An Introduction to the Representation Theory of Finite Groups

(Mathematical Subjects Classification 20-01)

Nathaniel S. Kuhn

2274 Amherst St.

Palo Alto, CA 94306

In this paper I will outline the theory of representations of finite groups, at a level which I believe is appropriate to the end of a full-year introductory course in abstract algebra or the beginning of a second-year course.

Why another exposition of this material, especially given the characteristically beautiful book of [Serre]? The reason is this: the introductory material in Serre's book was written for quantum chemists, and therefore minimized the use of abstract linear algebra. As with most of the standard treatments, the approach taken can seem computational and unmotivated, especially in the treatment of the orthogonality relations. Here we give an interpretation of the orthogonality relations as counting up G-isomorphisms—an interpretation which I first saw in a paper of [Bott] in the Lie-group context, and which deserves a wider audience. Readers already familiar with the theory of group representations may wish to skip directly to section 8, referring back as necessary to material in earlier sections (especially the material on Hom(V,W) in propositions 4.3 and 7.2, and the material on averaging in section 3). Finally, section 10 presents an introduction to the structure theory of semisimple rings as a series of exercises with detailed hints, which places the material on group representations in a more general context.

0. Prerequisites

Most of the prerequisites for this material are covered in a year's course in abstract algebra. The heaviest demands are in linear algebra, where Jordan canonical form is referred to several times in passing. Of course the basics of group theory are assumed, as are the elementary notions of field extensions. From ring theory, we assume the definitions of terms such as algebra, division ring and division algebra, zero divisor, and a few others. In the final section the reader will need to have a firm grasp on the basics of ideals and modules. A familiarity with Zorn's lemma arguments will help there, too.

The most important prerequisite which is generally not covered in a first year algebra course is the notion of a tensor product. Usually, these spaces are constructed by a wildly inifinite process using quotients of free modules; since we need them only in the context of finite dimensional vector spaces, we offer the following sketch of a development of the subject:

Definition 0.1: Suppose that V and W are finite dimensional vector spaces (over a field k). A **tensor product** of V and W is a vector space $V\otimes W$ with a bilinear map $\otimes: V\times W\to T$ which is *universal* for bilinear maps from $V\times W$ to any other vector space, in the following sense: if $\mu: V\times W\to U$ is a bilinear map, there is a *unique* linear map $\lambda: V\otimes W\to U$ such that $\mu(v,w) = \lambda(\otimes(v,w))$, i.e. $\mu = \lambda\circ\otimes$.

The tensor product gives us a tool for reducing questions about bilinear maps to questions about linear maps. Following standard notation, we write $v\otimes w$ instead of $\otimes(v,w)$. Rewriting the above definition, $\mu(v,w) = \lambda(v\otimes w)$. Another way of looking at this is that $\otimes$ allows us to "factor out" the bilinear part of μ. The space $V\otimes W$ and the map $\otimes$ are "big enough" to factor all bilinear maps, but "small enough" so that this factorization is unique.

Proposition 0.2: Tensor products are unique: If $V \otimes W$ and $\otimes$ are a tensor product of V and W in the sense above, and $V \otimes' W$ and $\otimes'$ are another tensor product, then there is a (unique) isomorphism $\phi: V \otimes W \to V \otimes' W$ so that $\otimes' = \phi \circ \otimes$.

Proof: (Anyone who is familiar with the standard argument for the uniqueness of a universal object could do this as an exercise. Anyone who is not should make lots of drawings with little arrows to follow this argument). First, since $\otimes'$ is bilinear and $\otimes$ has the universal property, we see that ϕ exists with $\otimes' = \phi \circ \otimes$. Symmetrically, since $\otimes$ is bilinear and $\otimes'$ has the universal property, there is a ϕ' such that $\otimes = \phi' \circ \otimes'$. It follows that $\otimes = \phi' \circ \otimes' = \phi' \circ (\phi \circ \otimes) = (\phi' \circ \phi) \circ \otimes$. Now by the universal property of $V \otimes W$ and $\otimes$, applied to the bilinear map $\otimes$ itself, there's only one map from $V \otimes W$ to itself such that $\otimes = \nu \circ \otimes$. But the identity map on $V \otimes W$, $1_{V \otimes W}$, and the map $\phi' \circ \phi$ both have this property, hence $\phi' \circ \phi = 1_{V \otimes W}$. By an exactly symmetrical argument, $\phi \circ \phi' = 1_{V \otimes' W}$. Thus ϕ is an isomorphism. By the universal property, it is unique. •

Since a tensor product of V and W is unique up to a unique isomorphism, we are justified in calling $V \otimes W$ "the" tensor product of V and W as long as it exists. The next proposition shows that we can always construct $V \otimes W$ when V and W are finite dimensional (in fact, finite dimensionality is not necessary, but the proof is more technical without it). The essential point is that a bilinear map is completely specified by where it takes pairs of basis elements, just as a linear map is specified by where it takes single basis elements.

Proposition 0.3: Let V and W be finite dimensional vector spaces. Then a tensor product $V \otimes W$ exists.

Proof: Let $\{v_1, \cdots, v_n\}$ be a basis for V, and $\{w_1, \cdots, w_m\}$ be a basis for W. Then we can take $V \otimes W$ to be the $n \times m$ dimensional vector space whose basis elements are the $n \times m$ formal symbols $v_i \otimes w_j$ for $1 \le i \le n$, $1 \le j \le m$. Define the bilinear map $\otimes$ by the following rule:

$$\left(\sum_{i=1}^{n} a_i v_i\right) \otimes \left(\sum_{j=1}^{m} b_j v_j\right) = \Sigma_{1\le i\le n,\ 1\le j\le m}\ a_i b_j\ v_i \otimes w_j.$$

It is a straightforward exercise to verify that this is bilinear. To prove that we have the universal mapping property, suppose that $\mu: V \times W \to U$ is bilinear. Then define $\lambda: V \otimes W \to U$ by $\lambda(v_i \otimes w_j) = \mu(v_i, w_j)$ on the basis elements, and extend it by linearity to the rest of $V \otimes W$. Clearly $\mu = \lambda \circ \otimes$ for pairs of basis elements, and it is not hard to show that they must be equal for any pair of elements. Similarly it is not hard to show that this is the only possible λ with the desired property. •

Exercise 0.4: Show that there is a natural isomorphism Φ between $V^* \otimes W$ and Hom(V,W). Hint: Show that there is a bilinear map from $V^* \times W$ to Hom(V,W) taking (ϕ, w) to the linear transformation $(v \mapsto \phi(v)w)$. Show that this map is surjective. Use dimension counting to show that it is an isomorphism. It may help in thinking about this to notice that if $\{v_1^*, \cdots, v_n^*\}$ is a basis for V^* and $\{w_1, \cdots, w_n\}$ is a basis for W, $\Phi(v_j^*, w_i)$ corresponds to the elementary matrix with a 1 in the (i,j) entry and zeros elsewhere.

Exercise 0.5: Show that the vector space $V^* \otimes V^*$ is naturally isomorphic to the vector space of all bilinear forms on V.

1. Representations, G-Maps, and Isomorphisms

Definition 1.1: Suppose that G is a finite group and V is a finite dimensional vector space over a field k. A **representation** of G on V is a linear left action of G on V. That is, we have a map $G \times V \to V$; the image of (g,v) under this map is denoted $g \cdot v$ or simply gv. To say that it is a left action means:

(1.1.i)	$e \cdot v = v$	(e the identity in G)
(1.1.ii)	$g(hv) = (gh)v$	$\forall\ g,h \in G,\ v \in V$

Linearity means:

(1.1.iii)	$g(v+w) = gv + gw$	$\forall\ g \in G,\ \forall\ v,w \in V$

(1.1.iv) $g(\alpha v) = \alpha gv$ $\forall\ g\in G,\ v\in V,\ \alpha\in k$

By abuse of language we will often refer to the vector space V as the representation, but we will always mean to include the action. We will also refer to representations as **G-modules**. Although it is not necessary to assume that representations are finite dimensional, we shall do so throughout as it makes some of the proofs simpler technically. When we wish to particularly stress the ground field k, we say that V is a representation of V **over k**.

Of course, it is possible to define representations using *right* actions, simply by writing the group elements on the right in (i)-(iv). While the theory is symmetrical, a subtle difference is introduced in changing (ii):

(1.1.ii') $(vg)h = v(gh)$.

In other words, (ii) says that in a *left* action the group element gh acts by applying h first, then g; (ii') says that in a *right* action it acts by applying g first, then h. While the latter might seem to make more sense *a priori* (since it allows one to read strings of operations in the usual left-to-right order), most mathematicians are so accustomed to writing functions on the left that we will stick to that convention. Thus we assume throughout that all actions are left actions unless otherwise specified.

A completely equivalent characterization of a representation involves homomorphisms of the group G. Given $g\in G$, we can define a linear map $T_V(g)$ from V to itself (an **endomorphism** of V) by $T_V(g)(v) = gv$; the linearity of $T_V(g)$ follows from (1.1.iii) and (1.1.iv). (We will often write just T(g) when the representation is clear from the context.) It is an easy exercise to show that $T_V(g)$ is invertible and in fact that $(T_V(g))^{-1} = T_V(g^{-1})$. Thus $T_V(g)$ is an element of the group of invertible endomorphisms of V, which we shall denote by GL(V). T_V is a map $G\rightarrow GL(V)$, and it follows from (ii) that this map is a group homomorphism. Conversely, it is a simple matter to reconstruct a linear action from such a homomorphism.

Our eventual goal is to understand all possible representations of a finite group G, up to a suitable notion of isomorphism.

Definition 1.2: Let V, W be representations of G. $\Phi:V\to W$ is a **G-map** if it is a linear map which preserves (commutes with) the G-action, i.e. $\Phi(gv) = g\Phi(v)$. A **G-isomorphism** is a G-map with an inverse which is also a G-map. Two representations are **isomorphic** if there is a G-isomorphism between them.

Exercise 1.3: If $\Phi:V\to W$ is a G-map which is invertible as a linear transformation, show that Φ^{-1} is a G-map, i.e. Φ is a G-isomorphism.

Exercise 1.4: If $\alpha\in k$, multiplication by α is a G-map of V to itself (i.e. a **G-endomorphism** of V). If $\alpha \neq 0$, it is a G-isomorphism.

2. Examples

Example 2.1: If G is any group and V any vector space, the **trivial representation** of G on V is given by gv = v. Show that this corresponds to the trivial homomorphism $G\to GL(V)$.

Example 2.2: Let G denote the cyclic group of order n, C_n, and let g_0 be a fixed generator of G. Let k be the field **C** of complex numbers, and let V be **C** considered as a one-dimensional vector space. We will find all possible representations of G on V, by finding all homomorphisms $G\to GL(V)$. A linear map of **C** to itself is simply multiplication by a complex scalar; it is easy to see that GL(V), the group of invertible self-maps of V must be isomorphic to **C***, the group of non-zero elements of **C** under multiplication. Since g_0 has order n, it must map to an nth root of unity. (Verify this.) From the n n-th roots of unity, we obtain n distinct representations by $T_k(g_0) = e^{2\pi ki/n}$, $k = 0,\dots,n-1$. Each T_k is extended to make it a homomorphism by taking $T_k(g_0{}^m) = (e^{2\pi ki/n})^m = e^{2\pi kmi/n}$.

Exercise 2.3: Show that no two of these n representations are isomorphic. Hint: show that if T and T' are isomorphic representations, T(g) and T'(g) are similar linear transformations and hence must have the same eigenvalues.

Example 2.4: The dihedral group of order 2n is given by generators a and b satisfying the relations $a^n = e$, $b^2 = e$, and $bab^{-1} = a^{-1}$. This is the group of symmetries of a regular n-gon, and as such it has a natural representation as rigid motions of the plane, with a representing a rotation by $2\pi/n$, and b representing a reflection along a line of symmetry through the origin. In fact, as above there is a family of representations on a two dimensional vector space (over **R** or **C**) given by the following matrices:

$$T_k(a^m) = \begin{bmatrix} \cos 2\pi km/n & -\sin 2\pi km/n \\ \sin 2\pi km/n & \cos 2\pi km/n \end{bmatrix} \qquad T_k(b) = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

Note that T_1 corresponds to the "natural" representation mentioned above. Also, while the action of this representation is closely linked to the previous example, T_1 cannot be thought of as a one-dimensional complex representation, since b would need to act by complex conjugation which is not **C**-linear.

Exercise 2.5: Verify that this is a representation by showing that these matrices satisfy the defining relations of the group. Show that the representations given by T_k and T_{n-k} are isomorphic.

Example 2.6: An arbitrary abelian group G is a direct product of cyclic groups, $G = C_1 \times C_2 \times \ldots \times C_m$, where C_i is of order n_i. Since everything is commutative, we can simply multiply the representations of the individual factors to get a representation of G:

$$T((g_1,\ldots,g_m)) = T_{k_1}(g_1) \ldots T_{k_m}(g_m) \qquad 0 \le k_i < n_i,$$

where T_i is as in example 2.2.

Exercise 2.7: Show that this gives $|G| = n_1 \ldots n_m$ non-isomorphic representations of G.

Example 2.8: Let $G = S_n$, the symmetric group on n objects, thought of as self-maps of the set $\{1,\ldots,n\}$. Let B be a vector space of dimension n with basis $v_1,\ldots,v_n$, and let S_n act on V by permutations $T_V(\pi)(v_i) = v_{\pi(i)}$, and extend linearly. This is called the **permutation representation.**

Exercise 2.9: Verify that this is a representation; the only thing that really needs checking is that this is a left action.

Note that S_n also has a non-trivial one-dimensional "sign representation" given by $\pi \rightarrow \operatorname{sgn} \pi$.

Example 2.10: Let G be any finite group and let k[G] denote the vector space over k of dimension $|G|$ whose basis consists simply of the elements of G (the reason for this notation is because k[G] is actually a ring, as we shall see in section 9). k[G] consists of formal linear combinations of elements of G, $\Sigma_{g\in G}\, a_g g$. Then G acts on k[G] by $h(\Sigma_{g\in G}\, a_g g) = \Sigma_{g\in G}\, a_g(hg)$. This is called the **regular** representation of G, and it shows that any group has a representation in which every non-identity element acts non-trivially. (Note that we could construct this same representation by first using Cayley's theorem to embed G in the symmetric group $S_{|G|}$, and then taking using the permutation representation of example 2.8.)

Exercise 2.11: Show that: (i) the representation k[G] is G-isomorphic to the vector space of functions $f\colon G \rightarrow k$ with the action $(hf)(g) = f(h^{-1}g)$; (ii) we could give a different but G-isomorphic action on these functions by $(hf)(g) =$

f(gh). What is the isomorphism with k[G] in this case? (iii) Show that the action on functions given by (hf)(g) = f(hg) is **not** a left action, but a right action.

Example 2.12: Let **H** be the algebra over **R** of quaterions, given by **H** = { a + bi + cj + dk | a, b, c, d ∈ **R** }, with multiplication given by the rule ij = k, jk = i, ki = j, and ji = -k, kj = -i , ik = -j. Verify that Q = { ±1, ±i, ±j, ±k } is a group of order eight (called the *quaternion group of order eight*) which acts on **H** by left multiplication, giving a four-dimensional representation over **R**. If we treat **H** as a vector space over **C** by letting **C** act by right multiplication, this gives a two-dimensional representation over **C**. (Why do we let **C** act by *right* multiplication?)

Exercise 2.13: Show that the center of Q is {±1}. Show that Q/{±1} is isomorphic to $C_2 \times C_2$ and use this to construct three non-trivial one-dimensional representations of Q.

3. Submodules, Invariants, and Averaging

Definition 3.1: If V is a G-module, W ⊆ V is a **submodule** if it is a subspace of V which is preserved by the G-action, i.e. gW ⊆ W ∀g∈G (in which case gW = W). A G-module V always has two **trivial** submodules, {0} and V itself (of course this is different from being a trivial representation). V is **irreducible** if it is non-zero and has no submodules other than the trivial ones. We shall see that irreducible representations are the building blocks out which all representations are made.

Exercise 3.2: Show that a trivial representation is irreducible if and only if it is one-dimensional.

Exercise 3.3: If Φ: $V \to W$ is a G-map, show that the kernel of Φ is a submodule of V and the image of Φ is a submodule of W.

Exercise 3.4: Consider the permutation representation of S_n on the n-dimensional vector space V of example 2.8. Show that the subspace V_0 spanned by the vector $v_1 + \cdots + v_n$ is a submodule. Show that the subspace $V_1 = \{ \sum_{i=1}^{n} a_i v_i \mid \sum_{i=1}^{n} a_i = 0 \}$ is a submodule.

Definition 3.5: Given a G-module V, the **invariants** of V under the action of G (denoted V^G) is the set of elements of V fixed by the G-action, i.e. $V^G = \{ v \in V \mid gv = v, \forall g \in G \}$. It is easy to see that V^G is a submodule of V.

Exercise 3.6: Show that if V is irreducible, $V^G = \{0\}$, unless V is the trivial representation (in which case V is one-dimensional).

Exercise 3.7: Show that V_0 of exercise 3.4 is equal to V^G.

Definition 3.8: Given a G-module V, the **averaging map** A_V: $V \to V$ is given by $A_V(v) = \frac{1}{|G|}\Sigma_{g \in G}\, gv$.

Geometrically, $A_V(v)$ is the center of mass of the G-orbit of v. Of course, we cannot take $\frac{1}{|G|}$ if the characteristic of k divides $|G|$. To avoid such problems, we'll assume from now on that k has characteristic zero, although for much of what follows it is sufficient to assume only that the characteristic of k does not divide $|G|$. Note that as an endomorphism of V, $A_V = \frac{1}{|G|}\Sigma_{g \in G}\, T_V(g)$.

The following properties of A_V follow directly from the definition:

(3.8.i) $gA_V(v) = A_V(v)$, i.e. $A_V(v) \in V^G$

(3.8.ii) $A_V(v) = v$ if $v \in V^G$

(3.8.iii) $A_V(gv) = A_V(v)$

Properties (i) and (ii) show that A_V is a projection of V onto V^G. From (i) and (iii) we see that $A_V(gv) = A_V(v) = gA_V(v)$, so A_V is a G-endomorphism of V.

Exercise 3.9: Verify properties (3.8.i-iii) of A_V.

Exercise 3.10: If V is irreducible, show that A_V is either 0 or the identity, and the latter obtains only if V is trivial.

Exercise 3.11: Returning to exercises 3.4 and 3.7, show that V_0 is the image of A_V and V_1 is its kernel.

4. New Representations from Old

If V and W are G-modules, we can make their direct sum $V \oplus W$ into a G-module in the obvious way: $g(v,w) = (gv,gw)$. Similarly let G act on the tensor product $V \otimes W$ by $g(v \otimes w) = (gv) \otimes (gw)$. (The existence of a linear map taking $v \otimes w$ to $(gv) \otimes (gw)$ follows from the universal mapping property of tensor products.)

It is slightly trickier to make the dual space V^* into a representation. If S and T are linear endomorphisms of V we have maps S^{tr} and T^{tr}: $V^* \to V^*$, but there is a "twist": $(S \circ T)^{tr} = T^{tr} \circ S^{tr}$. The result of this is that letting g act on V^* by $(T_V(g))^{tr}$ gives a right action rather than a left action (check this). We can correct this by letting g act on V^* by $(T_V(g^{-1}))^{tr}$—the inverse introduces a compensating twist.

Exercise 4.1: Show that this does give a G-action.

Exercise 4.2: Show that the space of k-valued functions on G, with the action $(hf)(g) = f(h^{-1}g)$ considered in exercise 2.11 (i) is isomorphic to $k[G]^*$ in a natural way.

Finally, if V and W are G-modules, we can make a G-module out of Hom(V,W), the space of all (k-)linear maps from $V \to W$, by introducing a similar twist. For $\Phi \in \text{Hom}(V,W)$, let $g\Phi = T_W(g) \circ \Phi \circ T_V(g^{-1})$. The crucial fact about this representation is this:

Proposition 4.3: $\Phi \in (\text{Hom}(V,W))^G$ if and only if Φ is a G-map.

Proof: $\Phi \in (\text{Hom}(V,W))^G \Leftrightarrow \forall\ g\in G\ g\Phi = \Phi \Leftrightarrow \forall\ g\in G\ T_W(g) \circ \Phi \circ T_V(g^{-1}) = \Phi \Leftrightarrow \forall\ g\in G\ T_W(g) \circ \Phi = \Phi \circ T_V(g) \Leftrightarrow \forall\ g\in G, v\in V\ g\Phi(v) = \Phi(gv)$ $\Leftrightarrow \Phi$ is a G-map. •

Exercise 4.4: Show that the natural isomorphism $V^*\otimes W \to \text{Hom}(V,W)$ of exercise 0.4 is compatible with the G-structures given above, that is, that the isomorphism is in fact a G-isomorphism.

Exercise 4.5: Show that if a G-map Φ: $V \to V$ is a projection, then V is G-isomorphic to a direct sum of the kernel of Φ and the image of Φ. Conclude that in exercise 3.11 $V = V_0 \oplus V_1$.

Exercise 4.6: Let G be the cyclic group with two elements {e,o} acting on the space V of real-valued functions of a real variable. Let o act by reflecting the graph of a function in the y-axis, (of)(x) = f(-x). Use the averaging map and the previous exercise to decompose any function as the sum of an even function and an odd function. [Do not worry about infinite dimensionality.]

5. Complete Reducibility

Here is a crucial piece of our story:

Definition 5.1: A G-module V is said to be **completely reducible** if it can be decomposed as a direct sum of irreducible representations.

Theorem 5.2: Every finite dimensional representation of a finite group G is completely reducible.

Although finite dimensionality is not a necessary hypothesis here, we assume it to simplify the proof—but see exercise 10.4.

Proof: Let V be a representation, and assume by induction that the result holds for all representations of dimension smaller than V. If V is irreducible, the theorem is trivially true. Otherwise it has a non-trivial submodule W. If we could find a submodule $W' \subseteq V$ with $V = W \oplus W'$ (that is, a *complement* to W), we would be done: non-triviality of W implies that both W and W' have dimension smaller than V, hence each is a direct sum of irreducibles, and the result follows. Thus we need only the following:

Lemma 5.3: Let V be a G-module, W a submodule. Then there is a submodule W' of V such that $V = W \oplus W'$

Proof: Let $\pi \in \mathrm{Hom}(V,W)$ be a projection of V onto W. The kernel of π is a complement to W as a subspace of V, but it need not be a submodule of V; our strategy is to turn π into a G-map by averaging, because the kernel of a G-map is always a submodule. Let $\pi' = A_{\mathrm{Hom}(V,W)}\pi$; π' is in $(\mathrm{Hom}(V,W))^G$, hence π' is a G-map by proposition 4.3. Since $\pi' \in \mathrm{Hom}(V,W)$, we need only show that it is the identity on W to show that π' is a projection onto W. Then we can complete the proof by taking $W' = \ker \pi'$, since exercise 4.5 tells us that $V = \mathrm{im}\, \pi' \oplus \ker \pi' = W \oplus W'$.

$$\text{For } w \in W, \pi'(w) = \frac{1}{|G|}\Sigma_{g\in G} (T_W(g)\circ\pi\circ T_V(g^{-1}))w$$
$$= \frac{1}{|G|}\Sigma_{g\in G} T_W(g)(\pi(T_V(g^{-1})w)).$$
$$= \frac{1}{|G|}\Sigma_{g\in G} g(\pi(g^{-1}w)).$$

If $w\in W$, $g^{-1}w \in W$, since W is a submodule. Then $\pi(g^{-1}w) = g^{-1}w$, since π is a projection onto W. Finally, $g(\pi(g^{-1}w)) = g(g^{-1}w) = w$, and so $\pi'(w) = \frac{1}{|G|}\Sigma_{g\in G}w = w$. Thus π' is a projection onto W and the proof is complete. (Note that the essential point of this last argument is that since restricted to W is the identity, it is a G-map. Thus is it in $(\mathrm{Hom}(V,W))^G$ by proposition 4.3, so that when we average it, it is still the identity map by (3.8.ii).) •

Note that this proof depends on averaging, and hence on our earlier assumption that k has characteristic zero. When the characteristic of k divides |G|, theorem 5.2 is not valid. On the other hand, the result does hold when G is a compact Lie group—because the same averaging process works.

Of course when we split V completely into irreducibles, some of the summands may be isomorphic to one another. Suppose that $V_1,\ldots,V_k$ are the non-isomorphic represenations occurring in V, and there are n_i copies of V_i, which we call $V_{i,1},\ldots V_{i,n_i}$. Then V is the direct sum $(V_{1,1}\oplus V_{1,2}\oplus\cdots\oplus V_{1,n_1})\oplus\cdots\oplus(V_{i,1}\oplus\cdots\oplus V_{i,n_i})\oplus\cdots\oplus(V_{k,1}\oplus\cdots\oplus V_{k,n_k})$. Up to isomorphism, we can write V as $n_1V_1\oplus\cdots\oplus n_kV_k$ or $\oplus_{i=1}^{k} n_iV_i$, indicating that V_i occurs in V with multiplicity n_i. The splitting into individual irreducibles $V_{i,j}$ is not uniquely determined. This is analogous to the situation in a vector space, where one can think of a basis as spliting the vector space into a direct sum of one-dimensional—i.e., irreducible—subspaces, but there are infinitely many choices of bases. However, the numbers n_i *are* uniquely determined (see corollary 8.7 for the proof), just as the dimension of a vector space (the number of one-dimensional subspaces given by the basis) is. Furthermore, the direct sum of all the irreducibles of each isomorphism type is also uniquely determined, that is, $V_{i,1}\oplus\cdots\oplus V_{i,n_i}$ is independent of any choices made in the decomposition (see exercise 6.9 for the proof).

Exercise 5.4: As a special case, verify the two facts metioned above for the trivial submodule. Let V_1 denote the trivial represenation, and show that $V_{1,1}\oplus V_{1,2}\oplus\cdots\oplus V_{1,n_1} = V^G$, hence this sum is independent of any choices. Similarly, $n_1 = \dim V^G$ is uniquely determined.

Exercise 5.5: Give an alternate proof of lemma 5.3 valid for $k = \mathbf{R}$, by proving the following:

Theorem 5.6: Let V be a representation of G over the field **R**. Then there is an inner product $< , >$ on V in which each G acts orthogonally, i.e. each T(g) is orthogonal, i.e. $<gv,gw> = <v,w>$.

Then complete exercise 5.5 by proving that under such an inner product the orthogonal complement of any submodule is a submodule. (Hint for theorem 5.6: Let $b \in V^*\otimes V^*$ be any positive definite inner product. Show that $A_{V^*\otimes V^*}b$ is a positive definite inner product with the desired property.)

Exercise 5.7: Modify this proof to work over $k = \mathbf{C}$, using Hermitian inner products and unitary operators.

Exercise 5.8: Let V be a representation of G over an algebraically closed field. Show that for every $g\in G$, $T_V(g)$ can be diagonalized and that the eigenvalues are n-th roots of unity, where $n = |G|$. (Here are hints outlining two proofs. Hint 1 (valid over **C** only): Use the previous exercise and the diagonalization theorem for unitary operators. Hint 2 (valid when characteristic of k does not divide $|G|$, in particular for characteristic zero): Use Jordan canonical form and the fact that T(g) is of finite order.)

Exercise 5.9: Prove that the n one-dimensional representations of the cyclic group C_n given in example 2.2 give all possible irreducible representations of

C_n over **C**. (Hint: Apply the preceding exercise to T(g) to show that any representation decomposes as a direct sum of the desired type.)

Exercise 5.10: Prove a result analogous to exercise 5.9 for an arbitrary finite abelian group, using the one-dimensional representations of example 2.6. (Hint: Prove the fact that if A and B are linear transformations which commute, then A preserves the eigenspaces of B.)

6. Schur's Lemma and its Consequences

The following simple lemma turns out to be crucial:

Lemma 6.1 (Schur's lemma): Let V, W be irreducible G-modules, and let $\Phi: V \to W$ be a G-map. Then either $\Phi = 0$ or Φ is a G-isomorphism.

Proof: Suppose $\Phi \neq 0$. Then $\ker \Phi \neq V$. But $\ker \Phi$ is a submodule, so by irreducibility $\ker \Phi = \{0\}$ and Φ is one-to-one. Similarly, since $\Phi \neq 0$, $\mathrm{im}\, \Phi \neq \{0\}$, so $\mathrm{im}\, \Phi = W$, and Φ is onto. Thus Φ is an invertible linear map, hence a G-isomorphism by exercise 3.1. •

If V is a G-module, the G-endomorphims of V, $\mathrm{End}_G(V)$ = {G-maps from V to itself} $\subseteq$ Hom(V,V), is a finite dimensional vector space over k. It contains the identity operator I and all scalar multiplication operators kI; if $V \neq \{0\}$ these maps form a one-dimensional subspace. Furthermore, $\mathrm{End}_G(V)$ is closed under composition and the elements of kI commute with every element, so it forms an algebra over k. Schur's lemma tells us that if V is irreducible, any non-zero element of $\mathrm{End}_G(V)$ is invertible. In other words, we have:

Corollary 6.2: If V is an irreducible representation of G, $\mathrm{End}_G(V)$ forms a division algebra over k.

The following theorem tells us that if k is algebraically closed, the situation is even simpler.

Theorem 6.3: Let V be an irreducible representation of G over an algebraically closed field k. Then $End_G(V) = kI$, i.e. the only G-endomorphisms of V are given by scalar multiplication.

This follows immediately from:

Lemma 6.4: Let D be a finite dimensional division algebra over an algebraically closed field k. Then D = k.

Proof: Let $d \in D$. The idea of the proof is that powers of d generate a finite dimensional extension field of k, which must be trivial since k is algebraically closed. By finite dimensionality, powers of D must be linearly dependent, hence d satisfies some polynomial $p(x) \in k[x]$ of smallest possible degree. If deg $p > 1$, p must factor non-trivially over k, i.e. $p(x) = p_1(x)p_2(x)$, since k is algebraically closed. Now $0 = p(d) = p_1(d)p_2(d)$. Since D is a division algebra it has no zero divisors, so $p_1(d) = 0$ or $p_2(d) = 0$ contradicting the minimality of p. Thus deg $p \leq 1$ and it follows that $d \in k$. •

From now on, unless explicitly stated otherwise, we assume that k is algebraically closed; take k = **C** if you like.

Propostion 6.5: Suppose that V and W are irreducible G-modules and that $\Phi: V \to W$ is a non-zero G-map (hence an isomorphism by Schur's lemma). If $\Phi': V \to W$ is another G-map, $\Phi' = \alpha\Phi$.

Proof: $\Phi^{-1} \circ \Phi' \in End_G(V)$; by the above theorem $\Phi^{-1} \circ \Phi' = \alpha I$, and hence $\Phi' = \alpha\Phi$. •

Corollary 6.6: Let V, W be irreducible G-modules. Then:

$$\dim \{\text{G-maps } V \to W\} = \begin{cases} 1 & \text{V isomorphic to W} \\ 0 & \text{otherwise} \end{cases}$$

Exercise 6.7: Show that if V is reducible, $End_G(V)$ contains zero-divisors, and hence it cannot be a division algebra; this is a converse of corollary 6.2.

Exercise 6.8: This problem demonstrates that the hypothesis that k be algebraically closed is necessary in theorem 6.4. Consider the quaternion group Q of order eight from example 2.12, and consider the representation **H** as a four-dimensional representation over the non-algebraically closed field **R**. Show that $\mathrm{End}_Q(\mathbf{H})$ is anti-isomorphic to the division algebra **H** (that is, there is a map ϕ which is an isomorphism except that $\phi(ab) = \phi(b)\phi(a)$. This twist occurs because we let **H** act by *right* multiplication). Now consider **H** as a two-dimensional complex representation. Why is the endomorphism group smaller?

The following exercise proves the results asserted in the discussion preceding exercise 5.4:

Exercise 6.9: (i) Suppose that W is an irreducible G-module, V is a G-module split into irreducible components, and Φ: W$\rightarrow$V is a G-map. Show that the image of Φ is contained in the sum of the irreducible components of V isomorphic to W. Hint: Apply Schur's lemma to the projection of Φ onto each irreducible component of V. (ii) Use the result in (i) to prove that the sum of the irreducibles in a particular isomorphism class is independent of the splitting of V.

7. Characters

This section and the following one make fundamental use of the finite dimensionality of our representation, since it relies heavily on traces of linear transformations. Similarly, it assumes that k = **C**, since complex conjugation is expoited throughout. Both of these assumptions will be used from now on.

Definition 7.1: Let V be a representation of G. The **character** of the representation V is a complex-valued function on G, χ_V: G$\rightarrow$**C**, given by $\chi_V(g) = \mathrm{tr}\, T_V(g)$.

Note that $\chi(hgh^{-1}) = \text{tr}\,(T(h)T(g)T(h^{-1})) = \text{tr}\,T(g) = \chi(g)$, that is, χ is constant on the conjugacy classes of G (a **class function**). Also, for a one-dimensional representation we can identify the character with the representation.

Proposition 7.2: (i) $\chi_V(e) = \dim V$

(ii) $\chi_V(g^{-1}) = \overline{\chi_V(g)}$

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

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

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

(vi) $\chi_{\text{Hom}(V,W)}(g) = \overline{\chi_V(g)}\,\chi_W(g)$

Proof: Exercise. (For (ii), use the fact that T(g) has finite order or use exercise 5.7 or 5.8 to show that any eigenvalue α lies on the unit circle. Then use $\alpha^{-1} = \overline{\alpha}$. For (iv), use the fact that if $\{v_i\}$ ($\{w_j\}$ repectively) form a basis for V (W respectively), then $v_i \otimes w_j$ forms a basis of $V\otimes W$. For (vi), use (iv), (v), and exercise 4.4.)

It follows that if $V = \oplus_{i=1}^{k} n_i V_i$, then $\chi_V = \sum_{i=1}^{k} n_i \chi_{V_i}$.

Exercise 7.3: Compute the character of the the representations of the dihedral group given in example 2.4.

Exercise 7.4: Show that the character of the regular representation (example 2.10) is

$$\chi_R(g) = \begin{cases} |G| & g = e \\ 0 & \text{otherwise} \end{cases}$$

Exercise 7.5: For the permutation representation of S_n (example 2.8), show that $\chi_V(\pi)$ is the number of fixed points of π. Using the notation and result of exercise 4.5, show that $\chi_{V_1}(\pi) = \chi_V(\pi) - 1$.

Exercise 7.6: Compute the character of the representation of Q on **H** of example 2.12, first as a representation over **R**, then over **C**.

Exercise 7.7: Show that the class functions on G form a subspace of all C-valued functions on G, whose dimension is the number of conjugacy classes in G.

8. Orthogonality Relations

We can use the character of the representation V to count the dimension of V^G (which equals the number of trivial irreducible components in V by exercise 5.4).

Lemma 8.1: $\dim V^G = \frac{1}{|G|}\Sigma_{g\in G}\, \chi_V(g)$

Proof:
$$\frac{1}{|G|}\Sigma_{g\in G}\, \chi_V(g) = \frac{1}{|G|}\Sigma_{g\in G}\, \mathrm{tr}\, T_V(g)$$
$$= \mathrm{tr}\, [\frac{1}{|G|}\Sigma_{g\in G}\, T_V(g)]$$
$$= \mathrm{tr}\, A_V$$

But A_V is a projection of V onto V^G, and the result follows from this exercise:

Exercise 8.2: Let π: $V\rightarrow W\subseteq V$ be a projection. Then $\mathrm{tr}\,\pi = \dim W$. (Hint: This is easy with the right basis for V.)

We can define a Hermitian inner product on C-valued functions on G by $<f_1,f_2> = \Sigma_{g\in G}\ \overline{f_1(g)}\, f_2(g)$. With respect to this inner product, the characters of irreducible representations are orthonormal, as stated in the following centrally important theorem:

Theorem 8.3 (Orthogonality Relations): Let V, W be irreducible representations of G. Then:
$$<\chi_V,\chi_W> = \begin{cases} 1 & \text{if V is isomorphic to W} \\ 0 & \text{otherwise} \end{cases}$$

Proof:
$$<\chi_V,\chi_W> = \frac{1}{|G|}\, \Sigma_{g\in G}\ \overline{\chi_V(g)}\, \chi_W(g)$$
$$= \frac{1}{|G|}\, \Sigma_{g\in G}\, \chi_{\mathrm{Hom}(V,W)}(g) \qquad \text{by proposition 7.2.vi,}$$

$$= \dim \mathrm{Hom}(V,W)^G \qquad \text{by lemma 8.1,}$$
$$= \begin{cases} 1 & \text{if V is isomorphic to W} \\ 0 & \text{otherwise} \end{cases} \qquad \text{by 6.6.} \quad \bullet$$

An immediate consequence of this theorem is that the characters of irreducible representations, being orthonormal, are linearly independent class functions. By exercise 7.7, the dimension of the space of class functions is equal to the number of conjugacy classes in G; thus we have the following result:

Corollary 8.4: The number of irreducible representations of a finite group G is less than or equal to the number of conjugacy classes of G.

We shall see in section 9 that these numbers are in fact equal.

Exercise 8.5: Use the inequality of corollary 8.4 to show that the irreducible representations over **C** of a finite abelian group given in example 2.6 are the only possible ones, and prove that in this particular case the inequality is an equality. Note how much easier this proof is than the one outlined in exercise 5.9.

Exercise 8.6: Show that the Hermitian inner product on functions given above is invariant under the G-action on functions given in exercise 2.11.

Using characters we can compute the multiplicity of an irreducible reperesentation in an arbitrary representation, and deduce a string of fundamental results:

Theorem 8.7: Let V be a representation of G, and W an irreducible representation. Then $<\chi_W,\chi_V>$ is equal to the multiplicity of W in V.

Proof: Suppose that $V = \oplus_{i=1}^{k} n_i V_i$ where the V_i are distinct irreducible representations. Then $<\chi_W,\chi_V> = <\chi_W,\sum_{i=1}^{k} n_i\chi_{V_i}> = \sum_{i=1}^{k} n_i<\chi_W,\chi_{V_i}>$. If W is not among the V_i, then $<\chi_W,\chi_{V_i}>$ is zero by the orthogonality relations, and the formula follows since the multiplicity of W is zero. If W is isomorphic to V_j, then $<\chi_W,\chi_{V_i}> = \delta_{ij}$, so $\sum_{i=1}^{k} n_i<\chi_W,\chi_{V_i}> = \sum_{i=1}^{k} n_i\delta_{ij} = n_j$, proving the formula. •

Corollary 8.8: The multiplicities of irreducible representations in V are independent of the splitting of V.

Corollary 8.9: A representation is determined up to isomorphism by its character.

Proposition 8.10: Let V be an irreducible representation of G. Then V occurs in the regular representation k[G] of G with multiplicity dim V. In particular, every irreducible representation of G occurs in k[G].

Proof: By theorem 8.7, the multiplicity is

$$
\begin{aligned}
<\chi_V,\chi_R> &= \frac{1}{|G|}\sum_{g\in G} \overline{\chi_V(g)}\, \chi_R(g) \\
&= \frac{1}{|G|}\, \overline{\chi_V(e)}\; |G| && \text{by exercise 7.4} \\
&= \overline{\chi_V(e)} \\
&= \dim V && \text{by proposition 7.2.i.} \quad \bullet
\end{aligned}
$$

Corollary 8.11: There are only finitely many isomorphism types of irreducible representations of G.

Corollary 8.12: Let $V_1,\ldots,V_h$ be the distinct irreducible representations of G, and let d_i be the dimension of V_i. Then $\sum_{i=1}^{h} d_i^2 = |G|$.

Proof: By the previous proposition, $k[G] = \oplus_{i=1}^{h} d_iV_i$. But $\dim k[G] = |G|$, and $\dim \oplus_{i=1}^{h} d_iV_i = \sum_{i=1}^{h} d_i^2$. •

Exercise 8.13: If $V = \oplus_{i=1}^{k} n_iV_i$ where the V_i are distinct irreducible representations, show that $<\chi_V,\chi_V> = \Sigma_{i=1} n_i^2$. Use this result to give another proof of corollary 8.12. Also, prove a partial converse to theorem 8.3: if $<\chi_V,\chi_V> = 1$, then V is irreducible.

Example 8.14: What are the irreducible representations over **C** of S_3?

S_3 has three conjugacy classes—the identity, the (three) transpositions, and the (two) 3-cycles. We know two irreducible representations—the trivial representation and the sign representation—whose respective characters χ_1 and χ_2 are:

$\chi_1(e) = 1$	$\chi_1((1\ 2)) = 1$	$\chi_1((1\ 2\ 3)) = 1$
$\chi_2(e) = 1$	$\chi_2((1\ 2)) = -1$	$\chi_2((1\ 2\ 3)) = 1$

Let χ_3 be the character of the representation V_1 of exercise 3.4. By exercise 7.5,

$\chi_3(e) = 2$	$\chi_3((1\ 2)) = 0$	$\chi_3((1\ 2\ 3)) = -1$

$<\chi_3,\chi_3> = (1/6)\ (1 \cdot 2^2 + 3 \cdot 0^2 + 2 \cdot (-1)^2) = 1$, so χ_3 is irreducible by the previous exercise. Since S_3 has three conjugacy classes, it has at most three irreducible representations (by corollary 8.4), and thus we have found them all.

Corollary 8.12 gives a numerical constraint on the dimensions of the irreducible representations which can be used in the computation of the table of characters of all the representations of a given finite group.

Exercise 8.15: Use corollary 8.12 and the orthogonality relations to find the character table of S_3 without using any knowledge of the representation V_1. You may use your knowledge of χ_1 and χ_2.

We will not prove another important numerical constraint: the dimension of an irreducible representation divides the order of the group (in fact, it divides the index of the center in the group). This follows from results on algebraic integers. See [Serre, p. 53].

Exercise 8.16: Compute the character table of the dihedral group of order 2n, using the representations constructed in example 2.4, some of which may be reducible. (Hint: It is easiest to understand the conjugacy classes in this group geometrically in terms of their action on an n-gon. Treat the cases n even and n odd separately.)

Exercise 8.17: Find all the representations over **C** of the quaternion group Q of order eight (See example 2.12 and exercise 7.6).

9. The Group Algebra

The previous section contained most of the basic results from representation theory. In this section, we will translate the results of our earlier work into the language of ring theory. We will use the structural information we gain to tie up a loose ends from before: theorem 9.9 shows that a finite group has as many irreducible representations as it has conjugacy classes, strengthening our earlier inequality in corollary 8.4.

From now on, all rings and algebras will be assumed to have a multiplicative identity. In particular, all direct sums of rings must be finite.

We can define a multiplication on the regular representation k[G] (example 2.10) simply by extending the group multiplication linearly: $(\Sigma_{g\in G}a_g g)(\Sigma_{g\in G}b_g g) = \Sigma_{g,h\in G}\, a_g b_h(gh)$. This algebra is known as the **group ring** or **group algebra**, and is denoted k[G] or simply kG. The element $1\cdot e$ acts as a multiplicative identity, and we shall simply write it as 1. Thus k[G] is the algebra generated over k by the elements of G with their group product, hence the notation.

Exercise 9.1: Verify that k[G] is in fact an algebra over k (i.e. a ring containing an isomorphic copy of k in its center), and that it is commutative if and only if the group G is commutative.

A representation can be turned into a left module over the group ring by defining $(\Sigma_{g\in G}a_g g)v = \Sigma_{g\in G}a_g gv$. Conversely, any left module over k[G] gives a representation just by restricting the action to group elements. In fact, this correspondence between representations and modules over k[G] is the origin of the terminology "G-modules," "submodules," etc., for representations of G.

Example 9.2: V is irreducible as a representation if and only if it is simple as a k[G] module. (A module over a ring R is **simple** if it has no submodules other than {0} and itself).

Example 9.3: The regular representation corresponds to k[G] acting on itself on the left.

Exercise 9.4: Prove that a module over k[G] is finitely generated if and only if the corresponding representation is finite dimensional. You will need to use the fact that G is finite.

Exercise 9.5: If $C \subseteq G$ is a conjugacy class, show that $\Sigma_{g \in C}\, g$ lies in the center of k[G], denoted Z(k[G]). Show that such elements form a basis for the center as a vector space over k, and that the dimension of the center is therefore equal to the number of conjugacy classes of G.

Exercise 9.6: Rewrite the multiplication law on k[G] in terms of functions on G (exercise 2.11). Note the similarity to "convolution" in Fourier analysis. Show also that the center of k[G], corresponds to the class functions, and thus that exercise 9.5 and exercise 7.7 state the same result.

Our first goal is to show k[G] can be split as a direct sum of matrix algebras. If V is a representations of G, every element of k[G] acts linearly on V, and hence we have a ring homomorphism $k[G] \to End_k(V)$. Of course, if V is n dimensional, $End_k(V)$ is simply isomorphic to the ring $Mat_n(k)$ of $n \times n$ matrices over k. If $V_1, \cdots, V_h$ are the distinct irreducible representations and V_i is of dimension d_i, we can combine these homomorphisms into a map $k[G] \to \oplus_{i=1}^{h} End_k(V_i) \approx \oplus_{i=1}^{h} Mat_{d_i}(k)$. This map reflects the action of k[G] on each irreducible representation of G. Using the results of section 8, which were derived for $k = \mathbf{C}$, we can prove the following theorem, which is proved for more general ground fields in exercise 9.17:

Theorem 9.7: The group algebra **C**[G] is isomorphic to the direct sum $\oplus_{i=1}^{h}\mathrm{End}_{\mathbf{C}}(V_i) \approx \oplus_{i=1}^{h}\mathrm{Mat}_{d_i}(\mathbf{C})$.

Proof: An element of $a \in \mathbf{C}[G]$ is in the kernel of the homomorphism into $\oplus_{i=1}^{h}\mathrm{End}_{\mathbf{C}}(V_i)$ exactly when it acts as multiplication by zero on any irreducible representation of G, and hence on any representation of G (in particular on the regular representation itself). Hence $a \cdot 1 = 0$, so $a = 0$, and the map is injective. The dimension of $\oplus_{i=1}^{h}\mathrm{End}_{\mathbf{C}}(V_i) \approx \oplus_{i=1}^{h}\mathrm{Mat}_{d_i}(\mathbf{C})$ is $\sum_{i=1}^{h} d_i^2$, which is $|G|$ by corollary 8.12. But $|G|$ is the dimension of k[G], and hence the map is surjective. •

Exercise 9.8: Let V be an irreducible representation of a finite group G. Then the linear transformations $\{T_V(g) \mid g \in G\}$ span the space of all k-linear self maps of V, $\mathrm{End}_k(V)$. We shall see later that we can prove this result as a corollary of Burnside's theorem 10.17.

We are finally in a position to prove the following result:

Theorem 9.9: There are as many irreducible representations of a finite group G as there are conjugacy classes in G.

Proof: Use the following two exercises to show that the number of conjugacy classes in G equals $\dim_k Z(k[G])$, which is the number of conjugacy classes in G by exercise 9.5.

Exercise 9.10: If $R_1, \cdots, R_h$ are rings, then the center $Z(\oplus_{i=1}^{h} R_i) = \oplus_{i=1}^{h} Z(R_i)$.

Exercise 9.11: The center $Z(\mathrm{Mat}_n(k))$ consists of the diagonal matrices. Hint: A matrix $Z \in Z(\mathrm{Mat}_n(k))$ must commute with the elementary matrix Δ_{ij} which has a 1 in the (i,j) entry and zeros elsewhere. Alternatively, die-hard computation-haters can prove this result in the following way: Regarding $Mat_n(k)$ as endomorphisms of k^n, show that for $v \in k^n$ there is an element of $Mat_n(k)$ whose zero-eigenspace consists solely of multiples of v. Then use the hint to exercise 5.10 to show that Z must preserve this subspace, that is, that v is an eigenvector of Z. Finally, prove that since every vector is an eigenvector of Z, Z must be a scalar multiplication.

The goal of the next few exercises is to give another proof that k[G] is isomorphic to $\oplus_{i=1}^{h} Mat_{d_i}(k)$ (part of theorem 9.7). The proof is based on computing the ring $End_{k[G]}(k[G])$ and showing that it is "almost isomorphic" to k[G]. Unlike the previous proof, it is independent of the material on characters developed in section 8; hence it is valid over an arbitrary algebraically closed field whose characteristic does not divide the order of G. In particular, the results 9.8-11 remain valid over such a field. Many of the results presented below will be used heavily in section 10.

Exercise 9.12: If R is a ring, show that n×n matrices with coefficients in R, $Mat_n(R)$ form a ring. (We are simply verifying that the commutativity of R is not essential in working with matrices over R).

Exercise 9.13: Let R be a ring, and let $E = E_1 \oplus \cdots \oplus E_n$ and $F = F_1 \oplus \cdots \oplus F_m$ be direct sum decompositions of modules over R. Show that $Hom_R(E,F)$ is naturally isomorphic to $\oplus_{1 \le i \le m,\ 1 \le j \le n} Hom_R(E_j, F_i)$ by decomposing $\phi \in Hom_R(E,F)$ as an m×n matrix $[\phi_{ij}]_{1 \le i \le m,\ 1 \le j \le n}$. Show that under this isomorphism, the composition of two maps (when defined) follows the usual law of matrix multiplication $[\phi' \circ \phi]_{ij} = \sum_{k=1}^{n} \phi'_{ik} \circ \phi_{kj}$.

Exercise 9.14: Show that if k is algebraically closed, $End_{k[G]}(k[G])$ is isomorphic to a direct sum $\oplus_{i=1}^{h} Mat_{d_i}(k)$, where d_i is the multiplicity in k[G] of the i-th irreducible representation. Hint: Use the decomposition of k[G] into irreducible G-modules, the previous exercise, and proposition 6.5.

Note that when k ≠ **C**, we have not proved the results of proposition 8.10: that d_i is equal to the dimension of the i-th irreducible representation, and that every irreducible representation occurs in k[G]. The first of these follows from exercise 9.17 and exercise 10.1.vi and vii, while the second can be deduced from exercise 10.15.

Exercise 9.15: (i) Let R be a ring. Show that the self-maps of R which commute with left multiplication (i.e. $\phi(ab) = a\phi(b)$) are all given by *right* multiplication by elements of R. Hint: Where does the map take the element 1? (This is an example of the principle that *left multiplications commute with right multiplications and nothing else*. Specifically, if L_x is the map which left multiplies by x, and R_y is the map which right multiplies by y, $L_xR_y = R_yL_x$—a formula which is equivalent to the associative law. Furthermore, any map which commutes with all L_x must be R_y for some y. This same principle can be used in the theory of Lie groups to show that left-invariant vector fields are infinitesimal *right* translations, since the flows they generate commute with all *left* translations.) (ii) Show that the ring of endomorphisms of R considered as a left R-module, $End_R(R)$, is anti-isomorphic to R (see exercise 6.8 for the concept of anti-isomorphism). (iii) Prove the result of exercise 6.8 using part (ii) applied to the division algebra **H**.

Exercise 9.16: Demonstrate the following:

(i) The composition of two anti-isomorphisms is an isomorphism; hence if a ring R is anti-isomorphic to R' and R'', R' must be ismorphic to R''.

(ii) The ring $Mat_n(k)$ is anti-isomorphic to itself (Hint: Use the transpose map).

Exercise 9.17: Prove k[G] is isomorphic to $\oplus_{i=1}^{h} Mat_{d_i}(k)$, using the results of the previous three exercises.

10. Structure of Semisimple Algebras

In the previous section we concluded with a theorem on the structure of the group algebra. In this section we will see (through a series of exercises) that this result is a special case of the structure theory of semisimple rings. This structure theorem 10.14 and Burnside's theorem 10.17 are the primary goals of this section.

Exercise 10.1: Let V be an n-dimensional vector space over k, (say k^n) with basis $\{v_1,\cdots,v_n\}$. Let R = $End_k(V)$ be the algebra of linear self-maps of V (essentially the algebra of n×n matrices over k). The purpose of this exercise is to give a picture of this algebra and its ideals. Prove the following assertions:

(i) Considered as an R module (with the action given by $\phi v = \phi(v)$ for $\phi \in R$, $v \in V$), V is a simple R-module.

(ii) If $W \subseteq V$ is a subspace, the set $\{ \Phi \in R \mid \text{image } \Phi \subseteq W \}$ is a right ideal of R. If W = span $\{v_1,\cdots,v_k\}$, this ideal consists of all matrices with zeros outside of the first k rows. Show that there is a one-to-one correspondence between right ideals and subspaces of V.

(iii) If $W \subseteq V$ is a subspace, the set $\{ \Phi \in R \mid W \subseteq \text{kernel } \Phi \}$ is a left ideal of R. If W = span $\{v_1,\cdots,v_k\}$, this ideal consists of all matrices with zeros in the first k columns. Show that there is a one-to-one correspondence between left ideals and subspaces of V.

(iv) The only two-sided ideals in R are {0} and R.

(v) Under the correspondence in (iii), minimal left ideals (a left ideal is **minimal** if it contains no left ideals other than {0} and itself, that is, it is simple as a module) correspond to n-1 dimensional subspaces of V in a one-to-one fashion.

(vi) Matrices with non-zero entries in only the i-th column form a minimal left ideal. What is the corresponding subspace? R can be written as a direct sum of n minimal left ideals.

(vii) A minimal left ideal in R is isomorphic to V as an R-module. Hint: If the ideal corresponds to a subspace W, consider the evaluation map on a fixed vector not in W.

(viii) Any two minimal left ideals are conjugate by an inner automorphism of R, i.e. if L_1 and L_2 are minimal left ideals, that there is a $\phi \in GL(V)$ such that $L_1 = \phi L_2 \phi^{-1}$.

Exercise 10.2: Let R be a ring. Prove the following assertions:

(i) There is a ring $\overline{R}$ which is anti-isomorphic to R. Hint: Take the same underlying set as R, but with a different multiplication operation. Show that $\overline{\overline{R}}$ is isomorphic to R. You can do this directly or use the fact that the composition of two anti-isomorphisms is an isomorphism (exercise 9.16.i).

(ii) Left R-modules correspond precisely to right $\overline{R}$ -modules. Similarly, right R modules correspond to left $\overline{R}$ -modules.

(iii) Exercise 9.15.ii can be interpreted as saying that $End_R(R)$ is isomorphic to $\overline{R}$; we have essentially turned right multiplication by elements of R into a left action of $\overline{R}$.

(iv) $Mat_n(R)$ is anti-isomorphic to $Mat_n(\overline{R})$. Hint: Use the transpose map.

(v) If R is commutative, $Mat_n(R)$ is anti-isomorphic to itself. (Use (iv). Note that this is essentially identical to exercise 9.16.ii.)

Exercise 10.3: The goal of this exercise is to show that linear algebra still works over a division ring D, that is, that commutativity is not a necessary hypothesis.

(i) Show that finitely generated modules over division rings have bases (i.e., are free modules) with a well-defined dimension. Hint: Convince yourself that the usual proof works. From now on, we will refer to finitely generated modules over D as finite dimensional D-vector spaces.

(ii) Show that linear transformations are represented by matrices. There are two ways to do this, and it is probably worth going through them both. The first is to develop the theory exactly as one does for vector spaces, taking care not to use commutativity. You will notice that if you try to write the matrices on the left, you will have to introduce a different law for matrix multiplication which switches the order of the factors. The second way is to use exercise 9.13 to reduce the problem to matrices, and exercise 10.2.iii to show that the entries lie in $\overline{D}$ (since $End_D(D)$ must act by *right* multiplication). Note that putting the entries in $\overline{D}$ solves the problem of switching the order of factors in the matrix multiplication law. Thus, for example, if V is an n-dimensional D-vector space, $End_D(V)$ is isomorphic to $Mat_n(\overline{D})$.

(iii) Show that the results of exercise 10.1 still hold over a division ring.

(iv) Let V be a finite-dimensional D-vector space, and let $R = End_D(V)$ be the ring of D-linear self-maps of V. Since V can also be regarded as an R-module, we may consider the ring $End_R(V)$ consisting of the self-maps of V which preserve addition and commute with the action of R. Prove this non-commutative analog of exercise 9.11: If $V \neq \{0\}$, then $End_R(V)$ is isomorphic to D. Hint: (a) Since every element of $R = End_D(V)$ is D-linear, left multiplication of V by elements of D commutes with the action of R. Thus we can find an isomorphic copy of D inside

$End_R(V)$. (b) Show that every element of $End_R(V)$ is of this form, by using elementary matrices in R.

Exercise 10.4: Let M be a module over a ring R. Show that the following conditions are equivalent:

(10.4.i) M is a direct sum of simple modules.

(10.4.ii) M is a sum of simple modules.

(10.4.iii) Every submodule of M has a complement in M. That is, if $E \subseteq M$ is a submodule, there is a submodule F with $E \oplus F = M$.

Under these circumstances, M is called a **semisimple** R-module. Note that theorem 5.3 proves that if G is a finite group with the characteristic of k not dividing |G|, then every module over k[G] is satisfies (iii), and hence is semisimple.

Hint: For (ii) ⇒ (iii), try the special case with M a *finite* sum of simple modules $\oplus M_i$, and build F inductively as a (direct) sum of a subset of the M_i. Then adapt your proof using a Zorn's lemma argument. (iii) ⇒ (i) is essentially theorem 5.2, but again you will need to use a Zorn's lemma argument. Readers who wish to avoid Zorn's lemma will have to use finite sums in (i) and (ii), and add a hypothesis in (iii) that R is a k-algebra and M is finite dimensional as a k-vector space.

Exercise 10.5: Show that quotient modules of semisimple modules are semisimple. Similarly for submodules.

Exercise 10.6: A ring R is called **semisimple** if it is semisimple as a module over itself. By exercise 10.4, R is semisimple if it is a direct sum of minimal left ideals. $End_k(V)$ is semisimple by exercise 10.1.vi, and $End_D(V)$ is semisimple by exercise 10.3.iii. k[G] is semisimple when the characteristic of k does not divide |G|, since *every* k[G] module is semisimple. Prove that this is the general case: if R is semisimple then every R-module is semisimple. Hint: Use the fact that every module is the quotient of a free module.

Exercise 10.7: Prove that a finite direct sum of semisimple rings is semisimple. In particular, $\oplus_{i=1}^{h} Mat_{d_i}(D_i)$ is semisimple.

The goal of the following exercises is a structure theorem 10.14 which says that every semisimple ring is of the form given in exercise 10.7.

Exercise 10.8: Prove that a semisimple ring is a *finite* direct sum of minimal left ideals. Hint: Look at the decomposition of the element 1 under the direct sum, and use the fact that elements of the direct sum can be non-zero in only finitely many places.

Exercise 10.9: Prove Schur's lemma: if M and M' are simple R-modules, then any non-zero element of $Hom_R(M,M')$ is an isomorphism. If M is simple, then $End_R(M)$ is a division ring.

Exercise 10.10: Let $L \subset R$ be a minimal left ideal, and M be a simple R-module. Then if L is not isomorphic to M as an R-module, $LM = \{0\}$. Hint: If $Lm \neq \{0\}$ for some $m \in M$, show that the map $L \to M$ given by right-multiplication by m is an isomorphism.

Exercise 10.11: Prove that the following conditions on a ring R are equivalent:

(10.11.i) R is semisimple and has only one isomorphism type of simple module.

(10.11.ii) R is semisimple and has only one isomorphism type of minimal left ideal.

(10.11.iii) R is a direct sum of isomorphic minimal left ideals.

A ring satisfying these conditions is called **simple**. Show that $End_k(V)$ is simple. Hint: (i) $\Rightarrow$ (ii) and (ii) $\Rightarrow$ (iii) are trivial. For (iii) $\Rightarrow$ (i), use exercise 10.10 and the decomposition of $1 \in R$.

Exercise 10.12: Prove this structure theorem for semisimple rings: a semisimple ring R is a finite direct sum of simple rings. Hint: Split R into minimal left ideals, and collect these into h isomorphism classes: $L_{1,1},\cdots,L_{1,n_1}$ through $L_{h,1},\cdots L_{h,n_h}$. Let R_i be the sum of the minimal ideals of the decomposition which fall in the i-th isomorphism class, $\oplus_{j=1}^{n_i} L_{i,j}$. Then $R = \oplus_{i=1}^{h} R_i$, and clearly R_i is a left ideal and hence closed under multiplication. Use exercise 10.10 to show that $R_iR_j = \{0\}$ if $i \neq j$, and conclude that R_i is a two-sided ideal. Finally, use the decomposition of the identity of R to prove that R_i has an identity element. R_i is simple by condition 10.11.iii.

Exercise 10.13: Prove this structure theorem for simple rings: a simple ring R is isomorphic to a matrix ring over a division ring. Here is an outline of a proof (Note the similarity to exercise 9.17, and also compare exercise 10.3.iv):

(i) By exercise 10.11, R is a direct sum of minimal left ideals $L_1 \oplus \cdots \oplus L_{n}$., each of which is isomorphic to the unique simple module for R.

(ii) Using exercise 10.9 (Schur's lemma), $End_R(L_i)$ is isomorphic to a fixed division ring, call it D. We can regard each L_i as a D-vector space, and note that by definition the elements of R act D-linearly.

(iii) Using exercise 9.13, $End_R(R)$ is isomorphic to the matrix ring $Mat_n(D)$.

(vi) Using exercise 10.2.iii and iv, show that R is isomorphic to $Mat_n(\overline{D})$.

Combining exercises 10.12 and 10.13, we have proved the following:

Theorem 10.14 (structure theorem for semisimple rings): Every semisimple rings can be written as $\oplus_{i=1}^{h}\mathrm{Mat}_{d_i}(D_i)$, a direct sum of matrix rings over division rings.

Exercise 10.15: Let R be a semi-simple ring, and let M be a simple R-module. Show that M is isomorphic to one of the minimal left ideals of R. Hint: Use exercise 10.10.

Exercise 10.16: Let Q be the quaternion group of order eight, of exercise 2.12. Exhibit **R**[Q] explicitly as a direct sum of matrix algebras over division rings.

No introduction to this subject would be complete without the following theorem, which was alluded to in exercise 9.8:

Exercise 10.17: Prove **Burnside's theorem**: let V be a finite-dimensional vector space over an algebraically closed field k, and let $R \subseteq \mathrm{End}_k(V)$ be a subalgebra such that V is a simple R-module. Then $R = \mathrm{End}_k(V)$.

(i) Show that $\mathrm{End}_R(V) = k$. Hint: This is essentially theorem 6.3.

Next, let $\{v_1,\cdots,v_n\}$ be a basis for V, and let $w_1,\cdots,w_n$ be arbitrary elements of V. To show that every element of $\mathrm{End}_k(V)$ lies in R, it is enough to show that R contains an element mapping the v_i to the w_i; do this by showing that the map $\Phi: R \to V^n$ given by $\Phi(r) = (rv_1,\cdots,rv_n)$ is surjective. Here is an outline:

(ii) Use the fact that V^n is a semisimple R-module to show that if $\mathrm{Im}\,\Phi \neq V^n$, there is an R-homomorphism $\pi: V^n \to V$ mapping the image of Φ to $\{0\}$. Hint: Note that from the proof of exercise 10.4, we can take the complement of any submodule in V^n to be some V^k.

(iii) Show that any R-map $\pi: V^n \to V$ is of the form $\pi(\alpha_1,\cdots,\alpha_n) = \sum_{i=1}^{n} k_i\alpha_i$ (here $\alpha_i \in V$) for fixed $k_i \in k$, $1 \leq i \leq n$.

(iv) By examining $\Phi(1)$ conclude the result from (ii) and (iii), using the linear independence of the v_i.

Exercise 10.18: Use the previous exercise to give an alternate proof of exercise 9.8.

Bibliography

Bott, Raoul. [A paper on representations of compact Lie groups which I saw in preprint form in 1978. I am in the process of tracking this reference down.]

Serre, Jean-Pierre. Représentations linéaires des groupes finies. Paris: Hermann, 1971. English translation: Linear Representations of Finite Groups. New York: Springer-Verlag, 1977.