Representations of the symmetric group

1 Conjugacy classes and Young diagrams

Let us recall what we know about irreducible representations of S_n so far: we have the two 1-dimensional representations $\mathbb{C}$ and $\mathbb{C}(\varepsilon)$, and an irreducible representation V of dimension $n-1$ which satisfies: $V \oplus \mathbb{C} \cong \mathbb{C}[S_n/S_{n-1}]$, where $\mathbb{C}[S_n/S_{n-1}]$ is the standard permutation representation of S_n coming from its action on $\{1, \dots, n\}$. There is also the irreducible representation $V \otimes \varepsilon$, which is not isomorphic to V once $n \geq 4$. Our goal in this set of notes will be to describe a construction of all of the irreducible representations of S_n.

We begin by recalling the usual description of the conjugacy classes in S_n. Every $\sigma \in S_n$ can be written as $\gamma_1 \cdots \gamma_k$, where the γ_i are pairwise disjoint n_i-cycles and the product is unique up to order. Here the identity 1 corresponds to the empty product ($k = 0$). We may as well reorder so that $n_1 \geq n_2 \geq \cdots \geq n_k$, so that the n_i form a non-increasing sequence of integers at least 2 and $\sum_{i=1}^{k} n_i \leq n$. We will refer to the sequence $(n_1, \dots, n_k)$ as the *cycle type* of σ. For example, an r-cycle has cycle type r. The element $(1,2)(3,4,5)$ has cycle type $(3,2)$. Two elements of S_n are conjugate $\iff$ they have the same cycle type.

It is convenient to rewrite this description of the conjugacy classes via partitions:

Definition 1.1. A partition λ of n, which we write symbolically as $\lambda \vdash n$, is a weakly decreasing (i.e. non-increasing) sequence of positive integers $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_\ell$, such that $\sum_{i=1}^{\ell} \lambda_i = n$.

Note that, given a cycle type, i.e. a non-increasing sequence of integers $n_1 \geq n_2 \geq \cdots \geq n_k$ at least 2 and $\sum_{i=1}^{\ell} n_i \leq n$, we can always enlarge the sequence to a partition by considering

$$n_1 \geq n_2 \geq \cdots \geq n_k \geq 1 \geq 1 \geq \cdots \geq 1,$$

where the number of terms equal to 1 that we add is $n - \sum_{i=1}^{k} n_i$. Conversely, given a partition λ, we obtain a cycle type by dropping off all of the terms at the end with $\lambda_i = 1$. Thus we see that the conjugacy classes of S_n are indexed by partitions of n. It is therefore reasonable to hope that the irreducible representations of S_n are also indexed by partitions.

Definition 1.2. Given a partition $\lambda \vdash n$, the *Young subgroup* $S_\lambda \le S_n$ is the subgroup of S_n defined by: $\sigma \in S_\lambda \iff \sigma(\{1, \dots, \lambda_1\}) = \{1, \dots, \lambda_1\}$, $\sigma(\{\lambda_1 + 1, \dots, \lambda_1 + \lambda_2\}) = \{\lambda_1 + 1, \dots, \lambda_1 + \lambda_2\}$, $\dots$, and more generally, for all i, $1 \le i \le \ell$,

$$\sigma\left(\left\{\sum_{j=1}^{i-1} \lambda_j + 1, \dots, \sum_{j=1}^{i} \lambda_j\right\}\right) = \left\{\sum_{j=1}^{i-1} \lambda_j + 1, \dots, \sum_{j=1}^{i} \lambda_j\right\}.$$

In other words, S_λ is the subgroup which preserves the first set of λ_1 consecutive elements of $\{1, \dots, n\}$, then the next set of λ_2 consecutive elements, and so on. Thus clearly

$$S_\lambda \cong S_{\lambda_1} \times \cdots \times S_{\lambda_\ell}.$$

Hence $\#(S_\lambda) = (\lambda_1)! \cdots (\lambda_\ell)!$ and $\#(S_n/S_\lambda) = n!/(\lambda_1)! \cdots (\lambda_\ell)!$.

For example, if $\lambda = (n)$, then $S_\lambda = S_n$, whereas if $\lambda = (1, 1, \dots, 1)$, then $S_\lambda = \{1\}$.

Let $M_\lambda = \mathbb{C}[S_n/S_\lambda] = \mathrm{Ind}_{S_\lambda}^{S_n} \mathbb{C}$. The basic idea will be to locate an irreducible subspace S^λ of M^λ satisfying certain properties. The representations S^λ will exactly be the irreducible representations of S_n up to isomorphism.

2 Young diagrams and Young tableaux

Definition 2.1. Given a partition $\lambda \vdash n$, its *Young diagram* is given by drawing n boxes in ℓ rows, flush left, with the i^{th} row having λ_i boxes.

For example, given $\lambda = (3, 2, 1, 1, 1) \vdash 8$ its Young diagram is

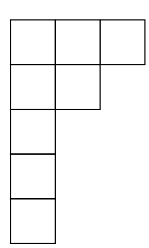

At the two extremes, for $\lambda = (n)$, the corresponding diagram is

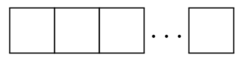

and for $\lambda = (1, 1, \ldots, 1)$, the corresponding diagram is

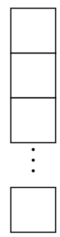

We define an operation of transpose (written $\lambda \mapsto \lambda^T$) on Young diagrams by switching rows and columns. For example, with $\lambda = (3, 2, 1, 1, 1) \vdash 8$ as before, the transpose diagram is

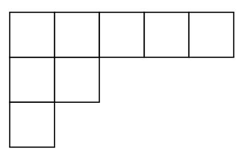

which corresponds to $\lambda^T = (5, 2, 1)$. Likewise $(n)^T = (1, 1, , \ldots, 1)$. Clearly $(\lambda^T)^T = \lambda$.

Next, we define a partial order on the set of all partitions:

Definition 2.2. Suppose that $\lambda, \mu \vdash n$, where $\lambda = (\lambda_1, \ldots, \lambda_\ell)$ and $\mu = (\mu_1, \ldots, \mu_m)$. Then λ *dominates* μ, written $\lambda \trianglerighteq \mu$, if, for all i,

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

Here, if $i > \ell$, we set $\lambda_i = 0$, and similarly if $i > m$ we set $\mu_i = 0$. The definition amounts to saying that, for every i, the first i rows of the Young diagram for λ contain at least as many boxes as the first i rows of the Young diagram for μ.

The relation $\trianglerighteq$ is only a partial order because not every two partitions are comparable. For example, $(5, 2, 1) \trianglerighteq (3, 4, 1)$, but $(5, 1, 1, 1)$ and $(3, 4, 1)$ are not comparable. For every partition λ, $(n) \trianglerighteq \lambda$ and $\lambda \trianglerighteq (1, 1, \ldots, 1)$.

The following lemma makes precise the sense in which $\trianglerighteq$ is a partial order:

Lemma 2.3. *With $\trianglerighteq$ defined as above, and for all $\lambda, \mu, \nu \vdash n$,*

(i) $\lambda \trianglerighteq \lambda$.

(ii) *If $\lambda \trianglerighteq \mu$ and $\mu \trianglerighteq \nu$, then $\lambda \trianglerighteq \nu$.*

(iii) *If $\lambda \trianglerighteq \mu$ and $\mu \trianglerighteq \lambda$, then $\lambda = \mu$.*

Proof. (i) and (ii) follow easily from the definition. As for (iii), note that by definition $\lambda_1 \geq \mu_1$ and $\mu_1 \geq \lambda_1$, hence $\lambda_1 = \mu_1$. Assume inductively that we have shown that $\lambda_k = \mu_k$, $k \leq i-1$. Then since $\lambda \trianglerighteq \mu$,

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

and hence $\lambda_i \geq \mu_i$. By symmetry, $\mu_i \geq \lambda_i$. Hence $\lambda_i = \mu_i$, completing the inductive step and hence the proof of (iii). □

Definition 2.4. Given $\lambda \vdash n$, a *λ-tableau t* or a *tableau of type λ* is a labeling of the n boxes of the Young diagram of λ by the elements of $\{1, \dots, n\}$, in other words a way to fill in the n boxes of the Young diagram with the elements of $\{1, \dots, n\}$, using each element exactly once. Hence, given λ, there are exactly $n!$ tableau of type λ. For example, given λ, the *basic λ-tableau* t_0 is obtained by filling in the boxes consecutively: for $\lambda = (3,2,1) \vdash 6$, the basic tableau t_0 is

1	2	3
4	5	
6		

Two λ-tableaux t_1 and t_2 are *equivalent*, written $t_1 \sim t_2$, if, for every i, the set of entries in the i^{th} row of t_1 is the same as the set of entries in the i^{th} row of t_1. In other words, t_2 is obtained from t_1 by permuting each row of t_1. For example,

3	2	6
4	1	
5		

and

6	2	3
1	4	
5		

are equivalent.

We write the equivalence class containing t as $[t]$. An equivalence class of λ-tableaux is called a *λ-tabloid* or a *tabloid of type* λ.

Clearly, S_n acts transitively on the set of λ-tableaux and preserves the equivalence relation $\sim$. Thus S_n acts transitively on the set of λ-tabloids. If t_0 is the basic λ-tableau, then the isotropy subgroup of t_0 is S_λ, the Young subgroup. Hence we can identify the set of all λ-tabloids with S_n/S_λ. In particular, there are $n!/(\lambda_1)!\cdots(\lambda_\ell)!$ λ-tabloids.

Given a λ-tableau t, we can define a λ^T-tableau t^T in the obvious way. Clearly, if $\sigma \in S_n$, then $(\sigma \cdot t)^T = \sigma \cdot (t^T)$. However, if $t_1 \sim t_2$, t_1^T and t_2^T are **not** in general equivalent.

As before, let $M_\lambda = \mathbb{C}[S_n/S_\lambda] = \operatorname{Ind}_{S_\lambda}^{S_n} \mathbb{C}$. We view M^λ as having a basis consisting of λ-tabloids. Our goal will be to find an irreducible subspace S^λ of M^λ with the property that $\dim \operatorname{Hom}^{S_n}(S^\lambda, M^\lambda) = 1$, i.e. that the multiplicity of S^λ in M^λ is 1, and such that $\operatorname{Hom}^{S_n}(S^\lambda, M^\mu) \neq 0 \implies \lambda \trianglerighteq \mu$.

Example 2.5. (1) If $\lambda = (n)$, then $S_\lambda = S_n$, $M^{(n)}$ is the trivial representation $\mathbb{C}$, and necessarily $S^{(n)} = \mathbb{C}$. Note that $(n) \trianglerighteq \mu$ for every partition $\mu \vdash n$, and also that the trivial representation occurs in M^μ for every μ because M^μ is a permutation representation.

(2) If $\lambda = (1, \dots, 1)$, then $M^{(1,\dots,1)}$ is the regular representation. We will see that $S^{(1,\dots,1)} = \mathbb{C}(\varepsilon)$. Note that $\mu \trianglerighteq (1, \dots, 1)$ for every $\mu \vdash n$. Likewise, $\operatorname{Hom}^{S_n}(S^\mu, M^{(1,\dots,1)}) \neq 0$ since every irreducible representation is isomorphic to a subspace of the regular representation.

(3) If $\lambda = (n-1, 1)$, then $S_\lambda \cong S_{n-1}$ and $M^{(n-1,1)}$ is the permutation representation of S_n acting on $\{1, \dots, n\}$. Hence $M^{(n-1,1)} \cong \mathbb{C} \oplus V$, where V is irreducible of dimension $n-1$. Correspondingly, if $\lambda \trianglerighteq (n-1, 1)$, then either $\lambda = (n)$ or $\lambda = (n-1, 1)$.

3 Row and column stabilizers; polytabloids

Definition 3.1. Let t be a λ-tableau with associated tabloid $[t]$. We define the *row stabilizer* R_t to be the subgroup of S_n consisting of all elements σ such that, for every i, σ preserves the set of elements in the i^{th} row of t. Equivalently, R_t is the isotropy subgroup of the associated tabloid $[t]$. For $t = t_0$, $R_{t_0} = S_\lambda$ is the Young subgroup. For a general t, R_t is conjugate to S_λ as we shall see shortly.

We likewise define the *column stabilizer* C_t to be the subgroup of S_n consisting of all elements σ such that, for every i, σ preserves the set of elements in the i^{th} **column** of t. Thus C_t is the isotropy subgroup of $[t^T]$, so that $C_t = R_{t^T}$, and hence C_t is conjugate to S_{λ^T}. However, C_t depends on the tableau t, not just on the tabloid $[t]$.

Lemma 3.2. *For all tableaux t and all $\sigma \in S_n$,*

(i) $R_t \cap C_t = \{1\}$.

(ii) $R_{\sigma \cdot t} = \sigma R_t \sigma^{-1}$ *and* $C_{\sigma \cdot t} = \sigma C_t \sigma^{-1}$.

Proof. (i) Let a be the $(i,j)^{\text{th}}$ entry of t, i.e. a lies in the i^{th} row and j^{th} column of t. If $\sigma \in R_t \cap C_t$, then $\sigma(a)$ is also in the i^{th} row and j^{th} column of t. Thus $\sigma(a) = a$ for all $a \in \{1, \dots, n\}$, so that $\sigma = 1$.
(ii) This is a general fact about isotropy subgroups for group actions. □

If t is a λ-tableau, we define the following element of the group algebra:

$$A_t = \sum_{\sigma \in C_t} \varepsilon(\sigma)\sigma \in \mathbb{C}[S_n].$$

Note that we sum over the **column stabilizer** C_t. Since the group algebra acts on all representations, given an S_n-representation V, we can view A_t as defining a linear map $V \to V$. In particular, A_t defines a linear map $M^\mu \to M^\mu$ for all $\mu \vdash n$.

Definition 3.3. Given a λ-tableau t, the *polytabloid* e_t associated to t is the element

$$e_t = A_t([t]) = \sum_{\sigma \in C_t} \varepsilon(\sigma)\sigma \cdot [t] = \sum_{\sigma \in C_t} \varepsilon(\sigma)[\sigma \cdot t] \in M^\lambda.$$

Remark 3.4. As we shall see in numerous examples, e_t depends on t, not just on $[t]$, because C_t depends on t and not just on $[t]$.

Lemma 3.5. *For all tableaux t, $e_t \neq 0$.*

Proof. Note first that, if $\sigma \in C_t$ and $\sigma \cdot [t] = [t]$, then $\sigma \in R_t$ and hence $\sigma \in R_t \cap C_t = \{1\}$. Likewise, if $\sigma_1, \sigma_2 \in C_t$ and $\sigma_1 \cdot [t] = \sigma_2 \cdot [t]$, then $\sigma_2^{-1}\sigma_1 = 1$ and hence $\sigma_1 = \sigma_2$. It follows that $e_t = \sum_{\sigma \in C_t} \varepsilon(\sigma)\sigma \cdot [t]$ is a sum of different basis vectors in M^λ, with coefficients ± 1, and hence $e_t \neq 0$. □

Lemma 3.6. *For all tableaux t and all $\sigma \in S_n$,*

(i) $\sigma \cdot A_t = A_{\sigma \cdot t}\sigma$ *as elements of the group algebra.*

(ii) $\sigma \cdot e_t = e_{\sigma \cdot t}$.

Proof. (i) We have $\sigma A_t = \sum_{\tau \in C_t} \varepsilon(\tau)\sigma\tau$. On the other hand,

$$\begin{aligned} A_{\sigma \cdot t}\sigma &= \sum_{\tau \in C_{\sigma \cdot t}} \varepsilon(\tau)\tau\sigma = \sum_{\tau \in \sigma C_t \sigma^{-1}} \varepsilon(\tau)\tau\sigma \\ &= \sum_{\tau \in C_t} \varepsilon(\sigma\tau\sigma^{-1})\sigma\tau\sigma^{-1}\sigma = \sum_{\tau \in C_t} = \varepsilon(\tau)\sigma\tau. \end{aligned}$$

Comparing, we see that $A_{\sigma \cdot t}\sigma = \sigma \cdot A_t$.

(ii) By definition,

$$\sigma \cdot e_t = \sigma A_t([t]) = A_{\sigma \cdot t}\sigma([t]) = A_{\sigma \cdot t}([\sigma \cdot t]) = e_{\sigma \cdot t}.$$

□

We now define the representation S^λ. The idea is as follows: let G be a finite group and V an irreducible G-representation. For a fixed vector $v \in V$, the span of the set

$$G \cdot v = \{\rho_V(g)(v) : g \in G\}$$

is clearly a nonzero G-invariant subspace of V. If moreover $v \in W$, where W is an irreducible subspace of V, then this span is a nonzero G-invariant subspace of W, hence must equal W.

Definition 3.7. Given $\lambda \vdash n$, define S^λ, the *Specht representation*, to be the span of the polytabloids e_t, where t is a λ-tableau. Since $\sigma \cdot e_t = e_{\sigma \cdot t}$, S^λ is an S_n-invariant subspace of M^λ, hence an S_n-representation.

Example 3.8. (1) If $\lambda = (n)$, then $M^{(n)} = \mathbb{C}$ with the trivial action of S_n. Here, there is only one tableau $[t]$, $C_t = \{1\}$, and $A_t = \mathrm{Id}$.

(2) Let $\lambda = (n-1, 1)$. Then every tableau t is of the form

$*$	$*$	$*$	$\ldots$	$*$
k				

for a unique $k = k(t)$, $1 \leq k \leq n$. Moreover, two tableaux t_1 and t_2 are equivalent $\iff k(t_1) = k(t_2)$. Hence the $(n-1,1)$-tabloids are indexed by $k \in \{1, \ldots, n\}$. Let $[k]$ denote the corresponding equivalence class. Clearly $\sigma \cdot [k] = [\sigma(k)]$. Thus $M^{(n-1,1)} \cong \mathbb{C}^n$, with basis vectors $[1], \ldots, [n]$, and the S_n-action is the same as the standard permutation representation. If $t \in [k]$, let the entry in the first row and column of t be i, so that t is of the form

$$\begin{array}{|c|c|c|c|c|} \hline i & * & * & \cdots & * \\ \hline k \\ \cline{1-1} \end{array}$$

Then $C_t = \{1, (ik)\}$. Note that C_t depends on t, not just $[t] = [k]$. Hence $A_t([t]) = [k] + \varepsilon((ik))(ik) \cdot [k] = [k] - [i]$. The vectors $[k] - [i]$ are not linearly independent, and their span in $M^{(n-1,1)} \cong \mathbb{C}^n$ is easily seen to be

$$\left\{ a_1[1] + \cdots + a_n[n] : \sum_{i=1}^{n} a_i = 0 \right\}.$$

Thus $S^{(n-1,1)} \cong V$, the standard irreducible representation of dimension $n-1$ of S_n.

(3) For $\lambda = (1, 1, \ldots, 1)$, $M^{(1,1,\ldots,1)} = \mathbb{C}[S_n]$, $C_t = S_n$ for every t, and $R_t = \{1\}$. The tableaux are the same as the tabloids, and correspond to elements $\sigma \in S_n$ via: t_σ is the $(1, 1, \ldots, 1)$-tableau whose entries going vertically are $\sigma(1), \sigma(2), \ldots, \sigma(n)$. Thus $t_1 = t_0$ is the basic tableau and $t_\sigma = \sigma \cdot t_1$. Then $A_t = \sum_{\sigma \in S_n} \varepsilon(\sigma)\sigma$ for every t, and

$$e_{t_1} = A_t([t_1]) = \sum_{\sigma \in S_n} \varepsilon(\sigma)[t_\sigma].$$

By a standard calculation in the group algebra, for every $\tau \in S_n$,

$$\tau \cdot A_t = A_t \cdot \tau = \varepsilon(\tau) A_t.$$

Thus $\tau \cdot e_{t_1} = e_{t_\tau} = A_t \tau([t_1]) = \varepsilon(\tau) e_{t_1} = \pm e_{t_1}$, so $S^{(1,1,\ldots,1)}$ is 1-dimensional, and $\tau(e_{t_1}) = \varepsilon(\tau) e_{t_1}$, so that $S^{(1,1,\ldots,1)} \cong \mathbb{C}(\varepsilon)$.

4 Proof of irreducibility

We begin with the following lemma:

Lemma 4.1 (Dominance lemma)**.** *Let $\lambda, \mu \vdash n$, let t be a λ-tableau and let s be a μ-tableau. Suppose that, for every i, if $a \neq b$ are two entries in the i^{th} row of s, then a and b lie in different columns of t. Then $\lambda \trianglerighteq \mu$.*

Proof. We first establish a claim which we shall also use:

Claim 4.2. *With hypotheses as above, there exists a $\sigma \in C_t$ such that, after replacing t by $\sigma \cdot t$, for every i, the elements in the first i rows of s all appear in the first i rows of t. Equivalently, if S_i is the set of elements in the i^{th} row of s and T_j is the set of elements in the j^{th} row of t, then, for every i, $\sigma(S_i) \subseteq \bigcup_{j \leq i} T_j$.*

First let us show that the claim implies the lemma. Assuming the claim, for every i there are $\mu_1 + \cdots + \mu_i$ elements in the first i rows of s. Since they all appear in the first i rows of t, the number of elements in the first i rows of t, namely $\lambda_1 + \cdots + \lambda_i$, has to be at least as large as $\mu_1 + \cdots + \mu_i$. In other words, for every i,

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

and hence $\lambda \trianglerighteq \mu$. □

Proof of the claim. Note that the hypotheses of the lemma are unchanged by applying column permutations to t, i.e. by replacing t by $\sigma \cdot t$ for $\sigma \in C_t$. We will give an inductive construct of an appropriate σ.

For $i = 1$, the entries of the first row of s are in different columns of t. In particular there are $\lambda_1 \geq \mu_1$ columns of t. Permute each column of t containing an element of the first row of s by moving the given element into the first row (for example, by a transposition if it is not already in the first row). This replaces t by $\sigma_1 \cdot t$ for some $\sigma_1 \in C_t$.

For the inductive step of the construction, suppose that we have found a $\sigma_i \in C_t$ such that, after replacing t by $\sigma_i \cdot t$, the elements in the first i rows of s all appear in the first i rows of t. Now consider the entries in the $(i+1)^{\text{st}}$ row of s. If any of these entries appear in one of the first $(i+1)^{\text{st}}$ rows of t, we leave the corresponding columns alone. If some entry a appears in the j^{th} row of t with $j > i+1$, suppose that a is also in the k^{th} column of t. Then no other entry in the $(i+1)^{\text{st}}$ row of s lies in the k^{th} column of t. Also, since a lies below the $(i+1)^{\text{st}}$ row of t, the k^{th} column of t has a nonempty intersection with $(i+1)^{\text{st}}$ row of t. Then we can permute the k^{th} column of t by switching the in the j^{th} row, namely a, with the entry in the $(i+1)^{\text{st}}$ row. This procedure doesn't affect the first i rows, and can be done independently for each entry in the $(i+1)^{\text{st}}$ row of s which lies in in the j^{th} row of t for some $j > i+1$, since these all lie in different columns. We thus modify $\sigma_i \cdot t$ by a column permutation, and hence t by a column permutation σ_{i+1}, so that $\sigma_{i+1} \cdot t$ has the desired properties. This completes the inductive step of the construction. □

Recall that M^λ has a basis consisting of the λ-tabloids $[t]$. We can introduce a positive definite Hermitian inner product $\langle \cdot, \cdot \rangle$ on M^λ by decreeing that this basis is unitary, i.e. that

$$\langle [t_1], [t_2] \rangle = \begin{cases} 1, & \text{if } [t_1] = [t_2]; \\ 0, & \text{otherwise.} \end{cases}$$

Since S_n acts on M^λ by permuting the basis vectors, this Hermitian inner product is S_n-invariant. In what follows, when we speak about orthogonality, it will be with respect to $\langle \cdot, \cdot \rangle$.

The key technical result we need now is:

Theorem 4.3 (Submodule theorem)**.** *Let V be an S_n-invariant subspace of M^λ. Then either $S^\lambda \subseteq V$ or $V \subseteq (S^\lambda)^\perp$.*

Proof. We start with the following lemma:

Lemma 4.4. *Suppose that $\lambda, \mu \vdash n$, that t is a λ-tableau and that s is a μ-tableau. If $A_t([s]) \neq 0$, then $\lambda \trianglerighteq \mu$. If moreover $\lambda = \mu$, then $A_t([s]) = \pm e_t$.*

Proof. First, we claim that, if $A_t([s]) \neq 0$, then for every i, if a and b are two elements in the i^{th} row of s, then a and b are in different columns of t. The dominance lemma then implies that $\lambda \trianglerighteq \mu$. To see the claim, suppose by contradiction that a and b are in the same column of t. Then $(ab) \in C_t$, and also $(ab) \in R_s$, so that $(ab) \cdot [s] = s$. Let $H = \{1, (ab)\} \leq C_t$ be the subgroup generated by (ab). We can then break C_t up into the left cosets for H: if $\sigma_1, \ldots, \sigma_N$ are a set of representatives for C_t/H, then

$$C_t = \{\sigma_1, \sigma_1 \cdot (ab), \ldots, \sigma_N, \sigma_N \cdot (ab)\}.$$

Then, writing the elements of C_t as above,

$$\begin{aligned} A_t([s]) = \sum_{\sigma \in C_t} \varepsilon(\sigma)\sigma \cdot [s] &= \sum_{i=1}^{N} (\varepsilon(\sigma_i)\sigma_i \cdot [s] + \varepsilon(\sigma_i \cdot (ab))\sigma_i \cdot (ab) \cdot [s]) \\ &= \sum_{i=1}^{N} (\varepsilon(\sigma_i)\sigma_i \cdot [s] - \varepsilon(\sigma_i)\sigma_i \cdot [s]) = 0, \end{aligned}$$

where we have used the fact that $(ab) \cdot [s] = [s]$ and that

$$\varepsilon(\sigma_i \cdot (ab)) = \varepsilon(\sigma_i)\varepsilon((ab)) = -\varepsilon(\sigma_i).$$

But this contradicts the assumption that $A_t([s]) \neq 0$.

Now suppose that $\lambda = \mu$ and that $A_t([s]) \neq 0$. As we have seen, the hypotheses of the dominance lemma hold. Then by Claim 4.2 there exists a $\tau \in C_t$ such that, if S_i is the set of elements in the i^{th} row of s and T_j is the set of elements in the j^{th} row of t, then, for every i, $\tau(S_i) \subseteq \bigcup_{j \le i} T_j$. We claim that this forces $\tau \cdot [s] = [t]$. First, $\tau(S_1) \subseteq T_1$, but since $\lambda_1 = \mu_1$, S_1 and T_1 have the same number of elements. Since τ is injective, $\tau(S_1) = T_1$. Suppose by induction that we have proved that $\tau(S_j) = T_j$ for all $j < i$. Then since τ is injective, the statement that $\tau(S_i) \subseteq \bigcup_{j \le i} T_j$ forces $\tau(S_i) \subseteq T_i$. Again by counting, since $\lambda_i = \mu_i$, $\tau(S_i) = T_i$. It follows that, t is obtained from $\tau \cdot s$ by some permutations of the rows. Thus $\tau \cdot [s] = [t]$. Then

$$\begin{aligned} A_t([s]) &= \sum_{\sigma \in C_t} \varepsilon(\sigma)\sigma \cdot [s] = \sum_{\sigma \in C_t} \varepsilon(\sigma)\sigma\tau^{-1} \cdot [t] \\ &= \varepsilon(\tau) \sum_{\sigma \in C_t} \varepsilon(\sigma \cdot \tau^{-1})\sigma\tau^{-1} \cdot [t] = \varepsilon(\tau) \sum_{\sigma \in C_t} \varepsilon(\sigma)\sigma \cdot [t] \\ &= \varepsilon(\tau) A_t([t]) = \varepsilon(\tau) e_t. \end{aligned}$$

Thus, if $A_t([s]) \neq 0$, then $A_t([s]) = \pm e_t$. □

Corollary 4.5. *If t is a λ-tableau, then $A_t(M^\lambda) = \mathbb{C} \cdot e_t$.*

Proof. We know that M^λ is spanned by the λ-tabloids $[s]$ and that $A_t([s])$ is either 0 or $\pm e_t$. Thus $A_t(M^\lambda) \subseteq \mathbb{C} \cdot e_t$. Finally, the image of A_t is $\mathbb{C} \cdot e_t$, as opposed to 0, since $A_t([t]) = e_t$. □

Returning to the proof of the submodule theorem, let V be an S_n-invariant subspace of M^λ. Then $\mathbb{C}[S_n](V) \subseteq V$ and hence $A_t(v) \in V$ for every λ-tableau t and every $v \in V$. As $A_t(M^\lambda) = \mathbb{C} \cdot e_t$, $e_t \in V$ as long as there exists a $v \in V$ such that $A_t(v) \neq 0$. In this case, since $\sigma \cdot e_t = e_{\sigma \cdot t}$ and V is S_n-invariant, $e_{\sigma \cdot t} \in V$ for all $\sigma \in S_n$. Since S_n acts transitively on the set of tableaux, $e_s \in V$ for every tableau s. As S^λ is generated by the e_s, $S^\lambda \subseteq V$.

Otherwise, $A_t(v) = 0$ for every tableau t and $v \in V$. Since the inner product $\langle \cdot, \cdot \rangle$ is S_n-invariant, $\langle \sigma(v), w \rangle = \langle v, \sigma^{-1}(w) \rangle$ for all $v, w \in M^\lambda$. Then, for all $v, w \in M^\lambda$, $\langle A_t(v), w \rangle = \langle v, A_t^*(w) \rangle$, where

$$A_t^* = \sum_{\sigma \in C_t} \varepsilon(\sigma)\sigma^{-1} = \sum_{\sigma \in C_t} \varepsilon(\sigma^{-1})\sigma^{-1} = \sum_{\sigma \in C_t} \varepsilon(\sigma)\sigma = A_t.$$

Thus $A_t(v) = 0$ for all $v \in V \implies \langle v, A_t(w)\rangle = 0$ for every $w \in M^\lambda$. Since the image of A_t is $\mathbb{C} \cdot e_t$, this implies that $\langle v, e_t\rangle = 0$ for every λ-tableau t. Since S^λ is the span of the e_t, $V \subseteq (S^\lambda)^\perp$ as claimed. □

Corollary 4.6. *S^λ is irreducible.*

Proof. Note that $S^\lambda \neq \{0\}$ as $e_t \neq 0$ for every t. If V is an S_n invariant subspace of S^λ, then by the submodule theorem either $S^\lambda \subseteq V$ or $V \subseteq (S^\lambda)^\perp$. In the first case, $V = S^\lambda$ since $V \subseteq S^\lambda$ and $S^\lambda \subseteq V$. In the second case, $V \subseteq S^\lambda \cap (S^\lambda)^\perp = \{0\}$. Thus every S_n-invariant subspace of S^λ is either S^λ or $\{0\}$, so that S^λ is irreducible. □

Corollary 4.7. *If $\operatorname{Hom}^{S_n}(S^\lambda, M^\mu) \neq 0$, then $\lambda \trianglerighteq \mu$. Moreover, if $\lambda = \mu$, then $\dim \operatorname{Hom}^{S_n}(S^\lambda, M^\lambda) = 1$. Thus the multiplicity of S^λ in M^λ is 1.*

Proof. If $F \neq 0$, then by Schur's lemma F is injective. Thus, for every λ-tableau t, $F(e_t) \neq 0$.

Since there is an S_n-invariant isomorphism

$$M^\lambda \cong S^\lambda \oplus (S^\lambda)^\perp,$$

we can extend F to an S_n-morphism $\widetilde{F} \colon M^\lambda \to M^\mu$ by setting $\widetilde{F} = F$ on S^λ and $\widetilde{F} = 0$ on $(S^\lambda)^\perp$. Since $\widetilde{F}$ is an S_n-morphism, it commutes with the action of $\mathbb{C}[S_n]$, so that $\widetilde{F} \circ A_t = A_t \circ \widetilde{F}$. But $A_t([t]) = e_t$, and hence

$$F(e_t) = \widetilde{F}(e_t) = \widetilde{F}(A_t([t])) = A_t(\widetilde{F}([t])).$$

We can write $\widetilde{F}([t])$ as a linear combination of μ-tabloids $[s]$. Since $F(e_t) \neq 0$, there must exist a μ-tabloid s such that $A_t([s]) \neq 0$. By Lemma 4.4, $\lambda \trianglerighteq \mu$. Moreover, if $\lambda = \mu$, then $A_t([s]) = \pm e_t$, so that $F(e_t) \in S^\lambda$ for all t. It follows that F is given by $i \circ G$, where $i \colon S^\lambda \to M^\lambda$ is the inclusion and $G \in \operatorname{Hom}^{S_n}(S^\lambda, S^\lambda)$. By Schur's lemma, $\operatorname{Hom}^{S_n}(S^\lambda, S^\lambda) = \mathbb{C}\,\mathrm{Id}$. Thus every S_n morphism from S^λ to M^λ is multiplication by a scalar, followed by inclusion, so that $\dim \operatorname{Hom}^{S_n}(S^\lambda, M^\lambda) = 1$. □

Corollary 4.8. *For all $\lambda, \mu \vdash n$, $S^\lambda \cong S^\mu$ as S_n-representations $\iff$ $\lambda = \mu$.*

Proof. Trivially, if $\lambda = \mu$, then $S^\lambda \cong S^\mu$. Conversely, suppose that $S^\lambda \cong S^\mu$. Then the composition of this isomorphism with the inclusion $S^\mu \subseteq M^\mu$ gives a nonzero element of $\operatorname{Hom}^{S_n}(S^\lambda, M^\mu) \neq 0$. The previous corollary then implies that $\lambda \trianglerighteq \mu$. By symmetry, $\mu \trianglerighteq \lambda$. Hence $\lambda = \mu$. □

5 Some concluding remarks

In this final section, we make some more remarks about the irreducible representations of S_n, mostly without proofs.

5.1 Rationality of the representations

As we have previously noted, if $\gcd(a, n!) = 1$ and $\sigma \in S_n$, then σ^a and σ are conjugate, and this implies that, for every representation V of S_n, the value of the character $\chi_V(\sigma)$ is an integer for every $\sigma \in S_n$. In fact, a stronger statement is true:

Theorem 5.1. *The irreducible representations S^λ of S_n are defined over $\mathbb{Q}$. Hence every representation of S_n can be defined over $\mathbb{Q}$.*

The main point of the proof is as follows. The representation M^λ is defined over $\mathbb{Q}$. In fact, $M^\lambda = \mathbb{C}[S_n/S_\lambda]$, with a basis consisting of the λ-tabloids $[t]$, and we can just take the corresponding $\mathbb{Q}$-vector space $M^\lambda_{\mathbb{Q}} = \mathbb{Q}[S_n/S_\lambda]$, with a $\mathbb{Q}$-basis consisting of the λ-tabloids $[t]$. Note that $\sigma \in S_n$ acts by permuting the tabloids, and hence the matrix corresponding to σ has rational entries, in fact every entry is either 0 or 1. The polytabloids e_t are also elements of $M^\lambda_{\mathbb{Q}}$, since they are linear combinations of certain tabloids with coefficients ± 1. Hence they span a vector subspace of M^λ which is also defined over $\mathbb{Q}$.

5.2 Explicit construction of some representations

We have already seen that the trivial representation $\mathbb{C}$ is isomorphic to $S^{(n)}$, that $\mathbb{C}(\varepsilon)$ is isomorphic to $S^{(1,\dots,1)}$, and that the standard representation V is isomorphic to $S^{(n-1,1)}$. Linear algebra can construct a few of the other irreducible representations directly. One basic linear algebra construction is exterior or alternating product: given a vector space U, we can construct a new vector space $\bigwedge^k U$, which is generated by expressions of the form $v_1 \wedge \cdots \wedge v_k$ which are multilinear in the v_i. For any collection $v_1, \dots, v_k$ of elements of U, we have the basic transformation law: for all $\sigma \in S_k$

$$v_{\sigma(1)} \wedge \cdots \wedge v_{\sigma(k)} = \varepsilon(\sigma) v_1 \wedge \cdots \wedge v_k.$$

If $u_1, \dots, u_d$ is a basis for U, then a basis for $\bigwedge^k U$ is given by:

$$\{u_{i_1} \wedge \cdots \wedge u_{i_k} : i_1 < \cdots < i_k\}.$$

In particular $\dim \bigwedge^k U = \binom{d}{k}$ for $k \leq d$, and $\dim \bigwedge^k U = 0$ for $k > d$. Then one can show:

Proposition 5.2. *For $k \leq n-1$, $\bigwedge^k V = \bigwedge^k S^{(n-1,1)}$ is an irreducible representation of S_n, and it is isomorphic to $S^{(n-k,1,\dots,1)}$.* □

An explicit proof is sketched in the HW.

Another construction of representations uses the symmetric product: given a vector space U, we can construct a new vector space $\operatorname{Sym}^k U$, which is generated by expressions of the form $v_1 \dots v_k$ which are multilinear in the v_i. For any collection $v_1, \dots, v_k$ of elements of U, we have the basic transformation law: for all $\sigma \in S_k$

$$v_{\sigma(1)} \dots v_{\sigma(k)} = v_1 \dots v_k.$$

If $u_1, \dots, u_d$ is a basis for U, then a basis for $\operatorname{Sym}^k U$ is given by:

$$\{u_{i_1} \cdots u_{i_k} : i_1 \leq \cdots \leq i_k\}.$$

In particular $\dim \operatorname{Sym}^k U = \binom{d+k-1}{k}$. It is then easy to check that, for $k \leq n/2$, there is an injective S_n-morphism $M^{(n-k,k)} \to \operatorname{Sym}^k V$. Hence $S^{(n-k,k)}$ is isomorphic to an S_n-invariant summand of $\operatorname{Sym}^k V$. For $n = 2$, it is easy to make this more explicit:

Proposition 5.3. $\operatorname{Sym}^2 V \cong \mathbb{C} \oplus V \oplus S^{(n-2,2)}$. □

In fact, one can identify the subspace $\mathbb{C} \oplus V$ explicitly as well and so give a concrete realization of $S^{(n-2,2)}$. Note that $\dim S^{(n-2,2)} = \dfrac{n(n-3)}{2}$.

5.3 Conjugate partitions

For every partition $\lambda \vdash n$, we have defined the transpose $\lambda^T \vdash n$, and $(\lambda^T)^T = \lambda$. Note that it is possible for $\lambda^T = \lambda$. For example, $(n)^T = (1, \dots, 1)$. For the representations S^λ, we have the following result, which generalizes $S^{(1,\dots,1)} = \mathbb{C}(\varepsilon)$:

Proposition 5.4. *$S^{\lambda^T} \cong S^\lambda \otimes \varepsilon$.* □

5.4 Representations of the alternating group

The alternating group A_n is a subgroup of S_n of index two, and so we can apply our general results about restrictions of irreducible representations to subgroups of index two:

Proposition 5.5. *Let $\lambda \vdash n$.*

(i) $\lambda = \lambda^T \iff S^\lambda \cong S^\lambda \otimes \varepsilon$. *In this case,*

$$\operatorname{Res}^{S_n}_{A_n} S^\lambda \cong \operatorname{Res}^{S_n}_{A_n}(S^\lambda \otimes \varepsilon) \cong W \oplus W',$$

where W and W' are two irreducible representations of A_n, with $\dim W = \dim W'$ and W, W' are not isomorphic.

(ii) $\lambda \neq \lambda^T \iff S^\lambda$ *and* $S^\lambda \otimes \varepsilon$. *In this case,*

$$\operatorname{Res}^{S_n}_{A_n} S^\lambda \cong \operatorname{Res}^{S_n}_{A_n}(S^\lambda \otimes \varepsilon)$$

is an irreducible representation of A_n.

Finally, every irreducible representation of A_n arises in this way.

Example 5.6. We consider the case $n = 5$. There are two 1-dimensional representations of S_5, $S^{(5)} \cong \mathbb{C}$ and $S^{(1,1,1,1,1)} \cong \mathbb{C}(\varepsilon)$. There are two 4-dimensional representations, the standard representation $V = S^{(4,1)}$ and $V \otimes \varepsilon = S^{(2,1,1,1)}$, where we have used the fact that $(4,1)^T = (2,1,1,1)$ and Proposition 5.4. Next, $S^{(3,2)}$ is an irreducible representation of dimension 5, and since $(3,2)^T = (2,2,1)$, we have $S^{(2,2,1)} \cong S^{(3,2)} \otimes \varepsilon$, also of dimension 5. Finally, $\bigwedge^2 V \cong S^{(3,1,1)}$ is irreducible of dimension 6. Since $(3,1,1)^T = (3,1,1)$, $\bigwedge^2 V \cong \bigwedge^2 V \otimes \varepsilon$, and this is the only irreducible representation up to isomorphism with this property.

As a check, we add up the sums of the squares of the irreducible representations constructed above:

$$1^2 + 1^2 + 4^2 + 4^2 + 5^2 + 5^2 + 6^2 = 120 = \#(S_5),$$

as expected.

We turn now to A_5. The representations $\mathbb{C}$ and $\mathbb{C}(\varepsilon)$ both restrict to the trivial representation of A_5. The representations V and $V \otimes \varepsilon$ both restrict to an irreducible representation of dimension 4, the restriction of the standard irreducible representation V to A_4. The representations $S^{(3,2)}$ and $S^{(2,2,1)} \cong S^{(3,2)} \otimes \varepsilon$ both restrict to an irreducible representation of dimension 5. Finally, the 6-dimensional representation $\bigwedge^2 V \cong S^{(3,1,1)}$ restricts on A_5 to $W \oplus W'$, where W and W' are two non-isomorphic irreducible representations of A_5. Finally, every irreducible representation of A_5 is one of these. As a check,

$$1^2 + 4^2 + 5^2 + 3^2 + 3^2 = 60 = \#(A_5).$$

With a little more effort, we can work out the character table for A_5. There are 5 conjugacy classes: all 3-cycles and products of two disjoint 2-cycles are conjugate in A_5, but there are two different conjugacy classes of 5-cycles (any two 5-cycles are conjugate in S_5, but not necessarily in A_5).

	1	$C((1,2,3))$	$C((12)(34))$	$C((12345))$	$C((21345))$
$\chi_{\mathbb{C}}$	1	1	1	1	1
χ_V	4	1	0	-1	-1
$\chi_{S^{3,2}}$	5	-1	1	0	0
χ_W	3	0	-1	$\frac{1+\sqrt{5}}{2}$	$\frac{1-\sqrt{5}}{2}$
$\chi_{W'}$	3	0	-1	$\frac{1-\sqrt{5}}{2}$	$\frac{1+\sqrt{5}}{2}$

Note: The images of the 3-dimensional representations W and W' can be realized as a subgroup of $SO(3)$, the *icosahedral group*. It is the group of symmetries of a regular dodecahedron, or equivalently of a regular icosahedron.

5.5 Further directions

There are many other questions one might ask about representations of S_n. Here are two:

Branching rules: The group S_n naturally contains S_{n-1} as a subgroup and in turn is naturally a subgroup of S_{n+1}. Given $\lambda \vdash n$ and the irreducible representation S^λ of S_n, we have the corresponding representation $\mathrm{Res}^{S_n}_{S_{n-1}} S^\lambda$ of S_{n-1} as well as the representation $\mathrm{Ind}^{S_{n+1}}_{S_n} S^\lambda$. These can both be described in terms of the Young diagram of λ.

Multiplication rules: Here, given $\lambda, \mu \vdash n$, the problem is to describe the the irreducible summands and their multiplicities of the representation $S^\lambda \otimes S^\mu$.

For a discussion of these and many other questions related to representations of S_n, we refer to the many books on S_n.