REPRESENTATION THEORY OF THE SYMMETRIC GROUP (FOLLOWING [Ful97])

MICHAEL WALTER

1. Diagrams and Tableaux

Diagrams and Tableaux. — A *(Young) diagram* λ is a partition of a natural number $n \leqslant 0$, which we often represent as a weakly descending tuple $(\lambda_1, \dots, \lambda_n)$ of positive integers summing to n, or, equivalently, as a diagram of n boxes with λ_i boxes in the i-th row. The length $|\lambda|$ of a diagram λ is defined to be the number n partition by λ, i.e., the total number of boxes. We also write $\lambda \vdash n$. We denote the number of rows by $l(\lambda)$.

A *filling* of a Young diagram is an assignment of numbers to each box, it is called a *numbering* if all entries have to be distinct.

A *(semistandard) (Young) tableau* is a filling which is (i) weakly increasing across each row and (ii) strictly increasing down each column. We denote by $\mathrm{Tab}(\lambda, m)$ the set of of tableaux of shape λ which are filled with $[m]$. Clearly, this set is nonempty if and only if $l(\lambda) \leqslant m$. Let us also denote by $\mathrm{Tab}(m) = \bigcup_\lambda \mathrm{Tab}(\lambda, m)$ the set of tableaux of arbitrary shape which are filled by $[m]$. The *content* $c(T)$ of a tableau $T \in \mathrm{Tab}(\lambda, m)$ is the m-tuple $(c_1, \dots, c_m)$ where c_k is the number of times k occurs in T.

A *standard (Young) tableau* is both a tableau with n boxes as well as a numbering by $[n] = \{1, \dots, n\}$. In particular, it is strictly increasing in both directions. We denote by $\mathrm{Tab}_{\mathrm{std}}(\lambda)$ the set of standard tableaux of shape λ.

Conjugate and Transpose. — The *conjugate* $\overline{\lambda}$ is defined by flipping a diagram λ over its main diagonal. Any filling F of λ determines a filling of the conjugate, called the *transpose* and denoted by F^T. The transpose of a standard tableau is again a standard tableau (but the analogue statement for general tableaux is false).

Order Relations. — There are two important order relations on the set of Young diagrams: (i) *Lexicographic order* $\leqslant$, where we consider diagrams as tuples of positive numbers, and (ii) *domination order* $\trianglelefteq$, which is defined as follows: λ is dominated by λ' if

$$\lambda_1 + \ldots + \lambda_k \leqslant \lambda'_1 + \ldots + \lambda'_k \quad (\forall k).$$

Lexicographic order is total, while domination order is not. Clearly, $\lambda \trianglelefteq \lambda'$ implies $\lambda \leqslant \lambda'$, while the converse does not hold in general.

Date: April 17, 2011.

Row Insertion. — Given a tableau T and a positive integer x, we construct a new tableau $T \leftarrow x$ by the following procedure (*row insertion*, or *row bumping*): If x is at least as large as all entries in the first row, place x in a new box at the end of that row; otherwise replace the first entry larger than x by x and repeat the above process on the second row.

Lemma. Row insertion is a well-defined operation on the set of tableaux, constructing a tableau T' with $|T| + 1$ boxes from a tableau with T boxes.

Moreover, it is *invertible* in the following sense: If we are given the resulting tableau T' together with the location of the box B that has been added, then by running the algorithm backwards[1] we recover the original tableau T and the element x that has been added. Similarly, if we apply inverse row insertion on some box B in a tableau T' and row-insert x back into the resulting tableau then we recover the tableau T' we started with.

Let us denote by $T' \rightarrow B$ the tuple (T, x) which we get by running *inverse row insertion* on box B in a tableau T'. Then the statement of the preceding Lemma can be summarized by the following formulas:

$$(T \leftarrow x) \rightarrow B = (T, x)$$
$$\leftarrow (T' \rightarrow B) = T'$$

We also record the following lemma, which is basic for the analysis of row insertions.

Lemma (Row Bumping Lemma, [Ful97, §1.1]**).** Let T be a tableau and x, x' positive integers. Denote by B and B' the new boxes in $(T \leftarrow x) \leftarrow x'$ arising from first inserting x and then x'[2]. Then:

(i) If $x \leqslant x'$, then B is strictly left of and weakly below of B'.

(ii) If $x > x'$, then B is weakly left of and strictly below of B'.

Product of Tableaux. — We can thus define the *product* $T \cdot U$ of two tableaux by row inserting each entry of U into X, from left to right and bottom to top. Clearly, the empty tableau is a unit with respect to the product. We will see later that this operation is associative, turning $\mathrm{Tab}(m)$ into a monoid.

As an example of how to handle the combinatorics of this product, we prove a combinatorial version of Pieri's famous formulas:

Example (Pieri's Formulas — Combinatorial Version). Let T be a tableau of arbitrary shape λ, and U a tableau of shape (n) (resp. (1^n)), that is, given by weakly ascending numbers $x_1 \leqslant \ldots \leqslant x_n$ (resp. strictly ascending numbers $x_n < \ldots < x_1$). Then the shape of the product tableau

$$T \cdot U = ((T \leftarrow x_1) \leftarrow \ldots) \leftarrow x_n$$

[1] Let y be the element in box B (note that B is always the right-most box of a certain row) and remove B from the diagram. Then find right-most entry x in the row above which is strictly smaller than y, replace it by y, and repeat the procedure with x. Stop after having replaced an element x in the first row. Then we have recovered the original tableau T, and x is the element that has been row-inserted.

[2] These are of course *not* necessarily the boxes which x and x' have been placed into.

is obtained by adding n boxes to λ, with no two boxes in the same column (resp. row).

Conversely, given the tableau T, any tableau T' of the latter shape can be *uniquely* factorized in the form $T' = T \cdot U$ where U is a tableau of shape (n) (resp. (1^n)).

Proof (Column version). The first assertion is immediate from the Row Bumping Lemma.

For the converse, denote by $B_1, \ldots, B_n$ the boxes in T' which are not present in T, order from left to right. Perform inverse row insertion on T' in the reverse order $B_n, \ldots, B_1$. Then we get positive integers $x_n, \ldots, x_1$ such that

$$T' = ((T \leftarrow x_1) \leftarrow \ldots) \leftarrow x_n.$$

Since this of course adjoins the boxes B_i in the order $B_1, \ldots, B_n$, i.e., strictly from left to right, we can conclude from the Row Bumping Lemma that $x_1 \leqslant \ldots \leqslant x_n$.

For uniqueness, observe that if $x_1 \leqslant \ldots \leqslant x_n$ is any ascending sequence of positive integers such that

$$T' = ((T \leftarrow x_1) \leftarrow \ldots) \leftarrow x_n$$

then by the Row Bumping Lemma these row insertions always adjoin the boxes B_1, ..., B_n in that order. Consequently, the x_i are uniquely defined because they can be recovered by inverse row insertion in the order $B_n, \ldots, B_1$. ☐

2. Words

Word monoid. — In this section we will give an algebraic description of the set of tableaux filled with $[m]$ as a quotient of the free monoid $F(m)$ of words in the alphabet $[m]$, the product is by juxtaposition.

To do so, let us first assign a word to every tableau T, its *(row) word* $w(T)$, by writing the entries of T from left to right and bottom to top.[3] This defines a map

$$w\colon \operatorname{Tab}(m) \hookrightarrow F(m).$$

Conversely, if w is a word *which comes from a tableau* T then, by breaking the word whenever one number is strictly greater than the next, we recover the tableau T. For example,

$$w_{\text{row}}\left(\begin{array}{|c|c|c|c|}\hline 1 & 2 & 4 & 5 \\ \hline 6 & 8 & 9 \\ \cline{1-3}\end{array}\right) = 6891245.$$

Knuth Equivalence. — We will now study the effect of row insertion on a tableau. The algorithm for constructing $T \leftarrow x$ can be formalized as follows: Inserting x into a row $\underline{r}$ amounts to (i) decomposing $\underline{r} = \underline{u}x'\underline{v}$ into words $\underline{u}$, $\underline{v}$ and a letter x such that all letters of $\underline{u}$ are less or equal to x and $x' > x$, (ii) replacing x' by x, and (iii) bumping x' into the next row (i.e., *in front of* the row $\underline{u}x\underline{v}$). That is,

$$\underline{u}x'\underline{v}x \rightsquigarrow x'\underline{u}x\underline{v} \quad \text{if } \underline{u} \leqslant x < x' \leqslant \underline{v}.$$

[3]One similarly defines the *column word* by writing the entries column by column, from bottom to top and left to right.

Let us decompose this transformation into more elementary ones:

$$\begin{aligned}\underline{u}x'\underline{v}x &= u_1 \dots u_k x' v_1 \dots v_l x \\ &\rightsquigarrow u_1 \dots u_k x' x v_1 \dots v_l \\ &\rightsquigarrow x u_1 \dots u_k x v_1 \dots v_l = x' \underline{u} x \underline{v},\end{aligned}$$

where we have (i) first moved x in front of all but the first entry larger than x and (ii) we have moved x' (which is characterized by sitting between two smaller elements) to the very left.

That is, by starting out with the string $w(T)x$ and applying either of the following transformations, dubbed *elementary Knuth transformation*,

$$\begin{aligned} abx &\mapsto axb \quad \text{if } x < a \leqslant b \\ axb &\mapsto xab \quad \text{if } a \leqslant b < x. \end{aligned}$$

we arrive at the word representing the tableau $T \leftarrow x$. If we define *Knuth equivalence* as the equivalence relation $\equiv$ generated by these steps then the above discussion shows the following:

Lemma. The induced map $w\colon \operatorname{Tab}(m) \to F(m)/\equiv$ preserves products, i.e.,

$$w(T)w(U) = w(T \cdot U).$$

One can show that this equivalence relation not only makes the induced map preserve products, but also that the induced map becomes bijective.

Theorem ([Ful97, p. 22]). Every word is Knuth equivalent to a single tableau. Consequently, the map from above, assigning to a tableau its row word, induces an isomorphism $\operatorname{Tab}(m) \cong F(m)/\equiv$.

In particular, the set of tableau $\operatorname{Tab}(m)$ filled with $[m]$ with the product operation defined above forms an associative monoid, called the *tableau monoid* (or *plactic monoid*).

Observe that it is a *consequence* of the above that the tableau associated to a word $w = x_1 \cdots x_n$ is given by

$$\boxed{x_1} \cdot \dots \cdot \boxed{x_n} = (((\boxed{x_1} \leftarrow x_2) \leftarrow \dots) \leftarrow x_n).$$

Robinson-Schensted Correspondence. — Although every word in a Knuth equivalence class determines the same tableau, we have seen that the row insertion algorithm is invertible (that is, we should recover the word we started from) if only we store the the order in which the boxes have been added when constructing the tableau from the word. We can formalize this idea as follows: Given a word $w = x_1 \cdots x_n = w(T)$, define the associated *recording tableau* by putting k in the box which has been adjoined in the k-th step, that is, when row-inserting x_k. Observe that recording tableaux are standard tableaux. Thus:

Lemma (Robinson-Schensted Correspondence). The map sending a word w to the pair (T, R) consisting of its associated tableau and the recording tableau defines a bijection

$$[m]^n \to \coprod_{\lambda \vdash n} \operatorname{Tab}(\lambda, m) \times \operatorname{Tab}_{\mathrm{std}}(\lambda), \quad w \mapsto (T, R)$$

between words of length n in the alphabet $[m]$ and pairs (P, R) of tableaux of equal shape $\lambda \vdash n$, where P is a tableau with entries in $[m]$ and where R is a standard tableau.

We record the special case of words with $n = m$ and no repetitions (that is, of permutations in $[n]$) for later reuse:

Corollary (Robinson Correspondence). The above map restricts to a bijection $S^n \cong \coprod_{\lambda \vdash n} \mathrm{Tab}_{\mathrm{std}}(\lambda)^2$.

Tableau Ring. — Given any monoid M, we can define its associated *monoid ring* $\mathbb{Z}[M]$ as the free $\mathbb{Z}$-module with basis M and product defined such that the inclusion $M \hookrightarrow \mathbb{Z}[M]$ is a homomorphism of monoids.

The *tableau ring* $\mathbb{Z}[\mathrm{Tab}(m)]$ then is the monoid ring associated to the tableau monoid $\mathrm{Tab}(m)$. It is an associative ring with unit represented by the empty tableaux, and non-commutative if $m > 2$.

For every shape λ we define $S_\lambda \in \mathbb{Z}[\mathrm{Tab}(m)]$ as the formal sum of all tableaux of shape λ. Observe that this element is nonzero if and only if $l(\lambda) \leqslant m$ (otherwise there are no such tableaux!).

Combinatorial assertions about tableaux can often be stated more succinctly using the tableau ring:

Example (Pieri's Formulas — Algebraic Version). We have the following relations in every $\mathbb{Z}[\mathrm{Tab}(m)]$:

$$S_\lambda \cdot S_{(n)} = \sum_\mu S_\mu$$
$$S_\lambda \cdot S_{(1^n)} = \sum_{\mu'} S_{\mu'}$$

where the sum is over all μ (resp. μ') that are obtained from λ by adding n boxes, with no two boxes in the same column (resp. row).

3. Schur Polynomials

Schur Polynomials. — We can assign to every tableau in $T \in \mathrm{Tab}(\lambda, m)$ a monomial of degree $|\lambda|$,

$$X^T := X^{c(T)} := \prod_{k=1}^{m} X_k^{\text{number of times } k \text{ occurs in } T} \in \mathbb{Z}[X_1, \ldots, X_m],$$

and by the universal property this defined a $\mathbb{Z}$-linear map

$$(-)^T \colon \mathbb{Z}[\mathrm{Tab}(m)] \to \mathbb{Z}[X_1, \ldots, X_m],$$

which in fact is a ring homomorphism, as $c(T \cdot U) = c(T) + c(U)$.

The image of the element S_λ under this map is called the *Schur polynomial* $s_\lambda = s_\lambda(X_1, \ldots, X_m)$.

For example, the Schur polynomial for $\lambda = (n)$ is just the n-th *completely symmetric polynomial* $h_n(X_1, \ldots, X_m)$, i.e., the sum of all monomials of degree n in m variables.

Dually, the Schur polynomial for $\overline{\lambda} = (1^n)$ is the n-th *elementary symmetric polynomial* $e_n(X_1, \ldots, X_m)$, i.e., the sum of all monomials of the form $X_{i_1} \cdots X_{i_n}$ where the (i_k) are strictly increasing numbers in $[m]$.

Polynomial Identities. — As $(-)^T$ is a ring homomorphism, every relation in the monoid of tableaux implies an identity between the corresponding polynomials. For example, the algebraic version of Pieri's formulas immediately implies the following:

Example (Pieri's Formulas — Polynomial Version). We have the following identities in the polynomial ring $\mathbb{Z}[X_1, \ldots, X_m]$:

$$s_\lambda(X_1, \ldots, X_m) h_n(X_1, \ldots, X_m) = \sum_\mu s_\mu(X_1, \ldots, X_m)$$
$$s_\lambda(X_1, \ldots, X_m) e_n(X_1, \ldots, X_m) = \sum_{\mu'} s_{\mu'}(X_1, \ldots, X_m)$$

where the sum is over all μ (resp. μ') that are obtained from λ by adding n boxes, with no two boxes in the same column (resp. row).

Kostka Numbers. — Let λ be a diagram and $c = (c_1, \ldots, c_m)$ a tuple of nonnegative integers. We define the *Kostka number* $K_{\lambda,c}$ as the number of tableaux of shape λ and content c.

Equivalently, $K_{\lambda,c}$ is the number of sequences of tableaux

$$\varnothing = \lambda^{(0)} \subseteq \lambda^{(1)} \subseteq \ldots \subseteq \lambda^{(m)} = \lambda$$

such that in each step $\lambda^{(i-1)} \rightsquigarrow \lambda^{(i)}$ one adds precisely c_i boxes, with no two in the same column.[4] In other words, we have the identity

$$K_{\lambda,c} = \sum_{\lambda^{(1)}} \cdots \sum_{\lambda^{(m)} = \lambda} 1$$

where the $\lambda^{(i)}$ are constrained as above.

Lemma. For all diagrams λ and weight tuples $c = (c_1, \ldots, c_m)$ we have

$$S_{(c_1)} \cdots S_{(c_m)} = \sum_\lambda K_{\lambda,c} \, S_\lambda \in \mathbb{Z}[\mathrm{Tab}(m)],$$

where λ runs over all tableaux (filled with $[m]$).

Proof. In view of the identity displayed above, the assertion follows from Pieri's first formula by induction over m. □

[4]The bijection assigns to a tableaux T of shape λ and content c the filtration defined by letting $\lambda^{(i)}$ be the sub-tableau of λ with entries no larger than i.

Symmetric Polynomials. — In particular, the preceding Lemma implies identities between the Schur polynomials and certain symmetric polynomials:

$$h_{c_1}(X_1, \ldots, X_m) \cdots h_{c_m}(X_1, \ldots, X_m) = \sum_{\lambda} K_{\lambda,c}\, s_\lambda(X_1, \ldots, X_m).$$

We can use these to show that the Schur polynomials are in fact symmetric:

Lemma. The Schur polynomials $s_\lambda(X_1, \ldots, X_m)$ are symmetric polynomials in $\mathbb{Z}[X_1, \ldots, X_m]$.

Proof. By varying the content vector c, we get a system of linear equations in the polynomial ring $\mathbb{Z}[X_1, \ldots, X_m]$. We would like to argue that the matrix $(K_{\lambda,c})$ is invertible over $\mathbb{Z}$ when choosing the content vectors carefully.

In fact, we will only consider content vectors μ which are diagrams with the same number of boxes as λ. Let us order these partitions lexicographically. We claim that $(K_{\lambda,\mu})$ is a square matrix such that (i) all non-zero entries satisfy $\mu \leqslant \lambda$ and (ii) all diagonal elements are equal to one. The assertion then follows from Gauss' algorithm over $\mathbb{Z}$. But these claims are true simply because the entries of a tableau have to increase strictly down each column. □

4. Representation Theory of S_n

Row Group and Column Group. — In this section we consider numberings by $[n]$ of Young diagrams of shape λ with n boxes.[5] Let us denote the set of all such diagrams by $\mathrm{Num}(\lambda)$. We let the symmetric group S_n act by replacing the entry of each box in a numbering by its permutation. For example,

$$(123) \cdot \begin{array}{|c|c|c|}\hline 1 & 2 & 4 \\ \hline 3 & 5 \\ \cline{1-2}\end{array} = \begin{array}{|c|c|c|}\hline 2 & 3 & 4 \\ \hline 1 & 5 \\ \cline{1-2}\end{array}.$$

The *row group* $R(T)$ of such a numbering then is the subgroup of permutations which permute the entries of each row among themselves. Clearly, $R(T)$ is isomorphic to a group of the form $S_{\lambda_1} \times \cdots \times S_{\lambda_{l(\lambda)}}$ (a *Young subgroup*), where λ is the shape of T. Likewise, the *column group* $C(T)$ is defined as the subgroup of permutations preserving the columns of the numbering T.

Tabloids. — A *tabloid* of shape $\lambda \vdash n$ is an equivalence class of numberings of λ by $[n]$, two such numberings being equivalent if the corresponding rows contain the same entries.[6] We write $\{T\}$ for the tabloid corresponding to a numbering T. For example,

$$\{\begin{array}{|c|c|}\hline 1 & 2 \\ \hline 3 \\ \cline{1-1}\end{array}\} = \{\begin{array}{|c|c|}\hline 2 & 1 \\ \hline 3 \\ \cline{1-1}\end{array}\}.$$

As permuting boxes and permuting entries commutes with each other, the action of the symmetric group on numberings descends to an action on the set of tabloids via $g \cdot \{T\} = \{g \cdot T\}$.

[5] That is, standard tableaux, but without the ordering condition.

[6] In other words, tabloids are cosets with respect to *another* canonical action of S_n on $\mathrm{Num}(\lambda)$ by permuting boxes, restricted to the Young subgroup corresponding to the diagram λ.

Consequently, the complex vector space M^λ with basis the tabloids of shape λ is a linear representation of S_n.

Young Symmetrizers and Specht Modules. — With a numbering T we can associate certain elements in the group algebra of the symmetric group,

$$r(T) := \sum_{\pi \in R(T)} \pi, \qquad c(T) := \sum_{\pi \in C(T)} \operatorname{sign}(\pi)\pi, \qquad , y(T) := c(T)r(T),$$

called *Young symmetrizers*. It is easy to check that, up to a non-zero scalar, $r(T)$ and $c(T)$ are idempotent.

We will associate to every numbering T its "column-wise anti-symmetrization"

$$v(T) := c(T) \cdot \{T\} = \sum_{\pi \in C(T)} \operatorname{sign}(\pi)\{\pi \cdot T\} \in M^\lambda.$$

Lemma. For all numberings $T \in \operatorname{Num}(\lambda)$ and all permutations $\pi \in S_n$, we have

$$\pi \cdot v(T) = v(\pi \cdot T)$$

Proof. This is immediate from the fact that $C(\pi \cdot T) = \pi C(T) \pi^{-1}$. □

Consequently, the linear span S^λ of all such vectors $v(T)$ is a representation of the symmetric group S_n, called the *Specht module*. Observe that $S^\lambda = \mathbb{C}[S_n] \cdot v(T)$ for any numbering T (an obvious consequence of the Lemma).

The following Lemma is the main combinatorial ingredient of the representation theory of the symmetric group:

Lemma. Let $\lambda, \lambda' \vdash n$ such that λ does not strictly dominate λ'. Then for any two numberings $T \in \operatorname{Num}(\lambda)$, $T' \in \operatorname{Num}(\lambda')$, exactly one of the following occurs:

(i) There exists a pair of distinct integers that occur in the same column of T and in the same row of T', and $c(T) \cdot \{T'\} = 0$.

(ii) The shapes λ and λ' are the same, there exist $\pi \in C(T)$, $\pi' \in R(T')$ such that $\pi \cdot T = \pi' \cdot T'$, and $c(T) \cdot \{T'\} = \pm v(T)$.

Proof. Let us first suppose that there exists such a pair of integers. Denote by τ the transposition permuting them, and consider

$$c(T) \cdot \{T'\} = (c(T)\tau) \cdot (\tau \cdot \{T'\}).$$

On the one hand, $\tau \cdot \{T'\} = \{T'\}$, because the pair of integers occurs in the same row of T'. On the other hand, $C(T)\tau = C(T)$, since it occurs in the same column of T, and thus $c(T)\tau = -c(T)$. Consequently, $c(T) \cdot \{T'\}$ vanishes.

Conversely, suppose that there exist no two such integers. Then the entries in the first (in fact, any) row of T' occur in different columns of T, and we can find a permutation in $C(T)$ moving them to the first row. Iterating this procedure with the other rows of T', we find a permutation $\pi \in C(T)$ such that every row of $\pi \cdot T$ contains (at least) the entries of the corresponding row of T'. Applying π does not change the shape of T, so that we have $\lambda \trianglerighteq \lambda'$. But, by assumption, λ does not strictly dominate λ', so both are in fact equal. Thus $\pi \cdot T$ and T' have the same shape and contain the same entries in each row. This means that we can find $\pi' \in R(T)$ such that $\pi \cdot T = \pi' \cdot T'$. In particular,

$$c(T) \cdot \{T'\} = c(T) \cdot \{\pi' T'\} = c(T) \cdot \{\pi T\} = c(T) \cdot \pi \cdot \{T\} = \operatorname{sign}(\pi) v(T). \qquad \square$$

Proposition. For each diagram $\lambda \vdash n$, the Specht module S^λ is an irreducible complex representation of the symmetric group S_n, and all these representations are non-isomorphic.

Conversely, every irreducible complex representation of S_n is of that form.

Proof. Pairwise non-isomorphic: Let $T \in \operatorname{Num}(\lambda)$. Then:

$$v(T) = c(T) \cdot \{T\} \propto c(T)^2 \cdot \{T\} = c(T) \cdot v(T) \subseteq c(T) \cdot S^\lambda \subseteq c(T) \cdot M^\lambda = \mathbb{C} v(T),$$

where the last identity is by the preceding Lemma, hence $c(T) \cdot S^\lambda = \mathbb{C} v(T)$.

On the other hand, if $\lambda \neq \lambda'$, say $\lambda < \lambda'$, then, as λ does not strictly dominate λ', this Lemma also shows that

$$c(T) \cdot S^{\lambda'} \subseteq c(T) \cdot M^{\lambda'} = 0.$$

That is, we can distinguish the representations S^λ, $S^{\lambda'}$ using elements of the group algebra. They are thus non-isomorphic.

Irreducibility: Now we show that S^λ is irreducible, or, equivalently, decomposable (we work over characteristic zero). Indeed, suppose that S^λ decomposes into two subrepresentations, say $S^\lambda = V \oplus W$. Then

$$\mathbb{C} v(T) = c(T) \cdot S^\lambda = (c(T) \cdot V) \oplus (c(T) \cdot W),$$

thus precisely one of the right-hand side summands is non-zero. Consequently, one of the subrepresentations contains $v(T)$, say V, and

$$S^\lambda = \mathbb{C}[S_n] \cdot v(T) \subseteq V \subseteq S^\lambda.$$

Hence $S^\lambda = V$, and it follows that S^λ is indecomposable.

Completeness: By elementary representation theory, the number of isomorphism classes of representations is equal to the number of conjugacy classes, which in turn is equal to the number of partitions of n. The last assertion follows thus. $\square$

We remark that the representations M^λ and S^λ could have been defined first over the rational numbers and then tensored with $\mathbb{C}$, so their characters are $\mathbb{Q}$-valued. This in particular implies that the irreducible representations S^λ are *self-dual*.

While the $v(T)$ span the irreducible representation S^λ, the following Lemma shows that restricting T to standard tableaux we get a basis of S^λ:

Lemma. The vectors $(v(T))$, where T varies over the standard tableaux of shape λ, form a basis of S^λ.

Proof. Linear independence: Let us totally order the numberings in $\operatorname{Num}(\lambda)$ by setting $T > T'$ if the largest numbers which is placed in two different boxes occurs earlier in the column word of T than in the column word of T'.

Observe that, if T is a *standard* tableau, then $\pi \cdot T > T$ for all $1 \neq \pi \in R(T)$ (since in each row of T, the entries are ordered strictly ascending from left to right). It follows that every tabloid is represented by at most one standard tableau, which is then the minimal element of the equivalence class.

Similarly, we have $\pi \cdot T < T$ for all $1 \neq \pi \in C(T)$. Consequently, among the standard tableau whose associated tabloid has non-zero coefficient in

$$v(T) = \sum_{\pi \in C(T)} \operatorname{sign}(\pi)\{\pi \cdot T\}$$

there exists a unique maximal one, namely $\{T\}$.[7] Thus a non-trivial linear combination $\sum_{T \in \mathrm{Tab}_{\mathrm{std}}(\lambda)} c_T v(T)$ can never be zero, since if T is maximal with $c_T \neq 0$ then the coefficient of $\{T\}$ is non-zero.

Completeness: A basic result in the representation theory of finite groups says that the square of the dimensions of all irreps sums to the group order. Consequently,

$$n! = \sum_{\lambda \vdash n} (\dim S^\lambda)^2 \geqslant \sum_{\lambda \vdash n} (\#\,\mathrm{Tab}_{\mathrm{std}}(\lambda))^2 = n!,$$

where the last identity is due to the Robinson Correspondence. We conclude that $\dim S^\lambda = \#\,\mathrm{Tab}_{\mathrm{std}}(\lambda)$ for all shapes λ in parallel. □

Examples. — We will now look at some examples of Specht modules (in particular, we will see all irreducible representations of S_3):

Example ($\lambda = (1^n)$). The representation S^λ is the trivial representation of S_n (the column groups are all trivial).

Example ($\lambda = (n)$). The representation S^λ is the signum representation of S_n. Indeed, the representation is one-dimensional, spanned by the vector v_T corresponding to the single standard tableau, and as $C(T) = S_n$, we have

$$\pi \cdot v_T = \operatorname{sign}(\pi)\, v_T \quad (\forall \pi).$$

Example ($\lambda = \begin{array}{|c|c|}\hline &\\\hline \\\cline{1-1}\end{array}$). The standard tableaux of shape λ are $\begin{array}{|c|c|}\hline 1&2\\\hline 3\\\cline{1-1}\end{array}$ and $\begin{array}{|c|c|}\hline 1&3\\\hline 2\\\cline{1-1}\end{array}$, thus S^λ is two-dimensional with basis

$$X := v\Big(\begin{array}{|c|c|}\hline 1&2\\\hline 3\\\cline{1-1}\end{array}\Big) = \Big\{\begin{array}{|c|c|}\hline 1&2\\\hline 3\\\cline{1-1}\end{array}\Big\} - \Big\{\begin{array}{|c|c|}\hline 2&3\\\hline 1\\\cline{1-1}\end{array}\Big\},$$
$$Y := v\Big(\begin{array}{|c|c|}\hline 1&3\\\hline 2\\\cline{1-1}\end{array}\Big) = \Big\{\begin{array}{|c|c|}\hline 1&3\\\hline 2\\\cline{1-1}\end{array}\Big\} - \Big\{\begin{array}{|c|c|}\hline 2&3\\\hline 1\\\cline{1-1}\end{array}\Big\},$$

and the action of S_3 is given by

$$(1\ 2) \cdot X = \Big\{\begin{array}{|c|c|}\hline 1&2\\\hline 3\\\cline{1-1}\end{array}\Big\} - \Big\{\begin{array}{|c|c|}\hline 1&3\\\hline 2\\\cline{1-1}\end{array}\Big\} = X - Y,$$
$$(1\ 2) \cdot Y = \Big\{\begin{array}{|c|c|}\hline 2&3\\\hline 1\\\cline{1-1}\end{array}\Big\} - \Big\{\begin{array}{|c|c|}\hline 1&3\\\hline 2\\\cline{1-1}\end{array}\Big\} = -Y,$$
$$(2\ 3) \cdot Y = \Big\{\begin{array}{|c|c|}\hline 1&2\\\hline 3\\\cline{1-1}\end{array}\Big\} - \Big\{\begin{array}{|c|c|}\hline 2&3\\\hline 1\\\cline{1-1}\end{array}\Big\} = X.$$

It is in fact straightforward to show that S^λ is isomorphic to the representation on

$$\{(x, y, z) : x + y + z = 0\}$$

[7] In fact, $\{T\}$ should be the only one, since if π

by permuting coordinates.[8]

A Glimpse at Schur-Weyl Duality. Every diagram λ induces a functor, the *Schur functor* $\mathbb{S}^\lambda$ sending a complex vector space V to the $\mathrm{GL}(V)$-representation

$$\mathbb{S}^\lambda(V) := V^{\otimes n} \otimes_{\mathbb{C}[S_n]} S^\lambda.$$

Proposition. We have

$$V^{\otimes n} \cong \sum_{\lambda \vdash n} \mathbb{S}^\lambda(V) \otimes S^\lambda$$

as GL_n–S_n–bimodules.

Proof. By the Peter-Weyl analogue for finite groups and self-duality of the Specht modules, we have $\mathbb{C}[S_n] \cong \sum_{\lambda \vdash n} S^\lambda \otimes S^\lambda$ as S_n–bimodules. Thus:

$$V^{\otimes n} \cong V^{\otimes n} \otimes_{\mathbb{C}[S_n]} \mathbb{C}[S_n] \cong \sum_{\lambda \vdash n} V^{\otimes n} \otimes_{\mathbb{C}[S_n]} S^\lambda \otimes S^\lambda \cong \sum_{\lambda \vdash n} \mathbb{S}^\lambda(V) \otimes S^\lambda. \qquad \square$$

In fact, the representations $\mathbb{S}^\lambda(V)$ for diagrams λ with at most $\dim V$ rows are all of the irreducible *polynomial* representations of $\mathrm{GL}(V)$ (and zero otherwise). See [Ful97, §8] for more details.

References

[Ful97] William Fulton, *Young Tableaux*, Student Texts, no. 35, London Mathematical Society, 1997.

[8] The basis vectors X and Y correspond to $(1, 0, -1)$ and $(1, -1, 0)$, respectively.