = x$ for x in $H \setminus G$ if and only if $xgx^{-1} \in H$, so to speak. The condition makes sense since the conjugacy class of xgx^{-1} depends only on Hx .

The most important fact about induced representations is **Frobenius reciprocity**. There are several versions.

10.2. Proposition. (Frobenius reciprocity I) *If σ is an irreducible representation of H and π is an irreducible representation of G , then π occurs as often in $\text{Ind}(\sigma)$ as σ occurs in the restriction of π to H .*

Proof. This amounts to showing that

$$\begin{aligned} \chi_\pi \bullet_G \chi_{\text{Ind}} &= \chi_\pi \bullet_H \chi_\sigma \\ \chi_\pi \bullet_G \chi_{\text{Ind}} &= \int_G \chi_\pi(g^{-1}) \sum_{\substack{x \in H \setminus G \\ xgx^{-1} \in H}} \chi_\sigma(xgx^{-1}) dg \end{aligned}$$

This is

$$\begin{aligned} \frac{1}{|G|} \sum_g \chi_\pi(g^{-1}) \sum_{\substack{x \in H \setminus G \\ xgx^{-1} \in H}} \chi_\sigma(xgx^{-1}) &= \frac{1}{|G|} \frac{1}{|H|} \sum_{x \in G} \sum_h \chi_\pi(xh^{-1}x^{-1}) \chi_\sigma(h) \\ &= \frac{1}{|H|} \sum_h \chi_\pi(h^{-1}) \chi_\sigma(h). \quad \color{orange}{\blacksquare} \end{aligned}$$

There is another, in some ways more enlightening way to see it. Suppose the restriction of (π, V) to H does contain (σ, U) . This means that we have a projection from V to U commuting with H . On the other hand, if π occurs in $\text{Ind}(\sigma)$ we have an embedding of V_π into $\text{Ind}(\sigma)$ commuting with G . We may ask two natural questions suggested by Frobenius reciprocity: (1) Given a G -embedding $V \hookrightarrow \text{Ind}(\sigma)$, what is the associated projection onto U ? (2) Given an H -projection $V_\pi \rightarrow U$, what is the corresponding G -embedding of V_π into $\text{Ind}(\sigma)$? These are both easily answered by a different proof of a generalization of Frobenius reciprocity.

No longer assume π or σ to be irreducible. Define

$$\Omega_\sigma: \text{Ind}(\sigma) \rightarrow U, \quad F \mapsto F(1),$$

which commutes with H .

10.3. Theorem. (Frobenius reciprocity II) *If π is a representation of G and σ is one of H , then $F \mapsto \Omega \circ F$ is a linear isomorphism*

$$\text{Hom}_G(V_\pi, \text{Ind}(\sigma | H, G)) \cong \text{Hom}_H(V_\pi, U_\sigma).$$

Proof. We need only to specify an inverse—given φ we take

$$\Phi_v(g) = \varphi(\pi(g)v). \quad \color{orange}{\blacksquare}$$

This illuminates a result we have seen before. If $H = \{1\}$ then $\text{Ind}(\mathbb{C} | H, G) = \mathbb{C}(G)$, and Frobenius reciprocity tells us an irreducible representation π occurs \dim_π times. But it tells us more—it gives us those \dim_π copies, parametrized by the dual space $\text{Hom}_\mathbb{C}(V_\pi, \mathbb{C})$. In fact, the product $G \times G$ acts on $\mathbb{C}(G)$ on left and right, and we have the **matrix coefficient** $G \times G$ -map

$$\widehat{V} \otimes V \rightarrow \mathbb{C}(G), \quad \widehat{v} \times v \mapsto \Phi_{\widehat{v}, v}(g) = \langle \widehat{v}, \pi(g)v \rangle,$$

where $\langle \widehat{v}, v \rangle$ is the value of the linear function \widehat{v} on v .

10.4. Corollary. *Every irreducible representation of H is contained in some irreducible representation of G .*

Proof. Given σ , take π to be an irreducible subrepresentation of $\text{Ind}(\sigma)$. ▣

Suppose H and K to be two subgroups of G . Suppose (ρ, U) to be a representation of H , (σ, V) one of K . How do we describe

$$\text{Hom}_G(\text{Ind}(\rho|_H, G), \text{Ind}(\sigma|_K, G)) ?$$

There are two ways to answer this question.

If f lies in $C(G, U)$ then

$$\bar{f}(g) = \sum_{x \in H} \rho^{-1}(x) f(xg)$$

lies in $\text{Ind}(\rho|_H, G)$, and $f \mapsto \bar{f}$ is a G -equivariant surjection. Given a G -equivariant map

$$\Phi: \text{Ind}(\rho|_H, G) \rightarrow \text{Ind}(\sigma|_K, G),$$

consider the composite map

$$f \mapsto \Omega_\sigma(\Phi(\bar{f}))$$

from $C(G, U)$ to V . It is necessarily of the form

$$\varphi \mapsto \sum_G \langle F(g), \varphi(g) \rangle$$

for some $F = F_\Phi$ with values in $\text{Hom}(U, V)$.

10.5. Proposition. *The map taking Φ to F_Φ is an isomorphism of*

$$\text{Hom}_G(\text{Ind}(\rho|_H, G), \text{Ind}(\sigma|_K, G))$$

with the space of all F in $C(G, \text{Hom}(U, V))$ such that

$$F(hgk) = \sigma(k)^{-1} F(g) \rho(h)^{-1}.$$

Proof. • On the one hand, the defining property of F is

$$\sum_G \langle F(g), \varphi(g) \rangle = \Omega_\sigma(\Phi(\bar{f}))$$

for all f in $C(G, U)$. On the other

$$\begin{aligned} \overline{R_g f} &= R_g \bar{f} \\ \Phi(R_g f) &= R_g \Phi(\bar{f}). \end{aligned}$$

Therefore

$$\begin{aligned} \sum_G \langle F(g), f(gk) \rangle &= \sum_G \langle F(gk^{-1}), f(g) \rangle \\ &= \sum_G \langle F(g), [R_k f](g) \rangle \\ &= \Omega_\sigma(\Phi(\overline{R_k f})) \\ &= \Omega_\sigma(\Phi(R_k \bar{f})) \\ &= \Omega_\sigma(R_k \Phi(\bar{f})) \\ &= \sigma(k) \Omega_\sigma(\Phi(\bar{f})) \\ &= \sigma(k) \sum_G \langle F(g), f(g) \rangle \end{aligned}$$

so that

$$\sigma(k)^{-1} \sum_G \langle F(gk^{-1}), f(g) \rangle = \sum_G \langle F(g), f(g) \rangle$$

for all f in $C(G, U)$. Hence

$$\sigma(k)F(gk) = F(g), \quad F(gk) = \sigma(k)^{-1}F(g)$$

for all k in K .

• What about $F(hg)$ for h in H ? We have

$$\begin{aligned} \sum_G \langle F(hg), f(g) \rangle &= \sum_G \langle F(g), [L_h f](g) \rangle \\ &= \Omega(\Phi(\overline{L_h f})) \end{aligned}$$

and also

$$\begin{aligned} \overline{L_h f}(g) &= \sum_{x \in H} \rho^{-1}(x) L_h f(xg) \\ &= \sum_{x \in H} \rho^{-1}(x) f(h^{-1}xg) \\ &= \sum_{y \in H} \rho^{-1}(hy) f(yg) \\ &= \sum_{y \in H} \rho^{-1}(y) \rho^{-1}(h) f(yg) \\ &= \overline{\rho^{-1}(h) f}(g) \end{aligned}$$

so that

$$F(hg) = F(g)\rho^{-1}(h)$$

for all h in H . This proves the necessity of the condition on F .

• It is easy to check that, conversely, these conditions guarantee that Φ is well defined and G -equivariant. ▣

Another approach to the same question depends on Frobenius reciprocity. This tells us that

$$\text{Hom}_G(\text{Ind}(\rho | H, G), \text{Ind}(\sigma | K, G)) \cong \text{Hom}_K(\text{Ind}(\rho | H, G), V)$$

so determining the space of G -equivariant maps from $\text{Ind}(\rho | H, G)$ to $\text{Ind}(\sigma | K, G)$ is equivalent to determining the restriction of $\text{Ind}(\rho | H, G)$ to K .

As a representation of K , the induced representation $\text{Ind}(\rho | H, G)$ is equal to the direct sum of the subspaces of functions with support on the double cosets HgK . If W is (for the moment) the space of all functions f in $C(HgK, U)$ such that $f(hx) = \rho(h)f(x)$ for all h in H , then restriction to K induces an isomorphism of W with the induced representation $\text{Ind}(\rho^g | g^{-1}Hg \cap K, K)$, where

$$\rho^g(h) = \rho(hg^{-1}).$$

10.6. Proposition. *The restriction of $\text{Ind}(\rho | H, G)$ to K is isomorphic to the direct sum over all double cosets HgK of the induced representations $\text{Ind}(\rho^g | g^{-1}Hg \cap K, K)$.*

By Frobenius reciprocity, the representation $\text{Ind}(\rho | H, G)$ is irreducible if and only if ρ occurs exactly once in its restriction to H , or equivalently if

$$\text{Hom}_H(\rho, \text{Ind}(\rho^g | g^{-1}Hg \cap H, H)) = 0$$

or $g \notin H$. Or, again by Frobenius reciprocity:

10.7. Corollary. *The representation $\text{Ind}(\rho | H, G)$ is irreducible if and only if*

$$\text{Hom}_{g^{-1}Hg \cap H}(\rho, \rho^g) = 0$$

for $g \notin H$.

11. Normal subgroups

Now suppose H to be a normal subgroup of G . The classification of irreducible representations of G is reduced, to some extent, to that of classifying those of the smaller groups H and $H \backslash G$.

Suppose π to be an irreducible representation of G whose restriction to H contains the irreducible representation σ of H . Then every conjugate σ^g is also contained in π . In fact conjugation by $\pi(g)$ induces an isomorphism

$$\mathrm{Hom}_H(\sigma, \pi) \cong \mathrm{Hom}_H(\sigma^g, \pi).$$

If the restriction of π to H contains σ , then by Frobenius reciprocity π is contained in $\mathrm{Ind}(\sigma | H, G)$. But the restriction of this induced representation to H contains only conjugates of σ . So the restriction of π to H is the direct sum of copies of conjugates of σ , each one occurring with the same multiplicity. In particular, if $Z(\sigma)$ is the subgroup of all g in G with $\sigma^g \cong \sigma$, then the restriction of $\mathrm{Ind}(\sigma | H, Z(\sigma))$ to H is the direct sum of $|Z(\sigma)/H|$ copies of σ . Now induction is transitive: inducing σ from H to G is equivalent to inducing first to $Z(\sigma)$ and then to G . Not much can be said in general about the first stage, but the second stage, at least, is easy to understand.

11.1. Proposition. *Suppose H to be normal in G , σ an irreducible representation of H .*

- (a) *If χ is an irreducible constituent of $I_\chi = \mathrm{Ind}(\sigma | H, Z(\sigma))$, then $\mathrm{Ind}(\chi | Z(\sigma), G)$ is irreducible;*
- (b) *if χ_1 and χ_2 are two constituents, then $I_{\chi_1} = \mathrm{Ind}(\chi_1 | Z, G)$ and $I_{\chi_2} = \mathrm{Ind}(\chi_2 | Z, G)$ are isomorphic if and only if $\chi_1 \cong \chi_2$.*

Proof. Let (χ_i, U_i) ($i = 1, 2$) be two (possibly equal) constituents. It is to be shown that

$$\mathrm{Hom}_G(I_{\chi_1}, I_{\chi_2}) = \begin{cases} \mathbb{C} & \text{if } \chi_1 \cong \chi_2 \\ 0 & \text{otherwise.} \end{cases}$$

I shall apply Proposition 10.5. Suppose Φ to be a function on G with values in $\mathrm{Hom}(U_1, U_2)$ such that

$$\Phi(z_1 g z_2) = \chi_2^{-1}(z_2) \Phi(g) \chi_1^{-1}(z_1).$$

Then for h in H

$$\begin{aligned} \Phi(hg) &= \Phi(g) \chi_1^{-1}(h) = \Phi(g \cdot g^{-1}hg) = \chi_2^{-1}(g^{-1}hg) \Phi(g) \\ \chi_2(g^{-1}hg) \Phi(g) &= \Phi(g) \chi_1(h) \end{aligned}$$

so that Φ intertwines χ_1 and $\chi_2^{g^{-1}}$ as H representations. Both χ_1 and χ_2 are direct sums of copies of σ , so if $\Phi(g) \neq 0$ then $g \in Z(\sigma)$. Thus Φ has support in $Z(\sigma)$. But then Φ on $Z(\sigma)$ is determined by $\Phi(1)$, and Φ intertwines χ_1 and χ_2 . Since they are irreducible, the dimension of possible Φ is either 0 or 1. ▮

For us, the main consequence of this, to be applied in a later section, is:

11.2. Corollary. *Suppose H to be a normal subgroup of G , χ a character of H whose centralizer in G is H itself. Then $\mathrm{Ind}(\chi | H, G)$ is irreducible, and its restriction to H is the direct sum of conjugates of χ , each with multiplicity 1.*

This is rather elementary, and can be proved directly, but it is useful to place it in a wider context.

Now suppose A to be an abelian normal subgroup of G , and suppose G is the semi-direct product of A and $A \backslash G$. This means there exists a subgroup K of G projecting isomorphically onto $A \backslash G$. The conjugation by K permutes the characters of A . Given a character χ of A , let $Z_G(\chi)$ be the subgroup of G fixing it. This is the semi-direct product of A and $Z_K(\chi) = K \cap Z_G(\chi)$. The character χ extends to a character of $Z_G(\chi)$. If σ is an irreducible representation of $Z_K(\chi)$, then $\chi \otimes \sigma$ is also an irreducible representation of $Z_G(\chi)$. By the previous Proposition, we get the first two assertions of:

11.3. Corollary. *In these circumstances:*

- (a) the representation $\text{Ind}(\chi \otimes \sigma | Z_G(\chi), G)$ is irreducible;
- (b) one of these is isomorphic to another if and only if the inducing representation of one is the conjugate of the other;
- (c) all representations of G are obtained in this way.

Proof. Suppose π to be any irreducible representation of G , and suppose that its restriction to A contains the character χ . It therefore embeds into $\text{Ind}(\chi | A, G)$. But induction is transitive, so $\text{Ind}(\chi | A, G)$ is the representation of G induced from the representation $\text{Ind}(\chi | A, Z_G(\chi))$ to G . The representation $\text{Ind}(\chi | A, Z_G(\chi))$ is equal to $\chi \otimes \mathbb{C}[A \setminus Z_G(\chi)]$. The second factor is isomorphic to $\mathbb{C}(Z_G(\chi))$. ▣

We get unique representatives if we fix a representative of each G -orbit of characters of A .

12. The Fourier transform on abelian groups

Representations of abelian groups fall into a category all their own. The subject even has its own rather special conventions. For example, if χ is a character of G I'll express the value of χ at g as $\langle \chi, g \rangle$.

Throughout this section, suppose G to be an abelian group. All of its irreducible representations are one-dimensional, and coincide with characters of G . If χ and ρ are two characters then their product is also a character, and the product of χ and χ^{-1} is the trivial character. The characters therefore make up a group, the abelian group G^* dual to G .

Since G is abelian, each element is its own conjugacy class, so by Proposition 9.1 the size of G^* is the same as the size of G , and the dual in turn of G^* also has the same size as G . Every g in G determines a character of G^* , the one taking χ to $\chi(g)$.

12.1. Proposition. *The canonical map from G to G^{**} is an isomorphism.*

Proof. Since G and G^{**} have the same number of elements, it suffices to show that if $\langle \chi, g \rangle = 1$ for all χ then $\chi = 1$. But this is a matter of definition! ▣

The groups G and G^* are of the same size, and in fact isomorphic, although there is no natural way to identify them. For example, if $G = \mathbb{Z}/N$ then G^* is naturally isomorphic to the multiplicative group of N -th roots of unity. This also is a cyclic group, but there is no truly distinguished generator. Therefore G and G^* are isomorphic if G is cyclic, but not in a distinguished way. Because any finite abelian group is a product of cyclic groups, the group G^* is always isomorphic to G , but not canonically.

The groups G and G^* are not naturally isomorphic, but there is a natural isomorphism of $\mathbb{C}(G)$ with $\mathbb{C}(G^*)$. The basic idea is quite simple—by Proposition 8.2 every function f in $\mathbb{C}(G)$ may be expanded as a linear combination of characters, and the coefficients are a function on the space of characters associated to f . However, there is a matter of normalization to be taken into account. If f is in $\mathbb{C}(G)$, its **Fourier transform** is the function

$$\hat{f}(\chi) = \frac{1}{\sqrt{|G|}} \sum_{g \in G} f(g) \chi^{-1}(g),$$

which lies in $\mathbb{C}(G^*)$. We shall see in a moment the reason for the factor $1/\sqrt{|G|}$.

12.2. Proposition. *The map taking f to \hat{f} is an isomorphism of $\mathbb{C}(G)$ with $\mathbb{C}(G^*)$, with inverse*

$$f(g) = \frac{1}{\sqrt{|G|}} \sum_{\chi \in G^*} \hat{f}(\chi) \chi(g).$$

Proof. This is a special case of a general theorem we have already seen, but I'll prove it again, beginning with a very elementary result:

12.3. Lemma. *If G is any group and χ a character of G then*

$$\sum_{x \in G} \chi(x) = \begin{cases} |G| & \text{if } \chi \text{ is the trivial character} \\ 0 & \text{otherwise.} \end{cases}$$

Proof of the Lemma. The first equation is immediate. As for the second, if χ is not the trivial character we can find g with $\chi(g) \neq 1$. Then a change of variables gives us

$$\sum_{x \in G} \chi(x) = \sum_{x \in G} \chi(xg) = \chi(g) \sum_{x \in G} \chi(x)$$

and then

$$(1 - \chi(g)) \sum_{x \in G} \chi(x) = 0.$$

So if $\chi \neq 1 \dots$ ▢

Now for the proof of Proposition 12.2. The Lemma will be applied to the group G^* , whose dual is G . We have

$$\begin{aligned} \frac{1}{\sqrt{|G|}} \sum_{\chi \in G^*} \chi(x) \widehat{f}(\chi) &= \frac{1}{|G|} \sum_{\chi \in G^*} \chi(x) \sum_{y \in G} \chi^{-1}(y) f(y) \\ &= \frac{1}{|G|} \sum_{y \in G} f(y) \sum_{\chi \in G^*} \chi(xy^{-1}) \\ &= f(x). \quad \text{▢} \end{aligned}$$

The reason for the factor $1/\sqrt{|G|}$ has yet to appear, but here it is:

12.4. Proposition. (Plancherel formula) *For f in $\mathbb{C}(G)$*

$$\sum_{x \in G} |f(x)|^2 = \sum_{\chi \in G^*} |\widehat{f}(\chi)|^2.$$

This theory is a finite analogue of the well known theorem about Fourier series: *If f is a smooth function on the circle*

$$\mathbb{S} = \{z \in \mathbb{C} \mid |z| = 1\}$$

then

$$f(z) = \sum_n c_n z^n$$

where

$$c_n = \frac{1}{2\pi} \int_{|z|=1} f(z) z^{-n} \frac{dz}{z}.$$

It is actually more than an analogue. We can embed the group μ_N of N -th roots of unity into \mathbb{S} . As N grows, these embedded subgroups get denser and denser in \mathbb{S} . It turns out that the Fourier transform on these finite groups can be used as an approximation of Fourier series. If we take $N = 2^n$ this observation is the basis of the **fast Fourier transform**, which has proved to be an extremely practical tool for signal analysis.

Exercise 12.1. This exercise and the next two are all part of a series.

Suppose $R = \mathbb{Z}/N$, let $\psi(m) = e^{2\pi im/N}$ for all m in \mathbb{Z}/N . For each f in $\mathbb{C}(R)$ define

$$\widehat{f}(m) = \frac{1}{\sqrt{N}} \sum_{k \in R} f(k) \psi(-km).$$

Prove directly that $\widehat{\widehat{f}}(k) = f(-k)$, and that

$$\sum_k |f(k)|^2 = \sum_m |\widehat{f}(m)|^2.$$

Exercise 12.2. A character of the multiplicative group R^\times in R is called **primitive** if and only if χ is not lifted from a projection $R = (\mathbb{Z}/N)^\times \rightarrow (\mathbb{Z}/M)^\times$ with $M < N$ a divisor of N . (a) List all primitive multiplicative characters of $\mathbb{Z}/16$. (b) If p is an odd prime, how many primitive characters of $(\mathbb{Z}/p^m)^\times$ are there? (c) Prove that χ is primitive if and only if the linear space of all functions f on R such that

$$f(ax) = \chi(a)f(x)$$

for all x in R , a in R^\times has dimension one.

Exercise 12.3. Assume χ to be a primitive multiplicative character of R and define the Gauss sum:

$$G_{\chi,\psi} = \sum_n \chi(n)\psi(n).$$

What is $|G_{\chi,\psi}|$?

13. Harmonic analysis on G

I recall that the regular representation of G on the space $\mathbb{C}(G)$ is defined to be the **right** regular representation

$$R_g F(x) = F(xg)$$

(as opposed to the left regular representation $F(g^{-1}x)$).

13.1. Corollary. *The decomposition of $\mathbb{C}(G)$ is*

$$\mathbb{C}(G) = \bigoplus \dim_\pi \cdot \pi.$$

This means that an irreducible representation of G occurs in the regular representation \dim_π times. As one consequence, the irreducible representations of G have dimensions n_1, \dots, n_r then

$$n_1^2 + \dots + n_r^2 = |G|.$$

This frequently provides a good check on computations.

Proof. Calculating the inner product of its character with the characters of irreducible representations. This is easy since the trace of R_g for any $g \neq 1$ is 0, and the trace of R_1 is $|G|$. ▮

We know something more explicit—if π is irreducible, the embedding of $\widehat{V}_\pi \otimes V_\pi$ gives \dim_π linearly independent copies of V_π in $\mathbb{C}(G)$, one for each element of a basis of \widehat{V} .

14. Examples

Example. A **dihedral group** is the semi-direct product of a cyclic group and a quotient of order two. If it has order $2n$, suppose it σ to be the non-trivial element of order two with $\sigma^2 = 1$, and with

$$\sigma \cdot k \cdot \sigma = -k$$

for m in \mathbb{Z}/n . If $n = 2m$ is even, it has $m + 3$ conjugacy classes, 4 characters, and $m - 1$ irreducible representations of dimension 2, in agreement with

$$1 + 1 + 1 + 1 + \left(\frac{m-1}{2}\right)^2 = 2n.$$

What if n is odd?

Example. Suppose $G = \mathfrak{S}_3$. It contains \mathfrak{A}_3 as a normal subgroup of index 2. If σ is either character of \mathfrak{A}_3 of order three, its conjugates by $\mathfrak{A}_3 \setminus \mathfrak{S}_3$ are σ and σ^{-1} . Hence $\text{Ind}(\sigma | \mathfrak{A}_3, \mathfrak{S}_3)$ is irreducible, of dimension 2.

Example. Let $G = \mathfrak{S}_4$. This is a bit trickier. First of all, if A is the group of order 4 containing $(12)(34)$, $(13)(24)$, and $(14)(23)$, it is normal with quotient \mathfrak{S}_3 . Therefore G has irreducible representations of dimensions 1, 1, and 2 lifted from \mathfrak{S}_3 . There are 5 conjugacy classes in G , so there are two representations missing, say of dimensions n_1 and n_2 . We know that

$$1^2 + 1^2 + 2^2 + n_1^2 + n_2^2 = 24$$

so of course $n_1 = n_2 = 3$. But how can we construct the missing two? Applying the remarks in the previous example, we get them as representations induced by characters of some subgroup of index 3 in G .

Example. Let G be the group

$$\begin{bmatrix} a & x \\ 0 & 1 \end{bmatrix}$$

with $a \neq 0, x$ in \mathbb{Z}/p .

Example. Let G be the 'Heisenberg' group in $\text{SL}_3(\mathbb{Z}/p)$ of all unipotent matrices

$$\begin{bmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{bmatrix}.$$

The center Z of this group is made up of the matrices with $x = y = 0$, hence isomorphic to \mathbb{Z}/p . Any irreducible representation of G must restrict to a single character of Z . There are p characters. The representations that restrict to the trivial character of Z are those lifted from the quotient $Z \setminus G$, which is isomorphic to $(\mathbb{Z}/p)^2$. For each non-trivial character χ of Z there exists exactly one irreducible representation of G , of dimension p , restricting to χ . Check:

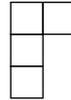
$$p^2 \cdot 1^2 + (p-1) \cdot p^2 = p^3.$$

Example. The trick of restricting to a normal subgroup by no means always works. For example, if $n \geq 5$ the only normal subgroup of \mathfrak{S}_n is \mathfrak{A}_n , which has **no** normal subgroups. Still, its representations and those of \mathfrak{S}_n can be found in terms of induced representations.

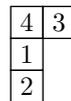
There is no obvious pattern to the construction of representations of \mathfrak{S}_3 and \mathfrak{S}_4 , so it is hard to see how to deal with the general symmetric groups, especially since \mathfrak{S}_n for $n \geq 5$ has only \mathfrak{A}_n as normal subgroup. However, there is a uniform and elegant method for finding all irreducible representations of all \mathfrak{S}_n . I'll state here without proof how it works.

There are as many irreducible representations of \mathfrak{S}_n as there are conjugacy classes, and as many conjugacy classes as partitions of n . Thus the partition $4 = 2 + 1 + 1$ corresponds to the conjugacy class $(12)(3)(4)$. I shall explain how to derive an irreducible representation of \mathfrak{S}_n from each partition.

To each partition $n = n_1 + \dots + n_k$ with $n_1 \geq n_2 \geq \dots$ corresponds what is called a **Young diagram**, with rows of boxes, the length of the i -th row down equal to n_i . For example, to the partition $4 = 2 + 1 + 1$ corresponds the diagram



A **Young tableau** is an assignment of numbers, 1 to n , of the boxes, like this:

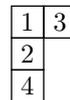
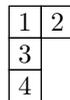


There will be $n!$ tableaux of a given shape, and the symmetric group acts on them— σ in \mathfrak{S}_n replaces i by $\sigma(i)$. To each tableau T corresponds a polynomial, the **discriminant** of the tableau. It is the product of all $x_i - x_j$ where i, j are in the same column of T , and i occurs further down than j . For example, the tableau above corresponds to

$$(x_2 - x_1)(x_2 - x_4)(x_1 - x_4).$$

The group \mathfrak{S}_n acts on the space of all such polynomials. *This representation is irreducible, different partitions give rise to different representations, and every representation of \mathfrak{S}_n is one of these.*

These polynomials are not linearly independent. For example, if one applies a permutation that just permutes the entries of each column the polynomial just changes by a sign. There is a very simple way to find a basis. Call a tableau **ordered** if the numbers in a row increase from left to right, and in a column from top to bottom. For example, here are all ordered tableaux associated to the diagram above:



The polynomials corresponding to ordered tableaux are a basis for the space of the representation associated to the partition.

This is not at all obvious, but at least for \mathfrak{S}_4 one can check it explicitly. The full story is quite interesting, and lays out an explicit algorithm to represent the polynomial associated to any tableau as a linear combination of those associated to ordered tableaux. This story is to be found in [Peel:1975].

15. References

1. M. H. Peel, 'Specht modules and symmetric groups', *Journal of Algebra* **36** (1975), 88–97.
2. J-P. Serre, **Représentations linéaires des groupes finis**, Hermann et Cie, Paris, 1967.

There is an English translation, titled **Linear representations of finite groups**. It was published by Springer-Verlag in 1977 and contains additional material.