Representations of Finite Groups
Course Notes

Katherine Christianson

Spring 2016

A Note to the Reader

These notes were taken in a course titled “Representations of Finite Groups,” taught by Professor Mikhail Khovanov at Columbia University in spring 2016. With the exception of the proofs of a few theorems in the classification of the McKay graphs of finite subgroups of SU(2) (at the end of Section 4.2.3), I believe that these notes cover all the material presented in class during that semester. I have not organized the notes in chronological order but have instead arranged them thematically, because I think that will make them more useful for reference and review. The first chapter contains all background material on topics other than representation theory (for instance, abstract algebra and category theory). The second chapter contains the core of the representation theory covered in the course. The third chapter contains several constructions of representations (for instance, tensor product and induced representations). Finally, the fourth chapter contains applications of the theory in Chapters 2 and 3 to group theory and also to specific areas of interest in representation theory.

Although I have edited these notes as carefully as I could, there are no doubt mistakes in many places. If you find any, please let me know at this email address: kchristianson29@gmail.com. If you have any other questions or comments about the notes, feel free to email me about those as well.

Good luck, and happy studying!

Chapter 1

Algebra and Category Theory

1.1 Vector Spaces

1.1.1 Infinite-Dimensional Spaces and Zorn's Lemma

In a standard linear algebra course, one often sees proofs of many basic facts about finite-dimensional vector spaces (for instance, that every finite-dimensional vector space has a basis). Many (though not all) of these results also hold for infinite-dimensional cases, but the proofs are often omitted in an introductory course because they are somewhat more technical than their finite-dimensional analogs. In this section, we will discuss the background of Zorn's Lemma, a crucial tool in dealing with infinite-dimensional vector spaces. Zorn's Lemma is also important for dealing with general modules, which need not be finitely generated; for instance, it is used heavily in the proof of Theorem 1.4.5.

We begin with a few definitions.

Definition 1.1.1. A *partially ordered set* P is a set together with a binary relation $\leq$ such that:

1. for all a in P, $a \leq a$ (this is called reflexivity);

2. for all a and b in P, if $a \leq b$ and $b \leq a$, then $a = b$ (this is called antisymmetry); and

3. for all a, b, and c in P, if $a \leq b$ and $b \leq c$, then $a \leq c$ (this is called transitivity).

In this case, the binary relation is called a *partial order* on the set P. A *totally ordered set* T is a set together with a binary relation $\leq$ which is reflexive and antisymmetric (i.e. satisfies items 1 and 2 in the definition of a partially ordered set above) and such that, for all a and b in T, either $a \leq b$ or $b \leq a$ (this property is called totality). In this case, the binary relation is called a *total order* on the set T.

Remark. One can check that antisymmetry and totality imply reflexivity. Thus, a totally ordered set is equivalent to a partially ordered set in which the binary relation is total.

With these definitions in hand, we can now state Zorn's Lemma.

Lemma 1.1.2 (Zorn's Lemma)**.** *Let P be a nonempty partially ordered set such that, for any totally ordered subset T of P, there exists some u in S with $t \leq u$ for all t in T. Then, there exists a* maximal element *m of P, i.e. one such that there is no element $x \neq m$ of P with $m \leq x$.*

Zorn's Lemma is really a fundamental set-theoretic statement. In fact, it turns out that Zorn's Lemma is equivalent to the Axiom of Choice, which states that, given $f : X \to Y$ a surjective map of sets, there exists some $g : Y \to X$ such that $f \circ g = \mathrm{id}_Y$. (Intuitively, the idea here is that, for every y in Y, we can "choose" an element x of X in $f^{-1}(y)$, so that g is defined by sending y to x.) The Axiom of Choice is assumed in the standard formulation of set theory that mathematicians use today, which is called ZFC: the "Z" and "F" stand for Zermelo and Fraenkel, who created this system, and the "C" stands for choice, as in the Axiom of Choice. Thus, since we have the Axiom of Choice, we also have Zorn's Lemma.

Typically, the way that one applies Zorn's Lemma in an algebraic setting is by picking a collection of objects of interest — for instance, a collection of ideals of a ring or submodules of a module which satisfy a certain property — and using inclusion as a partial order. As an example of this usage, we will conclude this section by proving that any infinite-dimensional vector space has a basis, a standard result for which the arguments in the finite-dimensional case do not work.

Theorem 1.1.3. *Any vector space V over a field k has a basis.*

Proof. Let A be the collection of all sets of linearly independent vectors in V. If $\dim V = 0$, then there is nothing to prove; otherwise, there exists some nonzero element v of V, and the set $\{v\}$ is an element of A. This implies that A is nonempty. Moreover, A is a partially ordered set, with the partial order given by inclusion of sets.

Let $T = \{S_i\}_{i \in I}$ be a totally ordered subset of A. Define $S = \cup_{i \in I} S_i$. Now, suppose that S is not linearly independent: in particular, suppose we have $\sum_{i=1}^n a_i s_i = 0$, where $a_i \in k$ and $s_i \in S$ for all i and the a_i are not all 0. Then, for each i, there exists some element A_i of T containing s_i. Since T is totally ordered, we can arrange the A_i in a chain in ascending order. Up to reordering the s_i and the A_i, we then have

$$A_1 \subseteq A_2 \subseteq \cdots \subseteq A_n.$$

This implies that s_i is in A_n for all i. But then, A_n is not linearly independent, which contradicts the fact that A_n is in A. This proves that S is linearly independent and hence an element of A. Moreover, it is clear that $S_i \subseteq S$ for all i.

Now, we have shown that the conditions of Zorn's Lemma are satisfied, so we may apply it to obtain some maximal element M of A. We claim that M is a basis for V. Since M is in A, it is linearly independent. Suppose that M does not span V; then, there is some v in V which is linearly independent of all the vectors in M. So, $M \cup \{v\} \supsetneq M$ is a linearly independent set and hence an element of A, contradicting maximality of M in A. So, M spans V, which proves that it is a basis. $\square$

1.1.2 Tensor Products of Vector Spaces

Suppose that V and W are vector spaces over a field k. We can do $V \oplus W$, which is a vector space of dimension $\dim V + \dim W$. Can we construct an another operation which produces a vector space of dimension $(\dim V)(\dim W)$? It turns out that the answer is yes. First, we define a vector space $\widetilde{V \otimes W}$ by taking as a basis elements of the form $v \otimes w$ for any $v \in V$ and $w \in W$. We then define a subspace S of $\widetilde{V \otimes W}$ which is spanned by elements of the form $(v_1+v_2) \otimes w - v_1 \otimes w - v_2 \otimes w$, $v \otimes (w_1 + w_2) - v \otimes w_1 - v \otimes w_2$, $\lambda(v \otimes w) - (\lambda v) \otimes w$, and $\lambda(v \otimes w) - v \otimes (\lambda w)$ (for any $v, v_i \in V$, $w, w_i \in W$, and $\lambda \in k$). Finally, we define the *tensor product* $V \otimes W$ to be the quotient $\widetilde{V \otimes W}/S$.

The following theorem gives us an explicit basis for the tensor product in terms of bases for the original two spaces, which allows us to compute the dimension of the tensor product in the finite case.

Theorem 1.1.4. *Given bases $\{v_i\}_{i \in I}$ and $\{w_j\}_{j \in J}$ of vector spaces V and W (respectively), $\{v_i \otimes w_j\}_{i \in I, j \in J}$ is a basis of $V \otimes W$.*

Proof. Let K be the subspace of $\widetilde{V \otimes W}$ with basis $\{v_i \otimes w_j\}$. We then define a map $\widetilde{V \otimes W} \to K$ by $v \otimes w \mapsto \sum_{i,j} a_i b_j (v_i \otimes w_j)$ for any $v = \sum_i a_i v_i$ and $w = \sum_j b_j w_j$. One checks that this is a linear map whose kernel is precisely S, so it induces an isomorphism $V \otimes W \cong K$. □

Corollary 1.1.5. *If V and W are two finite-dimensional vector spaces, then $\dim(V \otimes W) = (\dim V) \cdot (\dim W)$.*

Proof. Using the notation of the theoerm, the number of elements in the basis $\{v_i \otimes w_j\}$ of $V \otimes W$ is precisely $|I| \cdot |J| = (\dim V) \cdot (\dim W)$. □

The following proposition states without proof several useful properties of the tensor product of vector spaces.

Proposition 1.1.6. *Let V_1, V_2, and V_3 be vector spaces over a field k.*

1. *(Commutativity) There is a canonical isomorphism*
$$V_1 \otimes V_2 \cong V_2 \otimes V_1.$$
2. *(Associativity) There is a canonical isomorphism*
$$V_1 \otimes (V_2 \otimes V_3) \cong (V_1 \otimes V_2) \otimes V_3.$$
3. *(Distributivity) There is a canonical isomorphism*
$$(V_1 \oplus V_2) \otimes V_3 \cong (V_1 \otimes V_3) \oplus (V_2 \otimes V_3).$$

1.1.3 Dual Spaces

Definition 1.1.7. Let V be a vector space over a field k. Then, a *linear functional* is a linear map $V \to k$. We define the *dual space* V^* to be the set of linear functionals on V.

Given any vector space V over a field k, the dual space V^* is a vector space, with addition given by

$$(f_1 + f_2)(v) = f_1(v) + f_2(v)$$

for any $f_1, f_2 \in V^*$ and any $v \in V$ and scalar multiplication given by

$$(cf)(v) = cf(v)$$

for any $c \in k$, $f \in V^*$, and $v \in V$.

The following result helps us understand the dual basis explicitly by constructing a basis for it.

Proposition 1.1.8. *Let V be a finite-dimensional vector space over a field k, and let $\{v_1, \ldots, v_n\}$ be a basis for V. Then, a basis for V^* is given by $\{v_1^*, \ldots, v_n^*\}$, where v_i^* is the linear functional defined by sending v_i to 1 and v_j to 0 for all $i \neq j$ (in other words, v_i^* is the projection onto the subspace spanned by v_i). In particular,* $\dim V = \dim V^*$.

Proof. Given any element $f \in V^*$, one can check from the definitions that $f = \sum_{i=1}^n f(v_i)v_i^*$. This proves that the v_i^* span V^*. On the other hand, suppose that

$$\sum_{i=1}^n a_i v_i^* = 0,$$

Then, this linear functional is 0 on every element of V; so, plugging in v_i gives $a_i = 0$ for each i. This proves that the v_i^* are linearly independent, so that they form a basis for V^*. $\square$

Remark. Note that if V is not finite-dimensional, then a linear functional on V is specified by an infinite tuple of elements of the base field, which means that V^* actually has uncountable dimension. In our case, however, we are generally working with finite-dimensional vector spaces, so this shouldn't be an issue.

The following proposition shows that dualizing a vector space behaves well under direct sums and tensor products.

Proposition 1.1.9. *Let V and W be any two vector spaces over a field k.*

1. $(V \oplus W)^* \cong V^* \oplus W^*$.

2. *If one of V or W is finite-dimensional, then* $(V \otimes W)^* \cong V^* \otimes W^*$.

Proof. 1. We can define a map $\psi : (V \oplus W)^* \to V^* \oplus W^*$ by sending a linear functional $f \in (V \oplus W)^*$ to $(f|_V, f|_W)$. Clearly f is 0 if and only if $f|_V = 0$ and $f|_W = 0$, so this ψ is injective. On the other hand, given any $f \in V^*$ and $g \in W^*$, we can define an element $h \in (V \oplus W)^*$ by $h(v, w) = (f(v), g(w))$. One can check that this is a linear functional and that $h|_V = f$ and $h|_W = g$, so that $\psi(h) = (f, g)$. So, ψ is a surjection, which means it is an isomorphism.

2. Here, we define a map $\psi : V^* \otimes W^* \to (V \otimes W)^*$ by $\psi(v^* \otimes w^*) \to (v \otimes w)^*$. One then checks that this is an isomorphism (and in fact a canonical one).

□

1.2 Modules

1.2.1 Basics of Modules

Groups and group actions arise naturally as symmetries of mathematical objects. For instance, given a field k, the group $\mathrm{GL}(k, n)$ acts on k^{n^2}, and the symmetric groups S_n acts on a set X with $|X| = n$. We may break down group actions into orbits, which correspond bijectively to conjugacy classes of subgroups of the group in question. We can even phrase this scenario a little more generally (albeit at the cost of our nice orbit structure) by considering the action of a monoid on a set. The most common instance of this is when we take the monoid of endomorphisms on an object.

However, almost all deep structures in math require some form of addition embedded in their construction. Thus, we wish to "linearize" the notion of a group acting on a set. This leads to the idea of a ring acting on a module, which we will define below.

Remark. Throughout these notes, rings will be taken to have a 1 element, and we will generally assume that $1 \neq 0$ (which amounts to excluding the trivial ring from consideration).

Definition 1.2.1. Let R be a ring. a *left R-module*, or a *left module over R*, is an abelian group together with a map (which we call the *action of R on M*) $\cdot : R \times M \to M$ such that, for all $a, b \in R$ and $m \in M$:

1. $(a + b) \cdot m = a \cdot m + b \cdot m$;
2. $a \cdot (m + n) = a \cdot m + a \cdot n$;
3. $a \cdot (b \cdot m) = (ab) \cdot m$; and
4. $1 \cdot m = m$.

Remark. We may understand the first 2 conditions in the above definition as a bilinearity condition. The third condition is an associativity condition, and the fourth a unitality condition.

Remark. Right modules are defined in a fairly analogous manner. However, in a commutative ring (which covers many of the interesting cases of modules — in particular, most rings arising in topology and geometry, where we deal mainly with rings of continuous functions), left modules are right modules, so we need not worry about the distinction between left and right modules very much in practice. In the below, we will generally focus simply on modules (which, in the non-commutative case, means bi-modules, i.e. left modules which are also right modules).

The module is in fact a rather general structure. Below we give several examples of modules, which illustrate the breadth of their applications and give some insight into their general structure.

Example.

1. Let k be a field, let $R = \mathrm{Mat}(n, k)$, and let $M = k^n$. Then M is an R-module, with the action of R given by matrix multiplication of column vectors.

2. In the case where $R = \mathbb{Z}$, we claim that any abelian group is a $\mathbb{Z}$-module (the converse is trivially true). Given an abelian group G, we can define the action of $\mathbb{Z}$ on G by setting
$$n \cdot \underbrace{m = m + m + \cdots + m}_{n \text{ times}}$$
for any $n \in \mathbb{Z}$ and $m \in G$. Clearly $1 \cdot m = m$ for all m; associativity and bilinearity also easily follow from the definition of the action. Thus, G has the structure of a $\mathbb{Z}$-module.

3. In the case where $R = k$ is a field, k-modules are precisely vector spaces over k, as one can readily check by comparing the definitions of the two objects.

4. Given any ring R, we may consider R as a module over itself (with the action given by multiplication in the ring). We call this the *regular module* of R.

5. Given any ring R and any positive integer n, we may form the *free module of rank n* R^n consisting of n-tuples of elements of R, with $r \cdot (r_1, \ldots, r_n) = (rr_1, \ldots, rr_n)$ for any $r \in R$ and $(r_1, \ldots, r_n) \in R^n$.

6. Let $R = \mathbb{Z}/n$. Then, the definition of a module tells us that any R-module M is an abelian group with $n \cdot m = 0$ for all $m \in M$.

7. Let $R = k[x]$ for some field k. Then, R-modules are k-vector spaces, and the action of x on the vector space is given by a linear operator.

8. Let $R = k[x, x^{-1}]$. R-modules are k-vector spaces, where x acts as an invertible linear operator on the vector space (and of course, x^{-1} acts as the inverse operator).

9. Let $R = k[x, y]$. Then, R-modules are k-vector spaces, where x and y acts on the vector space via linear operators that commute with each other (since x and y commute in R).

In order to fully define the category of modules over a given ring, we must also define the concept of a morphism of modules.

Definition 1.2.2. Let R be a ring, and let M and N be R-modules. Then, $\varphi : M \to N$ is a *homomorphism of R-modules* if $\varphi(m_1+m_2) = \varphi(m_1)+\varphi(m_2)$ for all $m_1, m_2 \in \varphi$, and if $\varphi(rm) = r\varphi(m)$ for all $r \in R$ and $m \in M$. We say that φ is an *isomorphism of R-modules* if φ has an inverse homomorphism $\varphi^{-1} : N \to M$.

Remark. Recall that in the case of topological spaces, a continuous map is a homeomorphism if and only if it has a continuous inverse; in particular, this means that some bijective continuous maps are not homeomorphisms (because their inverses may not be continuous). We have no such problems in the case of modules, however: given a bijective homomorphism of R-modules $\varphi : M \to N$, φ has an inverse because it is bijective, and this inverse is in fact always a homomorphism. Thus, to check that a homomorphism is an isomorphism is equivalent to checking that it is bijective.

We can also define subobjects in the category of R-modules.

Definition 1.2.3. Let M be an R-module. Then, a subgroup N of M is an *R-submodule* if, for any $r \in R$ and $n \in N$, $r \cdot n \in N$. If N is a submodule with $N \neq 0$ and $N \neq M$, then we say that N is a *proper R-submodule.*

Example.

1. Let $R = \mathbb{Z}$. Above, we noted that any abelian group G is an R-module. In similar fashion, one may check that any subgroup H of an abelian group G is a $\mathbb{Z}$-submodule of G.

2. Let $R = k$ be a field. k-modules are k-vector spaces, so k-submodules are precisely k-subspaces.

3. Let $R = k[x]$. Then, R-submodules are k-subspaces which are stable under the linear operator representing the action of x.

4. Given any R-module M and any $m \in M$, we can define the submodule generated by m by $Rm = \{r \cdot m : r \in R\}$.

5. When we consider the regular module of a ring R (i.e. R as a module over itself), then the submodules of R are precisely the ideals of R.

We may define the kernel and image of module homomorphisms in much the same way that one defines them for any other algebraic object.

Definition 1.2.4. Let $\varphi : M \to N$ be a homomorphism of R-modules. Then, we define the *kernel* of φ to be $\ker \varphi = \{m \in M : \varphi(m) = 0\}$ and the *image* of φ to be $\operatorname{Im} \varphi = \{n \in N : \varphi(m) = n \text{ for some } m \in M\}$.

Remark. Given an R-module homomorphism $\varphi : M \to N$, $\ker\varphi$ is an R-submodule of M, and $\operatorname{Im}\varphi$ is an R-submodule of N.

Given an R-submodule N of M, we define the quotient M/N to be the abelian group given by the quotient of abelian groups M/N, with the action of R given by $r(m+N) = rm + N$ for any $m \in M$ and $r \in R$. (See Homework 1 for a proof that this action is well-defined and gives an R-module structure on M/N.)

Given an R-module M, we may decompose M into the sum of two modules in two different ways. First, given two R-submodules M_1 and M_2 of M, we can define the *internal sum* of M_1 and M_2 to be the R-submodule

$$M_1 + M_2 = \{m_1 + m_2 : m_i \in M_i\}.$$

Note that $M_1 + M_2$ is the smalle submodule of M containing bth M_1 and M_2. Alternately, given any two R-modules M_1 and M_2, we may define the *(external) direct sum* of M_1 and M_2 to be the R-module

$$M_1 \oplus M_2 = \{(m_1, m_2) : m_i \in M_i\},$$

with both addition and the action of R defined componentwise.

Example. In the cases where modules correspond to other algebraic objects, the direct sum of modules generally corresponds to the direct sum of those objects. For instance, when R is a field, so that R-modules are vector space, the direct sum of R-modules is the direct sum of vector spaces; and when $R = \mathbb{Z}$, $\mathbb{Z}$-modules are abelian groups, and the direct sum of R-modules is the direct sum of abelian groups.

Example. The free R-module of rank n, R^n, is isomorphic to a direct sum of R (as a module over itself) with itself n times:

$$R^n \cong \underbrace{R \oplus \cdots \oplus R}_{n \text{ times}}.$$

Example. Let $R = \mathrm{Mat}(n,k)$ for some field k, and define an R-module $V = \{(r_1, \ldots, r_n) : r_i \in k\}$, where the action of R on V is given by matrix multiplication (thinking of the elements of V as column vectors). We call V a *column module* over R. Moreover, we have

$$R \cong V^n \cong \oplus_{i=1}^n V.$$

(Here we understand this isomorphism as an isomorphism of R-modules, so that we are thinking of R as a module over itself.)

Now, given any two R-submodules M_1 and M_2 of a module M, we may consider either the direct sum or the internal sum of these modules. We can then relate these sums by a homomorphism $\varphi : M_1 \oplus_R M_2 \to M_1 + M_2$ which sends (m_1, m_2) to $m_1 + m_2$ for any $m_1 \in M_1$ and $m_2 \in M_2$. This homomorphism is surjective. Moreover, if $(m_1, m_2) \in \ker\varphi$, then we must have $m_1 + m_2 = 0$, or equivalently, $m_1 = -m_2$. This implies that $m_1 \in M_1 \cap M_2$ (it is in M_1 by definition and in

m_2 because it is the negative of an element in M_2, which is a subgroup of M). Thus, $\ker \varphi = M_1 \cap M_2$. Then, the First Isomorphism Theorem for modules (see Homework 2 for more information) gives us

$$M_1 \oplus_R M_2 / M_1 \cap M_2 \cong M_1 + M_2.$$

In particular, when M_1 and M_2 are disjoint, we have $M_1 \oplus_R M_2 \cong M_1 + M_2$.

1.2.2 Idempotents and Module Structure

We begin with a definition.

Definition 1.2.5. Let R be a ring. Then, an element $e \in R$ is *idempotent* if $e^2 = e$.

Example. Let $R = \mathrm{Mat}(n, k) \cong \mathrm{End}_k(V)$ for some field k. Idempotents of R are precisely projection maps. This is rather intuitive: projecting twice onto the same subspace is the same as projecting once.

In many situations, idempotents are not very natural; indeed, we generally expect that the only idempotents of a ring may be 0 and 1 (which are always idempotents). However, when we do have them, idempotents give rise to a very nice structure.

Given any ring R and idempotent e of R, we see that $1-e$ is also idempotent, as $(1-e)^2 = 1 - e - e + e^2 = 1 - 2e + e = 1 - e$. This means that idempotents come in pairs. Moreover, $e(1-e) = e^2 - e = 0$. As such, any $a \in Re \cap R(1-e)$ is a multiple of e and $(1-e)$ and hence of $e(1-e) = 0$, so $a = 0$. So, $Re + R(1-e)$ is a submodule of R which contains $e + (1-e) = 1$ and hence all of R. Thus, $R = Re + R(1-e)$. Since $Re \cap R(1-e) = \{0\}$, we have by the above that $R \cong Re \oplus R(1-e)$. In summary, idempotents come in pairs and correspond to direct sum decompositions of modules. (Indeed, this correspondence is an if and only if: given $R \cong M_1 \oplus M_2$ for any two R-modules M_1 and M_2, one easily checks that $e = (1,0)$ and $1-e = (0,1)$ are idempotents of $M_1 \oplus M_2$ and hence correspond to idempotents of R.)

Example. Let $R = \mathrm{Mat}(n, k)$ for some field k. For any $1 \leq i, j \leq n$, we define a matrix $e_{ij} \in R$ which has a 1 in the (i,j)th entry and a 0 in every other entry. We call the e_{ij} *elementary matrices.* Now, the e_{ij} are all idempotent, so in particular we have

$$R \cong Re_{11} + R(1 - e_{11}).$$

Notice that Re_{11} is the submodule consisting of all matrices of the form

$$\left(\begin{array}{c|c} a_{11} & \\ \vdots & 0 \\ a_{n1} & \end{array}\right),$$

while $R(1-e_{11})$ consists of all matrices of the form

$$\left(\begin{array}{c|ccc} & a_{12} & \cdots & a_{1n} \\ 0 & \vdots & \ddots & \vdots \\ & a_{n2} & \cdots & a_{nn} \end{array}\right).$$

So, our direct sum decomposition of R has essentially split all matrices into their first column (which is in Re_{11}) and the rest of their columns (which are in $R(1-e_{11})$). Now, e_{22} is an idempotent of $R(1-e_{11})$. By analogy with the e_{11} case, we see that the decomposition $R(1-e_{11}) \cong R(1-e_{11})e_{22} + R(1-e_{11})(1-e_{22})$ will split off the second column of each matrix. Repeating this argument for all the e_{ii}, we end up with a direct sum decomposition

$$R \cong Re_{11} \oplus \cdots \oplus Re_{nn}.$$

Drawing from the example, we may generalize the above argument to any finite number of idempotents. Let $e_1, \ldots, e_n \in R$ be idempotents with $e_i e_j = 0$ for all $i \neq j$ (we sometimes say that the e_i are *orthogonal* given this condition). Moreover, suppose that $\sum_{i=1}^n e_i = 1$. Then, we have a direct sum decomposition $R \cong \oplus_{i=1}^n Re_i$.

The behavior of our idempotents is even nicer when they commute with the rest of the ring.

Definition 1.2.6. Let R be a ring. We define the *center* of R to be $\mathrm{Z}(R) = \{z \in R : zr = rz \forall r \in R\}$. If $e \in R$ is an idempotent which lies in the center of R, we call e a *central idempotent*.

Example. Let $R = \mathrm{Mat}(n, k)$ for some field k. Then, $\mathrm{Z}(R) = k \cdot I$, where I is the $n \times n$ identity matrix. That is, elements of the center of R are simply multiples of the identity matrix.

Now, suppose we have a ring R and a central idempotent e of R. Then, Re is closed under multiplication: for any $r, r' \in R$, we have $(re)(r'e) = rr'e^2 = (rr')e^2 \in Re$. So, Re is a subring. It doesn't actually 1 in it, but it does have a multiplicative identity, which is e. These same statements all apply when replacing $1-e$ for e; thus, analog to the above direct sum decomposition, we can write

$$R \cong Re \times R(1-e).$$

Notice that, rather than the direct sum of *modules* which we had above, this is a direct product of *rings*.

Example. Let n and m be coprime positive integers, and consider the ring $\mathbb{Z}/mn$. By Bézout's lemma, there exist some integers a and b such that $an+bm=1$. Now, an is an idempotent:

$$(an)^2 - an = an(an-1) = -anbm \equiv 0 \pmod{nm},$$

which implies that $(an)^2 = an$. Since $bm = 1 - an$, we then get a decomposition of $\mathbb{Z}/mn$ as a direct product of rings:

$$\mathbb{Z}/mn \cong \mathbb{Z}/m \times \mathbb{Z}/n.$$

(Note that this the content of one version of the Chinese Remainder Theorem.)

As with the direct sum decomposition of the regular module above, one can generalize this direct product decomposition to any finite number of idempotents.

The assumption that our idempotent is central is crucial here. Suppose we attempt a similar decomposition with e a general idempotent of R. Then, we know that, as modules, we have $R \cong Re \oplus R(1-e)$. Because e does not commute with R, we in fact have $Re \cong eRe \oplus (1-e)Re$ as abelian groups. (The reason for this decomposition is essentially the same as for the direct sum decomposition of R, but with multiplication on the other side.) Given a similar decomposition of $R(1-e)$, we get that

$$R \cong eRe \oplus (1-e)Re \oplus eR(1-e) \oplus (1-e)R(1-e)$$

as abelian groups. As for multiplication in R, we have $(eRe)(eR(1-e)) = eR(1-e)$, so that the four abelian groups in the above decomposition of R multiply together like blocks of a matrix algebra. We may thus think of elements of R like 2×2 matrices of the form $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, where $a \in eRe$, $b \in eR(1-e)$, $c \in (1-e)Re$, and $d \in (1-e)R(1-e)$.

Central idempotents also create a very nice structure on all modules over a ring. Let R be a ring, M an R-module, and e a central idempotent of R. Then, we can write $R \cong Re \times R(1-e)$. Moreover, an analogous argument to the one for the regular module decomposition $R \cong Re \oplus R(1-e)$ gives us $M \cong eM \oplus (1-e)M$ for any R-module M. If we consider the action of $Re \times R(1-e)$ on $eM \oplus (1-e)M$, we see that Re acts on eM exactly how R acts on M and Re acts trivially on $(1-e)M$ (and similarly for $R(1-e)$). Thus, any R-module is a direct sum of an Re-module and an $R(1-e)$-module, and the action of R on such a module is essentially given componentwise: that is, with R and M as above, given any $r, r' \in R$ and any $m, m' \in M$, we have

$$(re \oplus r'(1-e)) \cdot (em \oplus (1-e)m') = erm \oplus (1-e)r'm'.$$

Now, we have seen that any central idempotent gives rise to a decomposition into a product of two rings. The converse is also true: given a product ring $R = R_1 \times R_2$, the element $(1,0)$ is a central idempotent of R (which in fact gives rise to precisely to the decomposition $R = R_1 \times R_2$). In light of this, we can take our above discussion of modules over rings with central idempotents to be a statement about modules over any product ring. Since this will be useful to us in the future, we summarize the results of the above discussion in the following proposition.

Proposition 1.2.7. *Let $R = R_1 \times R_2$ be a direct product of two rings. Then, any R-modules M can be written $M = M_1 \oplus M_2$, where M_1 is an R_1-module, M_2 is an R_2-module, and the action of R on M is given component-wise (i.e. by R_1 acting on M_1 and R_2 acting on M_2).*

	Ring	Module
Idempotent	$R = \left(\begin{array}{c\|c} eRe & eR(1-e) \\ \hline (1-e)Re & (1-e)R(1-e) \end{array}\right)$	$R \cong Re \oplus R(1-e)$
Central Idempotent	$R \cong Re \times R(1-e)$	$M \cong eM \oplus (1-e)M$

Table 1.1: Table of decompositions of rings and modules via idempotents. Throughout, R is a ring with idempotent e, and M is an R-module.

We have now considered four different circumstances in which idempotents give rise to various decompositions of rings and modules. Table 1.1 summarizes the outcomes of each various combination.

1.2.3 Tensor Products of Modules

Let M and N be R-modules. Since one can think of modules as vector spaces but over rings instead of over fields, it seems logical to try to generalize the definition of the tensor product of vector spaces to define a tensor product of modules $M \otimes_R N$. Just as with the definition of the tensor product of vector spaces, we can do this by quotienting the direct product $M \times N$ by the ideal generated by the relations for R-bilinearity. The linearity relations are essentially the same as before: we take elements of the form $(m + m') \otimes n - m \otimes n - m' \otimes n$ and $m \otimes (n + n') - m \otimes n - m \otimes n'$. As for the last relation, our first guess may be to take elements of the form $am \otimes n - m \otimes an$. In the commutative case, this works fine; however, setting $am \otimes n = m \otimes an$ implies that $abm \otimes n = bm \otimes an = m \otimes ban$, which in the non-commutative case is not quite what we would expect it to be. To fix this, we can take M to be a right R-module and N to be a left R-module and impose the relation $ma \otimes n = m \otimes an$. In words, we can move elements of R past the tensor product sign, but not past M or N (equivalently, the letters must stay in the same order, ignoring the tensor product sign). In summary, our definition of the tensor product of modules is given below.

Definition 1.2.8. Let M be a right R-module, and let N be a left R-module. Then, we define the *tensor product* $M \otimes_R N$ to be the quotient of the abelian group $M \times N$ by the subgroup generated by all elements of the forms $(m + m') \otimes n - m \otimes n - m' \otimes n$, $m \otimes (n + n') - m \otimes n - m \otimes n'$, and $ma \otimes n = m \otimes an$.

Because this definition requires us to worry about whether modules are left or right modules or even both, the following definition will also be useful to us.

Definition 1.2.9. Let R and S be rings. We say that M is an (R, S)-*bimodule* if M is a left R-module and a right S-module and, moreover, if the actions of R and S on M commute: that is, for any $r \in R$, $s \in S$, and $m \in M$, we have $(rm)s = r(ms)$. In the case where $R = S$, we call M a R-*bimodule.*

Remark. Recall that when R is commutative, any left module M is also a right module. One can check that these module structures do commute with each other, so M is in fact an R-bimodule.

In order to better understand tensor products, one natural question to ask is: what is the size of $M \otimes_R N$? In general, this is a difficult question; however, in concrete instances, it is often not too hard to answer.

Example. Consider $M = \mathbb{Z}/2$ and $N = \mathbb{Z}/3$ as $\mathbb{Z}$-modules. Then, consider $\mathbb{Z}/2 \otimes \mathbb{Z}/3$. In this module, $1 \otimes 1 = 3 \otimes 1 = 1 \otimes 3 = 1 \otimes 0 = 0$. But $1 \otimes 1$ generates all of $\mathbb{Z}/2 \otimes \mathbb{Z}/3$, so $\mathbb{Z}/2 \otimes \mathbb{Z}/3$ is trivial.

Exercise. Generalizing the previous example, prove that for any positive integers m and n,

$$\mathbb{Z}/n \otimes \mathbb{Z}/m \cong \mathbb{Z}/\gcd(m,n).$$

We can generalize the above example and exercise still further with the following proposition.

Proposition 1.2.10. *Let R be a commutative ring, and let I and J be ideals of R. Then, we have an isomorphism*

$$\begin{aligned} R/I \otimes_R R/J &\xrightarrow{\sim} R/(I+J) \\ a \otimes b &\mapsto ab \end{aligned}$$

We now give some basic properties of tensor products of modules.

Proposition 1.2.11. *Let M_1 and M_2 be right R-modules, and let N_1 and N_2 be left R-modules.*

1. *$R \otimes_R N_1 \cong N_1$, and likewise, $M_1 \otimes_R R \cong M_1$.*
2. *Assuming these tensor products are all defined (i.e. that M_2 is an R-bimodule), we have a canonical isomorphism*
$$(M_1 \otimes_R M_2) \otimes_R N_1 \cong M_1 \otimes_R (M_2 \otimes_R N_1).$$
3. *The tensor product distributes over the direct sum of modules: that is,*
$$(M_1 \oplus M_2) \otimes_R N_1 \cong (M_1 \otimes N_1) \oplus (M_2 \otimes N_1),$$
and
$$M_1 \otimes_R (N_1 \oplus N_2) \cong (M_1 \otimes N_1) \oplus (M_1 \otimes N_2).$$

Proof. We show only that $R \otimes_R N_1 \cong N_1$; for proofs of the other statements, see Dummit and Foote, Section 10.4. Define $\varphi : R \otimes_R N \to N$ by $\varphi(a \otimes n) = an$. One can check that this is a well-defined homomorphism, and moreover that its inverse is given by $n \mapsto 1 \otimes n$ (which essentially amounts to the statement that $a \otimes n = 1 \otimes an$). □

Corollary 1.2.12. *$R^n \otimes_R N \cong N^n$ for any R-module N and any n. Likewise, $M \otimes_R R^n \cong M^n$.*

Proof. Recall that R^n is the direct sum of R with itself n times. The result then follows from the distributivity of the tensor product over the direct sum. □

In general, the tensor product of two modules is only an abelian group. Since we are working with modules, we would like to require that it be a module. For this, we require a little added structure on our modules, as described in the following proposition.

Proposition 1.2.13. *If M is an (R, S)-bimodule and N is a left S-module, then $M \otimes_S N$ is a left R-module, with the action of R defined by*

$$r(m \otimes n) = (rm) \otimes n.$$

1.3 Algebras

1.3.1 Division Rings and Division Algebras

Sometimes, it is useful to consider objects with multiplicative inverses but which are not necessarily commutative (and hence are not fields). The following definitions provide some vocabulary for this scenario.

Definition 1.3.1. We say that a ring R is a *division ring* if every nonzero element of R has a multiplicative inverse in R.

Let k be a field. We say that a ring R is a *k-algebra* if $k \subset R$, $1_k = 1_R$, and elements of k commute with all elements of R (that is, $k \subset \mathrm{Z}(R)$). We say that a k-algebra A is a *division algebra* if every nonzero element of A has a multiplicative inverse in A.

Example. Define the *quaternions*, $\mathbb{H}$, to be the R-algebra generated by 1, i, j, and k, with the defining relations $i^2 = j^2 = ijk = -1$. Given an element $q = a + bi + cj + dk \in \mathbb{H}$ (here a, b, c, and d are in $\mathbb{R}$), we can define the conjugate of q to be $\overline{q} = a - bi - cj - dk$. One can check that $q\overline{q} = a^2 + b^2 + c^2 + d^2$. Assuming that q is not zero (i.e. that at least one of a, b, c, or d is nonzero), this gives

$$q^{-1} = \frac{\overline{q}}{a^2 + b^2 + c^2 + d^2} \in \mathbb{H},$$

which implies that $\mathbb{H}$ is a division algebra.

The main fields we will be interested in are $\mathbb{R}$ and $\mathbb{C}$, so we wish to classify division algebras over these fields. The following two propositions do just this.

Proposition 1.3.2. *Any division algebra over $\mathbb{C}$ which is finite-dimensional (as a vector space over $\mathbb{C}$) is isomorphic to $\mathbb{C}$.*

Proof. Let D be a finite-dimensional division algebra over $\mathbb{C}$. Suppose there exists some $d \in D \setminus \mathbb{C}$. Then, d acts on D via a $\mathbb{C}$-linear map via left multiplication; denote this map by ℓ_d. Now, the characteristic polynomial of ℓ_d is a polynomial of degree at least 1 (otherwise, $\dim D = 0$ implies that D has no elements not in $\mathbb{C}$, a contradiction) with complex coefficients. There is at least one root λ of this polynomial; since $\mathbb{C}$ is algebraically closed, we have $\lambda \in \mathbb{C}$. So, λ is an eigenvalue of ℓ_d, and it has a nonzero eigenvector b in D. So, we have $db = \lambda b$, or equivalently,

$(d-\lambda)b = 0$. However, d is not in $\mathbb{C}$, so $d \neq \lambda$. Thus, $d - \lambda \neq 0$, so there exists some e in D such that $e(d-\lambda) = 1$. But then $e(d-\lambda)b = b = e \cdot 0 = 0$, which is a contradiction. So, d cannot exist, which implies that $D = \mathbb{C}$. □

Remark. The same proof here goes through if we replace $\mathbb{C}$ by any algebraically closed field.

Proposition 1.3.3. *Up to isomorphism, $\mathbb{R}$, $\mathbb{C}$, and $\mathbb{H}$ are the only division algebras over $\mathbb{R}$.*

Notice that our classification of division algebras over $\mathbb{C}$ is limited only to finite-dimensional algebras. For instance, the function field $\mathbb{C}(x)$ consisting of rational functions in x with coefficients in $\mathbb{C}$ is another division algebra over $\mathbb{C}$ (in fact, it is a field extension of $\mathbb{C}$, i.e. a commutative division algebra). However, $\mathbb{C}(x)$ is not finitely generated over $\mathbb{C}$: for instance, one can check that $\{\frac{1}{x-\lambda}\}_{\lambda \in \mathbb{C}}$ is an uncountable set of linearly independent elements of $\mathbb{C}(x)$.

1.3.2 Tensor, Symmetric, and Exterior Algebras

Throughout this section, we will work over a field k of characteristic 0. Consider a vector space V over k. Then, S_n acts on $V^{\otimes n}$ by permuting factors in the natural way. For instance, consider S_2 on $V \otimes V$. The action of S_2 fixes elements of the form $v \otimes w + w \otimes v$ and hence the whole subspace W of $V \otimes V$ spanned by such elements. So, we define S^2V (or $\mathrm{Sym}^2 V$) to be W. One can check that, if $\{v_1, \ldots, v_n\}$ is a basis of V, then $\{v_i \otimes v_j + v_j \otimes v_i\}_{i \leq j}$ forms a basis of S^2V. To make it easier to talk about this space, for any v and w in V, we will write

$$v \cdot w = v \otimes w + w \otimes v,$$

and we will call S^2V the *2nd symmetric power of V*.

We might likewise pick out the subspace where the action of $1 \in S_2$ is given by id and the action of $(1,2) \in S_2$ is given by $-\mathrm{id}$ (i.e. where S_2 acts by the sign representation rather than the trivial representation). This subspace will be generated by elements of the form $v \otimes w - w \otimes v$. We call this space the *2nd exterior power of V*, denoted $\Lambda^2 V$, and we write

$$v \wedge w = v \otimes w - w \otimes v.$$

One can check direclty that $V^{\otimes 2} = S^2V \oplus \Lambda^2 V$. We call S^2V and Λ^2V the *isotypical components* of the action of S_2 on $V^{\otimes}$. (In fact, given the action of any group G on any vector space W, we can always decompose W as $\oplus_i W_i^{n_i}$, where the W_i are the isotypical components of the action of G. To get this decomposition, we can look at the augmentation representation of G and decompose it into irreducible representations, which each act on a separate piece of W; then, the irreducible decomposition gives us the decomposition of W into isotypical components. However, we will not need this argument in such generality, so we won't go through the details here.)

We can define the nth symmetric or exterior power of V much as we did with the 2nd powers above.

Definition 1.3.4. Let V be a vector space over k. The *nth symmetric power* of V, denoted S^nV, is the subspace of $V^{\otimes n}$ on which S_n acts by the identity. The *nth exterior power* of V, denoted Λ^nV, is the subspace of $V^{\otimes n}$ where any $\sigma \in S_n$ acts by $\operatorname{sgn}(\sigma)\operatorname{id}$ (in the language of representation theory, we say that S_n acts by the sign representation). By convention, when $n = 0$, we take $S^0V = \Lambda^0V = k$.

Remark. Notice that, for any vector space V, $S^1V = \Lambda^1V = V$.

Remark. There is also an alternative definition of S^nV and Λ^nV. The general statement is notationally cumbersome, so we will simply give the definition for $n = 2$. We can define S^2V to be the quotient of $V \otimes V$ by the subspace generated by elements of the form $v \otimes w - w \otimes v$ for any v and w in V. Similarly, we can define Λ^2V to be the quotient of $V \otimes V$ by the subspace generated by elements of the form $v \otimes w + w \otimes v$. Notice that, since $v \otimes w - w \otimes v = 0$ in S^2V, we have $v \otimes w = w \otimes v$, just like our above definitions require of S^2. Likewise, we see that $v \otimes w = -w \otimes v$ in Λ^2V. Using this, one can check that these definitions are equivalent to the ones given above. (This equivalent definition is where we need the characteristic of k to be 0: in general, these two definitions are not equivalent.)

Notice that, by definition, every element of S^nV is of the form $w_1 \cdot w_2 \cdots w_n$ for some $w_i \in V$ (these are the only sorts of elements that are fixed by all of S_n). Likewise, every element of Λ^nV is of the form $w_1 \wedge w_2 \wedge \cdots \wedge w_n$. Moreover, by definition, one sees that $\cdot$ is commutative and $\wedge$ is anticommutative. One interesting consequence of this is that, for any $v \in V$, we must have $v \wedge v = -(v \wedge v)$, so that $v \wedge v = 0$. In this way, the anticommutativity of $\wedge$ actually forces many elements of the exterior power to be 0. These cancellations are such that exterior powers actually have much smaller dimension than symmetric powers in general, as the following proposition demonstrates.

Proposition 1.3.5. *Let V be a finite-dimenional vector space over k. Pick a basis $\{v_1, \ldots, v_n\}$ for V. Then, the set*

$$S = \{v_{i_1} \cdot v_{i_2} \cdot \cdots \cdot v_{i_m}\}_{i_1 \leq i_2 \leq \cdots \leq i_m}$$

forms a basis for S^mV, and the set

$$L = \{v_{i_1} \wedge v_{i_2} \wedge \cdots \wedge v_{i_m}\}_{i_1 < i_2 < \cdots < i_m}$$

forms a basis for Λ^mV. In particular,

$$\dim(S^mV) = \binom{n+m-1}{m},$$

and

$$\dim(\Lambda^mV) = \binom{n}{m}.$$

Sketch of proof. One can check that both S and L are linearly independent, so it suffices to check that they span the symmetric and exterior powers (respectively).

Notice that $\cdot$ and $\wedge$ both distribute through addition (essentially because tensor products of elements do). Using this and the fact that the v_i form a basis, we can write a general element $w_1 \cdot w_2 \cdots w_m \in S^m V$ as a linear combination of elements of the form $v_{i_1} \cdot v_{i_2} \cdots v_{i_m}$ for some i_j (which are in no particular order); likewise, we can write a general element $x_1 \wedge x_2 \wedge \cdots \wedge x_m$ as a linear combination of elements of the form $v_{j_1} \wedge v_{j_2} \wedge \cdots \wedge v_{j_m}$. For the symmetric power, however, $\cdot$ is symmetric, so we can rearrange the i_j to be in increasing order and so get a linear combination of elements in S, which proves that S spans $S^m V$. As for the exterior power, $\wedge$ is anticommutative, so we can rearrange the j_i in increasing order up to a negative sign. However, this still results in a linear combination of elements in L, so that L spans $\Lambda^m V$. $\square$

Corollary 1.3.6. *Let V be a finite-dimensional vector space over k, and let $n = \dim V$. Then, $\Lambda^m V = 0$ for all $m > n$.*

Proof. We show that all the elements of the set L in the above proposition are 0; then, a basis for $\Lambda^m V$ consists only of 0, so the vector space must be trivial. Notice that for any element $v_{i_1} \wedge v_{i_2} \wedge \cdots \wedge v_{i_m}$ in L, $i_j = i_\ell$ for some j and ℓ by the pidgeonhole principle (since there are n choices of values for the i_j but $m > n$ different i_j's). But then, we can permute all the v_{i_j} to put v_{i_j} and v_{i_ℓ} next to each other; at most, this changes the original element only by a negative sign, which does not affect whether or not it is 0. But, as noted above, $v_{i_j} = v_{i_\ell}$ implies that $v_{i_j} \wedge v_{i_\ell} = 0$, which by associativity of $\wedge$ implies that the entire element is 0. $\square$

Remark. In light of the above corollary, given some vector space V of dimension $n < \infty$, we sometimes call $\Lambda^n V$ the *top exterior power* of V, since every higher exterior power is 0.

Now, instead of working with each $V^{\otimes n}$, we can actually combine them into one large algebraic object.

Definition 1.3.7. Let V be a vector space over k. We define the *tensor algebra* of V to be

$$T(V) = \bigoplus_{i=0}^{\infty} V^{\otimes i}.$$

The tensor algebra is clearly a k-vector space, and it is in fact an associative k-algebra with multiplication given by the tensor product. More specifically, given $x_1 \otimes \cdots \otimes x_m \in V^{\otimes m}$ and $y_1 \otimes \cdots \otimes y_n \in V^{\otimes n}$, we define the product of these elements in $T(V)$ to be

$$x_1 \otimes \cdots \otimes x_m \otimes y_1 \otimes \cdots \otimes y_n \in V^{\otimes (m+n)}.$$

It turns out that $T(V)$ is a free object on V in the category of associative algebras over k. Without worrying about the precise definition of a "free object," we will simply say that it behaves like other free objects, e.g. vector spaces over fields or free modules over rings. In particular, given our discussion at the end of Section

1.5, this gives us an adjunction between the functor T which sends vector spaces to their tensor algebras and the forgetful functor: that is,

$$\mathrm{Hom}_k(T(V), R) \cong \mathrm{Hom}_k(V, R)$$

(here the hom set on the left is of k-algebra homomorphisms, while the one on the right is of maps of k-vector spaces).

Now that we've combined all the tensor powers of a vector space to make an algebra, we can also combine all of their symmetric and exterior powers to form other algebras.

Definition 1.3.8. Let V be a vector space over k. We define the *symmetric algebra*, $S(V)$, to be the quotient of $T(V)$ by the subspace generated by the elements of the form $v \otimes w - w \otimes v$. Similarly, we define the *exterior algebra*, $\Lambda(V)$, to be the quotient of $T(V)$ by the subspace generated by elements of the form $v \otimes w + w \otimes v$.

One can check that, for any vector space V,

$$S(V) \cong \bigoplus_{i=1}^{\infty} S^i V$$

and

$$\Lambda(V) \cong \bigoplus_{i=1}^{\infty} \Lambda^i V.$$

Moreover, $S(V)$ and $\Lambda(V)$ are associative k-algebras, with multiplication given by $v \cdot w$ and $v \wedge w$, respectively.

We end this section by discussion an interesting connection between exterior powers and determinants. Let V be a 2-dimensional vector space, and suppose $\alpha : V \to V$ is a linear map. Pick a basis $\{v_1, v_2\}$ for V, and let $A = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ be the matrix representing α in the chosen basis. Then, we can define α as a linear map on $V \otimes V$ by setting $\alpha(v \otimes w) = \alpha(w) \otimes \alpha(w)$ and extending linearly. Then, for any v and w in V,

$$\alpha(v \wedge w) = \alpha(v \otimes w) - \alpha(w \otimes v) = \alpha(v) \otimes \alpha(w) - \alpha(w) \otimes \alpha(v) = \alpha(v) \wedge \alpha(w).$$

So, α takes generators of $\Lambda^2 V$ to other generators of $\Lambda^2 V$, which implies that it restricts to a linear map on $\Lambda^2 V$. Now, by Proposition 1.3.5, $\Lambda^2 V$ has a basis consisting of the single element $v_1 \wedge v_2$. Because $\wedge$ distributes over addition and moreover $v \wedge v = 0$ for all v in V, we get

$$\begin{aligned} \alpha(v_1 \wedge v_2) = \alpha(v_1) \wedge \alpha(v_2) &= (av_1 + cv_2) \wedge (bv_1 + dv_2) \\ &= av_1 \wedge dv_2 + cv_2 \wedge bv_1 \\ &= (ad - bc)(v_1 \wedge v_2) = (\det A)(v_1 \wedge v_2). \end{aligned}$$

Indeed, it turns out that this is not a quirk of the 2nd exterior power but in fact generalizes to all of them, as described in the following proposition.

Proposition 1.3.9. *Let V be a finite-dimensional vector space, and let $n = \dim V$. Then, for any linear map $\alpha : V \to V$, we can define α on $\Lambda^n V$ by applying α component-wise and noting that it will fix $\Lambda^n V$ under this definition (see above for the details in the dimension-2 case). Let $\{v_1, \ldots, v_n\}$ be a basis for V, so that $v_1 \wedge v_2 \wedge \cdots \wedge v_n$ generates $\Lambda^n V$, and let A be the matrix corresponding to α in this basis. Then,*

$$\alpha(v_1 \wedge v_2 \wedge \cdots \wedge v_n) = (\det A)(v_1 \wedge v_2 \wedge \cdots \wedge v_n).$$

1.4 Homomorphisms and Decompositions of Modules

1.4.1 Irreducible and Completely Reducible Modules

Modules in general can have a very complicated structure. In order to deal with this, we would like to be able to take modules and break them down into smaller, simpler modules. The goal of this section is to develop some machinery for doing just that.

Definition 1.4.1. Let N be an R-submodule of an R-module M. Then, a *complimentary submodule* N' is one such that $N + N' = M$ and $N \cap N' = 0$. We say that M is *completely reducible* if every submodule of M has a complimentary submodule.

Example. Let k be a field, and let V be a module over k (i.e. a k-vector space). Then, V is completely reducible: given any subspace W of V, a complementary submodule to W is $W^\perp$.

Intuitively, what we want to do with completely reducible modules is break them down into sums of modules which are in some sense as small as possible. To do this, we will first need to introduce some notions of "small" modules.

Definition 1.4.2. We say that an R-module M is *cyclic* if there exists $m \in M$ such that $M = Rm = \{rm : r \in R\}$.

For any R-module M and any $m \in M$, we may construct a cyclic R-submodule of M by taking $Rm \subset M$. We may also construct a morphism $\varphi : R \to M$ defined by sending 1 to m (and hence r to rm for any r). Clearly $\operatorname{Im} \varphi = Rm$. Moreover, $\ker \varphi$ is precisely the submodule $\operatorname{Ann}(m) = \{r \in R : rm = 0\}$ of R, which we call the *annihilator* of m. Then, the First Isomorphism Theorem for modules (see Homework 2) gives us

$$Rm \cong R/\operatorname{Ann}(m).$$

In particular, when $\operatorname{Ann}(m) = 0$, Rm is the free module of rank 1 over R, which is isomorphic to the regular module.

Now, cyclic modules are very nice, but in fact the condition for "small" modules that we want is a little bit stronger than being cyclic.

Definition 1.4.3. We say that a nonzero R-module M is *irreducible* of *simple* if M has no proper submodules.

Proposition 1.4.4. *Any irreducible R-module M is cyclic.*

Proof. Let m be a nonzero element of M. Then, Rm is a submodule of M which is nonzero. But since M is irreducible, this submodule cannot be proper, which implies that $M = Rm$. ☐

Example.

1. Recall that $\mathbb{Z}$-modules are precisely abelian groups. If we restrict our attention to finitely-generated modules, one can check (using, e.g., the Fundamental Theorem of Finitely Generated Abelian Groups) that the finitely generated $\mathbb{Z}$-modules are precisely cyclic groups of prime order.

2. Consider a vector space V over a field k. Notice that the subspace generated by any nonzero element v of V is a nonzero k-submodule. If V is to be irreducible, then this subspace must be all of V, so that $\dim V = 1$ (since it has $\{v\}$ as a basis).

With this, we are ready to state an equivalent characterization to completely reducible which will be useful to us.

Theorem 1.4.5. *Let M be an R-module. Then, M is completely reducible if and only if*

$$M \cong \bigoplus_{i \in I} L_i$$

for some index set I and some simple submodules L_i of M.

Proof. If $M \cong \oplus_{i \in I} L_i$, then Proposition 1.2.7 tells us that any submodule N of M is of the form $\oplus_{i \in I} K_i$, where K_i is a submodule of L_i. But since the L_i are simple, each K_i is either 0 or L_i, so that $N = \oplus_{j \in J} L_j$ for some subset J of I. But then, clearly a complimentary submodule to N is given by $\oplus_{i \notin J} L_i$.

Conversely, suppose that M is completely reducible. Let A be the collection of all submodules N of M such that $N = \oplus_{i \in I} L_i$, where the L_i are simple submodules of N. We wish to apply Zorn's Lemma to A. (Note that the partial order we are using on A is inclusion of submodules.) First, we show that A is nonempty. Let B be the set of proper submodules of M. If B is empty, then M is simple and hence is trivially a direct sum of simple submodules, so we are done. Otherwise, one can check that, given a totally ordered subset $\{N_i\}_{i \in I}$ of B, $\cup_{i \in I} N_i$ is again an element of B, and this element is clearly a maximal element for $\{N_i\}_{i \in I}$. Thus, we can apply Zorn's Lemma to B to get a maximal proper submodule K of M. Since M is completely reducible, we can write $M \cong K \oplus L$. Then, $L \cong M/K$, so if X is a proper submodule of L, then the preimage of X under the quotient by K is a proper submodule of M which contains K, contradicting maximality of K in B. Thus, L has no proper submodules, which implies that it is a simple submodule of M. L is in A by definition, so A is nonempty.

Now, given a totally ordered subset $\{N_i\}_{i \in I}$ of of A, $\cup_{i \in I} N_i$ is again a submodule of M. Moreover, for each i, since N_i is in A, we can write $N_i = \oplus_{j \in J_i} L_{i,j}$, where $L_{i,j}$ is a simple submodule of N_i. Then, for any i and j such that $N_i \subseteq N_j$, a similar argument to that for the other direction of this proof above tells us that the $L_{i,k}$ are a subset of the $L_{j,k}$. So, for any i, let $\{K_{i,j}\}_{j \in K_i}$ be the set of $L_{i,j}$ which are not of the form $L_{k,\ell}$ for some ℓ and some k such that $N_k \subseteq N_i$. Then, we claim that

$$\bigcup_{i \in I} N_i = \bigoplus_{i \in I} \bigoplus_{j \in K_i} K_{i,j}.$$

That the left-hand side contains the right-hand side is clear, since each $K_{i,j}$ is a submodule of N_i. Conversely, given any element n of $\cup_{i \in I} N_i$, n is in N_i for some i. But N_i is a direct sum of the $K_{i,j}$ along with some other simple submodules which are contained in N_k for some $k < i$. So, either n is in $K_{i,j}$ for some j, or it is in $K_{k,j}$ for some j and k. Either way, n is in the right-hand side of the above equation, which proves equality. This implies that $\cup_{i \in I} N_i$ is an element of A. Finally, we clearly have $N_i \subseteq \cup_{i \in I} N_i$ for all i.

Since A is nonempty and has upper bounds for totally ordered subsets, we can apply Zorn's Lemma to pick a maximal element M_0 of A. To complete the proof, we show that $M_0 = M$, which means that M is a direct sum of simple submodules. Suppose otherwise; then, there exists some element m of M such that $m \notin M_0$. Then, using a similar argument to our first application of Zorn's Lemma above (to prove that A is nonempty), we can find a maximal proper submodule K of Rm and then write $Rm \cong K \oplus L$, where L is a simple submodule of Rm. This implies that L is a simple submodule of M. Moreover, since $L \subset Rm$, we have $L \cap M_0 \subset Rm \cap M_0 = 0$. So, $M_0 + L \cong M_0 \oplus L$ is a submodule of M which is strictly larger than M_0 (since $L \not\subset M_0$) and which can be written as a direct sum of simple submodules of M (just take the direct sum decomposition of M_0 and add L to it). But this contradicts maximality of M_0 in A. So, we must have $M_0 = M$. □

Remark. If we are willing to assume that M is finitely generated in the above theorem, then the proof becomes much simpler. In this case, we can simply use the application of Zorn's Lemma which we used in the above proof to show that A was nonempty to find a simple submodule L of M and then decompose M as $M \cong L \oplus K$. We can then repeat the same argument with K, and continue this repetition until we have all of M as a direct sum of simple modules. The key observation here is that, if M is finitely generated by n elements, then K can be generated by at most $n - 1$ elements (otherwise it would not be proper); thus, we will only have to repeat our argument n times. However, in the general case, we would have to repeat the argument infinitely many times, which would be unsound, so we have to use Zorn's Lemma.

The above characterization of complete reducibility is very important. It implies that, in order to understand completely reducible modules, we just need to understand irreducible modules. These modules turn out to have a very nice structure, which we will explore a little bit in the following section.

1.4.2 Homomorphisms and Endomorphisms of Modules

With all of the algebraic tools we've built up, we can now discuss the structure of homomorphisms of modules. We denote by $\mathrm{Hom}_R(M, N)$ the set of all homomorphisms of R-modules from M to N. We call $\mathrm{Hom}_R(M, N)$ a *hom set*. Hom sets of modules have the structure of an abelian group, where for any $\varphi, \psi \in \mathrm{Hom}_R(M, N)$ we define the sum $\varphi + \psi$ to be the R-module homomorphism which sends $m \in M$ to $\varphi(m) + \psi(m)$.

A special case of the above setup is when $M = N$; in this case, we write $\mathrm{End}_R(M) = \mathrm{Hom}_R(M, M)$, which denotes the set of *endomorphisms* of the R-module M. $\mathrm{End}_R(M)$ has the group structure of any hom set, but in fact we may define another operation on $\mathrm{End}_R(M)$ by composition of endomorphisms. This endows $\mathrm{End}_R(M)$ with the structure of a ring, where addition is given by the addition on hom sets and multiplication is given by composition.

Example. Consider $\varphi \in \mathrm{Hom}_R(R, M)$. Suppose that $\varphi(1) = m$. Then, for any $r \in R$, we have by the definition of an R-module that $\varphi(r) = \varphi(r \cdot 1) = r\varphi(1) = rm$, so that φ is completely determined by where it sends 1. On the other hand, for any $m \in M$, we may define a map $\psi : R \to M$ by $\psi(r) = rm$. One can check that this satisfies the definition of an R-module homomorphism. Thus, we have precisely one homomorphism for each element of m, which means that $\mathrm{Hom}_R(R, M) = M$. (We have shown that this equality makes sense as a bijection of sets; however, one can also check that the given bijection is in fact an isomorphism of abelian groups.)

In turns out that, given some conditions on the modules in question, rings of endomorphisms can have even more structure. We first consider the case of simple modules.

Proposition 1.4.6. *Let M and N be simple R-modules. Then, any homomorphism $\varphi : M \to N$ of R-modules is either the zero map or an isomorphism.*

Proof. $\ker \varphi \subset M$ is a submodule of M, so either $\ker M = M$, in which case φ is the 0 map, or $\ker M = 0$, in which case φ is injective. Likewise, $\mathrm{Im}\, \varphi$ is a submodule of N, so either $\mathrm{Im}\, \varphi = 0$, in which case φ is the 0 map or $\mathrm{Im}\, \varphi = N$, in which case φ is surjective. Putting these two facts together, we see that φ is either the 0 map or bijective. □

The above proposition implies Schur's Lemma, which is a fundamental result to Representation Theory.

Lemma 1.4.7 (Schur's Lemma)**.** *If L is a simple R-module, then $\mathrm{End}_R(L)$ is a division ring.*

Proof. We know that $\mathrm{End}_R(L)$ is a ring, so it suffices to show that this ring has multiplicative inverses. Given any nonzero element φ of $\mathrm{End}_R(L)$, the above proposition implies that φ is an isomorphism, so it has an inverse $\psi \in \mathrm{End}_R(L)$. Since the multiplication operation in $\mathrm{End}_R(L)$ is composition, ψ is also a multiplicative inverse to φ. □

If our base ring has the extra structure of an algebra over a field, then our hom sets will also acquire some of this structure.

Proposition 1.4.8. *Let R be a k-algebra, and let M and N be R-modules. Then,* $\mathrm{Hom}_R(M, N)$ *is a k-vector space.*

Proof. We know that $\mathrm{Hom}_R(M, N)$ is an abelian group. For any $\lambda \in k$ and any $\varphi \in \mathrm{Hom}_R(M, N)$, we can define $\lambda\varphi$ by $(\lambda\varphi)(m) = \lambda \cdot \varphi(m)$. One can check that this is compatible with addition in the requisite ways, so that it makes $\mathrm{Hom}_R(M, N)$ into a vector space over k. □

If we combine our results about simple modules and the above proposition about algebras, we get even more structure on our hom sets.

Proposition 1.4.9. *Let L be a simple module over a k-algebra R. Then,* $\mathrm{End}_R(L)$ *is a division algebra over k.*

Proof. The above proposition implies that $\mathrm{End}_R(L)$ is a k-vector space. Moreover, it is easy to check that $k \cdot \mathrm{id} \subset \mathrm{End}_R(L)$ is isomorphic to k as a ring. Notice that $1 \cdot \mathrm{id} = \mathrm{id}$ is the identity in $\mathrm{End}_R(L)$ and also in k (under the isomorphism $k \cdot \mathrm{id} \cong k$). Finally, the identity commutes with everything, which implies that $k \cdot \mathrm{id} \subset Z(\mathrm{End}_R(L))$. This proves that $\mathrm{End}_R(L)$ is a k-algebra. Finally, Schur's Lemma implies that this k-algebra is also a division ring, which makes it a division algebra. □

At this point, we need to take a slight detour from hom sets to talk about opposite rings.

Definition 1.4.10. For a ring R, we define the *opposite ring* R^{op}, to be the ring which is isomorphic to R as an abelian group but in which the order of multiplication is reversed: that is, for any r and s in R, $r \cdot s$ in R^{op} is $s \cdot r$ in R. An *antihomomorphism* of rings $\varphi : R \to S$ is an isomorphism of abelian groups such that $\varphi(ab) = \varphi(b)\varphi(a)$.

Example. Let $R = \mathrm{Mat}_n(k)$ be the ring of $n \times n$ matrices with coefficients in a field k. Then, the map $\varphi : R \to R$ sending $A \mapsto A^T$ is an antihomomorphism: for any $A, B \in R$, we have $(A + B)^T = A^T + B^T$ but $(AB)^T = B^T A^T$.

Exercise. Let R be the subring of $\mathrm{Mat}_3(k)$ consisting of matrices of the form

$$\begin{pmatrix} * & * & 0 \\ 0 & * & 0 \\ 0 & * & * \end{pmatrix}$$

(here the stars denote any element of k). Show that we can express R^{op} as the subring of $\mathrm{Mat}_3(k)$ consisting of matrices of the form

$$\begin{pmatrix} * & 0 & 0 \\ * & * & * \\ 0 & 0 & * \end{pmatrix}.$$

Then, show that $R \not\cong R^{\mathrm{op}}$. (Hint: look at idempotents of R and R^{op}.)

With this background in hand, we are ready to discuss a very particular case of hom sets: namely, those of the form $\mathrm{Hom}_R(R, R)$.

Proposition 1.4.11. *For any ring R, $\mathrm{End}_R(R) \cong R^{\mathrm{op}}$.*

Proof. Given any homomorphism of R-modules $\varphi : R \to R$, $\varphi(1) = a$ for some $a \in R$, which implies that $\varphi(r) = ra$ for all r. Thus, φ is determined by where it sends 1; moreover, it is possible to define a homomorphism which sends 1 to any element of R. So, if we denote φ by φ_a, then we have $\mathrm{Hom}_R(R, R) = \{\varphi_a\}_{a \in R}$. Notice that $\varphi_a + \varphi_b = \varphi_{a+b}$, but

$$\varphi_b \circ \varphi_a(x) = \varphi_b(xa) = xab = \varphi_{ab}(x).$$

This proves that the map $\mathrm{Hom}_R(R, R) \to R^{\mathrm{op}}$ sending $\varphi_a \mapsto a$ is an isomorphism. $\square$

The following proposition is not hard to prove; it is left as an exercise to the reader.

Proposition 1.4.12. *For any R-modules M_1, M_2, and M_3,*

$$\mathrm{Hom}_R(M_1 \oplus M_2, M_3) \cong \mathrm{Hom}_R(M_1, M_3) \oplus \mathrm{Hom}_R(M_2, M_3).$$

Likewise,

$$\mathrm{Hom}_R(M_1, M_2 \oplus M_3) \cong \mathrm{Hom}_R(M_1, M_2) \oplus \mathrm{Hom}_R(M_1, M_3).$$

An immediate consequence of the above proposition is that, for any R-module M,

$$\mathrm{Hom}_R(M, M^n) \cong \mathrm{End}_R(M)^n = \underbrace{\mathrm{End}_R(M) \oplus \cdots \oplus \mathrm{End}_R(M)}_{n \text{ times}}.$$

Using this, one can prove the following proposition.

Proposition 1.4.13. *For any R-module M, define $S = \mathrm{End}_R(M)$. Then,*

$$\mathrm{End}_R(M^n) = \mathrm{Hom}_R(M^n, M^n) \cong \mathrm{Mat}_n(S).$$

1.4.3 Semisimple Rings

Definition 1.4.14. We say that a ring R is *semisimple* if the regular module (i.e. R considered as a module over itself) is completely reducible.

The main structural theorem about semisimple rings is the Artin–Wedderburn Theorem.

Theorem 1.4.15 (Artin–Wedderburn)**.** *Let R be a semisimple ring. Then,*

$$R \cong \prod_{i=1}^{k} \mathrm{Mat}_{n_i}(D_i),$$

where D_i is a division ring for all i. If, in addition, R is a k-algebra, then the D_i are division algebras over k.

Proof. Since R is semisimple, we can write $R = \oplus_{i \in I} L_i$ for some submodules L_i of R. This implies that 1 is in the direct sum of the L_i, i.e. that $1 = \sum_{i \in J} \ell_i$ for some $\ell_i \in L_i$ and some finite subset J of I. (Notice that J is finite because every element of the direct sum is 0 in all but finitely many components.) This implies that $R = \oplus_{i \in J} L_i$, so that we can take this direct sum to be over finitely many simple modules. By reindexing the L_i and collecting identical simple modules in the direct sum, then, we have $R = \oplus_{i=1}^k L_i^{n_i}$, where $L_i \neq L_j$ for all $i \neq j$.

Now, we have by Proposition 1.4.12:

$$\mathrm{End}_R(R) \cong \mathrm{End}_R(\oplus_{i=1}^k L_i^{n_i}) = \prod_{i=1}^k \mathrm{Mat}_{n_i}(\mathrm{End}_R(L_i)).$$

On the other hand, Proposition 1.4.11 tells us that $\mathrm{End}_R(R) \cong R^{\mathrm{op}}$. Now, $(R^{\mathrm{op}})^{\mathrm{op}} \cong R$, and applying op to a ring can easily be seen to distribute through direct products of rings. Moreover, for any ring S and any n, one can check that $\mathrm{Mat}_n(S)^{\mathrm{op}} = \mathrm{Mat}_n(S^{\mathrm{op}})$. Putting this all together, we can take the opposite ring of both sides of the above equation to both sides of the above equation to get

$$R \cong \oplus_{i=1}^k \mathrm{Mat}_{n_i}(D_i),$$

where $D_i = \mathrm{End}_R(L_i)^{\mathrm{op}}$ for all i. By Schur's Lemma (Lemma 1.4.7), $\mathrm{End}_R(L_i)$ is a division ring. The opposite ring of a division ring is again a division ring, so this implies that D_i is a division ring for all i. Finally, we note that, if R is a k-algebra, then all the L_i are also k-algberas, so by Proposition 1.4.9, $\mathrm{End}_R(L_i)$ is a division algebra over k. This implies that D_i is also a division algebra over k. $\square$

Now that we have a very nice decomposition of semisimple rings, we wish to classify modules over them. By Proposition 1.2.7, modules over product rings are simply direct sums of modules. Thus, by the Artin–Wedderburn Theorem, it suffices to classify modules over matrix rings over division rings. The following theorem helps do just that.

Theorem 1.4.16. *For any ring R, there is a bijection between isomorphism classes of modules over R and isomorphism classes of modules over* $\mathrm{Mat}_n(R)$.

Proof. We define a map φ from modules over R to modules over $\mathrm{Mat}_n(R)$ by sending an R-module M to M^n considered as a "column module" over $\mathrm{Mat}_n(R)$: that is, $\mathrm{Mat}_n(R)$ acts by matrix multiplication on elements of M^n, which are understood as $n \times 1$ "vectors." Conversely, given a module N over $\mathrm{Mat}_n(R)$, let e_{ii} be the matrix with 1 in the (i, i)th place and 0's everywhere else. We can then define $e_{ii}N$ as a module over R: an element r of R acts by matrix multiplication by ae_{ii}. So, we define a map ψ from modules over $\mathrm{Mat}_n(R)$ to modules over R by sending N to $e_{11}N$.

Now, given an R-module M, $\varphi(M) = M^n$ as a column module over $\mathrm{Mat}_n(R)$, so $e_{11}M^n$ is composed of $n \times 1$ vectors with 0 in all but the first component; moreover, the R-module structure of $e_{11}M^n$ is that an element r of R acts by multiplication on the first component of the column vector. So, simply forgetting about all but the first component of each column vector (since all the other components are 0, anyway) gives us an isomorphism from $e_{11}M^n \mapsto M$. This proves that $\psi \circ \varphi = \text{id}$. Conversely, given a $\text{Mat}_n(R)$-module N, $\psi(N) = e_{11}N$, and $\varphi(e_{11}N)$ gives us a column module which is $e_{11}N$ in every component. Comparing definitions, one sees that this is the same as $e_{11}N \oplus \cdots \oplus e_{nn}N$, where here $e_{ii}N$ is considered as a $\text{Mat}_n(R)$-module by viewing it as a submodule of N. It remains to prove that $N \cong e_{11}N \oplus \cdots \oplus e_{nn}N$. First, notice that, in $e_{11}N + \cdots + e_{nn}N$ (notice we are doing the internal sum here, not the direct sum), we have $e_{11} \cdot 1 + \cdots + e_{nn} \cdot 1 = \text{id} \cdot 1 = 1$. This implies that $1 \in e_{11}N + \cdots + e_{nn}N$, so that $N = e_{11}N + \cdots + e_{nn}N$. Moreover, for any $i \neq j$, we have $e_{ii}N \cap e_{jj}N = 0$. Suppose otherwise; then, there is some element of N which can be written as $e_{ii}n$ for some $n \in N$ and as $e_{jj}n'$ for some $n' \in N$. Then, multiplying the equation $e_{ii}n = e_{jj}n'$ by e_{ii} gives

$$e_{ii}^2 n = e_{ii}n = e_{ii}e_{jj}n' = 0n' = 0.$$

So, our element $e_{ii}n$ in $e_{ii}N \cap e_{jj}N$ is actually 0, which implies that 0 is the only element of $e_{ii}N \cap e_{jj}N$. Because of this, we have $e_{11}N + \cdots + e_{nn}N \cong e_{11}N \oplus \cdots \oplus e_{nn}N$, so that $N \cong e_{11}N \oplus \cdots \oplus e_{nn}N$, as desired. □

With this theorem, we see that modules over semisimple rings can be understood by considering modules over the division rings in the decomposition given by the Artin–Wedderburn Theorem. Thus, we have reduced the classification of modules over semisimple rings to the classification of modules over division rings. If our semisimple ring is in fact a k-algebra, then the problem is reduced even further to classifying modules of division algebras over k. As we have seen in Section 1.3.1, there are often not many division algebras over fields of interest to us in representation theory, so this is a very nice simplification of the problem.

For our purposes, we will be interested in completely reducible modules over semisimple rings. In this case, we can reduce our module into a direct sum of simple modules, which by our above discussion must in fact be modules over a matrix algebra in the decomposition of the semisimple ring given by Artin–Wedderburn. So, the problem of classifying such modules reduces to classifying simple modules over matrix algebras. In the case where the matrix algebra is over a field, this is actually not hard to do.

Proposition 1.4.17. *Any simple module over* $\text{Mat}_n(k)$ *for a field* k *is isomorphic to the column module* k^n.

Proof. Any simple module L over $\text{Mat}_n(k)$ corresponds, via Theorem 1.4.16, to a module V over k. So, V is a k-vector space. Now, suppose that W is some proper submodule of V. Then, using the map φ defined in the proof of Theorem 1.4.16, $W \subsetneq V$ implies that $\varphi(W) \subsetneq \varphi(V) = L$. However, $\varphi(W) \neq 0$, because $W \neq 0$ and φ is injective; so $\varphi(W)$ is a proper submodule of L, which is a contradiction. This proves that W cannot exist, so V is also simple. But as we've discussed, since vector spaces are completely reducible, V is simple if and only if it's 1-dimensional. So, we have $V \cong k$, whence $L = \varphi(V) = \varphi(k) = k^n$. □

In particular, suppose that R is a semisimple algebra over an algebraically closed field k. (For our purposes, we will mainly be interested in the case where $k = \mathbb{C}$.) Then, by Artin–Wedderburn, we can write $R = \prod_{i=1}^r \mathrm{Mat}_{n_i}(D_i)$, where D_i is a division algebra over k for all i. By Theorem 1.3.2, however, the only division algebra over k is k itself, so that $R = \prod_{i=1}^r \mathrm{Mat}_{n_i}(k)$. By our above discussion, any module M over R can be written as $M = \oplus_{i=1}^r M_i$, where M_i is a module over $\mathrm{Mat}_{n_i}(k)$ for all i. If M is completely reducible, then we can take all the M_i to be direct sums of simple modules over $\mathrm{Mat}_{n_i}(k)$. But by the above proposition, the only such simple module is k^{n_i}, so that $M_i \cong (k^{n_i})^{m_i} = k^{m_i n_i}$ for some postive integer m_i. Then, we have $M = \oplus_{i=1}^r k^{m_i n_i}$. This completely classifies completely reducible modules over semisimple algebras over algebraically closed fields. Although this may sound like a very specific scenario, it turns out to be the only important scenario for representation theory, which makes this classification very useful for our purposes.

1.5 Category Theory

Definition 1.5.1. A *category* $\mathscr{C}$ consists of a collection, which we call the set of *objects* of $\mathscr{C}$ and denote by $\mathrm{Ob}(\mathscr{C})$ or simply by $\mathscr{C}$, along with, for each pair of objects X and Y of $\mathscr{C}$, a set $\mathrm{Hom}_{\mathscr{C}}(X, Y)$ (or simply $\mathrm{Hom}(X, Y)$) of maps $X \to Y$, which we call the set of *morphisms* from X to Y in $\mathscr{C}$. We further require that, for any objects $X, Y, Z \in \mathscr{C}$, we have a binary operation

$$\mathrm{Hom}(X, Y) \times \mathrm{Hom}(Y, Z) \to \mathrm{Hom}(X, Z),$$

usually thought of as composition and denoted by $\circ$, which satisfies the following properties:

1. composition is associative: that is, for any $W, X, Y, Z \in \mathscr{C}$ and any $f \in \mathrm{Hom}(W, X)$, $g \in \mathrm{Hom}(X, Y)$, and $h \in \mathrm{Hom}(Y, Z)$,

$$h \circ (g \circ f) = (h \circ g) \circ f;$$

   and

2. for every object X of $\mathscr{C}$, there exists a morphism id_X, which we call the *identity on* X, such that for all $f \in \mathrm{Hom}(X, Y)$ for some Y and all $g \in \mathrm{Hom}(Z, X)$ for some Z, $f \circ \mathrm{id}_X = f$ and $\mathrm{id}_x \circ g = g$.

Although the definition of a category may sound complicated and abstruse, categories actually arise quite naturally in many disparate fields of math. Below we list some common examples of categories and introduce some notation for them.

Example.

1. Perhaps the most basic category is Set, in which objects are sets and morphisms are set maps.

2. We have a category Grp in which objects are groups and morphisms are group homomorphisms. This category also has a subcategory Ab consisting only of abelian groups (with morphisms still given by group homomorphisms).

3. Rings also form a category, which we call Ring. The objects are rings, and the morphisms are ring homomorphisms.

4. Given any ring R, we can form the category of R-modules, where morphisms are homomorphisms of R-modules. We denote this category by Mod_R. Likewise, given any field k, we can form the category of k-vector spaces, where morphisms are k-linear maps. We denote this category by Vect_k.

5. Categories don't just arise in algebra: we can also define the category Top consisting of topological spaces, where morphisms are continuous maps. In algebraic topology, in particular, we are often interested in maps only up to homotopy. To this end, one can define a separate category, called hTop, consisting of all topological spaces, where morphisms are homotopy classes of continuous maps.

Now that we've defined categories, we also want to understand what maps between categories look like. It turns out that there are actually two different forms of such maps, which we define now.

Definition 1.5.2. Let $\mathscr{C}$ and $\mathscr{D}$ be two categories. Then, a *covariant functor* $F : \mathscr{C} \to \mathscr{D}$ associates to each object X in $\mathscr{C}$ an object $F(X)$ in $\mathscr{D}$ and associates to each morphism $f \in \mathrm{Hom}_{\mathscr{C}}(X, Y)$ a morphism $F(f) \in \mathrm{Hom}_{\mathscr{D}}(F(X), F(Y))$ such that

1. $F(\mathrm{id}_X) = \mathrm{id}_{F(X)}$ for all $X \in \mathscr{C}$, and

2. $F(g \circ f) = F(g) \circ F(f)$ for any $f \in \mathrm{Hom}_{\mathscr{C}}(X, Y)$ and any $g \in \mathrm{Hom}_{\mathscr{C}}(Y, Z)$.

A *contravariant functor* $F : \mathscr{C} \to \mathscr{D}$ associates to each object X in $\mathscr{C}$ an object $F(X)$ in $\mathscr{D}$ and associates to each morphism $f \in \mathrm{Hom}_{\mathscr{C}}(X, Y)$ a morphism $F(f) \in \mathrm{Hom}_{\mathscr{D}}(F(Y), F(X))$ such that

1. $F(\mathrm{id}_X) = \mathrm{id}_{F(X)}$ for all $X \in \mathscr{C}$, and

2. $F(g \circ f) = F(f) \circ F(g)$ for any $f \in \mathrm{Hom}_{\mathscr{C}}(X, Y)$ and any $g \in \mathrm{Hom}_{\mathscr{C}}(Y, Z)$.

Remark. We often call covariant functors simply *functors*.

Example. We can define the abelianization functor $A : \mathrm{Grp} \to \mathrm{Ab}$ by $A(G) = G/[G, G]$. Then, given any group homomorphism $f : G \to H$, one can check that $f([G, G]) \subset [H, H]$, so that f induces a map $\widetilde{f} : G/[G, G] \to H/[H, H]$; so, we define $A(f) = \widetilde{f}$. This definition of A makes it a covariant functor (covariant because $A(f)$ goes from $A(G)$ to $A(H)$).

Example. Given any field k, we can define the dualization functor $* : \mathrm{Vect}_k \to \mathrm{Vect}_k$ by sending $V \mapsto V^*$ and $f : V \to W$ to the dual map $f^* : W^* \to V^*$. (See the very beginning of Section 3.3 for a precise construction of f^*.) One can check that this definition makes $*$ into a contravariant functor (contravariant because f^* goes from W to V, whereas f goes from V to W).

Now, given two functors $F : \mathscr{C} \to \mathscr{D}$ and $G : \mathscr{D} \to \mathscr{C}$, we are often interested in characterizing nice relationships between F and G. For instance, they may be "inverses" in a certain sense (this is made precise by the definition of a natural isomorphism). However, even when functors are not directly inverses, they can still interact nicely. We now define one such scenario.

Definition 1.5.3. Let $F : \mathscr{C} \to \mathscr{D}$ and $G : \mathscr{D} \to \mathscr{C}$ be two functors. We say that F and G are *adjoint* if, for every $X \in \mathscr{C}$ and $Y \in \mathscr{D}$, there is a natural isomorphism

$$\mathrm{Hom}_{\mathscr{D}}(F(X), Y) \cong \mathrm{Hom}_{\mathscr{C}}(X, G(Y)).$$

In this case, we say that F is a *left adjoint* to G (because it appears on the left in the Hom set in the above equation) and that G is a *right adjoint* to F (because it appears on the right in the Hom set).

Remark. The phrase "natural isomorphism" in this diagram merits some explaining. In the notation of the definition, suppose we have bijections $\varphi_{X,Y} : \mathrm{Hom}_{\mathscr{D}}(F(X), Y) \to \mathrm{Hom}_{\mathscr{C}}(X, G(Y))$. In order for the $\varphi_{X,Y}$ to be considered natural, we require that, for any X and X' in $\mathscr{C}$ and any morphism $\gamma : X \to X'$, the following diagram commutes:

$$\begin{tikzcd}
\mathrm{Hom}_{\mathscr{D}}(F(X), Y) \arrow[r, "\varphi_{X,Y}"] \arrow[d, "F(\gamma)\circ -"] & \mathrm{Hom}_{\mathscr{C}}(X, G(Y)) \arrow[d, "\gamma\circ -"] \\
\mathrm{Hom}_{\mathscr{D}}(F(X'), Y) \arrow[r, "\varphi_{X',Y}"] & \mathrm{Hom}_{\mathscr{C}}(X', G(Y))
\end{tikzcd}$$

Likewise, for any morphism $\eta \in \mathrm{Hom}_{\mathscr{D}}(Y, Y')$, we require that the following diagram commutes:

$$\begin{tikzcd}
\mathrm{Hom}_{\mathscr{D}}(F(X), Y) \arrow[r, "\varphi_{X,Y}"] \arrow[d, "\eta\circ -"] & \mathrm{Hom}_{\mathscr{C}}(X, G(Y)) \arrow[d, "G(\eta)\circ -"] \\
\mathrm{Hom}_{\mathscr{D}}(F(X), Y') \arrow[r, "\varphi_{X,Y'}"] & \mathrm{Hom}_{\mathscr{C}}(X, G(Y'))
\end{tikzcd}$$

Morally, the point of these diagrams is that the isomorphisms $\varphi_{X,Y}$ are defined generally, based on properties of the functors and the categories, rather than based on attributed specific to X of Y. In practice, this intuition is usually convincing enough, so one rarely checks these diagrams directly.

In order to make our discussion of adjoint functors more concrete, we describe a specific set of examples of adjoint functors. First, we need to define the functors we will use in these adjunctions. One class of functors we'll use are the so-called *forgetful functors*. These are functors which are defined by "forgetting" some of the structure of the objects in a category. For instance, we can define a forgetful functor Grp $\to$ Set which simply sends any group G to itself considered as a set (i.e. "forgetting" the multiplication on G) and any group homomorphism to itself (i.e. "forgetting" that it respects multiplication and thinking of it as just a set map). Likewise, one can define a forgetful functor Ring $\to$ Ab By "forgetting" the multiplicative structure on rings and ring homomorphisms and viewing them simply as abelian groups and group homomorphisms.

The second class of functors we will use are the *free object functors*. These are functors which sends an object X to some object when can be considered "free over X." For instance, we can define a functor F : Set $\to$ Vect$_k$ by sending a set X to the vector space $k[X]$ generated by elements of X. Then, any set map $f : X \to Y$ defines a mapping of a basis of $k[X]$ to a basis of $k[Y]$ and so defines a linear map $k[X] \to k[Y]$; setting $F(f)$ to be this linear map then defines F as a functor. Alternately, we could define G : Set $\to$ Grp by sending a set X to the free group with generators the elements of X. Again, morphisms of free groups are defined by what they do to generators, so a set map will give rise to a map on free groups, which shows that G is functorial.

Now, it turns out that free object functors are generally left adjoint to forgetful functors. The reason for this is essentially encoded in what we've already said about free object functors, namely: a map on the generators of a free object uniquely specifies a map on the free object. As a concrete example, let F : Set $\to$ Vect$_k$ be the free object functor and G : Vect$_k$ $\to$ Set be the forgetful functor. Then, we have a natural isomorphism

$$\mathrm{Hom}_{\mathrm{Vect}_k}(F(X), Y) \cong \mathrm{Hom}_{\mathrm{Set}}(X, G(Y)).$$

The isomorphism arises by taking any linear map $F(X) \to Y$, restricting it to the elements of X, and then considering it as a set map $X \to Y$. This is injective because X is a basis for $F(X)$, so a linear map is determined by what it does on X; and it is surjective because we can send the elements of X to any elements of Y and then extend linearly to all of $F(X)$ to define a linear map.

Chapter 2

Representation Theory of Finite Groups

2.1 Group Algebras and Representations

Definition 2.1.1. Now, let G be a group. Then, we can define an associative k-algebra by $k[G] = \{\sum_{i=1}^{n} a_i g_i : n \geq 0, a_i \in k, g_i \in G\}$, with addition defined by $a_1 g + a_2 g = (a_1 + a_2) g$ and multiplication defined by

$$(\sum_{i=1}^{m} a_i g_i)(\sum_{j=1}^{n} b_i h_i) = \sum_{i=1}^{m} \sum_{j=1}^{n} a_i b_j g_i h_j.$$

We call $k[G]$ a *group algebra*.

Remark. The same definition can be used to define a k-algebra even when G is a monoid. However, we will not have much reason to consider this scenario.

Given a group G and a field k, one can check that a basis for $k[G]$ as a k-vector space is given by the elements of G. In particular, this implies that, when G is a finite group, $k[G]$ is a finite-dimensional vector space over k.

Example. Let k be a field, and consider the group algebra $k[C_\infty]$. If g is a generator of C_∞, then elements of $k[G]$ are of the form $\sum_{i=-N}^{N} a_i g^i$, where $a_i \in k$ for all i. Because of this characterization, we have $k[C_\infty] \cong k[g, g^{-1}]$.

Now, suppose that V is a module over a group algebra $k[G]$. Then, V is a vector space over $k \subset k[G]$ (since modules over fields are vector spaces), and each element g of G acts on the vector space V by a k-linear map. Denote this k-linear map by $\varphi(g) : V \to V$. Then, for V to be a module, we must have $\varphi(g)\varphi(h) = \varphi(gh)$ for all g and h in G. So, φ defines a homomorphism $G \to \mathrm{GL}(V)$, where $\mathrm{GL}(V)$ is the *general linear group*, i.e. the group of invertible linear maps from V to itself.

Notice that the module structure of V is completely characterized by φ and makes no reference to group algebras. It is thus a very convenient way to conceptualize modules over group algebras. This leads us to the following definition.

Definition 2.1.2. Let V be a module over a group algebra $k[G]$, and let $\varphi : G \to \mathrm{GL}(V)$ be the homomorphism corresponding to V as constructed above. Then, we call V (or sometimes φ) a *representation of G over k*. We say that W is a *subrepresentation* of V if W is a $k[G]$-submodule of V, i.e. if W is a subspace of the vector space V which is fixed by the action of elements of G. The *dimension* or *degree* of a representation is its dimension as a vector space. We say that a representation $\varphi : G \to \mathrm{GL}(V)$ is *faithful* if φ is injective.

Now, with V, G, and φ as above, suppose that φ sends a general element g of G to the matrix A_g (that is, A_g represents the action of g on V in some given basis). Let B be a change-of-basis matrix for V. Then, G acts on $B(V)$ by $g \mapsto BAB^{-1}$, which one can check is isomorphic to our original representation (by which we mean isomorphic as $k[G]$-modules). So, the representation V does not depend on the basis we choose to specify φ.

Recall that all vector spaces are completely reducible. However, the same is not true of representations.

Example. Let $G = C_\infty$, and let V be a 2-dimensional vector space over a field k. Pick a generator g of G and a basis $\{v_1, v_2\}$ of V. Then, to define a representation of G, we just have to define a homomorphism $G \mapsto \mathrm{GL}(V)$. It suffices to pick where we send g in this map; so, we define a representation by sending $g \mapsto \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$. Then, we have $gv_1 = v_1$, so that $kv_1 \subseteq V$ is a $k[G]$-submodule of V which is fixed by the action of G, hence it is a subrepresentation. However, its complement in the vector space is kv_2, which is not fixed by G. So, the subrepresentation has no complimentary subrepresentation, which implies that V is not a completely reducible $k[G]$-module.

However, if we restrict our attention to representations of finite groups, it turns out that our representations will be completely reducible (with a minor assumption on the field over which we work).

Theorem 2.1.3 (Maschke's Theorem). *Let G be a finite group, and let k be a field such that* $\operatorname{char} k \neq |G|$. *Then, any representation of G over k is completely reducible.*

Proof. Let V be a representation of G over k, and let W be a proper suprepresentation of V (of course, if no proper subrepresentation exists, i.e. if V is simple, then there is nothing to prove). Now, W is a subspace of the vector space V, so we may choose a compliment $\widetilde{W}$ of W in this vector space. Notice that picking $\widetilde{W}$ is equivalent to picking a k-linear map $\widetilde{p} : V \to W$ which is the identity on W: given $\widetilde{p}$, we may set $\widetilde{W} = \ker \widetilde{p}$, which one can check is a complement to W; and conversely, given $\widetilde{W}$, we can write $V = W \oplus \widetilde{W}$, so we can define $\widetilde{p}$ to be the projection of any element onto W. So, given $\widetilde{W}$, we also get a projection map $\widetilde{p}$. Then, we define a map $p : V \to W$ by

$$p(v) = \frac{1}{|G|} \sum_{g \in G} g\widetilde{p}(g^{-1}v).$$

(Here is where the assumption that $\operatorname{char} k \neq |G|$ comes in: this implies that $|G| \neq 0$ in k, so that we can indeed divide by this scalar in the definition of p.) Now, p is a k-linear map, since $\widetilde{p}$ is and since every element of G commutes with k in $k[G]$. Moreover, for any element w of W and any element g of G, $g^{-1}w$ is in W because W is a subrepresentation of V (so it is fixed by the action of G), so we have

$$p(w) = \frac{1}{|G|}\sum_{g\in G} g\widetilde{p}(g^{-1}w) = \frac{1}{|G|}\sum_{g\in G} g(g^{-1}w) = \frac{1}{|G|}\sum_{g\in G} w = w.$$

So, p is the identity on W. Finally, for any element h of G and any v in V,

$$hp(v) = \frac{1}{|G|}\sum_{g\in G} hg\widetilde{p}(g^{-1}v) = \frac{1}{|G|}\sum_{f\in G} f\widetilde{p}(f^{-1}hv) = p(hv).$$

(Here the second equality follows from setting $f = hg$ and noting that, as g runs over all the elements of G, so will f.) Thus, p commutes with the action of G, which implies that, for any $z \in \ker p$, we have $p(hz) = hp(z) = 0$, so that $hz \in \ker p$ for all h in G. Thus, the kernel of p is fixed by the action of G, which makes it a subrepresentation of V. Finally, we note that, by our above discussion, $\ker p$ is a complement to W as a subspace of V, so $\ker p$ is a complementary subrepresentation to W. □

Because the representations considered in Maschke's Theorem are completely reducible, Theorem 1.4.5 implies that they can be written as direct sums of irreducible subrepresentations. The "averaging process" used to define p in the proof of Maschke's Theorem can also be useful for finding these decompositions into irreducible subrepresentations, as the following example demonstrates.

Example. Let $G = C_n$, let g be a generator of G, and let V be a representation of G over some field k with $\operatorname{char} k \neq n$. Then, given any element v of V, we can define

$$v_+ = \frac{v + g(v) + g^2(v) + \cdots + g^{n-1}(v)}{n}$$

and

$$v_- = \frac{v - g(v) + g^2(v) - \cdots + (-1)^{n-1}g^{n-1}(v)}{n}.$$

Then, we set $V_+ = \{v_+ : v \in V\}$ and $V_- = \{v_- : v \in V\}$. One can check that V_+ and V_- are subspaces of V. Notice that $gv_+ = v_+$, and $gv_- = -v_-$ for any v. This immediately implies that V_+ is a subrepresentation of V; as for V_-, we note that $-v_- = (-v)_-$, so that g also fixes V_- and V_- is a subrepresentation of V. Moreover, $V_+ \cap V_- = 0$, and $1_- + 1_+ = 1$, so that $V = V_+ + V_-$. Thus, we have decomposed our representation into a direct sum of two subrepresentations. However, we can break them down further. If $\{v_1, \ldots, v_m\}$ is a basis for V_+ and $\{w_1, \ldots, w_n\}$ is a basis for V_-, then the above discussion shows that kv_i and kw_j are subrepresentations of V_+ and V_- (respectively) for all i and j. So, our above decomposition becomes

$$V = \oplus_{i=1}^m kv_i \oplus (\oplus_{j=1}^n kw_j).$$

Since the kv_i and kw_j are 1-dimensional vector spaces, they are simple as vector spaces and hence simple as representations, so the above is a decomposition of V into irreducible representations. This implies that V is completely reducible, which confirms the statement of Maschke's Theorem in this case.

Exercise. It turns out that Maschke's Theorem is the strongest condition of this form that we can put on complete reducibility of representations. To see that we cannot strengthen it in the natural way, find a representation of C_2 over a field of characteristic 2 which is not completely reducible.

Maschke's Theorem is incredibly important to the representation theory of finite groups. To see its impact, notice that, for any finite group G and any field k of characteristic not equal to $|G|$, we can consider the *regular representation* of G, which is given by considering $k[G]$ as a module over itself. Then, Maschke's Theorem implies that this representation is completely reducible, i.e. that $k[G]$ is a semisimple ring. So, we can apply all of our results from Section 1.4.3 to the classification of representations of G. In particular, we will be most interested in the case where $k = \mathbb{C}$; in this case, our discussion about algebraically closed fields at the end of Section 1.4.3 implies that

$$\mathbb{C}[G] \cong \prod_{i=1}^{k} \mathrm{Mat}_{n_i}(\mathbb{C}).$$

With this information, we can already say a several interesting things about the structure of complex representations of G.

Theorem 2.1.4. *Let G be any finite group, and write*

$$\mathbb{C}[G] \cong \prod_{i=1}^{k} \mathrm{Mat}_{n_i}(\mathbb{C}).$$

Then, G has exactly k irreducible representations, and their dimensions are precisely the n_i. Moreover, k is the number of conjugacy classes of G, and

$$|G| = \sum_{i=1}^{k} n_i^2.$$

Proof. Any representation V of G is a module over $\mathbb{C}[G]$, so by Proposition 1.2.7, we can write $V = \oplus_{i=1}^{k} V_i$, where V_i is a module over the $\mathrm{Mat}_{n_i}(\mathbb{C})$. So, if V is simple, we see that $V_i = 0$ for all but one value of i (otherwise, throwing out one of the nonzero V_i gives a proper submodule of V); in other words, simple $\mathbb{C}[G]$-modules are just simple modules over $\mathrm{Mat}_{n_i}(\mathbb{C})$ for some i. But by Proposition 1.4.17, there is only one such module: the column module $\mathbb{C}^{n_i}$. So, there are exactly k irreducible representations of G over $\mathbb{C}$, corresponding to the $\mathbb{C}^{n_i}$ for any i. Clearly the dimension of $\mathbb{C}^{n_i}$ is n_i.

Notice that the dimension of $\mathbb{C}[G]$ (as a $\mathbb{C}$-vector space) is $|G|$ and the dimension of $\mathrm{Mat}_{n_i}(\mathbb{C})$ is n_i^2. Because $\mathbb{C}[G] \cong \prod_{i=1}^k \mathrm{Mat}_{n_i}(\mathbb{C})$, the dimensions of these two vector spaces must be equal, so we get

$$|G| = \sum_{i=1}^k n_i^2.$$

We would also like to compare the dimensions of the centers of these $\mathbb{C}$-algebras. We require two facts: first, for any rings R_1 and R_2, $\mathrm{Z}(R_1 \times R_2) = \mathrm{Z}(R_1) \times \mathrm{Z}(R_2)$; and second, the center of any matrix algebra consists only of multiple of the identity. (See Homework 2 for a proof of both of these facts.) Using this, we see that the dimension of the center of a matrix algebra is 1, so that

$$\dim \mathrm{Z}\left(\prod_{i=1}^k \mathrm{Mat}_{n_i}(\mathbb{C})\right) = \dim\left(\prod_{i=1}^k \mathrm{Z}(\mathrm{Mat}_{n_i}(\mathbb{C}))\right) = \sum_{i=1}^k \dim(\mathrm{Z}(\mathrm{Mat}_{n_i}(\mathbb{C}))) = k.$$

On the other hand, let x be an element of $\mathbb{C}[G]$. Then, we can write $x = \sum_{g \in G} a_g g$, where $a_g \in \mathbb{C}$ for all g. Now, $x \in \mathrm{Z}(\mathbb{C}[G])$ if and only if $xh = hx'$ for all h in G, i.e. if and only if

$$h\left(\sum_{g \in G} a_g g\right) h^{-1} = \sum_{g \in G} a_g (hgh^{-1}) = \sum_{g \in G} a_g g.$$

Comparing coefficients of each g, we see that this equation is equivalent to the statement that $a_{hgh^{-1}} = a_g$ for all g and h in G, i.e. that x has the same coefficient for every element of a given conjugacy class of G. Let $\mathcal{O}_1, \ldots, \mathcal{O}_m$ be the conjugacy classes of G, and define $g_i = \sum_{g \in \mathcal{O}_i} g$. Then, by what we've just said, x is in $\mathrm{Z}(\mathbb{C}[G])$ if and only if x is a linear combination of the g_i. One can check that the g_i are linearly independent in $\mathbb{C}[G]$, so they form a basis for $\mathrm{Z}(\mathbb{C}[G])$; since there is one conjugacy class for each g_i, we have $\dim(\mathrm{Z}(\mathbb{C}[G])) = m$ is the number of conjugacy classes of G. On the other hand, we determined above that this dimension is k, so k is the number of conjugacy classes of G. $\square$

Corollary 2.1.5. *The number of irreducible representations of any finite group G is finite.*

Proof. This follows from the above theorem along with the fact that the number of conjugacy classes of G is finite. $\square$

In the special case where G is abelian, the situation gets even nicer, as the following proposition shows.

Proposition 2.1.6. *Let G be a finite abelian group. Then, every irreducible complex representation of G is 1-dimensional, and the number of irreducible representations is $|G|$.*

Proof. Since G is abelian, $\mathbb{C}[G]$ is a commutative $\mathbb{C}$-algebra. On the other hand, we have

$$\mathbb{C}[G] = \prod_{i=1}^{k} \mathrm{Mat}_{n_i}(\mathbb{C}).$$

The only way for the right-hand side to be commutative is if $\mathrm{Mat}_{n_i}(\mathbb{C})$ is commutative for all i. But this is true if and only if $n_i = 1$ for all i. By the above theorem, the dimensions of the irreducible representations of G are precisely the n_i, so the irreducible representations must be 1-dimensional. Moreover, the above theorem tells us that

$$|G| = \sum_{i=1}^{k} n_i = \sum_{i=1}^{k} 1 = k.$$

This implies that the number of irreducible representations of G is $k = |G|$. □

We will assume from now on that all groups are finite and all representations are over $\mathbb{C}$ unless stated otherwise. In light of the above discussion, our main goal will be to understand the irreducible representations of groups, because direct sums of these form all other representations. We will develop many tools to do this along the way.

The first main strategy when understanding irreducible representations is to consider the 1-dimensional representations separately. (Notice that all 1-dimensional representations are automatically irreducible: any subrepresentation is a subspace and hence either has dimension 0, in which case it is trivial, or has dimension 1, in which case it is the entire space.) It turns out that these have a much nicer structure than the general irreducible represenations of a finite group, as the following proposition demonstrates.

Proposition 2.1.7. *Let $\varphi : G \to \mathrm{GL}_1(\mathbb{C}) \cong \mathbb{C}^\times$ be a 1-dimensional representation of a finite group G. Then, there is a unique map $\widetilde{\varphi} : G/[G,G] \to \mathbb{C}^\times$ such that the following diagram commutes:*

$$\begin{tikzcd} G \arrow[r, "\varphi"] \arrow[d, "q"] & \mathbb{C}^\times \\ G/[G,G] \arrow[ur, "\widetilde{\varphi}"] & \end{tikzcd}$$

Moreover, φ sends every element of G to a $|G|$th root of unity, and the number of 1-dimensional representations of G is $|G/[G,G]|$.

Proof. Recall that every element of $[G,G]$ is of the form $ghg^{-1}h^{-1}$ for some g and h in G. For any such element, we have

$$\varphi(ghg^{-1}h^{-1}) = \varphi(g)\varphi(h)\varphi(g)^{-1}\varphi(h)^{-1} = 1,$$

where the final equality holds because the target $\mathbb{C}^\times$ of φ is abelian. This implies that φ factors through a unique map $\widetilde{\varphi} : G/[G,G] \to \mathbb{C}^\times$.

Now, $\mathbb{C}^\times$ is isomorphic as a group to $S^1 \times \mathbb{R}_{>0}$ via the identification of the complex plane minus the origin with the union of all circles around the origin of positive radius. (Explicitly, the isomorphism $\mathbb{C}^\times \to S^1 \times \mathbb{R}_{>0}$ is given by $c \mapsto (c/|c|, |c|)$.) Let $|G| = n$. Then, for any g in G, $\varphi(g)^n = \varphi(g^n) = 1$, so $\varphi(g)$ is an element of multiplicative order n in $S^1 \times \mathbb{R}_{>0}$. The only such elements are those of the form $(\zeta, 1)$, where ζ is an nth root of unity. This point corresponds to the point $1 \cdot \zeta = \zeta$ in $\mathbb{C}^\times$, so that $\varphi(g) = \zeta$ is an nth root of unity.

Since $G/[G,G]$ is abelian, by the Fundamental Theorem of Finitely Generated Abelian Groups, we can write $G/[G,G] = \oplus_{i=1}^r C_{p_i}$. So, any homomorphism $\widetilde{\varphi} : G/[G,G] \to \mathbb{C}^\times$ is determined by where it sends a generator for each of the $\mathbb{Z}/p_i\mathbb{Z}$. Since this generator has order p_i in G, it must be sent to an element of order p_i in $\mathbb{C}^\times$, which by the above arguments must be a p_ith root of unity. Any p_ith root of unity will define a valid homomorphism, however, and there are exactly p_i of them, so the number of possible choices for $\widetilde{\varphi}$ is $\prod_{i=1}^r p_i$. Notice that this product is precisely $|G/[G,G]|$. Moreover, one can check that any $\widetilde{\varphi}$ determines a unique $\varphi : G \to \mathbb{C}^\times$. (This is actually part of a more general fact, namely that the inclusion functor of the category of abelian groups into the category of groups is adjoint to the abelianization functor; for a proof, see Homework 10.) Since there are exactly $|G/[G,G]|$ choices of $\widetilde{\varphi}$, there are exactly $|G/[G,G]|$ choices of φ. □

Example. Let $G = S_n$. Then, it is a fact from group theory that $[S_n, S_n] = A_n$. So, for any homomorphism $\varphi : S \to \mathbb{C}^\times$, $\varphi(A_n) = 1$. Moreover, $S_n/A_n \cong \mathbb{Z}/2$, so the above theorem tells us that there are precisely $|\mathbb{Z}/2| = 2$ 1-dimensional representations of S_n. The first is given by sending S_n/A_n identically to 1; since we must send every element of A_n to 1, this corresponds to sending every element of S_n to 1, which gives the trivial representation. The other 1-dimensional representation is given by sending the identity in S_n/A_n to 1 and the non-identity element to -1. This corresponds to sending elements of A_n to 1 and elements outside A_n to -1, or equivalently, sending each element σ of S_n to its sign $\mathrm{sgn}(\sigma)$. We call this representation the *sign representation* of S_n.

2.2 Intertwiners

Now that we've defined what a representation is, we may be interested in what maps between representations look like. The following definition describes this notion explicitly.

Definition 2.2.1. Let V and W be two representations of a group G over a field k. Then, we say that a map $\psi : V \to W$ is a *homomorphism of representations* or *intertwiner* if ψ is a homomorphism of $k[G]$-modules (or equivalently, a linear transformation which commutes with the action of G).

Given any intertwiner $\varphi : V \to W$, because φ commutes with the group action, one can check that the $\ker \varphi$ is a subrepresentation of V and $\operatorname{Im} \varphi$ is a subrepresentation of W.

Although there is not much to say about intertwiners in general, we will now establish a couple of basic results which will be of use to us later.

Proposition 2.2.2. *Let V and W be representations of a finite group G, and let $f : V \to W$ be a linear map. Then, $\frac{1}{|G|}\sum_{g\in G} gfg^{-1}$ is an intertwiner.*

Proof. The map in question is clearly linear, so we just need to check that it commutes with the action of G. For any h in G and v in V, we have

$$h\frac{1}{|G|}\sum_{g\in G} gf(g^{-1}v) = \frac{1}{|G|}\sum_{g\in G} hgf(g^{-1}v) = \frac{1}{|G|}\sum_{k\in G} kf(k^{-1}hv),$$

where the final equality here is given by the substitution $k = hg$. This proves that $h\frac{1}{|G|}\sum_{g\in G} gfg^{-1} = \frac{1}{|G|}\sum_{g\in G} gfg^{-1}h$, as desired. ☐

Remark. Notice that the construction of the intertwiner in this proposition closely resembles the proof of Maschke's Theorem.

Proposition 2.2.3. *Let V be an irreducible representation of a group G, and let $\varphi : V \to V$ be an intertwiner. Then, φ is a scalar multiple of the identity map.*

Proof. φ is a linear transformation of a finite-dimensional vector space over $\mathbb{C}$, so it has some eigenvalue λ. (This follows from the fact that the characteristic polynomial of φ must have a complex root.) Then, $\varphi - \lambda\,\mathrm{id}_V$ is again an intertwiner, since φ and id_V are intertwiners. Moreover, the kernel K of $\varphi - \lambda\,\mathrm{id}_V$ is a subrepresentation of V which is nonzero: it contains the span of the eigenvector corresponding to λ. Since V is irreducible, we must then have $K = V$, which implies that $\varphi = \lambda\,\mathrm{id}_V$. ☐

2.3 Inner Products and Complete Reducibility

Throughout this section, we will assume that all vector spaces are finite-dimensional over $\mathbb{C}$. We first recall a definition from linear algebra.

Definition 2.3.1. Let V be a vector space over $\mathbb{C}$. Then, an *inner product* on V is a map $\langle\cdot,\cdot\rangle : V \times V \to \mathbb{C}$ which is:

1. conjugate symmetric, i.e. $\langle v, w\rangle = \overline{\langle w, j\rangle}$ for all v and w in V;

2. $\mathbb{C}$-linear in the first coordinate, i.e. $\langle u+v, w\rangle = \langle u, w\rangle + \langle v, w\rangle$ and $\langle cu, v\rangle = c\langle u, v\rangle$ for all u, v, and w in V and all c in $\mathbb{C}$;

3. $\mathbb{C}$-antilinear in the second coordinate, i.e. $\langle u, v + w\rangle = \langle u, v\rangle + \langle u, w\rangle$ and $\langle u, cv\rangle = \overline{c}\langle u, v\rangle$ for all u, v, and w in V and any c in $\mathbb{C}$; and

4. positive-definite, i.e. $\langle v, v\rangle \geq 0$ for all v in V, with equality if and only if $v = 0$.

An *inner product space* is a pair $(V, \langle\cdot,\cdot\rangle)$ consisting of a vector space V and an inner product $\langle\cdot,\cdot\rangle$ on V.

Remark. One can show that the first two properties in the above definition actually imply the third. Thus, when checking that something is an inner product, one need not check for antilinearity.

Example. Given a vector space V, suppose that we have fixed a basis for V. Then, for any v and w in V, we can write $v = (v_i)$ and $w = (w_i)$. We can define an inner product on V by setting $\langle v, w \rangle = \sum_{i=1}^n v_i \overline{w_i}$. In the case where the basis is the standard orthonormal basis for $\mathbb{C}$, this inner product is simply the dot product on $\mathbb{C}$.

Given a vector space V and an inner product on V, by linearity of the inner product, it is specified by what it does to a basis $\{v_1, \ldots, v_n\}$ of V. Let $b_{ij} = \langle v_i, v_j \rangle$, and define B to be the matrix (b_{ij}). Then, by conjugate symmetry of the inner product, we have $B^T = \overline{B}$, or equivalently $B^* = B$, where $*$ denotes the conjugate transpose. This leads us to the following definition.

Definition 2.3.2. We say that an $n \times n$ matrix A is *Hermitian* if $A^* = A$.

Now, Recall that any (finite-dimensional) inner product space $(V, \langle \cdot, \cdot \rangle)$ has an orthonormal basis given by Gram–Schmidt orthonormalization. In this basis, one can check that, by linearity of the inner product, the inner product must look like the one described in the above example: that is, for any $v = (v_i)$ and $w = (w_i)$ in V, we have

$$\langle v, w \rangle = \sum_{i=1}^n v_i \overline{w_i} = v^T \overline{w}$$

(where in the last equality we are thinking of v and w as column vectors). In this context, we would like to consider linear transformations $A : V \to V$ which respect the inner product, i.e. such that $\langle Av, Aw \rangle = \langle v, w \rangle$ for all v and w in V. Combining this with the above equation, we have

$$\langle Av, Aw \rangle = (Av)^T \overline{Aw} = v^T A^T \overline{A} \overline{w} = \langle v, w \rangle = v^T \overline{w}.$$

This only way this can be true for all v and w is if $A^T \overline{A} = I$, or equivalently, if $A^* A = I$. This prompts the following definition.

Definition 2.3.3. We say that a linear transformation $T : V \to V$ of a vector space V is *unitary* if $\langle Tv, Tw \rangle = \langle v, w \rangle$ for all v and w in V. We say that an $n \times n$ matrix A is *unitary* if $A^* A = I$. Finally, we define the *unitary group* $\mathrm{U}(n)$ to be the group of all unitary $n \times n$ matrices (with operation given by multiplication).

Remark. By the above discussion, we have that a linear transformation is unitary if and only if the matrix which represents it in an orthonormal basis is unitary.

Example. In the 1-dimensional vector space $\mathbb{C}$, one can check using the standard inner product on $\mathbb{C}^n$ and the definition of unitary that $\mathrm{U}(1) \cong S^1$ is the subset of $\mathbb{C}$ consisting of all elements of norm 1.

We now wish to import this discussion of unitary transformations into the realm of representation theory. To do this, we require the following definition.

Definition 2.3.4. Suppose V is a representation of a group G, and fix some inner product on V. Then, we say that the representation is *unitary* if the action of G on V preserves the inner product: that is, for all g in G and all $v, w \in V$, $\langle gv, gw \rangle = \langle v, w \rangle$.

Remark. Thinking of a representation as a homomorphism $G \to \mathrm{GL}(V)$, the statement that a representation $\varphi : G \to \mathrm{GL}(V)$ is unitary is precisely the statement that the elements of g act by unitary matrices on V, or in other words, that $\varphi(G) \subseteq \mathrm{U}(V)$, the subgroup of $\mathrm{GL}(V)$ consisting of unitary matrices.

The main significance of a unitary representation for representation theory is given by the following theorem.

Theorem 2.3.5. *If a representation V of a group G is unitary with respect to some inner product $\langle \cdot, \cdot \rangle$, then φ is completely reducible.*

Proof. Given a subrepresentation W of V, consider the orthogonal complement $W^\perp$ of W. We know that $V \cong W \oplus W^\perp$ as vector spaces. Moreover, for any g in G, v in $W^\perp$, and w in W, we have

$$\langle gv, w \rangle = \langle g^{-1}gv, g^{-1}w \rangle = \langle v, g^{-1}w \rangle = 0,$$

where the first equality follows by the fact that V is unitary and the thir dequality follows from the fact that $g^{-1}w \in W$. This proves that $gv \in W^\perp$, so that $W^\perp$ is fixed by the action of G and hence is a subrepresentation of V. So, $W^\perp$ is the complimentary subrepresentation to W. □

In the case of complex representations of finite groups, the situation becomes quite simple: it turns out that every representation is unitary.

Theorem 2.3.6. *Let V be a complex representation of a finite group G. Then, V is unitary with respect to some inner product on $\mathbb{C}^n$.*

Proof. Fix any inner product $\langle \cdot, \cdot \rangle'$ on $\mathbb{C}^n$. Now, each element g of G defines an inner product $\langle \cdot, \cdot \rangle_g$ on $\mathbb{C}^n$ by setting $\langle v, w \rangle_g = \langle gv, gw \rangle$ for all v and w in V. We define $\langle v, w \rangle = \sum_{g \in G} \langle v, w \rangle_g$. One can check that $\langle \cdot, \cdot \rangle$ is an inner product on $\mathbb{C}^n$. Moreover, for any $h \in G$ and any v and w in V, we have

$$\langle hv, hw \rangle = \sum_{g \in G} \langle ghv, ghw \rangle' = \sum_{k \in G} \langle kv, kw \rangle' = \langle v, w \rangle.$$

(Here the penultimate inequality follows by setting $k = gh$.) So, the representation V is unitary with respect to $\langle \cdot, \cdot \rangle$. □

Notice that the above theorem gives another proof of Maschke's Theorem: any complex representation of a finite group is unitary, so by Theorem 2.3.5, it is completely reducible.

2.4 Character Theory

We know that representations can be thought of as group homomorphisms $G \to \mathrm{GL}(V)$. Since each element of $\mathrm{GL}(V)$ is a linear transformation of V, we are interested in characterizing these linear transformations somehow. If we choose a basis for V, we can write them as a matrix, but this representation depends on the basis. We would rather understand our linear transformations more intrinsically, through some notion that does not depend on a choice of basis. For this, we use a standard invariant from linear algebra: the trace of a matrix.

Definition 2.4.1. Let $G \to \mathrm{GL}(V)$ be a representation which sends an element g of G to the linear transformation corresponding to the matrix A_g (in some fixed basis of V). Then, we define the *character* of the representation to be the map $\chi_V : G \to \mathbb{C}$ given by $\chi_V(g) = \mathrm{tr}(A_g)$.

Example. Let G be any group, and consider the trivial representation V of G, i.e. the 1-dimensional representation in which every element of G acts by the identity, which in 1 dimension simply corresponds to multiplication by 1. So, we have $\chi_V(g) = \mathrm{tr}(1) = 1$ for all g in G.

Since the character of a representation does not depend on the basis of V, we can compute it in any basis we like. This allows us to pick whatever basis is most convenient to us. This technique will allows us to establish many basic properties of characters. First, however, we require a standard result from linear algebra, which we will prove using the techniques of representation theory.

Proposition 2.4.2. *Let $A \in \mathrm{GL}_n(\mathbb{C})$ such that $A^m = I$ for some m. Then, A is conjugate to a diagonal matrix.*

Proof. Define a representation V of the cyclic group C_m by picking a generator g of C_m and setting $g \mapsto A$. (One can check that this does, in fact, define a homomorphism $C_m \to \mathrm{GL}_n(\mathbb{C})$.) Now, by Maschke's Theorem, V is completely reducible. Moreover, by Proposition 2.1.6, the irreducible representations of C_m are all 1-dimensional. So, we can write $V = \oplus_{i=1}^n W_i$, where the $W_i \cong \mathbb{C} \cdot w_i$ are 1-dimensional representations of C_m. Changing bases from the one in which A is defined to $\{w_i\}_{i=1}^n$ (that this is a basis follows from the fact that the direct sum of the W_i is V) amounts to conjugating A by some element of $\mathrm{GL}_n(\mathbb{C})$; call the resulting matrix B. Then, the action of g fixes W_i for all i, so gw_i is some multiple of w_i, which implies that the ith column of B has 0's everywhere except in the ith row (i.e. the coefficient of w_j in gw_i is 0 for all $j \neq i$). This implies that B is diagonal. $\square$

We can now prove some basic properties of characters.

Proposition 2.4.3. *Let V be a complex representation of a (not necessarily finite) group G.*

1. *For all g and h in G,*
$$\chi_V(hgh^{-1}) = \chi_V(g).$$

In other words, the value of the character on an element g of G depends only on the conjugacy class of g.

2. *For any other complex representation W of G,*

$$\chi_{V\oplus W}(g) = \chi_V(g) + \chi_W(g)$$

for all g in G.

3. $\chi_V(1) = \dim V$.

4. *If $|G| = n$ is finite, then for any g in G,*

$$\chi_V(g) = \zeta_1 + \cdots + \zeta_n,$$

where ζ_i is an nth root of unity for all i. In particular, $|\chi_V(g)| \leq n$, with equality if and only if g acts by the identity on V.

5. *If G is finite, then for any g in G, $\chi_V(g^{-1}) = \overline{\chi_V(g)}$.*

Proof.

1. Let A and B be matrices corresponding to the actions of g and h (respectively) on V. Because conjugation preserves the trace of a matrix, we have

$$\chi_V(hgh^{-1}) = \operatorname{tr}(BAB^{-1}) = \operatorname{tr}(A) = \chi_V(g).$$

2. Fix a basis $B = \{v_i\}_{i=1}^m$ for V and a basis $B' = \{w_j\}_{j=1}^n$ for W. Then, $B \cup B'$ is a basis for $V \oplus W$ (here we are thinking of the v_i and w_j as lying inside of $V \oplus W$ via the inclusion maps $V \hookrightarrow V \oplus W$ and $W \hookrightarrow V \oplus W$). For any g in G, if S and T are matrices corresponding to the action of g on V and W (respectively) in the bases B and B' (respectively), then the matrix corresponding to the action of g on $V \oplus W$ in the basis $B \cup B'$ is

$$\left(\begin{array}{c|c} S & 0 \\ \hline 0 & T \end{array}\right).$$

Clearly the trace of this matrix is $\operatorname{tr}(S)+\operatorname{tr}(T)$, which implies that $\chi_{V\oplus W}(g) = \chi_V(g) + \chi_W(g)$.

3. Any homomorphism $G \to \mathrm{GL}(V)$ must send 1 to the identity matrix I. This means that 1 acts by the identity in every representation, so we have $\chi_V(1) = \operatorname{tr}(I) = \dim V$.

4. Suppose that $\varphi : G \to \mathrm{GL}(V)$ is the homomorphism corresponding to the representation V. Then, for any g in G, we have

$$\varphi(g)^n = \varphi(g^n) = \varphi(1) = I.$$

So, we may apply Proposition 2.4.2 to conclude that $\varphi(g)$ is conjugate to a diagonal matrix D. Since conjugation preserves the trace of a matrix, we may as well use D to compute $\chi_V(g)$. Now, conjugation does not affect multiplicative order, so $D^n = 1$. Since the product of diagonal matrices is simply the diagonal matrix consisting of the pairwise products of the diagonal elements, this implies that each element on the diagonal of D is an nth root of unity; call them $\zeta_1, \zeta_2, \ldots, \zeta_n$. Then, we have $\chi_V(g) = \mathrm{tr}(D) = \zeta_1 + \zeta_2 + \cdots + \zeta_n$. The fact that $|\zeta_1 + \zeta_2 + \cdots + \zeta_n| \leq n$, with equality if and only if $\zeta_i = 1$ for all i, follows from the properties of roots of unity; for a proof, see Homework 7. Finally, notice that $\zeta_i = 1$ for all i if and only if $D = I$, in which case g acts by the identity in the basis in which its action is given by D (and hence in every other basis, since I is fixed by conjugation).

5. By the arguments in point 4 above, g acts in some basis by a diagonal matrix D whose diagonal entries are some nth roots of unity $\zeta_1, \ldots, \zeta_n$. In this basis, then, g^{-1} acts by D^{-1}, which has as its diagonal entries $\zeta_1^{-1}, \ldots, \zeta_n^{-1}$. Using the fact that $\zeta^{-1} = \overline{\zeta}$ for any root of unity ζ, we have

$$\chi_V(g^{-1}) = \zeta_1^{-1} + \cdots + \zeta_n^{-1} = \overline{\zeta_1} + \cdots + \overline{\zeta_n} = \overline{\zeta_1 + \cdots + \zeta_n} = \overline{\chi_V(g)}.$$

□

In light of this proposition, we make the following definition.

Definition 2.4.4. Let G be a group, and let $f : G \to \mathbb{C}$ be a map such that, for and g and h in G, $f(hgh^{-1}) = f(g)$. Then, we say that f is a *class function* on G.

Notice that the character of any representation is a class function, by the above proposition. Moreover, we can consider the set of class functions as vectors in a $\mathbb{C}$-vector space, with addition and scalar multiplication given by

$$(f_1 + f_2)(g) = f_1(g) + f_2(g)$$

and

$$(cf)(g) = c \cdot f(g).$$

This vector space has a natural inner product defined by

$$\langle f_1, f_2 \rangle = \frac{1}{|G|} \sum_{g \in G} f_1(g)\overline{f_2(g)}.$$

It is often convenient to write this inner product in another form. Let $\mathcal{O}_1, \ldots, \mathcal{O}_m$ denote the conjugacy classes of G, and let g_i be a representative of $\mathcal{O}_i$ for all i. Then, G is the disjoint union of the $\mathcal{O}_i$, and any class functions f_1 and f_2 take the same value on all of $\mathcal{O}_i$, so we have

$$\langle f_1, f_2 \rangle = \frac{1}{|G|} \sum_{i=1}^{m} |\mathcal{O}_i| f_1(g_i)\overline{f_2(g_i)} = \sum_{i=1}^{m} \frac{f_1(g_i)\overline{f_2(g_i)}}{|C_G(g_i)|},$$

where $C_G(g_i)$ is the centralizer of g_i. (The last inequality here follows from the fact that $|\mathcal{O}_i||C_G(g_i)| = |G|$, which is essentially an application of Lagrange's Theorem.)

Now, irreducible representations have a very nice structure which we might hope to translate into constraints on their characters. In order to do this, it will be beneficial to consider the characters of irreducible representations as vectors in the vector space of class functions. The following theorem gives us our first nice result about these vectors.

Theorem 2.4.5. *Let G be a finite group, and let $\varphi_1 : G \to \mathrm{GL}(V)$ and $\varphi_2 : G \to \mathrm{GL}(W)$ be two irreducible representations of G. Let χ_i be the character corresponding to φ_i for $i \in \{1, 2\}$. If φ_1 and φ_2 are non-isomorphic, then*

$$\langle \chi_1, \chi_2 \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_1(g)\overline{\chi_2(g)} = 0.$$

If φ_1 and φ_2 are isomorphic, then

$$\langle \chi_1, \chi_2 \rangle = \langle \chi_1, \chi_1 \rangle = 1.$$

In other words, the set of irreducible representations of G forms an orthonormal set in the space of class functions.

Proof. Let $\dim V = m$ and $\dim W = n$, and define e_{ij} to be the $n \times m$ matrix with a 1 in the (i, j)th position and 0's everywhere else. First, suppose that V and W are non-isomorphic. Then, for all i and j, e_{ij} defines a linear map $V \to M$, so by Proposition 2.2.2, the matrix $f_{ij} = \frac{1}{|G|}\sum_{g \in G} \varphi_1(g)e_{ij}\varphi_2(g^{-1})$ corresponds to an intertwiner $V \to W$. By Schur's Lemma, this intertwiner is either the zero map or an isomorphism; since we've assumed that V and W are non-isomorphic, it must be the zero map. So, f_{ij} is the zero matrix for all i and j. In particular, this implies that the (i, j)th entry in f_{ij} is 0: that is,

$$0 = \frac{1}{|G|} \sum_{g \in G} \varphi_1(g)_{ii}\varphi_2(g^{-1})_{jj}.$$

Summing this equation together over all i and j, we get

$$\begin{aligned} 0 = \frac{1}{|G|} \sum_{g \in G} \sum_{i,j} \varphi_1(g)_{ii}\varphi_2(g^{-1})_{jj} &= \frac{1}{|G|} \sum_{g \in G} \sum_{j} \chi_1(g)\varphi_2(g^{-1})_{jj} \\ &= \frac{1}{|G|} \sum_{g \in G} \chi_1(g)\chi_2(g^{-1}) \\ &= \frac{1}{|G|} \sum_{g \in G} \chi_1(g)\overline{\chi_2(g)} = \langle \chi_1, \chi_2 \rangle. \end{aligned}$$

(Here the penultimate equality follows from the above proposition.)

Now, suppose that φ_1 and φ_2 are isomorphic. Then, there exists an isomorphism $V \cong W$ which respects the action of G, from which one can deduce that $\chi_1 = \chi_2$. Thus, we may as well just work with χ_1. This time, we define $f_{ij} = \frac{1}{|G|}\sum_{g \in G} \varphi_1(g)e_{ij}\varphi_1(g^{-1})$. f_{ij} is an intertwiner by Proposition 2.2.2, so by

Proposition 2.2.3, $f_{ij} = \lambda I$ for some $\lambda \in \mathbb{C}$. In particular, $\operatorname{tr}(f_{ij}) = m\lambda$. On the other hand, since trace is invariant under conjugation,

$$\operatorname{tr}(\varphi_1(g)e_{ij}\varphi_1(g^{-1})) = \operatorname{tr}(\varphi_1(g)e_{ij}\varphi_1(g)^{-1}) = \operatorname{tr}(e_{ij}).$$

By additivity of the trace, then,

$$m\lambda = \operatorname{tr}(f_{ij}) = \frac{1}{|G|}\sum_{g\in G}\operatorname{tr}(\varphi_1(g)e_{ij}\varphi_1(g^{-1})) = \frac{1}{|G|}\sum_{g\in G}\operatorname{tr}(e_{ij}) = \operatorname{tr}(e_{ij}) = \delta_{ij}.$$

So, when $i = j$, $\lambda = 1/m$ and $f_{ii} = (1/m)\operatorname{id}_V$. In particular, the (i,i)th entry of f_{ii} is $1/m$. Likewise, when $i \neq j$, $\lambda = 0$, and the (i,j)th entry of f_{ij} is 0. In equations:

$$\frac{1}{|G|}\sum_{g\in G}\varphi_1(g)_{ii}\varphi_2(g^{-1})_{jj} = \begin{cases} \frac{1}{m}, & i = j \\ 0, & i \neq j \end{cases}$$

Summing this equation over all i and j, we get

$$1 = \frac{1}{|G|}\sum_{g\in G}\sum_{i,j=1}^{n}\varphi_1(g)_{ii}\varphi_1(g^{-1})_{jj} = \frac{1}{|G|}\sum_{g\in G}\chi_1(g)\chi_1(g^{-1}) = \langle \chi_1, \chi_1\rangle.$$

$\square$

Because they are statements about the orthonormality of characters in the vector space of class functions, the equations given by the above theorem are often called the *orthogonality relations of the first kind.* They are useful to us for several different reasons. First, using some clever linear algebra arguments, we can obtain a similar set of equations, which are called the *orthogonality relations of the second kind.*

Theorem 2.4.6. *Let G be a finite group, let $V_1, \ldots, V_m$ be the irreducible representations of G, and let χ_i be the character of V_i for all i. Then, given any g and h in G and any i, we have*

$$\sum_{i=1}^{m}\chi_i(g)\overline{\chi_i(h)} = \begin{cases} |C_G(g)|, & g \text{ is conjugate to } h \\ 0, & \text{otherwise} \end{cases}.$$

Proof. Pick representatives $g_1, \ldots, g_m$ for the conjugacy classes of G. It suffices to prove that

$$\sum_{i=1}^{m}\chi_i(g_j)\overline{\chi_i(g_k)} = \delta_{jk}|C_G(g)|.$$

To see this, notice that there is some g_j representating the conjugacy class of g and some g_k representing the conjugacy class of h. So, the above equation along with the fact that $\chi_i(g_j) = \chi_i(g)$ and $\chi_i(g_k) = \chi_i(h)$ yields the desired result.

Now, define a matrix $C = (c_{ij})$ by setting

$$c_{ij} = \frac{\chi_i(g_j)}{\sqrt{|C_G(g_i)|}}.$$

Notice that the orthogonality relations of the first kind, along with the reformulation of the inner product we discussed above, give us

$$\langle \chi_i, \chi_j \rangle = \sum_{k=1}^{m} \frac{\chi_i(g_k)\overline{\chi_j(g_k)}}{|C_G(g_k)|} = \delta_{ij}.$$

On the other hand, the terms of this sum are of the form $c_{ik}\overline{c_{jk}}$, so that the sum itself is precisely the (i,j)th coordinate of CC^*. So, the above equation implies that $CC^* = I$, i.e. that C is a unitary matrix. Then, C^T is also a unitary matrix, since $C^T(C^T)^* = C^T\overline{C} = (C^*C)^T = I^T = I$. Checking the individual entries of this matrix equation gives

$$\frac{1}{\sqrt{|C_G(g_j)||C_G(g_k)|}} \sum_{i=1}^{m} \chi_i(g_j)\overline{\chi_i(g_k)} = \delta_{jk},$$

which implies the desired result. □

We can also use the orthogonality relations to prove that characters determine representations up to isomorphism.

Proposition 2.4.7. *Let G be a finite group, and let V and W be two representations of G. Then, $V \cong W$ if and only if $\chi_V = \chi_W$.*

Proof. Let $V_1, \ldots, V_k$ be the irreducible representations of G, and let χ_i be the character of V_i for all i. Then, we can write $V \cong \oplus_{i=1}^k V_i^{m_i}$ and $W \cong \oplus_{j=1}^k V_i^{n_i}$. So, Proposition 2.4.3 implies that $\chi_V = \sum_{i=1}^k m_i\chi_i$ and $\chi_W = \sum_{i=1}^r n_i\chi_i$. Now, if V and W are isomorphic, then their decompositions into irreducible representations must be the same, so that $m_i = n_i$ for all i. This implies that the above expressions for χ_V and χ_W are the same, so that $\chi_V = \chi_W$.

Conversely, suppose that $\chi_V = \chi_W$. Notice that, for any i, we have by linearity of the inner product

$$\langle \chi_V, \chi_i \rangle = \langle \sum_{j=1}^{k} m_j\chi_j, \chi_i \rangle = \sum_{j=1}^{k} m_j \langle \chi_j, \chi_i \rangle = m_i,$$

where the last inequality follows from the orthogonality relations of the first kind. An identical calculation implies that $\langle \chi_W, \chi_i \rangle = n_i$. Since $\chi_V = \chi_W$, we must have $\langle \chi_V, \chi_i \rangle = \langle \chi_W, \chi_i \rangle$ for all i, i.e. $m_i = n_i$ for all i. But this implies that $V \cong W$. □

This proposition encapsulates a lot of the reason that characters are useful for studying representations. There are many situations where using facts about characters can allow us to easily write down the values of characters of irreducible representations even without knowing what the representations are. By the above proposition, a representation is determined by its character, so writing down these characters does specify the irreducible representations of a group even if we don't know what the representations are explicitly.

There are a couple of other nice facts that we can obtain about inner products of characters. The first concerns the character of the regular representation $\mathbb{C}[G]$ of a finite group G (i.e. $\mathbb{C}[G]$ considered as a module over itself).

Proposition 2.4.8. *Let V be any representation of a finite group G. Then,*

$$\langle \chi_V, \chi_{\mathbb{C}[G]} \rangle = \dim V.$$

Proof. Notice that $\mathbb{C}[G]$ acts on itself by left multiplication. For any $g \neq 1$ in G, left multiplication on G by g is injective; since the elements of G form a basis of $\mathbb{C}[G]$ over G, this implies that $\chi_{\mathbb{C}[G]}(g) = 0$. On the other hand, 1 fixes every element of G, so that $\chi_{\mathbb{C}[G]}(1) = |G|$. So, for any representation V of G, we have

$$\langle \chi_V, \chi_{\mathbb{C}[G]} \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_V(g) \overline{\chi_{\mathbb{C}[G]}(g)} = \frac{1}{|G|}(\chi_V(1)|G|) = \chi_V(1) = \dim V.$$

□

We can also relate the inner product of two characters to the dimension of a Hom set, which will occasionally we useful to us.

Proposition 2.4.9. *Let V and W be any representations of a finite group G. Then,*

$$\langle \chi_V, \chi_W \rangle = \dim_{\mathbb{C}}(\mathrm{Hom}_G(V, W)).$$

(Here the set $\mathrm{Hom}_G(V, W)$ denotes the set of intertwiners $V \to W$, which is a $\mathbb{C}$-vector space by Proposition 1.4.8.)

Proof. Let $V_1, \ldots, V_k$ denote the irreducible representations of G, and let χ_i be the character of V_i for all i. Then, we can write $V = \oplus_{i=1}^k V_i^{m_i}$ and $W = \oplus_{j=1}^k V_j^{n_j}$. So, by linearity of the inner product along with the orthogonality relations of the first kind, we have

$$\langle \chi_V, \chi_W \rangle = \langle \sum_{i=1}^k m_i \chi_i, \sum_{j=1}^k n_j \chi_j \rangle = \sum_{i,j=1}^k m_i n_j \langle \chi_i, \chi_j \rangle = \sum_{i=1}^k m_i n_i.$$

On the other hand, by Proposition 1.4.12, we have

$$\mathrm{Hom}_G(V, W) = \mathrm{Hom}_G(\oplus_{i=1}^k V_i^{m_i}, \oplus_{j=1}^k V_j^{n_j}) \cong \oplus_{i,j=1}^k \mathrm{Hom}_G(V_i, V_j)^{m_i n_j}.$$

Now, by Schur's Lemma, $\mathrm{Hom}_G(V_i, V_j) = 0$ when $j \neq i$, since V_i and V_j are not isomorphic, and by Proposition 2.2.3, $\mathrm{Hom}_G(V_i, V_i) \cong \mathbb{C}$, since this Hom set consists of scalar multiples of the identity. So, we have $\mathrm{Hom}_G(V, W) \cong \oplus_{i=1}^k \mathbb{C}^{m_i n_i}$. This implies that

$$\dim_{\mathbb{C}}(\mathrm{Hom}_G(V, W)) = \sum_{i=1}^k m_i n_i = \langle \chi_V, \chi_W \rangle.$$

□

Finally, we have a result which allows us to easily identify irreducible representions by taking the inner product of their characters.

Proposition 2.4.10. *Let V be a representation of a finite group G. Then, V is irreducible if and only if $\langle \chi_V, \chi_V \rangle = 1$.*

Proof. Let $V_1, \ldots, V_k$ be the irreducible representations of G, and let χ_i be the character of V_i for all i. Write $V = \oplus_{i=1}^k V_i^{n_i}$. Then, we have

$$\langle \chi_V, \chi_V \rangle = \langle \sum_{i=1}^k n_i \chi_i, \sum_{j=1}^k n_j \chi_j \rangle = \sum_{i,j=1}^k n_i n_j \langle \chi_i, \chi_j \rangle = \sum_{i=1}^k n_i^2.$$

This sum is 1 if and only if $n_i = 1$ for some i and $n_j = 0$ for all $j \neq i$, which is true if and only if $V \cong V_i$ for some i, i.e. if and only if V is irreducible. □

2.5 Computing Character Tables

We have seen that all representations of a finite group G are direct sums of irreducible representations and, moreover, that representations are determined by their characters. These two facts together imply that the characters of the irreducible representations of G encapsulate the information about every complex representation of G. For this reason, it is often very useful to compute the characters of the irreducible representations of a group. We formulate these characters succinctly using a *character table*, which gives the value of the character of each irreducible representation on each conjugacy class of a group. (Recall that characters are constant on conjugacy classes, so this information determines the character everywhere.)

Perhaps the easiest way to describe this process is to give an example.

Example. Now, consider the cyclic group C_4. Let g be a generator of C_4. Then, C_4 is abelian, so each element constitutes its own conjugacy class. Moreover, by Proposition 2.1.6, every irreducible representation of C_4 is 1-dimensional. Each such representation is determined by where it sends g, and by Proposition 2.1.7, it must send g to a 4th root of unity. Any 4th root of unity will work, so each of the 4 4th roots of unity (1, -1, i, and $-i$) yields its own irreducible representation of C_4. By Proposition 2.1.6, this gives us all $|C_4| = 4$ irreducible representations of C_4. So, the character table of C_4 is:

	1	g	g^2	g^3
χ_0	1	1	1	1
χ_1	i	-1	$-i$	
χ_2	1	-1	1	-1
χ_3	1	$-i$	-1	i

In the above case, all of the irreducible representations were 1-dimensional, so our understanding of 1-dimensional representations via Propositions 2.1.6 and 2.1.7 suffices to determine all the irreducible representations. When this is not the case, we have to use more of the representation theory we've developed, as the following example demonstrates.

Example. Consider the group S_3. Recall that conjugacy classes in any symmetric group are determined by cycle type. So, S_3 has 3 conjugacy classes: they are represented by 1, (12), and (123). By the example at the very end of Section 2.1, S_3 has two 1-dimensional representations: the trivial representation and the sign representation. Let the characters of these be χ_0 and χ_1, respectively. From the definitions of these representations, one can immediately see that χ_0 is constantly 1 everywhere, while χ_1 sends (12) to -1 and 1 and (123) to 1.

Now, we know that there is one irreducible representation of S_3 for each conjugacy class, so we need to find one more irreducible representation. Let its character by χ_2. We know that $\chi_2(1)$ is the dimension of this last irreducible representation. Moreover, by Theorem 2.1.4, the sum of the squares of the dimensions of the irreducible representations is $|S_3| = 6$. Since the other two irreducible representations are 1-dimensional, this last irreducible representation must have dimension 2, so that $\chi_2(1) = 2$.

So far, our character table for S_3 is:

	1	(12)	(123)
χ_0	1	1	1
χ_1	1	-1	1
χ_2	2		

We can now complete the table using the orthogonality relations of the second kind. When applied to the conjugacy classes of 1 and (12), these relations imply that

$$\chi_0(1)\overline{\chi_0(12)} + \chi_1(1)\overline{\chi_1(12)} + \chi_2(1)\overline{\chi_2(12)} = 0,$$

since 1 and (12) are not in the same conjugacy class. Plugging in the values we know here:

$$1 - 1 + 2\overline{\chi_2(12)} = 0,$$

which implies that $\chi_2(12) = 0$. One can do likewise for the conjugacy classes of 1 and (123) to find $\chi_2(123)$; alternately, applying the orthogonality relations of the first kind to χ_1 and χ_2 gives

$$\begin{aligned} 0 = \frac{1}{|S_3|}\sum_{\sigma \in S_3} \chi_1(\sigma)\overline{\chi_2(\sigma)} &= \frac{1}{6}\left(\chi_1(1)\overline{\chi_2(1)} + 3\chi_1(12)\overline{\chi_2(12)} + 2\chi_1(123)\overline{\chi_2(123)}\right) \\ &= \frac{1}{6}\left(2 + 0 + 2\overline{\chi_2(123)}\right) \end{aligned}$$

From this, one can deduce that $\chi_2(123) = -1$. Putting everything together, then, the character table for S_3 is:

	1	(12)	(123)
χ_0	1	1	1
χ_1	1	-1	1
χ_2	2	0	-1

Notice the use of Theorem 2.1.4 as well as the orthogonality relations in this example to obtain information about the characters of representations. Even though we have not explicitly described the 2-dimensional irreducible representation of S_3, we were still able to compute its character using this information. Notice, moreover, the nice relationship between the orthogonality relations and the character table: the orthogonality relations of the first kind give rise to relations between the rows of the table, while the orthogonality relations of the second kind give rise to relations between the columns of the table.

The general procedure for finding irreducible representations of a finite group G is as follows:

1. Compute the conjugacy classes of G and pick representatives of them. We will use these to compute the value of characters of G.

2. Compute $[G, G]$ and use this along with Proposition 2.1.7 to find the 1-dimensional representations of G.

3. Use general knowledge of the structure of irreducible representations of G (for instance, Theorem 2.1.4) to find the dimensions of the remaining irreducible representations of G. This tells us the value of the characters of these representations at 1.

4. Use the orthogonality relations of the second kind (sometimes the first-kind relations are helpful too) to compute the remaining values of the characters of the non-1-dimensional representations.

This procedure is sufficient to compute the character tables of most groups which are small enough for the character table to be easily calculable by hand.

We end this section with one more example, which demonstrates the above procedure. For more examples, see Homeworks 6 and 7 (in which the relevant problems are divided into parts that demonstrate each of the steps of the above procedure).

Example. Consider the dihedral group of order 8, D_4. This group corresponds to the rotations and reflections of a square. A presentation of D_4 is given by $D_4 = \langle a, b \mid a^4 = b^2 = (ab)^2 = 1\rangle$, where a corresponds to a rotation of the square by $\pi/2$ and b corresponds to a reflection of the square about an axis perpendicular to two of its sides. With this presentation, the conjugacy classes of D_4 are represented by 1, a, a^2, b, and ab.

One can check that $[D_4, D_4] = \{1, a^2\}$. So, $|D_4/[D_4, D_4]| = 4$, which means we have 4 1-dimensional representations. In fact, $D_4/[D_4, D_4] \cong C_2 \times C_2$: this quotient is generated by images of a and b, since D_4 is generated by these elements, and the images of both a and b have order 2 in the quotient, so we must have the unique group of order 4 generated by 2 order-2 elements. This implies that any homomorphism $D_4/[D_4, D_4] \to \mathbb{C}^\times$ must send the images of both a and b to 2nd roots of unity, so that the corresponding 1-dimensional representation $D_4 \to \mathbb{C}^\times$ sends a and b to 2nd roots of unity. Both a and b can be sent to either of the 2nd root of unity, and the 4 possibilities here determine the 4 1-dimensional representations. The character table for these 4 representations is:

	1	a	a^2	b	ab
χ_0	1	1	1	1	1
χ_1	1	-1	1	-1	1
χ_2	1	-1	1	1	-1
χ_3	1	1	1	-1	-1

Now, since there are 5 conjugacy classes of D_4 and only 4 1-dimensional representations, there is one more irreducible representation we haven't found yet. By Theorem 2.1.4, this last irreducible representation must have dimension 2. Then, using the orthogonality relations of the second kind, one can fill in the rest of the character table for D_4. The result is:

	1	a	a^2	b	ab
χ_0	1	1	1	1	1
χ_1	1	-1	1	-1	1
χ_2	1	-1	1	1	-1
χ_3	1	1	1	-1	-1
χ_4	2	0	-2	0	0

2.6 Algebraic Integers and the Dimension Theorem

In this section, we use some results of number theory to prove a powerful theorem about the dimensions of irreducible representations. We then apply this theorem to recover some standard results from group theory. We begin with a definition.

Definition 2.6.1. The *algebraic integers*, denoted $\mathbb{A}$, are the subset of $\mathbb{C}$ consisting of all complex numbers which are the root of some monic polynomial with coefficients in $\mathbb{Z}$.

Example.

1. For any $\alpha \in \mathbb{A}$, we have that $-\alpha \in \mathbb{A}$: if $p(x)$ is a polynomial in $\mathbb{Z}[x]$ with α as a root, then $p(-x)$ is a polynomial in $\mathbb{Z}[x]$ with $-\alpha$ as a root.

2. Every nth root of unity is a root of the polynomial $z^n - 1$ and hence is in $\mathbb{A}$.

3. The eigenvalues of any element of $\mathrm{Mat}_n(\mathbb{Z})$ are in $\mathbb{A}$, since the determinant of such a matrix is a polynomial with integer coefficients.

We quickly go through some first properties of the algebraic integers.

Lemma 2.6.2. *Let $y \in \mathbb{C}$. Then, $y \in \mathbb{A}$ if and only if there exist $y_1, \ldots, y_t \in \mathbb{C}$ not all 0 such that $yy_i = \sum_{j=1}^{t} a_{ij} y_i$ for some $a_{ij} \in \mathbb{Z}$.*

Proof. Suppose $y \in \mathbb{A}$. Then, let $p(x) = x^d + a_{d-1}x^{d-1} + \cdots + a_1 x + a_0$ be a polynomial such that $p(y) = 0$ and $a_i \in \mathbb{Z}$ for all i. For all $1 \leq i \leq d$, define $y_i = y^i$, and set $y_0 = 1$j. Now, $yy_i = y_{i+1}$ for all $i < d-1$, and for $i = d-1$,

$$yy_i = y^d = -a_{d-1}y^{d-1} - \cdots - a_1 y - a_0 = -a_{d-1}y_{d-1} - \cdots - a_1 y_1 - a_0 y_0.$$

Conversely, given a_{ij} and y_i as in the proposition, let $A = (a_{ij})$ and $Y = (y_1, \ldots, y_t)$. Then, by assumption, $AY = yY$, which implies that y is an eigenvalue of A. But as discussed in the above example, such an eigenvalue is an element of $\mathbb{A}$. □

Proposition 2.6.3. *$\mathbb{A}$ is a ring.*

Proof. Fix y and z in $\mathbb{A}$. By the above lemma, we can find a_{ij} and y_i such that $yy_i = \sum_j a_{ij} y_i$ and b_{ij} and z_i such that $zz_i = \sum_j b_{ij} z_j$. Then, applying the above lemma with $\{y_i z_j\}_{i,j}$ as the set of y_i and $\{a_{ij}\}_{i,j} \cup \{b_{ij}\}_{i,j}$ as the set of a_{ij} proves that $y + z$ is in $\mathbb{A}$. Likewise, applying the lemma with $\{y_i z_j\}_{i,j}$ as the set of y_i and $\{a_{ij} + b_{i'j'}\}_{i,j,i',j'}$ as the set of a_{ij} proves that yz is in $\mathbb{A}$. □

Proposition 2.6.4. $\mathbb{A} \cap \mathbb{Q} = \mathbb{Z}$.

Proof. Clearly $\mathbb{Z} \subseteq \mathbb{A} \cap \mathbb{Q}$. Conversely, let q be an element of $\mathbb{A} \cap \mathbb{Q}$. Then, there is a monic polynomial $p(x) = x^d + a_{d-1}x^{d-1} + \cdots + a_0$ with $a_i \in \mathbb{Z}$ for all i and $p(q) = 0$. Write $q = m/n$, where $m, n \in \mathbb{Z}$ and $(m, n) = 1$. Then, we have

$$0 = n^{d-1}p(q) = \frac{m^d}{n} + a_{d-1}m^{d-1} + \cdots + n^{d-2}a_1 m + n^{d-1}a_0,$$

or equivalently,

$$\frac{m^d}{n} = -a_{d-1}m^{d-1} - \cdots - n^{d-2}a_1 m - n^{d-1}a_0.$$

The right-hand side of this equation is in $\mathbb{Z}$, since each term is. This implies that the left-hand side must be in $\mathbb{Z}$ as well. But since $(m, n) = 1$, this is only possible if $n = 1$, in which case $q = m \in \mathbb{Z}$. □

Remark. Recall from Galois theory that, given any $\alpha \in \mathbb{A}$ and $\sigma \in \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$, $\sigma(\alpha) \in \mathbb{A}$. (Note that here, $\overline{Q}$ is the Galois closure of $\mathbb{Q}$.) This is essentially because σ permutes the roots of the minimal polynomial of α.

We can relate the algebraic integers to representation theory by understanding when certain operations on values of characters produce algebraic integers.

Proposition 2.6.5. *Let G be a finite group and χ a complex character of G. Then, $\chi(g) \in \mathbb{A}$ for all $g \in G$.*

Proof. $\chi(g)$ is a sum of roots of unity, and we know that all of these are in $\mathbb{A}$ (see the above example). Since $\mathbb{A}$ is a ring, this means that $\chi(g) \in \mathbb{A}$. □

Theorem 2.6.6. *Let V be an irreducible representation over $\mathbb{C}$ of a finite group G. Suppose the dimension of V is d. Let g be an element of G and h be the size of the conjugacy class containing g. Then, $\frac{h\chi_V(g)}{d} \in \mathbb{A}$.*

Proof. Let $\mathcal{O}_1, \ldots, \mathcal{O}_s$ be the conjugacy classes of G, let $h_i = |\mathcal{O}_i|$, and for each i, fix some $g_i \in \mathcal{O}_i$. It suffices to prove that $h\chi_V(g_i)/d \in \mathbb{A}$ for all i: then, g is in the same conjugacy class as g_i for some i, which implies that $\chi_V(g) = \chi_V(g_i)$ and $h = h_i$, so the statement of the theorem follows.

Define $T_i = \sum_{g\in\mathcal{O}_i} g \in \mathbb{Z}[G] \subset \mathbb{C}[G]$. By the proof of Theorem 2.1.4, $\mathrm{Z}(\mathbb{C}[G])$ is a vector space over $\mathbb{C}$ with basis given by $\{T_i\}$. Likewise, $\mathrm{Z}(\mathbb{Z}[G])$ is a free $\mathbb{Z}$-module with generators given by $\{T_i\}$. Now, $\mathrm{Z}(\mathbb{Z}[G])$ is also a ring, so for any i and j, $T_iT_j \in \mathrm{Z}(\mathbb{Z}[G])$. This means that we can write

$$T_iT_j = \sum_{k=1}^{s} a_{ijk}T_k,$$

with $a_{ijk} \in \mathbb{Z}$ for all i, j, and k.

Now, let $T_i' : V \to V$ denote the action of T_i on V, so that $T_i'(v) = \sum_{g\in\mathcal{O}_i} gv$ for any $v \in V$. Then, by Proposition 2.2.3, $T_i' = \lambda_i I$ for some $\lambda_i \in \mathbb{C}$. The composition $T_i' \circ T_j'$ corresponds to acting by T_j and then by T_i, which is the same as acting by T_iT_j. So, the above expression for T_iT_j gives us

$$T_i' \circ T_j' = \sum_{k=1}^{s} a_{ijk}T_k'.$$

Plugging in $\lambda_\ell I$ for T_ℓ everywhere gives

$$\lambda_i\lambda_j I = \left(\sum_{k=1}^{s} a_{ijk}\lambda_k\right) I,$$

so that

$$\lambda_i\lambda_j = \sum_{k=1}^{s} a_{ijk}\lambda_k.$$

By Lemma 2.6.2, this implies that the λ_i are all algebraic integers.

If $\varphi : G \to \mathrm{GL}(V)$ is the homomorphism corresponding to the representation V, then we have

$$\mathrm{tr}(T_i') = \mathrm{tr}(\sum_{g\in\mathcal{O}_i} \varphi(g)) = \sum_{g\in\mathcal{O}_i} \mathrm{tr}(\varphi(g)) = \sum_{g\in\mathcal{O}_i} \chi_V(g) = h_i\chi_V(g_i).$$

On the other hand, $\mathrm{tr}(T_i') = \mathrm{tr}(\lambda_i I) = d\lambda_i$. Putting these together gives $\lambda_i = \frac{h_i\chi_i(g_i)}{d}$. Since we have already said that the λ_i are algebraic integers for all i, this implies that $\frac{h_i\chi_i(g_i)}{d} \in \mathbb{A}$ for all i. □

Corollary 2.6.7. *In the notation of the theorem, if h and d are coprime, then $\frac{\chi(g)}{d} \in \mathbb{A}$.*

Proof. We can write $ah + bd = 1$ for some $a, b \in \mathbb{Z}$. So, $\frac{ah}{d} + b = \frac{1}{d}$, or equivalently,

$$a\frac{h\chi(g)}{d} + b\chi(g) = \frac{\chi(g)}{d}.$$

Using the theorem, the previous proposition, and the fact that $\mathbb{A}$ is a ring, we see that everything on the left-hand side of this equation is an algebraic integer. This implies that the right-hand side of the equation is also an algebraic integer. □

The above theorem allows us to easily prove a strong theorem about the degree of irreducible representations of finite groups.

Theorem 2.6.8 (Dimension Theorem)**.** *Let V be a complex irreducible representation of degree d of a finite group G. Then, d divides $|G|$.*

Proof. Let $\mathcal{O}_1, \ldots, \mathcal{O}_s$ be the conjugacy classes of G, and let $h_i = |\mathcal{O}_i|$ for all i. By the orthogonality relations of the first kind,

$$1 = \langle \chi_V, \chi_V \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_V(g)\overline{\chi_V(g)}.$$

Multiplying by $\frac{|G|}{d}$ gives

$$\frac{|G|}{d} = \sum_{g \in G} \frac{\chi_V(g)}{d}\overline{\chi_V(g)} = \sum_{i=1}^{s} \left(\sum_{g \in \mathcal{O}_i} \frac{\chi_V(g)}{d}\overline{\chi_V(g)} \right) = \sum_{i=1}^{s} \frac{h_i \chi_V(g_i)}{d}\overline{\chi_V(g)},$$

The right-hand side of this equation is a sum of products of algebraic integers, so it is an algebraic integer. This implies that $\frac{|G|}{d} \in \mathbb{A}$. On the other hand, $\frac{|G|}{d} \in \mathbb{Q}$ since $|G|, d \in \mathbb{Z}$, so we have $\frac{|G|}{d} \in \mathbb{A} \cap \mathbb{Q} = \mathbb{Z}$. □

We can use this theorem to prove a couple results from group theory. While our proofs will certainly not be the most elementary ones for these results, they will show the power of the representation theory that we've built up.

Corollary 2.6.9. *Suppose G is a group such that $|G| = p^2$, where p is prime. Then, G is abelian.*

Proof. By the Dimension Theorem, the dimension of every irreducible representation of G must have order dividing p^2. Moreover, the sum of squares of these dimensions is $|G|$. So, we either have one irreducible representation of degree p or p^2 irreducible representations of degree 1. In the former case, G must have only one conjugacy class by Theorem 2.1.4; but since 1 is always in its own conjugacy class, this implies that every element of G is equal to 1, so that G is trivial, contradicting the fact that $|G| = p^2$. So, we must instead have p^2 irreducible representations all of degree 1. But then every irreducible representation of G is of degree 1, which implies that G is abelian (see Homework 5 for a proof of this fact). □

Corollary 2.6.10. *Let $p < q$ be prime, and suppose $q \not\equiv 1 \bmod p$. If $|G| = pq$, then G is abelian.*

Proof. Let $d_1, \ldots, d_s$ be the degrees of the irreducible complex representations of G. Then, d_i divides pq for all i, and $pq = d_1^2 + \cdots + d_s^2$. If $d_i = q$ for some i, then $d_i^2 = q^2 > pq$, a contradiction; so, the only possible values of the d_i are 1 and p. Let m and n denote the number of degree-1 and degree-p irreducible representations (respectively) of G. Then, we have $pq = m + np^2$. Since p divides pq and np^2, it must divide m. Moreover, by Proposition 2.1.7, $m = |G/[G, G]| =$

$|G|/|[G, G]|$ is a quotient of $|G|$, so m divides $|G| = pq$. $m = p$ or $m = pq$. If $m = p$, then $pq = p + np^2$, or equivalently, $q = 1 + np$. But this means that $q \equiv 1 \mod p$, contradicting our assumption. So we must have $m = pq$. Then, all of the irreducible representations of G are 1-dimensional, which implies that G is abelian (see Homework 5 for a proof of this fact). $\square$

Remark. We might wonder whether or not this last corollary is the strongest result of its form that we can prove. To this end, take the group of affine transformations, i.e. those which take $x \in \mathbb{R}$ to $ax + b \in \mathbb{R}$, with $a \neq 0$. We can represent these by matrices $\left(\begin{smallmatrix} a & b \\ 0 & 1 \end{smallmatrix}\right)$, which then sends a vector $(x, 1)$ to $(ax + b, 1)$ via matrix multiplication. We consider the subgroup G of these affine transformations in which $a \in (\mathbb{Z}/n)^\times$ and $b \in \mathbb{Z}/n$. Then, we have $|G| = n\varphi(n)$. Suppose $n = q$ is prime. Then, $|(\mathbb{Z}/q)^\times| = q - 1$, so $|G| = q(q - 1)$. We then fix a prime p dividing $q - 1$ and take the subgroup H of G consisting of matrices as above with the restriction that $a \in \mathbb{Z}/q$ and $a^p = 1$. One can check that H is, in fact, a subgroup of G and that it is non-abelian. On the other hand, $|H| = pq$. This proves that the condition $q \not\equiv 1 \mod p$ is necessary in the above corollary.

Chapter 3

Constructions of Representations

3.1 Tensor Products of Representations

In this section, we will consider two ways in which the tensor product on vector spaces can give rise to a tensor product of representations. The first, known as the external tensor product of representations, will take representations of two different groups and create a representation of the product group; the second, known as the internal tensor product of representations, will take two representations of the same group and create another representation of that group.

3.1.1 External Tensor Product

Definition 3.1.1. Suppose that groups G and H act on the vector spaces V and W, respectively. These group actions define V and W as representations of G and H (respectively). Then, we define the *external tensor product of representations* to be the vector space $V \otimes W$ with the action of $G \times H$ given by $(g, h)(v \otimes w) = gv \otimes hw$ for any $g \in G$, $h \in H$, $v \in V$, and $w \in W$.

Remark. We stated in Proposition **??** that tensor products of vector spaces are associative, commutative, and distributive through direct sums. Because the isomorphisms in all the statements of this proposition can be seen to respect the added structure of a representation, the same facts apply to the external tensor product of representations.

Now that we've defined a new sort of representation, the easiest way to try to understand it is to compute its character, which we do now.

Proposition 3.1.2. *Let V and W be representations of finite groups G and H (respectively). Then, for any $g \in G$ and $h \in H$,*

$$\chi_{V \otimes W}(g, h) = \chi_V(g)\chi_W(h).$$

Proof. Choose bases $\{v_i\}_i$ and $\{w_j\}_j$ of V and W respectively. Then, $\{v_i \otimes w_j\}_{i,j}$ is a basis for $V \otimes W$ by Theorem 1.1.4. Now, we can write $g(v_i) = \sum_{i'} a_{i'i} v_{i'}$ and

$h(w_i) = \sum_{j'} b_{j'j} w_{j'}$. By linearity of the tensor product, we have for any i and j

$$(g, h)(v_i \otimes w_j) = \left(\sum_{i'} a_{i'i} v_{i'} \right) \otimes \left(\sum_{j'} b_{j'j} w_{j'} \right) = \sum_{i', j'} a_{i'i} b_{j'j} (v_{i'} \otimes w_{j'}).$$

So, we have $\chi_V(g) = \sum_i a_{ii}$, $\chi_W(h) = \sum_j b_{jj}$, and $\chi_{V \otimes W}(g, h) = \sum_{i,j} a_{ii} b_{jj}$, whence the desired statement follows. ☐

Using this computation along with Proposition 2.4.10, we can understand when the external tensor product of representations is irreducible.

Theorem 3.1.3. *Let V and W be representations of finite groups G and H (respectively). Then, $V \otimes W$ is an irreducible representation of $G \times H$ if and only if V and W are irreducible representations of G and H.*

Proof. Applying the definition of the inner product as well as the above proposition gives

$$\begin{aligned} \langle \chi_{V \otimes W}, \chi_{V \otimes W} \rangle &= \frac{1}{|G \times H|} \sum_{g,h} \chi_{V \otimes W}(g, h) \overline{\chi_{V \otimes W}(g, h)} \\ &= \frac{1}{|G| \cdot |H|} \sum_{g,h} \chi_V(g) \chi_W(h) \overline{\chi_V(g) \chi_W(h)} \\ &= \left(\frac{1}{|G|} \sum_g \chi_V(g) \overline{\chi_V(g)} \right) \left(\frac{1}{|H|} \sum_h \chi_W(h) \overline{\chi_W(h)} \right) \\ &= \langle \chi_V, \chi_V \rangle \langle \chi_W, \chi_W \rangle. \end{aligned}$$

By Proposition 2.4.10, $V \otimes W$ is irreducible if and only if the left-hand side of this equation is 1. Since $\langle \chi_V, \chi_V \rangle$ and $\langle \chi_W, \chi_W \rangle$ are positive integers (this follows from the orthogonality relations; see, e.g., the proof of Proposition 2.4.10), the left-hand side of the above equation is 1 if and only if $\langle \chi_V, \chi_V \rangle = \langle \chi_W, \chi_W \rangle = 1$. But by Proposition 2.4.10, this is true if and only if V and W are irreducible. ☐

It is natural to wonder if the irreducible representations of $G \times H$ described in this theorem are the only irreducible representations.

Proposition 3.1.4. *For any two finite groups G and H, The only irreducible representations of $G \times H$ are those of the form $V \otimes W$, where V is an irreducible representation of G and W is an irreducible representation of H.*

Proof. Let $V_1, \ldots, V_m$ be the irreducible representations of G, let $W_1, \ldots, W_n$ be the irreducible representations of H, and set $d_i = \dim V_i$ and $t_j = \dim W_j$ for all i and j. Then, $\dim(V_i \otimes W_j) = d_i t_j$ for all i and j. So, the sum of the squares of the dimensions of the $V_i \otimes W_j$ is

$$\sum_{i,j} (d_i t_j)^2 = \left(\sum_i d_i^2 \right) \left(\sum_j t_j^2 \right) = |G||H| = |G \times H|,$$

where the penultimate equality follows from Theorem 2.1.4. However, the sum of the squares of the dimensions of *all* the irreducible representations of $G \times H$ is $|G \times H|$; so, if there were any irreducible representation other than $V_i \otimes W_j$ for some i and j, then its dimension would be at least 1, so the sum of the squares of the dimensions of irreducible representations of $G \times H$ would be at least $|G \times H| + 1$, a contradiction. □

3.1.2 Internal Tensor Product

Definition 3.1.5. Let G be a group which acts on two vector spaces V and W. This defines V and W as representations of G. Then, we define the *internal tensor product of representations* to be the vector space $V \otimes W$ with the action of G given by $g(v \otimes w) = gv \otimes gw$ for all g in G, v in V, and w in W.

Another way to think of the internal tensor product is as follows. Given two representations V and W of a group G, we can form the external tensor product, which is a representation $V \otimes W$ of $G \times G$. We can then restrict this representation to the diagonal subgoup of $G \times G$, i.e. to $\Delta = \{(g, g) : g \in G\} \subseteq G \times G$. One can check that $\Delta \cong G$ and moreover that the representation of G given in this way is the same as the internal tensor product $V \otimes W$ defined above.

Because of this relationship with the external tensor product, we might expect that the nice properties of that construction will translate over to the internal tensor product. Indeed, just as with the external tensor product, one can check that the internal tensor product is commutative, associative, and distributive through direct sums of representations. Moreover, the proof of Proposition 3.1.2 goes through again, which gives us:

Proposition 3.1.6. *Let V and W be any two representations of a finite group G. For all g in G,*

$$\chi_{V \otimes W}(g) = \chi_V(g)\chi_W(g)$$

On the other hand, if V and W are irreducible representations of a finite group G, we know that $V \otimes W$ is an irreducible representation of $G \times G$ by Theorem 3.1.3. However, it might not restrict to an irreducible representation of the diagonal, so the internal tensor product may not be irreducible.

Example. Recall from the example in Section 2.5 that the character table of S_3 is:

	1	(12)	(123)
χ_0	1	1	1
χ_1	1	-1	1
χ_2	2	0	-1

Let V_i be the irreducible representation of S_3 with character χ_i for $i \in \{0, 1, 2\}$. Then, $V_2 \otimes V_2$ is a representation of dimension $(\dim V_2)^2 = 4$. However, there are no irreducible representations of S_3 of dimension 4, so $V_2 \otimes V_2$ cannot be irreducible. Indeed, because $\chi_{V_2 \otimes V_2} = \chi_2^2$, we get $\chi_{V_2 \otimes V_2}(1) = 4$, $\chi_{V_2 \otimes V_2}(12) = 0$, and $\chi_{V_2 \otimes V_2}(123) = 1$. This implies that $\chi_{V_2 \otimes V_2} = \chi_0 + \chi_1 + \chi_2$, which by Proposition 2.4.7 gives us

$$V_2 \otimes V_2 \cong V_0 \oplus V_1 \oplus V_2.$$

On the other hand, one can check that $V_2 \otimes V_1$ has character $\chi_2 \cdot \chi_1 = \chi_2$, so that $V_2 \otimes V_1 \cong V_2$ is irreducible.

The above example seems to indicate that, in some circumstances, the internal tensor product of irreducible representations is irreducible. We describe the main such circumstance in the following proposition.

Proposition 3.1.7. *Let V and W be irreducible representations of a finite group G, and suppose that W is one-dimensional. Then, $V \otimes W$ is irreducible.*

Proof. Write $W = \mathbb{C}w$. Then, W corresponds to a homomorphism $\varphi : G \to \mathbb{C}^*$, so that $gw = \varphi(g)w$ for any $g \in G$. Now, we can define a new representation $W' = \mathbb{C}w'$ by a homomorphism $\varphi' : G \to \mathbb{C}^*$ such that $\varphi'(g) = \varphi(g)^{-1}$. (Because $\mathbb{C}^\times$ is abelian, one can check that φ' is, in fact, a homomorphism.) Then, the action of G on $W \otimes W'$ is given by $g(w \otimes w') = \varphi(g)w \otimes \varphi(g)^{-1}w' = \varphi(g)\varphi(g)^{-1}(w \otimes w') = w \otimes w'$, so that $W \otimes W'$ is isomorphic to the trivial representation on G. Then, we have

$$(V \otimes W) \otimes W' \cong V \otimes (W \otimes W') \cong V.$$

(Here we are using the fact that tensoring by the trivial representation does not change a representation, which follows from comparing characters and noting that the character of the trivial representation has value 1 everywhere.)

Now, suppose that U were some proper subrepresentation of $V \otimes W$. Then, $U \otimes W'$ would be a proper subrepresentation of $V \otimes W \otimes W' \cong V$ (that it's a proper subspace comes from comparing bases, and that it is a subrepresentation follows from the fact that U is a subrepresentation of $V \otimes W$). This contradicts the irreducibility of V. So, U cannot exist, which implies that $V \otimes W$ is irreducible. $\square$

Notice that, in the above proof, we constructed some representation W' of G such that $W \otimes W'$ is isomorphic to the trivial representation. Since this construction will work for any 1-dimensional representation of a finite group G, and since tensoring by the trivial representation does not change a representation, we see that the isomorphism classes of 1-dimensional representations (i.e. the set $\text{Hom}(G, \mathbb{C}^\times)$) form an abelian group, with the operation given by the (internal) tensor product and the identity given by the trivial representation. In fact, we can explicitly characterize this group, as the following proposition shows.

Proposition 3.1.8. *For any finite group G,*

$$\text{Hom}(G, \mathbb{C}^\times) \cong G/[G, G],$$

where the left-hand side is given a group structure via the internal tensor product.

Proof. By Proposition 2.1.7, we have $\text{Hom}(G, \mathbb{C}^\times) = \text{Hom}(G/[G, G], \mathbb{C}^\times)$. Now, because $G/[G, G]$ is a finite abelian group, we can write

$$G/[G, G] \cong C_{q_1} \times \cdots \times C_{q_r},$$

where the q_i are all prime powers. Then, by a similar argument to that of Proposition 1.4.12, we have

$$\mathrm{Hom}(G/[G,G], \mathbb{C}^\times) \cong \mathrm{Hom}\left(\prod_{i=1}^r C_{q_i}, \mathbb{C}^\times\right) \cong \prod_{i=1}^r \mathrm{Hom}(C_{q_i}, \mathbb{C}^\times).$$

For any i, fix a generator g_i of C_{q_i}. Then, any homomorphism $C_{q_i} \to \mathbb{C}^\times$ is determined by where it sends g_i and must send g_i to a q_ith root of unity. Any q_ith root of unity will work, so $\mathrm{Hom}(C_{q_i}, \mathbb{C}^\times)$ is in bijection with the q_ith roots of unity. Then, for any $\varphi, \psi \in \mathrm{Hom}(C_{q_i}, \mathbb{C}^\times)$, the tensor product of the representations given by φ and ψ sends g_i to $\varphi(g_i)\psi(g_i)$. This proves that in fact, the bijection from $\mathrm{Hom}(C_{q_i}, \mathbb{C}^\times)$ to the q_ith roots of unity respects multiplication in each of these groups, so it is an isomorphism. Since the q_ith roots of unity are isomorphic to C_{q_i}, we have $\mathrm{Hom}(C_{q_i}, \mathbb{C}^\times) \cong C_{q_i}$. So,

$$\mathrm{Hom}(G/[G,G], \mathbb{C}^\times) \cong \prod_{i=1}^r C_{q_i} \cong G/[G,G].$$

$\square$

Another consequence of Proposition 3.1.7 is that tensoring by any 1-dimensional representation V of a finite group G acts as a bijection on the isomorphism classes of all irreducible representations of G. To see that it is an injection, notice that if $W \otimes V \cong W' \otimes V$ for some irreducible representations W and W' of G, then there exists some 1-dimensional representation V' such that $V \otimes V'$ is the trivial representation, so

$$W \cong (W \otimes V) \otimes V' \cong (W' \otimes V) \otimes V' \cong W'.$$

Then, surjectivity follows from the pidgeonhole principle.

Example. As we've seen in the example in Section 2.5, S_3 has two 1-dimensional representations, so the group of 1-dimensional representations of S_3 is isomorphic to $\mathbb{Z}/2$.

Example. What happens if we take a representation V of a group G and tensor with the regular representation $\mathbb{C}[G]$? Because the character of the regular representation is 1 on 1 and 0 on everything else, we have for any $g \in G$

$$\chi_{V \otimes \mathbb{C}G}(g) = \chi_V(g)\chi_{\mathbb{C}G}(g) = \begin{cases} 0, & g \neq 1 \\ |G| \cdot \dim V, & g = 1 \end{cases}.$$

Notice that the character of $(\mathbb{C}G)^{\dim V}$ is the same as this character, which implies that $V \otimes \mathbb{C}G \cong (\mathbb{C}G)^{\dim V}$.

The following example shows how we can use the fact that tensoring by a 1-dimensional representation is a bijection on irreducible representations to help us compute the tensor products of representations.

Example. Recall from the end of Section 2.5 that the character table of D_4 is:

	1	a	a^2	b	ab
χ_0	1	1	1	1	1
χ_1	1	-1	1	-1	1
χ_2	1	-1	1	1	-1
χ_3	1	1	1	-1	-1
χ_4	2	0	-2	0	0

Let V_i be the irreducible representation of D_4 with character χ_i. We can now compute some tensor products of V_4 with itself. First, $\chi_{V_4 \otimes V_4} = \chi_{V_4}^2$, so $\chi_{V_4 \otimes V_4}(1) = \chi_{V_4 \otimes V_4}(a^2) = 4$ and $\chi_{V_4 \otimes V_4}$ is 0 everywhere else. One can check that this is precisely $\chi_0 + \chi_1 + \chi_2 + \chi_3$, so we have

$$V_4 \otimes V_4 \cong V_0 \oplus V_1 \oplus V_2 \oplus V_3.$$

Moreover, by distributivity of the tensor product over the direct sum, we get

$$V_4^{\otimes 3} \cong (V_0 \oplus V_1 \oplus V_2 \oplus V_3) \otimes V_4 \cong \oplus_{n=1}^4 V_n \oplus V_4 \cong V_4^{\oplus 4}.$$

Here the last equality follows from the fact that tensoring V_4 by a 1-dimensional representation must give us a 2-dimensional irreducible representation, of which the only one is V_4.

3.2 Permutation and Augmentation Representations

In this section, we will discuss a way to associate representations of groups to group actions. Because group actions often represent symmetries of objects in generality, such a construction is in practice very interesting. Moreover, we will see that the structure of the group action allows us to say more about the structure of the corresponding representation.

Definition 3.2.1. Let $\sigma : G \to S_X$ define the action of a group G on a finite set X (here S_X denotes the symmetric group on the elements of X). Then, we can define a vector space $\mathbb{C}[X] = \{\sum_{x \in X} c_x x \mid c_x \in \mathbb{C}\}$ (this is just the $\mathbb{C}$-vector space whose basis is given by the elements of x). We then define the *permutation representation* $\widetilde{\sigma} : G \to \mathrm{GL}(\mathbb{C}[X])$ corresponding to σ to be the map which extends the action of σ linearly to all of $\mathbb{C}[X]$. More explicitly, for any $v = \sum_{x \in X} c_x x \in \mathbb{C}[X]$ and any $g \in G$, we define

$$\widetilde{\sigma}(g)v = \sum_{x \in X} c_x \sigma(g)(x).$$

Remark. Throughout this section, given any group action $\sigma : G \to S_X$, we will assume that X is a finite set unless stated otherwise.

Remark. The permutation representation looks a lot like the regular representation in its definition. Indeed, one can check that if we define $\sigma : G \to S_G$ by sending g to the permutation $h \mapsto gh$, then $\widetilde{\sigma}$ is the regular representation.

The following lemma establishes some basic properties of permutation representations.

Lemma 3.2.2. *Let $\sigma : G \to S_X$ be a group action on a finite set X.*

1. *$\widetilde{\sigma}$ is a unitary representation with respect to the standard inner product on $\mathbb{C}[X]$.*

2. *For any $g \in G$,*
$$\chi_{\widetilde{\sigma}}(g) = |\operatorname{Fix}(g)|.$$

Proof.

1. Let $v = \sum_{x \in X} b_x x$ and $w = \sum_{x \in X} c_x x$ be two elements of $\mathbb{C}[X]$. Then, by definition of
$$\widetilde{\sigma}(g)(v) = \sum_{x \in X} b_{g^{-1}x} x$$
and likewise for w. So,
$$\langle v, w \rangle = \sum_{x \in X} b_x \overline{c_x} = \sum_{y \in G} b_{g^{-1}y} \overline{c_{g^{-1}y}} = \langle \widetilde{\sigma}(g)(v), \widetilde{\sigma}(g)(w) \rangle.$$

2. By definition, $\chi_{\widetilde{\sigma}}(g) = \operatorname{tr}(\widetilde{\sigma}(g))$. In the basis given by the elements of X, the matrix $\widetilde{\sigma}(g)$ permutes the basis elements, so it has a 1 on the diagonal in the row corresponding to $x \in X$ if and only if it fixes that basis element, i.e. if and only if $\sigma(g)(x) = x$. There are precisely $|\operatorname{Fix}(g)|$ such choices of x, so $\operatorname{tr}(\widetilde{\sigma}(g)) = |\operatorname{Fix}(g)|$.

□

In light of the above expression for the character of the permutation (and also because fixed points of group actions are interesting in their own right), we are led to define the equivalent of fixed points in representation theory.

Definition 3.2.3. Given a representation V of a group G, we define the subspace of *G-invariants* by
$$V^G = \{v \in V : gv = v \forall g \in G\}.$$

Example. Consider the standard action $\sigma : S_n \to S_X$, where $X = \{1, \ldots, n\}$. Let e_i be the basis vector of $\mathbb{C}[X]$ corresponding to $i \in X$. Then, $\widetilde{\sigma}$ is defined by letting the elements of S_n permute the e_i. The only element of $\mathbb{C}[X]$ that is fixed by all of S_n is the sum of all the e_i (or any multiple of this sum). So, the subspace of S_n-invariants of $\widetilde{\sigma}$ is
$$(\mathbb{C}^n)^{S_n} \cong \mathbb{C} \cdot \left(\sum_{i=1}^n e_i \right).$$

Notice that for any representation V of a group G, V^G is a subspace of V that is fixed by the action of G, so it is a subrepresentation of V. Moreover, using our understanding of the significance of inner products of characters for determining irreducible decompositions of representations, we can easily prove another such relation for V^G which will be useful to us.

Lemma 3.2.4. *Let G be a finite group, let V_0 be the trivial representation of G, and define $\chi_0 = V_0$. Then, for any representation V of G,*

$$\dim_{\mathbb{C}} V^G = \langle \chi_V, \chi_0 \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_V(G).$$

Proof. Let $\{v_1, \ldots, v_r\}$ be a basis for V^G, so that $\dim V^G = r$. Then, for all i, v_i is fixed by the action of G (since $v_i \in V^G$), which implies that $\mathbb{C} \cdot v_i$ is a subrepresentation of V^G isomorphic to V_0. So, $V^G \cong \oplus_{i=1}^r \mathbb{C} \cdot v_i \cong V_0^r$, and $\chi_{V^G} = \chi_0^r$. By linearity of the inner product and the orthonormality of irreducible characters, we then have

$$\langle \chi_V, \chi_0 \rangle = \langle \chi_0^r, \chi_0 \rangle = r \langle \chi_0, \chi_0 \rangle = r = \dim_{\mathbb{C}} V^G.$$

□

Now, it is natural to wonder what the subspace of G-invariants looks like for a permutation representation. One might expect that this subspace should be governed by the structure of the fixed points of the orbits of the group action. This is indeed the case, as the following proposition demonstrates.

Proposition 3.2.5. *Let $\sigma : G \to S_X$ be the action of a (not necessarily finite) group G on a finite set X. Denote the permutation representation of σ by V. Let $\mathcal{O}_1, \ldots, \mathcal{O}_m$ denote the orbits of σ, and for all i, let $v_i = \sum_{x \in \mathcal{O}_i} x \in V$. Then, $\{v_1, \ldots, v_m\}$ forms a basis for V^G.*

Proof. First, notice that the v_i all lie in V^G, since the action of any element of G permutes the elements of any orbit by definition. Now, let $v = \sum_x c_x x \in V^G$. For any elements x and y of X which are in the same orbit, then there is an element of g of G which takes $x \mapsto y$, which means that the coefficient of y in $g(v)$ is c_x. On the other hand, v is in V^G, so the action of g must not change v, which implies that c_x must be equal to the coefficient of y in v, i.e. c_y. Thus, $c_x = c_y$ for all x and y in the same orbit, from which it follows that $\sum_x c_x x$ is a linear combination of the v_i. This proves that the v_i span V^G. Finally, the v_i are all orthogonal (under the standard inner product): each of them is a sum of different basis elements of V. This implies that the v_i are linearly independent, so they form a basis of V^G. □

One consequence of the above proposition is that the permutation representation corresponding to any nontrivial group action (i.e. an action on any set with at least 2 elements) is never irreducible.

Corollary 3.2.6. *Let $\sigma : G \to S_X$ be the action of a group G on a finite set X. Suppose $|X| > 1$. Then, the permutation representation V corresponding to σ is reducible.*

Proof. If G fixes every element of X, then V is simply a direct sum of $|X| > 1$ copies of the trivial representation, so V is reducible. So, suppose that G does not fix some element of X. Then the number of orbits of the group action must be strictly less than $|X|$. By the above proposition, V^G has dimension equal to the number of orbits of σ, so this implies that $\dim V^G < |X| = \dim V$. On the other hand, Since there must be at least one orbit, $\dim V^G \geq 1$. This implies that V^G is a proper subrepresentation of V, so that V is not simple. $\square$

We now take a moment to prove a theorem from group theory that will be useful to us later.

Theorem 3.2.7 (Burnside). *Let $\sigma : G \to S_X$ be the action of a group G on a finite set X, and let m be the number of orbits of σ. Then,*

$$m = \frac{1}{|G|} \sum_{g \in G} |\operatorname{Fix}(g)|,$$

where $\operatorname{Fix}(g) = \{x \in X : gx = x\}$.

Proof. Recall that if $y \in X$ is in the orbit $\mathcal{O}_y$ of σ, then $|\mathcal{O}_y| = |G|/|\operatorname{Stab}(y)|$ by the Orbit-Stabilizer Theorem. So, we have

$$\begin{aligned}\frac{1}{|G|} \sum_{g \in G} |\operatorname{Fix}(g)| &= \frac{1}{|G|} \sum_{x \in X} |\operatorname{Stab}(x)| = \frac{1}{|G|} \sum_{\mathcal{O} \text{ an orbit}} \sum_{x \in \mathcal{O}} |\operatorname{Stab}(x)| \\ &= \frac{1}{|G|} \sum_{\mathcal{O} \text{ an orbit}} \sum_{x \in \mathcal{O}} \frac{|G|}{|\mathcal{O}|} = \frac{1}{|G|} \sum_{\mathcal{O} \text{ an orbit}} |G| = m.\end{aligned}$$

$\square$

Because there is not much to say about group actions in full generality, there isn't much to say about general permutation representations. As a result, we require some definitions of special types of group actions in order to continue our discussion of the relationship between group actions and representations.

Definition 3.2.8. We say that a group action $\sigma : G \to S_X$ is *transitive* if for all x and y in X, there exists some $g \in G$ such that $gx = y$. We define the *rank* of σ to be the number of orbits of the "diagonal action" of G on $X \times X$ (which is defined by $g(x, x') = (gx, gx')$). We say that σ is *2-transitive* if σ is transitive and if for all $x \neq y$ and $x' \neq y'$, there exists some $g \in G$ such that $gx = x'$ and $gy = y'$.

We now present a few equivalent conditions to 2-transitivity.

Lemma 3.2.9. *Let $\sigma : G \to S_X$ be a transitive group action. Then, the following are equivalent:*

1. *σ is 2-transitive;*

2. $\operatorname{rank} \sigma = 2$*; and*

3. *for all $x \in X$, the stabilizer G_x acts transitively on $X \setminus \{x\}$.*

Proof. We show that 1 and 2 are equivalent. (For a proof that 1 and 3 are equivalent, see Homework 8.) First, notice that the diagonal action of G on $X \times X$ has the diagonal $\Delta = \{(x, x) : x \in X\}$ as an orbit: by transitivity of σ, G can send any element of the diagonal to any other element, so Δ is contained in one orbit; but on the other hand, for any $g \in G$ and any $x \in X$, $g(x, x) = (gx, gx) \in \Delta$, so that nothing outside of Δ can be in its orbit. So, $\operatorname{rank} \sigma = 2$ if and only if everything outside of the diagonal of $X \times X$ is a single orbit, i.e. if and only if, for all $(x, y) \neq (x', y')$ with $x \neq y$ and $x' \neq y'$, there exists some g in G such that $g(x, y) = (x', y')$. But this is equivalent to the definition of 2-transitivity of σ given above. $\square$

The following characterization of the rank of a permutation will be useful to us soon.

Lemma 3.2.10. *Let $\sigma : G \to S_X$ be the action of a group on a finite set. Then,*

$$\operatorname{rank} \sigma = \langle \chi_{\widetilde{\sigma}}, \chi_{\widetilde{\sigma}} \rangle.$$

Proof. Notice that for any $g \in G$, under the diagonal action of G on $X \times X$, g fixes precisely the pairs $(x, x') \in X \times X$ for which g fixes both x and x'. So, g fixes $|\operatorname{Fix}(g)|^2$ points of $X \times X$ (here by $\operatorname{Fix}(g)$ we mean the number of fixed points of g under the action given by σ). So, we have

$$\operatorname{rank} \sigma = \frac{1}{|G|} \sum_{g \in G} |\operatorname{Fix}(g)|^2 = \langle \chi_{\widetilde{\sigma}}, \chi_{\widetilde{\sigma}} \rangle,$$

where the first equality follows from Burnside's Theorem and the second from our description of $\chi_{\widetilde{\sigma}}$ above. $\square$

If $\sigma : G \to S_X$ is a transitive group action, then the only orbit is all of X. By Proposition 3.2.5, then, V^G is a 1-dimensional, trivial subrepresentation of $V = \mathbb{C}[X]$ generated by $v = \sum_{x \in X} x$, so that $V^G \cong \mathbb{C}v$. Now, given any $g \in G$ and $w \in (\mathbb{C}v)^\perp$, because g^{-1} acts via a unitary matrix by Lemma 3.2.2 and fixes v,

$$\langle gw, v \rangle = \langle g^{-1}gw, g^{-1}v \rangle \langle w, v \rangle = 0.$$

So, $gw \in (\mathbb{C}v)^\perp$, which implies that $(\mathbb{C}v)^\perp$ is fixed under the action of G and so is a subrepresentation of V.

Definition 3.2.11. Let $\sigma : G \to S_X$ be a transitive group action, and let v be a generator of V^G (which, as discussed above, is 1-dimensional). Then, we define the *augmentation representation* corresponding to σ to be the subrepresentation of $\widetilde{\sigma}$ defined by

$$\operatorname{Aug}(\widetilde{\sigma}) = (\mathbb{C}v)^\perp = (V^G)^\perp.$$

The following is the main result we will prove about augmentation representations.

Theorem 3.2.12. *Let $\sigma : G \to S_X$ be a transitive group action. Then, σ is 2-transitive if and only if* $\operatorname{Aug}(\widetilde{\sigma})$ *is irreducible.*

Proof. Let V be the permutation representation of σ. By the above discussion, we have $V = \mathrm{Aug}(\widetilde{\sigma}) \oplus V^G$, where V^G is isomorphic to the trivial representation V_0 of G. Define $\chi_0 = \chi_{V_0}$ and $\chi_A = \chi_{\mathrm{Aug}(\widetilde{\sigma})}$. Then, our direct sum decomposition of V implies that $\chi_{\widetilde{\sigma}} = \chi_A + \chi_0$, or equivalently, $\chi_A = \chi_{\widetilde{\sigma}} - \chi_0$. So, by linearity of the inner product,

$$\langle \chi_A, \chi_A \rangle = \langle \chi_{\widetilde{\sigma}} - \chi_0, \chi_{\widetilde{\sigma}} - \chi_0 \rangle = \langle \chi_{\widetilde{\sigma}}, \chi_{\widetilde{\sigma}} \rangle - \langle \chi_{\widetilde{\sigma}}, \chi_0 \rangle - \langle \chi_0, \chi_{\widetilde{\sigma}} \rangle + \langle \chi_0, \chi_0 \rangle.$$

Now, $\langle \chi_{\widetilde{\sigma}}, \chi_0 \rangle$ is the multiplicity of the trivial representation in $V \cong \mathrm{Aug}(\widetilde{\sigma}) \oplus V^G$. On the other hand, if $w \in \mathrm{Aug}(\widetilde{\sigma}) = (V^G)^\perp$ generates a copy of the trivial representation in V, then we have $gw = w$ for all $g \in G$. This implies that $w \in V^G$. So, $w \in V^G \cap (V^G)^\perp = \{0\}$, which implies that $w = 0$, contradicting the assumption that it generates a 1-dimensional subrepresentatoin. This implies that there are no copies of the trivial representation in $\mathrm{Aug}(\widetilde{\sigma})$, so the multiplicity of the trivial representation in V is precisely 1. In other words, $\langle \chi_{\widetilde{\sigma}}, \chi_0 \rangle = \overline{\langle \chi_0, \chi_{\widetilde{\sigma}} \rangle} = 1$. Plugging this into the above equation gives

$$\langle \chi_A, \chi_A \rangle = \langle \chi_{\widetilde{\sigma}}, \chi_{\widetilde{\sigma}} \rangle - 1 - 1 + 1 = \langle \chi_{\widetilde{\sigma}}, \chi_{\widetilde{\sigma}} \rangle - 1.$$

Now, the augmentation representation is trivial if and only if the above equation is equal to 1, i.e. if and only if $\langle \chi_{\widetilde{\sigma}}, \chi_{\widetilde{\sigma}} \rangle = 2$. But by Lemma 3.2.10, this inner product is equal to $\operatorname{rank}\sigma$, and by Lemma 3.2.9, $\operatorname{rank}\sigma = 2$ if and only if σ is 2-transitive. □

We end this section with an example that shows how knowledge of permutation representations (as well as general representation theory) can help us understand irreducible representations and fill in character tables.

Example. Consider S_n acting on $\{1, \ldots, n\}$ in the natural way for $n \geq 2$. One can check that this action is 2-transitive. (See Homework 8 for a proof.) In particular, we can apply this to determine the character table for S_4. We know from the example at the end of Section 2.1 that we have the trivial and sign representations; call them V_0 and V_1, respectively. We also have the augmentation representation by the above theorem; call it V_2. Let $\chi_i = \chi_{V_i}$ for all i. Then, since we know what all these representations are, we can write down their characters on representatives of conjugacy classes of S_4:

	1	(12)	(123)	(1234)	(12)(34)
χ_0	1	1	1	1	1
χ_1	1	-1	1	-1	1
χ_2	3	1	0	-1	-1

Now, by Proposition 3.1.7, tensoring any irreducible representation by the sign representation gives us another irreducible representation. By comparing characters, one sees that $V_0 \otimes V_1 \cong V_1$, so this doesn't help us here. However, $V_2 \otimes V_1$ has a different character than V_0, V_1, and V_2, so it must be a new irreducible representation. Define $V_3 = V_2 \otimes V_1$ and $\chi_3 = \chi_{V_3}$. Then, adding χ_3 to our table gives:

	1	(12)	(123)	(1234)	(12)(34)
χ_0	1	1	1	1	1
χ_1	1	-1	1	-1	1
χ_2	3	1	0	-1	-1
χ_3	3	-1	0	1	-1

Because S_4 has 5 conjugacy classes, there is one more representation of V_4. Call its character χ_4. Using the orthogonality relation on the column of 1 with itself (or, equivalently, the equation in Theorem 2.1.4), we get that $\chi_4(1) = 2$. Then, the orthogonality relation of the column of 1 with all the other columns gives us the rest of the values of this character.

	1	(12)	(123)	(1234)	(12)(34)
χ_0	1	1	1	1	1
χ_1	1	-1	1	-1	1
χ_2	3	1	0	-1	-1
χ_3	3	-1	0	1	-1
χ_4	2	0	-1	0	2

3.3 Dual Representations

Let V and W be vector spaces over a field k, and let $\alpha : V \to W$ is a linear map. then, α induces a map $\alpha^* : W^* \to V^*$ by sending a linear functional $f : W \to k$ to the linear functional $f \circ \alpha : V \to k$. We can give a more explicit description of α^* by fixing bases. Let $\{v_1, \ldots, v_m\}$ and $\{w_1, \ldots, w_n\}$ be bases for V and W, respectively. Write

$$\alpha(v_j) = \sum_{i=1}^{n} a_{ij} w_i$$

for all j, and let $A = (a_{ij})$, so that A is the matrix representing α in the chosen bases. Then, one can check that

$$\alpha^*(w_i^*) = \sum_{j=1}^{n} a_{ij} v_j^*.$$

So, the matrix representing α^* is precisely A^T.

We can frame this discussion in the language of category theory as follows. We can define a map $* : \mathrm{Vect}_k \to \mathrm{Vect}_k$ from the category of vector spaces over k to itself by sending $V \mapsto V^*$ and $\alpha : V \to W$ to $\alpha^* : W \to V$. Given any linear maps $\alpha : V \to W$ and $\beta : W \to X$, if A and B are the matrices representating α and β (respectively) in some fixed bases, then $(AB)^T = B^T A^T$ implies that $(\beta \circ \alpha)^* = \alpha^* \circ \beta^*$. Moreover, $I^T = I$ implies that $\mathrm{id}^* = \mathrm{id}$. This proves that $*$ is a contravariant functor.

Now, suppose that we have a representation $\varphi : G \to \mathrm{GL}(V)$. Then, we would like to have G act on V^*. Since G acts by linear maps, the above discussion suggests that we should set $g \in G$ to act by $\varphi(g)^T$. However, transposing reverses the order of the product, so this does not define a homomorphism. To fix this, we define a dual representation $\varphi^* : G \to \mathrm{GL}(V)$ by $\varphi^*(g) = \varphi(g^{-1})^T$. Then, for any g and h in G, we have

$$\varphi^*(gh) = \varphi((gh)^{-1})^T = \varphi(h^{-1}g^{-1})^T = \varphi(g^{-1})^T \varphi(g^{-1})^T = \varphi^*(g)\varphi^*(h).$$

On the level of functions, given an an element f of V^* and an element g of G, gf is the linear functional which sends $v \in V$ to $f(g^{-1}v)$.

From the definition, we can immediately compute the character of V^*:

$$\chi_{V^*}(g) = \mathrm{tr}(A(g^{-1})^T) = \mathrm{tr}(A(g^{-1})) = \chi_V(g^{-1}).$$

In the case where G is finite (or, more generally, if the action of G on V is unitary, so that $A(g^{-1})^T = \overline{A(g)}$), we have $\chi_V(g^{-1}) = \overline{\chi_V(g)}$, which so that

$$\chi_{V^*}(g) = \overline{\chi_V}(g).$$

From this computation, we can already establish several properties of dual representations.

Proposition 3.3.1. *Let V and W be representations of a group G.*

1. $(V \oplus W)^* \cong V^* \oplus W^*$.
2. $(V \otimes W)^* \cong V^* \otimes W^*$.

Proof. We have already seen in Proposition 1.1.9 that the isomorphisms in the proposition are true as isomorphisms of vector spaces. So, one could directly check that the isomorphisms in that proposition respect the group action. Alternatively, we can compare characters: for all $g \in G$, we have

$$\chi_{(V\oplus W)^*}(g) = \chi_{V\oplus W}(g^{-1}) = \chi_V(g^{-1}) + \chi_W(g^{-1}) = \chi_{V^*}(g) + \chi_{W^*}(g) = \chi_{V^*\oplus V^*}(g),$$

and likewise,

$$\chi_{(V\otimes W)^*}(g) = \chi_{V\otimes W}(g^{-1}) = \chi_V(g^{-1})\chi_W(g^{-1}) = \chi_{V^*}(g)\chi_{W^*}(g) = \chi_{V^*\otimes W^*}(g).$$

□

Proposition 3.3.2. *Let V be a representation of a finite group G. Then, V is self-dual (i.e. $V \cong V^*$ as representations) if and only if χ_V takes real values everywhere.*

Proof. $V \cong V^*$ if and only if $\chi_V = \chi_{V^*} = \overline{\chi_V}$. □

Proposition 3.3.3. *Let V be a representation of a finite group G. Then, V is irreducible if and only if V^* is.*

Proof. Using the above expression for the character of V^*,

$$\langle \chi_{V^*}, \chi_{V^*} \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_{V^*}(g)\overline{\chi_{V^*}(g)} = \frac{1}{|G|} \sum_{g \in G} \overline{\chi_V(g)}\chi_V(g) = \langle \chi_V, \chi_V \rangle.$$

By Proposition 2.4.10, V is irreducible if and only if the above equation is 1, which in turn is true if and only if V^* is irreducible. □

One interesting consequence of the above proposition is that dualization always permutes the irreducible representations of a finite group. We now give a few examples to illustrate the above properties of dual representations.

Example. Consider $G = C_3$. If g is a generator of G and ζ is a primitive 3rd root of unity, then the character table for G is:

	1	g	g^2
χ_0	1	1	1
χ_1	1	ζ	ζ^2
χ_2	1	ζ^2	ζ

Since χ_0 takes real values everywhere, $V_0^* \cong V_0$ by the above proposiion. Recalling that $\overline{\zeta} = \zeta^2$, we see that $V_2^* \cong V_1$, and $V_1^* \cong V_2$. So, we see here that dualization does in fact permute the irreducible representations, as we noted above.

Example. Let $G = S_n$, and let V be an irreducible representation of G. Then, $\chi_{V^*}(\sigma) = \chi_V(\sigma^{-1})$ for all σ in S_n. On the other hand, σ and σ^{-1} have the same cycle type and hence are conjugate. So, $\chi_V(\sigma^{-1}) = \chi_V(\sigma)$, which implies that $\chi_{V^*} = \chi_V$ and $V^* \cong V$.

A generalization of the argument from the above example gives us the following theorem.

Theorem 3.3.4. *Let G be a finite group. Then, all the complex representations of G are self-dual if and only if conjugacy classes of G are self-dual: that is, if and only if, for every conjugacay class $\mathcal{O}$ of G, $\mathcal{O} = \mathcal{O}^*$, where $\mathcal{O}^* = \{g^{-1} : g \in \mathcal{O}\}$.*

Proof. We prove only one direction; the other is somewhat more involved. Suppose that the conjugacy classes of G are self-dual. Then, for any complex representation V of G and any $g \in G$, g and g^{-1} are conjugate, so $\chi_V(g) = \chi_V(g^{-1})$. Then,

$$\chi_{V^*}(g) = \chi_V(g^{-1}) = \chi_V(g),$$

so $\chi_{V^*} = \chi_V$ and $V \cong V^*$.

□

3.4 Induced and Restricted Representations

Let R and S be commutative rings, and fix an (R, S)-bimodule M. Then, by Proposition 1.2.13, for any S-module N, $M \otimes_S N$ is an R-module. One can check that, given $\alpha : N_1 \to N_2$ a homomorphism of S-modules, we get a morphism $1 \otimes \alpha : M \otimes N_1 \to M \otimes N_2$ (defined by $m \otimes n \mapsto m \otimes \alpha(n)$). One can check that this satisfies the definition of a functor, so that $M \otimes_S - : \mathrm{Mod}_S \to \mathrm{Mod}_R$ is a functor from the category of S-modules to the category of R-modules.

We can apply this functor in representation theory as follows. Let H be a subgroup of a group G. Then, for a given field k, $S = k[H]$ is a unital subring of $R = k[G]$. So, we take $M = R$, which is a left R-module via left multiplication by R and a right S-module via left multiplication by S. In this case, the functor $R \otimes -$ takes $k[H]$-modules (i.e. representations of H) to $k[G]$-modules (i.e. representations of G). We call $R \otimes -$ the *induction* functor, and we denote it Ind_H^G, or, when the group and subgroup in question are clear from context, simply Ind.

Remark. Given a representation V of H, we sometimes write $\chi_V \uparrow G$ as shorthand for $\chi_{\mathrm{Ind}(V)}$.

Now that we have a functor $\mathrm{Ind}_H^G : \mathrm{Mod}_{k[H]} \to \mathrm{Mod}_{k[G]}$, we may wonder if we have a functor in the other direction. This one is somewhat easier to construct: given any representation V of G, we can restrict the action of G on V to H to get a representation of H. Likewise, given an intertwiner $V \to W$ of representations of G, if we restrict the action of G on both V and W to H, we get an intertwiner of representations of H. This defines a functor, which we call the *restriction functor* and denote $\mathrm{Res}_H^G : \mathrm{Mod}_{k[G]} \to \mathrm{Mod}_{k[H]}$ (or simply Res, when the group and subgroup are clear from contest). The restriction functor behaves very nicely: for instance, from the definition, we immediately see that $\chi_{\mathrm{Res}(V)} = \chi_V$ for any representation V of any group G. So, we will focus our efforts on understanding the induction functor, especially as it relates to the restriction functor. We begin by characterizing the dimension of an induced representation.

Lemma 3.4.1. *Let H be a subgroup of a group G. For any representation V of H,*

$$\dim(\mathrm{Ind}_H^G(V)) = [G : H] \dim V.$$

Proof. Pick representatives $g_1, \ldots, g_m$ for the cosets of H, so that $G = \sqcup_{i=1}^m g_i H$. Then, as a $\mathbb{C}[H]$-module, $\mathbb{C}[G]$ is free with basis $\{g_i\}_{i=1}^m$. So, by Theorem 1.1.4, $\mathrm{Ind}(V)$ has a basis given by $\{g_i \otimes v_j\}$, where $\{v_j\}$ is a basis of V. Noting that there are $m = [G : H]$ g_i's and $\dim V$ v_j's then gives the desired result. $\square$

Example. Let $G = S_3$, $H = S_2 = \{1, (1,2)\}$. Then, picking coset representatives for H, we have $G = H \sqcup (1,3)H \sqcup (2,3)H$. Let V be the sign representation of H, defined by $V = \mathbb{C}v$ and $(1,2)v = -v$. Then, by the above lemma, $\dim(\mathrm{Ind}(V)) = [G : H] \cdot 1 = 3$. As discussed in the proof of the lemma, a basis for $\mathrm{Ind}(V)$ is given by $\{g_1 \otimes v, g_2 \otimes v, g_3 \otimes v\}$, where $g_1 = 1$, $g_2 = (1,3)$, and $g_2 = (2,3)$. Then, by definition of the $\mathbb{C}[G]$-module structure of $\mathrm{Ind}(V)$, we have for any $\sigma \in S_3$,

$$\sigma(g_i \otimes v) = (\sigma g_i) \otimes v.$$

Example. Let H be a subgroup of a group G, and let $V_0 \cong \mathbb{C} \cdot v$ denote the trivial representation of H. Pick representatives $g_1, \ldots, g_m$ for the cosets of H. Then, $\{g_i \otimes v\}_{i=1}^m$ is a basis for $\mathrm{Ind}(V)$. So, we can define a linear map of $\mathbb{C}$-vector spaces $\alpha : \mathrm{Ind}(V) \to \mathbb{C}[G/H]$ by sending $g_i \otimes v \mapsto g_i H$. Since this maps a basis of $\mathrm{Ind}(V)$ to a basis of $\mathbb{C}[G/H]$, it is an isomorphism of vector spaces. Moreover, for any i and any $g \in G$,

$$\alpha(g(g_i \otimes v)) = (g g_i)H = g(g_i H) = g\alpha(g_i \otimes v),$$

which suffices to prove that α is an intertwiner and hence an isomorphism of representations. So,

$$\text{Ind}(V) \cong \mathbb{C}[G/H],$$

where the action of G on $\mathbb{C}[G/H]$ is given by multiplication by elements of G (which permutes cosets of H).

Example. For any G and H, we have

$$\text{Ind}(\mathbb{C}[H]) = \mathbb{C}[G] \otimes_{\mathbb{C}[H]} \mathbb{C}[H] \cong \mathbb{C}[G],$$

where the last isomorphism follows from Proposition 1.2.11.

As with all of our constructions of representations, we wish to know what the character of the induce representation is. Using the sorts of coset arguments we've used above, this is not such a bad computation, as the following propositin demonstrates.

Proposition 3.4.2. *Let H be a subgroup of a finite group G. Pick representatives $g_1, \ldots, g_m$ for the cosets of H. Then, for any representation V of H,*

$$\text{Ind}_H^G(V) = \sum_{i=1}^m \widetilde{\chi_V}(g_i^{-1} g g_i) = \frac{1}{|H|} \sum_{f \in G} \widetilde{V}(f^{-1} g f),$$

where

$$\widetilde{\chi_V}(g) = \begin{cases} 0, & g \notin H \\ \chi_V(g), g \in H \end{cases}.$$

Proof. Fix a basis $\{v_i\}_{i=1}^d$ for V, and let $g \in G$. Then, $g(g_i \otimes v_j) = g_k \otimes h v_j$, where g_k represents the coset containing $g g_i$. If $k = i$, then $g g_i = g_i(g_i^{-1} g g_i) \in g_i H$, which implies that $g_i^{-1} g g_i \in H$. Then, using the definition of the tensor product,

$$g(g_i \otimes v_j) = g_i(g_i^{-1} g g_i) \otimes v_j = g_i \otimes (g_i^{-1} g g_i) v_j.$$

So, fixing i and summing over all j, we see that this choice of i contributes $\chi_V(g_i^{-1} g g_i)$ to the trace. On the other hand, if $k \neq i$, then the contribution of this basis element to the trace of the matrix representing the action of g is 0. Moreover, we have $g_i^{-1} g g_i \notin H$, so that $\widetilde{\chi_V}(g_i^{-1} g g_i) = 0$. So, in either case, the contribution of a given i, after summing over all j, is $\widetilde{\chi_V}(g_i^{-1} g g_i)$. Summing this over all i then gives

$$\chi_{\text{Ind}(V)}(g) = \sum_{i=1}^m \widetilde{\chi_V}(g_i^{-1} g g_i).$$

Now, for any i and any $f \in g_i H$, we have $f = g_i h$ for some $h \in H$. Then, $f^{-1} g f = h^{-1}(g_i^{-1} g g_i) h$, which is conjugate to g_i. So, $\chi_V(f^{-1} g f) = \chi_V(g_i^{-1} g g_i)$. Moreover, if g does not fix g_i, so that $g_i^{-1} g g_i \notin H$, then g also does not fix f (which is in the same coset as g_i), so that $f^{-1} g f \notin H$. This implies that $\widetilde{\chi_V}(f^{-1} g f) = \widetilde{\chi_V}(f^{-1} g f)$, so that

$$\widetilde{\chi_V}(g_i^{-1} g g_i) = \frac{1}{|H|} \sum_{f \in g_i H} \widetilde{\chi_V}(f^{-1} g f).$$

Plugging this into the above expression for $\chi_{\text{Ind}(V)}(g)$, we get

$$\chi_{\text{Ind}(V)}(g) = \sum_{i=1}^{m} \sum_{f \in g_i H} \frac{1}{|H|} \widetilde{\chi_V}(f^{-1}gf) = \frac{1}{|H|} \sum_{f \in G} \widetilde{\chi_V}(f^{-1}gf).$$

□

While the above equation is somewhat messy, it can actually be not so bad to compute by hand. If we pick H to be small relative to G, then $\widetilde{\chi_V}$ will be 0 on a lot of values, so we expect $\chi_{\text{Ind}(V)}(g)$ to be 0 for many g given an arbitrary representation V of H. This helps reduce our computations. As an extreme case of this, when $H = \{1\}$, the only representation of H is the trivial representation V_0. We saw in an example above that $\text{Ind}(V_0) \cong \mathbb{C}[G]$, and we know that $\chi_{\mathbb{C}[G]}$ has value 0 everywhere except on the identity of G.

Example. Using $S_2 \subset S_3$ and V the sign representation on S_2, as in the above example, one can use the above formula for the character of $\text{Ind}(V)$ to check that $\chi_{\text{Ind}(V)}(1,2) = -1$ and $\chi_{\text{Ind}(V)}(1,2,3) = 0$. Moreover, we know from Lemma 3.4.1 that $\chi_{\text{Ind}(V)}(1) = \dim(\text{Ind}(V)) = 3 \cdot \dim V = 3$. Now, recall that the character table of S_3 is:

	1	$(1,2)$	$(1,2,3)$
χ_0	1	1	1
χ_1	1	-1	1
χ_2	2	0	-1

Let V_i be the irreducible representation of S_3 with character χ_i for all i. Then, one sees that $\chi_{\text{Ind}(V)} = \chi_1 + \chi_2$, whence $\text{Ind}(V) \cong V_1 \oplus V_2$.

We now state few nice properties of Ind.

Proposition 3.4.3. *Let H be a subgroup of a finite group G, and let V_1 and V_2 be two representations of H. Then,*

$$\text{Ind}_H^G(V_1 \oplus V_2) \cong \text{Ind}_H^G(V_1) \oplus \text{Ind}_H^G(V_2).$$

Proof. One can use Proposition 3.4.2 to show that the characters of the two representations in question are equal, which means the representations must be isomorphic. □

Proposition 3.4.4. *Let H be a subgroup of a finite group G. Then, any irreducible representation V of G is a direct summand of* $\text{Ind}(W)$ *for some irreducible representation W of H.*

Corollary 3.4.5. *If G has an abelian subgroup of index n, then any irreducible representation of G has dimension at most n.*

Proof. Suppose H is an abelian subgroup of G with $[G : H] = n$. By the proposition, any irreducible representation V of G must be a direct summand of $\mathrm{Ind}(W)$ for some irreducible representation W of H. Since H is abelian, $\dim W = 1$. So, we have by Lemma 3.4.1,

$$\dim V \leq \dim \mathrm{Ind}(W) = [G : H] \dim W = n \cdot 1 = n.$$

□

With this rather powerful corollary, we can actually understand much of the structure of the irreducible representations of D_n.

Corollary 3.4.6. *Let $D_n = \langle a, b | a^n = b^2 = (ab)^2 = 1 \rangle$ be the nth dihedral group. Then, all the irreducible representations of D_n are either 1- or 2-dimensional. Moreover, when n is odd, there are 2 1-dimensional irreducible representations of D_n and $\frac{n-1}{2}$ 2-dimensional irreducible representations, and when n is even, there are 4 1-dimensional representations and $\frac{n-2}{2}$ 2-dimensional irreducible representations of D_n.*

Proof. Applying the above corollary to the abelian subgroup generated by a, which has index 2 in D_n, shows that the irreducible representations of D_n are at most 2-dimensional. Now, one can check that

$$D_n/[D_n, D_n] \cong \begin{cases} C_2 \times C_2, & n \text{ even} \\ C_2, & n \text{ odd} \end{cases}.$$

So, when n is even, we have $|C_2 \times C_2| = 4$ 1-dimensional irreducible representations; using Theorem 2.1.4 then implies that there are $\frac{2n-4}{4} = \frac{n-2}{2}$ 2-dimensional irreducible representations. Likewise, when n is odd, we have 2 1-dimensional irreducible representations and hence $\frac{2n-2}{4} = \frac{n-1}{2}$ 2-dimensional irreducible representations. □

The following result tells us one of the nicest relationships that we have between the induction and restriction functors, namely that they are adjoints. Although we mainly care about this in the context of representations, we will prove it for general rings instead of group algebras.

Theorem 3.4.7 (Fröbenius reciprocity)**.** *Let S be a unital subring of a ring R. Then, for any S-module M and any R-module N, we have a natural isomorphism*

$$\mathrm{Hom}_R(\mathrm{Ind}(M), N) \cong \mathrm{Hom}_S(M, \mathrm{Res}(N)).$$

In other words, Res *is a right adjoint to* Ind.

Proof. Given $\beta \in \mathrm{Hom}_S(M, \mathrm{Res}(N))$, we wish to define a map $\alpha : R \otimes_S M \to N$. For this, we define $\alpha(a \otimes m) = a\beta(m)$. One can check that this is a well-defined homomorphism of R-modules. So, we define $\psi_{M,N} : \mathrm{Hom}_S(M, \mathrm{Res}(N)) \to \mathrm{Hom}_R(\mathrm{Ind}(M), N)$ by setting $\psi_{M,N}(\beta) = \alpha$. Conversely, given $\alpha \in \mathrm{Hom}_R(\mathrm{Ind}(M), N)$, we define $\beta \in \mathrm{Hom}_S(M, \mathrm{Res}(N))$ by $\beta(m) = \alpha(1 \otimes m)$. Again, one checks that this is a well-defined homomorphism of R-modules. So, we can define $\varphi_{M,N}$: $\mathrm{Hom}_R(\mathrm{Ind}(M,N)) \to \mathrm{Hom}_S(M, \mathrm{Res}(N))$ by setting $\varphi_{M,N}(\alpha) = \beta$. Using these definitions, one can check that ψ and φ are inverses, so that ψ provides the desired bijection.

To check that our bijection are natural, we need to know that, given any $\gamma : N \to N'$, we have the following commutative diagram:

$$\begin{array}{ccc} \mathrm{Hom}_R(\mathrm{Ind}(M), N) & \xrightarrow{\varphi_{M,N}} & \mathrm{Hom}_S(M, \mathrm{Res}(N)) \\ \Big\downarrow{\scriptstyle \gamma\circ *} & & \Big\downarrow{\scriptstyle \gamma\circ *} \\ \mathrm{Hom}_R(\mathrm{Ind}(M), N') & \xrightarrow{\varphi_{M,N'}} & \mathrm{Hom}_S(M, \mathrm{Res}(N')) \end{array}$$

(We also need to check a similar diagram for a morphism $M' \to M$.) One can check this diagram correctly, but intuitively, it is clear that it will commute: our definition of $\varphi_{M,N}$ real assumptions about the structure of M or N, so it will respect homomorphisms of these objects. $\square$

Remark. We discussed in Secion 1.5 a general adjunction between free object functors and forgetful functors. We can think of the adjunction between Ind and Res sort of like those types of adjunctions. Intuitively, the restriction functor behaves like a forgetful functor: we forget about the structure of the module over the whole ring and only remember the module structure for the subring. Conversely, as we noted above, for any subgroup H of a group G, $k[G]$ is a free $k[H]$-module with basis given by coset representatives of H. Intuitively, then, tensoring by $k[G]$ over $k[H]$ (i.e. applying Ind) uses that free module structure to extend a representation of H to a representation of G, so that Ind behaves like a free object functor as well.

In the case of a subgroup H of G, Fröbenius reciprocity tells us that

$$\mathrm{Hom}_G(\mathrm{Ind}(W), V) \cong \mathrm{Hom}_H(W, \mathrm{Res}(V))$$

for any representations V of G and W of H. In particular, the dimensions of these two hom sets must be equal. Writing these dimensions as inner products using Proposition 2.4.9 and equating them, we have

$$\langle \chi_{\mathrm{Ind}(W)}, \chi_V \rangle = \langle \chi_W, \chi_{\mathrm{Res}(V)} \rangle.$$

Now, let $V_1, \ldots, V_m$ be the irreducible representations of G, let $\chi_{V_i} = \psi_i$, let $W_1, \ldots, W_n$ be the irreducible representations of H, and let $\chi_{W_i} = \chi_i$. Then, the above gives

$$\langle \chi_j \uparrow G, \psi_i \rangle = \langle \chi_j, \psi_i \downarrow H \rangle.$$

The left-hand side here is the multiplicity of V_i in $\mathrm{Ind}(W_j)$, while the right-hand side is the multiplicity of W_j in $\mathrm{Res}(V_i)$. The above equation thus gives us a useful relationship between the multiplicities of irreducible representations in induced and restricted representations. We can express this relationship slightly more succinctly, as summarized in the following proposition.

Proposition 3.4.8. *Let H be a subgroup of a finite group G, let $V_1, \ldots, V_m$ be the irreducible representations of G and let $W_1, \ldots, W_n$ be the irreducible representations of H. For every j, write*

$$\operatorname{Ind}(W_j) = \oplus_{i=1}^{n} V_i^{a_{ij}},$$

so that a_{ij} is the multiplicity of V_i in $\operatorname{Ind}(W_j)$; likewise, for every i, write

$$\operatorname{Res}(V_i) = \oplus_{j=1}^{m} W_j^{b_{ij}}.$$

if $A = (a_{ij})$ and if $B = (b_{ij})$, then

$$A = B^T.$$

We end this section with an example that shows the power of Fröbenius reciprocity in computing induced and restricted representations.

Example. Consider $S_2 \subset S_3$. Let V_0 and V_1 denote the trivial and sign representation, respectively, of S_2. Let W_0 and W_1 denote the trivial and sign representation of S_3, respectively, and let W_2 denote the unique reducible representation of S_3 of dimension 2. One can check that $\operatorname{Ind}(V_0)$ is the permutation representation on $\mathbb{C}^3$ and is isomorphic to $W_0 \oplus W_1$. Likewise, we showed in an example above that $\operatorname{Ind}(V_1) \cong W_1 \oplus W_2$. We can then immediately get the irreducible decompositions of the restrictions of the W_i using Fröbenius reciprocity: $\operatorname{Res}(W_0) \cong V_0$, $\operatorname{Res}(W_1) \cong V_1$, and $\operatorname{Res}(W_2) \cong V_0 \oplus V_1$.

3.5 Tensor Algebra Representations

Let G be a group acting on a vector space V. Then, G acts on $V^{\otimes k}$ by acting on each component: that is, for any $g \in G$ and any $v_1 \otimes \cdots \otimes v_k$, we define

$$g(v_1 \otimes \cdots \otimes v_k) = gv_1 \otimes \cdots \otimes gv_k.$$

In particular, we will focus on the case $k = 2$. (Note, however, that what we say here can be generalized to any value of k.) Thinking of S^2V as the subspace of $V \otimes V$ generated by elements of the form $v \cdot w = v \otimes w + w \otimes v$, we see that, for all $g \in G$,

$$g(v \cdot w) = gv \otimes gw + gw \otimes gv = gv \cdot gw,$$

which is another generator of S^2V. This implies that G fixes S^2V, so that S^2V is a subrepresentation of $V \otimes V$. By an analogous argument, $\Lambda^2 V$ is also a subrepresentation of $V \otimes V$.

Using out definitions of the symmetric and tensor algebras, we can compute the characters of these representations.

Proposition 3.5.1. *Let G be a finite group, and let V be a representation of G. Then, for any $g \in G$,*

$$\chi_{S^2V}(g) = \frac{\chi_V(g)^2 + \chi_V(g^2)}{2},$$

and

$$\chi_{\Lambda^2V}(g) = \frac{\chi_V(g)^2 - \chi_V(g^2)}{2}.$$

Proof. Fix $g \in G$, and pick a basis $\{v_1, \ldots, v_n\}$ for V in which g acts by a diagonal matrix on V (we know that such a basis exists by Proposition 2.4.2). Then, for all i, g acts by some $|G|$th root of unity λ_i on v_i (here λ_i is the ith element of the diagonal of the matrix representating the action of g). So, we have

$$\chi_V(g) = \lambda_1 + \cdots + \lambda_n,$$

and

$$\chi_V(g^2) = \lambda_1^2 + \cdots + \lambda_n^2.$$

Now, by Proposition 1.3.5, $\{v_i \cdot v_j\}_{i \leq j}$ is a basis for S^2V. For any i and j, we have

$$g(v_i \cdot v_j) = gv_i \cdot gv_j = \lambda_i v_i \otimes \lambda_j v_j + \lambda_j v_j \otimes \lambda_i v_i = \lambda_i \lambda_j (v_i \otimes v_j) + \lambda_i \lambda_j (v_j \otimes v_i) = \lambda_i \lambda_j (v_i \cdot v_j).$$

This implies that

$$\chi_{S^2V}(g) = \sum_{i \leq j} \lambda_i \lambda_j = \sum_{i<j} \lambda_i \lambda_j + \sum_{i=1}^n \lambda_i^2.$$

On the other hand, using our expressions for $\chi_V(g)$ and $\chi_V(g^2)$ above,

$$\chi_V(g)^2 = (\lambda_1 + \cdots + \lambda_n)^2 = \sum_{i=1}^n \lambda_i^2 + 2\sum_{i<j} \lambda_i \lambda_j = 2\chi_{S^2V}(g) - \sum_{i=1}^n \lambda_i^2 = 2\chi_{S^2V}(g) - \chi_V(g^2).$$

Rearranging, we get

$$\chi_{S^2V}(g) = \frac{\chi_V(g)^2 + \chi_V(g^2)}{2}.$$

As for Λ^2V, Proposition 1.3.5 tells us that a basis is given by $\{v_i \wedge v_j\}_{i<j}$. For any $i \neq j$, then, much as for the symmetric algebra above, we have

$$g(v_i \wedge v_j) = (gv_i) \wedge (gv_j) = \lambda_i \lambda_j (v_i \wedge v_j),$$

So, that

$$\chi_{\Lambda^2V}(g) = \sum_{i<j} \lambda_i \lambda_j.$$

Using the above expression for $\chi_V(g)^2$ gives

$$\chi_V(g)^2 = \chi_V(g^2) + 2\chi_{\Lambda^2V}(g),$$

which implies that

$$\chi_{\Lambda^2V}(g) = \frac{\chi_V(g)^2 - \chi_V(g^2)}{2}.$$

□

Corollary 3.5.2. *For any finite group G and any representation V of G,*

$$V \otimes V \cong S^2V \oplus \Lambda^2V.$$

Proof. Using the above proposition and comparing characters, we see that, for all g in G,

$$\chi_{\Lambda^2 V}(g) + \chi_{S^2 V}(g) = \frac{\chi_V(g)^2 - \chi_V(g^2)}{2} + \frac{\chi_V(g)^2 + \chi_V(g^2)}{2} = \chi_V(g)^2 = \chi_{V \otimes V}(g).$$

□

Example. Consider the case $G = S_3$. The character table for S_3 is:

	1	(12)	(123)
χ_0	1	1	1
χ_1	1	-1	1
χ_2	2	0	-1

Comparing characters, one sees that

$$V_2^{\otimes 2} \cong V_0 \oplus V_1 \oplus V_2.$$

Using the character formula for $\Lambda^2 V$ in the above proposition, we get $\Lambda^2 V \cong V_1$. Then, by either using the character formula for $S^2 V$ or the above corollay, we can conclude that $S^2 V \cong V_0 \oplus V_2$.

Chapter 4

Applications

4.1 Cyclotomic Extensions and Finite Simple Groups

In this chapter, we establish some basic properties of cyclotomic field extensions and use these along with representation theory to prove some interesting results about the possible orders of finite simple groups. We begin by recalling some Galois theory. Define $\omega_n = \exp(\frac{2\pi i}{n})$. Then, we have $\mathbb{Q} \subset \mathbb{Q}(\omega_n) \subset \mathbb{C}$. This field is the splitting field of

$$x^n - 1 = \prod_{i=0}^{n-1}(x - \omega_n^k) = \prod_{\zeta^n=1, \zeta \text{ primitive}} (x - \zeta) = \prod_{d|n} \Phi_d(x),$$

where $\Phi_d(x)$ is the dth cyclotomic polynomial.

Example. One can prove that, for any p prime, the pth cyclotomic polynomial is $\Phi_p(x) = x^{p-1} + \cdots + x + 1$.

Example. As an example of the above factorization, we have

$$x^4 - 1 = (x^2 + 1)(x + 1)(x - 1).$$

One can then check that $x - 1 = \Phi_1(x)$, $x + 1 = \Phi_2(x)$, and $x^2 + 1 = \Phi_4(x)$, so that the above factorization is precisely the product of the Φ_d for d dividing 4.

The following proposition is standard in the theory of cyclotomic polynomials.

Proposition 4.1.1. *Let $n > 0$.*

1. *$\Phi_n(x)$ has degree $\varphi(n)$, where φ is the Euler totient function.*

2. *$\Phi_n(x)$ is a monic polynomial in $\mathbb{Z}[x]$.*

3. *$\Phi_n(x)$ is irreducible.*

Proof. See Dummit and Foote, Section 13.6, p. 554 (Lemma 40 and Theorem 41). □

The next proposition follows from Galois theory.

Proposition 4.1.2.

1. *For any* $\alpha \in \mathbb{Q}(\omega_n)$*,* $\alpha \in \mathbb{Q}$ *if and only if* $\sigma(\alpha) = \alpha$ *for all* $\sigma \in \mathrm{Gal}(\mathbb{Q}(\omega_n)/\mathbb{Q})$*.*
2. *For any* $\alpha \in \mathbb{Q}(\omega_n)$*,*
$$\prod_{\sigma \in \mathrm{Gal}(\mathbb{Q}(\omega_n)/\mathbb{Q})} \sigma(\alpha) \in \mathbb{Q}.$$

With these preliminaries in hand, we can return to discussing representation theory.

Theorem 4.1.3. *Let* G *be a group of order* n*,* $\mathcal{O}$ *be a conjugacy class of* G*, and* $\varphi : G \to \mathrm{GL}(d, \mathbb{C})$ *be an irreducible representation. Suppose that* $h = |\mathcal{O}|$ *is coprime to* d*. Then, either*

1. *there exists* $\lambda \in \mathbb{C}^\times$ *such that* $\varphi(g) = \lambda I$ *for all* $g \in \mathcal{O}$*, or*
2. $\chi_p(g) = 0$ *for all* $g \in \mathcal{O}$*.*

Proof. Let $\chi = \chi_\varphi$. Recall that $\chi(g) = \lambda_1 + \cdots + \lambda_d$ for some dth roots of unity λ_i. So, $|\chi(g)| = d$ if and only if $\lambda_i = \lambda_j$ for all i and j, in which case condition (1) in the statement of the theorem must hold. So, suppose that condition (1) does not hold; then, we must have $|\chi(g)| < d$, and the λ_i cannot all be equal.

Now, by Corollary 2.6.7, we have that $\alpha = \chi(g)/d$ is in $\mathbb{A}$. For any $\sigma \in \Gamma = \mathrm{Gal}(\mathbb{Q}(\omega_n)/\mathbb{Q}) \cong (\mathbb{Z}/n)^\times$, we have by linearity of σ that
$$\sigma(\chi(g)) = \sigma(\lambda_1) + \cdots + \sigma(\lambda_d).$$
Since σ is an automorphism, it must send each λ_i to another dth root of unity. However, by injectivity of σ along with the fact that the λ_i are not all equal, we see that the $\sigma(\lambda_i)$ are not all equal. By our above arguments, this implies that $|\sigma(\chi(g))| < d$, so that
$$|\sigma(\alpha)| = \frac{1}{d}|\sigma(\lambda_1) + \cdots + \sigma(\lambda_d)| < 1.$$
Now, define
$$\beta = \prod_{\sigma \in \Gamma} \sigma(\alpha).$$
Notice that β is a product of elements with norm less than 1, so $|\beta| < 1$. By Proposition 4.1.2, $\beta \in \mathbb{Q}$; moreover, $\alpha \in \mathbb{A}$ implies that $\sigma(\alpha) \in \mathbb{A}$ for all σ, so that $\beta \in \mathbb{A}$. Then, $\beta \in \mathbb{Q} \cap \mathbb{A} = \mathbb{Z}$, and $|\beta| < 1$, which means that $\beta = 0$. This implies that $\sigma(\alpha) = 0$ for some σ. By injectivity of σ, we then have $\alpha = 0$ and therefore $\chi(g) = 0$. ☐

We would like to use this intersection of Galois theory and representation theory to study the group-theoretic property of being a simple group. To this end, we have the following lemma.

Lemma 4.1.4. *Let G be a finite non-abelian group. Suppose that there exists a conjugacy class $\mathcal{O} \neq \{1\}$ of G such that $|\mathcal{O}| = p^t$ for some p prime and $t > 0$. Then, G is not simple.*

Proof. Assume G is simple, and let $\varphi_1, \ldots, \varphi_s$ be the irreducible representations of G, where φ_1 is the trivial representation. Define $\chi_i = \chi_{\varphi_i}$ and $d_i = \deg \varphi_i$ for all i. Now, the commutator subgroup is always normal, and since G is non-abelian, it is nontrivial. So, we must have $G = [G, G]$, as otherwise $[G, G]$ would be a proper normal subgroup, contradicting simplicity of G. By Proposition 2.1.7, this implies that G has only one 1-dimensional representation (namely, the trivial one).

Now, for any $i \neq 1$, φ_i is not the trivial representation, so $\ker \varphi_i \neq G$. On the other hand, the kernel of any group homomorphism is a normal subgroup, so by simplicity of G, we must have $\ker \varphi_i = \{1\}$ (in other words, φ_i is a faithful representation). If there exists some $\lambda \in \mathbb{C}^\times$ such that $\varphi_i(g) = \lambda I$ for all $g \in \mathcal{O}$, then φ_i takes the same value on at least $|\mathcal{O}| \geq 2$ elements, which contradicts the fact that φ_i is faithful. This implies that condition 1 of Theorem 4.1.3 is not satisfied. So, there are two possibilities: either $|\mathcal{O}|$ and d_i are not coprime, so that we cannot apply the theorem; or, we can apply the theorem to conclude that $\chi_i(g) = 0$ for all $g \in \mathcal{O}$. Note that since the only prime divisor of $|\mathcal{O}|$ is p, the former case is true if and only if p divides d_i.

Consider the regular representation $\mathbb{C}[G] = d_1\varphi_1 \oplus \cdots \oplus d_s\varphi_s$. Then, for any $g \in \mathcal{O}$, we have $\chi_{\mathbb{C}[G]}(g) = \sum_{i=1}^{s} d_k \chi_k(g)$. On the other hand, $\chi_{\mathbb{C}[G]}(g)$ is the number of fixed points of multiplication by g, which is 0 since $g \neq 1$. Noting that $d_1 = \chi_1(g) = 1$, we have

$$0 = 1 + \sum_{i=2}^{s} d_i \chi_i(g) = 1 + \sum_{p|d_i} d_i \chi_i(g) + \sum_{p\nmid d_i} d_i \chi_i(g).$$

Now, for all i such that $p \nmid d_i$, by the above, we must be in the case where $\chi_i(g) = 0$. This implies that the second sum in the above equation vanishes. On the other hand, for all i such that $p|d_i$, we can write $d_i = pd_i'$ for some d_i'. Then, the above equation gives

$$\sum_{p|d_i} pd_i'\chi_i(g) = -1,$$

so that

$$\sum_{p|d_i} d_i'\chi_i(g) = -\frac{1}{p}.$$

The sum on the left-hand side here is in $\mathbb{A}$, since each term is, and it is also in $\mathbb{Q}$, so it is in $\mathbb{Z}$. But $-\frac{1}{p}$ is not in $\mathbb{Z}$, contradicting our original assumption that G is simple. $\square$

We can now prove our main result, which rules out many possible orders of finite, simple, non-abelian groups.

Theorem 4.1.5. *If G is finite, simple, and non-abelian, then $|G| \neq p^a q^b$ for any p and q prime and any $a, b > 0$.*

Proof. Assume there is some finite, simple, non-abelian group G such that $|G| = p^a q^b$ for some p and q prime and some $a, b > 0$. Then, there exists a Sylow q-subgroup H of G of order q^b. It is a fact from group theory that any group of prime power order has nontrivial center (this follows from the conjugacy class equation). This means that we can pick some $h \in Z(H) \setminus \{1\}$. Then, $\mathrm{C}_G(h) \supset H$, so $|\mathrm{C}_G(h)|$ is at least q^b and so must be equal to $p^c q^b$ for some c. But then $|\mathcal{O}_h||\mathrm{C}_G(h)| = |G|$ implies that $|\mathcal{O}_h| = p^{a-c}$. So, $\mathcal{O}_h$ satisfies the conditions of the above lemma, which means that G is not simple, a contradiction. $\square$

4.2 Lie Groups and McKay Graphs

4.2.1 Finite Subgroups of SO(3)

We first define the orthogonal groups, which are essentially real analogs of the unitary groups.

Definition 4.2.1. We define the nth *orthogonal group* to be

$$\mathrm{O}(n) = \{A \in \mathrm{Mat}_n(\mathbb{R}) : AA^T = I\}.$$

We define the nth *special orthogonal group*, SO(3), to be the subgroup of O(3) consisting of matrices of determinant 1.

We can think of $\mathrm{O}(n)$ as the group of symmetries of the inner product space that is $\mathbb{R}^n$ with the standard inner product. We can also endow $\mathrm{O}(n)$ with a topology inherited as a subspace of $\mathrm{Mat}_n(\mathbb{R}) \cong \mathbb{R}^{n^2}$. With this topology, $\mathrm{O}(n)$ becomes a *topological group*, i.e. a group with a topology on it such that the multiplication map $(g, h) \mapsto gh$ and the inversion map $g \mapsto g^{-1}$ are continuous. (In fact, these maps are smooth under the manifold structure of $\mathbb{R}^{n^2}$, which makes $\mathrm{O}(n)$ a *Lie group*.)

Because the elements of $\mathrm{O}(n)$ represent linear transformations $\mathbb{R}^n \to \mathbb{R}^n$ (in the standard orthonormal basis), we can consider what these linear transformations look like geometrically. It turns out that they are precisely the subset set of dstiance-preserving automorphisms of $\mathbb{R}^n$ which are symmetries of S^{n-1}. With this view, $\mathrm{SO}(n)$ is the set of distance-preserving automorphisms of $\mathbb{R}^n$ that fix S^{n-1} and, moreover, preserve its orientation; in other words, it is the group of roataions of S^{n-1}.

Example. Consider the case where $n = 2$. Then, by what we've said about, O(2) consists of the symmetries of S^1, i.e. to all rotations about the origin and all reflections about lines through the origin. By contrast, SO(2) is simply the group of rotations about the origin. Because $\mathbb{R}^2 \cong \mathbb{C}$, we can think of these as rotations of the complex plane. In this setting, the elements of SO(2) correspond precisely to multiplication by $e^{i\theta}$ for any θ. Under this correspondence, we see that

$$\mathrm{SO}(2) \cong S^1 \subset \mathbb{C}.$$

We will primearily be interested in SO(3). Our goal will be to classify the finite subgroups of this orthogonal group. To do this, we require a few preliminary results.

Proposition 4.2.2. *If n is odd, then* $\mathrm{O}(n) \cong \mathrm{SO}(n) \times C_2$.

Proof. Note that by definition, every element of $\mathrm{O}(n)$ has determinant equal to either 1 or -1. Now, $-I \in \mathrm{O}(n)$ has determinant $(-1)^n = -1$, so multiplying by $-I$ will send elements of $\mathrm{SO}(n)$ to $\mathrm{O}(n) \setminus \mathrm{SO}(n)$ and vice versa. From this, one can check that I and $-I$ are representatives of the cosets of $\mathrm{SO}(n)$, so that $\mathrm{O}(n) \cong \mathrm{SO}(n) \times \{\pm I\} \cong \mathrm{SO}(n) \times C_2$. □

Proposition 4.2.3. *Any finite subgroup of* $\mathrm{O}(2)$ *is either cyclic or dihedral.*

Proof. Let G be a finite subgroup of $\mathrm{O}(2)$. Consider $H = G \cap \mathrm{SO}(2) \subseteq \mathrm{SO}(2)$. As we discussed in the above example, $\mathrm{SO}(2) \cong S^1$. So, H is a finite subgroup of S^1, which means it must be isomorphic to C_n for all n. If $H = G$, then G is cyclic, and we are done. Otherwise, there must exist some $s \in G$ such that $s \notin \mathrm{SO}(2)$. By the above example, then, s is a reflection about some line through the origin.

We can consider s in terms of what it does to the unit circle, which we will consider in $\mathbb{C}$ via the isomorphism $\mathbb{R}^2 \cong \mathbb{C}$. Suppose s corresponds to reflection about the line through $e^{i\theta}$. Then, any $e^{i\tau} \in S^1$ will be sent, under s, to $\exp(i(\theta - \tau + \theta)) = \exp(i(2\theta - \tau))$. This is the same thing $e^{i\tau}$ is sent to upon first reflecting about the real axis (which sends $e^{i\tau} \mapsto e^{-i\tau}$) and then rotating by 2θ. Abusing notation so that $e^{i\theta}$ stands for a rotation by θ, we write $s = e^{i\theta}x$, where x is the reflection about the x axis.

Now, we claim that $H \cdot s$ contains every element of G not in $\mathrm{SO}(2)$. Let $t \in G$ such that $t \notin \mathrm{SO}(2)$. Then, t corresponds to a reflection about a line through $e^{i\tau}$ for some τ. By our above argument for s, we have $t = e^{i\tau}x$. Then, the element ts^{-1} is an element of G satisfying $(ts^{-1})s = t$. On the other hand, we have

$$e^{i(\tau-\theta)}s = e^{i(\tau-\theta)}e^{i\theta}x = e^{i\tau}x = t.$$

Since there is a unique element of $\mathrm{O}(2)$ that sends s to t under left multiplication, we must have $\exp(i(\tau - \theta)) = ts^{-1} \in G$. But $\exp(i(\tau - \theta))$ is a rotation, so it is an element of $\mathrm{SO}(2)$ and hence an element of H. This proves that $t \in H \cdot s$, so that $G = H \cup H \cdot s \cong C_n \cup C_n \cdot s$. G is then the group of symmetries of a regular $2n$-gon (which is generated by C_n, the group of rotations, along with one reflection), so it is isomorphic to D_n. □

We are now ready to classify the finite subgroups of SO(3).

Theorem 4.2.4. *Let G be a finite subgroup of* $\mathrm{SO}(3)$. *Then, one of the following holds.*

1. *$G \cong C_n = \langle a \mid a^n = 1 \rangle$ is the cyclic group of order n for some n. In other words, the elements of G are precisely the rotations about some axis by $2k\pi/n$ for $1 \leq k \leq n$.*

2. $G \cong D_n = \langle a, b \mid a^n = b^2 = (ab)^2 = 1 \rangle$ *is the dihedral group of order* $2n$ *for some* n. *In other words, the elements of* G *are precisely the symmetries of some regular* n*-gon.*

3. $G \cong A_4 = \langle a, b | a^2 = b^3 = (ab)^3 = 1 \rangle$ *is the tetrahedral group. In other words, the elements of* G *are the symmetries of the regular tetrahedron.*

4. $G \cong S_4 = \langle a, b | a^2 = b^3 = (ab)^4 = 1 \rangle$ *is the octahedral group. In other words, the elements of* G *are the symmetries of the regular octahedron.*

5. $G \cong A_5 = \langle a, b | a^2 = b^3 = (ab)^5 = 1 \rangle$ *is the dodecahedral group. In other words, the elements of* G *are the rotations of the regular dodecahedron (icosahedron).*

Proof. Let X be the subset of S^2 consisting of the points which are fixed by some element of $G \setminus \{1\}$. We claim that the action of G fixes X. For any $g \in G$ and $x \in X$, x is fixed by some $h \neq 1$ in G, so we have

$$ghg^{-1}(gx) = ghx = gx.$$

This implies that gx is fixed by $ghg^{-1} \neq 1$ and so is an element of X. So, the action of G on S^2 restricts to a group action on X.

As discussed above, elements of G are symmetries of the sphere $S^2 \subset \mathbb{R}^3$. This means that each element h of G is a rotation about some line through the origin, so $h \in G \setminus \{1\}$ fixes precisely 2 points of S^2 (the 2 points where the sphere intersects the axis of rotation), whereas 1 fixes all of X. In particular, then, X is finite, so we can apply Burnside's Theorem (Theorem 3.2.7) to get

$$|X/G| = \frac{1}{|G|}(|X| + 2(|G| - 1)).$$

(Here X/G is the set of orbits of G.) Let $N = |X/G|$, and choose representatives $x_1, \ldots, x_N$ for the orbits of G on X. Then, we have $X = \cup_{i=1}^N Gx_i$, so the above equation becomes

$$N = \frac{1}{|G|}\left(\sum_{i=1}^N |Gx_i| + 2(|G| - 1)\right). \tag{4.1}$$

Now, the Orbit-Stabilizer Theorem tells us that, for any i, $|Gx_i|/|G| = |G_{x_i}|$. Applying this and rearranging the above equation, we get

$$\frac{2(|G|-1)}{|G|} = N - \sum_{i=1}^N \frac{|Gx_i|}{|G|} = N - \sum_{i=1}^N \frac{1}{|G_{x_i}|} = \sum_{i=1}^N \left(1 - \frac{1}{|G_{x_i}|}\right). \tag{4.2}$$

The x_i are fixed by at least one element of G and hence also the inverse of that element, so that $a_i \geq 2$ for all i. The right-hand side of the above is then at least $\sum_{i=1}^N 1 - \frac{1}{2} = \frac{N}{2}$, which gives us

$$N \leq \frac{4(|G|-1)}{|G|} < 4,$$

so that $N \leq 3$. We now consider each possibility for N in turn.

1. Suppose $N = 1$. Then, $|Gx_1| = |G|$, (4.1) gives
$$1 = \frac{1}{|G|}\left(|G| + 2(|G| - 1)\right) = 3 - \frac{2}{|G|}.$$
Solving for $|G|$ then gives $|G| = 1$, so that $G \cong C_1$.

2. Suppose $N = 2$. Then, equation (4.1) gives
$$2 = \frac{1}{|G|}\left(|Gx_1| + |Gx_2| + 2(|G| - 1)\right) = 2 + \frac{|Gx_1| + |Gx_2| + 2}{|G|},$$
so that
$$|Gx_1| + |Gx_2| + 2 = 0.$$
Since $|Gx_i| \geq 1$ for all i, this is only possible if $|Gx_1| = |Gx_2| = 1$. Then, we have $|X| = |Gx_1| + |Gx_2| = 2$. This implies that every nonidentity element of G must be a rotation about the same axis (so that they all fix the same 2 points). If we take the two fixed points to be the north and south poles of S^2, then the group of all rotations about the axis through these points is the group of rotations of the equator. So, the group of all rotations about this axis is isomorphic to the rotation group of the circle, which, as dicussed in the above example, is $\mathrm{SO}(2) \cong S^1$. We then have that G is a finite subgroup of S^1, so that $G \cong C_n$ is cyclic.

3. Suppose $N = 3$. Then, (4.2) gives
$$2 - \frac{2}{|G|} = 3 - \frac{1}{|G_{x_1}|} - \frac{1}{|G_{x_2}|} - \frac{1}{|G_{x_3}|}.$$
Letting $a_i = |G_{x_i}|$ for all i and rearranging the above equation, we get
$$\frac{1}{a_1} + \frac{1}{a_2} + \frac{1}{a_3} = 1 + \frac{2}{|G|} > 1. \tag{4.3}$$
By reordering the x_i if necessary, we may take $a_1 \leq a_2 \leq a_3$.

It turns out that there are relatively few triples which satisfy the inequality in (4.3). Indeed, if $a_i \geq 3$ for all i, then
$$\frac{1}{a_1} + \frac{1}{a_2} + \frac{1}{a_3} \leq \frac{1}{3} + \frac{1}{3} + \frac{1}{3} = 1.$$
So, we must have $a_i < 3$ for some i and, in particular, $a_1 < 3$. On the other hand, if $a_1 = 1$, then the above equation gives us
$$1 + \frac{1}{a_2} + \frac{1}{a_3} = 1 + \frac{2}{|G|},$$
so that
$$\frac{1}{a_2} + \frac{1}{a_3} = \frac{2}{|G|}.$$

Since $a_2, a_3 \leq |G|$, this can only be true if $a_2 = a_3 = |G|$. Then, every element of G fixes every element of X, and there are 3 elements of X. But then either G itself is trivial (which we have already considered above) or G contains some nonidentity element, which only fixes 2 points of S^2, a contradiction. So, we cannot have $a_1 = 1$, which implies that $a_1 = 2$.

Now, plugging this value of a_1 into (4.3) gives

$$\frac{1}{a_2} + \frac{1}{a_3} > \frac{1}{2}.$$

Notice that if a_2 and a_3 are both at least 4, then we have

$$\frac{1}{a_2} + \frac{1}{a_3} \leq \frac{1}{4} + \frac{1}{4} = \frac{1}{2},$$

a contradiction. In particular, then, $a_2 < 4$. First, suppose that $a_2 = 2$. Then, the above inequality is satisfied by any value of a_3. In the case where $a_3 = 2$, the equation in (4.3) gives us $|G| = 4$. So, either $G \cong C_4$ or $G \cong C_2 \times C_2$. For any i, G_{x_i} is a finite subgroup of the group of rotations about the axis through x_i, so it is cyclic. By the Orbit-Stabilizer Theorem, $|Gx_i| = |G|/a_i = 2$ for all i, so each G_{x_i} is a unique copy of C_2 in G. But C_4 has only one copy of C_2 in it, so we must have $G \cong C_2 \times C_2 \cong D_2$. (Indeed, one can check in this case that G acts as the symmetry group of a regular 2-gon with vertices on the sphere.) Suppose instead that $a_3 = n$ for some $n > 2$. Then, (4.3) gives us $|G| = 2n$. As noted above, G_{x_3} is a cyclic subgroup of G consisting of rotations about the axis through x_3. It must be generated by a rotation g by $\frac{2\pi}{n}$. On the other hand, the Orbit-Stabilizer Theorem gives us $|Gx_3| = |G|/a_3 = 2$. Since x_3 and $-x_3$ must have stabilizers of the same size, they must be the two elements of Gx_3. So, there exists some $s \in G$ such that $sx_3 = -x_3$. Then, g fixes $-x_3$ as well, so $g^i s$ does not fix x_3 for any $1 \leq i < n$. Thus, $g^i s$ gives us n elements of $G \setminus G_{x_3}$. But since $|G_{x_3}| = n$, this accounts for all $2n$ elements of G. So, $G = C_n \cup C_n \cdot s$, and since C_n is a rotation group and s is a reflection, we get $G \cong D_{2n}$.

Finally, we are left with the case where $a_2 = 3$. Then, we must have $a_3 \in \{3, 4, 5\}$. We demonstrate just the case where $a_3 = 3$; the other two are similar. By (4.3), $|G| = 12$, so the Orbit-Stabilizer Theorem gives $|Gx_3| = 4$. Now, $G_{x_3} \cong C_3$ is generated by a rotation g about the line through z and $-z$ by $2\pi/3$. Let u be some element of Gx_3 other than x_3; then, gu and $g^2 u$ must be the other two elements of this orbit. Now, one can check that u, gu, and $g^2 u$ are the vertices of an equilateral triangle. Moreover, gu, $g^2 u$, and x_3 are also the vertices of the equilateral triangle. So, the convex hull of these 4 points is a regular tetrahedron, which means that G contains the symmetries of the tetrahedron, i.e. a copy of A_4. But G has exactly 12 elements, and $|A_4| = 12$, so these are all the elements of G, and $G \cong A_4$.

$\square$

4.2.2 Finite Subgroups of SU(2)

In this section, we will use our classification of the finite subgroups of SO(3) to classify the finite subgroups of SU(2). The latter subgroups provide a rich setting for many ideas in representation theory. First, we define SU(2).

Definition 4.2.5. The *nth unitary group*, $\mathrm{U}(n)$, is the subgroup of $\mathrm{GL}(n, \mathbb{C})$ consisting of matrices such that $AA^* = I$. The *nth special unitary group*, $\mathrm{SU}(n)$, is the subgroup of $\mathrm{U}(n)$ consisting of matrices of determinant 1.

We begin by trying to understand conjugacy classes in $\mathrm{U}(n)$.

Theorem 4.2.6. *Any matric A in $\mathrm{U}(n)$ has eigenvalues of magnitude 1 and is conjugate in $\mathrm{U}(n)$ to a diagonal matrix.*

Proof. The proof is similar to Theorem 2.3.6 in Steinberg (p. 9). For the full details, see Homework 11. □

Remark. Note that the above theorem does not hold for all of $\mathrm{GL}(n, \mathbb{C})$ or $\mathrm{GL}(n, \mathbb{R})$: for instance, given any $\lambda \neq \pm 1$, $\left(\begin{smallmatrix} \lambda & 1 \\ 0 & \lambda \end{smallmatrix}\right) \in \mathrm{GL}(n)$ is not diagonalizable.

Given the above theorem, it is easy to see when two elements of $\mathrm{U}(n)$ are conjugate. Since each is conjugate to a diagonal matrix, we can just check if the diagonal matrices are conjugate. On the other hand, one can check that the only conjugation we can do to diagonal matrices is permute their diagonal. So, every conjugacy class contains a unique diagonal matrix up to permutation of the diagonal.

We wish to translate this understanding of conjugacy classes to $\mathrm{SU}(n)$. For any A in $\mathrm{SU}(n)$, by the above theorem, we can write $A = BDB^{-1}$, where $D \in \mathrm{U}(n)$ is diagonal and $B \in \mathrm{U}(n)$. Notice that $BB^* = I$ implies that $(\det B)(\overline{\det B}) = 1$, so that $\det B \in S^1 \subset \mathbb{C}$. Fix some nth root λ of $\det B$, and define $B' = \lambda^{-1} B$. Then, $B' \in \mathrm{U}(n)$, and

$$\det B' = \det(\lambda^{-1} I B) = \det(\lambda^{-1} I) \det(B) = \lambda^{-n}(\det B) = (\det B)^{-1}(\det B) = 1,$$

so that $B' \in \mathrm{SU}(n)$. Moreover,

$$(B')^{-1} A B' = \lambda B^{-1} A (\lambda^{-1} B) = \lambda\lambda^{-1} B^{-1} A B = D.$$

So, A is in fact conjugate by a matrix in $\mathrm{SU}(n)$ to a diagonal matrix (which is also in $\mathrm{SU}(n)$ since both A and B' are). Thus, our same arguments for $\mathrm{U}(n)$ above apply to $\mathrm{SU}(n)$ and show that each conjugacy class of $\mathrm{SU}(n)$ contains a unique diagonal matrix up to permutation of the diagonal.

In the case of SU(2), this classification of diagonal matrices is particularly nice, as the following proposition demonstrates.

Proposition 4.2.7. *Conjugacy classes in* $\mathrm{SU}(2)$ *are classifed by the trace function* $\mathrm{tr} : \mathrm{SU}(2) \to [-2, 2]$.

Proof. By the above, it suffices to show that each diagonal matrix with different diagonal entries has a different trace. In SU(2), each diagonal matrix D has elements elements λ and $\overline{\lambda}$ on its diagonal for some λ such that $|\lambda| = 1$. Then, we have $\operatorname{tr} D = \lambda + \overline{\lambda} = 2\operatorname{Re}(\lambda) \in [-2, 2]$. On the other hand, notice that only two elements in S^1 have real part equal to $\operatorname{Re}(\lambda)$: λ and $\overline{\lambda}$. So, any diagonal matrix D' has trace $2\operatorname{Re}(\lambda)$ if and only if it has λ or $\overline{\lambda}$ on its diagonal, and hence also the conjugate of this. But this is true if and only if D' is identical to D up to possibly permuting the diagonal elements, which as noted above is true if and only if D and D' are in the same conjugacy class. $\square$

We now shift gears to define a representation of $\mathrm{U}(n)$ that will help us better understand the unitary group. Let $W_+ \subset \operatorname{Mat}(n, \mathbb{C})$ be the set of Hermition matrices (i.e. those satisfying $A^* = A$), and let $W_- \subset \operatorname{Mat}(n, \mathbb{C})$ be the set of anti-Hermitian matrices (i.e. those satisfying $A^* = -A$). Let W_+^0 and W_-^0 be the subsets of W_+ and W_- (respectively) consisting of matrices of trace 0. Now, we can define a representation of $\mathrm{U}(n)$ by having it act on $\operatorname{Mat}(n, \mathbb{C})$ by left multiplication. Then, for any $U \in \mathrm{U}(n)$ and any A in either W_+ or W_-, we have

$$UA(UA)^* = UAA^*U^* = U(\pm I)U^* = \pm I,$$

so that UA is again in W_+ or W_-. Thus, both W_+ and W_- are subrepresentations of $\operatorname{Mat}_n(\mathbb{C})$. Indeed, one can check that we have an intertwiner $W_+ \to W_-$ defined by multiplication by i, which has an inverse defined by multiplication by $-i$. So, $iW_+ \cong W_-$ as representations. We can use these representations to give a decomposition of the representation $\operatorname{Mat}_n(\mathbb{C})$ into a direct sum of irreducible representations.

Proposition 4.2.8. *As real representations of* $\mathrm{U}(n)$*,*

$$\operatorname{Mat}_n(\mathbb{C}) \cong W_+^0 \oplus \mathbb{R} \cdot I \oplus W_-^0 \oplus i\mathbb{R} \cdot I.$$

Moreover, all of the representations on the right-hand side of this isomorphism are irreducible.

Proof. One can check that the map

$$\begin{aligned} \operatorname{Mat}_n(\mathbb{C}) &\xrightarrow{\sim} W_+ \oplus W_- \\ B &\mapsto \left(\frac{B + B^*}{2}, \frac{B - B^*}{2}\right) \end{aligned}$$

is an isomorphism of vector spaces which respects the action of $\mathrm{U}(n)$, so it is an isomorphism of representations. Likewise, one can check that the map

$$\begin{aligned} W_+ &\xrightarrow{\sim} W_+^0 \oplus \mathbb{R} \cdot I \\ B &\mapsto \left(B - \frac{\operatorname{tr} B}{nI}, \frac{\operatorname{tr} B}{nI}\right) \end{aligned}$$

is an isomorphism of representations, as is the map

$$\begin{aligned} W_- &\xrightarrow{\sim} W_-^0 \oplus i\mathbb{R} \cdot I \\ B &\mapsto \left(B - \frac{\operatorname{tr} B}{nI}, \frac{\operatorname{tr} B}{nI}\right) \end{aligned}$$

Putting this all together, we see that

$$\mathrm{Mat}_n(\mathbb{C}) \cong W_+ \oplus W_- \cong W_+^0 \oplus \mathbb{R} \cdot I \oplus W_-^0 \oplus i\mathbb{R} \cdot I.$$

$\mathbb{R} \cdot I$ and $i\mathbb{R} \cdot I$ are irreducible because they are 1-dimensional. One can prove that W_+^0 and W_-^0 are also irreducible. □

Now, we will specialize to consider SU(2). For any matrix A in SU(2), the unitary condition gives us

$$\begin{pmatrix} a & b \\ -\overline{b} & \overline{a} \end{pmatrix},$$

and the determinant-1 condition tells us that $|a|^2 + |b|^2 = 1$. Thus, elements of SU(2) are in bijection with elements on the 3-sphere $S^3 \subset \mathbb{R}^4 \cong \mathbb{C}^2$. In fact, one can check that this gives S^3 the structure of a topological group. It's worth noting that the isomorphism $\mathrm{SU}(2) \cong S^3$ is rather unique: the only other spheres which are topological groups are S^1 and $S^0 \cong C_2$.

Notice that, in the above expression of A, we have $\operatorname{tr} A = a + \overline{a} = 2x_1$, where (x_1, x_2, x_3, x_4) represents (a, b) in $\mathbb{R}^4$. So, the traceless elements of SU(2), which constitute one conjugacy class of SU(2) by Proposition 4.2.7, are precisely those of the form

$$\begin{pmatrix} ix_2 & b \\ -\overline{b} & ix_2 \end{pmatrix}.$$

Every element in this conjugacy class is conjugate to iI, and since conjugacy preserves order, this means that every matrix in the conjugacy class has order 4. In fact, one can show that the elements of order 4 are precisely those that have trace 0, so this completely classifies order-4 elements of SU(2). (For a proof, see Homework 11.) We can also completely classify the elements of order 2 in SU(2): they are conjugate to a diagonal matrix D, and the condition $D^2 = I$ implies that the diagonal elements of D are 2nd root of unity; the determinant-1 condition implies that both are ± 1, so the fact that D is not the identity (which has order 1) implies that $D = -I$. Thus, $-I$ is the only element of order 2 in SU(2).

We wish to use the representation of U(2) by $\mathrm{Mat}_n(\mathbb{C})$ that we defined above to help understand the structure of SU(2). To do this, we need to consider the interactions of this matrix algebra with the quaternions. Recall that $\mathbb{H} \cong \mathbb{R}^4$ (as vector spaces); moreover, if $\mathbb{H}_1$ denotes the quaternions of norm 1, then the map

$$\begin{aligned} \mathbb{H}^\times &\xrightarrow{\sim} \mathbb{H}_1^\times \times \mathbb{R}_{>0} \\ q &\mapsto (\frac{q}{|q|}, |q|).) \end{aligned}$$

is an isomorphism. Now, recall that we have a representation $\rho : Q_8 \to \mathrm{Mat}(2, \mathbb{C})$ defined by setting $\psi(1) = I$,

$$\rho(i) = \underline{i} = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix},$$

$$\rho(j) = \underline{j} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix},$$

and

$$\rho(k) = \underline{k} = \begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}$$

(see Homework 6.) Then, one can check that ρ is faithful, so if we extend ρ linearly to a homomorphism $\psi : \mathbb{H} \to \mathrm{Mat}_2(\mathbb{C})$, then ψ will be injective. This implies that ψ is an isomorphism between $\mathbb{H}$ and $\operatorname{Im}\psi = \mathbb{R} \cdot I \oplus \mathbb{R} \cdot \underline{i} \oplus \mathbb{R} \cdot \underline{j} \oplus \mathbb{R} \cdot \underline{k}$. On the other hand, one can check that

$$W_-^0 = \mathbb{R}\underline{i} \oplus \mathbb{R}\underline{j} \oplus \mathbb{R}\underline{k}.$$

So, ψ furnishes an isomorphism $\mathbb{H} \cong \mathbb{R} \cdot I \oplus W_-^0$. Explicitly, this isomorphism is defined by

$$\psi(a + bi + cj + dk) = \begin{pmatrix} a + bi & c + di \\ -c + di & a - bi \end{pmatrix} = \begin{pmatrix} z & w \\ -\overline{w} & \overline{z} \end{pmatrix},$$

where $z = a+bi$ and $w = c+di$. Moreover, if we assume that $|a+bi+cj+dk|^2 = 1$ in the above, we get $|z|^2 + |w|^2 = 1$, so that $\psi(\mathbb{H}_1^\times) = S^3 = \mathrm{SU}(2)$. If we now consider the set $\mathbb{H}_0$ of purely imaginary quaterions (i.e. the ones with no real component), then we see from the definition of ψ that $\psi(\mathbb{H}_0) = \mathbb{R}\underline{i} \oplus \mathbb{R}\underline{j} \oplus \mathbb{R}\underline{k} = W_-^0$. (Note that $\mathbb{H}_0$ and W_-^0 are only vector subspaces, not subalgebras, so the isomorphism given by ψ on these is only an isomorphism of vector spaces.)

Now, consider conjugating $\mathbb{H}_0$ by elements of $\mathbb{H}_1^\times$. We claim that this defines a group action. To see this, we just need to check that it fixes $\mathbb{H}_0$. So, for any $q_1 = ai + bj + ck \in \mathbb{H}_0$ and any $q_2 = d + ei + fj + gk \in \mathbb{H}_1^\times$, we have

$$\begin{aligned} \operatorname{Re}(q_2 q_1 q_2^{-1}) &= \operatorname{Re}(q_2 q_1 \overline{q_2}) = \operatorname{Re}(q_2(ai + bj + ck)(d - ei - fj - gk)) \\ &= \operatorname{Re}(q_2[-(ae + bf + cg) - (ad + bg - cf)i - (-ag + bd + ce)j - (af - be + cd)k]) \\ &= -d(ae + bf + cg) + e(ad + bg - cf) + f(-ag + bd + ce) + g(af - be + cd) \\ &= 0 \end{aligned}$$

This proves that $q_2 q_1 q_2^{-1} \in \mathbb{H}_0$, so that conjugation does define an action of $\mathbb{H}_1^\times$ on $\mathbb{H}_0$. Then, using the isomorphisms $\mathrm{SU}(2) \cong \mathbb{H}_1^\times$ and $W_-^0 \cong \mathbb{H}_0$, we get an action of $\mathrm{SU}(2)$ on W_-^0, which gives rise to a homomorphism $\varphi : \mathrm{SU}(2) \to \mathrm{GL}(W_-^0) \cong \mathrm{GL}(3, \mathbb{R})$. We now define an inner product on W_-^0 by setting

$$\langle X, Y \rangle = -\frac{1}{2} \operatorname{tr}(XY).$$

Under this inner product, one can check that the elements $\underline{i}$, $\underline{j}$, and $\underline{k}$ form an orthonormal basis. Considered on $\mathbb{H}_0$, the inner product is given by $\langle q_1, q_2 \rangle = \operatorname{Re}(q_1 \overline{q_2})$. Now, one can check that the action of $\mathbb{H}_1^\times$ preserves this inner product on $\mathbb{H}_0$, which implies that the image of ϕ lies inside $\mathrm{SO}(3)$.

In fact, one can prove that φ is surjective, so that the First Isomorphism Theorem implies:

$$\mathrm{SO}(3) \cong \mathrm{SU}(2)/(\ker \varphi).$$

Now, suppose $A \in \mathrm{SU}(2)$ is in the kernel of φ. Then, A acts by conjugation on the matrices of W_-^0, so for all $B \in W_-^0$, we have $ABA^{-1} = B$, or equivalently,

$AB = BA$. On the other hand, we showed above that $\psi(\mathbb{H}) = \mathbb{R} \cdot I \oplus W_-^0$; A commutes with all of W_-^0, and it commute swith I, so it commutes with all of $\psi(\mathbb{H})$ and in particular with $\psi(\mathbb{H}_1^\times) = \mathrm{SU}(2)$. However, the only elements in the center of SU(2) are multiples of the identity, and the only such multiples which have determinant 1 are $\pm I$. So, we have shown that $B = \pm I$ and $\ker \varphi = \{\pm I\}$. Summarizing all of this in one theorem, we have:

Theorem 4.2.9. *Let $\varphi : \mathrm{SU}(2) \to \mathrm{SO}(3)$ be the homomorphism defined above. Then, φ is surjective, and $\ker \varphi = \{\pm I\}$. In particular,*

$$\mathrm{SO}(3) \cong \mathrm{SU}(2)/\{\pm I\}.$$

Remark. Under our identification of $SU(2)$ with S^3, the above quotient amounts to identifying x with $-x$ on S^3, so that $SO(3) \cong \mathbb{RP}^2$.

It's worth pointing out just how much work the quaternions are doing for us here. For general m and n, there is no relationship between $\mathrm{SU}(n)$ and $\mathrm{SO}(m)$. Indeed, the only other scenario where we do have a relationship is:

$$SU(2) \times SU(2)/\{(1,1),(-1,-1)\} \cong SO(4).$$

But this isomorphism is actually also proven using the quaternions (but this time using the action given by left or right multiplication instead of conjugation). In short, these relationships are unique to low-dimensional unitary/orthogonal groups and arise only because the quaternions behave so nicely with these groups.

With all of these results built up, we are now ready to classify the finite subgroups of SU(2).

Theorem 4.2.10. *Suppose G is a finite subgroup of $SU(2)$. Then, one of the following holds.*

1. *$G \cong C_n = \langle a \mid a^n = 1\rangle$ is the cyclic group of order n for some n.*
2. *$G \cong D_n^* = \langle a, b \mid a^2 = b^2 = (ab)^n\rangle$ is the binary dihedral group of order n for some n.*
3. *$G \cong A_4^* = \langle a, b \mid a^2 = b^3 = (ab)^3\rangle$ is the binary tetrahedral group.*
4. *$G \cong S_4^* = \langle a, b \mid a^2 = b^3 = (ab)^4\rangle$ is the binary octahedral group.*
5. *$G \cong A_5^* = \langle a, b \mid a^2 = b^3 = (ab)^5\rangle$ is a binary dodecahedral group.*

Proof. Write $H = \varphi(G) \subset \mathrm{SO}(3)$, where φ is the map $\mathrm{SU}(2) \to \mathrm{SO}(3)$ defined above. Then, H is a finite subgroup of SO(3), so we can use our classification of these (Theorem 4.2.4) to determine what it is. We will go through each possibility in turn.

1. Suppose $H \cong C_n$ for some n. Then, there are two possibilities. If the order of G is odd, then G cannot contain $-I$ (since $-I$ has order 2, which does not divide $|G|$). So, quotienting SU(2) by $\{\pm I\}$ does not affect G, which

implies that $G \cong H \cong C_n$. Otherwise, the order of G is even, so G must contain an element of order 2. The only such element in SU(2) is $-I$, so $-I \in G$. Then, $G/\{\pm 1\} \cong H$ implies that $G \cong H \times \{\pm I\} \cong C_n \times C_2 \cong C_{2n}$. In either case, one can check that G must be conjugate to the subgroup

$$\left\{ \begin{pmatrix} \zeta & 0 \\ 0 & \zeta^{-1} \end{pmatrix}^k \mid 0 \leq k < n \right\},$$

where $\zeta = \exp(2\pi i/n)$.

2. Suppose $H \cong D_n$ for some n. Then, we have $|G| = 2|H| = 4n$ is even, so $-I \in G$. Now, D_n contains a copy of C_n as a normal subgroup; since n is even, we have by our arguments in case 1 above that $\varphi^{-1}(C_n) \cong C_{2n} \subseteq G$. Then, $D_n = C_n \sqcup r \cdot C_n$ for some s implies that $G = \varphi^{-1}(H) \cong C_{2n} \sqcup s \cdot C_{2n}$. Now, r has order 2 in D_n, so s has order 4 in SU(2). Then, $s^2 = -I$ is the unique order-2 element of SU(2). Now, (up to conjugation), we have already said that C_{2n} is generated by $c = \left(\begin{smallmatrix} \zeta & 0 \\ 0 & \zeta^{-1} \end{smallmatrix}\right)$ for some primitive $2n$th root of unity ζ. Moreover, one can check that we can replace s with sc^m for some m and so take s to be $\left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$. (The idea is that the elements of order 4 of SU(2) form a 2-sphere in $S^3 \cong \mathrm{SU}(2)$, and s acts as a rotation by π in H; from this, it follows that s is just some rotation of $\left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$ by some power of c.) So, up to conjugation, we have shown that $G \cong \gamma^{-1}(H) = D_n^*$, where D_n^* is the binary dihedral group of order $4n$, which is generated by the two matrices c and s defined above. Letting $a = s$ and $b = cs$, one can see that $a^2 = b^2 = (ab)^n$ (they are all equal to $-I$), which defines the presentation of the group D_n^*. Notice also that $D_n^* \not\cong D_{2n}$: the latter has many elements of order 2, but the former has only 1.

3. The remaining three cases are $H \cong A_4$, $H \cong S_4$, and $H \cong A_5$. All three of these groups have even order, which means that G does as well. By much the same arguments as for the case where $H \cong D_n$, one can explicitly compute $G = \varphi^{-1}(H)$. Once again, one checks that these are not isomorphic to obvious groups that we already know: for instance, A_4^* has order $2|A_4| = |S_4|$, but A_4^* is not isomorphic to S^4, since the former has 1 element of order 4 while the latter has several.

□

4.2.3 McKay Graphs

Now that we know what the finite subgroups of SU(2) are, we wish to understand the representation theory of these subgroups. First, notice that SU(2) has a fundamental representation defined by letting the matrices of SU(2) act as linear operators on $\mathbb{C}^2$. Given any finite subgroup G of SU(2), we can restrict this fundamental representation to G to get a representation V of G. Notice that the fundamental representation is faithful by definition, which means that its restriction, V, is as well. It turns out that V enjoys some nice properties, as the following two propositions show.

Proposition 4.2.11. *Let G be a finite subgroup of* SU(2)*, and let V be the restriction of the fundamental representation of* SU(2) *to G. Then, V is reducible if and only if $G \cong C_n$ for some n.*

Proof. Suppose V is reducible. Then, we can write $V = U_1 \oplus U_2$. Since V is 2-dimensional, the U_i have dimension 1. So, the image of G in $\mathrm{GL}(V) \cong \mathrm{GL}(2, \mathbb{C})$ lies in $\mathrm{GL}(1) \times \mathrm{GL}(1) \cong \mathbb{C}^* \times \mathbb{C}^*$. But since this representation is faithful, G is then isomorphic to a subgroup of $\mathbb{C}^* \times \mathbb{C}^*$, which means it is abelian. But by Theorem 4.2.10, the only finite abelian subgroups of $SU(2)$ are the cyclic groups, so $G \cong C_n$ for some n. Conversely, if $G \cong C_n$, then G is abelian, so all of its irreducible representations are 1-dimensional. But V has dimension 2, so it must be reducible. □

Proposition 4.2.12. *Let G be a finite subgroup of* SU(2)*, and let V be the restriction of the fundamental representation of* SU(2) *to G. Then, V is a self-dual representation of G.*

Proof. Every element g of G is an element of SU(2), so it is conjugate in SU(2) to a diagonal matrix with diagonal entries λ and λ^{-1}, where $\lambda^{-1} = \overline{\lambda}$. This implies that $\chi_V(g) = \lambda + \overline{\lambda} = \overline{\chi_V(g)} = \chi_{V^*}(g)$. So, we have $\chi_V = \chi_{V^*}$ and $V \cong V^*$. □

The fact that V is self-dual is actually very important for our purposes. Notice that, for any finite-dimensional vector spaces U, V, and W over a field k, we have

$$\mathrm{Hom}_k(U \otimes V, W) \cong \mathrm{Hom}_k(U, V^* \otimes W).$$

In particular, taking $U = k$ gives $\mathrm{Hom}(V, W) \cong V^* \otimes W$. If now we suppose that U, V, and W are representations of a group G, then the same hom set equality is true, but this time with hom sets consisting of intertwiners rather than linear maps. Comparing dimensions of both sides, of this equality, we have

$$\dim(\mathrm{Hom}_G(U \otimes V, W)) = \langle \chi_U \chi_V, \chi_W \rangle = \dim(\mathrm{Hom}_G(U, V^* \otimes W)) = \langle \chi_U, \chi_{V^*} \chi_W \rangle.$$

Because V is self-dual, $\chi_V = \chi_{V^*}$. So, we have effectively moved χ_V from the first factor of the inner product to the second without changing the inner product.

Now, for any finite subgroup G of SU(2), let $V_1, \ldots, V_m$ be the irreducible representations of G, and let $\chi_i = \chi_{V_i}$ for all i. We can define a graph Γ by making one vertex for each V_i, $i \in \{1, \ldots, m\}$, and, for every i and j, making $m(i, j) = \langle \chi_{V_i}, \chi_V \chi_{V_j} \rangle$ edges between vertices corresponding to V_i and V_j. (Note that this inner product is precisely the multiplicity of V_i in the irreducible decomposition of $V \otimes V_j$.) Now, because V is self-dual,

$$m(i, j) = \langle \chi_{V_i} \chi_V, \chi_j \rangle = \langle \chi_{V_i}, \chi_V \chi_j \rangle = m(j, i),$$

so that we have an unambiguous description of the graph Γ. To complete our definition of Γ, we make it into a weighted graph, where the vertex corresponding to V_i is $\dim V_i$. Notice that this construction will work for any G and any self-dual representation V. We call it the *McKay graph* of G.

Example. Take $G \cong C_n$, and let g be a generator of G. Now, g acts on V by the matrix it represents, which is diagonal in some basis. This diagonal matrix has order dividing $|G| = n$ and is unitary, so its diagonal elements are $|G|$th roots of unitary that are conjugate to each other. Suppose they are ζ and $\overline{\zeta}$. Notice that g has order n in G, so ζ must have order n and so be a primitive root of unity.

Now, we can define representations $V_0, \ldots, V_{n-1}$ of G by having g act on $V_i \cong \mathbb{C}$ by multiplication by ζ^i for all i. As i runs from 0 to $n-1$, g will then act by each different nth root of unity, so the V_i are precisely the 1-dimensional representations of G. Notice that from the above, g acts by ζ in one coordinate and $\overline{\zeta} = \zeta^{-1} = \zeta^{n-1}$ in the other. This implies that, for all i, g^i acts by ζ^i in one coordinate and ζ^{n-i} in the other coordinate. So, we have

$$V \cong V_1 \oplus V_{n-1}.$$

Because tensor products distribute through direct sums, we have for all j

$$V_j \otimes V \cong V_j \otimes (V_1 \oplus V_{n-1}) \cong (V_j \otimes V_1) \oplus (V_j \otimes V_{n-1}) \cong V_{j+1} \oplus V_{j-1},$$

where the last isomorphism can be obtained by compared characters. This implies that $m(j, j+1) = m(j, j-1) = 1$ and $m(j, i) = 0$ for all $i \neq j \pm 1$. So, the McKay graph of G has n vertices, each with one edge to the previous vertex and one edge to the next vertex. One can draw this as a regular n-gon.

As the above example illustrates, McKay graphs can encapsulate much of the representation theory of these groups. For this reason, we will work to understand the McKay graphs of all the finite subgroups of SU(2). We begin by investigating some of the properties of these graphs.

Theorem 4.2.13. *Let G be any finite subgroup of* SU(2)*. The McKay graph of G is connected.*

Proof. Let $V_0, \ldots, V_{n-1}$ denote the irreducible representations of G, and let $\chi_i = \chi_{V_i}$ for all i. We show that V_i and V_0 are in the same connected component for all i¿ Notice that V_i and V_0 are in the same connected component if and only if $V_i \subset V^{\otimes n}$ for some n, i.e. if and only if $\langle \chi_i, \chi_V^n \rangle \neq 0$ for some n. Since we have seen in the above example that the McKay grph of C_n is connected, we can restrict ourselves to the case $G \not\cong C_n$, in which case the order of G is even and $-I \in G$.

Now, we have for all n

$$\begin{aligned}\frac{1}{2^n}\langle \chi_i, \chi_V^n \rangle &= \frac{1}{|G|2^n} \sum_{g \in G} \chi_i(g)\overline{\chi_V(g)^n} \\ &= \frac{1}{|G|}\left(\frac{\chi_i(I)\chi_V(I)^n + \chi_i(-I)\chi_V(-I)^n}{2^n} + \sum_{g \neq \pm I} \frac{\chi_i(g)\chi_V(g)^n}{2^n} \right)\end{aligned}$$

(here we have used that $\chi_V = \chi_{V^*}$ by Proposition 4.2.12). Now, by Proposition 2.4.3, we have for all g that $|\chi_V(g)| \leq 2$, with equality if and only if g is a multiple of the identity. But the only multiples of the identity in SU(2) are $\pm I$, so we see that $|\chi_V(g)| < 2$ for every term of the sum in the above equation, which implies that this sum goes to 0 in the limit as n goes to infinity. On the other hand, we have $\chi_V(I) = \operatorname{tr} I = 2$ and $\chi_V(-I) = \operatorname{tr}(-I) = -2$, so that

$$\lim_{n\to\infty} \frac{1}{2^n}\langle \chi_i, \chi_V^n\rangle = \frac{1}{|G|}\lim_{n\to\infty} \frac{\chi_i(I)2^n + \chi_i(-I)(-2)^n}{2^n} + 0 = \frac{1}{|G|}\lim_{n\to\infty} \chi_i(I) + (-1)^n\chi_i(-I).$$

Now, $\chi_i(I) = \dim V_i$ and $\chi_i(-I) = \pm \dim V_i$ ($-I$ either by I or by $-I$ in V_i), so the above limit becomes

$$\lim_{n\to\infty} \frac{1}{2^n}\langle \chi_i, \chi_V^n\rangle = \dim V_i \lim_{n\to\infty} 1 \pm (-1)^n.$$

This limit does not exist, so in particular it does not converge to 0. This implies that there exists some n such that

$$\frac{1}{2^n}\langle \chi_i, \chi_V^n\rangle \neq 0$$

(otherwise, the limit would converge to 0). But then, we have $\langle \chi_i, \chi_V^n\rangle \neq 0$, which as noted above implies that V_0 and V_i are in the same connected component. □

Remark. Although we did use some properties of V in the above proof, one can still generalize the argument in the proof to prove the result when G is any finite group and V is any faithful representation of G. For a proof, see Homework 11.

Proposition 4.2.14. *Let G be a finite subgroup of* $\mathrm{SU}(2)$*, and let Γ be the McKay graph of G. Then, Γ is bipartite if and only if $G \not\cong C_n$ for any n odd. Moreover, if this is the case, a bipartition of the vertices of Γ is given by splitting the irreducible representations of G into the set of those in which $-I$ acts by the identity and the set of those in which $-I$ acts by $-I$.*

Proof. First, notice that the only odd-ordered finite subgroups of SU(2) are the cyclic groups of odd order. Since any even subgroup must contain an element of order 2, and since the only element of order 2 in SU(2) is $-I$, we see that G contains $-I$ if and only if $G \not\cong C_n$ for any n odd. Now, by inspecting the McKay graph of C_n for n odd (which we found explicitly in the above example), one can see that the graph is not bipartite. So, it remains to prove that when G contains $-I$, then Γ is bipartite in the desired way.

Notice that for any representation W of G, if $-I$ acts by a matrix A on W, then $(-I)^2 = I$ acts by A^2, so we must have $A^2 = I$. Because $A \in \mathrm{SU}(2)$, we get $A = \pm I$. If W is in fact irreducible, then

$$\chi_{V\otimes W}(-I) = \chi_V(-I)\chi_W(-I) = \operatorname{tr}(-I)\operatorname{tr}(A) = \begin{cases} 2, & A = I \\ -2, & A = I \end{cases}.$$

Since $-I$ must act by either I or $-I$ on $V \otimes W$, we see that $-I$ acts by $-I$ on $V \otimes W$ (and hence on all irreducible subrepresentations of this representation) if $A = I$ and by I if $A = -I$. Thus, however $-I$ acts on W, it will act differently on every irreducible representation of G which is connected to W in Γ. This proves that Γ is bipartite with bipartition given in the proposition statement. □

The above proposition allows us to relate the irreducible representations of finite subgroups of $\mathrm{SU}(2)$ to those of the finite subgroups of $\mathrm{SO}(3)$. Suppose G is any non-cyclic finite subgroup of $\mathrm{SU}(2)$, so that G contains $-I$. Let $\varphi : G \to \mathrm{GL}(W)$ be an irreducible representation of G, and suppose that $\varphi(-I) = I$. Then, φ factors through a homomorphism $\overline{\varphi} : G/\{\pm I\} \to \mathrm{GL}(W)$. But under the isomorphism $\mathrm{SO}(3) \cong \mathrm{SU}(2)/\{\pm I\}$, $G/\{\pm I\}$ is isomorphic to a finite subgroup H of $\mathrm{SO}(3)$, so $\overline{\varphi}$ is a representation of H. One can check that given any proper subrepresentation X of $\overline{\varphi}$, the preimage of X under the quotient map gives a proper subrepresentation of φ. So, the fact that φ is irreducible implies that $\overline{\varphi}$ is irreducible as well. On the other hand, if $\varphi(-I) = -I$, then φ does not factor through any such map, so it does not induce an irreducible representation of $G/\{\pm I\}$. In summary, we have proven the following corollary to the above proposition:

Corollary 4.2.15. *Let G be a non-cyclic finite subgroup of $\mathrm{SU}(2)$. Then, the McKay graph of G has a bipartition given by splitting the irreducible representations of G into the set of those that descend to irreducible representations of the group $G/\{\pm I\}$ (which is a finite subgroup of $\mathrm{SO}(3)$ via the isomorphism $\mathrm{SO}(3) \cong \mathrm{SU}(2)/\{\pm I\}$) and the set of those that do not descend to $G/\{\pm I\}$.*

We now establish a few more nice relationships between the McKay graphs of finite subgroups of $\mathrm{SU}(2)$ and the representation theory of these groups.

Proposition 4.2.16. *Let G be any finite subgroup of $\mathrm{SU}(2)$, and let $V_1, \ldots, V_n$ denote the irreducible representations of G. Define $d_i = \dim V_i$. Then, for any i,*

$$2d_i = \sum_{j=1}^{n} m(i,j) d_j.$$

Proof. By definition of the $m(i,j)$,

$$V_i \otimes V = \bigoplus_{j=1}^{n} V_j^{m(i,j)}$$

for all i. Comparing dimensions of both sides of this equation, we get

$$(\dim V_i)(\dim V) = 2d_i = \sum_{j=1}^{n} m(i,j) \dim V_j = \sum_{j=1}^{n} m(i,j) d_j.$$

□

Proposition 4.2.17. *Let G be a finite subgroup of $\mathrm{SU}(2)$, and let Γ be the McKay graph of G. Then, Γ has a loop if and only if $G \cong C_1$, and Γ has multiple edges between the same two vertices if and only if $G \cong C_2$.*

Proof. First, if $G \cong C_1$, then the one irreducible representation of G is the trivial representation, V_0. Moreover, since V is the 2-dimensional representation defined by letting the one element of G act by the identity, we have $V_0 \otimes V \cong V \cong V_0 \oplus V_0$.

So, Γ one vertex with one loop on it. Conversely, if Γ has a loop, then it cannot be bipartite, which by Proposition 4.2.14 implies that $G \cong C_n$ for some n odd. But we have seen in the above example that the McKay graph of C_n does not have a loop for $n > 1$, so we must have $n = 1$ and $C \cong C_1$.

Let $V_1, \ldots, V_n$ be the irreducible representations of G. If $G \cong C_2$, then by our above example, Γ looks like

In particular, Γ has more than one edge between its two vertices. Conversely, suppose that $m(i,j) \geq 2$ for some i and j. Then, by Proposition 4.2.16 above, we have

$$2d_i = m(i,j)d_j + \sum_{k \neq j} m(i,k)d_k$$

and likewise

$$2d_j = m(i,j)d_i + \sum_{k \neq i} m(j,k)d_k.$$

Manipulating the first of these equations and then substituting in the second, we get

$$\begin{aligned} 2d_i &= 2d_j + (m(i,j) - 2)d_j + \sum_{k \neq j} m(i,k)d_k \\ &= m(i,j)d_i + \sum_{k \neq i} m(j,k)d_k + (m(i,j) - 2)d_j + \sum_{k \neq j} m(i,j)d_k \end{aligned}$$

Subtracting $2d_i$ from both sides of this equation gives

$$(m(i,j) - 2)d_i + \sum_{k \neq i} m(j,k)d_k + (m(i,j) - 2)d_j + \sum_{k \neq j} m(i,j)d_k = 0.$$

On the other hand, since $m(i,j) \geq 2$, every term on the left-hand side of this equation is nonnegative, so they must all be 0. Thus, $m(i,j) = 2$, $m(j,k) = 0$ for all $k \neq i$, and $m(i,k) = 0$ for all $k \neq j$. If there are any vertices other than those corresponding to V_i and V_j, then Γ must be disconnected, contradicting Theorem 4.2.13. So, V_i and V_j must correspond to the only vertices of Γ. Then, G has two irreducible representations, hence precisely two conjugacy classes. The only such subgroup of $\mathrm{SU}(2)$ is C_2. $\square$

Proposition 4.2.18. *Let G be any finite non-cyclic subgroup of* $\mathrm{SU}(2)$*, and let V_0 denote the trivial representation of G. Then, the vertex representing V_0 in the McKay graph of G has only one edge coming from it, and that edge goes to the vertex representing V. In other words, the graph has a portion of the form*

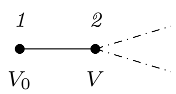

Proof. Since $G \ncong C_n$, Proposition 4.2.11 implies that V is irreducible. Notice that $V_0 \otimes V \cong V$, so that the only irreducible representation in $V_0 \otimes V$ is V itself, which occurs with multiplicity one. Thus, V_0 has no edge to any vertex except for V, and there is precisely one edge from V_0 to V. $\square$

We now introduce a general framework to help classify the McKay graphs. For any graph Γ with vertices $\{e_1, \ldots, e_n\}$, we can define an inner product space $V(\Gamma)$ by taking the vector space $\oplus_{i=1}^n \mathbb{R} \cdot e_i$ with basis vectors corresponding to vertices of Γ, setting

$$\langle e_i, e_j \rangle = \begin{cases} 2, & i = j \\ -1, & i \neq j, i \text{ and } j \text{ connected} \\ 0, & \text{otherwise} \end{cases},$$

and extending $\langle \cdot, \cdot \rangle$ linearly to all of $V(\Gamma)$. In order to talk about the properties of these inner product spaces and their associated graphs, we first need to define some terminology.

Definition 4.2.19. Let Γ be a connected graph. We say that Γ is *simply laced* if it has no multiple edges. We say that a simply laced graph is an *affine graph* if we can assign positive weights $d_1, \ldots, d_n$ to each of the n vertices of Γ such that

$$2d_i = \sum_{i-j} d_j.$$

(Here "$i - j$" means "all vertices j which are connected to vertex i.") We say that a simply laced graph is a *finite Dynkin diagram* if it is a proper subgraph of an affine graph. We say that Γ is *indefinite* if it properly contains an affine graph.

The connection between the above properties of graphs and their associated inner product spaces is given by the following proposition.

Proposition 4.2.20. *Let Γ be a simply laced graph.*

1. *If Γ is finite, then the inner product on $V(\Gamma)$ is positive definite.*

2. *If Γ is affine, then the inner product on $V(\Gamma)$ is positive semi-definite, and its null space is spanned by the element $w_0 = \sum_{i=1}^n d_i e_i$, where $e_1, \ldots, e_n$ are the vertices of Γ and $d_1, \ldots, d_n$ are the weights on the e_i which make Γ affine.*

3. *If Γ is indefinite, then the inner product on $V(\Gamma)$ is indefinite.*

Proof. See Qi You's notes (available on the course website). $\square$

By Propositions 4.2.16 and 4.2.17, the McKay graphs of all finite subgroups of SU(2) are affine. In fact, one can completely classify the affine graphs and show that each of them corresponds to precisely one finite subgroup of SU(2). Unfortunately, I have not had time to type up the proofs of these powerful facts, so I will simply state them here. For the full proofs, see Qi You's notes (available on the course website).

Theorem 4.2.21. *The following is a complete list of affine graphs.*

1. *For $n \geq 2$, we define the graph $\widetilde{A_n}$ to be the graph with $n+1$ vertices with edges forming an $n+1$-gon, like so:*

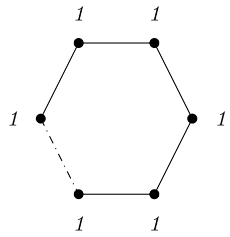

(Note that the numbers next to each vertex denote the weights which make the graph affine.)

2. *For $n \geq 5$, we define the graph $\widetilde{D_n}$ to be the graph with $n+1$ vertices in the form:*

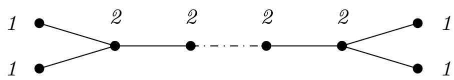

3. *We define the graph $\widetilde{E_6}$ to be*

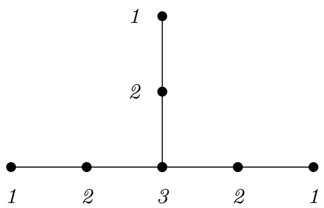

4. *We define the graph $\widetilde{E_7}$ to be*

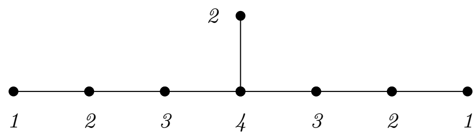

5. *We define the graph $\widetilde{E_8}$ to be*

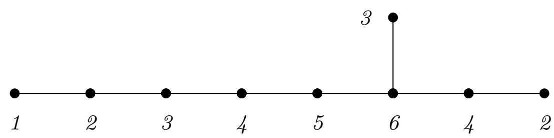

Theorem 4.2.22 (McKay Correspondence)**.** *The finite subgroups of* SU(2) *other than C_1 and C_2 are in 1-to-1 correspondence with the affine graphs, with the correspondence given by taking the McKay graph of a given subgroup. More precisely:*

1. the McKay graph of C_n for $n > 2$ is $\widetilde{A_{n-1}}$;

2. the McKay graph of D_n^* is $\widetilde{D_{n+2}}$;

3. the McKay graph of A_4^* is $\widetilde{E_6}$;

4. the McKay graph of S_4^* is $\widetilde{E_7}$; and

5. the McKay graph of A_5^* is $\widetilde{E_8}$.

Although we have not explicitly found the representations of the finite non-cyclic subgroups of SU(2), the results we've developed actually encapsulate a lot of the representation theory of these subgroups. We end this section with a couple examples that demonstrate just how much information we can glean about these subgroups using the results in this section.

Example. Consider the group A_4^*, which by the McKay correspondence has McKay graph $\widetilde{E_6}$. Now, one of the weight-1 vertices in $\widetilde{E_6}$ must correspond to the trivial representation V_0 of A_4^*. Because of the symmetry of the graph, we can pick V_0 to correspond to any of the weight-1 vertices. Moreover, by Proposition 4.2.14 above, because A_4^* is non-cyclic, $\widetilde{E_6}$ is bipartite. If we mark all the vertices of $\widetilde{E_6}$ which descend to irreducible representations of $A_4^*/\{\pm I\} \cong A_4$ with an "x," then, noting that $-I$ acts by I on V_0, the graph of $\widetilde{E_6}$ looks like:

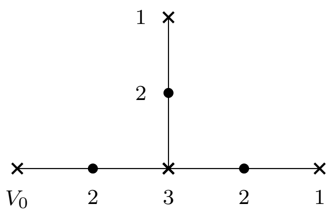

Let's discuss this correspondence with irreducible representations of A_4 from an algebraic perspective. One can check that the elements of $[A_4, A_4]$ are all products of 2 disjoint 2-cycles, which implies that $[A_4, A_4] \cong C_2 \times C_2$. Then, $A_4/[A_4, A_4] \cong C_3$, so that A_4 has 3 1-dimensional representations. Now, because all of the weight-1 vertices in the above diagram have "x"s on them (i.e. all the 1-dimensional representations of A_4^* descend to A_4), we have a surjective map

$$A_4^*/[A_4^*, A_4^*] \to A_4/[A_4, A_4].$$

On the other hand, by looking at the diagram $\widetilde{E_6}$, we see that A_4^* has exactly 3 1-dimensional representations, so in fact this surjective map is between finite groups of the same size, which implies that it is injective and hence an isomorphism.

Now, we can also understand what representations correspond to the various vertices of the above graph. First, notice that by Proposition 4.2.18, the weight-2 vertex adjacent to V_0 must correspond to V. In order to understand the weight-3 vertex, then, we must consider $V \otimes V$. Recall that $V \otimes V \cong S^2V \oplus \Lambda^2 V$. Now, $\Lambda^2 V$ consists of the top forms of V, since $\dim V = 2$. Any element g of A_4^* acts by the matrix it represents on V, so by Proposition 1.3.9, g acts by scalar multiplication by $\det g = 1$ on $\Lambda^2 V$. This implies that $\Lambda^2 V$ is the trivial representation and hence that $V \otimes V \cong S^2 V \oplus V_0$. Now, the weight-3 vertex in $\widetilde{E_6}$ is connected to V, so there must be some 3-dimensional irreducible subrepresentation of $V \otimes V$, hence of $S^2 V$ (notice that for any subrepresentation W of $S^2 V$, $W \oplus V_0$ is not irreducible). But $S^2 V$ has dimension 3, so this irreducible subrepresentation must be all of $S^2 V$. So, we can now label 2 more vertices in our graph:

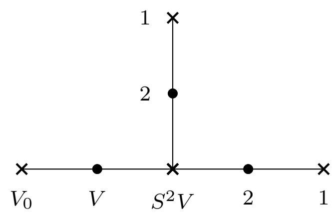

Let V_1 and V_2 be the nontrivial 1-dimensional representations of A_4^*. Now, tensoring by a 1-dimensional representation permutes irreducible representations but preserves their dimension. Since V is one of the irreducible representations of dimension 2, we see that $V \otimes V_1$ and $V \otimes V_2$ must be the other 2. Notice that none of these 2-dimensional irreducible representations can be isomorphic. For instance, suppose $V \otimes V_2 \cong V \otimes V_1$. Then, the group structure on 1-dimensional representations implies that there exists some V_i such that $V_2 \otimes V_i \cong V_0$, and moreover, $V_1 \otimes V_i \not\cong V_0$. Tensoring by V_i, we get

$$V \cong V \otimes V_2 \otimes V_i \cong V \otimes V_1 \otimes V_i.$$

So, tensoring V by $V_1 \otimes V_i$ gives us back V, which by comparing characters implies that $V_1 \otimes V_i \cong V_0$, a contradiction. So, $V \otimes V_2 \not\cong V \otimes V_1$.

The only remaining question is which weight-2 vertex in the graph corresponds to $V \otimes V_1$ and which to $V \otimes V_2$. For any $i \in \{1, 2\}$, we have

$$V \otimes (V \otimes V_i) \cong (V \otimes V_0) \otimes V_i \cong (S^2 V \oplus V_0) \otimes V_i \cong (S^2 V \otimes V_i) \oplus V_i.$$

This implies that $V \otimes V_i$ is connected to V_i. So, we can now label every vertex of $\widetilde{E_6}$ with the representation corresponding to it:

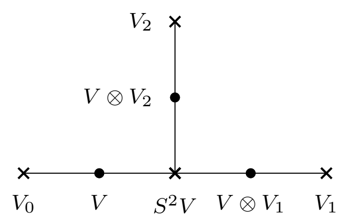

Using the fact that the group of 1-dimensional representations here has order 3 and so is isomorphic to C_3, one can check that tensoring by any 1-dimensional representation here corresponds to permuting the three branches of the graph stemming from $S^2 V$. Indeed, a generalization of our argument that $V \otimes V_1$ is connected to V_1 implies that tensoring any McKay graph by any 1-dimensional representation constitutes an automorphism of the graph, which gives us a nice, visual interpretation of this tensor operation.

Example. Consider $\widetilde{E_7}$, which corresponds to S_4^*. We use much the same arguments as in the above example. First, we know that the trivial representation V_0 is adjacent to V in $\widetilde{E_7}$. Using our bipartition of $\widetilde{E_7}$ and noting that $-I$ acts by I on V_0, our graph of $\widetilde{E_7}$ looks like

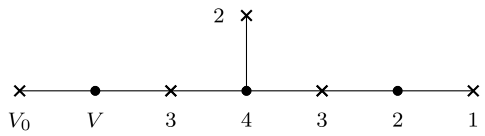

where an "x" signifies that the representation descends to $S_4^*/\{\pm I\} \cong S_4$. Notice that, because the nontrivial 1-dimensional representation, V_1, descends to a nontrivial 1-dimensional representation of S_4, it must descend to the sign representation on S_4. By much the same arguments as in the above example, S^2V is irreducible and adjacent to V, $V \otimes V_1$ is adjacent to V_1, and $S^2V \otimes V_1$ is adjacent to $V \otimes V_1$. So, we have:

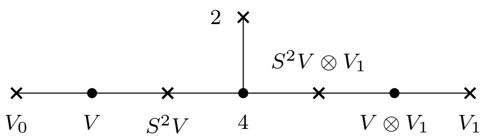

Finally, one can check that the 4-dimensional irreducible representation is S^3V, while the remaining unlabelled 2-dimensional representation corresponds to the unique 2-dimensional irreducible representation of S_3 via the chain $S_4^* \to S_4 \to S_3$.

4.3 Representations of the Symmetric Group

We begin by proving a result about hom sets. Though we only need this result in a very specific, representation-theoretic setting, we may as well prove it in the general setting.

Proposition 4.3.1. *Let R be a ring with idempotents e_1 and e_2. Then,*

$$\mathrm{Hom}_R(Re_1, Re_2) \cong e_1Re_2$$

as abelian groups.

Proof. We can define a map $\alpha : \mathrm{Hom}_R(Re_1, Re_2) \to e_1Re_2$ as follows. For any $f \in \mathrm{Hom}_R(Re_1, Re_2)$, because f is a homomorphism of R-modules, it is determined by where it sends e_1. Suppose $f(e_1) = a_f$. Then, because e_1 is an element of R and f respects the action of R, we have

$$a_f = f(e_1) = f(e_1^2) = e_1 f(e_1) = e_1 a_f.$$

So, $a_f \in e_1Re_2$. We define $\alpha(f) = a_f$. One can check that α defines a homomorphism of R-modules. Conversely, we define a map $\beta : e_1Re_2 \to \mathrm{Hom}_R(Re_1, Re_2)$ by sending $b \in e_1Re_2$ to the homomorphism $f : Re_1 \to Re_2$ defined by $f(re_1) = rb$. By much the same computation as for a_f above, one checks that f is in fact a homomorphism of R-modules. Finally, one checks that β is a homomorphism and, moreover, that α and β are inverses. $\square$

Now, let H be a subgroup of a finite group G, and let $\underline{\mathbb{C}}$ denote the trivial representation of H. Then, by an example at the beginning of Section 3.4, we have
$$\mathrm{Ind}_H^G(\underline{\mathbb{C}}) \cong \mathbb{C}[G/H].$$
Define $e_H = \frac{1}{|H|}\sum_{h \in H} h$. Then, e_H is an idempotent in $\mathbb{C}[H] \subset \mathbb{C}[G]$. Fixing representatives $g_1, \ldots, g_m$ for the cosets of H, one can check that the following map is an isomorphism:
$$\begin{aligned} \mathbb{C}[G/H] &\xrightarrow{\sim} \mathbb{C}[G]e_H. \\ g_iH &\mapsto g_ie_H \end{aligned}$$
With a little more work, this isomorphism and the above hom set result yield the following proposition.

Proposition 4.3.2. *Let H_1 and H_2 be two subgroups of a finite group G, and let $\underline{\mathbb{C}}_i$ be the trivial representation on H_i. For any i, define $e_{H_i} = \frac{1}{|H_i|}\sum_{h \in H_i} h$. Then,*
$$\mathrm{Hom}_R(\mathrm{Ind}_{H_1}^G(\underline{\mathbb{C}}_1), \mathrm{Ind}_{H_2}^G(\underline{\mathbb{C}}_2) \cong e_{H_1}\mathbb{C}[G]e_{H_2}.$$
Moreover, $e_{H_1}\mathbb{C}[G]e_{H_2}$ has a basis of double coset representatives, i.e. a basis of the form $\{e_{H_1}g_ie_{H_2}\}_{i=1}^m$, where the g_i and g_j are in the same coset of either H_1 or H_2 if and only if $i = j$.

Proof. By the above, we have $\mathrm{Ind}_{H_i}^G(\underline{\mathbb{C}}_i) \cong \mathbb{C}[G]e_{H_i}$ for any i. So, the desired isomorphism follows directly from Proposition 4.3.1. Notice that double coset representatives are precisely representatives of the orbits of the action of $H_1 \times H_2$ on G defined by $(h_1, h_2)g = h_1gh_2^{-1}$. We can always partition the elements of a a set into orbits, so there exists some $g_1, \ldots, g_m$ such that $G = \sqcup_{i=1}^m H_1g_iH_2$. Now, a basis for $e_{H_1}\mathbb{C}[G]e_{H_2}$ is given by $\{e_{H_1}ge_{H_2}\}_{g \in G}$. On the other hand, for any $g \in G$, we can write $g = h_1g_ih_2$ for some i. Then, we have $e_{H_1}h_1 = e_{H_1}$, since multiplication by h_1 permutes the elements of H_1. Likewise, $h_2e_{H_2} = e_{H_2}$. Putting these together, we see that
$$e_{H_1}ge_{H_2} = e_{H_1}h_1gh_2e_{H_2} = e_{H_1}g_ie_{H_2}.$$
This proves that the set $\{e_{H_1}g_ie_{H_2}\}_{i=1}^m$ spans $e_{H_1}\mathbb{C}[G]e_{H_2}$. On the other hand, all of the elements in this set are distinct, because they are elements of distinct orbits of a group action. This implies that the set is linearly independent, so that it forms a basis for $e_{H_1}\mathbb{C}[G]e_{H_2}$. □

Now, we turn to the representation theory of the symmetric group S_n. Given any element $\sigma \in S_n$, we can form a tuple $\lambda = (\lambda_1, \ldots, \lambda_k)$, where λ_i is the number of i-cycles in σ. Because conjugacy classes of S_n are classified by cycle type, two elements of S_n are conjugate if and only if they give rise to the same λ in this way. Notice that $\sum_i \lambda_i = n$, so that λ in fact constitutes a (non-negative) partition of n. We can represent such partitions succinctly via diagrams, as the following example demonstrates.

Example. Consider the case where $n = 4$. Take, for instance, the partition $(3, 1)$. We can represent this partition diagrammatically as:

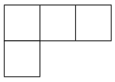

Here, the first row has 3 squares, which corresponds to the fact that the first number in the partition is 3, and the second row has 1 square, which corresponds to the fact that the second number in the partition is 1. We call such a diagram a *Young diagram.* Note that by convention, we always draw the rows of Young diagrams by starting with the longest row at the top and then going in descending order of length as one goes down the rows. With this convention, we can write out the Young diagram corresponding to all the rest of the partitions of 4.

(4) : 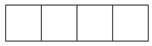

$(2, 2)$: 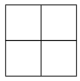

$(2, 1, 1)$: 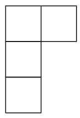

$(1, 1, 1, 1)$: 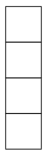

We now discuss a couple of notational aspects of partitions that will help us talk about them. First, if λ is a partition of n, we write $\lambda \vdash n$. We will also write $\lambda_1, \ldots, \lambda_k$ to denote each term of a partition λ, so that $\lambda = (\lambda_1, \ldots, \lambda_k)$. We will generally write partitions in descending order, e.g. $(3, 1)$ rather than $(1, 3)$. The only exception is when we are explicitly considering their relation to cycle types of elements of S_n, in which case the order of the partition elements matters. (As we will see, however, which cycle length occurs in which frequency is rarely a relevant consideration for our purposes.) Moreover, when elements in a partition are repeated, we sometimes denote this with exponentiation for short. For instance, we write $(2, 1, 1) = (2, 1^2)$ and $(1, 1, 1, 1) = (1^4)$.

Now, given some partition λ of n, we associate λ to the subgroup

$$S_\lambda = S_{\lambda_1} \times S_{\lambda_2} \times \cdots \times S_{\lambda_k}$$

of S_n. Let $\underline{\mathbb{C}}$ denote the trivial representation of S_λ. Then, we associate to λ a representation

$$M_\lambda = \mathrm{Ind}_{S_\lambda}^{S_n}(\underline{\mathbb{C}}) \cong \mathbb{C}[S_n/S_\lambda] \cong \mathbb{C}[S_n]e_\lambda,$$

where, much as above, we define $e_\lambda = e_{M_\lambda} = \frac{1}{|S_\lambda|}\sum_{\sigma \in S_\lambda} \sigma$. Notice that $|S_\lambda| = \lambda_1! \cdots \lambda_k!$, so that

$$\dim M_\lambda = \frac{n!}{\lambda_1! \cdots \lambda_k!}.$$

Example. If $\lambda = (n)$, then $S_\lambda = S_n$, and $M_\lambda = \mathbb{C}[S_n]e_{S_n}$ has dimension $\frac{n!}{n!} = 1$. Notice that every element of S_n fixes e_{S_n}, so that $M_{(n)}$ is the trivial representation of S_n.

If instead $\lambda = (1)^n$, then $S_\lambda = S_1 \times \cdots \times S_1$ is trivial, so that $e_\lambda = 1$. This implies that $M_{(1^n)} \cong \mathbb{C}[S_n] \cdot 1 = \mathbb{C}[S_n]$ is the regular representation of S_n.

There is another visualization of elements of S_n that will be useful to us. We can view a permutation σ of S_n as a braid, i.e. an intertwining of several strands, where we number the strands from 1 to n and their relative order after the intertwining defines the permutation. For instance, a braid corresponding to $(1, 2, 3) \in S_4$ would look like:

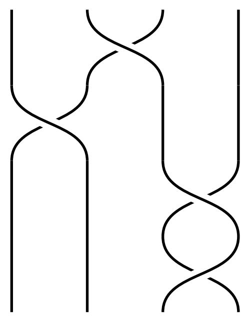

Notice that the 1st strand (from the left) ends in the 2nd position, the 2nd in the 3rd position, and the 3rd in the 1st position, and the 4th in the 4th position. This corresponds to the fact that $(1, 2, 3)$ sends 1 to 2, 2 to 3, 3 to 1, and fixes 4. Now, this braid is not the only one representing $(1, 2, 3)$: for instance, because the two crossings of the 3rd and 4th strand do not affect the final order of the strands, we may remove these crossing to get another diagram for $(1, 2, 3)$:

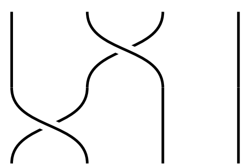

This process of removing extraneous crossings, if repeated, can reduce any braid diagram to one with a minimal number of crossings. This minimal diagram will be characterized by the fact that no two strands cross more than once. One can check that such a minimal diagram is unique (if you are unconvinced, play around with a few simple examples).

With the help of these braid diagrams, the following definition becomes more visual and intuitive.

Definition 4.3.3. Let $\sigma \in S_n$. We define the *length* of σ, $\ell(\sigma)$, to be the number of inversions in σ, i.e. the number of pairs (i, j) such that $i < j$ but $\sigma(i) > \sigma(j)$.

Notice that the length of any permutation σ is the number of crossings in a minimal braid diagram for σ. Since a unique such minimal diagram exists, we conclude that the length of a permutation is well-defined.

Exercise. Show that, for any σ and τ in S_n,

$$\ell(\sigma\tau) \leq \ell(\sigma) + \ell(\tau).$$

Now, given two partitions λ and μ of n, our above results suggest that we will be interested in giving braid diagrams to represent elements of $M_\lambda = \mathbb{C}[S_n]e_\lambda$ or elements of $e_\mu\mathbb{C}[S_n]e_\lambda$. To this end, we first note that

$$e_\lambda = \frac{1}{|S_\lambda|}\sum_{\sigma \in S_\lambda} \sigma = \frac{1}{|S_{\lambda_1}| \cdots |S_{\lambda_k}|} \sum_{(\sigma_1, \ldots, \sigma_k) \in S_{\lambda_1} \times \cdots \times S_{\lambda_k}} (\sigma_1, \ldots, \sigma_k) = \prod_{i=1}^k \frac{1}{|S_{\lambda_i}|} \sum_{\sigma \in S_{\lambda_i}} \sigma = \prod_{i=1}^k e_{(\lambda_i)}.$$

In words, we can understand e_λ, which is an "average" over all elements of S_λ, as actually being a product of averages over each of the S_{λ_i}. Now, we think of S_{λ_i} as permuting the subset $\{\lambda_{i-1} + 1, \lambda_{i-1} + 2, \ldots, \lambda_{i-1} + \lambda_i\} \subseteq \{1, \ldots, n\}$. In terms of braid diagrams, then, the $(\lambda_{i-1} + 1)$st through $(\lambda_{i-1} + \lambda_i)$th strands are the ones whose crossings denote elements of S_{λ_i}. Now, we can't easily draw a diagram corresponding to e_{λ_i}, since it is an average of elements of S_{λ_i}, so instead we will represent this element with a box over the strands corresponding to S_{λ_i} in our braid diagrams (the idea being that the box signifies taking equal parts of all possible braids we could put in its place). By the above equation, we will need one box for each λ_i, since we can average over each of these separately to get all of e_λ.

As an example of this setup, take the partition $\lambda = (2, 2, 1)$ of 5 and the element $\sigma = (1, 2, 3)(4, 5)$ of S_5. Then, one braid diagram representation of σe_λ is

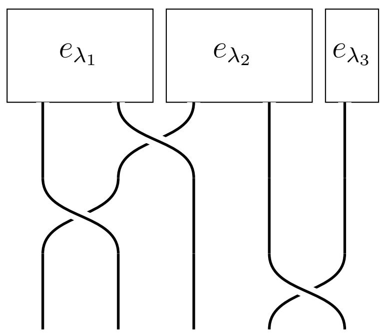

Notice that the braid diagram is read from top to bottom: since e_λ is to the right of σ, we do that permutation (or rather, "average of permutations") first, which is why the boxes representing the e_{λ_i} are at the top of the diagram. Likewise, if

$\mu = (3,1,1)$, then the braid diagram for $e_\mu \sigma e_\lambda$ is

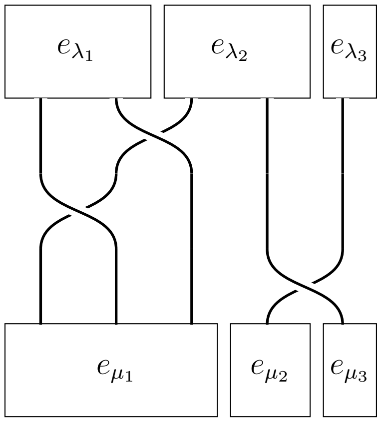

In addition to removing the crossings of strands that don't affect the endpoints of the strands, we have another form of simplification that we can do to diagrams with boxes in them. If we have a crossing of two strands which either come from or subsequently go to the same box, then we can absorb the crossing into the box. For instance, take the crossing between strands 1 and 2 in the above diagram. This corresponds to the element $(1,2) \in S_{\mu_1}$. Since multiplication by such an element permutes the elements of S_{μ_1}, one sees that $e_{\mu_1}(1,2) = e_{\mu_1}$, so that we can remove the crossing of strands 1 and 2 without affecting the element in question. The resulting diagram is then:

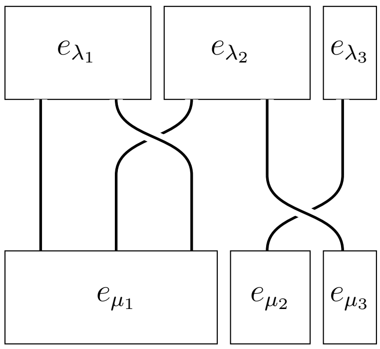

We could repeat this process with the crossing of strands 2 and 3 in this diagram to reduce the diagrm further. However, we can not uncross strands 4 and 5, since these come from and go to different boxes. Notice that each of these simplifications preserves the double coset of σ: for instance, removing the crossing $(1,2)$ from our example σ above corresponds to multiplying by $(1,2)^{-1} \in S_{\mu_1} \subset S_\mu$, and $S_\mu \sigma S_\lambda = S_\mu((1,2)^{-1}\sigma)S_\lambda$. Moreover, one can see from the braid diagrams that $\ell((1,2)^{-1}\sigma) \leq \ell(\sigma)$, so that we can repeatedly perform these simplifications to reduce to a diagram of minimal crossings (i.e. an element in the same double coset as σ with minimal length). This immediately implies the following proposition.

Proposition 4.3.4. *For any partitions λ and μ of n, each double coset of S_λ and S_μ as subgroups of S_n has a unique representative τ of minimal length, which is characterized by the following property: for any 2 strands that enter the same box* *at any point in the braid diagram of $e_\mu \tau e_\lambda$, the strands must be uncrossed on the side of the box at which they enter.*

Now, notice that from the outset, we have defined M_λ using the trivial representation. We might wonder whether we can use the other 1-dimensional representation of S_n, namely, the sign representation. It turns out that this scenario also plays a vital role in the representation theory of the symmetric group. To this end, given any $\lambda \vdash n$, we define

$$e_\lambda^- = \frac{1}{|S_\lambda|} \sum_{\sigma \in S_\lambda} \mathrm{sgn}(\sigma)\sigma,$$

and

$$M_\lambda^- = \mathrm{Ind}_{S_\lambda}^{S_n}(\mathbb{C}_s),$$

where $\mathbb{C}_s$ is the sign representation of S_n. By much the same arguments as for M_λ above, one can show that

$$M_\lambda^- \cong \mathbb{C}[S_n]e_\lambda^-.$$

We will draw box diagrams for elements of the form σe_λ^- just as we draw them for σe_λ above, with a box denoting the "sign-weighted averaging" of e_λ^-. However, our ability to simplify crossings by absorbing them into boxes relies on the fact that e_λ is fixed under multiplication by the 2-cycle representing the crossing. In this case, multiplying e_λ^- by a 2-cycle swaps each term with a term of opposite sign, so we get $(i,j)e_\lambda^- = -e_\lambda^-$ for any 2-cycle $(i,j) \in S_n$. This implies that when we absorb crossings into boxes for e_λ^-, the diagram we get is actually the negative of the element the diagram represents in the group algebra. So long as one keeps track of these negatives, however, this type of simplification still allows us to reduce to diagrams with minimal crossings.

We are nearly ready to use all of this machinery to classify the irreducible representations of S_n. First, however, we need just a couple more definitions relating to partitions.

Definition 4.3.5. Let λ be a partition of n. We define the *dual partition* $\lambda^* \vdash n$ by setting λ_i^* to be the number of choices of j for which $\lambda_j \geq i$. Equivalently, λ^* corresponds to the Young diagram obtained from the Young diagram of λ by reflecting along the diagonal from the top left corner of the diagram (so that, e.g., the left column becomes the top row).

We define a (total) order on the set of partitions of n as follows: for any $\lambda, \mu \vdash n$, we say that $\mu > \lambda$ if there exists $k \geq 1$ such that $\mu_i = \lambda_i$ for all $1 \leq i < k$ and $\mu_k > \lambda_k$. (In other words, we use the lexicographical ordering.)

The following theorem is in some ways the crux of our classification of the irreducible representations of S_n.

Theorem 4.3.6.

1. *For any partitions λ and μ of n such that $\mu > \lambda^*$,* $\mathrm{Hom}_{\mathbb{C}[S_n]}(M_\lambda, M_\mu^-) = 0$.
2. *For any partition λ of n,* $\mathrm{Hom}_{\mathbb{C}[S_n]}(M_\lambda, M_{\lambda^*}^-) \cong \mathbb{C}$.

Proof. First, let λ and μ be any partitions of n. By Proposition 4.3.2, we have

$$\mathrm{Hom}_{\mathbb{C}[S_n]}(M_\lambda, M_\mu^-) \cong e_\lambda \mathbb{C}[S_n] e_\mu^-.$$

Consider the braid diagram corresponding to $e_\lambda \sigma e_\mu^-$ for any $\sigma \in S_n$. Suppose there are two strands in this diagram which both come from S_{λ_i} for some i and both go to S_{μ_j} for some j. Then, by absorbing any crossings of these two strands into the box for e_{λ_i}, we may assume that the two strands do not cross each other. Now, we can "pull a crossing out" from the box for e_{μ_j} to cross the two strands and pick up a negative sign. But then, we can absorb this cross back into the box for e_{λ_i} to get the original diagram we started with, but now with a minus sign. Using diagrams:

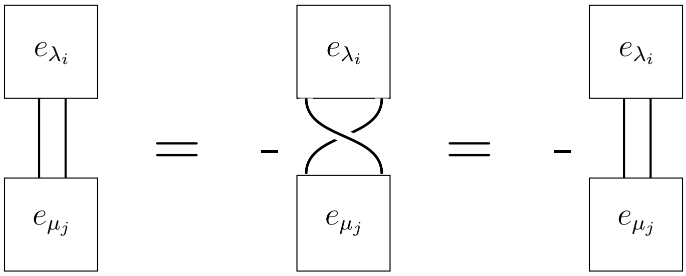

The element this diagram represents in the group algebra is thus equal to its own additive inverse, so it must be 0. So, in order for our element $e_\lambda \sigma e_\mu^-$ to be nonzero, the braid diagram of this element must send every pair of strands coming from the same box (i.e. from the same S_{μ_i}) to different boxes (i.e. to different S_{λ_j}'s).

Now, suppose that $\mu > \lambda^*$. In particular, this implies that $\mu_i \geq \lambda_i^*$ for all i. For any σ in S_n, we have μ_1 strands coming from S_{μ_1} in the braid diagram of $e_\lambda \sigma e_\mu^-$ and λ_1^* distinct S_{λ_j}'s. if $\mu_1 > \lambda_1^*$, then by the pidgeonhole principle, two strands from S_{μ_1} must go to the same S_{λ_j} for some j, so that $e_\lambda \sigma e_\mu^- = 0$ by our above argument. Otherwise, we have $\mu_1 = \lambda_1^*$, so that each S_{λ_j} has precisely 1 strand coming from S_{μ_1}. We can then repeat this argument for S_{μ_i} in place of S_{μ_1} for all i. For each i, if we have not already shown that $e_\lambda \sigma e_\mu^- = 0$, then every S_{λ_j} with at least $i-1$ strands has 1 strand coming from each of the S_{μ_k} for $1 \leq k < i$, for a total of $i-1$ strands accounted for in each such S_{λ_j}. Then, we need to put each of the μ_i strands from S_{μ_i} into all of the S_{λ_j} which have at least i strands (so that they have room for at least one more strand to in them). But there are precisely λ_i^* values of j for which S_{λ_j} has at least i strands. So, either $\mu_i > \lambda_i^*$, in which case $e_\lambda \sigma e_\mu^- = 0$ as above, or $\mu_i = \lambda_i^*$, in which case each of the λ_i^* available S_{λ_j}'s gets precisely 1 strand from S_{μ_i}. However, because $\mu > \lambda^*$, we know that we will get $\mu_i > \lambda_i^*$ at some point in this process, so at this step we will conclude that $e_\lambda \sigma e_\mu^- = 0$. Since this holds for all σ in S_n, we conclude that

$$\mathrm{Hom}_{\mathbb{C}[S_n]}(M_\lambda, M_\mu^-) \cong e_\lambda \mathbb{C}[S_n] e_\mu^- = 0.$$

If now instead we set $\mu = \lambda^*$, going through the same process as above, we will get $\mu_i = \lambda_i^*$ for all i, so that it is possible to make no 2 strands come from and go to the same box. This implies that there is some nonzero element of $e_\lambda \mathbb{C}[S_n] e_{\lambda^*}^-$. However, one can check that, up to uncrossing two strands that come from or go to the same box (i.e. up to at most a minus sign), there is only one way to make no 2 strands come from and go to the same box. This implies that

$$e_\lambda \mathbb{C}[S_n] e^-_{\lambda^*} \cong \mathbb{C} \cdot (e_\lambda \sigma e^-_{\lambda^*}) \cong \mathbb{C}$$

for some $\sigma \in S_n$, which gives us $\mathrm{Hom}_{\mathbb{C}[S_n]}(M_\lambda, M^-_{\lambda^*}) \cong \mathbb{C}$. □

Now, recall that, for any partition λ of n, M_λ and $M^-_{\lambda^*}$ are representations of S_n. So, by the proof of Proposition 2.4.9, the fact that $\mathrm{Hom}_{\mathbb{C}[S_n]}(M_\lambda, M^-_{\lambda^*}) \cong \mathbb{C}$ implies that M_λ and M_{λ^*} share precisely one irreducible subrepresentation. We define L_λ to be this irreducible representation. The following proposition and theorem show that these L_λ in fact give us all of the irreducible representations of S_n.

Proposition 4.3.7. *Let $\lambda \neq \mu$ be two distinct partitions of n. Then, $L_\lambda \ncong L_\mu$.*

Proof. Suppose that $L_\lambda \cong L_\mu$. Since the lexicographical order is total, either $\mu^* > \lambda^*$ or $\lambda^* > \mu^*$ (we cannot have equality, since that would imply that $\lambda = \mu$); without loss of generality, suppose $\mu^* > \lambda^*$. Then, by Theorem 4.3.6, we have $\mathrm{Hom}(M_\lambda, M^-_{\mu^*}) = 0$. On the other hand, $L_\lambda \subset M_\lambda$ by definition, and $L_\lambda \cong L_\mu \subseteq M^-_{\mu^*}$ by assumption. So, M_λ and $M^-_{\mu^*}$ share an irreducible representation, which means that there is at least one nontrivial intertwiner $M_\lambda \to M^-_{\mu^*}$ (namely, the one that maps $L_\lambda \to L_\mu$ via an isomorphism and sends all the other irreducible subrepresentations of M_λ to 0). This contradicts the fact that $\mathrm{Hom}(M_\lambda, M^-_{\mu^*}) = 0$. □

Theorem 4.3.8. *The set $\{L_\lambda\}_{\lambda \vdash n}$ is precisely the set of irreducible representations of S_n.*

Proof. Recall that the conjugacy classes of S_n are in bijection with the partitions of n, with the bijection sending a conjugacy class of S_n to the cycle type of its elements. By the above proposition, the set $I = \{L_\lambda\}_{\lambda \vdash n}$ contains a unique irreducible representation of S_n for each partition λ of n, which means that I is in bijection with the set of conjugacy classes of S_n. But the number of conjugacy classes of S_n is precisely the number of irreducible representations of S_n by Theorem 2.1.4, so the I must contain precisely as many irreducible representations as S_n has. □

This theorem is actually pretty remarkable. We've constructed a canonical bijection between conjugacy classes of S_n and irreducible representations of S_n. Although the cardinalities of these two sets are equal for any group, it is rare that such a bijection exists. Even in the very simple case of C_n, where we can easily understand all the irreducible representations, we do not have a canonical bijection (there is an easy bijection to consider, but it requires a choice of a generator of C_n and/or a primitive nth root of unity). What's more, our arguments here can even be generalized to classify representations of S_n over any field of characteristic 0.

With this, we have completely classified the irreducible representations of the symmetric group. However, we can strill improve our classification a little bit by understanding what the L_λ look like more explicitly. In a couple of simple cases, we can even compute these irreducible representations directly.

Proposition 4.3.9. *$L_{(n)}$ is isomorphic to the trivial representation of S_n, and $L_{(1^n)}$ is isomorphic to the sign representation.*

Proof. Let V_0 and V_1 denote the trivial and sign representations of S_n, respectively. Then, as discussed in an above example, $M_{(n)} \cong V_0$ and $M_{(1^n)} \cong \mathbb{C}[S_n]$. Likewise, one can check that $M^-_{(n)} \cong V_1$ and $M^-_{(1^n)} \cong V_0$. Noting that (n) and (1^n) are dual partitions, we see that $M_{(n)}$ and $M^-_{(1^n)}$ have a copy of the trivial representation in common, so that $L_{(n)} \cong V_0$, and $M_{(1^n)}$ and $M^-_{(n)}$ have a copy of the sign representation in common, so that $L_{(1^n)} \cong V_1$. □

The following proposition, which we state without proof, can also be useful for computing the L_λ.

Proposition 4.3.10. *λ be any partition of n, and let V_1 denote the sign representation of S_n. Then,*

$$L_\lambda \otimes V_1 \cong L_{\lambda^*}.$$

Example. Now, we can compute everything when n is small. Pick $n = 3$, and let V_0, V_1, and V_2 denote the trivial, sign, and irreducible 2-dimensional representations of S_3, respectively. Then, $M_{(3)} \cong V_0$ and $M^-_{(1^3)} \cong V_0 \oplus V_1 \oplus V_2^2$, so that $L_{(3)} \cong V_0$. Likewise, $M_{(2,1)} \cong V_0 \oplus V_2$ and $M^-_{(2,1)} \cong V_1 \oplus V_2$, so that $L_{(2,1)} \cong V_2$. Finally, $M_{(1^3)} \cong \mathbb{C}[S_3]$, and $M^-_{(3)} \cong V_1$, so that $L_{(1)^3} \cong V_1$. So, we do in fact recover all 3 irreducible representations of S_3. In similar fashion, one can compute the irreducible representations of S_4 with these methods.

We end this section by deriving explicit expression for L_λ.

Proposition 4.3.11. *Let λ be a partition of n, and pick σ such that $e_\lambda \sigma e^-_{\lambda^*}$ generates $e_\lambda \mathbb{C}[S_n] e^-_{\lambda^*}$. (Note that this module is isomorphic to $\mathbb{C}$ by Theorem 4.3.6 and so is generated by 1 element.) Then,*

$$L_\lambda \cong \mathbb{C}[S_n] e_\lambda \sigma e^-_{\lambda^*}.$$

Proof. By the proof of Proposition 4.3.1, the element $e_\lambda \sigma e^-_{\lambda^*} \in e_\lambda \mathbb{C}[S_n] e^-_{\lambda^*}$ corresponds, via the isomorphism $\operatorname{Hom}_{\mathbb{C}[S_n]}(M_\lambda, M^-_{\lambda^*}) \cong e_\lambda \mathbb{C}[S_n] e^-_{\lambda^*}$, to the intertwiner

$$\begin{aligned} \varphi_\lambda : M_\lambda &\to M^-_{\lambda^*} \\ \tau e_\lambda &\mapsto \tau e_\lambda \sigma e^-_{\lambda^*} \end{aligned}$$

This intertwiner must restrict to an intertwiner $L_\lambda \to L_\lambda$ and send the rest of M_λ to 0, so that $\operatorname{Im} \varphi_\lambda = L_\lambda$. On the other hand,

$$\operatorname{Im} \varphi_\lambda = M_\lambda \sigma e^-_{\lambda^*} = \mathbb{C}[S_n] e_\lambda \sigma e^-_{\lambda^*}.$$

□

Remark. In the notation of the above proposition, notice that

$$\mathbb{C}[S_n] e_\lambda \sigma e^-_{\lambda^*} \cong \mathbb{C}[S_n] \sigma^{-1} e_\lambda \sigma e^-_{\lambda^*}.$$

The element $Y = \sigma^{-1} e_\lambda \sigma e^-_{\lambda^*}$ is sometimes called the *Young idempotent*. (Note, however, that this is element is not an precisely an idempotent; rather, there exists some rational number q such that qY is idempotent.)