The Gelfand-Tsetlin Basis

(Too Many Direct Sums, and Also a Graph)

David Grabovsky

June 13, 2018

Abstract

The symmetric groups S_n, consisting of all permutations on a set of n elements, naturally contain each other like Matryoshka dolls. (S_{n-1} simply fixes the nth element permuted by S_n.) We will explore the hope that the representation theory of S_n is also inductive. Along the way, we will develop a tool called the *branching graph* to help us organize the way that the irreducible representations of S_n decompose into those of lower S_k. This decomposition results in the canonical *Gelfand-Tsetlin basis* for each irreducible representation of S_n. Finally, we will construct the *Gelfand-Tsetlin algebra* and prove that its spectrum uniquely identifies elements from the Gelfand-Tsetlin basis.

Contents

N.B. Throughout, all groups are finite and all vector spaces are finite-dimensional over $\mathbb{C}$.

1 Review of Representation Theory

The basic idea of representation theory is to study groups by “representing” them concretely as linear operators on vector spaces, thereby using methods from linear algebra to simplify otherwise intractible questions.

1.1 Basic Definitions

Definition 1.1. A *representation* of a group G is a homomorphism $\rho\colon G \to \mathrm{GL}(V)$. Each $g \in G$ is therefore represented by an invertible linear map $\rho(g)\colon V \to V$.

Another way to say the same thing is that G acts linearly on V, with $g \cdot v = \rho(g)(v)$. This gives V the structure of a G-module:

Definition 1.2. A *G-module* is an abelian group V, together with an action of G on V that respects its linear structure: explicitly, for $g, h \in G$ and $v, v' \in V$,
(a) $1_G \cdot v = v$;
(b) $g \cdot (h \cdot v) = (gh) \cdot v$;
(c) $g \cdot (v + v') = g \cdot v + g \cdot v'$.

The G-module structure of V packages together all of the information about the representation ρ it carries, so we will henceforth write V^ρ and speak of it indiscriminately as a “G-module” or a “representation.” The G-module perspective is useful because it allows us to speak of “G-linear” maps that respect extra action structure:

Definition 1.3. A *G-morphism* is a linear G-equivariant map of vector spaces $\phi\colon V^\rho \to V^\sigma$; that is, for every scalar $c \in \mathbb{C}$ and vectors $v, v' \in V^\rho$,
(a) $\phi(cv) = c\phi(v)$;
(b) $\phi(v + v') = \phi(v) + \phi(v')$;
(c) $\phi(g \cdot v) = g \cdot \phi(v) \iff \phi\big(\rho(g)(v)\big) = \sigma(g)\phi(v)$.

We will denote the set of all such maps by $\mathrm{Hom}_G(V^\rho, V^\sigma)$. Note that since we can add and scale G-morphisms, $\mathrm{Hom}_G(V^\rho, V^\sigma)$ acquires the structure of a vector space.

1.2 Irreducible Representations

Definition 1.4. Let V be a representation of G. A *subrepresentation* of G is a subspace $W \subset V$ such that for all $w \in W$ and $g \in G$, $g \cdot w \in W$. (“What happens in W stays in W.”)

Definition 1.5. An *irreducible* representation (sometimes “irrep”) V of G is one with no nontrivial proper subrepresentations, i.e. no G-invariant subspaces. We denote the set of irreducible representations of G by $\widehat{G}$, and write $\rho \in \widehat{G}$ to say that ρ is an irrep of G.

One can show that V decomposes as a direct sum of subrepresentations: if $W \subset V$ is a subrepresentation, then so is $W^\perp$, and moreover $V = W \oplus W^\perp$. (Here the orthogonal complement $W^\perp$ is defined using a G-equivariant inner product on V.) By continuing to split off subrepresentations until we reach the irreducible ones, we obtain the following result:

Theorem 1.6. *Every representation ρ of a finite group G on a finite-dimensional complex vector space V^ρ decomposes as a direct sum of irreducible representations V^σ.*

Some of the irreps in the decomposition of V^ρ may be isomorphic; that is, each irrep V^σ may appear more than once as a factor in V^ρ. Denoting this *multiplicity* by $m_\sigma \in \mathbb{N}$,

$$V^\rho = \bigoplus_{\sigma \in \widehat{G}} (V^\sigma)^{m_\sigma}. \tag{1.1}$$

As it turns out, there's a formula for m_σ. To prove it, we'll need the extremely useful

Lemma 1.7 (Schur)**.** *Let V^σ, V^ρ be irreducible representations of G, and let $\phi\colon V \to W$ be a G-morphism. Then ϕ is either 0 or an isomorphism acting by scalars, $\phi = \lambda I$. Thus,*

$$\mathrm{Hom}_G(V^\sigma, V^\rho) = \begin{cases} \mathbb{C}, & \text{if } V^\sigma \cong V^\rho; \\ 0, & \text{otherwise.} \end{cases} \tag{1.2}$$

Proposition 1.8. *The multiplicity of $\sigma \in \widehat{G}$ in a G-module V^ρ is $m_\sigma = \dim \mathrm{Hom}_G(V^\sigma, V^\rho)$.*

Proof. By Thm. 1.6, V^ρ decomposes as a direct sum over $\widehat{G}$:

$$V^\rho = \bigoplus_{\sigma' \in \widehat{G}} V^{\sigma'} = (V^\sigma)^{m_\sigma} \oplus \Big(\bigoplus_{\sigma' \neq \sigma} (V^{\sigma'})^{m_{\sigma'}} \Big). \tag{1.3}$$

Now Hom commutes with direct sums, so we have

$$\mathrm{Hom}_G(V^\sigma, V^\rho) = \mathrm{Hom}_G(V^\sigma, V^\sigma)^{m_\sigma} \oplus \Big(\bigoplus_{\sigma' \neq \sigma} \mathrm{Hom}_G(V^\sigma, V^{\sigma'}) \Big). \tag{1.4}$$

By Schur's lemma, $\mathrm{Hom}_G(V^\sigma, V^{\sigma'})$ is $\mathbb{C}$ if $\sigma = \sigma'$ and $\{1\}$ otherwise. Therefore we find that $\mathrm{Hom}_G(V^\sigma, V^\rho) = \mathbb{C}^{m_\sigma} \oplus \{1\} = \mathbb{C}^{m_\sigma} \implies m_\sigma = \dim \mathrm{Hom}_G(V^\sigma, V^\rho)$, as desired. ∎

2 The Branching Graph

2.1 An Inductive Chain of Groups

Our eventual goal is to study the representation theory of S_n. We begin by noting that the symmetric groups contain each other, $S_1 < S_2 < S_3 < \cdots$, and that the natural embedding $S_{n-1} \hookrightarrow S_n$ simply fixes the nth element permuted by S_n. The idea that the symmetric groups form an inductive chain motivates the hope that their representation theory might also be inductive, i.e. $\widehat{S}_n$ should depend on $\widehat{S}_{n-1}$. Therefore to set the stage, consider a chain of finite groups $\{1\} = G_0 < G_1 < G_2 < \cdots$. A natural question to ask is what happens to the irreps $\widehat{G}_n$ when we restrict them to G_{n-1} and whether we can glean any new insight by continuing to restrict all the way down to G_0.

Definition 2.1. The *restriction* of a representation V^ρ to a subgroup $H < G$ is defined by $\mathrm{Res}_H^G V^\rho = V^{\rho|_H}$, i.e. the same representation ρ, but restricted to H.

Now let $\lambda \in \widehat{G}_n$, and note that while V^λ is irreducible, its restriction $\mathrm{Res}_{G_{n-1}}^{G_n} V^\lambda$ may not be. It therefore decomposes into a direct sum of irreps $\mu \in \widehat{G}_{n-1}$ with multiplicities $m_\mu = \dim \mathrm{Hom}_G(V^\mu, V^\lambda)$:

$$V^\lambda = \bigoplus_{\mu \in \widehat{G}_{n-1}} (V^\mu)^{m_\mu}. \tag{2.1}$$

We can further decompose each V^μ into irreps from $\widehat{G}_{n-2}$, and so on inductively until we reach $\widehat{G}_0 = \{\bullet\}$, which yields 1-dimensional irreps $V^\bullet \cong \mathbb{C}$. As we are about to see, there is an elegant way of organizing this decomposition.

2.2 The Graph and the GZ Basis

Definition 2.2. The *branching graph* or *Bratteli diagram* is the directed multigraph whose vertices are elements of $\bigsqcup_{k\geq 0} \widehat{G}_k$, with $\widehat{G}_n$ called the nth *level*. Two vertices $\mu \in \widehat{G}_{n-1}$ and $\lambda \in \widehat{G}_n$ are connected by k dirrected edges if $k = \dim \mathrm{Hom}_{G_{n-1}}(V^\mu, V^\lambda)$.

We write $\mu \nearrow \lambda$ if $k \geq 1$; that is, if μ and λ are connected. What this really means is that V^μ is a factor in the decomposition (2.1) of V^λ. If all of the multiplicities k are 0 or 1, then diagram becomes a graph and we say that *the branching is simple.* This turns out to be the case for S_n, but we will not be able to prove this just yet. In the case of simple branching, (2.1) reduces to a direct sum over connected irreps on the $(n-1)$th level:

$$V^\lambda = \bigoplus_{\substack{\mu \in \widehat{G}_{n-1} \\ \mu \nearrow \lambda}} V^\mu \tag{2.2}$$

We further decompose each V^μ into spaces V^ρ for irreps $\rho \in \widehat{G}_{n-2}$ with $\rho \nearrow \mu$, and so on until we finally break V^λ down into one-dimensional irreps $V^\bullet = \mathbb{C} =: V_T$, one for each *increasing path*, i.e. a chain $T = \{\bullet = \lambda_0 \nearrow \lambda_1 \nearrow \cdots \nearrow \lambda_n = \lambda\}$, where each $\lambda_i \in \widehat{G}_i$. Thus:

Theorem 2.3. *In the case of simple branching, we have the canonical decomposition*

$$V^\lambda = \bigoplus_T V_T \tag{2.3}$$

into G_0-modules $V_T \cong \mathbb{C}$, indexed by all possible increasing paths T from $\bullet$ to λ.

Next, recall that like every representation, V^λ has a G_n-equivariant inner product $\langle \cdot, \cdot \rangle$. With respect to this inner product, we may choose a unit vector $v_T \in V_T$ in every factor of V^λ in (2.3). This defines, up to a complex multiple of unit norm, a basis $\{v_T\}$ for V^λ called the *Gelfand-Tsetlin basis* or *GZ basis.*

3 The Gelfand-Tsetlin Algebra

3.1 Group Algebras

We now step back to introduce one more algebraic structure—the eponymous one:

Definition 3.1. An *algebra* is a vector space A equipped with a bilinear product. Specifically, for all $v_1, v_2, w \in A$ and $c_1, c_2 \in \mathbb{C}$, there is a "multiplication" operation satisfying
(a) $w(v_1 + v_2) = wv_1 + wv_2$;
(b) $(v_1 + v_2)w = v_1w + v_2w$;
(c) $(c_1v_1)(c_2v_2) = (c_1c_2)v_1v_2$.

For example, for any vector space V, the space $\mathrm{Hom}(V, V)$ is an algebra. Addition and scalar multiplication of linear maps are defined as usual, and multiplication is given by the commutator, $A \cdot B := [A, B] = AB - BA$, whose bilinearity is easy to verify.

Definition 3.2. The *group algebra* $\mathbb{C}[G]$ is the algebra over $\mathbb{C}$ generated by the group G (the elements $g \in G$ comprise a basis). Addition is given by formal linear combinations of the group elements, while multiplication is defined by extending the given group law (defined on the basis) by linearity to the rest of $\mathbb{C}[G]$, e.g. $(c_1g_1 + c_2g_2)h = c_1(g_1h) + c_2(g_2h)$.

The group algebra $\mathbb{C}[G]$ can be thought of as a generalization of G that gives each group element a "weighting factor" in $\mathbb{C}$. There are two main reasons for introducing group algebras. The first is that every representation V^ρ is not only a G-module, but in fact more naturally a $\mathbb{C}[G]$-module: indeed, V^ρ is acted on not only by G, but also by linear combinations of group elements, which means that the entire group algebra is represented on V^ρ. The second reason is the following formula, which says that $\mathbb{C}[G]$ should be identified with the sum of spaces of *linear operators* on the irreducibles V^ρ.

Proposition 3.3. $\mathbb{C}[G] = \bigoplus_{\rho \in \widehat{G}} \mathrm{End}(V^\rho)$.

Proof. Omitted; see Serre, §6.2 for an elegant proof. ■

3.2 The GZ Algebra

We are almost ready to construct the Gelfand-Tsetlin algebra. We need one more definition:

Definition 3.4. The *center* of an algebra A, denoted $Z(A)$, is the set of $\alpha \in A$ that commute with all of A: $Z(A) := \{\alpha \in A \mid \forall \beta \in A,\ \alpha\beta = \beta\alpha\}$.

Returning now to our inductive situation, consider the chain $\mathbb{C}[G_1] \subset \mathbb{C}[G_2] \subset \cdots$, and let Z_i denote the center $Z(\mathbb{C}[G_i])$ of each group algebra for $1 \le i \le n$.

Definition 3.5. The *Gelfand-Tsetlin algebra* GZ_n is the algebra generated by the centers $Z_1, Z_2, ..., Z_n \subset \mathbb{C}[G_n]$; that is, $\mathrm{GZ}_n = \mathrm{span}\{Z_i\}_{i=1}^n$.

4 A Few Results

Theorem 4.1 (Prop. 1.1)**.** GZ_n *is the algebra of all operators on* V^λ *diagonal in the GZ basis* $\{v_T\}$*. Moreover, it is a maximal commutative subalgebra of* $\mathbb{C}[G_n]$*.*

Proof. Let V^ρ be a $\mathbb{C}[G]$-module. Any element $\alpha \in \mathbb{C}[G]$ determines a map $\alpha\colon V^\rho \to V^\rho$ by $\alpha(v) = \alpha v$. If $\alpha \in Z(\mathbb{C}[G])$, then the map α respects the $\mathbb{C}[G]$-module structure: therefore central elements of $\mathbb{C}[G]$ are $\mathbb{C}[G]$-morphisms of V^ρ. If in addition such a map $P \in Z(\mathbb{C}[G])$ satisfies $P^2 = P$, then it is a projection, and we call such elements *central idempotents.*

Now let $P_{\lambda_i} \in Z_i$ be the central idempotents projecting onto the irrep $\lambda_i \in \widehat{G}_i$. For each increasing path T, denote by $P_T \in \mathrm{GZ}_n$ the product $P_{\lambda_1} P_{\lambda_2} \cdots P_\lambda$. By construction, this operator is a projection onto V_T; running over all increasing paths T, we see that GZ_n contains projections onto all of the V_T. Using these, we can build all operators diagonalizable with eigenbasis $\{v_T\}$. Thus GZ_n contains the set D all operators diagonal in the GZ basis.

It remains to show that D also contains GZ_n. But D is a maximal commutative subalgebra of $\mathbb{C}[G_n]$: if $A \in D$ and $AB = BA$, then $B \in D$ as well. (Any matrix commuting with a diagonal matrix must also be diagonal.) Since GZ_n is commutative (it's generated by central elements) and D is maximal, we must have $\mathrm{GZ} \subset D$. ∎

Corollary 4.2. *Any* $v \in \{v_T\}$ *is uniquely determined, up to a scalar factor, by the eigenvalues of the elements of* GZ_n*.*

Proof. Any operator $A \in \mathrm{GZ}_n$ is diagonal in the basis $\{v_T\}$ by Thm. 4.1; along the diagonal lie its eigenvalues. Acting by A on v therefore scales it by an eigenvalue that selects which v was acted upon. So up to a scalar multiple, the eigenvalues of A identify v uniquely. ∎

We conclude with a criterion for the simple branching that obtains (as we claimed) when $G_n = S_n$. To make life easier, we will state it for *semisimple* ("nice") algebras, which our V^ρ are. Therefore let M be a semisimple finite-dimensional $\mathbb{C}$-algebra, and $N \subset M$ a subalgebra.

Definition 4.3. The *centralizer* $Z_N(M)$ of N in M is the set of elements of M that commute with N. That is, $Z_N(M) := \{\alpha \in M \mid \forall \beta \in N, \alpha\beta = \beta\alpha\}$. Note that $Z_M(M) = Z(M)$.

Theorem 4.4 (Prop. 1.4)**.** *The centralizer* $Z_N(M)$ *is commutative if and only if, for any* $\rho \in \widehat{M}$*, the restriction* $\mathrm{Res}^M_N V^\rho$ *of an irrep of* M *to* N *has simple multiplicities.*

Proof. We prove both implications, the second one by showing the contrapositive.

$(\implies)$ Assume that $Z_N(M)$ is commutative, and let V^μ and V^λ be irreps of M and N, respectively. Consider the M-module $\mathrm{Hom}_N(V^\mu, V^\lambda)$; it is also an irreducible $Z_N(M)$-module. Because $Z_N(M)$ is commutative, $\mathrm{Hom}_N(V^\mu, V^\lambda)$ becomes an irreducible representation of the abelian $Z_N(M)$, and must therefore be one-dimensional (or zero) by Schur's lemma. By Prop. 1.8, the multiplicity of μ in $\mathrm{Res}^M_N V^\lambda$ must be simple.

$(\impliedby)$ Conversely, suppose that $Z_N(M)$ is *not* commutative. Then there exists an irrep of $Z_N(M)$ of dimension more than one, which is to say $\mathrm{Hom}_N(V^\mu, V^\lambda)$ has dimension more than one as well. Hence the multiplicity of μ in $\mathrm{Res}^M_N V^\rho$ is not simple. ∎