2 Characters and The Direct Sum

2.1 Review

Last time, we defined linear representations of a group as homomorphisms $\rho : G \to \mathrm{GL}(V)$. We saw that by choosing a basis, we can rewrite linear representations as matrix representations, or homomorphisms $R : G \to \mathrm{GL}_n(\mathbb{C})$. We also introduced the **character** of a representation, which we'll discuss more today (and again in future lectures).

2.2 Characters

Recall that the character of a representation ρ is defined as the function χ_ρ on G where

$$\chi_\rho(g) = \operatorname{Tr} \rho(g)$$

for each $g \in G$. In other words, if we choose a basis and write ρ_g as a $n \times n$ matrix $R_g = (a_{ij})$, then $\chi_\rho(g) = \sum_{i=1}^{n} a_{ii}$.

At first glance, this definition appears to require us to work with *matrix* representations and to specify a basis, since the character is defined as the trace of a matrix. But thankfully, the character does not actually depend on the basis of $\mathrm{GL}(V)$ used to turn a linear representation into a matrix representation — a key property of the trace is that $\operatorname{Tr}(AB) = \operatorname{Tr}(BA)$ for any two matrices A and B, and so $\operatorname{Tr}(A) = \operatorname{Tr}(P^{-1}AP)$ for any invertible matrix P. So the characters of conjugate matrix representations coincide, and therefore the character of a linear representation is well-defined.

The fact that trace is invariant under conjugation also gives us an important property of the character — for any $g, x \in G$ we have

$$\chi_\rho(xgx^{-1}) = \chi_\rho(g),$$

since $\rho(xgx^{-1})$ and $\rho(g)$ are conjugate matrices and therefore have the same trace. So $\chi_\rho(g)$ depends only on the conjugacy class of g (meaning that χ_ρ evaluated at two conjugate elements of G will give the same result). Functions which only depend on the conjugacy class are known as **class functions**; so the character of any representation is a class function, and in order to compute the character, it's enough to compute its value on one representative of each conjugacy class.

Example 2.1
Consider the permutation representation of S_3 on $\mathbb{C}^3$, which we denote by ρ.

The conjugacy classes of S_n are described by cycle type — two elements are in the same conjugacy class if and only if the cycles in their cycle decomposition have the same lengths. So there are three conjugacy classes in S_3, with representatives (1), (12), and (123), respectively.

Clearly the permutation (1) maps to the identity matrix, which has trace 3, and therefore $\chi_\rho(1) = 3$. Meanwhile, the remaining two representatives map to the matrices

$$(12) \mapsto \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \ (123) \mapsto \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix},$$

so $\chi_\rho(12) = 1$ and $\chi_\rho(123) = 0$. So χ_ρ has the following description:

Conjugacy class	(1)	(12)	(123)
χ_ρ	3	1	0

Note that in a n-dimensional representation, the identity element of G must always map to the identity matrix, which consists of n entries of 1 on the diagonal — so $\chi_\rho(1) = \dim(\rho)$ for any representation ρ.

The simplest example of a character is a one-dimensional character (a character of a one-dimensional representation); such characters have fairly nice properties. If $\dim(\rho) = 1$, then each element ρ_g is a 1×1 invertible matrix, which can be thought of as a single nonzero number — in other words, $\mathrm{GL}_1(\mathbb{C}) = \mathbb{C}^\times$. Then $\chi_\rho(g)$ is just that number, so loosely speaking, we have $\chi_\rho(g) = \rho(g)$.

So in this case, $\chi : G \to \mathbb{C}^\times$ is a homomorphism, meaning that $\chi(gh) = \chi(g)\chi(h)$ for all $g, h \in G$. (This is *not* generally true for representations of higher dimension.) In particular, when G is finite, every $g \in G$ has finite order, and if $\text{ord}(g) = k$, then $\chi(g)$ must be a kth root of unity — written out explicitly, this is because $\chi(g)^k = \chi(g^k) = \chi(1) = 1$.

Example 2.2
Consider the group $\mathbb{Z}/n\mathbb{Z}$. Any one-dimensional representation of $\mathbb{Z}/n\mathbb{Z}$ is determined by the image of $\overline{1}$. So if we let $\zeta_n = e^{2\pi i/n}$, then we must have

$$\rho(\overline{1}) = \zeta^a = \cos\frac{2\pi a}{n} + i\sin\frac{2\pi a}{n}$$

for some integer a. Then the rest of the representation is given by $\overline{x} \mapsto \zeta_n^{ax}$ for each $\overline{x} \in \mathbb{Z}/n\mathbb{Z}$. So the character of this representation is also $\chi_\rho(\overline{x}) = \zeta_n^{ax}$.

2.3 Direct Sums

Now we'll focus on how various representations of a group relate to each other.

Guiding Question
Given two representations of a group, can we combine them to create a new representation?

One way to combine two representations is by taking their *direct sum*; this will turn out to be an important construction.

Definition 2.3 (Direct Sum)
Let $\psi : G \to \text{GL}(V)$ and $\eta : G \to \text{GL}(W)$ be two representations of the same group. Then their **direct sum** $\rho = \psi \oplus \eta$ is the representation $\rho : G \to \text{GL}(U \oplus W)$ where for each $g \in G$,

$$\rho_g(u, w) = (\psi_g(u), \eta_g(w)).$$

To describe this construction in terms of matrices, we can obtain a basis of $U \oplus W$ by appending the bases of U and W. In this basis, $\rho = \psi \oplus \eta$ will consist of the block diagonal matrices

$$\rho_g = \left(\begin{array}{c|c} \psi_g & 0 \\ \hline 0 & \eta_g \end{array}\right).$$

In particular, if $\rho = \psi \oplus \eta$, then we have

$$\chi_\rho = \chi_\psi + \chi_\eta.$$

Guiding Question
Given a representation, can it be split as a direct sum of smaller representations?

Clearly, if a representation ρ is given in a form where all the matrices ρ_g are block diagonal with blocks of the same dimensions, then it is possible to decompose ρ as a direct sum. The tricky part is when ρ is given by matrices which are not in block diagonal form, but may *become* block diagonal after a change of basis.

Example 2.4
Consider the group $\mathbb{Z}/2\mathbb{Z}$, and take the two-dimensional representation given by

$$\overline{1} \mapsto \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Clearly, in the standard basis, this matrix. However, we can instead take the basis consisting of $v_1 = (1, 1)^t$ and $v_2 = (1, -1)^t$. Then the matrix can be written as the diagonal matrix

$$\overline{1} \mapsto \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix},$$

and so the representation can in fact be written as a direct sum of two one-dimensional representations — the trivial representation on $\text{Span}(v_1)$, and the representation on $\text{Span}(v_2)$ where $\overline{1} \mapsto \begin{bmatrix} -1 \end{bmatrix}$.

In fact, $\mathbb{Z}/m\mathbb{Z}$ is relatively simple to analyze in general.

Example 2.5
Every n-dimensional representation of $\mathbb{Z}/m\mathbb{Z}$ can be split as the sum of n one-dimensional representations.

Proof. A representation ρ of $\mathbb{Z}/m\mathbb{Z}$ is determined by the matrix A corresponding to $\overline{1}$, since $\overline{1}$ generates $\mathbb{Z}/m\mathbb{Z}$; this matrix must satisfy the constraint $A^m = 1$.

So showing that ρ can be split as the direct sum of one-dimensional representations is equivalent to showing that A can be diagonalized — if A can be diagonalized, then we can find a basis $\{v_1, \ldots, v_n\}$ of V consisting of eigenvectors of A. Then we can split V as a direct sum of the spans of these eigenvectors. If v_i corresponds to the eigenvalue λ_i for each i, then we can write

$$\rho = \psi_1 \oplus \cdots \oplus \psi_n,$$

where ψ_i is the representation given by $\overline{1} \mapsto \lambda_i$, acting on $\text{Span}(v_i)$.

But it's possible to show that any matrix of finite order is in fact diagonalizable, by rewriting it in Jordan normal form and then showing that there can be no Jordan blocks of size greater than one. Since we know A has finite order, it is then possible to diagonalize A, and therefore to decompose ρ as a sum of one-dimensional representations. □

2.4 Irreducible Representations

Now that we've seen how to build new representations out of smaller ones, we would like to describe the "building blocks" of representations.

Definition 2.6
Let $\rho : G \to \text{GL}(V)$ be a representation of G. Then a subspace $W \subset V$ is called ***G*-invariant** if for each $g \in G$, we have $\rho_g(w) \in W$ for all $w \in W$.

In other words, W is G-invariant if it is taken to itself under the action of each element in G. In 18.701, we saw the concept of a T-invariant subspace for a linear operator T (a subspace taken to itself under the map T); in this situation, a subspace is G-invariant if it is invariant with respect to every one of the operators ρ_g.

Definition 2.7
A representation $\rho : G \to \text{GL}(V)$ is **irreducible** if the only G-invariant subspaces of V are 0 and V.

Note that if a representation $\rho : G \to \text{GL}(V)$ can be decomposed as a direct sum $\rho = \psi \oplus \eta$, where ψ annd η act on nonzero subspaces U and W with $V = U \oplus W$, then U is a G-invariant subspace of V (here U consists of the elements $(u, 0)$ in $V = U \oplus W$). So a direct sum of representations is always reducible (that is, not irreducible).

Meanwhile, if a representation $\rho : G \to \mathrm{GL}(V)$ is reducible, then by definition it has an invariant subspace $W \subset V$. Then we can restrict ρ to W. Initially each ρ_g is an automorphism of V, but since ρ_g preserves W, we can also think of ρ_g as an endomorphism of W (a linear map from W to itself). In fact, ρ_g must be an *automorphism*[¶] of W, since $\rho_{g^{-1}}$ restricted to W is still the inverse of ρ_g restricted to W. So then by restricting each ρ_g to W, we get a *subrepresentation* of ρ acting on W. This means any reducible representation V has a smaller representation W sitting inside it, in some sense.

In matrix form, we can let $\{v_1, \ldots, v_m\}$ be a basis of the invariant subspace W, and complete it to a basis $\{v_1, \ldots, v_n\}$ of V. With respect to this basis, the matrices in ρ are the block matrices

$$\rho_g = \left[\begin{array}{c|c} \psi_g & * \\ \hline 0 & \eta_g \end{array}\right],$$

where ψ is the representation formed by restricting ρ to W, the top-right entries $*$ are "junk," and η is the *quotient representation* $G \to \mathrm{GL}(V/W)$.

We would like to split ρ as a direct sum of ψ and another smaller representation. To do this, we'd like to pick a basis that turns the "junk" into a block of zeros for each $g \in G$. Then if $U = \mathrm{Span}(v_{m+1}, \ldots, v_n)$, we can split $V = W \oplus U$, and since W and U are both G-invariant, this would let us split $\rho = \psi \oplus \eta$.

It isn't immediately clear whether it's always possible to choose the basis in such a way. But as we'll see next class, for complex representations of *finite* groups, it is always possible! More precisely, we'll see the following result:

Theorem 2.8
Given a finite-dimensional complex representation of a finite group, each invariant subspace has a corresponding invariant complementary subspace.

This states that not only does a reducible representation have an invariant subspace $W \subset V$ (which we already know exists by definition), but there is also another invariant subspace $U \subset V$ for which $V = U \oplus W$. This will imply that a representation is reducible if and only if it can be broken down as a direct sum of nonzero representations, leading to the following result:

Theorem 2.9 (Maschke's Theorem)
Every finite-dimensional complex representation of a finite group can be written as a direct sum of irreducible representations.

This will turn out to be an incredibly useful result. Note that the proof requires the group to be *finite* — the theorem will not generally be true for infinite groups, although there is a generalization to compact groups (which we won't discuss).

[¶] An invertible endomorphism

Resource: Algebra II Student Notes
Spring 2022
Instructor: Roman Bezrukavnikov
Notes taken by Sanjana Das and Jakin Ng