COMBINATORIAL REPRESENTATION THEORY

ROHAN PUTHUKUDY

Abstract. In this paper, we explore representation theory from a combinatorial perspective. In particular, we first review basic concepts from representation theory and then use these to develop an understanding of the representations of certain objects that have a combinatorial flavor to them. More specifically, we will determine the representations of the symmetric group $\mathfrak{S}_n$.

Contents

1. Basics of Representation Theory

1.1. Representations and Modules.

Our focus will be towards representations of finite groups as we will later apply this theory to the case of the symmetric group on n letters, which we denote by $\mathfrak{S}_n$. As mentioned earlier, representations allow us to view algebraic structures, like groups, under a more concrete lens. Namely, we can view their elements as linear transformations of a vector space. This act of "viewing" is captured formally in the following.

Definition 1.1. Suppose G is a finite group and V is a vector space over the field $\mathbb{C}$. Then a *linear representation* of G in V is a homomorphism

$$\rho\colon G \to \mathrm{GL}(V)$$

Moreover, we say that V is a *representation space* of G or *G-module* and together they form a *representation*, (ρ, V).

Note that because group homomorphisms from G to $\mathrm{GL}(V)$ and linear group actions of G on V are in one-to-one correspondence, we can define G-modules in general, to be vectors spaces over the field $\mathbb{C}$ that are acted on by G linearly.

Although we present representations as a tuple containing both the linear map and the G-module, we may abuse the term *representation* to refer either one of its components individually. This is because each component uniquely determines the other due to the general definition of G-modules.

In the above definition, $\mathrm{GL}(V)$ denotes the general linear group or set of linear transformations of V. In the course of this paper, we will assume that V is finite-dimensional. This is a useful assumption as in this case, each element of $\mathrm{GL}(V)$ can be expressed as an invertible $n \times n$ matrix with complex entries, where $n = \dim(V)$. Linear representations with this view of $\mathrm{GL}(V)$ are called *matrix representations*. Moreover, we say that the *degree* of the representation, denoted $\deg V$, is n.

Definition 1.2. Let G be a finite group. Suppose ϕ and ψ are linear representations of G with corresponding representation spaces V and W. Then ϕ and ψ are *isomorphic*, denoted $\phi \cong \psi$, if there exists a linear isomorphism $\tau \colon V \to W$ that satisfies

$$\tau \circ \phi(g) = \psi(g) \circ \tau$$

for all $g \in G$.

Note that if $\phi \cong \psi$ as defined above, then clearly $V \cong W$ as vector spaces. However, the other direction does not necessarily hold. This is because G-modules have additional structure beyond being vector spaces. So when we discuss the isomorphism of G-modules, we need to ensure a kind of compatibility between their G-actions. We will formally specify this when G-homomorphisms are introduced in the next section. Before we proceed, however, we will first discuss a few important examples of representations.

Example 1.2.1 (Trivial Representation)**.** Let V be any 1-dimensional vector space over $\mathbb{C}$. For example, take $V = \mathbb{C}$. Now let G be any finite group. Then let us define $1_G \colon G \to \mathrm{GL}(V)$ as the map $1_G(g) \coloneqq \mathrm{id}_V$ for every $g \in G$.

It follows that this is a linear representation of degree 1 and we refer to it as the *trivial representation of G*. Note that such a representation exists for any finite group and corresponds to trivially viewing the elements of the group as elements of $\mathbb{C}$, hence the name. $\square$

As we see in the following example, there are also non-trivial representations of degree 1.

Example 1.2.2 (Sign Representation)**.** Let $G = \mathfrak{S}_n$ and V be a 1-dimensional vector space over $\mathbb{C}$, such as $\mathbb{C}$. Note that $\mathrm{GL}(V) \cong \mathbb{C}$. So the sign map, $\mathrm{sgn} \colon \mathfrak{S}_n \to \mathbb{C}$, constitutes a linear representation of $\mathfrak{S}_n$.

To see why, first recall that if $\pi \in \mathfrak{S}_n$, then $\mathrm{sgn}(\pi) \coloneqq (-1)^k$ where π can be decomposed into k transpositions. Note that sgn is well-defined and so it follows that it is indeed a homomorphism.

Thus sgn is a degree 1 linear representation of $\mathfrak{S}_n$, which we call the *sign representation* $\square$

These two representations will be used later when we discuss characters and character tables.

Example 1.2.3 (Regular Representation)**.** Let G be any finite group and suppose that $k = |G|$. Then let V be a k-dimensional vector space over $\mathbb{C}$ and let $(\boldsymbol{v}_g)_{g \in G}$ be a basis of V, indexed by the elements of G. Then we can construct a linear representation $\rho \colon G \to \mathrm{GL}(V)$ as $\rho(g) \coloneqq \rho_g \in \mathrm{GL}(V)$ for each $g \in G$. It now suffices to define ρ_g for each g.

Let us fix some $g \in G$. Then note that because $\rho_g \in \mathrm{GL}(V)$, it suffices to define ρ_g's behavior on each $\boldsymbol{v}_h$. In particular, we define $\rho_g(\boldsymbol{v}_h) \coloneqq \boldsymbol{v}_{gh}$.

Hence ρ is a degree k linear representation of G, which we call the *regular representation*. $\square$

Notice that in defining the regular representation, we used the multiplicative structure of the group to permute the basis vectors of V. This is strongly connected to the fact that left multiplication by a fixed element of a group constitutes a group action of the group on itself, leading to the following generalization.

Example 1.2.4 (Permutation Representation)**.** Let G be any finite group. Moreover, let X be a finite set with $k = |X|$, such that G acts on X via the action $g \cdot x$. Now let V be a k-dimensional vector space over $\mathbb{C}$ and let $(\boldsymbol{v}_x)_{x\in X}$ be a basis of V indexed by X. Then we construct the linear representation $\rho\colon G \to \mathrm{GL}(V)$ as $\rho(g) \coloneqq \rho_g$ for each $g \in G$. It now suffices to define ρ_g for each g.

Let us fix $g \in G$. Then recall that it suffices to define ρ_g's behavior on each basis vector $\boldsymbol{v}_x$. In particular, we define $\rho_g(\boldsymbol{v}_x) \coloneqq \boldsymbol{v}_{g\cdot x}$.

Hence ρ is a degree k linear representation of G that generalizes the regular representation, and we call it the *permutation representation.* $\square$

Earlier we mentioned how we can equivalently view representations as either linear representations or as the spaces themselves. This raises a question concerning the substructure of representations. Since our representation spaces aren't just vector spaces, but are also associated with a homomorphism that maps the group to automorphisms of our space (i.e., the associated linear representation), do vector subspaces constitute "subrepresentations"? As our intuition may tell us, being a vector subspace of a representation is not sufficient to constitute a representation space, but instead we also require the notion of G-invariance.

Definition 1.3. Suppose $\rho\colon G \to \mathrm{GL}(V)$ is a linear representation and W is a subspace of V. Then we say that W is a *subrepresentation (space)* if it is *G-invariant*.

We say that W is G-invariant if it is invariant under the action of G, where the action is specified by ρ. Concretely, W is G-invariant if for all $\boldsymbol{w} \in W$, we have $\rho(g)(\boldsymbol{w}) \in W$ holds for all $g \in G$.

In this setting, the restriction of $\rho(g)$ to W, $\rho(g)^W$, belongs to $\mathrm{GL}(W)$ and $\rho(gh)^W = \rho(g)^W \cdot \rho(h)^W$. So $\rho^W : G \to \mathrm{GL}(W)$ is a linear representation of G in W making the term "subrepresentation" meaningful.

Example 1.3.1 (Trivial Subrepresentation)**.** Let (ρ, V) be a representation of a finite group G. Then consider the following subspace of V.

$$W \coloneqq \{\boldsymbol{v} \in V \mid \rho(g)(\boldsymbol{v}) = \boldsymbol{v}, \forall g \in G\}$$

Note that W is a subrepresentation of V because it is G-invariant, by definition. More explicitly, let us fix some $g \in G$ and $w \in W$. Then we have $\rho(g)(w) = w \in W$. Hence W is G-invariant.

Now before we proceed we will recall some important constructions from linear algebra.

Definition 1.4. If we are given two vector spaces, say V_1 and V_2, then we can form a new vector space called their *direct sum*, defined as

$$V = V_1 \oplus V_2 \coloneqq \{(\boldsymbol{v}_1, \boldsymbol{v}_2) \mid \boldsymbol{v}_1 \in V_1, \boldsymbol{v}_2 \in V_2\}$$

Moreover, we have a mapping that sends every vector $\boldsymbol{v} \in V$ to its component $\boldsymbol{v}_1 \in V_1$. This mapping is called the *projection* of V onto V_1. Importantly, the kernel of this map is V_2. There exists a similar projection onto V_2.

This construction lifts itself to representations in a natural way. For G-modules, V_1 and V_2, their direct sum is also a G-module by defining our G-action component-wise. In particular, if $(v_1, v_2) \in V_1 \oplus V_2$ and $g \in G$, then our action is defined to be

$$g \cdot (v_1, v_2) \coloneqq (g \cdot v_1, g \cdot v_2)$$

Hence, the direct sum, $V_1 \oplus V_2$, is also a G-module when V_1 and V_2 are G-modules. In this case, we say that V_1 and V_2 are *complements.*

Additionally, this G-module also has an associated linear representation. This is because every group action induces a homomorphism from the group to the group of transformations on the object acted upon. In this context, our group action from before induces a homomorphism $\rho\colon G \to \mathrm{GL}(V_1 \oplus V_2)$. This is a linear representation by definition.

This notion of a direct sum will prove pivotal in our study of representations. As we will see later, we can decompose any representation into a direct sum of finitely many subrepresentations, much like how we can factor positive integers into products of prime numbers. Before we can even state this result however, we will begin with a more modest claim. Namely, we see that we can split any representation into a direct sum if provided with a subrepresentation.

Theorem 1.5. *Suppose $\rho\colon G \to \mathrm{GL}(V)$ is a linear representation and let W be a G-invariant subspace of V. Then there exists a complement W^0 of W in V which is also G-invariant.*

Proof. There are actually two approaches to this proof. We will present one way involving averaging a projection map, but we remark that there is an alternate approach that uses a Hermitian inner product on V to construct an orthogonal complement of V.

We begin with an arbitrary complement [1], W', of W in V. Associated with this complement, is a projection $\pi\colon V \to W$. Consider the following averaged map.

$$\pi^0 \coloneqq \frac{1}{|G|} \sum_{g \in G} \rho(g) \circ \pi \circ \rho(g)^{-1}$$

First we note that π^0 is in fact a projection of V onto W. This is because $\rho(g)$ preserves W for all $g \in G$ by G-invariance and π is a projection. In particular, we have that $\pi \circ \rho(g)^{-1} = \rho(g)^{-1}$ for all $g \in G$, giving us $(\rho(g) \circ \pi \circ \rho(g)^{-1})(\boldsymbol{w}) = (\rho(g) \circ \rho(g)^{-1})(\boldsymbol{w}) = \boldsymbol{w}$ for all $g \in G$ and $\boldsymbol{w} \in W$.

Since π^0 is a projection, it induces a complement, say W^0, of W in V. Now it suffices for us to show that W^0 is G-invariant. To this end, let $\boldsymbol{w} \in W^0$ and $h \in G$. We then want to show that $\rho(h)(\boldsymbol{w}) \in W^0$. Recall that $W^0 = \ker \pi^0$. This means if suffices to show that $\pi^0(\rho(h)(\boldsymbol{w})) = 0$. It also tells us that $\pi^0(\boldsymbol{w}) = 0$. Before we use this fact, note that we have

[1] Recall from linear algebra that we can always construct a complement given any subspace and parent vector space through a recursive selection of linearly independent vectors until their span is a complement. An alternate argument using Zorn's lemma exists when we lack finite dimensionality, but we assume finite-dimensionality in this paper.

the following property of π^0 because any element of G permutes the rest of the group.

$$\begin{aligned}\rho(h) \circ \pi^0 \circ \rho(h)^{-1} &= \frac{1}{|G|} \sum_{g \in G} \rho(h) \circ \rho(g) \circ \pi \circ \rho(g)^{-1} \circ \rho(h)^{-1} \\ &= \frac{1}{|G|} \sum_{g \in G} \rho(hg) \circ \pi \circ \rho(hg)^{-1} \\ &= \frac{1}{|G|} \sum_{k \in G} \rho(k) \circ \pi \circ \rho(k)^{-1} \\ &= \pi^0\end{aligned}$$

This gives us our desired equality as follows.

$$(\pi^0 \circ \rho(h))(\boldsymbol{w}) = (\rho(h) \circ \pi^0)(\boldsymbol{w}) = \rho(h)(\boldsymbol{w}) = 0$$

Hence W^0 is a G-invariant complement of W in V and we are finished. □

1.2. Reducibility and Schur's Lemma.

We will now work towards the result mentioned earlier. We begin by defining our "building blocks" for representations.

Definition 1.6. Let $\rho\colon G \to \mathrm{GL}(V)$ be a linear representation. Then we say that it is *irreducible* if V is not 0 and the only G-invariant subspaces of V are 0 and V itself.

Example 1.6.1. Let V be a 1-dimensional representation. Then it is *trivially* irreducible as its only *subspaces*, let alone submodules, are 0 and V itself.

Note that by Theorem 2.5, a G-module, V, is irreducible precisely when it cannot be written as a direct sum of subrepresentations. This leads us to our first major result in the representation theory of finite groups.

Theorem 1.7 (Maschke's Theorem)**.** *Every representation is the direct sum of irreducible representations.*

Proof. Let $\rho\colon G \to \mathrm{GL}(V)$ be a linear representation of a finite group G and let $n = \dim(V)$. We will prove the theorem via induction on n.

Base Case: Let $n = 1$. Then V is trivially irreducible as the only subspaces, let alone G-invariant subspaces, of V are 0 and V itself. So we are finished as $V = W_1 \oplus \cdots \oplus W_k$ with $k = 1$ and $W_1 = V$.

Induction Step: Suppose that the theorem holds for all G-modules of dimension less than n. Now it suffices to show that it holds for an arbitrary G-module of dimension n. Let V be such a G-module as above. Then if V is itself irreducible, then we are finished. So suppose that V is not irreducible. This implies that there exists some subrepresentation W of V that is not 0, nor V. So by Theorem 2.5, we have that $V = W \oplus W^0$. Moreover, we have that $0 < \dim W, \dim W^0 < n$. So by the inductive hypothesis, there exists irreducible G-modules $W_1, \ldots, W_k, W_1', \ldots, W_\ell'$ such that $W = W_1 \oplus \cdots \oplus W_k$ and $W^0 = W_1' \oplus \cdots \oplus W_\ell'$. Hence we obtain

$$V = W_1 \oplus \cdots \oplus W_k \oplus W_1' \oplus \cdots \oplus W_\ell'$$

so V is the direct sum of irreducible representations of G. □

This is a crucial result in the study of representations as the task of understanding all the representations of a given group can now be simplified into understanding all of its irreducible representations. This is considerably simpler for two reasons. First, irreducible representations are more specific due to their property of being irreducible. Second, every group has finitely many irreducible representations.

We remark that this result holds due to some of our underlying assumptions. Namely, that our groups are finite and our vector spaces are over $\mathbb{C}$. This result can be generalized to a broader scope, but it does not hold in full generality.

A natural question that may arise now is "Is this decomposition unique?" We will see that the answer to this is indeed yes[2], but in order to answer this question, we will finally formally define what it means for two G-modules to be isomorphic.

Definition 1.8. Let V and W be representations of a finite group G. Then a *G-homomorphism* is a linear transformation $\theta\colon V \to W$ such that

$$\theta(g \cdot \boldsymbol{v}) = g \cdot \theta(\boldsymbol{v})$$

for all $g \in G$ and $\boldsymbol{v} \in V$. Moreover, we say that θ is a *G-isomorphism* if it is bijective and can then write $V \cong W$.

Now that we have defined isomorphisms of representations, we present the following result which helps us better understand when two representations are isomorphic.

Lemma 1.9 (Schur's Lemma)**.** *Let V and W be two irreducible representations of a finite group G. If $\theta\colon V \to W$ is a G-homomorphism, then either*

(1) θ is a G-isomorphism
(2) θ is the zero map

This is a crucial result because it provides us with uniqueness (up to isomorphism and with multiplicity) for the decompositions obtained via Maschke's Theorem (Theorem 2.7). More specifically, given a G-module V, there exists distinct, irreducible G-modules $V_1, \ldots, V_k$ and multiplicities $m_1, \ldots, m_k$ such that $V = V_1^{\oplus m_1} \oplus \cdots \oplus V_k^{\oplus m_k}$ and this decomposition is unique. With that said, the proof of Schur's lemma is as follows.

Proof. Let V and W be irreducible G-modules and let $\theta\colon V \to W$ be a G-homomorphism. We will first argue that $\ker\theta$ is a submodule of V and $\operatorname{Im}\theta$ is a submodule of W. This is sufficient to show our claim, as the irreducibility of V and W means that this implies that $\ker\theta$ is either 0 or V. So θ is either the zero map or injective. Similarly, we have that $\operatorname{Im}\theta$ is either 0 or W. So θ is either the zero map or surjective. Combining these, we obtain that θ is either the zero map or an isomorphism.

We will begin by showing that $\ker\theta$ is a submodule of V. Note that it is a subspace as θ is a linear transformation, so it suffices to show that it is G-invariant. To that end, let $\boldsymbol{v} \in \ker\theta$. Then for all $g \in G$, we have $\theta(g \cdot \boldsymbol{v})) = g \cdot (\theta(\boldsymbol{v})) = g \cdot 0 = 0$. Hence $g \cdot \boldsymbol{v} \in \ker\theta$. So $\ker\theta$ is indeed a submodule of V.

Now we will consider $\operatorname{Im}\theta$. Let $\boldsymbol{w} \in \operatorname{Im}\theta$ and $g \in G$. Let us pick some $\boldsymbol{v} \in V$ such that $\boldsymbol{w} = \theta(\boldsymbol{v})$. Then by the definition of G-homomorphisms, we have

$$g \cdot \boldsymbol{w} = g \cdot \theta(\boldsymbol{v}) = \theta(g \cdot \boldsymbol{v}) \in \operatorname{Im}\theta$$

[2]The decompositions are not unique in a strict sense, but they are unique up to isomorphism and with certain multiplicities.

Hence Im θ is G-invariant and thus a submodule of W as it is clearly also a subspace of W. □

1.3. Characters.

As we now understand that representations have unique decompositions into irreducible parts, we would like to be able to find these decompositions when given a representation. More specifically, if we are given a G-module, V, we would like to determine the distinct irreducible subrepresentations $W_1, \ldots, W_k$ of V such that $V = W_1^{\oplus m_1} \oplus \cdots \oplus W_k^{\oplus m_k}$. In order to do this, we will need to develop tools to answer the following questions.

- How many distinct irreducible subrepresentations appear in a given representation's decomposition? (Determine k)
- Can we develop a criterion for determining if a subrepresentation is irreducible? (Determine $W_1, \ldots, W_k$)
- If we are given a representation, say V, and one of its irreducible subrepresentations, say W, how many times does W occur in V? (Determine $m_1, \ldots, m_k$)

In order to answer these questions, we will shift our focus towards the *trace* of linear representations, as we will see that this characterizes a representation up to isomorphism and the theory of these so-called *characters* will contain our the tools to answer our desired questions.

First, we recall the notion of a trace from linear algebra.

Definition 1.10. Let V be a finite dimensional vector space over $\mathbb{C}$ and $a\colon V \to V$ be a linear transformation with matrix form (a_{ij}). Then we define the *trace* of a to be the following.

$$\operatorname{Tr}(a) \coloneqq \sum_i a_{ii}$$

In addition to this definition, we can also view the $\operatorname{Tr}(a)$ as the sum of the eigenvalues of a, due to the existence of the Jordan canonical form.

Now recall that we assume that our G-modules are finite dimensional so we can extend the notion of trace to linear representations, motivating the following definition.

Definition 1.11. Let $\rho\colon G \to \mathrm{GL}(V)$ be a linear representation. Then we define the *character* of ρ to be the function $\chi_\rho\colon G \to \mathbb{C}$ defined by

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

for each $g \in G$.

Alongside this definition, we let characters inherit much of the terminology from their associated representation. I.e., irreducible characters are characters of irreducible representations.

Immediately from this definition, we already obtain some important properties of characters.

Proposition 1.12. *Let V be a representation of a finite group G, of degree n and with character χ. Then*

(1) $\chi(\epsilon) = n$.
(2) $\chi(g^{-1}) = \overline{\chi(g)}$.
(3) $\chi(hgh^{-1}) = \chi(g)$.

(4) If K is a conjugacy class of G, then

$$g, h \in K \implies \chi(g) = \chi(h)$$

where $\overline{z}$ is the complex conjugate of $z \in \mathbb{C}$.

Proof. Let V be a G-module with character χ, where $\dim V = n$. Moreover, let $\rho \colon G \to \mathrm{GL}(V)$ be the linear representation associated with V.

(1) Note that $\rho(\epsilon)(\boldsymbol{v}) = \boldsymbol{v}$ for all $\boldsymbol{v} \in V$, so the matrix form of $\rho(\epsilon)$ is the identity matrix I_n. It follows from this that $\chi(\epsilon) = \mathrm{Tr}(I_n) = n$.

(2) Let $g \in G$ and let us denote $\rho(g)$ as ρ_g. Note that ρ_g has finite order because g has finite order. Thus, so do its eigenvalues, $\lambda_1, \dots, \lambda_n$. Hence, they must be roots of unity, implying that $1 = |\lambda_i| = \lambda_i \overline{\lambda_i}$. So we have

$$\overline{\chi(g)} = \overline{\mathrm{Tr}(\rho_g)} = \sum \overline{\lambda_i} = \sum \lambda_i^{-1} = \mathrm{Tr}(\rho_g^{-1}) = \mathrm{Tr}(\rho_{g^{-1}}) = \chi(g^{-1})$$

(3) We have this property due to the fact that $\rho(h^{-1}) = \rho(h)^{-1}$ and $\mathrm{Tr}(AB) = \mathrm{Tr}(BA)$ as follows.

$$\chi(hgh^{-1}) = \mathrm{Tr}(\rho(h)\rho(g)\rho(h)^{-1}) = \mathrm{Tr}(\rho(g)) = \chi(g)$$

(4) This follows from part (3).

□

Property (4) is can be generalized to arbitrary functions on groups as illustrated in the following definition.

Definition 1.13. A *class function* on a group G is a function $f \colon G \to \mathbb{C}$ that is constant over conjugacy classes. I.e., $g \equiv h \implies f(g) = f(h)$. We denote the set of all class functions on G by $R(G)$ and note that this is actually a vector space over $\mathbb{C}$ with dimension equal to the number of conjugacy classes of G.

Since characters are class functions and finite groups have finitely many conjugacy classes, we can completely describe a character by its value on each conjugacy class. We typically depict this information in a tabular format, ranging over every irreducible character of a group.

Definition 1.14. Let G be a group. Then the *character table of* G is a table containing the values of each irreducible character of G's value on each of G's conjugacy classes, denoted χ_K where χ is an irreducible character and K is a conjugacy class of G.

$$\begin{array}{c|ccc} & \cdots & K & \cdots \\ \hline \vdots & & \vdots & \\ \chi & \cdots & \chi_K & \cdots \\ \vdots & & \vdots & \end{array}$$

Computing character tables by hand is indeed possible, see Sagan. There are also computational techniques for determining the character tables for small size groups (they also work for large groups, but will take longer).

Example 1.14.1. Here we use GAP to compute the character table of $\mathfrak{S}_3$. GAP contains a character table library (ctbllib) which can compute irreducible character tables using the Dixon-Schneider algorithm.

```
gap> LoadPackage("ctbllib");
    true
gap> S3:= SymmetricGroup(3);;  SetName(S3, "S3");
gap> Irr(CharacterTable(S3));
    [ Character( CharacterTable( S3 ), [ 1, -1, 1 ] ),
      Character( CharacterTable( S3 ), [ 2, 0, -1 ] ),
      Character( CharacterTable( S3 ), [ 1, 1, 1 ] ) ]
```

This tell us that the character table of $\mathfrak{S}_3$ is

	K_1	K_2	K_3
χ_1	1	-1	1
χ_2	1	1	1
χ_3	2	0	-1

Remark: When displaying the character table, observe that GAP does not tell us which conjugacy classes each of the characters' values corresponds to. This is because in performing this computation, GAP does not necessarily determine conjugacy classes explicitly. However, we can ask GAP for the conjugacy classes of a group as follows.

```
gap> ConjugacyClasses(S3);
    [ ()^G, (1,2)^G, (1,2,3)^G ]
```

Note: the order of these conjugacy classes does not necessarily correspond to their order in the character table from earlier. However, the values of the character table uniquely determine both the characters themselves, as well as which conjugacy class corresponds to each K_i for each i.

In particular, we have that χ_2 is the trivial representation (Example 2.2.1), χ_1 is the sign representation and χ_3 is the regular representation.

In the above example and in the definition of character tables, we focused on irreducible characters. It turns out we can use Maschke's theorem to determine the character of a reducible representation because the characters of direct sums behave naturally.

Proposition 1.15. *Let G be a finite group with representations V_1 and V_2, and let χ_1 and χ_2 be their respective characters. Then*

(1) The character χ of the direct sum $V_1 \oplus V_2$ is equal to $\chi_1 + \chi_2$.
(2) The character χ of the tensor product $V_1 \otimes V_2$ is equal to $\chi_1 \cdot \chi_2$.

Proof. Let $\rho_1 \colon G \to \mathrm{GL}(V_1)$ and $\rho_2 \colon G \to \mathrm{GL}(V_2)$ be two finite dimensional linear representations of G. Since ρ_1 and ρ_2 are finite dimensional, for each $g \in G$, we can associate $\rho_i(g)$ with a matrix M_i^g.

(1) Note we have that the representation $\rho_1 \oplus \rho_2$ is associated with the following matrix, for each $g \in G$.

$$\begin{bmatrix} M_1^g & 0 \\ 0 & M_2^g \end{bmatrix}$$

From this, we directly obtain

$$\begin{aligned}\chi(g) &= \mathrm{Tr}((\rho_1 \oplus \rho_2)(g)) \\ &= \mathrm{Tr}(M_1^g) + \mathrm{Tr}(M_2^g) \\ &= \mathrm{Tr}(\rho_1(g)) + \mathrm{Tr}(\rho_2(g)) \\ &= \chi_1(g) + \chi_2(g)\end{aligned}$$

(2) By definition we have that

$$\chi_1(g) = \sum_i m_{1,i}^g \text{ and } \chi_2(g) = \sum_i m_{2,i}^g$$

where $m_{j,k}^g$ denotes the k-th diagonal entry of M_j^g. Then we have

$$\chi(g) = \sum_{i,j} m_{1,i}^g \cdot m_{2,j}^g = \chi_1(g) \cdot \chi_2(g)$$

from the definition of tensor products.

□

We will now define an inner and scalar product for class functions and see that they are actually equivalent when we restrict to characters. This product will be important as it will be used to define a criterion for the irreducibility of characters and consequently of representations.

Definition 1.16. Let G be a finite group. We define the *inner product* of $R(G)$ to be

$$\langle \phi, \psi \rangle \coloneqq \frac{1}{|G|} \sum_{g \in G} \phi(g)\psi(g^{-1})$$

where $\phi, \psi \in R(G)$.

Note that we have $\langle \phi, \psi \rangle = \langle \psi, \phi \rangle$ and this product is linear in both ϕ and ψ. This product is equivalent to the following orthogonality relation for characters.

Definition 1.17. Let G be a finite group. We can then define the following *scalar product* of $R(G)$.

$$(\phi|\psi) \coloneqq \frac{1}{|G|} \sum_{g \in G} \phi(g)\overline{\psi(g)}$$

where $\phi, \psi \in R(G)$.

The equivalence of $\langle \phi, \psi \rangle$ and $(\phi|\psi)$ comes from Proposition 2.12 part (2). Namely, that $\psi(g^{-1}) = \overline{\psi(g)}$ when ψ is a character.

Proposition 1.18. *Let G be a finite group and $\phi, \psi \in R(G)$. Then let $\hat{\psi}$ be defined by $\hat{\psi}(g) = \overline{\psi(g^{-1})}$. Then we have*

$$(\phi|\psi) = \frac{1}{|G|} \sum_{g \in G} \phi(g)\hat{\psi}(g^{-1}) = \langle \phi, \hat{\psi} \rangle$$

Proof. □

In particular, if ϕ and ψ are characters, then we have $\psi = \hat{\psi}$ by Proposition 2.12. So it follows that $\langle \phi, \psi \rangle = (\phi|\psi)$ in this setting.

Before we proceed, we will present an important corollary of Schur's lemma in the context of character theory.

Corollary 1.19. *Let $\rho_1 \colon G \to \mathrm{GL}(V_1)$ and $\rho_2 \colon G \to \mathrm{GL}(V_2)$ be two linear representations of a group G, let h be a linear mapping of V_1 into V_2, and*

$$h^0 := \frac{1}{|G|} \sum_{g \in G} \rho_2(g)^{-1} h \rho_1(g)$$

Then we have

(1) If $\rho_1 \not\simeq \rho_2$, then $h^0 = 0$
(2) If $\rho_1 \simeq \rho_2$, then $h^0 = (1/\dim(V_1))\mathrm{Tr}(h) id$.

Proof. Note that $\rho_2(s)h^0 = h^0\rho_1(s)$ for any $s \in G$ as

$$\begin{aligned} \rho_2(s)^{-1}h^0\rho_1(s) &= \frac{1}{|G|} \sum_{g \in G} \rho_2(s)^{-1}\rho_2(g)^{-1}h\rho_1(g)\rho_1(s) \\ &= \frac{1}{|G|} \sum_{g \in G} \rho_2(gs)^{-1}h\rho_1(gs) \\ &= h^0 \end{aligned}$$

By Schur's lemma applied to h^0, if $\rho_1 \not\simeq \rho_2$, then h^0 must be the zero map. Otherwise, it will have to be some scalar multiple of the identity, say λ. In this latter case, we have

$$\mathrm{Tr}(h^0) = \frac{1}{|G|} \sum_{g \in G} \mathrm{Tr}(\rho_2^{-1} h \rho_1) = \mathrm{Tr}(h)$$

Note that $\mathrm{Tr}(\lambda I_n) = n \cdot \lambda$, so we obtain $\lambda = \frac{1}{n}\mathrm{Tr}(h)$, where $n = \dim(V_1)$. □

Now we can use this corollary to prove that our scalar product $(\phi|\psi)$ is indeed an orthogonality relation for characters.

Lemma 1.20. *Let ϕ and ψ be irreducible characters of a finite group G. Then*

$$(\phi|\psi) = \delta_{\phi,\psi}$$

Proof. Let ϕ and ψ be two irreducible characters of a finite group G. First suppose that they are equal and our underlying representation is associated to the matrix M^g with entries $m_{i,j}$ for each $g \in G$. Then we have $\phi(g) = \psi(g) = \sum_i m_{i,i}(g)$. This implies that

$$(\phi|\psi) = (\phi|\phi) = \langle \phi, \phi \rangle = \sum_{i,j} \langle m_{i,i}, m_{j,j} \rangle$$

By Corollary 2.19, with the linear representations being replaced by their matrix forms, we have that $\langle m_{i,i}, m_{j,j} \rangle = \delta_{i,j}/n$ where n is the degree of the representation associated with ϕ. So we have

$$(\phi|\psi) = (\phi|\phi) = 1$$

In the case where ϕ and ψ are not equal, we can apply Corollary 2.19 in a similar way to obtain that $(\phi|\psi) = 0$, completing the proof. □

We can now use what we have developed in character theory to answer one of our representation theory questions. In particular, we are now able to determine the "multiplicities" of irreducible subrepresentations in the decompostions given by Maschke's theorem.

Lemma 1.21. *Let V be a representation of a finite group G, with character ϕ. Suppose that V decomposes into the following direct sum of irreducible representations.*

$$V = W_1 \oplus \cdots \oplus W_k$$

Moreover, let W be an irreducible representation of G with character χ. Then the number of W_i's that are isomorphic to W is equal to $(\phi|\chi)$. We call this number "the number of times that W occurs in V"

Proof. Note that

$$\phi = \chi_1 + \cdots + \chi_k$$

by Proposition 2.15. Thus we have $(\phi|\chi) = (\chi_1|\chi) + \cdots + (\chi_k|\chi)$. Finally, by the previous lemma, we know that $(\chi_i|\chi) = 1$ precisely when $W_i \simeq W$. □

This lemma helps justify the name of *characters* for these class functions as we can now prove that they do indeed characterize representations.

Corollary 1.22. *Two representations with the same character are isomorphic*

Proof. By Lemma 2.21, we know that if two representations have the same character, then they contain the same number of copies of W, for all irreducible representations W. Hence, they must have the same decomposition, up to isomorphism, and thus are isomorphic. □

Finally, we obtain the following irreducibility criterion for representations.

Corollary 1.23. *Let V be a representation of a finite group G, with character χ. Then V is irreducible if and only if $(\chi|\chi) = 1$.*

Proof. Suppose $V = a_1W_1 \oplus \dots a_kW_k$, where the W_i are irreducible subrepresentations with characters χ_i. Then we have $\chi = \sum_i a_i\chi_i$. Recall from Lemma 2.21 that $a_i = (\chi|\chi_i)$. This implies

$$(\chi|\chi) = \sum_i m_i^2$$

Note that this is only equal to 1 in the case that one of the m_i's equals 1 and all the others equal 0. In other words, $(\chi|\chi) = 1$ precisely when $V \simeq W_i$ for some i and is thus irreducible. □

2. The Symmetric Group

We will now begin the combinatorial part of this paper, by focusing our attention on the symmetric group. Specifically, we want to fully understand the representations of the symmetric group. As we have seen through Maschke's Theorem (Theorem 2.7) and character theory, this amounts to understanding the irreducible representations of $\mathfrak{S}_n$.

Recall that the number of irreducible representations is precisely the number of conjugacy classes of the group. Moreover, for $\mathfrak{S}_n$ this is just the number of partitions of n. This observation motivates the approach that we take in this chapter. Specifically, we will try to associate an irreducible representation of $\mathfrak{S}_n$ with each partition of n. To do this, we will first construct subgroups for each partition and then construct a representation of $\mathfrak{S}_n$ from one on each of these subgroups.

2.1. Young Tableaux, Subgroups, and Tabloids.

We begin by recalling the definition of a partition of an integer.

Definition 2.1. Suppose $\lambda = (\lambda_1, \dots, \lambda_k) \in \mathbb{N}^k$ satisfies $\lambda_1 + \cdots + \lambda_k = n$ and $\lambda_1 \geq \cdots \geq \lambda_k$, then we say that λ is a *partition* of n, written $\lambda \vdash n$. Moreover, we use the notation $|\lambda| \coloneqq \sum_{i=1}^{k} \lambda_i$.

Example 2.1.1. Let $n = 3$. Then all of the partitions of n are listed here:

- $n = 1 + 1 + 1 \implies (1, 1, 1) \vdash n$
- $n = 2 + 1 \implies (2, 1) \vdash n$
- $n = 3 \implies (3) \vdash n$

Note that in general, we have that $(n) \vdash n$ and $(1^n) \coloneqq (1, \dots, 1) \vdash n$ where we use (a^n) to denote the tuple of n copies of a.

There is a nice way to visualize each of these partitions called a Young diagram.

Definition 2.2. Suppose $\lambda = (\lambda_1, \dots, \lambda_k) \vdash n$. Then the *Young diagram*, or *shape*, of λ is an array of boxes with k left-justified rows each containing λ_i boxes for $1 \leq i \leq k$.

Example 2.2.1. Let $n = 3$ as in Example 3.1.1. Then the Young diagrams associated with the partitions from earlier are below in the same order.

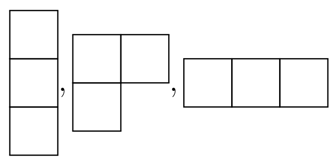

As mentioned earlier, we want to associate a subgroup of $\mathfrak{S}_n$ with each of these partitions. We typically view $\mathfrak{S}_n$ as the group of permutations of a set of n elements. This motivates the notation $\mathfrak{S}_A \coloneqq \mathfrak{S}_{|A|}$ for sets A.

Since a partition of n corresponds to a splitting of this implicit set, it is naturally associated to the following subgroup of $\mathfrak{S}_n$

Definition 2.3. Suppose $\lambda = (\lambda_1, \dots, \lambda_k) \vdash n$. Then we define the *Young subgroup* of $\mathfrak{S}_n$, associated with λ, to be the following.

$$S_\lambda \coloneqq \mathfrak{S}_{\{1,\dots,\lambda_1\}} \times \mathfrak{S}_{\{\lambda_1+1,\dots,\lambda_1+\lambda_2\}} \times \cdots \times \mathfrak{S}_{\{n-\lambda_k+1,\dots,n\}}$$

Note that

$$S_\lambda \cong \mathfrak{S}_{\lambda_1} \times \cdots \times \mathfrak{S}_{\lambda_k}$$

Now that we have defined this subgroup, we will construct a representation of $\mathfrak{S}_n$ associated with this partition. In representation theoretic-terms, we will be inducting on the trivial representation of S_λ to obtain a representation of $\mathfrak{S}_n$. However, here we present this from a combinatorial perspective by directly defining this representation from a partition.

First, we define a class of objects that are obtained from a given partition that will be used to generate our representation.

Definition 2.4. Suppose $\lambda \vdash n$. Then we define a *Young tableau of shape* λ, or a λ-*tableau*, to be a Young diagram where the numbers $1, \dots, n$ are written into the boxes bijectively. Moreover we denote one as t^λ where $t_{i,j}$ is the entry of t in position (i, j).

Example 2.4.1. Let $\lambda = (3, 1) \vdash 4 = n$. Then there are $n! = 4! = 24$ total Young tableaux of shape $\lambda = (3, 1)$. Here are a few.

$$t_1 = \begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline \multicolumn{1}{|c|}{4} \\ \cline{1-1}\end{array}, t_2 = \begin{array}{|c|c|c|}\hline 3 & 1 & 2 \\ \hline \multicolumn{1}{|c|}{4} \\ \cline{1-1}\end{array}, t_3 = \begin{array}{|c|c|c|}\hline 4 & 2 & 3 \\ \hline \multicolumn{1}{|c|}{1} \\ \cline{1-1}\end{array}$$

Note that in the above example, t_1 and t_2 have the same rows, just with their entries permuted within them. Viewing tableaux in this way actually constitutes an equivalence relation, leading us to the following definition.

Definition 2.5. Two λ-tableau, t_1 and t_2, are *row-equivalent*, denoted $t_1 \sim t_2$, if each of the corresponding rows share the same set of elements. A *tabloid of shape* λ, or λ-*tabloid* is then defined to be

$$\{t\} := \{t' \mid t' \sim t\}$$

where t has shape λ.

Example 2.5.1. As mentioned earlier, we have

$$\begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline \multicolumn{1}{|c|}{4} \\ \cline{1-1}\end{array} \sim \begin{array}{|c|c|c|}\hline 3 & 1 & 2 \\ \hline \multicolumn{1}{|c|}{4} \\ \cline{1-1}\end{array}$$

So these two tableaux belong to the same tabloid. Moreover, we know that this tabloid contains $3! = 6$ elements as each element is just a permutation of the top row. Visually, we can draw this tabloid as

$$\begin{array}{ccc}\hline 1 & 2 & 3 \\ \hline 4 \\ \cline{1-1}\end{array}$$

When depicting tabloids in this way, we typically order the entries of rows in increasing order as a convention.

We remark that a similar notion can be developed for column equivalence, but the two perspectives are ultimately the same as we can simply rotate column-equivalent tableaux by $90°$ clock-wise and reflect it by the vertical axis to obtain row-equivalent tableaux over the same partition.

In the above example, we saw that the tabloid corresponding to $\lambda = (3, 1)$ contained $3! = 6$ elements. This idea can be generalized to arbitrary tabloids. Suppose we have a partition $\lambda = (\lambda_1, \dots, \lambda_k) \vdash n$. Then every arbitrary λ-tabloid contains $\lambda_1! \cdots \lambda_k!$ many elements. Since there are $n!$ total λ-tableaux of n, this implies that there are $n!/(\lambda_1! \cdots \lambda_k!)$ λ-tabloids.

Now we make a key observation about tabloids and the symmetric group. Suppose $\lambda \vdash n$. Notice that since every λ-tableau contains the integers from 1 to n in specific entries of the Young diagram corresponding to λ, we can let $\mathfrak{S}_n$ act on the set of λ-tableaux by permuting the location of these integers in the λ-tableau's Young diagram.

Example 2.5.2. As an example, let $\lambda = (3, 1) \vdash 4 = n$ and note that $(123) \in \mathfrak{S}_n$. By the group action defined above, we have

$$(123) \cdot \begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline \multicolumn{1}{|c|}{4} \\ \cline{1-1}\end{array} = \begin{array}{|c|c|c|}\hline 2 & 3 & 1 \\ \hline \multicolumn{1}{|c|}{4} \\ \cline{1-1}\end{array}$$

Note that (123) preserves the tabloid

$$t = \begin{array}{ccc} \hline 1 & 2 & 3 \\ \hline 4 & & \\ \cline{1-1} \end{array}$$

However, (14) does not. Notice, however, that (14) sends every element of t to an element of

$$t' = \begin{array}{ccc} \hline 2 & 3 & 4 \\ \hline 1 & & \\ \cline{1-1} \end{array}$$

In general, we have for all $g \in \mathfrak{S}_n$ that $g \cdot t_1 \sim g \cdot t_2$ whenever $t_1 \sim t_2$. So we can extend our group action from the level of tableaux to the level of tabloids in a natural way.

This new group action of $\mathfrak{S}_n$ on the set of *tabloids* gives rise to a $\mathfrak{S}_n$-module (as in Example 2.2.4) which we define formally as follows.

Definition 2.6. Suppose $\lambda \vdash n$. Then we define the *permutation module corresponding to* λ as

$$M^\lambda := \mathbb{C}\{\boldsymbol{\{t_1\}}, \ldots, \boldsymbol{\{t_k\}}\}$$

where $\{t_1\}, \ldots, \{t_k\}$ are all the λ-tabloids.

Note: Recall that we use the boldface to denote vectors as here we are defining a k-dimensional vector space over $\mathbb{C}$.

Example 2.6.1. Let us take $\lambda = (3, 1) \vdash 4 = n$ as we have been doing. Then we should have that M^λ is a representation of $\mathfrak{S}_3$. First recall that there are $4!/(3!1!) = 4$ total λ-tabloids. By the nature of λ we can quickly observe that each tabloid simply corresponds to which number is in the bottom most row. Let t_i denote the tabloid with i in the bottom-most row for $i = 1, \ldots, 4$. More explicitly, we have

$$\boldsymbol{\{t_1\}} = \begin{array}{ccc} \hline 2 & 3 & 4 \\ \hline 1 & & \\ \cline{1-1} \end{array}, \boldsymbol{\{t_2\}} = \begin{array}{ccc} \hline 1 & 3 & 4 \\ \hline 2 & & \\ \cline{1-1} \end{array}, \boldsymbol{\{t_3\}} = \begin{array}{ccc} \hline 1 & 2 & 4 \\ \hline 3 & & \\ \cline{1-1} \end{array}, \boldsymbol{\{t_4\}} = \begin{array}{ccc} \hline 1 & 2 & 3 \\ \hline 4 & & \\ \cline{1-1} \end{array}$$

So our definition from above tells us that

$$M^\lambda = \mathbb{C}\left\{\boldsymbol{\{t_1\}}, \boldsymbol{\{t_2\}}, \boldsymbol{\{t_3\}}, \boldsymbol{\{t_4\}}\right\} \cong \mathbb{C}^4$$

This is indeed a representation of $\mathfrak{S}_4$ as $\mathfrak{S}_4$ acts on it linearly by the definition of a permutation representation.

For a concrete example of this group action, consider $\boldsymbol{v} = 3\boldsymbol{\{t_1\}} - 2\boldsymbol{\{t_3\}} \in M^\lambda$ and let $g = (123) \in \mathfrak{S}_4$. Then we have

$$\begin{aligned} g \cdot \boldsymbol{v} &= g \cdot (3\boldsymbol{\{t_1\}} - 2\boldsymbol{\{t_3\}}) \\ &= 3(g \cdot \boldsymbol{\{t_1\}}) - 2(g \cdot \boldsymbol{\{t_3\}}) \\ &= 3\boldsymbol{\{t_2\}} - 2\boldsymbol{\{t_1\}} \in M^\lambda \end{aligned}$$

An important property of these permutation modules is that they are cyclic.

Definition 2.7. A G-module, M, is said to be *cyclic* if there exists $\boldsymbol{v} \in M$ such that

$$M = \mathbb{C}(G\boldsymbol{v})$$

where $G\boldsymbol{v} = \{g\boldsymbol{v} \mid g \in G\}$. In this case, we say that M is generated by $\boldsymbol{v}$.

For a permutation module M^λ, any of the tabloids will suffice as a candidate for this $\boldsymbol{v}$. This is because for any tabloids $\{\boldsymbol{t_1}\}$ and $\{\boldsymbol{t_2}\}$, there exists some $g \in \mathfrak{S}_n$ such that $\{\boldsymbol{t_1}\} = g \cdot \{\boldsymbol{t_2}\}$.

At the beginning of this chapter, we stated that we are interested in producing representations of the symmetric group by inducting on the trivial representation of Young subgroups, S_λ. In constructing these M^λ, we have actually arrived at these representations in an alternate way. We capture this formally via the following.

Lemma 2.8. *Let $V^\lambda = 1 \uparrow_{S_\lambda}^{\mathfrak{S}_n}$. Then we have that $M^\lambda \cong V^\lambda$.*

Proof. Let $\{\pi_1, \dots, \pi_k\}$ be a *transversal* for S_λ. I.e., we have

$$\mathfrak{S}_n = \dot{\bigcup_i} \pi_i S_\lambda$$

We can now define a map $\theta \colon V^\lambda \to M^\lambda$ on each element of the transversal as $\theta(\pi_i S_\lambda) := \{\pi_i t^\lambda\}$ for each $i = 1, \dots, k$. Moreover, suppose θ extends linearly. Then one can verify that θ is indeed a $\mathfrak{S}_n$ isomorphism. $\square$

2.2. Ordering Tableaux.

Now that we have representations associated with each partition of n, it would be ideal if these were all irreducible. This, however, is not always the case. Our goal in the subsequent sections is to use these M^λ to arrive at the set of irreducible representations in the following way.

First we will order the partitions of n as $\lambda_1, \lambda_2, \dots,$ such that we obtain the following property. We will have that M^{λ_1} is irreducible and we denote it as S^{λ_1}. Then for all $i > 1$, we have that M^{λ_i} will be decomposable into a direct sum of S^{λ_j}'s for $1 \leq j < i$ and a new irreducible module which we will denote by S^{λ_i}.

Once we have this ordering, then we have that the irreducible representations of $\mathfrak{S}_n$ are simply the set of all of these S^{λ_i}.

In this section, we focus on defining this ordering and in the later sections we will describe the ireducibles more explicitly.

We begin with an ordering of partitions.

Definition 2.9. Let $\lambda = (\lambda_1, \dots, \lambda_k)$ and $\mu = (\mu_1, \dots, \mu_\ell)$ be two partitions of n. Then λ *dominates* μ, written as $\lambda \trianglerighteq \mu$, if

$$\lambda_1 + \cdots + \lambda_i \geq \mu_1 + \cdots + \mu_i$$

for all i. Note: if $i > k$, then we define $\lambda_i = 0$. Similarly, if $i > \ell$, then $\mu_i = 0$.

Example 2.9.1. Let $\lambda = (3, 1), \mu = (2, 2), \nu = (2, 1, 1)$ be partitions of n. Then visually, we have

$\lambda =$ 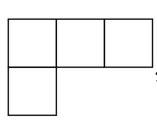, $\mu =$ 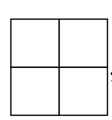, $\nu =$ 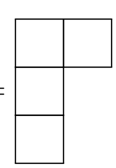

Note that $\lambda_1 = 3 \geq 2 = \mu_1$ and $\lambda_1 + \lambda_2 = 4 \geq 4 = \mu_1 + \mu_2$. So $\lambda \trianglerighteq \mu$. Similarly, $\mu_1 = 2 \geq 2 = \nu_1$, $\mu_1 + \mu_2 = 4 \geq 3 = \nu_1 + \nu_2$, and $\mu_1 + \mu_2 + \mu_3 = 4 \geq 4 = \nu_1 + \nu_2 + \nu_3$. So $\mu \trianglerighteq \nu$.

Note that $\lambda \trianglerighteq \nu$ because the transitivity of $\geq$ and the way we defined $\trianglerighteq$ imply that $\trianglerighteq$ is also transitive.

So visually, we have the following ordering.

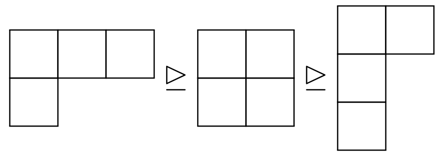

Observe that the wider and shorter Young diagrams appear to dominate the skinnier and taller ones. This is a general phenomenon that captures the nature of this ordering in a visual way.

Now that we have defined this ordering for partitions, we can establish an important connection between tableaux and dominance.

Lemma 2.10 (Dominance Lemma)**.** *Let t^λ and s^μ be tableaux. If, for each index i, every element in the i-th row of s^μ is in a different column of t^λ, then $\lambda \trianglerighteq \mu$*

Proof. Fix some arbitrary i. Then let us assume, by way of the hypothesis, that the elements of the j-th row of s^μ are in different columns of t^λ for $j = 1, \ldots, i$. Then we can sort the entries of each column of t^λ such that the elements of the first i rows of s^μ all occur in the first i rows of t^λ. Then we have

$$\begin{aligned} \lambda_1 + \cdots + \lambda_i = &\text{ the number of elements in the first } i \text{ rows of } t^\lambda \\ \geq &\text{ the number of elements of } s^\mu \text{ in the first } i \text{ rows of } t^\lambda \\ = &\ \mu_1 + \cdots + \mu_i \end{aligned}$$

Since i was arbitrary, this implies that $\lambda \trianglerighteq \mu$. $\square$

This will be the ordering that we will use for our permutation modules, but we remark that there are "finer" orderings that will also allow for the same result[3].

In particular, this implies that our first irreducible module will be $M^{(n)}$. We easily verify that this is indeed irreducible as there is only one tabloid of this shape, so this module has dimension 1 and thus must be irreducible.

2.3. Polytabloids and Specht Modules.

We now intend to more explicitly describe these irreducible modules, which we will do through the development of more combinatorial objects.

Definition 2.11. Let t be a tableau with rows $R_1, \ldots, R_a$ and columns $C_1, \ldots, C_b$. Then we define

$$R_t \coloneqq \mathfrak{S}_{R_1} \times \cdots \times \mathfrak{S}_{R_a}$$

to be the *row-stabilizer* and

$$C_t \coloneqq \mathfrak{S}_{C_1} \times \cdots \times \mathfrak{S}_{C_a}$$

to be the *column-stabilizer* of t.

[3] More specifically, we can define a *lexicographic* ordering, say $\geq$, and then we would be able to establish the same result by listing the permutation modules in dual lexicographic order. However, we have that if $\lambda \trianglerighteq \mu$, then $\lambda \geq \mu$ (lexicographic ordering is a *refinement* of dominance ordering) so our approach that uses dominance ordering leads to a stronger conclusion.

Note that these are both subgroups of $\mathfrak{S}_n$ when t is a tableau with n boxes. Note that the row-stabilizer provides another way for us to understand tabloids as we have $\{t\} = R_t t$. For the other subgroup, we use the column-stabilizer to define our next combinatorial object.

Definition 2.12. Let t be a tableau. Then its associated *polytabloid* is given by

$$\boldsymbol{e_t} := \kappa_t\{t\}$$

where $\kappa_t := \sum_{\pi \in C_t} \operatorname{sgn}(\pi)\pi$.

Example 2.12.1. Let $\lambda = (3, 2) \vdash 5 = n$ and let us define the following λ-tableau.

$$t = \begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline 4 & 5 & \multicolumn{1}{c}{} \\ \cline{1-2}\end{array}$$

Then the column stabilizer of t is $C_t = \{(), (14), (25), (14)(25)\}$. So the polytabloid associated with t expands into the following vector in M^λ.

$$\begin{aligned}
\boldsymbol{e_t} &= \kappa_t \cdot \begin{array}{ccc}\hline 1 & 2 & 3 \\ \hline 4 & 5 & \\ \cline{1-2}\end{array} \\
&= \left(\sum_{\pi \in C_t} \operatorname{sgn}(\pi)\pi\right) \cdot \begin{array}{ccc}\hline 1 & 2 & 3 \\ \hline 4 & 5 & \\ \cline{1-2}\end{array} \\
&= \sum_{\pi \in C_t} \operatorname{sgn}(\pi) \left(\pi \cdot \begin{array}{ccc}\hline 1 & 2 & 3 \\ \hline 4 & 5 & \\ \cline{1-2}\end{array}\right) \\
&= \begin{array}{ccc}\hline 1 & 2 & 3 \\ \hline 4 & 5 & \\ \cline{1-2}\end{array} - \begin{array}{ccc}\hline 4 & 2 & 3 \\ \hline 1 & 5 & \\ \cline{1-2}\end{array} - \begin{array}{ccc}\hline 1 & 5 & 3 \\ \hline 4 & 2 & \\ \cline{1-2}\end{array} + \begin{array}{ccc}\hline 4 & 5 & 3 \\ \hline 1 & 2 & \\ \cline{1-2}\end{array}
\end{aligned}$$

These objects interact with the symmetric group in some important ways, captured in the following lemma.

Lemma 2.13. *Let t be a tableau and π be a permutation. Then*

(1) $R_{\pi t} = \pi R_t \pi^{-1}$
(2) $C_{\pi t} = \pi C_t \pi^{-1}$
(3) $\kappa_{\pi t} = \pi \kappa_t \pi^{-1}$
(4) $\boldsymbol{e_{\pi t}} = \pi \boldsymbol{e_t}$

Proof.

(1) We have

$$\begin{aligned}
\sigma \in R_{\pi t} &\iff \sigma\{\pi t\} = \{\pi t\} \\
&\iff \pi^{-1}\sigma\pi\{t\} = \{t\} \\
&\iff \pi^{-1}\sigma\pi \in R_t \\
&\iff \sigma \in \pi R_t \pi^{-1}
\end{aligned}$$

(2) Let $[t]$ denote the equivalence class containing t, with respect to the equivalence relation of column similarity. Then similarly to (1), we have

$$\begin{aligned}\sigma \in C_{\pi t} &\iff \sigma[\pi t] = [\pi t]\\ &\iff \pi^{-1}\sigma\pi[t] = [t]\\ &\iff \pi^{-1}\sigma\pi \in C_t\\ &\iff \sigma \in \pi C_t \pi^{-1}\end{aligned}$$

(3) Recall that $\kappa_t = \sum_{\pi \in C_t} \operatorname{sgn}(\pi)\pi$. So by (2), we have

$$\begin{aligned}\kappa_{\pi t} &= \sum_{\sigma \in C_{\pi t}} \operatorname{sgn}(\sigma)\sigma\\ &= \sum_{\sigma \in C_t} \operatorname{sgn}(\pi\sigma\pi^{-1})\pi\sigma\pi^{-1}\\ &= \pi\left(\sum_{\sigma \in C_t} \operatorname{sgn}(\sigma)\sigma\right)\pi^{-1}\\ &= \pi\kappa_t\pi^{-1}\end{aligned}$$

(4) We have

$$\begin{aligned}\boldsymbol{e_{\pi t}} &= \kappa_{\pi t}\{\pi t\}\\ &= \pi\kappa_t\pi^{-1}\{\pi t\}\\ &= \pi\kappa_t\{t\}\\ &= \pi\boldsymbol{e_t}\end{aligned}$$

□

We can now define the irreducible representations of $\mathfrak{S}_n$.

Definition 2.14. Let $\lambda \vdash n$. Then the submodule of M^λ spanned by the polytabloids $\boldsymbol{e_t}$ (where t has shape λ) is called a *Specht module* and denoted by S^λ.

We are left now with two tasks. First we must verify that these constructed modules are indeed irreducible. Next, we show that they form *all* of the irreducible representations of $\mathfrak{S}_n$.

2.4. The Submodule Theorem.

Before we prove the irreducibility of Specht modules, we present some important properties of alternating groups sums. Note that if $H \leqslant \mathfrak{S}_n$ is a subgroup, then we define the alternating group sum of H to be $H^- = \sum_{h \in H} \operatorname{sgn}(h)h$. We are interested in these sums because they are a generalization of the κ_t term we used to define polytabloids.

In order to fully state the following lemma, we will also make use of an inner product on M^λ. Namely,

$$\langle \boldsymbol{\{t\}}, \boldsymbol{\{s\}} \rangle \coloneqq \delta_{\{t\},\{s\}}$$

Lemma 2.15 (The Sign Lemma). *Let $H \leqslant \mathfrak{S}_n$ be a subgroup.*

(1) If $\pi \in H$, then

$$\pi H^- = H^-\pi = \operatorname{sgn}(\pi)H^-$$

(2) For any $\boldsymbol{u}, \boldsymbol{v} \in M^\mu$,

$$\langle H^- \boldsymbol{u}, \boldsymbol{v} \rangle = \langle \boldsymbol{u}, H^- \boldsymbol{v} \rangle$$

(3) If $(b, c) \in H$, then there exists $k \in \mathbb{C}[\mathfrak{S}_n]$ such that

$$H^- = k(\epsilon - (b, c))$$

(4) If $(b, c) \in H$ and t is a tableau with b and c in the same row, then

$$H^- \{\boldsymbol{t}\} = 0$$

Proof.

(1) Let $\pi \in H$. Then we have the following.

$$\begin{aligned} \pi H^- &= \sum_{\sigma \in H} \operatorname{sgn}(\sigma) \pi\sigma \\ &= \sum_{\tau \in H} \operatorname{sgn}(\pi) \operatorname{sgn}(\tau) \tau \\ &= \operatorname{sgn}(\pi) H^- \end{aligned}$$

We have that $H^- \pi = \operatorname{sgn}(\pi) H^-$ by a similar argument where $\sigma\pi$ is substituted by an arbitrary τ.

(2) Since the sign of a permutation is the same as that of its inverse, we have the following.

$$\begin{aligned} \langle H^- \boldsymbol{u}, \boldsymbol{v} \rangle &= \sum_{\pi \in H} \langle (\operatorname{sgn}\pi) \pi \boldsymbol{u}, \boldsymbol{v} \rangle \\ &= \sum_{\pi \in H} \langle \boldsymbol{u}, (\operatorname{sgn}\pi) \pi^{-1} \boldsymbol{v} \rangle \\ &= \sum_{\pi \in H} \langle \boldsymbol{u}, (\operatorname{sgn}\pi) \pi \boldsymbol{v} \rangle \\ &= \langle \boldsymbol{u}, H^- \boldsymbol{v} \rangle \end{aligned}$$

(3) Let $K = \{\epsilon, (b, c)\}$ be a subgroup of H. Then there is a transversal such that we can express $H = \cup_i k_i K$. Then we have that $H^- = (\sum_i k_i^-)(\epsilon - (b, c))$.

(4) Let t be a tableau with b and c in the same row. Then $(b, c)\{\boldsymbol{t}\} = \{\boldsymbol{t}\}$. So by (3), we have

$$H^- \{\boldsymbol{t}\} = k(\epsilon - (b - c))\{\boldsymbol{t}\} = k\{\boldsymbol{t}\} - k\{\boldsymbol{t}\} = 0$$

□

We are now able to prove the important submodule theorem which will be used to complete our study of the representations of the symmetric group.

Theorem 2.16 (The Submodule Theorem). *Let U be a submodule of M^μ, then*

$$U \supseteq S^\mu \text{ or } U \subseteq S^{\mu\perp}$$

In particular, the S^μ are irreducible when we work over $\mathbb{C}$.

Proof. Let $\boldsymbol{u} \in U$ and let t be a μ-tableau. First we argue that if s is a μ-tableau, then

$$\kappa_t \{\boldsymbol{s}\} = \pm \boldsymbol{e_t}$$

This holds because we have that $\boldsymbol{\{s\}} = \pi\boldsymbol{\{t\}}$, which implies

$$\kappa_t\boldsymbol{\{s\}} = \kappa_t\pi\boldsymbol{\{t\}} = \mathrm{sgn}(\pi)\kappa_t\boldsymbol{\{t\}} = \pm\boldsymbol{e_t}$$

From this result, we obtain that $\kappa_t\boldsymbol{u}$ is a multiple of $\boldsymbol{e_t}$ because $\boldsymbol{u} = \sum_i c_i\boldsymbol{\{s_i\}}$ where the s_i are μ-tableaux. So we can write $\kappa_t\boldsymbol{u} = f\boldsymbol{e_t}$. We now have two cases for our submodule U.

For the first case, suppose there exists $\boldsymbol{u} \in U$ and a tableau t such that $f \neq 0$. Then there exists f^{-1} so we have $\boldsymbol{e_t} = f^{-1}\kappa_t\boldsymbol{u} \in U$. Thus, $S^\mu \subset U$.

Now let us consider the other case, where we always have $\kappa_t\boldsymbol{u} = 0$. Then we have the following by part (2) of the sign lemma.

$$\langle\boldsymbol{u}, \boldsymbol{e_t}\rangle = \langle\boldsymbol{u}, \kappa_t\boldsymbol{\{t\}}\rangle = \langle\kappa_t\boldsymbol{u}, \boldsymbol{\{t\}}\rangle = \langle 0, \boldsymbol{\{t\}}\rangle = 0$$

Since the $\boldsymbol{e_t}$ span S^μ, we obtain that $\boldsymbol{u} \in S^{\mu\perp}$. □

As a consequence to this theorem, we obtain a complete description of *all* the irreducible representations of the symmetric group.

Theorem 2.17. *The set* $\{S^\lambda\colon \lambda \vdash n\}$ *constitutes a complete set of irreducible representations of* $\mathfrak{S}_n$ *over* $\mathbb{C}$.

Proof. These modules are irreducible by the submodule theorem and because $S^\lambda \cap S^{\lambda\perp} = 0$. Moreover, we have the maximal number of irreducible modules so this is indeed the full set of irreducibles. □

3. References

- William Fulton, Joe Harris. *Representation Theory*. Springer-Verlag, 1991.
- Bruce E. Sagan. *The Symmetric Group*. Springer New York, NY, 2001.
- Jean-Pierre Serre. *Linear Representations of Finite Groups*. Springer-Verlag, 1977.