Lecture 3-29: Representations of finite groups

March 29, 2024

So far our attention has been focussed on similarity classes of matrices. We now in effect pass to similarity classes of certain sets of matrices arising from groups, trying to understand both groups and matrices better by exploiting the connections between them.

Given a group G and a finite-dimensional vector space V over a field k, we define a representation of G to be a homomorphism π from G to the group $GL(V)$ of invertible linear transformations from V to itself. The dimension of V is called the degree (or dimension) of π. While this is the official dimension of representation, in practice one usually prefers to work with the vector space V rather than the homomorphism π; accordingly it too is often called a representation of G. We can also call V a G-module, since G acts linearly on it: given $g \in G, v \in V$, we define $gv = \pi(g)v$.

For example, the group G of (length and angle-preserving) symmetries of any subset S of $\mathbb{R}^n$ with centroid at the origin turns out to act linearly on $\mathbb{R}^n$, so that this action gives rise to an n-dimensional (real) representation of G.

Two representations $\pi : G \to GL(V), \pi' : G \to GL(W)$ are called equivalent if the G-modules V and W are isomorphic, so that there is a there is an invertible linear map σ mapping V onto W such that $\sigma(\pi(g)v) = \pi'(g)\sigma(v)$ for all $g \in G$. (Such a map σ is often called an intertwining operator.) If $V = W$, this says exactly that there is an invertible linear map P on V with $\pi'(g) = P\pi(g)P^{-1}$ for all $g \in G$.

If V is a representation of G and W is a subspace of V stable under the action of G, then we call W a subrepresentation of V. If V is the direct sum of two subrepresentations V_1, V_2, then we say that V is decomposable. More generally, given two representations V_1, V_2 of G, we make their direct sum $V_1 \oplus V_2$ a representation of G in the obvious way; the degree of $V_1 \oplus V_2$ is the sum of the degrees of V_1 and V_2.

If G is infinite, then it typically has additional structure; at the very least it will almost always be a topological group. In this case we usually insist that a representation π preserve this structure, so that we restrict to continuous π if G is a topological group, to smooth π if G is a Lie group, and to holomorphic π if G is a complex Lie group. In this course we will always assume that G is finite, in which case no additional restrictions are placed on π.

Example

We begin with the simplest possible group G, namely a cyclic group C_n of order n with g as generator. Clearly any representation π of G is determined by the single matrix $\pi(g)$; given the above definition of equivalence, we see that equivalence classes of representations of G of degree m are in natural bijection to similarity classes of $m \times m$ matrices M with $M^n = I$.

Example

Now you already know that the number of such classes depends heavily on how the polynomial $x^n - 1$ factors in $k[x]$. If for example k is the complex field, or more generally any algebraically closed field of characteristic not dividing n, then this polynomial splits over k into distinct linear factors. If instead $k = \mathbb{Q}$, then this polynomial is the product of irreducible cyclotomic polynomials $\Phi_d(x)$ as d runs over the divisors of n, the degree of Φ_d being $\phi(d)$, the Euler ϕ function of d.

The situation is very different if the characteristic of k divides n, whether or not k is algebraically closed, since then the polynomial $x^n - 1$ fails to have distinct roots. From our work on canonical forms we get the following result.

Theorem

Over an algebraically closed field k of characteristic not dividing n, any representation V of C_n is the direct sum of one-dimensional representations, on each of which the generator g of C_n acts by an nth root of 1 in k. If $n = p$ is prime and k has characteristic p, then any representation of C_p is a direct sum of representations $k[x]/((x-1)^\ell$ for various $\ell \leq p$, in each of which the matrix $\pi(g)$ representing the generator g is similar to an $\ell \times \ell$ Jordan block with 1s on the diagonal.

Of course the picture is much more complicated for general finite groups G. We can however say at least that over any algebraically closed field k of characteristic not dividing the order of G, this characteristic also fails to divide the order of any element of G, so that if π is a representation of G, the matrix $\pi(g)$ is diagonalizable for every $g \in G$ with eigenvalues that are roots of 1 in $\mathbb{C}$.

Example

Now consider the next simplest case, where $G = C_2 \times C_2$ is the Klein four-group, and assume that k has characteristic 2. Here G is generated by commuting elements x, y of order 2. For every integer $n \geq 1$ we will define a representation W_n of G of degree $2n + 1$ as follows. A basis of W_n is given by $w_0, \ldots, w_n, v_1, \ldots, v_n$. The generators x and y act trivially on each w_i (that is, they fix it). Then $xv_i = v_i + w_{i-1}, yv_i = v_i + w_i$ for $i \geq 1$. A picture of this representation is given by

$$\begin{array}{ccccccccc} w_0 & & w_1 & & w_2 & & \ldots & & w_n \\ & v_1 & & v_2 & & \ldots & & v_n & \end{array}$$

Example

Here x acts on the v_i by moving to the left, while y acts on the v_i by moving to the right. You can check directly that the actions of x and y commute and that acting by x twice or by y twice is the identity. Now it is not difficult to show that the representation W_n is indecomposable. In fact, it remains indecomposable if we define v_i, w_i for all $i \in \mathbb{Z}$ and define the actions of x and y on the v_i and w_i as above. Thus G admits indecomposable representations of arbitrarily large degree, and even one of infinite degree.