THE REPRESENTATIONS OF THE SYMMETRIC GROUP

LAUREN K. WILLIAMS

Abstract. In this paper we classify the irreducible representations of the symmetric group S_n and give a proof of the hook formula for the dimension of each irreducible.

1. The Irreducible Representations of S_n

We construct the irreducible representations of the symmetric group. The number of irreducible representations of S_n is the number of conjugacy classes of S_n, which is the number of partitions of n. Recall that a partition λ of n is $\lambda = (\lambda_1, \ldots, \lambda_k)$ where $n = \lambda_1 + \cdots + \lambda_k$ and $\lambda_1 \geq \cdots \geq \lambda_k \geq 1$.

To a partition $\lambda = (\lambda_1, \ldots, \lambda_k)$ we associate a *Young diagram* with λ_i boxes in the ith row, the rows of boxes left-justified. We define a *tableau* on a given Young diagram to be a numbering of the boxes by the integers $1, 2, \ldots, n$, and we will call it *standard* if the rows and columns are increasing sequences. Below are examples of a Young diagram, a standard tableau, and a non-standard tableau, for $n = 8$, and $\lambda = (3, 3, 2)$.

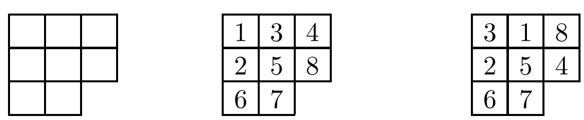

Young diagram   standard tableau   non-standard tableau

Figure 1

Given a tableau, we define two subgroups of the symmetric group as follows.

$$P = P_\lambda = \{g \in S_n \mid g \text{ preserves each row}\};$$
$$Q = Q_\lambda = \{g \in S_n \mid g \text{ preserves each column}\}.$$

Then we introduce two elements in the group algebra $\mathbb{C}[S_n]$ corresponding to the two subgroups:

$$a_\lambda = \sum_{g \in P} e_g \text{ and } b_\lambda = \sum_{g \in Q} \text{sign}(g) e_g.$$

We let $c_\lambda = a_\lambda b_\lambda$—this is called the *Young symmetrizer*.

After making some observations about the objects we have defined, we will be able to obtain each irreducible representation V_λ of S_n from the corresponding Young symmetrizer c_λ.

Notice that $P \cap Q = \{1\}$, so an element of S_n can be written in at most one way as a product pq for $p \in P$ and $q \in Q$. So $c = \sum \pm e_g$ where the sum is over all g that can be written as pq, and the coefficient of e_g is $\text{sign}(g)$.

Date: December 16, 1998.

Lemma 1.1. *For any $p \in P$ and $q \in Q$, we have $pc(\text{sign}(q)q) = c$.*

Proof. It's clear that for $p \in P$, we have $pa = ap = a$. Similarly, for $q \in Q$, we have $(\text{sign}(q)q)b = b(\text{sign}(q)q) = b$. Hence $pc(\text{sign}(q)q) = pab(\text{sign}(q)q) = ab = c$. □

Lemma 1.2. *Up to scalar multiples, c is the only element in $\mathbb{C}[S_n]$ with the above property.*

Proof. Suppose that $\sum n_g e_g \in \mathbb{C}[S_n]$ satisfies the property that for any $p \in P$ and $q \in Q$, $p(\sum n_g e_g)\,\text{sign}(q)q = \sum n_g e_g$. Then when we see what happens to the term $n_g e_g$ on the left-hand side, we get $pn_g e_g\,\text{sign}(q)q = \text{sign}(q)n_g e_{pgq}$, and comparing this to the corresponding term $n_{pgq}e_{pgq}$ on the right-hand side, we see that for all g, p, q,

$$n_{pgq} = \text{sign}(q)n_g. \tag{1}$$

In particular, $n_{pq} = \text{sign}(q)n_1$. So it is sufficient to show that $n_g = 0$ if $g \notin PQ$.

Claim 1.3. *For $g \notin PQ$, there is a transposition t such that $p = t \in P$ and $q = g^{-1}tg \in Q$.*

If this claim holds, then we'll get $g = pgq$ and by (1), $n_g = -n_g$. This implies that $n_g = 0$. So it remains to show the claim.

Note that if $T' = gT$ is the tableau obtained by letting g act on each entry of T, then the column stabilizer of T' is $Q' = gQg^{-1}$. And so the claim means that there are two distinct integers in the same row of T and the same column of T'; t is their transposition.

To verify the claim, suppose that there are not two such integers. Then all of the elements in the first row of T appear in different columns in T'—we can thus choose $q'_1 \in Q' = gQg^{-1}$ such that the first row of T and the first row of q'_1T' have the same elements—further, we can choose $p_1 \in P$ such that p_1T and q'_1T' have precisely the same first row. Repeating this process on the rest of the tableau, we get $p \in P$ and $q' \in Q'$ such that $pT = q'T'$. Then $pT = q'T$ so $p = q'g$ and $g = pq$ where $q = g^{-1}(q')^{-1}g \in Q$. Hence $g \in PQ$ and the claim is proved.

□

Corollary 1.4. *For any $x \in \mathbb{C}[S_n]$, $c_\lambda x c_\lambda$ is a scalar multiple of c_λ. In particular, $c_\lambda c_\lambda = n_\lambda c_\lambda$ for some $n_\lambda \in \mathbb{C}$.*

We will order partitions lexicographically, saying that $\lambda > \mu$ if the first nonvanishing $\lambda_i - \mu_i$ is positive.

Lemma 1.5. *If $\lambda > \mu$ then for all $x \in A$ we have $a_\lambda x b_\mu = 0$. In particular, if $\lambda > \mu$, then $c_\lambda c_\mu > 0$.*

Proof. It is sufficient to prove this for $x = g \in S_n$. Note that if T' is the tableau used to construct b_μ, then $gB_\mu g^{-1}$ is the "b-element" constructed from gT'. Also note that our results are dependent only on the partition chosen, not the tableau. So since $a_\lambda(gb_\mu g^{-1}) = 0$ implies that $a_\lambda g b_\mu = 0$, it suffices to show that $a_\lambda b_\mu = 0$. One can check that if $\lambda > \mu$ then there are two integers in the same row of T and the same column of T'. If t is the transposition of these integers, then $t \in P_\lambda$ and $t \in Q_\mu$. Hence $a_\lambda t = a_\lambda$ and $tb_\mu = -b_\mu$ so $a_\lambda b_\mu = a_\lambda t t b_\mu = -a_\lambda b_\mu$ which implies that $a_\lambda b_\mu = 0$. □

We can now obtain all irreducible representations of S_n.

Theorem 1.6. *The image of c_λ by right multiplication on $\mathbb{C}[S_n]$ is an irreducible representation V_λ of S_n, and every irreducible representation of S_n can be obtained in this way for a unique partition λ.*

For example, when $\lambda = (n)$, we get $P_\lambda = S_n$, $Q_\lambda = \{1\}$, $b_\lambda = 1$, and $c_\lambda = a_\lambda = \sum_{g \in S_n} e_g$, so $V_\lambda = \mathbb{C}[S_n] \sum_{g \in S_n} e_g = \mathbb{C} \sum_{g \in S_n} e_g$ is the trivial representation.

Proof. First we show that each V_λ is an irreducible representation of S_n. Since $V_\lambda = (\mathbb{C}[S_n])c_\lambda$, it is clear that V_λ is a $\mathbb{C}[S_n]$-invariant subspace, i.e., a representation. Now note that by Corollary 1.4, $c_\lambda V_\lambda \subset \mathbb{C}c_\lambda$. If $W \subset V_\lambda$ is a subrepresentation, then either $c_\lambda \in c_\lambda W$, in which case $c_\lambda W = \mathbb{C}c_\lambda$, or $c_\lambda \notin c_\lambda W$, in which case $c_\lambda W = 0$. If the former is true, then since W is a subrepresentation, $c_\lambda W \subset W$, and so in particular, $c_\lambda \in W$. Then we have $V_\lambda = (\mathbb{C}[S_n])c_\lambda \subset W$ and so $W = V_\lambda$. If the latter is true, then $WW \subset (\mathbb{C}[S_n])c_\lambda W = 0$ but this implies that $W = 0$. Hence V_λ is irreducible.

Now we will show that if $\lambda \neq \mu$, then V_λ and V_μ are not isomorphic. Without loss of generality we assume that $\lambda > \mu$. Then by Corollary 1.4, $c_\lambda V_\lambda = \mathbb{C}c_\lambda \neq 0$, but by Lemma 1.5, $c_\lambda V_\mu = c_\lambda A c_\mu = 0$, so they cannot be isomorphic $\mathbb{C}[S_n]$-modules.

Finally, note that these are all of the irreducible representations because we have as many irreducible representations as conjugacy classes. □

2. The Dimension of V^λ

We sketch a proof that the dimension of V^λ is equal to the number of standard tableaux on λ.

First note that for $\lambda = (\lambda_1, \dots, \lambda_k)$, a partition of n, we have a subgroup $S_\lambda = S_{\lambda_1} \times \cdots \times S_{\lambda_k}$ of S_n. We let U_λ be the representation of S_n obtained by inducing from the trivial representation of S_λ—this is equal to $\mathbb{C}[S_n]a_\lambda$. Hence V_λ appears in U_λ, as we have a surjection given by $x \mapsto xb_\lambda$. Observe that for $\lambda = (1, \dots, 1), U_\lambda$ is the regular representation, which contains every irreducible with multiplicity equal to its dimension. This motivates our study of U_λ.

The character ψ_λ of U_λ is straightforward to compute. If C_i is the conjugacy class in which each element consists of i_1 1-cycles, … , i_n n-cycles, then $\psi_\lambda(C_i)$ is the coefficient of the monomial $X^\lambda = x_1^{\lambda_1} \cdots x_k^{\lambda_k}$ in the polynomial

$$P^{(i_1,\dots,i_n)} = (x_1 + \cdots + x_k)^{i_1} \cdots (x_1^n + \cdots + x_k^n)^{i_n}.$$

Now for independent variables $x_1, \dots, x_k$, we define the power sums $P_j(x)$, and the discriminant $\Delta(x)$ by

$$P_j(x) = x_1^j + \cdots + x_k^j,$$
$$\Delta(x) = \prod_{i<j}(x_i - x_j).$$

We also define the coefficient

$$\omega_\lambda(i) = \left[\Delta P^{(i)}\right]_l, l = (\lambda_1 + k - 1, \lambda_2 + k - 2, \dots, \lambda_k)$$

which we will find is equal to $\chi_\lambda(C_i)$, the character of V_λ.

To show this, we use an identity which holds for any symmetric polynomial P,

$$[P]_\lambda = \sum_\mu K_{\mu\lambda}[\Delta P]_{(\mu_1+k-1,\dots,\mu_k)},$$

where the coefficients $K_{\mu\lambda}$ are the *Kostka numbers*. Given partitions λ and μ of n, the Kostka number $K_{\mu\lambda}$ is defined combinatorially as the number of ways to fill the boxes of the Young diagram for μ with λ_1 1's, up to λ_k k's, such that each row is a nondecreasing sequence and each column is an increasing sequence. In other words, $K_{\mu\lambda}$ is the number of *semistandard tableaux on* μ *of type* λ. In particular, $K_{\mu(1,\dots,1)}$ is the number of standard tableaux on λ.

One can show that

$$\psi_\lambda(C_i) = \sum_\mu K_{\mu\lambda}\omega_\lambda(i) = \omega_\mu(i) + \sum_{\mu>\lambda} K_{\mu\lambda}\omega_\mu(i) \tag{2}$$

and that the coefficients ω_λ, regarded as functions on the conjugacy classes of S_n, satisfy the same orthogonality relations as the irreducible characters of S_n. Using an inductive argument, it follows from (2) that $\omega_\lambda = \chi_\lambda$.

Corollary 2.1. *The integer $K_{\mu\lambda}$ is the multiplicity of the irreducible representation V_μ in the induced representation U_λ:*

$$U_\lambda \cong V_\lambda \oplus \bigoplus_{\mu>\lambda} K_{\mu\lambda} V_\mu.$$

When $\lambda = (1,\dots,1)$, it is clear that U_λ is the regular representation, and so $K_{\mu(1,\dots,1)}$, which is the number of standard tableaux on λ, is the dimension of V_μ.

3. The Hook Formula

We use the previous result that the number of standard Young tableaux of a given shape λ is the degree of the corresponding representation V_λ. We denote the dimension of V_λ by f^λ; the hook formula is an explicit formula for this dimension, and involves objects called *hooks*. We give a probabilistic proof of the hook formula, due to Greene, Nijenhuis, and Wilf [2].

Definition 3.1. *If $v = (i,j)$ is a node in the diagram of λ, then it has* hook

$$H_v = H_{i,j} = \{(i,j') \mid j' \geq j\} \cup \{(i',j) \mid i' \geq i\}$$

with corresponding hooklength $h_v = h_{i,j} = |H_{i,j}|$.

For example, in the diagram, the dotted cells are the hook $H_{1,2}$ with hooklength $h_{1,2} = 4$.

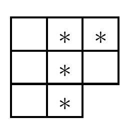

Figure 2

Theorem 3.2 (Hook Formula)**.** *If λ is a partition of n, then*

$$f^\lambda = \frac{n!}{\prod_{(i,j)\in\lambda} h_{i,j}}.$$

To prove this theorem, we will find an algorithm that always produces a standard λ-tableau. Moreover, it produces any given standard λ-tableau P with equal probability, $\frac{\prod h_{i,j}}{n!}$. And hence the number of standard λ-tableau is equal to the reciprocal of that probability.

The algorithm is as follows. Note that by an *inner corner* we mean a node of the tableau which has hooklength 1.

- Pick a random node $v \in \lambda$ (with probability $\frac{1}{n}$).
- While v is not an inner corner, repeat the following:
  - Pick a random node $\bar{v} \in H_v \setminus v$ (with probability $\frac{1}{h_v - 1}$).
  - Replace v with $\bar{v}$.
- Give the label n to the inner corner cell v that you've reached.
- Go back to the first step, replacing λ with $\lambda \setminus \{v\}$, and n with $n-1$, and repeat the loop until all cells are labeled.

We'll refer to the sequence of nodes generated by one pass through the algorithm (to reach an inner corner cell) as a *trial*. Notice that the x-coordinates of the nodes in a trial form a monotonically increasing sequence, as do the y-coordinates.

It's clear that our algorithm produces a standard labeling of λ. To show that all standard labelings have the same probability, we induct on n.

Consider a particular λ-tableau T. Let w be the cell containing n and let $\bar{T}$ be T with w deleted. Then $\text{prob}(T) = \text{prob}(w)\,\text{prob}(\bar{T})$ where $\text{prob}(w)$ is the probability that a trial ends at w. For our induction step, it will clearly suffice to show that

$$\text{prob}(w) = \frac{\frac{1}{n!}\prod_{v\in\lambda} h_v}{\frac{1}{(n-1)!}\prod_{\bar{v}\in\bar{\lambda}} h_{\bar{v}}}.$$

And everything in the above expression cancels except for terms corresponding to the following subset of λ:

$$S = \{v \neq w \mid w \in H_v\}.$$

Each hooklength in S is decreased by 1 so we can rewrite the previous expression as

$$\text{prob}(w) = \frac{1}{n}\prod_{v\in S}\frac{h_v}{h_v - 1} = \frac{1}{n}\prod_{v\in S}\left[1 + \frac{1}{h_v - 1}\right].$$

Now let w have coordinates (α, β) and notice that $v \in S$ means that v is in the same row or column as w. So if we let $a_i = h_{i,\beta} - 1$, and $b_j = h_{\alpha,j} - 1$, then we can rewrite the preceding expression as

$$\text{prob}(w) = \frac{1}{n}\prod_{i=1}^{\alpha-1}\left(1 + \frac{1}{a_i}\right)\prod_{j=1}^{\beta-1}\left(1 + \frac{1}{b_j}\right) = \frac{1}{n}\sum_{I,J}\prod_{i\in I}\frac{1}{a_i}\prod_{j\in J}\frac{1}{b_j}$$

where I ranges over the subsets of $[\alpha - 1]$ and J ranges over the subsets of $[\beta - 1]$. (Note that we use $[n]$ to denote the set $\{1, \ldots, n\}$.)

We can now interpret each term above in the following way.

Definition 3.3. *Given a trial ending at (α, β), we let the* horizontal projection *of the trial be*

$$I = \{i \neq \alpha \mid v = (i, j) \text{ for some } v \text{ on the trial}\}.$$

The *vertical projection* is defined analogously. Let $\text{prob}_{I,J}(\alpha, \beta)$ denote the sum of the probabilities of all trials ending at (α, β) with horizontal projection I and vertical projection J. So $\text{prob}(\alpha, \beta) = \sum_{I,J}\text{prob}_{I,J}(\alpha, \beta)$ where I ranges over the subsets of $[\alpha-1]$ and J ranges over the subsets of $[\beta-1]$. Now, to prove the theorem, it is sufficient to prove the following lemma:

Lemma 3.4. *Given (α,β), I a subset of $[\alpha-1]$, and J a subset of $[\beta-1]$,*

$$\text{prob}_{I,J}(\alpha,\beta)=\frac{1}{n}\prod_{i\in I}\frac{1}{a_i}\prod_{j\in J}\frac{1}{b_j}.$$

Proof. We use induction on $|I\cup J|$. First, assume that I or J is empty. WLOG $J=\emptyset$. Then if we order the terms of I in an increasing sequence $(i_1,i_2,\dots)$, the only possible trial is the sequence $\{(i_1,\beta),(i_2,\beta),\dots\}$. The probability of randomly choosing (i_1,β) according to the algorithm is $\frac{1}{n}$. Similarly, the probability of choosing (i_2,β) is $\frac{1}{a_{i_1}}$, and so on. So

$$\text{prob}_{I,J}(\alpha,\beta)=\frac{1}{na_{i_1}a_{i_2}\cdots}$$

as desired.

Now assume that neither I nor J is empty. Then we must have $v_1=(i_1,j_1)$ and there are exactly two choices for v_2—(i_2,j_1) or (i_1,j_2)—each of which is chosen with probability $\frac{1}{h_{i_1,j_1}}$. Let $\bar{I}=I\setminus\{i_1\}$ and $\bar{J}=J\setminus\{j_1\}$. Then

$$\begin{aligned}\text{prob}_{I,J}(\alpha,\beta)&=\frac{1}{h_{i_1,j_1}-1}[\text{prob}_{\bar{I},J}(\alpha,\beta)+\text{prob}_{I,\bar{J}}(\alpha,\beta)]\\&=\frac{1}{h_{i_1,j_1}-1}\left[\frac{1}{n\widehat{a_{i_1}}a_{i_2}\cdots b_{j_1}b_{j_2}\cdots}+\frac{1}{na_{i_1}a_{i_2}\cdots\widehat{b_{j_1}}b_{j_2}\cdots}\right]\\&=\frac{a_{i_1}+b_{j_1}}{h_{i_1,j_1}-1}\left[\frac{1}{na_{i_1}a_{i_2}\cdots b_{j_1}b_{j_2}\cdots}\right]\end{aligned}$$

where $\widehat{\cdot}$ means that the factor is deleted. But $h_{i_1,j_1-1}=(h_{i_1,\beta}-1)+(h_{\alpha,j_1}-1)=a_{i_1}+b_{j_1}$, so we are done. □

References

1. W. Fulton and J. Harris, *Representation theory: A first course*, Graduate Texts in Mathematics, no. 129, Springer-Verlag, New York, US, 1991.
2. Bruce Sagan, *The symmetric group*, Cole Mathematics Series, Wadsworth & Brooks, New York, US, 1991.

Department of Mathematics, Harvard University, Cambridge, MA 02138

Current address: 1624 Cataluna Place, Palos Verdes Estates, CA 90274

E-mail address: `lkwill@fas.harvard.edu`