Modular Representations of Symmetric Groups

Charlotte Chan

Sunday, August 7, 2011

1 Introduction

Over the last few weeks, we have seen lots of things in representation theory—we began by looking at an onslaught of examples to see explicitly what representations look like, we discussed character theory and looked at character tables, we talked about representations from a module-theoretic perspective, and then we saw how we can model the representation theory of the symmetric group via the combinatorics of Young tableaux. But most of our discussion has been about the representation theory of finite groups over the complex numbers. With Maschke's theorem in mind, it seems natural to ask what happens when the hypotheses of of this theorem fail. That is, what happens to the representations of a finite group G if we wish to work over a field of characteristic dividing the order of the group?

(As a side comment: It seems like a shame that when we have $\operatorname{char} K \mid |G|$, KG is not a semisimple algebra. But the perspective we should actually take is that it is fantastically *miraculous* that KG is semisimple when $\operatorname{char} K \nmid |G|$. For instance, if we were to think about some analogous statement for groups, it's ridiculously false! I mean, it's not even true that abelian groups are direct sums of simple groups! So we shouldn't be depressed about the times when Maschke's theorem fails. We should just be ecstatic that we have Maschke's theorem!)

In the last couple of lectures, we have looked at the specific case of when $G = S_n$, the symmetric group. Jeremy talked about how (miraculously!) we can parameterize representations of S_n over $\mathbb{C}$ (or an algebraically closed field of characteristic not dividing $n! = |S_n|$) by the partitions of n, which form a transversal for the conjugacy classes of S_n. We saw how the Specht modules S^λ, for partitions λ of n, form a transversal for the isomorphism classes of irreducible representations of S_n. In this lecture, we will try to obtain an analogous transversal for the case when the characteristic of the ground field of our representation space divides $n!$, the order of S_n.

2 Modular Representation Theory: A First Examination

We begin with a brief revision of some previous notions we've had in our discussion of representation theory and which results carry over.

Recall:

Definition 1. A *representation* of a group G over a field K is a group homomorphism

$$\rho : G \to \mathrm{GL}_n(K).$$

Equivalently, we may think of a representation as a finite-dimensional vector space V equipped with a linear G-action. Yet another way of thinking of a representation is as a module over the group algebra KG.

Definition 2. The *character* χ_V afforded by a representation (V, ρ) is the corresponding trace map:

$$\chi : G \to \mathbb{C}, \qquad \chi_V(g) := \mathrm{tr}(\rho(g)).$$

Note that since the trace of a linear transformation is well defined (i.e. it is independent of the choice of basis) and since ρ is a group homomorphism, χ is constant on conjugacy classes of G. As a warning, note that in general, χ is *not* a homomorphism.

We now introduce some notation that will be fixed for the remainder of this lecture, unless otherwise noted.

G is a finite group.
F is an algebraically closed field where $\mathrm{char}(F) \nmid |G|$.
K is an algebraically closed field where $\mathrm{char}(K) \mid |G|$.
$p := \mathrm{char}(K)$.
S_n is the symmetric group, the set of permutations on n objects.
We pick $n \in \mathbb{N} \cup \{0\}$ such that $\mathrm{char}(K) = p \leq n < \mathrm{char}(F)$.

Recall, from both Ian's and Jason's lectures, that every FG-module can be written as a direct sum of simple FG-modules. (This is just Maschke's theorem, or perhaps a slight generalization of Maschke's theorem.) This fails, unfortunately, when we consider KG-modules. An easy yet important example that we showed in the first lecture of this seminar is that if G is a p-group, the only simple KG-module is the trivial module. (We actually only proved this for when $K = \mathbb{Z}/p\mathbb{Z}$, but the idea is the same for when you prove this result for any characteristic p field.) From this example, we can see that the elements that "matter" in determining the simple KG-modules are the ones whose order are prime to p. These elements are called the *p-regular* elements of G. (To elaborate on this... We can, in some sense, "pull out" the p-group part of G (think about Sylow p-groups) and then study the action of the p-regular elements of G on vector spaces.) This gives us some intuition as to why you should believe the following fact that I will not prove:

Theorem 1. *The number of isomorphism classes of simple KG-modules is exactly the number of p-regular conjugacy classes of G.*

This is a generalization to the theorem we had from before that the number of isomorphism classes of simple FG-modules is exactly the number of conjugacy classes of G. Recall that this equality allowed us to form a square matrix encoding all the information of the characters afforded by irreducible representations and their values on the conjugacy classes of G. Because of Theorem 1, we would expect to be able to construct some analogous table for the "characters" afforded by simple KG-modules and their values on p-regular conjugacy classes. But first we must try to make some sense of what we mean by the word "character."

We can first try to naively define a character in the same way as we did for $\mathbb{C}G$-modules, i.e. take the trace map in the respective field. Now consider a p-dimensional vector space over K and let G act on K trivially. Then the described representation is simply

$$\rho : G \to \mathrm{GL}_p(K), \qquad g \mapsto \mathrm{Id}_p,$$

where Id_p is the $p \times p$ identity matrix. Taking the trace of this in K, we get a function on G that is identically zero. In fact, we get this same result if the dimension of the vector space at hand is any multiple of p. Hence if we were to define a character as the trace function in the ground field of the representation, then not only would the character not uniquely define a representation (up to isomorphism) as we had for $\mathbb{C}G$-modules, but an infinite set of KG-modules could share the same character!

We may salvage this description of a character in the following way. Write $|G| = p^k m$ where $k \in \mathbb{N}$ and $p \nmid m$. We can construct a ring homomorphism $\varphi : \mathbb{Z}[\zeta] \to K$ where ζ is a primitive mth root of unity by sending ζ to a primitive mth root of unity in K. Now, for a representation $\rho : G \to \mathrm{GL}_n(K)$, we may restrict this map to the set of p-regular elements of G, which we will denote by G_{reg}. Let $g \in G_{\mathrm{reg}}$. Then $\rho(g) \in \mathrm{GL}_n(K)$ is a matrix of order $l := |\langle g \rangle|$ in $\mathrm{GL}_n(K)$ and since $\mathrm{char}(K) = p \nmid l$ and $K = \overline{K}$, then $\rho(g)$ is diagonalizable. (This follows from Maschke's theorem and Schur's lemma.) Hence all the eigenvalues of $\rho(g)$ are lth roots of unity and hence also mth roots of unity. We may lift these eigenvalues to $\mathbb{Z}[\zeta] \subseteq \mathbb{C}$ via the ring homomorphism φ and then sum the eigenvalues in $\mathbb{C}$. The resulting sum is the *Brauer character* of ρ, and it is a function (in fact, a class function) $G_{\mathrm{reg}} \to \mathbb{C}$. It turns out that the Brauer characters of two irreducible representations are equal if and only if the representations are isomorphic, and hence Brauer characters give us the modular analogue of ordinary characters. Another fact: Any Brauer character can be written as a $\mathbb{Z}$-linear combination of irreducible Brauer characters. But unlike the case of ordinary characters, here, we allow for negative coefficients.

We will leave this discussion for now, and only return to it briefly at the end of the lecture when we compute the Brauer character table for S_3 in characteristic 2 and 3. This will come after we discuss the behavior of KS_n-modules and examine what happens to the Specht modules, which Jeremy proved were irreducible in characteristic 0.

3 Modular Representations of S_n

We have seen that the Specht modules S^λ for partitions λ of n form a transversal for the irreducible representations of S_n over $\mathbb{C}$ (or any algebraically closed field of characteristic 0). We used the following proposition to show that these are simple:

Proposition 1 (Submodule Lemma)**.** *Let $U \subseteq M^\lambda$ be a submodule of the permutation module M^λ. Then either $U \supseteq S^\lambda$ or $U \subseteq (S^\lambda)^\perp$, where the orthogonal complement is taken with respect to the inner product on M^λ defined by $\langle \{s\}, \{t\} \rangle = \delta_{\{s\},\{t\}}$ and then extended linearly to all of M^λ.*

(Note: The notation $\{t\}$ means the λ-tabloid associated to the λ-tableau t, i.e. the equivalence class of λ-tableaux with respect to the relation that entries in individual rows may be permuted.)

In characteristic 0, $S^\lambda \cap (S^\lambda)^\perp = 0$, and hence the above proposition allows us to conclude that the Specht modules are simple. (Since $\langle \cdot, \cdot \langle$ is a G-invariant inner product, then if S^λ is a representation of S_n, then so must $(S^\lambda)^\perp$. Hence $S^\lambda \cap (S^\lambda)^\perp$ is a representation of S_n.) However, in characteristic p, we may have a nontrivial intersection. Here is an example of when this happens.

Example 1. Take, for instance, $n = 5$, and consider the partition $\lambda = (3, 2)$. Then a basis for S^λ is composed of the standard λ-polytabloids (i.e. $e_t := \sum_{\pi \in C_t} \operatorname{sgn}(\pi)\{\pi t\}$, where t is a standard λ-tableau). So we have the following:

$$
\begin{aligned}
e_1 &= \left\{\begin{array}{|c|c|c|}\hline 1&2&3\\\hline 4&5\\\cline{1-2}\end{array}\right\} - \left\{\begin{array}{|c|c|c|}\hline 4&2&3\\\hline 1&5\\\cline{1-2}\end{array}\right\} - \left\{\begin{array}{|c|c|c|}\hline 1&5&3\\\hline 4&2\\\cline{1-2}\end{array}\right\} + \left\{\begin{array}{|c|c|c|}\hline 4&5&3\\\hline 1&2\\\cline{1-2}\end{array}\right\},\\
e_2 &= \left\{\begin{array}{|c|c|c|}\hline 1&2&4\\\hline 3&5\\\cline{1-2}\end{array}\right\} - \left\{\begin{array}{|c|c|c|}\hline 3&2&4\\\hline 1&5\\\cline{1-2}\end{array}\right\} - \left\{\begin{array}{|c|c|c|}\hline 1&5&4\\\hline 3&2\\\cline{1-2}\end{array}\right\} + \left\{\begin{array}{|c|c|c|}\hline 3&5&4\\\hline 1&2\\\cline{1-2}\end{array}\right\},\\
e_3 &= \left\{\begin{array}{|c|c|c|}\hline 1&2&5\\\hline 3&4\\\cline{1-2}\end{array}\right\} - \left\{\begin{array}{|c|c|c|}\hline 3&2&5\\\hline 1&4\\\cline{1-2}\end{array}\right\} - \left\{\begin{array}{|c|c|c|}\hline 1&4&5\\\hline 3&2\\\cline{1-2}\end{array}\right\} + \left\{\begin{array}{|c|c|c|}\hline 3&4&5\\\hline 1&2\\\cline{1-2}\end{array}\right\},\\
e_4 &= \left\{\begin{array}{|c|c|c|}\hline 1&3&4\\\hline 2&5\\\cline{1-2}\end{array}\right\} - \left\{\begin{array}{|c|c|c|}\hline 2&3&4\\\hline 1&5\\\cline{1-2}\end{array}\right\} - \left\{\begin{array}{|c|c|c|}\hline 1&5&4\\\hline 2&3\\\cline{1-2}\end{array}\right\} + \left\{\begin{array}{|c|c|c|}\hline 2&5&4\\\hline 1&3\\\cline{1-2}\end{array}\right\},\\
e_5 &= \left\{\begin{array}{|c|c|c|}\hline 1&3&5\\\hline 2&4\\\cline{1-2}\end{array}\right\} - \left\{\begin{array}{|c|c|c|}\hline 2&3&5\\\hline 1&4\\\cline{1-2}\end{array}\right\} - \left\{\begin{array}{|c|c|c|}\hline 1&4&5\\\hline 2&3\\\cline{1-2}\end{array}\right\} + \left\{\begin{array}{|c|c|c|}\hline 2&4&5\\\hline 1&3\\\cline{1-2}\end{array}\right\}.
\end{aligned}
$$

Taking all possible inner products, we have

$$
\begin{array}{lllll}
\langle e_1, e_1\rangle = 4, & \langle e_2, e_2\rangle = 4, & \langle e_3, e_3\rangle = 4, & \langle e_4, e_4\rangle = 4, & \langle e_5, e_5\rangle = 4,\\
\langle e_1, e_2\rangle = 2, & \langle e_2, e_3\rangle = 2, & \langle e_3, e_4\rangle = 1, & \langle e_4, e_5\rangle = 2, &\\
\langle e_1, e_3\rangle = 1, & \langle e_2, e_4\rangle = 2, & \langle e_3, e_5\rangle = 2, & &\\
\langle e_1, e_4\rangle = 1, & \langle e_2, e_5\rangle = 1, & & &\\
\langle e_1, e_5\rangle = 1. & & & &
\end{array}
$$

Now suppose that we have some element $\alpha = a_1e_1+a_2e_2+\cdots a_5e_5 \in S^\lambda$ such that $\langle \alpha, e_i\rangle = 0$ for all $i = 1, \ldots, 5$. Then $\alpha \in (S^\lambda)^\perp$. Taking all these dot products, we get the following:

$$\begin{aligned}
\langle \alpha, e_1\rangle &= 4a_1 + 2a_2 + a_3 + a_4 - a_5,\\
\langle \alpha, e_2\rangle &= 2a_1 + 4a_2 + 2a_3 + 2a_4 + a_5,\\
\langle \alpha, e_3\rangle &= a_1 + 2a_2 + 4a_3 + a_4 + 2a_5,\\
\langle \alpha, e_4\rangle &= a_1 + 2a_2 + a_3 + 4a_4 + 2a_5,\\
\langle \alpha, e_5\rangle &= -a_1 + a_2 + 2a_3 + 2a_4 + 4a_5.
\end{aligned}$$

So we would like to solve the following equation:

$$\begin{pmatrix} 4 & 2 & 1 & 1 & -1\\ 2 & 4 & 2 & 2 & 1\\ 1 & 2 & 4 & 1 & 2\\ 1 & 2 & 1 & 4 & 2\\ -1 & 1 & 2 & 2 & 4 \end{pmatrix} \begin{pmatrix} a_1\\ a_2\\ a_3\\ a_4\\ a_5 \end{pmatrix} = 0.$$

Notice that the 5×5 matrix is just the Gram matrix $(\langle a_i, a_j\rangle)_{i,j=1,\ldots 5}$. Now, in $\mathbb{C}$, this matrix has full rank, which means that the only solution is the trivial solution, which would force $\alpha = 0$. So this shows us explicitly that $S^\lambda \cap (S^\lambda)^\perp = 0$ when S^λ is a $\mathbb{C}S_n$-module. If change the field from $\mathbb{C}$ to something of characteristic 3, for instance, then our Gram matrix becomes:

$$\begin{pmatrix} 1 & 2 & 1 & 1 & 2\\ 2 & 1 & 2 & 2 & 1\\ 1 & 2 & 1 & 1 & 2\\ 1 & 2 & 1 & 1 & 2\\ 2 & 1 & 2 & 2 & 1 \end{pmatrix}.$$

This has rank 1, which means that the $S^\lambda \cap (S^\lambda)^\perp$, when viewed as a vector space, has dimension 4, so there are lots of things in this intersection! To give an explicit example, consider $\alpha = e_1+e_3+e_4$. Hence we have shown explicitly, in the case of $n = 5$ and $\lambda = (3, 2)$ that the Specht module S^λ is not irreducible.

Let us now return to Proposition 1, the submodule lemma. Let $U \subsetneq S^\lambda$ be a maximal submodule. By the submodule lemma, we know that $U \subseteq (S^\lambda)^\perp$, so $U \subseteq S^\lambda \cap (S^\lambda)^\perp$. If $S^\lambda \cap (S^\lambda)^\perp \neq S^\lambda$, then this means that $U = S^\lambda \cap (S^\lambda)^\perp$, and we have proven the following proposition.

Proposition 2. *If $S^\lambda \cap (S^\lambda)^\perp \neq S^\lambda$, then $S^\lambda \cap (S^\lambda)^\perp$ is the unique maximal submodule of S^λ.*

Now define

$$D^\lambda := S^\lambda/(S^\lambda \cap (S^\lambda)^\perp).$$

By the previous proposition, D^λ is simple whenever $S^\lambda \cap (S^\lambda)^\perp \neq S^\lambda$, and it is 0 otherwise.

Now the question is: For which partitions λ of n is D^λ simple? In order to answer this question, we first define some notation.

Definition 3. Let $\mathcal{P}_n$ denote the set of partitions of n. We adopt an alternative notation for a partition $\lambda \in \mathcal{P}_n$. We write

$$\lambda = (n^{m_n}, \ldots, 1^{m_1}),$$

where m_i denotes the number of parts of size i in λ. We say that λ is p-regular if $m_i < p$ for all $i = 1, \ldots, n$. We denote the set of all p-regular partitions of n by $\mathcal{P}_{n,p}$.

Despite what the terminology may suggest, the p-regular partitions of n *do not* correspond to the cycle shape of the p-regular conjugacy classes of S_n. Perhaps to justify the naming, we have the following.

Proposition 3. *The number of p-regular conjugacy classes is $|\mathcal{P}_{n,p}|$.*

Proof. Consider the following function:

$$\frac{\prod_{i\geq 1}(1-x^{pi})}{\prod_{i\geq 1}(1-x^i)}. \qquad (*)$$

Cancelling out all multiplicands $(1-x^{pi})$ for $i \geq 1$ in the bottom product, we get

$$(*) = \frac{1}{\prod_{\substack{p\nmid i \\ i\geq 1}}(1-x^i)} = \prod_{\substack{p\nmid i \\ i\geq 1}} \frac{1}{(1-x^i)} = \prod_{\substack{p\nmid i \\ i\geq 1}} (1 + x^i + x^{2i} + \cdots).$$

The coefficient of the x^n term is the number of partitions of n such that no part is divisible by p. Now, a conjugacy class is p-regular if and only if the cycle decomposition of each element in that conjugacy class has no cycle of length divisible by p. Phrased in terms of partitions, we have that $\lambda \in \mathcal{P}_n$ corresponds to a p-regular partition if and only if no part of λ is divisible by p. Hence the coefficient of the x^n term is exactly the number of p-regular conjugacy classes of S_n.

We can manipulate $(*)$ in a different way. If we first divide $1 - x^i$ into $1 - x^{pi}$, we get

$$(*) = \prod_{i\geq 1} (1 + x^i + x^{2i} + \cdots + x^{(p-1)i}),$$

and then the coefficient of the x^n term is the number of partitions with at most $p-1$ parts of size i for each $i = 1, \ldots, n$. In other words, the coefficient of the x^n term is the number of p-regular partitions.

Therefore, equating coefficients on both sides, we have that the number of p-regular conjugacy classes of S_n is equal to the number of p-regular partitions of n. □

We would now like to determine when $S^\lambda \cap (S^\lambda)^\perp = S^\lambda$. The main ingredient in the description of these partitions λ is the following lemma.

Lemma 1. *Let $\lambda = (n^{m_n}, \ldots, 1^{m_1}) \in \mathcal{P}_n$.*

a. If t and s are λ-tableau, then

$$\prod_{j=1}^{n} (m_j)! \,\big|\, \langle e_t, e_s \rangle.$$

b. If t is a λ-tableau and $\tilde{t}$ is the λ-tableau obtained by reversing the entries in each row, then

$$\langle e_t, e_{\tilde{t}} \rangle = \prod_{j=1}^{n} (m_j!)^j.$$

Proof. We begin by defining an equivalence relation $\sim$ on the tableaux of shape λ. If t_1 and t_2 are λ-tableaux, we say that $t_1 \sim t_2$ if and only if t_1 is obtained by permuting rows of t_2. Note that there are $\prod_{j=1}^{n} m_j!$ elements in each equivalence class.

Here is the key observation to this proof:

Key Observation. If $\{t_1\}$ appears in e_t, then so must $\{t_2\}$ for all t_2 with $t_2 \sim t_1$. Now, the coefficients of $\{t_1\}$ and $\{t_2\}$ are either the same or are of opposite sign, and whichever case occurs is dependent only on $\{t_1\}$ and $\{t_2\}$. In particular, this is independent of the choice of λ-tableau t.

Now assume that $\{t_1\}$ appears in e_t and e_s with the same coefficient (either $+1$ or -1). Then by the Key Observation, all $\prod_{j=1}^{n} m_j!$ elements of the equivalence class of $\{t_1\}$ must occur with the same coefficient in e_t and e_s. By the linearity of the inner product $\langle \cdot, \cdot \rangle$ on M^λ, the contribution of the equivalence class of $\{t_1\}$ to $\langle e_t, e_s \rangle$ is $\prod_{j=1}^{n} m_j!$. Similarly, if $\{t_1\}$ occurs in e_t and e_s with opposite sign, then the contribution is $-\prod_{j=1}^{n} m_j!$. Hence $\langle e_t, e_s \rangle$ is the sum of some number of $\pm \prod_{j=1}^{n} m_j!$, so

$$\prod_{j=1}^{n} m_j! \,\big|\, \langle e_t, e_s \rangle,$$

and this completes the proof of a.

For b, let $C \leq C_t$ be the subgroup of C_t consisting of all permutations $\pi \in C_t$ such that for all $i \in \{1, \ldots, n\}$, i and $\pi(i)$ are in rows of equal length in the λ-tableau t. Then

$$C \cong \prod_{j=1}^{n} (S_{m_j})^j,$$

since for each $j = 1, \ldots, n$, we may permute the m_j elements in each of the j columns in the rows of length j. So we have $|C| = \prod_{j=1}^{n} (m_j!)^j$. Now let $\tilde{t}$ be the λ-tableau obtained

by reversing the entries of each row of the λ-tableau t. If $\{t_1\}$ occurs in e_t and $e_{\tilde{t}}$, then necessarily $\{t_1\} = \{\pi t\}$ for some $\pi \in C$ and the coefficient of $\{t_1\}$ in e_t and $e_{\tilde{t}}$ must be the same. Thus

$$\langle e_t, e_{\tilde{t}} \rangle = \prod_{j=1}^{n} (m_j!)^j,$$

as desired. □

From this, we easily get the following proposition:

Proposition 4. *Let $\lambda = (n^{m_n}, \ldots, 1^{m_1}) \in \mathcal{P}_n$. Then $D^\lambda \neq 0$ if and only if $\lambda \in \mathcal{P}_{n,p}$.*

Proof. Suppose $\lambda \in \mathcal{P}_{n,p}$. Then for any λ-tableau t, let $\tilde{t}$ be the λ-tableau obtained from reversing the entries in each row of t. Then

$$\langle e_t, e_{\tilde{t}} \rangle = \prod_{j=1}^{n} (m_j!)^j \neq 0,$$

where the first equality holds from part b of the lemma and the second statement holds since λ is p-regular and hence $p \nmid \prod_{j=1}^{n} (m_j!)^j$. This means that $S^\lambda \not\subseteq (S^\lambda)^\perp$ so $S^\lambda \cap (S^\lambda)^\perp \neq S^\lambda$ and hence $D^\lambda \neq 0$.

Conversely, suppose that $\lambda \notin \mathcal{P}_{n,p}$. Then for all λ-tableaux t and s, we have

$$\langle e_t, e_s \rangle \equiv 0 \pmod{\prod_{j=1}^{n} (m_j)!},$$

but since λ is not p-regular, this means that $m_i \geq p$ for some $i = 1, \ldots, n$, so $p \,\big|\, \prod_{j=1}^{n} (m_j)!$. Therefore $\langle e_t, e_s \rangle = 0$ in K and $S^\lambda \subseteq (S^\lambda)^\perp$. It follows that $D^\lambda = 0$, and this completes the proof. □

We in fact have something more. We first state a fact without proof.

Fact. Consider $\lambda \in \mathcal{P}_{n,p}, \mu \in \mathcal{P}_n$. Let U be a KS_n-module of M^μ and let $\varphi \in \mathrm{Hom}_{KS_n}(S^\lambda, M^\mu/U)$. If $\varphi \neq 0$, then $\mu \trianglelefteq \lambda$, where $\trianglelefteq$ is the dominance order.

Proposition 5. *Let $\lambda, \mu \in \mathcal{P}_{n,p}$. Then $D^\lambda \cong D^\mu$ if and only if $\lambda = \mu$.*

Proof. It is clear that if $\lambda = \mu$, then $D^\lambda \cong D^\mu$. For the reverse direction, let $\psi : D^\lambda \to D^\mu$ be a KS_n-isomorphism. Define φ as the composition

$$\varphi : S^\lambda \twoheadrightarrow D^\lambda \xrightarrow{\psi} D^\mu := S^\mu/(S^\mu \cap (S^\mu)^\perp) \hookrightarrow M^\mu/(S^\mu \cap (S^\mu)^\perp).$$

Note that $\varphi \neq 0$. By the fact, we have $\mu \trianglelefteq \lambda$. Similarly, replacing ψ by ψ^{-1}, we have $\lambda \trianglelefteq \mu$. Hence $\lambda = \mu$ and this proves the proposition. □

From the results we have proven above, we can easily prove the following theorem.

Theorem 2. *The KS_n-modules D^λ for $\lambda \in \mathcal{P}_{n,p}$ form a transversal for the isomorphism classes of simple KS_n-modules.*

Proof. From Proposition 5, we know that the D^λ's are non-isomorphic, and from the remark following Proposition 2 and Proposition 4, we know that there are $|\,\mathcal{P}_{n,p}\,|$ such D^λ's. But by Proposition 3 and Theorem 1, we know that $|\,\mathcal{P}_{n,p}\,|$ is exactly the number of p-regular conjugacy classes of S_n and this is exactly the number of isomorphism classes of simple KS_n-modules. Therefore $\{D^\lambda : \lambda \in \mathcal{P}_{n,p}\}$ is indeed a complete set of representatives and we are done. □

Even though we have, in some sense, fully described the simple KS_n-modules, we have also, in some sense, done nothing at all. That is, we've constructed these objects in an abstract way, but do we really have an understanding, an intuition, of what is happening? Perhaps some people do, but I certainly don't understand this material enough to say I really know what is going on. As a start, though, we could ask what the dimension of D^λ is.

We in fact already answered this question in the example illustrating the non-simplicity of S^λ, $\lambda = (3,2)$ in characteristic 3. Recall that in that example, we computed the dimension of the vector space $S^\lambda \cap (S^\lambda)^\perp$ by looking at the dimension of the kernel of the linear transformation given by the Gram matrix $(\langle e_i, e_j\rangle)_{i,j=1,\dots,k}$, where $e_1, \dots, e_k$ is a basis for S^λ. This actually comes from a general fact:

Proposition 6. *Let F be any field and let V be a finite-dimensional F-vector space equipped with a non-degenerate form $\langle\cdot,\cdot\rangle : V \times V \to F$. Let W be a subspace of V with F-basis $\{e_1, \dots, e_k\}$. Then*

$$\dim_F(W/W \cap W^\perp) = \mathrm{rk}((\langle e_i, e_j\rangle)_{i,j=1,\dots,k}),$$

where $(\langle e_i, e_j\rangle)_{i,j=1,\dots,k}$ is the Gram matrix of W with respect to the basis $\{e_1, \dots, e_k\}$.

Proof. Let $W^* := \mathrm{Hom}_F(W, F)$ and let $\{e_1^*, \dots, e_m^*\}$ be the basis of W^* dual to $\{e_1, \dots, e_m\}$. We define an F-linear map $\varphi : W \to W^*$, $w \mapsto \varphi_w$, where $\varphi_w(u) = \langle w, u\rangle$. Then

$$\varphi e_i = \sum_{j=1}^{k} \varphi e_i(e_j) e_j^* = \sum_{j=1}^{k} \langle e_i, e_j\rangle e_j^*.$$

So the matrix corresponding to φ with respect to $\{e_1, \dots, e_k\}$ and $\{e_1^*, \dots, e_k^*\}$ is exactly the Gram matrix $(\langle e_i, e_j\rangle)_{i,j=1,\dots,k}$. Since $\ker(\varphi) = W \cap W^\perp$, then

$$\dim(W/W \cap W^\perp) = \dim(\Im(\varphi)) = \mathrm{rk}((\langle e_i, e_j\rangle)_{i,j=1,\dots,k}).$$

□

(Note that this is just a "high-brow" way of saying exactly why we did what we did in the example we worked through earlier in this lecture.)

We conclude this lecture with an example.

4 Ordinary and Brauer Characters of S_3

Recall that we have three irreducible representations of S_3 in $\mathbb{C}$:

- The trivial representation, T. This is a one-dimensional vector space on which S_3 acts trivially.
- The sign representation, S. This is a one-dimensional vector space where $\pi \in S_3$ acts by multiplication of $\text{sgn}(\pi)$.
- The standard representation, V. This is the two-dimensional vector subspace $\{(x, y, z) \in \mathbb{C}^3 : x + y + z = 0\} \subseteq \mathbb{C}^3$ where S_3 acts by permuting the coordinates. Decomposing $\mathbb{C}^3$ as a vector space, we have

$$\mathbb{C}^3 = \{(x, y, z) \in \mathbb{C}^3 : x + y + z = 0\} \oplus \{(x, y, z) \in \mathbb{C}^3 : x = y = z\}.$$

  Now, since both summands are G-invariant, then this means that this decomposition is also a decomposition of the permutation representation $\mathbb{C}^3$ as direct sum of representations. Now, the character of the permutation representation is just the number of fixed points, hence for any $g \in S_3$, we have $\text{Fix}(g) = \chi_V(g) + \chi_T(g) = \chi_V(g) + 1$.

From the above information, we can fill out the (ordinary) character of S_3.

	1	(12)	(123)
T	1	1	1
S	1	-1	1
V	2	0	-1

We can find out which of these representations is associated to which Specht modules. It is clear that the trivial representation T is just $S^{(3)}$. Similarly, it is clear that the sign representation is just $S^{(1,1,1)}$. (We can determine these two by looking at the dimension of the Specht modules and then looking at the action of S_3 on the standard λ-polytabloids.) This means that $S^{(2,1)}$ is the standard representation. (It may be a helpful exercise to go through exactly how the description of the Specht module $S^{(2,1)}$ is exactly the same as our description of the standard representation. You can start by writing down a basis for $S^{(2,1)}$.)

Now let's look at the Brauer characters. To get the Brauer characters of S_3 in, say, characteristic 3, we just need to restrict to the 3-regular conjugacy classes and see what happens. In this case, we get

	1	(12)
T	1	1
S	1	-1
V	2	0

Here, we see that the restriction of χ_V to G_{reg}, denoted $\chi_V|_{G_{\text{reg}}}$, is just the sum of the restrictions of χ_T and χ_S. Since T and S are both one-dimensional as vector spaces, they must be irreducible, and hence the Bruaer character table of S_3 in characteristic 3 is just

	1	(12)
T	1	1
S	1	-1

(We can see that $\chi_T|_{G_{\text{reg}}} + \chi_S|_{G_{\text{reg}}} = \chi_V|_{G_{\text{reg}}}$.

From our results in the preceding section, one of these must correspond to $D^{(3)}$ and the other must correspond to $D^{(2,1)}$. (Note that $D^{(1,1,1)} = 0$ in characteristic 3 since $(1,1,1)$ is not p-regular. We can see this in a more explicit way as well. The Specht module $S^{(1,1,1)}$ is one-dimensional has basis

$$e := \left\{\begin{array}{|c|}\hline 1\\\hline 2\\\hline 3\\\hline\end{array}\right\} - \left\{\begin{array}{|c|}\hline 1\\\hline 3\\\hline 2\\\hline\end{array}\right\} - \left\{\begin{array}{|c|}\hline 2\\\hline 1\\\hline 3\\\hline\end{array}\right\} + \left\{\begin{array}{|c|}\hline 2\\\hline 3\\\hline 1\\\hline\end{array}\right\} - \left\{\begin{array}{|c|}\hline 3\\\hline 2\\\hline 1\\\hline\end{array}\right\} + \left\{\begin{array}{|c|}\hline 3\\\hline 1\\\hline 2\\\hline\end{array}\right\},$$

and $\langle e, e\rangle = 6$, which is 0 in a field of characteristic 3.) Now, $D^{(3)}$ must be one-dimensional since it is nonzero and is the quotient of a one-dimensional space, so this means that $S^{(3)} = D^{(3)}$, so the restriction of χ_T is the Brauer character of $D^{(3)}$. This means that the restriction of the sign representation S takes on the Brauer character of $D^{(2,1)}$.

Now let's look at what happens in characteristic 2. Omitting the conjugacy classes of S_3 that are not 2-regular, we get

	1	(123)
T	1	1
S	1	1
V	2	-1

As we can see, the restriction of χ_T and χ_S coincide, so what is left of the character table is just

	1	(123)
T	1	1
V	2	-1

We would like to show that the Brauer character $\chi_V|_{G_{\text{reg}}}$ is irreducible. Now, we know that the irreducible representations are $D^{(3)}$ and $D^{(2,1)}$. We also know, by the previous discussion about the characteristic 3 case, that $D^{(3)}$ is the trivial representation. Now let us use Proposition 6 to compute the dimension of $D^{(2,1)}$. We start with a basis for the Specht module $S^{(2,1)}$:

$$e_1 := \left\{\begin{array}{|c|c|}\hline 1 & 2\\\hline 3\\\cline{1-1}\end{array}\right\} - \left\{\begin{array}{|c|c|}\hline 3 & 2\\\hline 1\\\cline{1-1}\end{array}\right\}$$

$$e_2 := \left\{\begin{array}{|c|c|}\hline 1 & 3\\\hline 2\\\cline{1-1}\end{array}\right\} - \left\{\begin{array}{|c|c|}\hline 2 & 3\\\hline 1\\\cline{1-1}\end{array}\right\}.$$

We compute the inner products:

$$\langle e_1, e_1\rangle = 2, \qquad \langle e_1, e_2\rangle = \langle e_2, e_1\rangle = 1, \qquad \langle e_2, e_2\rangle = 2,$$

and putting this information into our Gram matrix, we get:

$$\begin{pmatrix} \langle e_1, e_1\rangle & \langle e_1, e_2\rangle \\ \langle e_2, e_1\rangle & \langle e_2, e_2\rangle \end{pmatrix} = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} \equiv \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \pmod 2.$$

The rightmost matrix has rank 2, and hence $D^{(2,1)}$ has dimension 2 by Proposition 6. Since $S^{(2,1)}$ also has dimension 2, then this means that $D^{(2,1)} \cong S^{(2,1)}$, which means that the restriction of χ_V is indeed irreducible, and this proves that the above Brauer character table is indeed the table we want.

5 The Moral of the Story

The previous section illustrated that even for a small group like S_3, it takes quite a bit of work to try to compute the Brauer characters, even though we described earlier what the irreducibles are. On the other hand, we found that the p-regular partitions give us a way of parametrizing all of the KS_n-modules! So there are several things at work here: On the one hand, we have seen that when we lose the semisimplicity of the group algebra, we also lose many of the results we discussed in "classical" representation theory. In this lecture, we only saw the most basic of modular representation theory, and already it is quite difficult to really understand what is going on! On the other hand, we have also seen that even when the group algebra is not semisimple, we can get results that resemble what we had in the "classical" theory. I think it's really beautiful how everything is interconnected and even when you branch away from the representation theory of finite groups either towards modular representation theory or towards the representation theory of infinite groups, we still have results that resemble the results we learn in a first course on representation theory.

To conclude, here are some readings related to what we have done in this seminar:

- Linear Representations of Finite Groups, J.P. Serre (This is a really beautiful book. It starts out quite gently and ends with modular representation theory. He's also quite explicit in how he treats representations, so it feels relatively hands-on.)
- Representation Theory: A First Course, W. Fulton and J. Harris (I don't know this book, but Jeremy really likes it, and I trust his opinion. In terms of content, there's a massive amount of stuff in this book. It might be the most comprehensive book on this list.)
- Representations and Characters of Groups, G.D. James and M.W. Liebeck (This is the most gentle book you'll find on representation theory. It was the first book I

read on the subject, and it was great. It feeds things to you slowly and explicitly and there are lots of examples of character tables. It's great if you want some math to read before bed and it's also great if you're trying to learn things on your own without much algebra background.)

- Local Representation Theory, J.L. Alperin (This is an incredibly well-written book, but it is really quite dense and sparse. It also takes a "high-brow" module-theoretic approach, so it might feel too abstract if you've never seen the material before.)
- The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions, B.E. Sagan (There's a lot in this book about the combinatorial approach to studying representation theory of the symmetric group.)

It's been fun, I love representation theory as much as ever, and I hope that my fellow counselors have found this to be an interesting seminar!