Chapter 7. Representation Theory of the Symmetric Group

We have already built three irreducible representations of the symmetric group: the trivial, alternating and $n-1$ dimensional representations in Chapter 2. In this chapter we build the remaining representations and develop some of their properties.

To motivate the general construction, consider the space X of the unordered pairs $\{i, j\}$ of cardinality $\binom{n}{2}$. The symmetric group acts on these pairs by $\pi\{i, j\} = \{\pi(i), \pi(j)\}$. The permutation representation generated by this action can be described as an $\binom{n}{2}$ dimensional vector space spanned by basis vectors $e_{\{i,j\}}$. This space splits into three irreducibles: A one-dimensional trivial representation is spanned by $\overline{v} = \Sigma\ e_{\{i,j\}}$. An $n-1$ dimensional space is spanned by $v_i = \Sigma_j\ e_{\{i,j\}} - c\overline{v}, 1 \leq i \leq n$, with c chosen so v_i is orthogonal to $\overline{v}$. The complement of these two spaces is also an irreducible representation. A direct argument for these assertions is given at the end of Section A. The arguments generalize. The following treatment follows the first few sections of James (1978) quite closely. Chapter 7 in James and Kerber (1981) is another presentation.

A. Construction of the irreducible representations of the symmetric group.

There are various definitions relating to diagrams, tableaux, and tabloids. Let $\lambda = (\lambda_1, \ldots, \lambda_r)$ be a partition of n. Thus, $\lambda_1 \geq \lambda_2 \ldots \geq \lambda_r$ and $\lambda_1 + \ldots + \lambda_r = n$. The *diagram* of λ is an ordered set of boxes with λ_i boxes in row i. If $\lambda = (3, 3, 2, 1)$, the diagram is

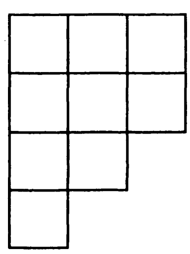

If λ and μ are partitions of n we say λ *dominates* μ, and write $\lambda \trianglerighteq \mu$, provided that $\lambda_1 \geq \mu_1, \lambda_1 + \lambda_2 \geq \mu_1 + \mu_2, \ldots$, etc. This partial order is widely used in various areas of mathematics. It is sometimes called the order of *majorization.* There is a book length treatment of this order by Marshall and Olkin (1979). They show that $\lambda \trianglerighteq \mu$ if and only if we can move from the diagram of λ to the diagram of μ by moving blocks from the right hand edge upward, one at a time, such that at each stage the resulting configuration is the diagram of a partition. Thus, $(4, 2) \triangleright (3, 3)$, but $(3, 3)$, and $(4, 1, 1)$ are not comparable. See Hazewinkel and Martin (1983) for many novel applications of the order.

A λ-*tableau* is an array of integers obtained by placing the integers from 1 through n into the diagram for λ. Clearly there are $n!$ λ-tableaux.

The following lemma is basic:

LEMMA 0. *Let λ and μ be partitions of n, suppose that t_1 is a λ-tableau and t_2 is a μ-tableau. Suppose that for each i the numbers from the ith row of t_2 belong to different columns of t_1. Then $\lambda \trianglerighteq \mu$.*

Proof. Since the numbers in the first row of t_2 are in different columns of t_1, $\lambda_1 \geq \mu_1$. The numbers in the second row of t_2 are in distinct columns of t_1, so no column of t_1 can have more than two of the numbers in the first or second row of t_2. Imagine "sliding these numbers up to the top of the columns of t_1." They fit in the first two rows, so $\lambda_1 + \lambda_2 \geq \mu_1 + \mu_2$. In general, no column of t_1 can have more than i of the numbers from the first i rows of t_2. $\square$

If t is a tableau, its *column-stabilizer* C_t is the subgroup of S_n keeping the columns of t fixed. For example, when $t = \begin{matrix} 1 & 2 & 4 & 5 \\ 3 & 6 & 7 & \\ 8 & & & \end{matrix}$, $C_t \cong S_{\{138\}} \times S_{\{26\}} \times S_{\{47\}} \times S_{\{5\}}$. The notation $S_{\{i,j,\ldots,k\}}$ means the subgroup of S_n permuting only the integers in brackets.

Define an equivalence relation on the set of λ-tableaux by considering $t_1 \sim t_2$ if each row of t_1 contains the same numbers as the corresponding row of t_2. The *tabloid* $\{t\}$ is the equivalence class of t. Think of a tabloid as a "tableau with unordered row entries." The permutations operate on the tabloids in the obvious way. The action is transitive and the subgroup stablizing the tabloid with $1, \ldots, \lambda_1$ in the first row, $\lambda_1 + 1, \ldots, \lambda_1 + \lambda_2$, in the second row, etc., is

$$S_{\{1,2,\ldots,\lambda_1\}} \times S_{\{\lambda_1+1,\ldots,\lambda_1+\lambda_2\}} \times \cdots.$$

It follows that there are $n!/\lambda_1! \ldots \lambda_r!$ λ-tabloids.

Define the *permutation representation* associated to the action of S_n on tabloids as a vector space with basis $e_{\{t\}}$. It is denoted M^λ. This representation is reducible but contains the irreducible representation we are after. To get this, define for each tableau t a *polytabloid* $e_t \in M^\lambda$ by

$$e_t = \sum_{\pi \in C_t} \operatorname{sgn}(\pi) e_{\pi\{t\}}.$$

Check that $\pi e_t = e_{\pi t}$, so the subspace of M^λ spanned by the $\{e_t\}$ is invariant under S_n (and generated as an "S_n module" by any e_t). It is called the *Specht module* S^λ. The object of the next collection of lemmas is to prove that S^λ is irreducible and that all the irreducible representations of S_n arise this way. These lemmas are all from Section 4 of James.

LEMMA 1. *Let λ and μ be partitions of n. Suppose that t is a λ tableau and s is a μ-tableau. Suppose that*

$$\sum_{\pi \in C_t} sgn(\pi) e_{\pi\{s\}} \neq 0.$$

Then $\lambda \trianglerighteq \mu$, *and if* $\lambda = \mu$, *the sum equals* $\pm e_t$.

Proof. Suppose for some a, b that a and b are in the same row of s and in the same column of t. Then

$$(\mathrm{id} - (ab))e_{\{s\}} = e_{\{s\}} - e_{\{s\}} = 0.$$

Since a and b are in the same column of t, the group $<\mathrm{id}, (ab)>$ is a subgroup of C_t. Let $\sigma_1, \ldots, \sigma_k$ be coset representatives, so $\sum_{\pi \in C_t} \mathrm{sgn}(\pi) e_{\pi\{s\}} = \sum_{i=1}^{k} \mathrm{sgn}(\sigma_i)\sigma_i\{\mathrm{id} - (ab)\}e_{\{s\}} = 0$. This is ruled out by hypothesis, so the numbers in the ith row of s are in different columns of t. Lemma 0 implies that $\lambda \trianglerighteq \mu$.

Suppose $\lambda = \mu$, and the sum does not vanish; then, again, numbers in the ith row of s appear in different columns of t. It follows that for a unique $\pi^* \in C_t, \pi^*\{t\} = \{s\}$ and this implies that the sum equals $\pm e_t$ (replace $\{s\}$ by $\pi^*\{t\}$ in the sum). □

LEMMA 2. *Let* $\mu \in M^{\mu}$, *and let* t *be a* μ *tableau. Then for some scalar* c

$$\sum_{\pi \in C_t} sgn(\pi)\pi u = c\ e_t.$$

Proof. u is a linear combination of $e_{\{s\}}$. For $u = e_{\{s\}}$, Lemma 1 gives the result with $c = 0$ or ± 1. □

Now put an inner product on M^{μ} which makes $e_{\{s\}}$ orthonormal: $< e_{\{s\}}, e_{\{t\}} > = 1$ if $\{s\} = \{t\}$ and 0 otherwise. This is S_n invariant. Consider the "operator" $A_t = \sum_{\pi \in C_t} \mathrm{sgn}(\pi)\pi$. For any $u, v \in M^{\mu}$, $< A_t u, v > = \Sigma\, \mathrm{sgn}\, \pi < \pi u, v > = \Sigma\, \mathrm{sgn}\, \pi < u, \pi^{-1} v > = < u, A_t v >$. Using this inner product we get:

LEMMA 3. *(Submodule Theorem). Let* U *be an invariant subspace of* M^{λ}. *Then either* $U \supset S^{\lambda}$ *or* $U \subset S^{\lambda\perp}$. *In particular,* S^{λ} *is irreducible.*

Proof. Suppose $u \in U$ and t is a λ-tableau. By Lemma 2, $A_t u$ is a constant times e_t. If we can choose u and t such that this constant is non-zero, then $e_t \in U$ and, since $\pi e_t = e_{\pi t}$, $S^{\lambda} \subset U$. If $A_t u = 0$ for all t and u, then $0 = < A_t u, e_{\{t\}} > = < u, A_t\ e_{\{t\}} > = < u, e_t >$. So $U \subset S^{\lambda\perp}$. □

At this stage we have one irreducible representation for each partition λ of n. The number of irreducible representations is the same as the number of conjugacy classes: see Theorem 7 of Chapter 2. This number is also the number of partitions of n as explained at the beginning of Chapter 2D. Hence, if we can show that the S^{λ} are all inequivalent, we have finished determining all of the irreducible representations of S_n.

LEMMA 4. *Let $T: M^\lambda \to M^\mu$ be a linear map that commutes with the action of S_n. Suppose that $S^\lambda \not\subset \ker T$. Then $\lambda \trianglerighteq \mu$. If $\lambda = \mu$, then the restriction of T to S^λ is a scalar multiple of id.*

Proof. By lemma 3, Ker $T \subset S^{\lambda\perp}$. Thus for any $t, 0 \neq Te_t = TA_t e_{\{t\}} = A_t Te_{\{t\}}$. But $Te_{\{t\}}$ is a linear combination of μ-tabloids $e_{\{s\}}$ and for at least one such $e_{\{s\}}$, $A_t\, e_{\{s\}} \neq 0$. By Lemma 1, $\lambda \trianglerighteq \mu$. If $\lambda = \mu$, then $Te_t = c\, e_t$ by the same argument. □

LEMMA 5. *Let $T: S^\lambda \to S^\mu$ be a linear map that commutes with the action of S_n. If $T \neq 0$, $\lambda \trianglerighteq \mu$.*

Proof. Any such T can be extended to a linear map from M^λ to M^μ by defining T to be 0 on $S^{\lambda\perp}$. The extended map commutes with the action of S_n. If $T \neq 0$, then Lemma 4 implies $\lambda \trianglerighteq \mu$. □

Theorem 1. *The S^λ are all of the irreducible representations of S_n.*

Proof. If S^λ is equivalent to S^μ, then, using Lemma 5 in both directions, $\lambda = \mu$. □

Remark. The argument for Lemma 4 shows that the irreducible representations in M^μ are S^μ (once) and some of $\{S^\lambda: \lambda \trianglerighteq \mu\}$ (possibly with repeats). In fact, S^λ occurs in M^μ if and only if $\lambda \trianglerighteq \mu$.

To complete this section, here is a direct proof of the decomposition of $M^{n-2,2}$ discussed in the introductory paragraph to this chapter. We begin with a broadly applicable result.

A USEFUL FACT.

Let G be a finite group acting on a set X. Extend the action to the product space X^k coordinatewise. The number of fixed points of the element $s \in G$ is $F(s) = |\{x: sx = x\}|$. For any positive integer k:

(1) $\frac{1}{|G|} \sum_s F(s)^k = |$ orbits of G acting on $X^k|$.

(2) Let R, V be the permutation representation associated to X. Thus V has as basis δ_x and $R_s(\delta_x) = \delta_{sx}$. The character of this representation is $\chi_R(s) = F(s)$. If R decomposes into irreducibles as $R = m_1\rho_1 \oplus \ldots \oplus m_n\rho_n$; then

$$\Sigma m_i^2 = |\text{ orbits of } G \text{ acting on } X^2|.$$

Proof. For (1) we have the action of G on X^k given by $s(x_1, \ldots, x_k) = (sx_1, \ldots, sx_k)$. Let C_i be a decomposition of X^k into G orbits. Then

$$\sum_s f(s)^k = \sum_s \sum_{x_1} \delta_{sx_1}(x_1) \ldots \sum_{x_k} \delta_{sx_k}(x_k) = \sum_{X^k} \sum_s \delta_{s\underset{\sim}{x}}(\underset{\sim}{x})$$
$$= \sum_i \sum_{\underset{\sim}{x} \in C_i} \sum_s \delta_{s\underset{\sim}{x}}(\underset{\sim}{x}).$$

The innermost sum is the cardinality of the stabilizer of G at $\underset{\sim}{x}$: $|N_{\underset{\sim}{x}}|$ with $N_{\underset{\sim}{x}} = \{s: s\underset{\sim}{x} = \underset{\sim}{x}\}$. Observe $N_{s\underset{\sim}{x}} = s\ N_{\underset{\sim}{x}} s^{-1}$. In particular, the size of $N_{\underset{\sim}{x}}$ doesn't depend on the choice of $\underset{\sim}{x}$ in a given orbit. Since $|G| = |N_{\underset{\sim}{x}}|\ |C_i|$ the inner sum equals $|G|/|C_i|$. The sum over $\underset{\sim}{x} \in C_i$ multiplies this by $|C_i|$. The final sum yields $|G| \cdot |\text{Orbits}|$ as required. To prove (2), we use the orthogonality of characters: $\chi_R = m_1\chi_1 + \ldots + m_n\chi_n$ so $< \chi_R|\chi_R >= m_1^2 + \ldots + m_n^2$. On the other hand, it is clear $\chi_R(s) = F(s)$ and $F(s^{-1}) = F(s)$, so $< \chi_R|\chi_R >= \frac{1}{|G|}\Sigma F(s)^2$. □

REMARKS AND APPLICATIONS

(a) With $k = 1$, part (1) is called Burnside's lemma. It is at the heart of Serre's exercise 2.6 which we have found so useful. It also forms the basis of the Polya-Redfield "theory of counting." See e.g., de Bruijn (1964).

(b) If G acts doubly transitively on X, then there are two orbits of G acting on $X \times X$: $\{(x,x)\}$ and $\{(x,y): y \neq x\}$. It follows that V decomposes into two irreducible components: One of these is clearly spanned by the constants. Thus its complement $\{v: \Sigma v_i = 0\}$ is irreducible.

(c) When G acts on itself we get back the decomposition of the regular representation.

(d) There is an amusing connection with probability problems. If G is considered as a probability space under the uniform distribution U, then $F(s)$ is a "random variable" corresponding to "pick an element of G at random and count how many fixed points it has." When $G = S_n$ and $X = \{1, 2, \ldots, n\}$, $F(\pi)$ is the number of fixed points of π. We know that this has an approximate Poisson distribution with mean 1. Part (1) gives a "formula" for all the moments of $F(g)$.

EXERCISE 1. Using (1), prove that the first n moments of $F(\pi)$ equal the first n moments of Poisson(1), where π is chosen at random on S_n.

(e) Let us decompose $M^{n-2,2}$. The space X is the set of unordered pairs $\{i,j\}$ with $\pi\{i,j\} = \{\pi(i), \pi(j)\}$. The permutation representation has dimension $\binom{n}{2}$. There are 3 orbits of S_n acting on $X \times X$ corresponding to pairs $\{i,j\}, \{k,\ell\}$ with 0, 1, or 2 integers in common. Thus, clearly S_n acts transitively on the set of pairs $(\{i,j\},\{i,j\})$. Also for $(\{i,j\},\{j,\ell\})$ $\ell \neq i,j$ and for $(\{i,j\},\ \{k,\ell\})$ with $\{k,\ell\} \cap \{i,j\} = \phi$. It follows that V splits into 3 irreducible subspaces. These are, the 1-dimensional space spanned by $\overline{v} = \Sigma e_{\{ij\}}$, the $n-1$-dimensional space spanned by $\overline{v}_i = \Sigma_j e_{\{ij\}} - c\overline{v}$ $1 \leq i \leq n$, and the complement of these two spaces. Clearly, the space spanned by $\overline{v}$ gives the trivial representation and the space spanned by $\overline{v}_i$ gives the $n-1$ dimensional representation. What is left is $n(n-3)/2$ dimensional. If we regard the permutation representation as the set of all functions on X with $sf(x) = f(s^{-1}x)$, then the trivial and $n-1$ dimensional representations are the set of functions of form $f\{i,j\} = f_1(i) + f_1(j)$.

EXERCISE 2. Show that for fixed j, $0 \leq j \leq n/2$, $M^{n-j,j}$ splits into $j+1$ distinct irreducible representations, the ith having dimension $\binom{n}{i} - \binom{n}{i-1}$. Hint: use the useful fact and induction, (e) above is the case $j = 2$.

We can build some new irreducible representations directly by tensoring the representation we know about with the alternating representation. Tensoring the alternating representation with the $n-1$ dimensional representation always gives a different irreducible representation. For $n = 4$ we already have all irreducible representations: 2 of 1 dimension, 2 of 3 dimensions and 1 of dimension $n(n-3)/2 = 2$. The sum of squares adds to 24. For $n > 4$ (but not $n = 4$) the $n(n-3)/2$ dimensional representation yields a new irreducible representation of the same dimension. For $n = 5$ this gives all the irreducible representations but 1. We can build this by considering the action of S_n on ordered pairs (i, j). That is, $M^{3,1,1}$.

B. More on representations of S_n

The books by James and James-Kerber are full of interesting and useful facts. Here is a brief description of some of the most useful ones, along with pointers to other work on representations of S_n.

(1) *The Standard Basis of* S^λ. We have defined S^λ as the representation of M^λ generated by elements e_t. There are $n!$ different e_t and the dimension of S^λ can be quite small. For example, if $\lambda = (n-1, 1)$, we know S^λ is $n-1$ dimensional. It turns out that a few of the e_t generate S^λ. Define t to be a *standard tableau* if the numbers increase along the rows and down the columns. Thus $\begin{array}{|c|c|c|}\hline 1 & 3 & 5 \\ \hline 2 & 4 \\ \cline{1-2}\end{array}$ is a standard $[3, 2]$ tableau. There is only 1 standard (n) tableau. There are $n-1$ standard $(n-1, 1)$ tableaux. In Section 8, James proves that $\{e_t | t$ is a standard λ-tableau$\}$ is a basis for S^μ. This is a beautiful result, but not so helpful in "really understanding S^λ." What one wants is a set of objects on which S_n acts that are comprehensible. The graphs in Section 5 of James are potentially very useful in this regard for small n. As far as I know, a "concrete" determination of the representations of S_n is an open problem. See (6) below.

(2) *The Dimension of* S^λ. There are a number of formulas for the dimension (and other values of the character) of the representation associated to λ. The dimensions get fairly large; they are bounded by $\sqrt{n!}$ of course, but they get quite large:

n	2	3	4	5	5	7	8	9	10
max dim	1	2	3	6	16	35	90	216	768

We know that $\dim(S^\lambda)$ equals the number of ways of placing the numbers $1, \ldots, n$ into the Young diagram for λ in such a way that the numbers increase along rows and down columns. From this follows bounds like the following which was so useful in Chapter 3:

$$\dim(S^\lambda) \leq \binom{n}{\lambda_1} \sqrt{(n-\lambda_1)!}\,.$$

There is a classical determinant formula (James, Corollary 19.5)

$$\dim(S^\lambda) = n!\ \mathrm{Det}|1/(\lambda_i - i + j)!|, \text{ where } 1/r! = 0 \text{ if } r < 0.$$

Finally, there is the hook formula for dimensions. Let λ be a partition of n. The (i, j) *hook* is that part of the Young diagram that starts at (i, j) and goes as far as it can either down or to the right. Thus, if $\lambda = (4,3,2)$,

	x	x	
	x		

the (2,2) hook is indicated by $x's$. The length of the (i,j) hook is the number of boxes in the hook. Using this terminology, the hook length formula says

$$\dim(S^{\lambda}) = n!/\text{product of hook lengths in } \lambda.$$

For example, when $\lambda = (4,3,2)$, the hook lengths are

6	5	3	1
4	3	1	
2	1		

The dimension is $9!/6! \cdot 3 = 168$.

The dimension of $S^{(n-1,1)}$ is $n-1$. The dimension of $S^{11\cdots 1}$ is 1.

Greene, Nijenhuis, and Wilf (1979, 1984) give an elegant, elementary proof of the hook length formula involving a random "hook walk" on a board of shape λ.

Hooks come into several other parts of representation theory – in particular, the Murnaghan-Nakayama rule for calculating the value of a character (section 21 of James).

(3) *Characters of the Symmetric Group.* To begin, we acknowledge a sad fact: there is no reasonable formula for the character $\chi_\lambda(\mu)$ where λ is a partition of n, χ_λ the associated irreducible character of S_n, and μ stands for a conjugacy class of S_n. This is countered by several facts.

(a) For small n (≤ 15) the characters have been explicitly tabulated. James-Kerber (1981) give tables for $n \leq 10$ and references for larger n.

(b) For large n there are useful asymptotic results for $\chi_\lambda(\mu)$. These are clearly explained in Flatto, Odlyzko, and Wales (1985).

(c) For any specific λ and μ there is an efficient algorithm for calculating the character called the Murnaghan-Nakayama rule. Section 21 of James (1978) or Theorem 2.4.7 of James-Kerber (1981) give details.

EXERCISE 3. Define a probability Q on S_n as follows: with probability p_n choose the identity with probability $1-p_n$ choose a random n-cycle. Determine the rate of convergence. How should p_n be chosen to make this as fast as possible? Hint: See (d) below.

(d) For "small" conjugacy classes μ, and arbitrary n and λ, there are formulas like Frobenius' formula used in Chapter 3D. Ingram (1950) gives useful references. See also Formula 2.3.17 in James-Kerber (1981).

(e) For some special shapes of λ, such as hooks $\lambda = (k,1,1,\ldots,1)$, closed form formulas are known, see e.g. Stanley (1983) or Macdonald (1979).

(f) There are also available rather intractable generating functions for the characters due to Frobenius. This analytic machinery is nicely presented in Chapter 1.7 of Macdonald (1979).

(4) *The Branching Theorem* (Section 9 of James). Consider ρ the $n-1$ dimensional representation of S_n. Let S_{n-1} be considered as a subgroup of S_n (all permutations that fix 1). Then ρ is a representation of S_{n-1}, "by restriction" James writes $S^{(n-1,1)} \downarrow S_{n-1}$. Observe that ρ restricted to S_{n-1} is reducible. If we choose the basis $e_1 - e_2, e_1 - e_3, \ldots, e_1 - e_n$; then the sum of the basis elements generates a one-dimensional invariant subspace. Since S_{n-1} operates doubly transitively on the basis elements, we have $\rho \downarrow S_{n-1}$ splitting into two irreducible subspaces; one of dimension 1 and one of dimension $n-2$.

The branching theorem gives the general result on how $S^{\mu} \downarrow S_{n-1}$ decomposes: there is one irreducible representation for each way of removing a "box" from the right hand side of the Young diagram for μ in such a way that the resulting configuration is a diagram for a partition. Thus, the diagram for $[n-1,1]$ can be reduced to $(n-1)$ or $(n-2,1)$ and these are the two irreducible components. The branching theorem is used to give a fast Fourier transform for computing all $\hat{f}(\rho)$ in Diaconis and Rockmore (1988).

EXERCISE 4. (Flatto, Odlyzko, Wales). Let ρ be an irreducible representation of S_n. Show that ρ restricted to S_{n-1} splits in a multiplicity free way. Using this, show that if P is a probability on S_n that is invariant under conjugation by S_{n-1} (so $P(\pi) = P(\sigma\pi\sigma^{-1})$ for $\sigma \in S_{n-1}$), then $\hat{P}(\rho)$ is diagonal for an appropriate basis which does not depend on P.

(5) *Young's Rule.* This gives a way to determine which irreducible subspaces occur in the decomposition of M^{λ}. It will be extremely useful in Chapter 8 in dealing with partially ordered data "in configuration λ." For example, data of the form "pick the best m of n" can be regarded as a vector in $M^{(n-m,m)}$, the components being the number of people who picked the subset corresponding to the second row of the associated tabloid. The decomposition of $M^{(n-m,m)}$ into irreducibles gives us a spectral decomposition of the frequencies and a nested sequence of models. See Chapter 8B and 9A.

Young's rule depends on the notion of semi-standard tableaux. This allows repeated numbers to be placed in a diagram. Let λ and μ be partitions of n. A *semi-standard tableau* of *shape* λ and *type* μ is a placement of integers $\leq n$ into a Young tableau of shape λ, with numbers nondecreasing in rows and strictly increasing down columns, such that the number i occurs μ_i times. Thus, if $\lambda = (4,1)$ and $\mu = (2,2,1)$, there are two tableaux of shape λ and type μ:

$$\begin{array}{cccc} 1 & 1 & 2 & 2 \\ 3 & & & \end{array} \qquad \begin{array}{cccc} 1 & 1 & 2 & 3 \\ 2 & & & \end{array}$$

Young's Rule: The multiplicity of S^{λ} in M^{μ} equals the number of semi-standard λ tableaux of type μ. As an example, consider, for $m \leq n/2$, $\mu = (n-m, m)$. We are decomposing M^{μ}. The possible shapes λ are

$$\begin{array}{c} \overbrace{1\,1\ldots1}^{n-m}\overbrace{2\ldots2}^{m},\overbrace{1\,1\ldots1}^{n-m}\overbrace{2\ldots2}^{m-1},\ldots \\ \hline 2 \end{array} \begin{array}{cccccccc} 1 & 1 & \ldots & 1 & 1 & 1 & \ldots & 1 \\ 2 & 2 & \ldots & 2 & 2 & & & \end{array}$$

Each occurs once only. Thus $M^{(n-m,m)} = S^{(n)} \oplus S^{(n-1,1)} \oplus S^{(n-2,2)} \oplus \ldots \oplus S^{(n-m,m)}$. By induction dim $S^{(n-m,m)} = \binom{n}{m} - \binom{n}{m-1}$. When we translate this decomposition into an interpretation for "best m out of n" data, the subspaces $S^{(n-m,m)}$ have interpretations:

S^n – The grand mean or # of people in sample.

$S^{n-1,1}$ – The effect of item i, $1 \leq i \leq n$.

$S^{n-2,2}$ – The effect of items $\{i, j\}$ adjusted for the effect of i and j.

$\vdots$

$S^{n-k,k}$ – The effect of a subset of k items adjusted for lower order effects.

Remarks. Many further examples of Young's rule appear in Chapter 8. Young's rule does not give an algorithm for decomposing M^{μ} or interpreting the S^{μ}. It just says what pieces appear. Section 17 of James (1978) solves both of these problems in a computationally useful way. This remark is applied in Chapter 8C below.

Young's rule is a special case of the Littlewood-Richardson rule which describes how a given representation of S_n restricts to the subgroup $S_k \times S_{n-k}$. See James and Kerber (1981, Sec. 2.8).

(6) *Kazhdan-Lusztig Representations.* The construction of the irreducible representations given in Section A constructs S^{λ} as a rather complicated subspace of the highly interpretable M^{λ}. Even using the standard basis ((1) above), S^{λ} is spanned by the mysterious Young symmetrizers e_t. It is desirable to have a more concrete combinatorial object on which the symmetric group acts, with associated permutation representation isomorphic to S^{λ}. An exciting step in this direction appears in Kazhdan and Lusztig (1979). They construct graphs on which S_n acts to give S^{λ}. For $n \leq 6$, these graphs are available in useful form. Kazhdan and Lusztig construct these representations as part of a unified study of Coxeter groups. The details involve an excursion into very high-powered homology. Garsia and McLarnan (1988) gives as close to an "in English" discussion as is currently available, showing the connections between Kazdahn and Lusztig's representations and Young's natural representation as developed in Chapter 3 of James-Kerber.

(7) *The Robinson-Schensted-Knuth (RSK) Correspondence.* There is a fascinating connection between the representation theory of S_n and a host of problems of interest to probabilists, statisticians, and combinatorialists centered about the R-S-K correspondence. The connected problems include sweeping generalizations of the ballot problem: if one puts λ_1-ones, λ_2-twos, ..., $\lambda_k - k$'s into an urn and draws without replacement, where $\lambda_1 \geq \lambda_2 \ldots \geq \lambda_k$ is a partition of n, then the chance that # ones $\geq$ # two $\geq \ldots \geq$ # k's at each stage of the drawing equals $f(\lambda)/n!$ where $f(\lambda) =$ dim(S^{λ}) discussed in (2) above. This links into formulas for the coverage of Kolmogorov-Smirnov tests, the distribution of the longest increasing subsequence in a permutation, and much else.

The connection centers around a 1-1 onto map $\pi \to (P, Q)$ between S_n and pairs of standard Young tableaux of the same shape. Since there are $f(\lambda)$ of these tableaux of shape λ, we have an explicit interpretation of the formula $n! = \Sigma_\lambda f(\lambda)^2$.

One route to accessing this material starts with Section 5.1.4 in Knuth (1975). Then try Stanley (1971), then some of the papers in Kung (1982). Narayana (1979) gives pointers to some statistical applications. Kerov and Virshik (1985) give applications to statistical analysis of other aspects of random permutations. White (1983) discusses the connection between the R-S-K correspondence and the Littlewood-Richardson rule.