

Nilpotent Orbits and Weyl Group Representations, I

O.S.U. Lie Groups Seminar

April 12, 2017

1. INTRODUCTION

Today, G will denote a complex Lie group and \mathfrak{g} the complexification of its Lie algebra. Moreover, to keep matters simple, and in particular amenable to finite computations, we may as well restrict G to be the prototypical class of simple complex Lie groups. Then the relatively simple linear algebraic object \mathfrak{g} determines G up to isogeny (a choice of covering groups). In fact, an even simpler datum determines \mathfrak{g} up to isomorphism; that is its corresponding root system Δ ; which in turn can be reduced to a even simpler combinatorial datum; a set Π of mutually obtuse linearly independent vectors in a real Euclidean space such that

$$\langle \alpha^\vee, \beta \rangle \equiv \frac{2 \langle \alpha, \beta \rangle}{\langle \alpha, \alpha \rangle} \in \begin{cases} \{2\} & \text{if } \alpha = \beta \\ \{0, -1, -2, -3\} & \text{if } \alpha \neq \beta \end{cases}, \quad \alpha, \beta \in \Pi$$

Such *simple systems* are then reduced to simple weighted planar diagrams constructed by the following algorithm:

- each vertex corresponds to a particular *simple root* $\alpha \in \Pi$
- if $\langle \alpha^\vee, \beta \rangle \neq 0$, then an edge is drawn between α and β . These edges are drawn between vertices whenever $\langle \alpha, \beta \rangle \neq 0$ according to the following rules
 - There is an undirected edge of weight 1 connecting α and β whenever $\langle \alpha^\vee, \beta \rangle = \langle \beta^\vee, \alpha \rangle = -1$
 - There is a directed edge of weight 2 from α to β whenever $\langle \alpha^\vee, \beta \rangle = -1$ and $\langle \beta^\vee, \alpha \rangle = -2$
 - There is a directed edge of weight 3 from α to β whenever $\langle \alpha^\vee, \beta \rangle = -1$ and $\langle \beta^\vee, \alpha \rangle = -3$
 (The edge possibilities listed above are actually the only possibilities that occur when $\langle \alpha, \beta \rangle \neq 0$)

If we demand further that for each $\alpha \in \Pi$, there is also a $\beta \in \Pi$ such that $\langle \alpha^\vee, \beta \rangle \in \{-1, -2, -3\}$ (so that we consider only connected Dynkin diagrams), then one has only a set of five exceptional possibilities (G_2, F_4, E_6, E_7 and E_8) and four infinite families (A_n, B_n, C_n, D_n) of simple systems.

Let me stress again the point of this. If (isomorphism classes of) simple complex Lie groups are determined by their simple systems, then so is their representation theory and, in fact, all constructs of their group and representation theory should be pre-determined by their simple systems. In this talk, I am going to show how about two seemingly disparate constructs, nilpotent orbits and Weyl group representations, are pinned together by this simple combinatorial/automata point of view. The experts are invited to contrast the point of view presented here with the seemingly deeper (and at any rate, more sophisticated) connections furnished by the Springer correspondence. I should also point out things that the domain of this point of view can be extended to that of reductive algebraic groups and the real reductive Lie groups forms (although in this extension the basic combinatorial datum determining the group must be enhanced to that of a *root datum*, but even so, the possibilities here are dictated firstly by the choice of a simple system.)

1.1. Weyl groups. Let $\Pi = \{\alpha_1, \dots, \alpha_r\}$ be a simple system inside a vector space V . Set

$$c_{ij} = \frac{2 \langle \alpha_i, \alpha_j \rangle}{\langle \alpha_i, \alpha_j \rangle} \in \{0, 2, -1, -2, -3\}$$

and define the linear transformations

$$s_i : V \rightarrow V \quad : \quad v \longmapsto v - \frac{2 \langle v, \alpha_i \rangle}{\langle \alpha_i, \alpha_i \rangle} \alpha_i \quad , \quad i = 1, \dots, r \quad .$$

Such transformations are *reflections* in V ; and as such generate a reflection subgroup W_Π of $O(V)$. The group W_Π is, in fact, a finite group; it is the *Weyl group of Π* .

Let G_Π be the simply connected complex Lie group corresponding to the simple system Π (which, in fact, can be constructed directly from Π). The maximal torii of G_Π are fundamental objects in the representation theory of G_Π (e.g., the maximal torii are in fact critical to parameterizations both finite and infinite dimensional representations of G_Π). Let T be a maximal torus in G_Π and let N be the normalizer of T in G . Then

$$W_{analytic} \equiv N/T \approx W_\Pi$$

Note that the Weyl group is, on the one hand, in very close proximity to the combinatorial data that determines a simply connected, simple, complex Lie group and, on the other hand, also arises *way-down-the-road* as a finite group associated with the representation theory of the group.

1.2. Nilpotent Orbits. The notion of a nilpotent orbit is crucial idea in modern representation theory. These objects arise as follows (or at least, the following seems to be a minimal method for constructing them). A simple system Π determines a simple Lie algebra $\mathfrak{g} = \mathfrak{g}_\Pi$, which in turn exponentiates to the group G which in turn acts on \mathfrak{g} via its adjoint representation. If $x \in \mathfrak{g}$ is such that ad_x acts nilpotently on \mathfrak{g} , then the same is true for every $Ad(g) \cdot x$. Sets of the form

$$\mathcal{O}_x = G \cdot x \quad , \quad x \text{ nilpotent}$$

are called *nilpotent orbits*

There's a number of ways nilpotent orbits enter modern representation theory.

First of all, there's a subprogram of the *Unitary Dual Problem* called the *orbit philosophy*. In this program, the problem of parameterizing all unitary representations of a given Lie group G , is regarded as the problem of "quantizing" coadjoint orbits. You see, the coadjoints orbits are always symplectic manifolds and as such are interpretable as the *phase space* of a classical mechanical system. A *quantization* of such a coadjoint orbit \mathcal{O} , would then be a construction that attaches to \mathcal{O} an irreducible unitary representation. The goal of the orbit philosophy is to parameterize the unitary dual via such constructions. Unfortunately, for each semisimple group there remains a finite list of coadjoint orbits for which no uniform construction of a unitary representation from the orbit is known. These problematic orbits are always nilpotent orbits.

Secondly, nilpotent orbits arise as an important invariant of irreducible representations. This invariant, the *associated variety of a representation* is actually interpretable as a *dequantization* of a representation and as such it retains a lot of intuitive information about the original representation (like its relative *size*, *K*-type structure, etc). The nilpotent orbits also provide a basis (double entendre intended) for more sophisticated invariants like the *associated cycle* of a representation.

Below I'll be denoting by \mathcal{N} the set of nilpotent orbits.

1.3. The Springer Correspondence. T.A. Springer's 1978 paper *A construction of Weyl group representations* established a remarkably tight connection between Weyl group representations and nilpotent orbits. His construction goes as follows. Let \mathfrak{g} be a complex semisimple Lie algebra, and let \mathfrak{B} be the flag variety of \mathfrak{g} , that is to say, the variety of all Borel subalgebras of \mathfrak{g} . \mathfrak{B} is a homogeneous space for $G_{ad} = \exp(\pi_{ad})$. Let x be a nilpotent element of \mathfrak{g} and let \mathfrak{B}_x be the subvariety of all Borel subgroups containing x . G_{ad}^x acts on \mathfrak{B}_x , and in fact it is not so hard to show that the induced action of $(G_{ad}^x)^o$ on $H^*(\mathfrak{B}_x, \mathbb{C})$ is trivial. Thus, $A(\mathcal{O}_x) = G_{ad}^x / (G_{ad}^x)^o$ acts on $H^*(\mathfrak{B}_x, \mathbb{C})$. Springer showed:

- There is also a natural action of W on $H^*(\mathfrak{B}_x, \mathbb{C})$ and this action commutes with that of $A = A(\mathcal{O}_x)$.
- Regard $X \equiv H^{top}(\mathfrak{B}_x, \mathbb{C})$ as a $W \times A$ module. As a representation of a finite group it has a unique decomposition (up to ordering and other trivialities)

$$(1) \quad U_x = \bigoplus_{\mu \in \widehat{A}} \pi_{x,\mu} \otimes V_{x,\mu}$$

- In (1), $\widehat{\pi}_\mu$ is either 0 or an irreducible representation of W and V_μ is an μ -isotypical representation of A .
- Each irreducible representation $\pi \in \widehat{W}$ occurs exactly once in such a decomposition; i.e., it occurs exactly once as a $\pi_{x,\mu}$ for a particular (conjugacy class of) x and a particular $\mu \in A(\mathcal{O}_x)$. This fact allows us to parameterize Weyl group representations by their Springer data

$$\pi \in \widehat{W} \leftrightarrow (\mathcal{O}_x, \mu)$$

This is called the *Springer parameterization of \widehat{W}* . Note, however, for a given orbit \mathcal{O}_x not every $\mu \in A(\mathcal{O}_x)$ need occur. Put another way, given a nilpotent element x and a representation $\mu \in A(\mathcal{O}_x)$, the μ -isotypical element of U_x is either empty or carries an irreducible representation $\pi_{x,\mu}$ of W unique for μ and \mathcal{O}_x .

- The trivial isotypical component (where A acts by the identity) is always non-zero in U_x . The corresponding irreducible Weyl group representation $\pi_{x,1}$ is called the *Springer representation* attached to the nilpotent orbit of x .

1.3.1. *Significance of the Springer correspondence.* Given the relative simplicity of Weyl group representation to nilpotent orbits, it is quite bizarre to see such an elaborate construction leading to a parameterization of Weyl group representations via nilpotent orbit data. And moreover, it's so sophisticated, it's hard to see how this parameterization can be put to practical use. But, in fact, what makes the Springer correspondence so important is that, to a large extent, it preserves the inter-relationships amongst nilpotent orbits and Weyl group representations. For example, the partial ordering of nilpotent orbits by inclusion is reflected by a natural partial ordering of Weyl group representations - and Spaltenstein duality amongst orbits is consistent with a corresponding "duality" for Weyl group representations.

In short, while technically obtuse, the Springer correspondence *organizes as it coordinatizes \widehat{W}* and the set of nilpotent orbits in a very useful way. The goal of this talk is to display the Springer relationship between \widehat{W} and \mathcal{N} in terms of similar combinatorial parameterizations for \widehat{W} and \mathcal{N} .

2. COMBINATORIAL PARAMETERIZATIONS OF \mathcal{N} AND \widehat{W}

2.1. Combinatorial Bala-Carter.

2.1.1. *Inclusion and Induction of nilpotent orbits.* Let \mathfrak{g} be a semisimple Lie algebra and let \mathfrak{l} be a Levi subalgebra of \mathfrak{g} . There are two fundamental ways of constructing a nilpotent orbit of a semisimple Lie algebra \mathfrak{g} from a nilpotent orbit $\mathcal{O}_\mathfrak{l}$ of \mathfrak{l} . The first and simplest method is **orbit inclusion**, where one simply constructs its *G-saturation* as

$$\text{inc}_\mathfrak{l}^\mathfrak{g}(\mathcal{O}_\mathfrak{l}) = \{X \in \mathfrak{g} \mid X = \text{Ad}(g)(x) \text{ for some } g \in G_{ad} \text{ and } x \in \mathcal{O}_\mathfrak{l}\}$$

The second method is called **orbit induction**. Here one picks an extension $\mathfrak{p} = \mathfrak{l} + \mathfrak{n}$ of \mathfrak{l} to a parabolic subalgebra of \mathfrak{g} , and then constructs the *G-saturation* of $\mathcal{O}_\mathfrak{l} + \mathfrak{n}$. Inside this *G-saturation* there will be a unique dense orbit (there will be several orbits in general, but a unique orbit of maximal dimension). We set

$$\text{ind}_\mathfrak{l}^\mathfrak{g}(\mathcal{O}_\mathfrak{l}) = \text{unique dense orbit in } G \cdot (\mathcal{O}_\mathfrak{l} + \mathfrak{n})$$

This orbit turns out to be independent of the choice of \mathfrak{n} (the extending nilradical).

An important special case of the latter construction is that of **Richardson orbits**. Here one begins with the trivial orbit $\mathbf{0}_\mathfrak{l}$ of a Levi subalgebra and then induces that up to a nilpotent orbit of \mathfrak{g}

$$\text{Richardson}(\mathfrak{l}) = \text{ind}_\mathfrak{l}^\mathfrak{g}(\mathbf{0}_\mathfrak{l})$$

A thing to note about Richardson orbits (which actually extends to induced orbits in general), the smaller the Levi \mathfrak{l} , the larger the nilradical \mathfrak{n} and so the larger the resulting Richardson orbit. Indeed, if we take

$\mathfrak{l} = \mathbf{0}$ (the zero Levi), then

Richardson ($\mathbf{0}$) = $d_{\mathbf{0}}^{\mathfrak{g}}$ ($\mathbf{0}_{\mathbf{0}}$) = unique dense orbit in $G \cdot \mathbf{n}_{\mathfrak{b}}$ = the principal (maximal) nilpotent orbit while

$$\textit{Richardson}(\mathfrak{g}) = \mathbf{0}_{\mathfrak{g}} \text{ the trivial (minimal) nilpotent orbit}$$

2.1.2. The Bala-Carter Theorem.

Definition 2.1. A nilpotent element x (or its associated nilpotent orbit \mathcal{O}_x) of \mathfrak{g} is said to be distinguished if it does not reside in any proper Levi subalgebra of \mathfrak{g} . (If this is true for x , it will also be true for any G -conjugate of x).

Theorem 2.2 (Bala-Carter). There is a 1 : 1 correspondence between nilpotent orbits of \mathfrak{g} and G -conjugacy classes of pairs $(\mathfrak{l}, \mathcal{O}_{\mathfrak{l}})$ where \mathfrak{l} is a Levi subalgebra of \mathfrak{g} and $\mathcal{O}_{\mathfrak{l}}$ is a distinguished orbit in \mathfrak{l} .

How this works is relatively easy to describe, the proof of course lies in the details. Start with a representative nilpotent $x \in \mathcal{O}$, if it is not distinguished, it lies in some proper Levi $\mathfrak{l} \subsetneq \mathfrak{g}$, if x is distinguished in \mathfrak{l} then it lives inside some proper subalgebra \mathfrak{l}' of \mathfrak{l} . Eventually, we find a minimal Levi \mathfrak{l}_{\min} for which x is distinguished. It turns out that all such minimal Levis are G -conjugate and

$$\mathcal{O} = G \cdot x = \text{inc}_{\mathfrak{l}_{\min}}^{\mathfrak{g}} (L_{\min} \cdot x) \quad .$$

shows how to construct \mathcal{O} from a distinguished orbit inside a Levi.

2.1.3. *Combinatorial parameters.* The problem now is to figure out a way of counting conjugacy classes of Levi subalgebras and distinguished orbits inside Levi subalgebras. The former is easy. Every Levi subalgebra is G -conjugate to a *standard Levi subalgebra* attached to a subset Γ of the simple roots for \mathfrak{g} and a Cartan-Weyl basis for \mathfrak{g} :

$$\mathfrak{l}_{\Gamma} = \mathfrak{t} \oplus \bigoplus_{\alpha \in \Delta(\Gamma)} \mathfrak{g}_{\alpha}$$

where $\Delta(\Gamma)$ is the root subsystem generated by Γ . If $\Gamma, \Gamma' \subset \Pi$ are two sets of simple roots, then

$$\mathfrak{l}_{\Gamma} \text{ is } G\text{-conjugate to } \mathfrak{l}_{\Gamma'} \iff \Gamma \text{ is } W\text{-conjugate to } \Gamma'$$

Thus,

$$\{\text{Levi subalgebras}\} / G \approx 2^{\Pi} / W$$

(the power set of Π modulo W).

To parameterize distinguished orbits of \mathfrak{l}_{Γ} , one first shows that a distinguished orbit is necessarily a Richardson orbit. Moreover, the inducing Levi $\mathfrak{l}_{\gamma} \subset \mathfrak{l}_{\Gamma}$ has to satisfy a certain combinatorial condition: viz.,

$$\begin{aligned} \#\Delta_{\gamma} + \#\Gamma &= \#\{\alpha \in \Delta_{\Gamma}^+ \mid \alpha = \alpha_1 + \alpha_2 \text{ with } \alpha_1 \in \Delta_{\gamma}, \alpha_2 \in \Gamma \setminus \gamma\} \\ &= \text{number of positive roots in } \Delta_{\Gamma} \text{ representable as a sum of} \\ &\quad \text{a root in } \Delta_{\gamma} \text{ and a simple root of } \Gamma \text{ not in } \gamma \end{aligned} \quad (2)$$

Accordingly, we call such a subset $\gamma \subset \Gamma$ a *distinguished subset* of Γ . It turns out

$$\text{ind}_{\mathfrak{l}_{\gamma}}^{\mathfrak{l}_{\Gamma}}(\mathbf{0}) = \text{ind}_{\mathfrak{l}_{\gamma'}}^{\mathfrak{l}_{\Gamma}}(\mathbf{0}) \iff \gamma' = w \cdot \gamma \text{ for some } w \in W_{\Gamma}$$

Remark 2.3. There's an easy way to determine if two subsets of simple roots are W -conjugate. Form the generalized Coxeter elements in W

$$c_{\gamma} = \prod_{\alpha \in \gamma} s_{\alpha}$$

Then γ is W -conjugate to γ' if and only if c_{γ} and $c_{\gamma'}$ are conjugate in W (and so, e.g. share the same standard reduced expression).

Putting this all together:

Theorem 2.4. *The nilpotent orbits of \mathfrak{g} are in a 1 : 1 correspondence with (equivalence classes of) ordered pairs (γ, Γ) of subsets of the simple roots of \mathfrak{g} where*

- $\gamma \subset \Gamma \subset \Pi$
- $\#\Delta_\gamma + \#\Gamma = \#\{\alpha \in \Delta_\Gamma^\pm \mid \alpha = \alpha_1 + \alpha_2 \text{ with } \alpha_i \in \Delta_\gamma, \alpha_2 \in \Gamma \setminus \gamma\}$
(Here γ is up to a W_Γ equivalent subset, and Γ is up to a $W_\mathfrak{k}$ equivalent subset. Using the properties of generalized Coxeter elements (cf. remark above), it is easy to fix unique standard representatives for γ and Γ).

Explicitly, the correspondence is

$$(\gamma, \Gamma) \longmapsto inc_{\Gamma}^{\mathfrak{g}} \left(ind_{\mathfrak{l}_\gamma}^{\mathfrak{l}_\Gamma} (\mathbf{0}_{\mathfrak{l}_\gamma}) \right)$$

2.2. Truncated induction of Weyl Group representations. Given a finite group G and a representation π_H of a normal subgroup H , one has a standard construction

$$ind_H^G(\pi_H)$$

of a representation of G . This construction, however, rarely takes irreducible representations of H to irreducible representations. In the Weyl group situation, however, there is a modification of this procedure that does take irreducibles to irreducibles. To set it up, we need a couple preliminaries.

Let \mathfrak{h} be a Cartan subalgebra of \mathfrak{g} . The Weyl group acts naturally by reflections on \mathfrak{h}^* and hence induces a representation of W on $S(\mathfrak{h})$ the symmetric algebra of \mathfrak{h} . Thinking of this as an action on polynomials, it preserves degrees. Moreover, for each irreducible representation π of W , there is a unique minimal integer n for which π appears as a W -isotypical component of $S^n(\mathfrak{h})$. This integer n is called the *fake degree* of π (as it can be regarded as a simplification of a more complicated *generic degree* of π).

A W -representation π is called *univalent* if the multiplicity of π in $S^{\det(\pi)}(\mathfrak{h})$ is 1. For type A_n all irreducible representations of W are univalent. For more general G , there's typically a small handful of non-univalent representations. One representation that is always univalent (and which has fake degree equal to number of positive roots) is the sign representation of W which is the 1-dimensional representation of W given by

$$sgn(w) = \begin{cases} +1 & \text{if } w \text{ has a reduced expression consisting of an even number of simple factors} \\ -1 & \text{if } w \text{ has a reduced expression consisting of an odd number of simple factors} \end{cases}$$

Theorem 2.5 (Macdonald, Lusztig, Spaltenstein). *Let Π be a simple system and let Γ be a subset of Π . Suppose π is a univalent representation of W_Γ of fake degree n and set*

$$X = Ind_{W_\Gamma}^W(\pi)$$

Then inside X there is a unique irreducible summand of X of minimal fake degree (which is also n).

The representation determined by the theorem above is called the *representation of W obtained from π by truncated induction* (“truncated” because we end up ignoring the summands of X then involve representations of higher fake degree). We shall denote this representation by

$$j_{W_\Gamma}^W(\pi)$$

An important property of truncated induction is that it preserves the property of univalence (and so truncated induction can be applied repeatedly). On the other, when π is not univalent, one no longer has a unique summand of minimal fake degree in $X = ind_{W_\Gamma}^W(\pi)$.

Of particular interest is the case when π is the sign representation. The irreducible representation

$$j_{W_\Gamma}^W(sgn_{W_\Gamma})$$

is called the *Macdonald representation corresponding to Γ* . Such representations are analogous to Richardson orbits, in that they really only depend on a (W -conjugacy class of a) Levi subsystem $\Gamma \subset \Pi$. In fact, the

following diagram holds:

$$(3) \quad \Gamma \subset \Pi \quad \begin{array}{c} \nearrow j_{W_\Gamma}^W(\varepsilon_\Gamma) \\ \downarrow \text{Springer} \\ \searrow \text{ind}_{\Gamma}^{\mathfrak{g}}(\mathbf{0}_t) \end{array}$$

where the Springer correspondence is the assignment to the nilpotent orbit $G \cdot x = j_{W_\Gamma}^W(\text{sgn}_{W_\Gamma})$ the unique irreducible W -representation occurring carried by the $A(\mathcal{O}_x)$ -trivial summand of $H^{\text{top}}(\mathfrak{B}_x, \mathbb{C})$. Other words, the Springer map above is the correspondence $\mathcal{N} \longleftrightarrow \widehat{W}_{\text{Springer}}$ that sends an orbit \mathcal{O} to the W -representation with Springer parameters $(\mathcal{O}, \mathbf{1})$. The diagram above exemplifies the theme announced in my introduction: both Macdonald representations and Richardson orbits are constructible for a simple combinatorial datum Γ (a choice of subset of the simple roots), and the Springer correspondence verifies that these are the “right way” to parameterize Macdonald representations and Richardson orbits.

(Intermission)

3. SPECIAL REPRESENTATIONS AND SPECIAL ORBITS

We now seek to extend the combinatorial correspondence of diagram (3) to a wider classes of W -reps and orbits. The first wider classes will be the *special representations* and the *special orbits*. The most expedient way to define these classes is as follows

Definition 3.1. Let d_1, \dots, d_r be the degrees of the basic polynomial invariants of W (exponents of Π plus 1) and let ϕ be the character of an irreducible representation of W . The **generic degree polynomial** of ϕ is

$$P_\phi(t) = \frac{1}{|W|} \prod_{i=1}^r (1 - t^{d_i}) \sum_{w \in W} \frac{\phi(w)}{\det(1 - tw)}$$

The **generic degree of ϕ** is the degree of the leading term of $P_\phi(t)$.

Remark 3.2. In general, the generic degree of $\pi \in \widehat{W}$ is \leq its (fake) degree.

Definition 3.3. An irreducible representation of W is called **special** if the degree of its generic degree polynomial is the same as its (fake) degree.

Definition 3.4. A nilpotent orbit is called **special** if the Springer correspondence $\mathcal{O} \rightarrow \pi_{(\mathcal{O}, \mathbf{1})}$ maps \mathcal{O} to a special representation of W .

Remark 3.5. Macdonald representations are always special and Richardson orbits are always special.

Remark 3.6. Although our definition of special orbits has been pinned to that of special representations of W , there is an alternative way of defining them that is independent of the Springer correspondence. Let π be an irreducible (\mathfrak{g}, K) module of regular integral infinitesimal character and let \mathcal{O}_π be the unique dense orbit in the associated variety of the annihilator of π in $U(\mathfrak{g})$. Then \mathcal{O}_π is a special orbit and all special orbits arise this way.

4. DUALITIES

If \mathfrak{g} is of type A_{n-1} then both its orbits and W -reps can be parameterized in terms of partitions of n . In this case as well, every W -representation and every nilpotent orbit is special. Moreover, there is a natural involution on both \mathcal{N} and \widehat{W} where an orbit (respectively, W -rep) parameterized by a partition \mathbf{p} is sent to the orbit (respectively, W -rep) parameterized by the partition transpose \mathbf{p}^t of \mathbf{p} . (This transpose can be

carried out by rotating the Young diagram corresponding to \mathbf{p} by 90° .) This orbit duality for A_{n-1} was generalized by Spaltenstein as follows:

Theorem 4.1 (Spaltenstein). *Let \mathfrak{g} be a simple Lie algebra. Then there is a unique map $d : \mathcal{N}_{\mathfrak{g}} \rightarrow \mathcal{N}_{\mathfrak{g}}$ such that*

- $d^2(\mathcal{O}) \leq \mathcal{O}$ (partial order by inclusion of closures)
- $d(\text{inc}_{\mathfrak{l}}^{\mathfrak{g}}(\mathcal{O}_{\text{prin}})) = \text{ind}_{\mathfrak{l}}^{\mathfrak{g}}(\mathbf{0}_{\mathfrak{l}})$
- $d(\mathcal{O})$ is always a special orbit.

Remark 4.2. While there were known algorithms for carrying out the action of d on the classical groups, the existence and uniqueness of d on the exceptional groups had to be done by hand. Barbasch and Vogan then provided a souped up version of the duality idea usually general representation-theoretical maps:

Theorem 4.3 (Barbasch-Vogan, 1985). *Let \mathfrak{g} be a semisimple Lie algebra and let \mathfrak{g}^{\vee} be its dual Lie algebra (short roots \longleftrightarrow long roots). Consider the map $\eta_{\mathfrak{g}} : \mathcal{N}_{\mathfrak{g}} \rightarrow \mathcal{N}_{\mathfrak{g}^{\vee}}$ defined by*

$$\begin{aligned} \mathcal{O} \ni x \rightarrow \{x, h, y\} &\rightarrow \frac{1}{2}h = \mu_{\mathcal{O}} \in (\mathfrak{h}^{\vee})^* \rightarrow J_{\mathcal{O}} = \max\{J \in \text{Prim}(U(\mathfrak{g}^{\vee})) \text{ with inf. char } \mu_{\mathcal{O}}\} \\ &\rightarrow \text{AssocVar}(U(\mathfrak{g}^{\vee})/J_{\mathcal{O}}) \xrightarrow{\text{unique dense orbit}} \eta_{\mathfrak{g}}(\mathcal{O}) \in \mathcal{N}_{\mathfrak{g}^{\vee}} \end{aligned}$$

Then

- If $\mathcal{O}_2 \subset \overline{\mathcal{O}_1}$ then $\eta_{\mathfrak{g}}(\mathcal{O}_1) \subset \overline{\eta_{\mathfrak{g}}(\mathcal{O}_2)}$ (order reversing)
- $\eta_{\mathfrak{g}} \circ \eta_{\mathfrak{g}^{\vee}} \circ \eta_{\mathfrak{g}} = \text{id}_{\mathfrak{g}}$
- $\text{Image}(\eta_{\mathfrak{g}}) = \{\text{special nilpotent orbits}\}$

In addition, Barbasch and Vogan proved the following theorem which exhibits our two methods of constructing nilpotent orbits from nilpotent orbits of Levis as being dual with respect to $\eta_{\mathfrak{g}}$:

Theorem 4.4. *If $\mathcal{O}_{\Gamma^{\vee}}$ is a nilpotent orbit in a Levi subalgebra of \mathfrak{g}^{\vee} of Γ^{\vee} , then*

$$\eta_{\mathfrak{g}^{\vee}}(\text{inc}_{\Gamma^{\vee}}^{\mathfrak{g}^{\vee}}(\mathcal{O}_{\Gamma^{\vee}})) = \text{ind}_{\Gamma}^{\mathfrak{g}}(\eta_{\Gamma^{\vee}}(\mathcal{O}_{\Gamma^{\vee}}))$$

5. VARIATIONS ON A FORMULA OF BARBASCH AND VOGAN

Recall that nilpotent orbits can be parameterized by (W -conjugacy classes of) certain *distinguished* subsets γ of subsets Γ of simple roots. For a given \mathfrak{g} , below I'll denote the parameter set by $\mathcal{B}_{\mathfrak{g}}$

$$\mathcal{B}_{\mathfrak{g}} = \{(\gamma, \Gamma) \mid \gamma \subset \Gamma \subseteq \Pi ; \gamma \text{ distinguished for } \Gamma\} / W$$

(cf. (2)).

5.1. An intrinsic characterization of special orbits. Recall our definition of special orbits was tied to that of special representations of W via the Springer correspondence. Let $\mathcal{S}_{\mathfrak{g}} \subset \mathcal{N}_{\mathfrak{g}}$ denote the set of special nilpotent orbits and let Φ be the map from $\mathcal{B}_{\mathfrak{g}^{\vee}}$ to $\mathcal{S}_{\mathfrak{g}}$ constructed as follows:

$$\begin{aligned} \Phi(\gamma^{\vee}, \Gamma^{\vee}) &= \eta_{\mathfrak{g}^{\vee}}(\text{inc}_{\mathfrak{l}_{\Gamma^{\vee}}}^{\mathfrak{g}^{\vee}}(\text{ind}_{\mathfrak{l}_{\gamma^{\vee}}}^{\mathfrak{g}^{\vee}}(\mathbf{0}_{\mathfrak{l}_{\gamma^{\vee}}})) \\ &= \text{ind}_{\Gamma}^{\mathfrak{g}}(\eta_{\Gamma^{\vee}}(\text{ind}_{\mathfrak{l}_{\gamma^{\vee}}}^{\mathfrak{g}^{\vee}}(\mathbf{0}_{\mathfrak{l}_{\gamma^{\vee}}})) \end{aligned}$$

Now the dual of a Richardson orbit corresponding to a Levi $\mathfrak{l}_{\gamma^{\vee}} \subset \mathfrak{l}_{\Gamma^{\vee}}$ is always equal to inclusion of the principal orbit of \mathfrak{l}_{γ} in \mathfrak{l}_{Γ} . Thus,

$$\mathcal{S}_{\mathfrak{g}} = \text{image}(\Phi) = \left\{ \text{ind}_{\Gamma}^{\mathfrak{g}}(\text{inc}_{\mathfrak{l}_{\gamma}}^{\mathfrak{g}}(\mathcal{O}_{\mathfrak{l}_{\gamma}, \text{prin}})) \mid (\gamma^{\vee}, \Gamma^{\vee}) \in \mathcal{B}_{\mathfrak{g}^{\vee}} \right\}$$

Note this characterization of $\mathcal{S}_{\mathfrak{g}}$, as orbits induced from principal orbits of distinguished Levis within Levis, uses the dual parameter set $\mathcal{B}_{\mathfrak{g}^{\vee}}$.

5.2. Duality for Weyl group representations. Let's look at the analog of the Barbasch-Vogan formula (3) via the following transcriptions

$$\begin{aligned} \mathcal{O}_{prin} &\longleftrightarrow \mathbf{I}_W \\ \mathbf{0}_{\mathfrak{g}} &\longleftrightarrow \text{sgn}(W) \\ \text{ind}_{\mathfrak{l}}^{\mathfrak{g}}(\cdot) &\longleftrightarrow \text{TrInd}_{W_{\mathfrak{g}}}^W \\ \eta_{\mathfrak{g}} &\longleftrightarrow \varepsilon_W \quad (\text{Lusztig's involution w/ twist by sign rep}) \end{aligned}$$

Then

$$\Phi : \mathcal{B}_{\mathfrak{g}^\vee} \rightarrow \mathcal{S}_{\mathfrak{g}} : (\gamma^\vee, \Gamma^\vee) \mapsto \eta_{\mathfrak{g}^\vee} \left(\text{inc}_{\mathfrak{e}_{\Gamma^\vee}}^{\mathfrak{g}^\vee} \left(\text{ind}_{\mathfrak{l}_{\Gamma^\vee}}^{\Gamma^\vee} \left(\mathbf{0}_{\Gamma^\vee} \right) \right) \right)$$

translates to

$$\Psi : \mathcal{B}_{\mathfrak{g}^\vee} \rightarrow \widehat{W}_{spec} : (\gamma^\vee, \Gamma^\vee) \mapsto j_{W_{\Gamma}}^{W_{\mathfrak{g}}} \left(\varepsilon_{W_{\Gamma^\vee}} \left(j_{W_{\gamma^\vee}}^{W_{\Gamma^\vee}} (\text{sgn}(W_{\gamma^\vee})) \right) \right)$$

Giving us a direct construction of special representations (Note: $W_{\mathfrak{g}} = W_{\mathfrak{g}^\vee}$, but nevertheless, the use of dual parameters is essential, because e.g. $W_{\mathfrak{l}}$ need not be conjugate to $W_{\mathfrak{l}^\vee}$)

5.3. Construction of Springer (orbit) representations.

Definition 5.1. Let Π_e be the simple roots of \mathfrak{g} together with the lowest root of \mathfrak{g} . Then any proper subset Γ of Π_e will provide a simple subsystem of \mathfrak{g} , as well as a corresponding semisimple subalgebra \mathfrak{l}_{Γ} of \mathfrak{g} . Π_e is called the **extended simple roots** of \mathfrak{g} .

Definition 5.2. Let

$$\mathcal{B}_{e,\mathfrak{g}} = \mathcal{B}_{\mathfrak{g}} = \{(\gamma, \Gamma) \mid \gamma \in \Gamma \subseteq \Pi_e ; \gamma \text{ distinguished for } \Gamma\} / W$$

We refer to $\mathcal{B}_{e,\mathfrak{g}}$ as the **extended Bala-Carter parameters** for \mathfrak{g} .

Theorem 5.3. The image of the map

$$\Psi_e : \mathcal{B}_{e,\mathfrak{g}^\vee} \rightarrow \widehat{W} : (\gamma^\vee, \Gamma^\vee) \mapsto j_{W_{\Gamma}}^{W_{\mathfrak{g}}} \left(\varepsilon_{W_{\Gamma^\vee}} \left(j_{W_{\gamma^\vee}}^{W_{\Gamma^\vee}} (\text{sgn}(W_{\gamma^\vee})) \right) \right)$$

the set of Springer representations of W .

Remark 5.4. Here the use of dual parameters is critical for another reason: a representation of W may be Springer for \mathfrak{g} but not for \mathfrak{g}^\vee (if only because $\#\mathcal{N}_{\mathfrak{g}} \neq \#\mathcal{N}_{\mathfrak{g}^\vee}$).

5.4. The telescoping of the map Ψ .

$$\begin{array}{ccccc} \{\text{Richardson orbits}\} & \subseteq & \{\text{Special Orbits}\} & \subseteq & \{\text{Nilpotent Orbits}\} \\ \{\text{Macdonald Reps}\} & \subseteq & \{\text{Special Reps}\} & \subseteq & \{\text{Orbit Reps}\} \\ \uparrow \Psi & & \uparrow \Psi & & \uparrow \Psi_e \\ 2^\Pi / W \quad (\gamma = \{\}) & \subseteq & \mathcal{B}_{\mathfrak{g}^\vee} & \subseteq & \mathcal{B}_{e,\mathfrak{g}} \end{array}$$

6. CELLS

In view of the telescoping of the map Ψ , it is natural to ask if there a natural extension of the map Ψ so that it produces a class of W -representations that contains the orbit representations. Of course, in such a picture, one would have to also widen the notion of orbits to something like local systems on nilpotents orbits (to accommodate Weyl group representations whose Springer parameters have a non-trivial $\widehat{A}_{\mathcal{O}}$ component). Alas, I haven't been able to do this.

What appears to be a more useful generalization in applications is to relax our focus on irreducible representations of W and instead look at a certain family of reducible representations; the *cell representations* of W .

Cell representations arise several different ways. First of all, there is a direct construction of the representations via induction from "parabolic subgroups" of W (I note here that the native terminology for Weyl groups is a bit askew - a parabolic subgroup of a Coxeter group (S, W) is a subgroup generated by a subset of S - so more like a Levi subgroup in representation theory terminology). Secondly, cell representations arise in the classification theory of primitive ideals in the enveloping algebra of $U(\mathfrak{g})$. And thirdly, cell representations arise special quotients of the W -graphs of Kazhdan-Lusztig theory and in Kazhdan-Lusztig-Vogan theory as the coherent continuation representation of W acting on families of irreducible admissible representations with the same associated variety.

6.1. Lusztig's construction of cell representations. Here is a relatively simple inductive procedure for producing all the cell representations of W . We first need is a modification of the operation of truncated induction. Recall

$$j_{W_\Gamma}^W(\pi_\Gamma) = \text{unique irreducible representation in } \text{ind}_{W_\Gamma}^W(\pi_\Gamma) \text{ with fake degree} = \text{deg}(\pi_\Gamma)$$

Set

$$\tilde{j}_{W_\Gamma}^W(\pi_\Gamma) = \text{direct sum of irreducible representations in } \text{ind}_{W_\Gamma}^W(\pi_\Gamma) \text{ with generic degree} = \text{generic degree of } \pi_\Gamma$$

Then the cell representations of W can be systematically constructed via an inductive procedure that goes as follows:

- The trivial representation $\mathbf{1}$ of W is a cell.
- If $\Gamma \subset \Pi$, and ϕ is a cell representation of W_J then $\tilde{j}_{W_\Gamma}^W(\phi)$ and $\tilde{j}_{W_\Gamma}^W(\phi) \otimes \text{sgn}_W$ are cells of W .