

NOTES FOR MATH 207A

STEVEN V SAM

CONTENTS

1. Introduction	3
1.1. Basic definitions	3
1.2. Basic operations	5
1.3. Representations of \mathfrak{sl}_2	5
1.4. Roadmap	7
2. Enveloping algebras	7
2.1. Definition	7
2.2. PBW basis	9
2.3. Induction	9
3. Back to \mathfrak{sl}_2	10
3.1. Verma modules	10
3.2. Spherical harmonics	11
4. Semisimple Lie algebras	12
4.1. Basic terminology	12
4.2. Classical series	13
4.3. Roots and weights	14
4.4. Root systems	17
4.5. Formal characters	20
4.6. Borel subalgebras	21
4.7. Highest weight representations	21
4.8. Definition of category \mathcal{O}	22
5. The center of $U(\mathfrak{g})$	24
5.1. Central characters	24
5.2. Linked weights	25
5.3. Harish-Chandra's theorem	26
5.4. Application: \mathcal{O} is artinian	27
5.5. Grothendieck group	28
6. Block decomposition	29
6.1. Yoneda Ext	29
6.2. Blocks	30
6.3. Subcategories \mathcal{O}_χ	31
6.4. Dominant and antidominant weights	32
6.5. Duality functor	33
7. Projectives and injectives	36
7.1. Projective objects	36

Date: December 8, 2023.

7.2. Projective covers	38
7.3. Injective objects	41
7.4. Standard filtrations	41
8. BGG resolution	43
8.1. General discussion of resolutions	44
8.2. (Relative) Chevalley–Eilenberg complex	44
8.3. Weights in the relative Chevalley–Eilenberg complex	45
8.4. Bott’s theorem	50
8.5. Complements on the BGG complex	52
8.6. Translation functors	53
9. Kazhdan–Lusztig polynomials	54
9.1. Coxeter groups	54
9.2. Hecke algebras	55
9.3. R-polynomials	55
9.4. Kazhdan–Lusztig polynomials	56
10. Abstract highest weight categories	57
10.1. Highest weight categories	57
10.2. Path algebras	60
10.3. Quasi-hereditary algebras	61
References	65

1. INTRODUCTION

The main object of study in this course is the “category \mathcal{O} ” introduced by Bernstein, Gelfand, and Gelfand in the study of representations of semisimple Lie algebras. This is a certain category of well-behaved representations which includes all of the finite-dimensional representations and makes use of the fact that semisimple Lie algebras have a certain “triangular structure”. I’m not going to assume any specific background knowledge other than algebraic fluency at the level of Math 200. I’ll start with examples and try to maintain doing concrete examples throughout.

Otherwise, these notes are not meant to be polished, but rather they are written in a way that makes it easy for me to remember what should be discussed during each lecture. They should still be readable, though I can’t guarantee it will be enjoyable.

1.1. Basic definitions. First, we’ll almost exclusively be dealing with algebras over the complex numbers \mathbf{C} . A **Lie algebra** consists of a vector space \mathfrak{g} together with a bilinear operation known as a “Lie bracket”

$$[\cdot, \cdot]: \mathfrak{g} \times \mathfrak{g} \rightarrow \mathfrak{g}$$

which is:

- skew-symmetric ($[x, y] = -[y, x]$ for all $x, y \in \mathfrak{g}$) and
- satisfies the **Jacobi identity**:

$$[x, [y, z]] + [y, [z, x]] + [z, [x, y]] = 0$$

for all $x, y, z \in \mathfrak{g}$.

Remark 1.1.1. There’s an important way to rephrase the Jacobi identity (subject to skew-symmetry). Given $x \in \mathfrak{g}$, let $\text{ad}_x: \mathfrak{g} \rightarrow \mathfrak{g}$ be the linear map $\text{ad}_x(y) = [x, y]$. Then ad_x is a *derivation*: for any $y, z \in \mathfrak{g}$, we have

$$\text{ad}_x([y, z]) = [\text{ad}_x(y), z] + [y, \text{ad}_x(z)].$$

Using skew-symmetry, you can see that requiring that ad_x is a derivation for all x is equivalent to the Jacobi identity. \square

Here’s an important example:

Example 1.1.2. Let A be any associative \mathbf{C} -algebra. Define $[x, y] = xy - yx$ (this is called a **commutator**). Then this is a Lie bracket on the underlying vector space of A . The check that the Jacobi identity holds is straightforward but not interesting to spell out. \square

For each integer $n \geq 0$, we let $\mathfrak{gl}_n(\mathbf{C})$ (or just \mathfrak{gl}_n) denote the set of $n \times n$ complex matrices, which we consider as a Lie algebra using the example above, i.e., $[x, y] = xy - yx$ where the product on the right side is the usual matrix multiplication.

Without appealing to bases, if V is a vector space (not necessarily finite-dimensional), let $\mathfrak{gl}(V)$ denote the set of linear operators $X: V \rightarrow V$, which is again a Lie algebra with the bracket $[X, Y] = XY - YX$.

Given Lie algebras $\mathfrak{g}, \mathfrak{g}'$, a **homomorphism** $\varphi: \mathfrak{g} \rightarrow \mathfrak{g}'$ is a linear map such that

$$[\varphi(x), \varphi(y)] = \varphi([x, y])$$

for all $x, y \in \mathfrak{g}$. This lets us define a **representation** of \mathfrak{g} as a homomorphism $\varphi: \mathfrak{g} \rightarrow \mathfrak{gl}(V)$ for some vector space V . Usually, V denotes the representation and the information φ is understood from context.

For $x \in \mathfrak{g}$ and $v \in V$, we usually just write xv instead of $\varphi(x)(v)$.

Example 1.1.3. For any \mathfrak{g} , the linear map $\text{ad}: \mathfrak{g} \rightarrow \mathfrak{gl}(\mathfrak{g})$ is a representation, called the **adjoint representation**. \square

Example 1.1.4. Let $\mathfrak{g} = \mathfrak{gl}_n(\mathbf{C})$ and let $E_{i,j}$ denote the matrix with a 1 in row i , column j , and 0's elsewhere. Take $V = \mathbf{C}[z_1, \dots, z_n]$, the space of polynomials in n variables and define $\varphi: \mathfrak{g} \rightarrow \mathfrak{gl}(V)$ by

$$\varphi(E_{i,j}) = z_i \partial_j$$

where ∂_j means partial derivative with respect to the variable z_j and z_i here means the operation of multiplying by z_i . In other words, given $f \in V$, we have

$$(z_i \partial_j) f = z_i \frac{\partial f}{\partial z_j}.$$

Since the $E_{i,j}$ are a basis, this defines a linear map. To verify that this is a Lie algebra homomorphism, by linearity it is enough to show that

$$\varphi([E_{i,j}, E_{k,\ell}]) = [\varphi(E_{i,j}), \varphi(E_{k,\ell})]$$

for all i, j, k, ℓ . Here are the relevant formulas. First,

$$[E_{i,j}, E_{k,\ell}] = \delta_{j,k} E_{i,\ell} - \delta_{i,\ell} E_{k,j}$$

where δ is the Kronecker delta, i.e., $\delta_{j,k} = \begin{cases} 1 & \text{if } j = k \\ 0 & \text{otherwise} \end{cases}$.

Second, to compute commutators between $z_i \partial_j$, we make a few observations: multiplying by variables is commutative, as is taking partial derivatives. Second, ∂_j and z_k commute if $j \neq k$. Otherwise, we have

$$\partial_i z_i - z_i \partial_i = \text{id}$$

as can be verified by applying both sides to an arbitrary polynomial. This can be more suggestively rewritten as

$$\partial_j z_k = z_k \partial_j + \delta_{k,j} \text{id}.$$

Putting this together gives

$$[z_i \partial_j, z_k \partial_\ell] = \delta_{j,k} z_i \partial_\ell - \delta_{i,\ell} z_k \partial_j.$$

Next, note that each operator $z_i \partial_j$ preserves degree: if $f \in V$ is homogeneous of degree d , then so is $z_i \partial_j f$. Since V is the symmetric algebra on \mathbf{C}^n , we will also denote this by $\text{Sym}(\mathbf{C}^n)$. Let $\text{Sym}^d(\mathbf{C}^n) \subset V$ be the subspace of homogeneous polynomials of degree d . Then each $\text{Sym}^d(\mathbf{C}^n)$ is invariant under the action of \mathfrak{g} and we have representations $\varphi_d: \mathfrak{g} \rightarrow \mathfrak{gl}(\text{Sym}^d(\mathbf{C}^n))$ given by the same formulas. \square

Finally, an annoying point: \mathfrak{gl}_n is natural to work with, but it is not a semisimple Lie algebra (to be defined later). Instead it has a (simple) subalgebra \mathfrak{sl}_n which is the subspace of matrices whose trace is 0 (note that $\text{Tr}(xy - yx) = 0$ for any $n \times n$ matrices x, y , so this is actually a subalgebra).

1.2. **Basic operations.** There are a bunch of standard algebraic notions.

- If V is a representation of \mathfrak{g} , then a subspace $W \subseteq V$ is a **subrepresentation** if, for all $x \in \mathfrak{g}$ and for all $v \in W$, we have $xv \in W$. In that case, the quotient space V/W has a natural structure of a representation called the **quotient representation**.

An **ideal** of \mathfrak{g} is a subrepresentation of the adjoint representation. In other words, a subspace $\mathfrak{a} \subseteq \mathfrak{g}$ such that for all $x \in \mathfrak{g}$ and $y \in \mathfrak{a}$, we have $[x, y] \in \mathfrak{a}$.

- Given a representation V , the **dual representation** on the linear dual V^* is defined by ($x \in \mathfrak{g}$, $f \in V^*$, $v \in V$)

$$(xf)(v) = -f(xv).$$

- Let V, W be representations of \mathfrak{g} .
 - The **direct sum representation** on $V \oplus W$ is given by $x(v, w) = (xv, xw)$.
 - The **tensor product representation** on $V \otimes W$ is given by

$$x\left(\sum_i v_i \otimes w_i\right) = \sum_i (xv_i \otimes w_i + v_i \otimes xw_i).$$

- The space of linear maps $V \rightarrow W$ is denoted $\text{Hom}(V, W)$. Given $\varphi: V \rightarrow W$ and $x \in \mathfrak{g}$, we define an action by

$$(xf)(v) = x(f(v)) - f(xv).$$

If $W = \mathbf{C}$ is trivial, this agrees with the dual representation on V^* . If V is finite-dimensional, this is isomorphic to $V^* \otimes W$. Finally, $f: V \rightarrow W$ is a homomorphism if and only if $xf = 0$ for all $x \in \mathfrak{g}$.

1.3. **Representations of \mathfrak{sl}_2 .** Now let's take a look at the representation $V = \text{Sym}^d(\mathbf{C}^n)$ more carefully in the case $n = 2$. In this case, there is a standard naming convention for basis elements of \mathfrak{sl}_2 :

$$Y = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}, \quad H = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad X = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}.$$

The obvious basis for V is $\{z_1^d, z_1^{d-1}z_2, \dots, z_1z_2^{d-1}, z_2^d\}$. Let's consider the action of the diagonal matrix H . Under φ , we have

$$H = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \mapsto z_1\partial_1 - z_2\partial_2.$$

Each basis element is in fact an eigenvector for $\varphi(H)$ since

$$(z_1\partial_1 - z_2\partial_2)z_1^{d-i}z_2^i = (d - 2i)z_1^{d-i}z_2^i.$$

We'll express this by saying that each $z_1^{d-i}z_2^i$ is a **weight vector** and that its **weight** is $d - 2i$ (I'll give the general setup and definitions later).

Next, consider the strictly upper-triangular matrix:

$$X = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \mapsto z_1\partial_2.$$

This moves between the weight vectors:

$$(z_1\partial_2)z_1^{d-i}z_2^i = iz_1^{d-(i-1)}z_2^{i-1}.$$

We can interpret this as saying that X increases the weight. Similarly, the strictly lower-triangular matrix Y decreases the weight. We say that d is the “highest weight” of this representation.

Now consider a general representation V of \mathfrak{sl}_2 and for each complex number α let V_α be the subspace of eigenvectors for H with eigenvalue α .

To understand what properties we might expect, let’s start with some useful formulas:

$$[H, X] = 2X, \quad [H, Y] = -2Y, \quad [X, Y] = H.$$

Lemma 1.3.1. *Let β be a weight. If $v \in V_\beta$, then $Xv \in V_{\beta+2}$ and $Yv \in V_{\beta-2}$.*

Proof. We have $(HX - XH)v = 2Xv$ since $[H, X] = 2X$. Since $v \in V_\beta$, this simplifies to $HXv = (\beta + 2)Xv$, so $Xv \in V_{\beta+2}$. The other statement is similar. \square

Note that this implies that if $v \in V_\alpha$ and α is a highest weight, then $Xv = 0$. Let’s formalize this:

Definition 1.3.2. A weight vector $v \in V$ is a **highest weight vector** if $Xv = 0$. \square

Define the partial ordering on weights: $\beta \leq \alpha$ means that $\alpha - \beta \in 2\mathbf{Z}_{\geq 0}$.

We’ll say that V is a **highest weight representation** of highest weight α if

- $\dim V_\alpha = 1$,
- α is the unique maximal weight of V , i.e., if $V_\beta \neq 0$ then $\beta \leq \alpha$, and
- V_α generates V , i.e., any subrepresentation containing V_α must be all of V .

All of these properties are satisfied by $\text{Sym}^d(\mathbf{C}^2)$ with $\alpha = d$.

This is the kind of phenomenon we’re trying to capture with a “triangular structure”: we want to decompose our algebra into 3 pieces: (strictly) upper-triangular, diagonal, and (strictly) lower-triangular and we want a good theory of highest weight representations.

If V is a highest weight representation with highest weight α , then V_α is spanned by highest weight vectors, but there could be other highest weight vectors (the terminology is a little confusing, but fairly standard, so keep this in mind).

Now V is a highest weight representation of highest weight α and pick nonzero $v_0 \in V_\alpha$ and define $v_i = Y^i v_0$ for all i . We already know that $v_i \in V_{\alpha-2i}$ (but v_i may be 0).

Lemma 1.3.3. *We have $Xv_i = i(\alpha - i + 1)v_{i-1}$.*

Proof. By induction on i : if $i = 0$, then $Xv_0 = 0$ by assumption that $V_{\alpha+2} = 0$. In general, we have

$$\begin{aligned} Xv_i &= XY^i v_0 = YXY^{i-1} v_0 + HY^{i-1} v_0 \\ &= YXv_{i-1} + (\alpha - 2i + 2)v_{i-1} \\ &= (i - 1)(\alpha - i + 2)Yv_{i-2} + (\alpha - 2i + 2)v_{i-1} \\ &= i(\alpha - i + 1)v_{i-1}. \end{aligned} \quad \square$$

Here are a few consequences of this statement:

- V is spanned by v_0, v_1, v_2, \dots
- If $\alpha \in \mathbf{C} \setminus \mathbf{Z}_{\geq 0}$, then V is irreducible, and $v_i \neq 0$ for all $i \geq 0$. First, the expression $i(\alpha - i + 1)$ is never 0 for $i > 0$, so any subrepresentation containing a nonzero weight vector must be all of V . Furthermore, every nonzero subrepresentation contains a nonzero weight vector.

This is just a general statement about eigenvectors, so suppose $w = w_1 + \cdots + w_n$ is a sum where each w_i is an eigenvector of an operator H with eigenvalue λ_i and all λ_i are distinct. Then

$$(H - \lambda_1)(H - \lambda_2) \cdots (H - \lambda_{n-1})w = (\lambda_n - \lambda_1) \cdots (\lambda_n - \lambda_{n-1})w_n.$$

- On the other hand, if $\alpha \in \mathbf{Z}_{\geq 0}$, then the span of $\{v_i \mid i \geq \alpha + 1\}$ is a subrepresentation V' of V and $v_i \neq 0$ for $i \leq \alpha$. Actually there are two possibilities here. As we've seen with $\text{Sym}^\alpha(\mathbf{C}^2)$, the space of polynomials of degree α , we could have $V' = 0$. But we can also have $V' \neq 0$, and we'll construct an example shortly. In the latter case, $v_{\alpha+1}$ gives a highest weight vector of weight $-\alpha - 2$, which is not itself the highest weight in V . Furthermore, V' is irreducible since $-\alpha - 2 \notin \mathbf{Z}_{\geq 0}$.

1.4. **Roadmap.** Here's the plan for what we should discuss (not necessarily in this order):

- Definition and classification of semisimple Lie algebras (omitting proofs)
- Generalization of above setup for \mathfrak{sl}_2 : highest weight representations and weight spaces. Construction of "free" highest weight representations (= Verma modules).
- Enveloping algebras used in previous step, so need to discuss basic properties.
- Classification of finite-dimensional representations using highest weight theory (omitting proofs), can realize all as quotients of Verma modules.
- The theory has many rich combinatorial gadgets (roots, weights, Weyl groups, etc.) and I'll spell this out as concretely as possible for classical series of semisimple Lie algebras.
- Category \mathcal{O} : Verma modules aren't finite-dimensional, so would be good to have a well-behaved category of representations that contains both them and finite-dimensional representations (studying all representations is too unwieldy).

After that we might do some more advanced topics depending on time left over and interest from audience.

2. ENVELOPING ALGEBRAS

2.1. **Definition.** For this definition, let \mathfrak{g} be an arbitrary Lie algebra (over any field \mathbf{k}). Let $T(\mathfrak{g})$ denote the tensor algebra, i.e.,

$$T(\mathfrak{g}) = \bigoplus_{n \geq 0} \mathfrak{g}^{\otimes n}$$

($\mathfrak{g}^{\otimes 0} = \mathbf{k}$). This is an algebra with the concatenation product, i.e., given $x_i, y_j \in \mathfrak{g}$, we define the product on simple tensors by

$$(x_1 \otimes \cdots \otimes x_m)(y_1 \otimes \cdots \otimes y_n) = x_1 \otimes \cdots \otimes x_m \otimes y_1 \otimes \cdots \otimes y_n$$

and extend linearly for arbitrary elements. This is a graded algebra with $T(\mathfrak{g})_n = \mathfrak{g}^{\otimes n}$ generated by its degree 1 piece $T(\mathfrak{g})_1$.

The **enveloping algebra** of \mathfrak{g} , denoted $U(\mathfrak{g})$, is the quotient of $T(\mathfrak{g})$ by the 2-sided ideal generated by

$$x \otimes y - y \otimes x - [x, y]$$

for all choices of $x, y \in \mathfrak{g}$. Here note that $x \otimes y - y \otimes x \in T(\mathfrak{g})_2$ and $[x, y] \in T(\mathfrak{g})_1$, so the relations are not homogeneous in general (except for the very special case when $[x, y] = 0$ for all x, y , see next example), so $U(\mathfrak{g})$ generally has no obvious structure of a graded algebra.

Example 2.1.1. Consider the special case of an **abelian** Lie algebra \mathfrak{g} where the bracket is always 0. Then the relations in $U(\mathfrak{g})$ simply say that the generators in $T(\mathfrak{g})_1$ all commute with each other, and we can identify it with the symmetric algebra on \mathfrak{g} . If we pick a basis $\{x_i\}$ for \mathfrak{g} , then $U(\mathfrak{g})$ is also isomorphic to the polynomial ring with variables x_i . \square

$U(\mathfrak{g})$ satisfies a universal property. Suppose that A is an associative \mathbf{k} -algebra and $f: \mathfrak{g} \rightarrow A$ is a linear map satisfying

$$f([x, y]) = f(x)f(y) - f(y)f(x).$$

Then this extends uniquely to an algebra homomorphism $f: U(\mathfrak{g}) \rightarrow A$.

In particular, suppose that V is a representation of \mathfrak{g} . Then we can take $A = \mathfrak{gl}(V)$ which is an associative algebra with the usual product. Then the representation extends to an algebra homomorphism $U(\mathfrak{g}) \rightarrow \mathfrak{gl}(V)$, which gives V the structure of a left $U(\mathfrak{g})$ -module. This gives the following important fact (after verifying some routine details):

Proposition 2.1.2. *The above correspondence defines an equivalence of categories between the category of \mathfrak{g} -representations and the category of left $U(\mathfrak{g})$ -modules.*

Remark 2.1.3. If you're familiar with group representations, then there is a similar result where $U(\mathfrak{g})$ is taking the role of the group algebra. \square

Remark 2.1.4. Consider the linear map $f: \mathfrak{g} \rightarrow U(\mathfrak{g}) \otimes U(\mathfrak{g})$ given by $f(x) = x \otimes 1 + 1 \otimes x$. Define an algebra structure on $U(\mathfrak{g}) \otimes U(\mathfrak{g})$ by $(x_1 \otimes y_1)(x_2 \otimes y_2) = x_1x_2 \otimes y_1y_2$ and extend linearly. Then we can check that

$$f([x, y]) = f(x)f(y) - f(y)f(x)$$

so we get an algebra homomorphism $U(\mathfrak{g}) \rightarrow U(\mathfrak{g}) \otimes U(\mathfrak{g})$ which we denote by Δ and call comultiplication. This gives $U(\mathfrak{g})$ the structure of a Hopf algebra and is responsible for the formula for tensor product representations. \square

The algebra $U(\mathfrak{g})$ has an important filtration: for $n \geq 0$, let F^n denote the image of $\bigoplus_{i=0}^n \mathfrak{g}^{\otimes i}$ under the quotient map $T(\mathfrak{g}) \rightarrow U(\mathfrak{g})$ (and set $F^n = 0$ for $n < 0$). If $x \in F^m$ and $y \in F^n$, then $xy \in F^{m+n}$, which means that the associated graded space

$$\text{gr } F^\bullet = \bigoplus_{n \geq 0} F^n / F^{n-1}.$$

inherits an algebra structure. It can be checked (e.g., by induction) that if $x \in F^m$ and $y \in F^n$, then $xy - yx \in F^{m+n-1}$, which means that $\text{gr } F^\bullet$ is a commutative algebra. Furthermore, it is generated by the image of F^1 and $F^0 = \mathbf{k}$. So if we pick a basis $\{x_i\}_{i \in I}$ for \mathfrak{g} , then we get a surjective map from the polynomial ring

$$\mathbf{k}[x_i \mid i \in I] \rightarrow \text{gr } F^\bullet.$$

This gives the following important result.

Proposition 2.1.5. *If $\dim \mathfrak{g} < \infty$, then $U(\mathfrak{g})$ is left-noetherian (and also right-noetherian).*

Proof. First, given a left ideal J of $U(\mathfrak{g})$, we define a filtration on it by $F^n J = F^n \cap J$. I'll leave these checks as exercises:

- (1) $\text{gr } F^\bullet J := \bigoplus_{n \geq 0} F^n J / F^{n-1} J$ is an ideal of $\text{gr } F^\bullet$.
- (2) If $J \subsetneq J'$ is a strict inclusion of left ideals in $U(\mathfrak{g})$, then we get a strict inclusion $\text{gr } F^\bullet J \subsetneq \text{gr } F^\bullet J'$.

If $U(\mathfrak{g})$ were not left-noetherian, there is an infinite chain of strict inclusions of left ideals. This leads to one inside $\text{gr } F^\bullet$, but that violates the Hilbert basis theorem. You can also do the above with “left” replaced by “right”. \square

Remark 2.1.6. Curiously, it is not known whether $U(\mathfrak{g})$ being left-noetherian implies that $\dim \mathfrak{g} < \infty$. In other words, if \mathfrak{g} is infinite-dimensional, we might expect that $U(\mathfrak{g})$ is not left-noetherian. This is known for various classes of Lie algebras, but the general problem is open. The above argument breaks down but that doesn’t say anything about the problem. \square

2.2. PBW basis. Pick a basis $\{x_i\}_{i \in I}$ for \mathfrak{g} along with a total ordering on the index set I (for our applications, I is going to be finite, but we don’t need to make that restriction here).

First note that the simple tensors $x_{i_1} \otimes \cdots \otimes x_{i_n}$ form a basis for the tensor algebra $T(\mathfrak{g})$ and hence give a spanning set for $U(\mathfrak{g})$. For the moment, let’s call these elements “monomials” and we’ll call it an “ordered monomial” if $i_1 \leq \cdots \leq i_n$. Note that every monomial can be written as a linear combination of ordered monomials: if $i > j$, then we have

$$\cdots x_i \otimes x_j \cdots = \cdots x_j \otimes x_i \cdots + \cdots [x_i, x_j] \cdots .$$

The second term has strictly smaller degree while the first term has less inversions (pairs, not necessarily consecutive, that are out of order) in its index set, so this process will eventually terminate.

Theorem 2.2.1 (Poincaré, Birkhoff, Witt; **PBW theorem**). *The set of ordered monomials is a basis for $U(\mathfrak{g})$.*

We will not discuss the proof.

Corollary 2.2.2. *The natural map $\mathfrak{g} \rightarrow U(\mathfrak{g})$ given by $x \mapsto x$ is injective.*

Proof. A basis for \mathfrak{g} maps to a linearly independent set of $U(\mathfrak{g})$. \square

Corollary 2.2.3. *If $\mathfrak{a} \subset \mathfrak{g}$ is a Lie subalgebra, then the natural map $U(\mathfrak{a}) \rightarrow U(\mathfrak{g})$ is injective, and $U(\mathfrak{g})$ is free as a left (or right) $U(\mathfrak{a})$ -module.*

Proof. Pick a basis $\{x_i\}_{i \in I}$ for \mathfrak{a} and extend it to a basis $\{x_i\}_{i \in J}$ for \mathfrak{g} . If we order them so that $i < j$ for all $i \in I$ and $j \in J \setminus I$, then the PBW theorem tells us that the set of increasing monomials using the basis vectors indexed by $J \setminus I$ give a basis for $U(\mathfrak{g})$ as a left $U(\mathfrak{a})$ -module. For the statement about right modules, we can instead order J so that $i > j$ for all $i \in I$ and $j \in J \setminus I$. \square

2.3. Induction. Let $\mathfrak{a} \subset \mathfrak{g}$ be a Lie subalgebra of \mathfrak{g} and let M be a representation of \mathfrak{a} . This is equivalently a left $U(\mathfrak{a})$ -module. Also, $U(\mathfrak{g})$ is naturally a right $U(\mathfrak{a})$ -module, so we can define the tensor product:

$$\text{Ind}_{\mathfrak{a}}^{\mathfrak{g}}(M) := U(\mathfrak{g}) \otimes_{U(\mathfrak{a})} M.$$

Recall that this just means the usual tensor product (of abelian groups) $U(\mathfrak{g}) \otimes M$ modulo the relations $xa \otimes m = x \otimes am$ for all $x \in U(\mathfrak{g})$, $a \in U(\mathfrak{a})$, and $m \in M$. This has the structure of a left $U(\mathfrak{g})$ -module via

$$x(y \otimes m) = xy \otimes m$$

for all $x, y \in U(\mathfrak{g})$ and $m \in M$ (this is well-defined since $U(\mathfrak{g})$ is associative). The resulting $U(\mathfrak{g})$ -module (or \mathfrak{g} -representation) is called the **induction** of M from \mathfrak{a} to \mathfrak{g} .

The PBW theorem lets us write down a basis for the induction as follows. First, pick a basis $\{z_k\}_{k \in K}$ for M . Next, pick a basis $\{x_i\}_{i \in I}$ for \mathfrak{a} and extend it to a basis $\{x_i\}_{i \in J}$ for \mathfrak{g} . Order J so that $i > j$ for all $i \in I$ and $j \in J \setminus I$. Then the set

$$\{x_{j_1} \otimes \cdots \otimes x_{j_n} \otimes z_k \mid j_1 \leq \cdots \leq j_n \in J \setminus I, k \in K\}$$

is the desired basis for $\text{Ind}_{\mathfrak{a}}^{\mathfrak{g}}(M)$. It can be a bit unwieldy to compute the action of \mathfrak{g} on this basis, but we have a well-defined procedure: given $y \in \mathfrak{g}$ and a basis element $x_{j_1} \otimes \cdots \otimes x_{j_n} \otimes z_k$ as above, we first rewrite $y \otimes x_{j_1} \otimes \cdots \otimes x_{j_n}$ as a linear combination of increasing monomials. In any such monomial, say $x_{j'_1} \otimes \cdots \otimes x_{j'_d}$, if $j'_r, j'_{r+1}, \dots, j'_d \in I$ and $j'_{r-1} \notin I$, then we replace the term

$$x_{j'_1} \otimes \cdots \otimes x_{j'_d} \otimes z_k$$

with

$$x_{j'_1} \otimes \cdots \otimes x_{j'_{r-1}} \otimes x'_{j'_r} \cdots x'_{j'_d} z_k.$$

3. BACK TO \mathfrak{sl}_2

Now let's take $\mathfrak{g} = \mathfrak{sl}_2$. Recall the basis we're using

$$Y = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}, \quad H = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad X = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}.$$

3.1. Verma modules. Let \mathfrak{h} be the span of H , which is an abelian subalgebra of \mathfrak{g} and let \mathfrak{b} be the span of H, X . Note that we have a quotient map $\mathfrak{b} \rightarrow \mathfrak{h}$ (since the span of X is an ideal). So any representation of \mathfrak{h} can be pulled back to give a representation of \mathfrak{b} (on which X acts by 0).

Going back to our definitions, given a representation V of \mathfrak{sl}_2 , a vector $v \in V$ is a highest weight vector if $Xv = 0$ and v is a weight vector, say $v \in V_{\alpha}$. Note that this is equivalent to the span of v being a \mathfrak{b} -subrepresentation of V .

To formalize this, pick $\alpha \in \mathbf{C}$. Then we get a one-dimensional representation $\mathfrak{h} \rightarrow \mathfrak{gl}_1$ which sends H to the 1×1 matrix (α) . Let's call this representation \mathbf{C}_{α} and let $z \in \mathbf{C}_{\alpha}$ denote 1 to avoid confusing notation. This can also be made into a \mathfrak{b} -representation on which X acts by 0 by the above comments, and we'll continue to use the notation \mathbf{C}_{α} .

Then the highest weight vectors of V with weight α are parameterized by the space

$$\text{Hom}_{\mathfrak{b}}(\mathbf{C}_{\alpha}, V).$$

By hom-tensor adjunction, this is the same as

$$\text{Hom}_{\mathfrak{g}}(\text{Ind}_{\mathfrak{b}}^{\mathfrak{g}} \mathbf{C}_{\alpha}, V).$$

Let's define $M(\alpha) = \text{Ind}_{\mathfrak{b}}^{\mathfrak{g}} \mathbf{C}_{\alpha}$ and call it a **Verma module**. In particular, if V is a highest weight representation of highest weight α , then it is a quotient of $M(\alpha)$.

Now let's analyze $M(\alpha)$. First, it is itself a highest weight representation of highest weight α since it's generated by the vector $1 \otimes z$ and $X(1 \otimes z) = 1 \otimes Xz = 0$.

Using the notation from last time, we define

$$v_0 = 1 \otimes z, \quad v_i = Y^i \otimes z \quad (i \geq 1).$$

For uniformity, write Y^0 for 1. We see that the v_i form a basis for $M(\alpha)$ (in fact, the PBW basis if we order the basis vectors $Y < H < X$).

Now let's compute the action of \mathfrak{g} on $M(\alpha)$. We've seen how this works before, but let's do it again using this new context to practice the PBW rewriting rule.

First, the easy one:

$$Yv_i = Y(Y^i \otimes z) = Y^{i+1} \otimes z = v_{i+1}$$

and no rewriting needs to be done.

Now consider the action of H . If $i = 0$, no rewriting is needed and we have $H(1 \otimes z) = 1 \otimes Hz = \alpha(1 \otimes z)$. For $i > 0$, we have $HY^i = YHY^{i-1} - 2Y^i$ since $[H, Y] = -2Y$. In other words, each time the H moves past a Y , it creates -2 copies of v_i , so we have

$$H(Y^i \otimes z) = Y^i \otimes Hz - 2i(Y^i \otimes z) = (\alpha - 2i)(Y^i \otimes z).$$

Finally, the action of X is a little more complicated. First, $X(1 \otimes z) = 0$ as we saw. For $i > 0$, we have (using what we just computed)

$$X(Y^i \otimes z) = YXY^{i-1} \otimes z + HY^{i-1} \otimes z = YXY^{i-1} \otimes z + (\alpha - 2i + 2)(Y^{i-1} \otimes z).$$

Iterating this $i - 1$ times will give us what we had before:

$$X(Y^i \otimes z) = i(\alpha - i + 1)(Y^{i-1} \otimes z).$$

Let's summarize a few points:

- If $\alpha \notin \mathbf{Z}_{\geq 0}$, then the only highest weight vector in $M(\alpha)$ is v_0 . In particular,

$$\mathrm{Hom}_{\mathfrak{g}}(M(\beta), M(\alpha)) = \begin{cases} 0 & \text{if } \beta \neq \alpha \\ \mathbf{C} & \text{if } \beta = \alpha \end{cases}.$$

In this case, $M(\alpha)$ is an irreducible representation.

- If $\alpha \in \mathbf{Z}_{\geq 0}$, then $v_{\alpha+1}$ is a highest weight vector of weight $-\alpha - 2$, and this is the only new one besides v_0 (ignoring scalar multiples). The \mathfrak{g} -submodule generated by $v_{\alpha+1}$ is the span of v_i for $i \geq \alpha + 1$. This is also the image of the map

$$M(-\alpha - 2) \rightarrow M(\alpha).$$

By construction, this map is given by $Y^i \otimes z \mapsto Y^{i+\alpha+1} \otimes z$, so it is injective by PBW. The quotient by this subrepresentation (or the cokernel of this map) can be shown to be isomorphic to the space of homogeneous polynomials in 2 variables of degree α (which we discussed earlier). Furthermore, $M(-\alpha - 2)$ is irreducible (either using the previous case or the explicit analysis here).

Next we'd like to generalize this setup for general semisimple Lie algebras, so we'll work on introducing what they are after the next example.

3.2. Spherical harmonics. Here's another example to illustrate highest weight representations of \mathfrak{sl}_2 .

Consider the polynomial ring $P = \mathbf{C}[z_1, \dots, z_n]$ and the quadratic polynomial

$$q = z_1^2 + z_2^2 + \dots + z_n^2.$$

The group of invertible linear change of coordinates that preserves q is an orthogonal group $\mathbf{O}_n(\mathbf{C})$ (think of it as a complexified version of the usual real orthogonal group). We'll be interested in operators that commute with the action of $\mathbf{O}_n(\mathbf{C})$. We will think of q as the operator of multiplication by q .

First, multiplication by q commutes with $\mathbf{O}_n(\mathbf{C})$ by definition. Second, so does the Laplacian

$$\Delta = \partial_1^2 + \partial_2^2 + \dots + \partial_n^2.$$

Now define the Euler operator

$$E = z_1\partial_1 + \cdots + z_n\partial_n.$$

Note that if f is homogeneous, then $Ef = (\deg f)f$, so this also commutes with $\mathbf{O}_n(\mathbf{C})$. Here are their commutator relations:

$$[E, q] = 2q, \quad [E, \Delta] = -2\Delta, \quad [\Delta, q] = 4E + 2nI,$$

where I just means the identity operator.

To check these, we can reduce to the case $n = 1$ because distinct variables and derivatives commute with each other. In the $n = 1$ case, we can write out the commutators explicitly and then repeatedly use the relation $z\partial = \partial z - I$ (to do this systematically, we want every term to have all ∂ before z , this is similar to the PBW basis idea). An alternative is to just check that both sides of each identity do exactly the same thing to any polynomial f .

Comparing this with the relations for \mathfrak{sl}_2 (recall they are $[H, X] = 2X$, $[H, Y] = -2Y$, $[X, Y] = H$), we can define a representation $\rho: \mathfrak{sl}_2 \rightarrow \mathfrak{gl}(P)$ by

$$\rho(X) = \frac{1}{2}\Delta, \quad \rho(H) = -E - \frac{n}{2}I, \quad \rho(Y) = -\frac{1}{2}q.$$

Let's consider the module structure for $n = 1$ (I'll write some exercises for the $n > 1$ case). In this case, P is clearly a direct sum of P_{odd} , the span of z, z^3, z^5, \dots , and P_{even} , the span of $1, z^2, z^4, \dots$.

Then P_{odd} is a highest weight representation with the highest weight being spanned by z . We have

$$Hz = -Ez - \frac{1}{2}z = -\frac{3}{2}z.$$

So we see that $P_{\text{odd}} \cong M(-\frac{3}{2})$, the Verma module with highest weight $-\frac{3}{2}$. Similarly, $P_{\text{even}} \cong M(-\frac{1}{2})$. In this case, $\mathbf{O}_1(\mathbf{C}) \cong \{1, -1\}$ and -1 acts on P_{odd} by -1 and on P_{even} by 1 .

In general, $\mathbf{O}_n(\mathbf{C})$ has more complicated representations, which helps to organize P (see exercises).

4. SEMISIMPLE LIE ALGEBRAS

This will not be a systematic treatment of the theory since that could be its own course. I'll just summarize important points; you can read [H2] for a self-contained development.

4.1. Basic terminology. For this section we're just working with finite-dimensional Lie algebras over the complex numbers.

Let \mathfrak{g} be such a Lie algebra. We make some definitions.

First, \mathfrak{g} is **simple** if it is non-abelian and contains no nonzero proper ideals.

Recall that we have the adjoint representation

$$\text{ad}: \mathfrak{g} \rightarrow \mathfrak{gl}(\mathfrak{g})$$

where $(\text{ad } x)(y) = [x, y]$. The **Killing form** κ is the symmetric bilinear form

$$\begin{aligned} \kappa: \mathfrak{g} \times \mathfrak{g} &\rightarrow \mathbf{C} \\ \kappa(x, y) &= \text{trace}((\text{ad } x)(\text{ad } y)). \end{aligned}$$

For all $x, y, z \in \mathfrak{g}$, this satisfies

$$\kappa([x, y], z) = -\kappa(y, [x, z]).$$

This can be proven using the fact that the trace of a product is invariant under cyclic permutations (i.e., $\text{trace}(ABC) = \text{trace}(BCA)$). If we identify bilinear forms with elements of $\mathfrak{g}^* \otimes \mathfrak{g}^*$, the identity is equivalent to saying that κ has a trivial action of \mathfrak{g} .

Example 4.1.1. Consider \mathfrak{sl}_2 , where we have an eigenspace decomposition in terms of H . Our choice for an eigenbasis has been $\{X, H, Y\}$. The above identity says that $\kappa(v, w) = 0$ unless the eigenvectors for v and w add to 0. In terms of this basis, here are the matrices for the adjoint representation:

$$\text{ad}(X) = \begin{bmatrix} 0 & -2 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}, \quad \text{ad}(H) = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -2 \end{bmatrix}, \quad \text{ad}(Y) = \begin{bmatrix} 0 & 0 & 0 \\ -1 & 0 & 0 \\ 0 & 2 & 0 \end{bmatrix}.$$

So:

$$\kappa(X, Y) = \kappa(Y, X) = 4, \quad \kappa(H, H) = 8$$

and all other pairings are 0 by the above observation on eigenvalues. \square

Theorem 4.1.2. *The following properties are equivalent:*

- (1) \mathfrak{g} is isomorphic to a direct product of simple Lie algebras.
- (2) The Killing form κ is nondegenerate.
- (3) \mathfrak{g} does not contain any nonzero abelian ideals.

A Lie algebra is **semisimple** if it satisfies any (equivalently, all) of the above conditions.

Their finite-dimensional representations satisfy “complete reducibility” in the following sense.

Theorem 4.1.3 (Weyl). *If V is a finite-dimensional representation of a semisimple Lie algebra, then V is isomorphic to a direct sum of irreducible representations.*

Notably, this will not apply to infinite-dimensional representations in general, which will be relevant for us later.

Example 4.1.4. You can already see this with \mathfrak{sl}_2 : check for yourself that when $\alpha \in \mathbf{Z}_{\geq 0}$, $M(-\alpha - 2)$ is the only irreducible subrepresentation of $M(\alpha)$, and hence $M(\alpha)$ is not a direct sum of irreducible representations. \square

Fortunately, we can classify the semisimple Lie algebras. The first condition tells us that it’s enough to classify simple Lie algebras, and we’ll spend some time describing them.

4.2. Classical series. Let’s start by describing the “classical series”.

Example 4.2.1. Let $n \geq 2$ be an integer. Then \mathfrak{sl}_n is a simple Lie algebra. Note that \mathfrak{gl}_n is not: the span of the identity matrix gives a nonzero abelian ideal. \square

Let V be an n -dimensional vector space equipped with a bilinear form β . We say that $x \in \mathfrak{gl}(V)$ stabilizes β if, for all $v, w \in V$, we have

$$\beta(x(v), w) + \beta(v, x(w)) = 0.$$

We can show that the subspace of x that preserve β (also called the **stabilizer** of β) is a Lie subalgebra of $\mathfrak{gl}(V)$. As before, if we identify β as an element of $V^* \otimes V^*$, this means that $x\beta = 0$.

To state this concretely, pick a basis v_1, \dots, v_n for V . The **Gram matrix** of β is the $n \times n$ matrix

$$M_\beta = \begin{bmatrix} \beta(v_1, v_1) & \beta(v_1, v_2) & \cdots & \beta(v_1, v_n) \\ \vdots & & & \\ \beta(v_n, v_1) & \beta(v_n, v_2) & \cdots & \beta(v_n, v_n) \end{bmatrix}$$

Then we have

$$\beta(v, w) = v^T M_\beta w$$

where in the right side, v and w are written as column vectors with respect to the basis v_1, \dots, v_n . The condition that β be nondegenerate is just equivalent to $\det M_\beta \neq 0$. Then x stabilizes β if and only if

$$x^T M = -Mx.$$

Example 4.2.2. If β is symmetric and nondegenerate, then its stabilizer is called a **special orthogonal Lie algebra**, denoted $\mathfrak{so}(V, \beta)$. Since \mathbf{C} is algebraically closed, any two nondegenerate symmetric bilinear forms β, β' differ by a change of basis, which means that $\mathfrak{so}(V, \beta) \cong \mathfrak{so}(V, \beta')$, so we usually just write $\mathfrak{so}(V)$, or just \mathfrak{so}_n (but the different choices of forms can be more or less convenient depending on what calculations we want to do).

If we diagonalize the form, i.e., work with an orthonormal basis e_1, \dots, e_n , then the Gram matrix is just the identity, so we can identify \mathfrak{so}_n with the subspace of skew-symmetric matrices (you can check directly that this is a Lie subalgebra).

I'll leave it to you to check that if $n \geq 3$, then \mathfrak{so}_n is semisimple (and is simple as long as $n \neq 4$). As a bonus, you can try checking that $\mathfrak{so}_3 \cong \mathfrak{sl}_2$, $\mathfrak{so}_4 \cong \mathfrak{sl}_2 \times \mathfrak{sl}_2$, and $\mathfrak{so}_6 \cong \mathfrak{sl}_4$. \square

Example 4.2.3. If β is skew-symmetric and nondegenerate (i.e., β is a symplectic form), then its stabilizer is called a **symplectic Lie algebra**, denoted $\mathfrak{sp}(V, \beta)$. Again, any two symplectic forms differ by a change of basis, so we usually just write $\mathfrak{sp}(V)$ or \mathfrak{sp}_n . Note that $n = 2m$ must be **even** for a symplectic form to exist.

There isn't really an "easy" choice of basis to identify what \mathfrak{sp} is like with the previous example. One popular choice is to pick a "symplectic basis" e_1, \dots, e_{2m} , i.e., one so that M_β is a block direct sum of m copies of $\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$. This is convenient for some purposes, and we'll see a different choice in the next section.

If $m \geq 1$, then \mathfrak{sp}_{2m} is simple. Furthermore, $\mathfrak{sp}_2 \cong \mathfrak{sl}_2$ and $\mathfrak{sp}_4 \cong \mathfrak{so}_5$. \square

That's almost all of them, actually. There are 5 more simple Lie algebras which are not one of the above 3 examples, called the "exceptional" Lie algebras, and that exhausts the entire list. They take a bit of effort to describe precisely and we won't focus much on it, so I'm not going to go into detail.

4.3. Roots and weights. Now let \mathfrak{g} be a semisimple Lie algebra. An element $x \in \mathfrak{g}$ is **semisimple** if $\text{ad } x \in \mathfrak{gl}(\mathfrak{g})$ is diagonalizable. A subalgebra of \mathfrak{g} is **toral** if all of its elements are semisimple. All toral subalgebras are abelian.

A maximal toral subalgebra is called a **Cartan subalgebra**, and is usually denoted by \mathfrak{h} .

Theorem 4.3.1. *Any two Cartan subalgebras $\mathfrak{h}, \mathfrak{h}'$ of \mathfrak{g} are conjugate: there exists an automorphism φ of \mathfrak{g} such that $\varphi(\mathfrak{h}) = \mathfrak{h}'$.*

In particular, all Cartan subalgebras are isomorphic. Since they are abelian, the only important information is their dimension, and $\dim(\mathfrak{h})$ is called the **rank** of \mathfrak{g} .

Next, recall that if x, y are diagonalizable endomorphisms of the same (finite-dimensional) vector space and they commute, then they can be simultaneously diagonalized: there exists a basis e_1, \dots, e_n so that each e_i is an eigenvector for both x and y . (If this doesn't sound familiar, prove it as an exercise.) More generally, this is true for any collection of diagonalizable endomorphisms which all pairwise commute.

We can apply this to a Cartan subalgebra \mathfrak{h} . For each $x, y \in \mathfrak{h}$, we know that $\text{ad } x$ and $\text{ad } y$ are diagonalizable (by definition) and commute with each other. That means there exists a basis e_1, \dots, e_n for \mathfrak{g} such that each e_i is an eigenvector for $\text{ad } x$ for all $x \in \mathfrak{h}$. In particular, for each i , there is a function $\lambda_i: \mathfrak{h} \rightarrow \mathbf{C}$ such that $(\text{ad } x)e_i = \lambda_i(x)e_i$. Since ad is linear, so is λ_i , so $\lambda_i \in \mathfrak{h}^*$. We call these functions λ_i **roots** when they are nonzero, and denote the set of roots by Φ .

First, let's describe them for the classical examples.

Example 4.3.2. Let $\mathfrak{g} = \mathfrak{sl}_n$. Then the space of diagonal matrices is a Cartan subalgebra \mathfrak{h} (we aren't going to prove all of the properties) so $\text{rank}(\mathfrak{sl}_n) = n - 1$.

For $i = 1, \dots, n$, let $\varepsilon_i: \mathfrak{h} \rightarrow \mathbf{C}$ be the function that returns the i th diagonal entry. Then we can identify \mathfrak{h}^* with the hyperplane of \mathbf{C}^n defined by $\varepsilon_1 + \dots + \varepsilon_n = 0$. The roots are $\varepsilon_i - \varepsilon_j$ ranging over all choices of $i \neq j$: the corresponding eigenvector is the matrix unit $E_{i,j}$ which is 1 in row i and column j and 0 elsewhere. \square

Example 4.3.3. For \mathfrak{so}_N , it will help to distinguish between the cases when N is even and odd. First, let's consider the even case $N = 2n$ and use a different symmetric form. All $2n \times 2n$ matrices will be written in block form, with all blocks of size $n \times n$. We choose β so that

$$M_\beta = \begin{bmatrix} 0 & I_n \\ I_n & 0 \end{bmatrix}.$$

Then $\begin{bmatrix} A & B \\ C & D \end{bmatrix}$ belongs to \mathfrak{so}_{2n} if and only if

$$B = -B^T, \quad C = -C^T, \quad D = -A^T.$$

For a Cartan subalgebra \mathfrak{h} , we can take the matrices where $B = C = 0$ and A is diagonal (and $D = -A^T$), i.e., diagonal matrices whose entries are of the form $x_1, \dots, x_n, -x_1, \dots, -x_n$. Then $\text{rank}(\mathfrak{so}_{2n}) = n$. For $i = 1, \dots, n$, let $\varepsilon_i: \mathfrak{h} \rightarrow \mathbf{C}$ be the function that takes the i th diagonal entry.

Again, letting $E_{i,j}$ denote the matrix unit with a 1 in row i and column j and 0's elsewhere, we have eigenvectors (corresponding to nonzero roots) as follows:

- $E_{i,j} - E_{n+j,n+i}$ for $1 \leq i, j \leq n$ and $i \neq j$ with root $\varepsilon_i - \varepsilon_j$, and
- $E_{i,n+j} - E_{j,n+i}$ for $1 \leq i, j \leq n$ and $i \neq j$ with root $\varepsilon_i + \varepsilon_j$, and
- $E_{n+i,j} - E_{n+j,i}$ for $1 \leq i, j \leq n$ and $i \neq j$ with root $-\varepsilon_i - \varepsilon_j$. \square

Example 4.3.4. For \mathfrak{sp}_{2n} , we can do something similar. All $2n \times 2n$ matrices will be written in block form, with all blocks of size $n \times n$. We choose β so that

$$M_\beta = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix}.$$

Then $\begin{bmatrix} A & B \\ C & D \end{bmatrix}$ belongs to \mathfrak{sp}_{2n} if and only if

$$B = B^T, \quad C = C^T, \quad D = -A^T.$$

For a Cartan subalgebra \mathfrak{h} , we can take the matrices where $B = C = 0$ and A is diagonal (and $D = -A^T$), i.e., diagonal matrices whose entries are of the form $x_1, \dots, x_n, -x_1, \dots, -x_n$. Then $\text{rank}(\mathfrak{sp}_{2n}) = n$. For $i = 1, \dots, n$, let $\varepsilon_i: \mathfrak{h} \rightarrow \mathbf{C}$ be the function that takes the i th diagonal entry.

Again, letting $E_{i,j}$ denote the matrix unit with a 1 in row i and column j and 0's elsewhere, we have eigenvectors (corresponding to nonzero roots) as follows:

- $E_{i,j} - E_{n+j,n+i}$ for $1 \leq i, j \leq n$ and $i \neq j$ with root $\varepsilon_i - \varepsilon_j$, and
- $E_{i,n+j} + E_{j,n+i}$ for $1 \leq i, j \leq n$ with root $\varepsilon_i + \varepsilon_j$, and
- $E_{n+i,j} + E_{n+j,i}$ for $1 \leq i, j \leq n$ with root $-\varepsilon_i - \varepsilon_j$.

The main difference between the previous case is that we allow $i = j$ in the last two cases. \square

Example 4.3.5. Finally, we discuss \mathfrak{so}_{2n+1} . All $(2n+1) \times (2n+1)$ matrices will be written in 3×3 block form, with the blocks of rows/columns having sizes $n, n, 1$. We choose β so that

$$M_\beta = \begin{bmatrix} 0 & I_n & 0 \\ I_n & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Then $\begin{bmatrix} A & B & E \\ C & D & F \\ G & H & J \end{bmatrix}$ belongs to \mathfrak{so}_{2n+1} if and only if

$$B = -B^T, \quad C = -C^T, \quad D = -A^T, \quad E = -H^T, \quad F = -G^T, \quad J = 0.$$

For a Cartan subalgebra \mathfrak{h} , we can take the matrices where $B = C = E = F = G = H = J = 0$ and A is diagonal (and $D = -A^T$), i.e., diagonal matrices whose entries are of the form $x_1, \dots, x_n, -x_1, \dots, -x_n, 0$. Then $\text{rank}(\mathfrak{so}_{2n+1}) = n$. For $i = 1, \dots, n$, let $\varepsilon_i: \mathfrak{h} \rightarrow \mathbf{C}$ be the function that takes the i th diagonal entry.

Again, letting $E_{i,j}$ denote the matrix unit with a 1 in row i and column j and 0's elsewhere, we have eigenvectors (corresponding to nonzero roots) as follows:

- $E_{i,j} - E_{n+j,n+i}$ for $1 \leq i, j \leq n$ and $i \neq j$ with root $\varepsilon_i - \varepsilon_j$, and
- $E_{i,n+j} - E_{j,n+i}$ for $1 \leq i, j \leq n$ and $i \neq j$ with root $\varepsilon_i + \varepsilon_j$, and
- $E_{n+i,j} - E_{n+j,i}$ for $1 \leq i, j \leq n$ and $i \neq j$ with root $-\varepsilon_i - \varepsilon_j$, and
- $E_{i,2n+1} - E_{2n+1,n+i}$ for $1 \leq i \leq n$ with root ε_i , and
- $E_{n+i,2n+1} - E_{2n+1,i}$ for $1 \leq i \leq n$ with root $-\varepsilon_i$. \square

We can actually do the same thing for any finite-dimensional representation $\varphi: \mathfrak{g} \rightarrow \mathfrak{gl}(V)$. It turns out that if $\text{ad } x$ is diagonalizable, then so is $\varphi(x)$ (we won't explain why, but it is part of the theory of abstract Jordan decomposition). Then the same thing applies: we can find a basis of V consisting of eigenvectors. This leads to a decomposition

$$V = \bigoplus_{\lambda \in \mathfrak{h}^*} V_\lambda$$

where

$$V_\lambda = \{v \in V \mid \varphi(x)(v) = \lambda(x)v\}.$$

The λ such that $V_\lambda \neq 0$ are called the **weights** of V . This generalizes the notion of roots, except that we don't consider 0 a root as a matter of convention (though it is a weight).

One special property about the adjoint representation is that for each root α , we have $\dim \mathfrak{g}_\alpha = 1$.

4.4. **Root systems.** Here's a few things we need to know about the set of roots Φ in a semisimple Lie algebra \mathfrak{g} :

- Φ spans \mathfrak{h}^* .
- If we pick a basis for \mathfrak{h}^* consisting of elements from Φ , then every other root can be written as a *rational* linear combination of these basis elements. Let $E_{\mathbf{Q}}$ be the \mathbf{Q} -vector space spanned by Φ and let $E = E_{\mathbf{Q}} \otimes_{\mathbf{Q}} \mathbf{R}$.
- The restriction of the Killing form κ to \mathfrak{h} is nondegenerate. In particular, it gives an isomorphism $\mathfrak{h} \cong \mathfrak{h}^*$, so we can also transfer the form to \mathfrak{h}^* . Furthermore, the restriction to E is positive-definite.

So, in particular, given a semisimple Lie algebra and a choice of Cartan subalgebra, we get the following data: a finite-dimensional real inner product space E and a finite collection of vectors $\Phi \subset E$. They satisfy some more conditions, and for the sake of completeness, let's axiomatize it.

First, let's define some notation. We let $(,)$ denote the inner product on E . Given $\alpha, \beta \in E$, define

$$\langle \beta, \alpha \rangle = \frac{2(\beta, \alpha)}{(\alpha, \alpha)}.$$

This follows the notation of [H2]. Warning: in [H3], this is instead denoted by $\langle \beta, \alpha^\vee \rangle$.

Given $\alpha \in E$, define the reflection $s_\alpha \in \mathbf{GL}(E)$ by

$$s_\alpha(v) := v - \langle v, \alpha \rangle \alpha = v - \frac{2(v, \alpha)}{(\alpha, \alpha)} \alpha.$$

Definition 4.4.1. Let E be a finite-dimensional real inner product space. A finite collection of nonzero vectors $\Phi \subset E$ is called a **root system** if it satisfies the following conditions:

- (1) Φ spans E .
- (2) Let $\alpha \in \Phi$. Then $-\alpha \in \Phi$, and if $c\alpha \in \Phi$ for some scalar c , then $c = \pm 1$.
- (3) For all $\alpha \in \Phi$, s_α preserves Φ , i.e., if $v \in \Phi$, then $s_\alpha(v) \in \Phi$.
- (4) For all $\alpha, \beta \in \Phi$, we have $\langle \beta, \alpha \rangle \in \mathbf{Z}$. □

Definition 4.4.2. The **root lattice** Λ_r is the \mathbf{Z} -span of Φ . The **weight lattice** is

$$\Lambda = \{ \beta \in \mathfrak{h}^* \mid \langle \beta, \alpha \rangle \in \mathbf{Z} \text{ for all } \alpha \in \Phi \}. \quad \square$$

Definition 4.4.3. A subset $\Delta \subset \Phi$ is called a **base** if:

- (1) Δ is a basis for E , and
- (2) If $\beta \in \Phi$ and $\beta = \sum_{\alpha \in \Delta} c_\alpha \alpha$, then either $\{c_\alpha\} \subset \mathbf{Z}_{\geq 0}$ or $\{-c_\alpha\} \subset \mathbf{Z}_{\geq 0}$. In other words, every root can be written as an integer linear combination of elements of Δ , and all nonzero coefficients have the same sign.

If a base Δ is fixed, then its elements are called **simple roots**. Roots which are non-negative (respectively, non-positive) linear combinations of Δ are called **positive roots** (respectively, **negative roots**), and the set is denoted Φ^+ , (respectively, Φ^-).

Remember that this is all depends on the choice of Δ , and is not intrinsic to Φ . □

The subgroup of $\mathbf{GL}(E)$ generated by the elements s_α for $\alpha \in \Phi$ is called the **Weyl group** of Φ , and usually denoted by W .

Proposition 4.4.4. *W is a finite group.*

Some facts: First, bases do exist, and W acts transitively on the set of all bases. Second, given any base Δ , W is generated by s_α for $\alpha \in \Delta$. The s_α are called **simple reflections**. In particular, every $w \in W$ is a product of simple reflections, and the minimal length of such a product is called the **length** of w , and denoted $\ell(w)$. Alternatively, $\ell(w)$ is the number of positive roots that become negative after applying w :

$$\ell(w) = |w(\Phi^+) \cap \Phi^-|.$$

Here's an important special case of this equivalence:

Proposition 4.4.5. *If $\alpha \in \Delta$, then s_α sends $\Phi^+ \setminus \{\alpha\}$ to itself and $s_\alpha(\alpha) = -\alpha$.*

First, let's examine the root systems for the classical series.

Example 4.4.6. For \mathfrak{sl}_n , we can take

$$E = \{(x_1, \dots, x_n) \in \mathbf{R}^n \mid x_1 + \dots + x_n = 0\}$$

with the usual dot product. Letting $\varepsilon_i \in \mathbf{R}^n$ denote the i th standard basis vector, we have $\Phi = \{\varepsilon_i - \varepsilon_j \mid i \neq j\}$. For a base, we can take

$$\Delta = \{\varepsilon_i - \varepsilon_{i+1} \mid i = 1, \dots, n-1\}.$$

Then $\varepsilon_i - \varepsilon_j$ is positive if and only if $i < j$, and negative otherwise.

The reflection $s_{\varepsilon_i - \varepsilon_j}$ is the transposition that swaps the i th and j th coordinates, so we see that W is the n th symmetric group. Furthermore, if $i < j$, then $w(\varepsilon_i - \varepsilon_j)$ is negative if and only if $w(i) > w(j)$. So $\ell(w)$ is the number of inversions of w : by definition, the number of pairs $i < j$ such that $w(i) > w(j)$.

For reference, this is the **type A root system**, and is specifically denoted A_{n-1} (the subscript is for the rank).

The root lattice is just the integer vectors:

$$\Lambda_r = \{(x_1, \dots, x_n) \in \mathbf{Z}^n \mid x_1 + \dots + x_n = 0\}.$$

The weight lattice consists of vectors whose sum is 0 and such that the difference between any two entries is an integer. This forces the entries to belong to $\frac{1}{n}\mathbf{Z}$ and Λ_r is a subgroup of index n inside Λ .

For instance, for \mathfrak{sl}_2 , Λ is the \mathbf{Z} -multiples of $(\frac{1}{2}, -\frac{1}{2})$. To translate to our previous notation, the weight $(x, -x)$ corresponds to the eigenvalue $2x$. \square

Example 4.4.7. For \mathfrak{sp}_{2n} , we can take $E = \mathbf{R}^n$ (thinking of (x_1, \dots, x_n) as the diagonal matrix of \mathfrak{sp}_{2n} with entries $x_1, \dots, x_n, -x_1, \dots, -x_n$) with the usual dot product. Again, let $\varepsilon_i \in \mathbf{R}^n$ be the i th standard basis vector. Then

$$\Phi = \{\varepsilon_i - \varepsilon_j \mid i \neq j\} \cup \{\pm(\varepsilon_i + \varepsilon_j) \mid i \leq j\}.$$

For a base, we can take

$$\Delta = \{\varepsilon_i - \varepsilon_{i+1} \mid i = 1, \dots, n-1\} \cup \{2\varepsilon_n\}.$$

Then

$$\Phi^+ = \{\varepsilon_i - \varepsilon_j \mid i < j\} \cup \{\varepsilon_i + \varepsilon_j \mid i \leq j\}.$$

As before, $s_{\varepsilon_i - \varepsilon_{i+1}}$ swaps the i th and $(i+1)$ st coordinates, while $s_{2\varepsilon_n}$ negates the n th coordinate. So the corresponding Weyl group W is the group of $n \times n$ signed permutation matrices, i.e., $n \times n$ matrices where each row and column has exactly 1 nonzero entry, and that entry is ± 1 . This is also known as the **hyperoctahedral group**. The length function is kind of

messy to describe, so we'll omit it. There is a discussion in [BB, §8.1], but be warned that a change of coordinates is needed.

This is the **type C root system**, denoted C_n .

The root lattice is

$$\Lambda_r = \{(x_1, \dots, x_n) \in \mathbf{Z}^n \mid x_1 + \dots + x_n \in 2\mathbf{Z}\}$$

and the weight lattice is $\Lambda = \mathbf{Z}^n$. In particular, Λ_r is an index 2 subgroup of Λ . \square

Example 4.4.8. For \mathfrak{so}_{2n+1} , we can take $E = \mathbf{R}^n$ (thinking of (x_1, \dots, x_n) as the diagonal matrix of \mathfrak{sp}_{2n} with entries $x_1, \dots, x_n, -x_1, \dots, -x_n, 0$) with the usual dot product. Again, let $\varepsilon_i \in \mathbf{R}^n$ be the i th standard basis vector. Then

$$\Phi = \{\varepsilon_i - \varepsilon_j \mid i \neq j\} \cup \{\pm(\varepsilon_i + \varepsilon_j) \mid i < j\} \cup \{\pm\varepsilon_i\}.$$

For a base, we can take

$$\Delta = \{\varepsilon_i - \varepsilon_{i+1} \mid i = 1, \dots, n-1\} \cup \{\varepsilon_n\}.$$

Then

$$\Phi^+ = \{\varepsilon_i - \varepsilon_j \mid i < j\} \cup \{\varepsilon_i + \varepsilon_j \mid i < j\} \cup \{\varepsilon_i\}.$$

As before, $s_{\varepsilon_i - \varepsilon_{i+1}}$ swaps the i th and $(i+1)$ st coordinates, while s_{ε_n} negates the n th coordinate. So the corresponding Weyl group W is again the hyperoctahedral group.

This is the **type B root system**, denoted B_n .

The root lattice is $\Lambda_r = \mathbf{Z}^n$ and the weight lattice is

$$\Lambda = \mathbf{Z}^n \cup \left(\mathbf{Z}^n + \left(\frac{1}{2}, \dots, \frac{1}{2} \right) \right).$$

In particular, Λ_r is an index 2 subgroup of Λ . \square

Example 4.4.9. For \mathfrak{so}_{2n} , we can take $E = \mathbf{R}^n$ (thinking of (x_1, \dots, x_n) as the diagonal matrix of \mathfrak{sp}_{2n} with entries $x_1, \dots, x_n, -x_1, \dots, -x_n$) with the usual dot product. Again, let $\varepsilon_i \in \mathbf{R}^n$ be the i th standard basis vector. Then

$$\Phi = \{\varepsilon_i - \varepsilon_j \mid i \neq j\} \cup \{\pm(\varepsilon_i + \varepsilon_j) \mid i < j\}.$$

(it looks like the symplectic case, but we don't allow $i = j$ in the second set). For a base, we can take

$$\Delta = \{\varepsilon_i - \varepsilon_{i+1} \mid i = 1, \dots, n-1\} \cup \{\varepsilon_{n-1} + \varepsilon_n\}.$$

Then

$$\Phi^+ = \{\varepsilon_i - \varepsilon_j \mid i < j\} \cup \{\varepsilon_i + \varepsilon_j \mid i < j\}.$$

As before, $s_{\varepsilon_i - \varepsilon_{i+1}}$ swaps the i th and $(i+1)$ st coordinates, while $s_{\varepsilon_{n-1} + \varepsilon_n}$ swaps the $(n-1)$ st and n th coordinates while also negating them. So the corresponding Weyl group W is a subgroup of $n \times n$ signed permutation matrices consisting of matrices where -1 appears an *even* number of times. This is an index 2 subgroup of the hyperoctahedral group. Again, we'll skip making the length function explicit.

This is the **type D root system**, denoted D_n .

The root lattice is

$$\Lambda_r = \{(x_1, \dots, x_n) \in \mathbf{Z}^n \mid x_1 + \dots + x_n \in 2\mathbf{Z}\}$$

and the weight lattice is

$$\Lambda = \mathbf{Z}^n \cup \left(\mathbf{Z}^n + \left(\frac{1}{2}, \dots, \frac{1}{2} \right) \right).$$

In particular, Λ_r is an index 4 subgroup of Λ . \square

Remark 4.4.10. There is an asymmetry here: the orthogonal cases get 2 cases depending on parity, while the symplectic case only gets 1 case. Technically, there is an additional “odd symplectic” case. The way to do this consistently is to enlarge our setting to allow Lie *superalgebras* and the corresponding root system ends up being the union $B_n \cup C_n$, sometimes called BC_n . This, however, violates the second axiom for root systems (it is sometimes referred to as a *non-reduced root system*). \square

4.5. Formal characters. A representation V of \mathfrak{g} is a **weight representation** if it decomposes into weight spaces for \mathfrak{h} (i.e., it is a semisimple representation of \mathfrak{h}):

$$V = \bigoplus_{\alpha \in \mathfrak{h}^*} V_\alpha.$$

This is automatic if $\dim V < \infty$. The elements of each V_α are called **weight vectors**.

The representations we deal with will satisfy $\dim V_\alpha < \infty$ for all α . In that case, it is convenient to record these dimensions as a **formal character** which can be thought of in some different ways.

First, we have a function $\text{ch}_V: \mathfrak{h}^* \rightarrow \mathbf{Z}_{\geq 0}$ defined by $\text{ch}_V(\alpha) = \dim V_\alpha$. These can be added in the obvious way. Given two functions $f, g: \mathfrak{h}^* \rightarrow \mathbf{Z}$, we can attempt to multiply them using a convolution product:

$$(f * g)(\alpha) = \sum_{\beta \in \mathfrak{h}^*} f(\beta)g(\alpha - \beta).$$

But this is generally not well-defined since there is no guarantee that the sum is finite.

To deal with this issue, let’s be more precise. Fix a base Δ for our root system Φ . Let Γ denote the set of non-negative linear combinations of elements in Δ . Given a function $f: \mathfrak{h}^* \rightarrow \mathbf{Z}$, define its support to be $\text{Supp}(f) = \{\alpha \in \mathfrak{h}^* \mid f(\alpha) \neq 0\}$. We define

$$\mathcal{X} = \{f: \mathfrak{h}^* \rightarrow \mathbf{Z} \mid \text{there exists } \lambda_1, \dots, \lambda_N \text{ such that } \text{Supp}(f) \subseteq \bigcup_{i=1}^N (\lambda_i - \Gamma)\}.$$

I’ll leave it as an exercise to check that if $f, g \in \mathcal{X}$, then $f * g$ is well-defined, so \mathcal{X} is a commutative ring. Note that ch_V need not belong to \mathcal{X} ; we will later restrict to studying representations where this holds.

A slightly different perspective on this: given $\alpha \in \mathfrak{h}^*$, introduce a formal symbol e^α , so that we can instead interpret the formal character as the sum

$$\sum_{\alpha \in \mathfrak{h}^*} (\dim V_\alpha) e^\alpha.$$

Instead of considering the convolution product, we instead use the rules $e^\alpha e^\beta = e^{\alpha+\beta}$ and “distribute” in the obvious way. Of course, the product is not always well-defined, so we need to restrict to linear sums corresponding to elements in \mathcal{X} .

The second notation can be helpful since we can treat them like formal power series. For instance, we can write down identities like

$$\frac{1}{1 - e^\alpha} = \sum_{n \geq 0} e^{n\alpha}$$

since $(1 - e^\alpha) \sum_{n \geq 0} e^{n\alpha} = e^0$, and e^0 is the multiplicative identity.

In any case, given two representations M, N , we have (whenever they are defined)

$$\text{ch}_M + \text{ch}_N = \text{ch}_{M \oplus N}, \quad \text{ch}_M * \text{ch}_N = \text{ch}_{M \otimes N}.$$

4.6. Borel subalgebras. Use the notation from before: \mathfrak{g} is a semisimple Lie algebra, \mathfrak{h} is a Cartan subalgebra, and Δ is a choice of simple roots for the corresponding root system Φ . The **Borel subalgebra** \mathfrak{b} associated to this data is the span of \mathfrak{h} together with the root spaces of its positive roots:

$$\mathfrak{b} = \mathfrak{h} \oplus \bigoplus_{\alpha \in \Phi^+} \mathfrak{g}_\alpha.$$

There is also an **opposite Borel subalgebra** \mathfrak{b}^- defined using negative roots:

$$\mathfrak{b}^- = \mathfrak{h} \oplus \bigoplus_{\alpha \in \Phi^-} \mathfrak{g}_\alpha.$$

Remark 4.6.1. This is not the usual way to define Borel subalgebras, so let's remark on that. Given a Lie algebra \mathfrak{a} , define $\mathfrak{a}^{(0)} = \mathfrak{a}$ and $\mathfrak{a}^{(i+1)} = [\mathfrak{a}^{(i)}, \mathfrak{a}^{(i)}]$ (the span of all brackets of elements in $\mathfrak{a}^{(i)}$); call \mathfrak{a} **solvable** if $\mathfrak{a}^{(N)} = 0$ for some N . Without reference to Cartan subalgebras, Borel subalgebras can equivalently be defined to be the maximal solvable Lie subalgebras of \mathfrak{g} . (There's a lot of work going into showing that these definitions are equivalent.) \square

In this case, we will write $\mathfrak{n} = [\mathfrak{b}, \mathfrak{b}]$. This is called the **nilpotent radical** of \mathfrak{b} . We note that \mathfrak{n} is an ideal of \mathfrak{b} and $\mathfrak{h} \cong \mathfrak{b}/\mathfrak{n}$ (prove it as an exercise). In particular, we have a quotient map $\mathfrak{b} \rightarrow \mathfrak{h}$, and we can choose this so the composition $\mathfrak{h} \rightarrow \mathfrak{b} \rightarrow \mathfrak{h}$ is the identity. We also define $\mathfrak{n}^- = [\mathfrak{b}^-, \mathfrak{b}^-]$.

4.7. Highest weight representations. Let V be a nonzero weight representation of \mathfrak{g} . A weight vector $v \in V$ is a **highest weight vector**¹ if $xv = 0$ for all $x \in \mathfrak{n}$. Finally, V is a **highest weight representation** if it is generated by a highest weight vector, i.e., there is a highest weight vector v^+ and the smallest subrepresentation containing v^+ is V itself. In that case, any 2 highest weight vectors are scalar multiples of each other, so if $v^+ \in V_\lambda$, then we call λ the highest weight of V .

Given $\lambda \in \mathfrak{h}^*$, we get a 1-dimensional representation \mathbf{C}_λ of \mathfrak{h} . We can restrict along the quotient map $\mathfrak{b} \rightarrow \mathfrak{h}$ to get a corresponding 1-dimensional representation of \mathfrak{b} , which we also denote \mathbf{C}_λ . The associated **Verma module** is

$$M(\lambda) = \text{Ind}_{\mathfrak{b}}^{\mathfrak{g}} \mathbf{C}_\lambda.$$

As a vector space (or representation of \mathfrak{h}), this is $U(\mathfrak{n}^-) \otimes_{\mathbf{C}} \mathbf{C}_\lambda$, so we have the following result.

Proposition 4.7.1. *We have*

$$\text{ch}_{M(\lambda)} = \frac{e^\lambda}{\prod_{\alpha \in \Phi^-} (1 - e^\alpha)}.$$

In particular, $\text{ch}_{M(\lambda)} \in \mathcal{X}$.

¹Humphreys uses the terminology “maximal vector”.

Proof. The main point is that by the PBW theorem, we have $U(\mathfrak{n}^-) \cong \text{Sym}(\mathfrak{n}^-)$ as an \mathfrak{h} -representation, where Sym denotes symmetric algebra. Breaking this down further, we have $\mathfrak{n}^- \cong \bigoplus_{\alpha \in \Phi^-} \mathbf{C}_\alpha$ and hence

$$\text{Sym}(\mathfrak{n}^-) \cong \bigotimes_{\alpha \in \Phi^-} \text{Sym}(\mathbf{C}_\alpha)$$

and

$$\text{ch}_{\text{Sym}(\mathbf{C}_\alpha)} = \sum_{n \geq 0} e^{n\alpha} = \frac{1}{1 - e^\alpha}.$$

For the second statement, the formula implies that the support of $\text{ch}_{M(\lambda)}$ is $\lambda - \Gamma$. \square

Then $M(\lambda)$ is a highest weight representation of weight λ and $\dim M(\lambda)_\lambda = 1$. Pick a highest weight vector $v \in M(\lambda)$. Then the space of homomorphisms $\varphi: M(\lambda) \rightarrow V$ is isomorphic to the space of highest weight vectors of V of weight λ , namely $\varphi(v)$ is the corresponding highest weight vector of V (prove this; use hom-tensor adjunction).

Finally, if N and N' are both subrepresentations of $M(\lambda)$ which do not contain $M(\lambda)_\lambda$, then neither does their sum $N + N'$. Hence there is a maximal proper submodule $N(\lambda)$ of $M(\lambda)$; denote the quotient by $L(\lambda)$. Then:

- $L(\lambda)$ is also a highest weight representation of highest weight λ and $L(\lambda)$ is irreducible.
- Any other irreducible highest weight representation of highest weight λ is isomorphic to $L(\lambda)$.

I'll leave these as exercises.

Just to make it easy to see, here is how they are related in a short exact sequence:

$$0 \rightarrow N(\lambda) \rightarrow M(\lambda) \rightarrow L(\lambda) \rightarrow 0.$$

So now we're caught up with terminology from the \mathfrak{sl}_2 case.

4.8. Definition of category \mathcal{O} . We now fix the usual notation (I likely won't repeat myself, but this notation will generally be reserved and fixed)

- \mathfrak{g} is a semisimple complex Lie algebra,
- $\mathfrak{h} \subset \mathfrak{g}$ is a Cartan subalgebra,
- $\Delta \subset \Phi$ is a choice of simple roots,
- \mathfrak{b} and \mathfrak{b}^- are the associated Borel subalgebra and opposite Borel subalgebra.
- $\mathfrak{n} = [\mathfrak{b}, \mathfrak{b}]$ and $\mathfrak{n}^- = [\mathfrak{b}^-, \mathfrak{b}^-]$.
- Λ_r is the set of integer linear combinations of Δ (the root lattice).
- Γ is the subset of Λ_r consisting of non-negative integer linear combinations.
- We equip \mathfrak{h}^* with a partial ordering: $\lambda \leq \mu$ if $\mu - \lambda \in \Gamma$.

Recall that we have a vector space decomposition

$$\mathfrak{g} = \mathfrak{n}^- \oplus \mathfrak{h} \oplus \mathfrak{n}.$$

Definition 4.8.1. Category \mathcal{O} is the full subcategory of the category of (left) $U(\mathfrak{g})$ -modules consisting of objects M satisfying these conditions:

- (1) M is a finitely generated $U(\mathfrak{g})$ -module,
- (2) M is a semisimple \mathfrak{h} -representation, i.e., M is a direct sum of its weight spaces,
- (3) M is locally \mathfrak{n} -finite: for every $v \in M$, the $U(\mathfrak{n})$ -submodule generated by v is finite-dimensional. \square

Here “full subcategory” just means that the homomorphisms between modules in category \mathcal{O} are just the usual $U(\mathfrak{g})$ -module homomorphisms. Given $M, N \in \mathcal{O}$, we’ll use $\text{Hom}_{\mathcal{O}}(M, N)$ to denote this space of homomorphisms.

About (3): if v is a weight vector of weight λ , then every weight μ appearing in $U(\mathfrak{n}) \cdot v$ satisfies $\mu \geq \lambda$. Hence Verma modules satisfy (3) and so also belong to \mathcal{O} (the first 2 conditions we’ve already discussed).

It follows from our previous discussions that every finite-dimensional representation belongs to \mathcal{O} .

Remark 4.8.2. The letter “O” was chosen since it is the first letter of the Russian word for “basic”. \square

Given $M \in \mathcal{O}$, let $\Pi(M) = \{\lambda \in \mathfrak{h}^* \mid M_\lambda \neq 0\}$.

Proposition 4.8.3. *Let $M \in \mathcal{O}$. Then:*

- (1) $\dim M_\lambda < \infty$ for all $\lambda \in \mathfrak{h}^*$.
- (2) There exist finitely many $\lambda_1, \dots, \lambda_N$ such that $\Pi(M) \subseteq \bigcup_i \lambda_i - \Gamma$.

In particular, $\text{ch}_M \in \mathcal{X}$.

Proof. Since M is finitely generated and a direct sum of its weight spaces, it has a finite generating set consisting of weight vectors. To prove the proposition, it suffices to consider the case when M is generated by a single weight vector v .

Let V be the $U(\mathfrak{n})$ -submodule generated by v ; V is finite-dimensional by axiom (3). Furthermore, V is invariant under \mathfrak{h} since $U(\mathfrak{n})$ is a sum of \mathfrak{h} -eigenspaces. So by PBW, we have a surjection of \mathfrak{h} -representations $U(\mathfrak{n}^-) \otimes V \rightarrow M$. So $\Pi(M) \subseteq \Pi(U(\mathfrak{n}^-) \otimes V)$, and

$$\text{ch}_{U(\mathfrak{n}^-) \otimes V} = \frac{\text{ch}_V}{\prod_{\alpha \in \Phi^-} (1 - e^\alpha)}.$$

from which it follows that $\dim(U(\mathfrak{n}^-) \otimes V)_\lambda < \infty$ for all λ and $\Pi(U(\mathfrak{n}^-) \otimes V) = \bigcup_i \lambda_i - \Gamma$ where λ_i are all of the weights of V . \square

Proposition 4.8.4. *Let $M \in \mathcal{O}$.*

- (1) M satisfies the ascending chain condition, i.e., every chain of submodules of M must be eventually constant. Equivalently, every submodule of M is finitely generated.

This property is saying that \mathcal{O} is a noetherian category.

- (2) Every submodule and quotient module of M also belongs to \mathcal{O} .
- (3) If $M_1, \dots, M_r \in \mathcal{O}$, then $M_1 \oplus \dots \oplus M_r \in \mathcal{O}$.
- (4) M is finitely generated as a $U(\mathfrak{n}^-)$ -module.
- (5) If N is a finite-dimensional representation, then $M \otimes N \in \mathcal{O}$; tensoring with N is an exact functor on \mathcal{O} .
- (6) M has a finite filtration $0 = M^0 \subset M^1 \subset \dots \subset M^r = M$ such that M^i/M^{i-1} is a highest weight representation for all i .

Proof. (1) We’ve already proven that $U(\mathfrak{g})$ is left-noetherian.

(2) Axioms (2) and (3) are clear from the definitions; axiom (1) follows from part (1).

(3) Clear from definitions.

(4) This can be extracted from the previous proof.

(5) Let v_1, \dots, v_r be generators for M . Then the finite-dimensional subspace $\sum_{i=1}^r v_i \otimes N$ generates $M \otimes N$ by direct inspection of how tensor product representations work. Similarly, since finite-dimensional representations are automatically \mathfrak{h} -semisimple, (2) also holds.

Finally, for $\sum_i v_i \otimes w_i \in M \otimes N$, we have

$$U(\mathfrak{n}) \cdot \sum_i v_i \otimes w_i \subseteq \sum_i U(\mathfrak{n}) \cdot v_i \otimes N$$

and the latter is finite-dimensional since M satisfies axiom (3).

(6) Assume $M \neq 0$. From the previous result, $\Pi(M)$ has a maximal element λ . Any $v \in M_\lambda$ must be a highest weight vector; let M^1 be the $U(\mathfrak{g})$ -submodule that v generates. Then $M/M^1 \in \mathcal{O}$; if it is nonzero, then it also has a nonzero highest weight vector: let M^2 be the preimage in M of the submodule that it generates. Repeating this, we obtain an increasing chain of submodules, which then must stabilize after finitely many steps by (1), and it satisfies the property we claimed. \square

Remark 4.8.5. Even though $\text{ch}_{M \otimes N}$ is well-defined for any $M, N \in \mathcal{O}$, in general, $M \otimes N$ may not belong to \mathcal{O} since it can fail to be a finitely generated $U(\mathfrak{n}^-)$ -module. This is easiest to see for $\mathfrak{g} = \mathfrak{sl}_2$ where $U(\mathfrak{n}^-) \cong \mathbf{C}[z]$ is a polynomial ring in 1 variable. Then Verma modules are free modules of rank 1 and the tensor product of any 2 (remember, the tensor product is over \mathbf{C}) is not finitely generated (for example, it grows too fast). \square

5. THE CENTER OF $U(\mathfrak{g})$

5.1. Central characters. We let $Z(\mathfrak{g})$ denote the center of $U(\mathfrak{g})$. To get a better understanding of Verma modules, we'll need to understand this ring.

Proposition 5.1.1. *Let M be a highest weight representation of highest weight λ . There exists an algebra homomorphism $\chi_\lambda: Z(\mathfrak{g}) \rightarrow \mathbf{C}$, depending only on λ , such that $zv = \chi_\lambda(z)v$ for all $z \in Z(\mathfrak{g})$ and $v \in M$.*

Proof. Since M is a quotient of the Verma module $M(\lambda)$ it suffices to assume that $M = M(\lambda)$. This also proves that the function χ_λ only depends on λ once we show that it exists.

Let $v^+ \in M_\lambda$ be a highest weight vector. For any $h \in \mathfrak{h}$ and $z \in Z(\mathfrak{g})$, we have

$$h(zv^+) = zhv^+ = \lambda(h)zv^+$$

which shows that $zv^+ \in M_\lambda$. Since $\dim M_\lambda = 1$, there exists a scalar $\chi_\lambda(z)$ such that $zv^+ = \chi_\lambda(z)v^+$. Since v^+ generates M , every element of M can be written as fv^+ for some $f \in U(\mathfrak{g})$. Since z is central, we have $zfv^+ = fzv^+ = \chi_\lambda(z)fv^+$.

It's immediate that χ_λ is an algebra homomorphism. \square

The homomorphism $\chi_\lambda: Z(\mathfrak{g}) \rightarrow \mathbf{C}$ is called the **central character** associated to λ .

With respect to the decomposition $\mathfrak{g} = \mathfrak{n}^- \oplus \mathfrak{h} \oplus \mathfrak{n}$, we get a linear projection map $\text{pr}: \mathfrak{g} \rightarrow \mathfrak{h}$ which also extends to a linear projection

$$\text{pr}: U(\mathfrak{g}) \rightarrow U(\mathfrak{h})$$

using the PBW basis (pick an ordering where $\mathfrak{n}^- < \mathfrak{h} < \mathfrak{n}$): any basis element that uses \mathfrak{n}^- or \mathfrak{n} is sent to 0.

Given $\lambda \in \mathfrak{h}^*$, we also extend it to an algebra homomorphism $\lambda: U(\mathfrak{h}) \rightarrow \mathbf{C}$. Concretely, since \mathfrak{h} is abelian, $U(\mathfrak{h}) \cong \text{Sym}(\mathfrak{h})$ as an algebra, and if we pick a basis x_1, \dots, x_n for \mathfrak{h} , then the elements are polynomials in the x_i , and

$$\lambda(f(x_1, \dots, x_n)) = f(\lambda(x_1), \dots, \lambda(x_n)).$$

The restriction of pr to $Z(\mathfrak{g})$ is denoted ξ :

$$\xi: Z(\mathfrak{g}) \rightarrow \text{Sym}(\mathfrak{h}).$$

Proposition 5.1.2. *For $z \in Z(\mathfrak{g})$, we have*

$$\chi_\lambda(z) = \lambda(\xi(z)).$$

In particular, ξ is an algebra homomorphism.

Proof. Use the notation from the previous proof: v^+ is a highest weight vector of $M(\lambda)$.

Using our PBW basis, we can write z as a linear combination of elements of the form $z_1 z_2 z_3$ where $z_1 \in U(\mathfrak{n}^-)$, $z_2 \in U(\mathfrak{h})$, and $z_3 \in U(\mathfrak{n})$ are all themselves PBW monomials. If $z_3 \neq 1$, then $z_1 z_2 z_3 v^+ = 0$. Otherwise if $z_3 = 1$, then $z_1 z_2 v^+ = z_1 \lambda(z_2) v^+$. Now if $z_1 \neq 1$, this element does not belong to $M(\lambda)_\lambda$, so it will cancel with some other terms (since we know that $z v^+ \in M(\lambda)_\lambda$). Hence only the terms with $z_1 = 1$ and $z_3 = 1$ will contribute, and the sum of these terms is precisely $\xi(z)$, which gives the claimed formula.

By definition, ξ is linear, so we need to check multiplication. Pick $z, z' \in Z(\mathfrak{g})$. Then for every λ , we have

$$\begin{aligned} \lambda(\xi(z z') - \xi(z)\xi(z')) &= \lambda(\xi(z z')) - \lambda(\xi(z))\lambda(\xi(z')) \\ &= \chi_\lambda(z z') - \chi_\lambda(z)\chi_\lambda(z') = 0 \end{aligned}$$

since χ_λ is a homomorphism. Since this holds for all λ , this means that $\xi(z z') - \xi(z)\xi(z')$ is a polynomial that is identically 0 on \mathfrak{h}^* , and hence it is the 0 polynomial. \square

ξ is called the **Harish-Chandra homomorphism**.

5.2. Linked weights. If $L(\mu)$ is a subquotient of $M(\lambda)$, then necessarily $\chi_\lambda = \chi_\mu$. We will try to understand when this equality holds.

Recall that for $\lambda, \alpha \in \mathfrak{h}^*$, we define $\langle \lambda, \alpha \rangle = 2(\lambda, \alpha) / (\alpha, \alpha)$.

As before, let v^+ denote a highest weight vector of weight λ in $M(\lambda)$.

Proposition 5.2.1. *Let $\lambda \in \mathfrak{h}^*$ be a weight, $\alpha \in \Delta$, and pick a nonzero $y \in \mathfrak{g}_{-\alpha}$. If $n := \langle \lambda, \alpha \rangle \in \mathbf{Z}_{\geq 0}$, then $y^{n+1}v^+$ is a nonzero highest weight vector of weight $\lambda - (n+1)\alpha$.*

In particular, there is a nonzero homomorphism

$$M(\lambda - (n+1)\alpha) \rightarrow M(\lambda).$$

This involves some calculations which I'll outline in the exercises.

In particular, using the notation above, we have

$$\lambda - (n+1)\alpha = \lambda - \langle \lambda, \alpha \rangle \alpha - \alpha = s_\alpha(\lambda) - \alpha.$$

We can phrase this another way that's more convenient. First, define $\rho \in \mathfrak{h}^*$ by

$$\rho = \frac{1}{2} \sum_{\alpha \in \Phi^+} \alpha.$$

This is another important notation and ρ will always have this meaning.

For $\alpha \in \Delta$, we stated before that s_α permutes $\Phi^+ \setminus \{\alpha\}$ and $s_\alpha(\alpha) = -\alpha$, so

$$s_\alpha(\rho) = \rho - \alpha.$$

In particular,

$$s_\alpha(\lambda) - \alpha = s_\alpha(\lambda + \rho) - \rho.$$

We define the **shifted Weyl group action** (or **dotted Weyl group action**) of W on \mathfrak{h}^* by

$$w \bullet \lambda = w(\lambda + \rho) - \rho.$$

Its significance is the following:

Proposition 5.2.2. *If $\mu = w \bullet \lambda$ for some $w \in W$, then $\chi_\mu = \chi_\lambda$.*

Proof. First suppose that $\lambda \in \Lambda$, pick $\alpha \in \Delta$, and let $n = \langle \lambda, \alpha \rangle$. We claim that $\chi_{s_\alpha \bullet \lambda} = \chi_\lambda$. If $n \geq 0$, we already saw this above. If $n = -1$, then $s_\alpha \bullet \lambda = \lambda$, so there is nothing to show. Otherwise, if $n \leq -2$, let $\mu = s_\alpha \bullet \lambda$. Then $\langle \mu, \alpha \rangle = -n - 2 \geq 0$, and the first case applied to μ shows that $\chi_\mu = \chi_{s_\alpha \bullet \mu}$. Since $\lambda = s_\alpha \bullet \mu$, we're done.

Next, the s_α generate W , and W stabilizes Λ , so we conclude that $\chi_{w \bullet \lambda} = \chi_\lambda$ for all $w \in W$. This translates to the statement that $\lambda(\xi(z)) = (w \bullet \lambda)(\xi(z))$ for all $z \in Z(\mathfrak{g})$. If we fix $z \in Z(\mathfrak{g})$ and $w \in W$, then the function

$$\lambda \mapsto \lambda(\xi(z)) - (w \bullet \lambda)(\xi(z))$$

is a polynomial function on \mathfrak{h}^* . We just said that this polynomial is 0 on Λ . Since Λ is dense in \mathfrak{h}^* under the Zariski topology, we conclude that it is identically 0. More concretely, if we pick a basis for Λ , then we can identify $\Lambda \cong \mathbf{Z}^n$ and $\mathfrak{h}^* \cong \mathbf{C}^n$. We're just saying that any polynomial that has every integer point as a solution must be identically 0 (left as an exercise). \square

This tells us that the image of ξ is invariant under a shifted action of W . We'll describe this more carefully.

Example 5.2.3. Just for reference, here is what ρ looks like for the classical root systems using the same notation from the earlier section:

$$\mathfrak{sl}_n: \frac{1}{2}(n-1, n-3, \dots, -n+3, -n+1).$$

$$\mathfrak{so}_{2n+1}: \frac{1}{2}(2n-1, 2n-3, \dots, 3, 1).$$

$$\mathfrak{sp}_{2n}: (n, n-1, \dots, 2, 1).$$

$$\mathfrak{so}_{2n}: (n-1, n-2, \dots, 1, 0). \quad \square$$

5.3. Harish-Chandra's theorem. We can identify $\text{Sym}(\mathfrak{h})$ with the space of polynomial functions on \mathfrak{h}^* . Using this, define $\tau: \text{Sym}(\mathfrak{h}) \rightarrow \text{Sym}(\mathfrak{h})$ by $(\tau f)(\lambda) = f(\lambda - \rho)$.

We now define the **twisted Harish-Chandra homomorphism** $\psi: Z(\mathfrak{g}) \rightarrow \text{Sym}(\mathfrak{h})$ by $\psi = \tau \circ \xi$. In particular, for every $\lambda \in \mathfrak{h}^*$ and $z \in Z(\mathfrak{g})$, we have

$$\chi_\lambda(z) = \lambda(\tau^{-1}\psi(z)) = (\lambda + \rho)(\psi(z)).$$

The linear action of W on \mathfrak{h} extends to an action of W on $\text{Sym}(\mathfrak{h})$ by algebra automorphisms. We claim that each $\psi(z)$ is invariant under this action: given $w \in W$ and $\lambda \in \mathfrak{h}^*$, we have

$$\lambda(w^{-1}(\psi(z))) = (w\lambda)(\psi(z)) = \chi_{w(\lambda)-\rho}(z) = \chi_{w \bullet (\lambda-\rho)}(z) = \chi_{\lambda-\rho}(z) = \lambda(\psi(z)).$$

Again, since this is true for all λ , we have $w^{-1}(\psi(z)) = \psi(z)$. In particular, the image of ψ lies in the space of W -invariant polynomials $\text{Sym}(\mathfrak{h})^W$.

Theorem 5.3.1 (Harish-Chandra). *ψ gives an isomorphism from $Z(\mathfrak{g})$ to $\text{Sym}(\mathfrak{h})^W$.*

We won't prove this, but here are some consequences.

Corollary 5.3.2. (1) *The Harish-Chandra map ξ is injective.*

- (2) If $\chi_\lambda = \chi_\mu$, there exists $w \in W$ such that $\mu = w \bullet \lambda$.
- (3) Every algebra homomorphism $Z(\mathfrak{g}) \rightarrow \mathbf{C}$ is of the form χ_λ .

Proof. (1) Immediate from Harish-Chandra's theorem since τ is invertible.

(2) Suppose that μ is not in the shifted W -orbit of λ ; equivalently, assume that $\mu + \rho$ and $\lambda + \rho$ are not in the same W -orbit. First, given two finite subsets S and T of \mathfrak{h}^* , there exists a polynomial $f \in \text{Sym}(\mathfrak{h})$ such that $f(s) = 0$ for all $s \in S$ and $f(t) = 1$ for all $t \in T$ (omitted). Apply this with S being the W -orbit of $\mu + \rho$ and T being the W -orbit of $\lambda + \rho$. Now define $g = \frac{1}{|W|} \sum_{w \in W} wf$ to get a W -invariant polynomial with the same property.

By Harish-Chandra's theorem, there exists z such that $\psi(z) = g$. But then $\chi_\lambda(z) = (\lambda + \rho)(\psi(z)) = g(\lambda + \rho) = 1$ while $\chi_\mu(z) = 0$, so we see that $\chi_\lambda \neq \chi_\mu$.

(3) Given an algebra homomorphism $\varphi: Z(\mathfrak{g}) \rightarrow \mathbf{C}$ define an algebra homomorphism $\chi: \text{Sym}(\mathfrak{h})^W \rightarrow \mathbf{C}$ via $\chi = \varphi \circ \psi^{-1}$. Then $\mathfrak{m} = \ker \chi$ is a maximal ideal of $\text{Sym}(\mathfrak{h})^W$. Next, $\text{Sym}(\mathfrak{h})^W \subset \text{Sym}(\mathfrak{h})$ is an integral extension (in the sense of commutative algebra) and hence (by the going-up theorem) there exists a maximal ideal $\mathfrak{m}' \subset \text{Sym}(\mathfrak{h})$ such that $\mathfrak{m}' \cap \text{Sym}(\mathfrak{h})^W = \mathfrak{m}$. Next, by Hilbert nullstellensatz, $\text{Sym}(\mathfrak{h})/\mathfrak{m}' \cong \mathbf{C}$, and so the quotient map $\text{Sym}(\mathfrak{h}) \rightarrow \mathbf{C}$ can be identified with evaluation at some point $\lambda \in \mathfrak{h}^*$. So χ has the same description, and this translates to $\varphi = \chi_{\lambda - \rho}$. \square

Remark 5.3.3. Actually, $\text{Sym}(\mathfrak{h})^W$ is known to be isomorphic to a polynomial ring itself. We just give the examples in the classical cases.

- For \mathfrak{sl}_n , recall that $\mathfrak{h} = \{(c_1, \dots, c_n) \in \mathbf{C}^n \mid c_1 + \dots + c_n = 0\}$ and W is the symmetric group on n letters permuting the coordinates. Let x_i be the i th standard basis vector. Then the generators for $\text{Sym}(\mathfrak{h})^W$ are the elementary symmetric polynomials $e_r(x)$ for $r = 2, 3, \dots, n$, which are the sum of all squarefree monomials of degree r :

$$e_r(x) = \sum_{i_1 < \dots < i_r} x_{i_1} \cdots x_{i_r}.$$

These are well-known to be algebraically independent.

- The cases \mathfrak{so}_{2n+1} and \mathfrak{sp}_{2n} are the same: we have $\mathfrak{h} = \mathbf{C}^n$ and W is the group of $n \times n$ signed permutation matrices. Again, letting x_i be the i th standard basis vector, the generators for $\text{Sym}(\mathfrak{h})^W$ are $e_r(x^2)$ for $r = 1, 2, \dots, n$:

$$e_r(x^2) = \sum_{i_1 < \dots < i_r} x_{i_1}^2 \cdots x_{i_r}^2.$$

- Finally, for \mathfrak{so}_{2n} , we have $\mathfrak{h} = \mathbf{C}^n$ and W is the group of $n \times n$ signed permutation matrices that have an even number of -1 . Using the notation above, the generators for $\text{Sym}(\mathfrak{h})^W$ are $e_r(x^2)$ for $r = 1, 2, \dots, n - 1$ together with $e_n(x)$. \square

5.4. Application: \mathcal{O} is artinian. We recall some definitions. A module M is **artinian** if it satisfies the descending chain condition, i.e., every descending chain of submodules is eventually constant. If the module is also finitely generated, this implies that it has a **composition series**: a finite chain of submodules $0 = M^0 \subset M^1 \subset \dots \subset M^r = M$ such that each M^i/M^{i-1} is irreducible.

Recall that every submodule and quotient module is artinian. Similarly, if M' is an artinian submodule of M such that M/M' is artinian, this implies that M is also artinian.

Proposition 5.4.1. *If $M \in \mathcal{O}$, then M is artinian.*

Proof. We already saw that M has a filtration whose quotients are highest weight modules, so it suffices to show that highest weight modules are artinian. Since they are quotients of Verma modules, we may also assume that $M = M(\lambda)$ for some $\lambda \in \mathfrak{h}^*$. Define $V = \sum_{w \in W} M(\lambda)_{w \bullet \lambda}$. Then $\dim V < \infty$ since each weight space is finite-dimensional and W is finite.

Now suppose that $N \subsetneq N'$ are submodules of M . Then $z \in Z(\mathfrak{g})$ acts on N'/N by the scalar $\chi_\lambda(z)$. Since $N'/N \neq 0$, it has a highest weight vector, say of weight μ . So this means $\chi_\lambda = \chi_\mu$ and hence $\mu = w \bullet \lambda$ for some $w \in W$. In particular, $N \cap V \neq N' \cap V$. This tells us that any descending chain of distinct submodules of M must have length at most $\dim V$, so M satisfies the descending chain condition. \square

So every $M \in \mathcal{O}$ has a composition series $M^0 \subset \cdots \subset M^r$. The irreducible representations are of the form $L(\lambda)$ for $\lambda \in \mathfrak{h}^*$. We let $[M : L(\lambda)]$ denote the number of i such that $L(\lambda) \cong M^i/M^{i-1}$. This does not depend on the choice of composition series by the Jordan–Hölder theorem.

Corollary 5.4.2. *If $M, N \in \mathcal{O}$, then $\dim \text{Hom}_{\mathcal{O}}(M, N) < \infty$.*

Proposition 5.4.3. *If $[M(\lambda) : L(\mu)] \neq 0$, then $\mu \leq \lambda$.*

Proof. If $[M(\lambda) : L(\mu)] \neq 0$, then $M(\lambda)_\mu \neq 0$. But every weight μ appearing in $M(\lambda)$ satisfies $\mu \leq \lambda$ since $M(\lambda) \cong U(\mathfrak{n}^-) \otimes \mathbf{C}_\lambda$ as an \mathfrak{h} -representation. \square

5.5. Grothendieck group. The **Grothendieck group** of \mathcal{O} , denoted $K(\mathcal{O})$ (here K is the first letter of the German word for “class”), is defined to be the free abelian group whose basis consists of the objects of \mathcal{O} (given M , let $[M]$ denote the corresponding basis element) modulo the relation $[B] = [A] + [C]$ for every short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$. (This is not technically correct due to set-theoretic issues, but they are easy to fix and not important for understanding what we will do with $K(\mathcal{O})$.)

This is a standard definition which you can make for any category with some notion of short exact sequences.

Example 5.5.1. Let Vec be the category of finite-dimensional (complex) vector spaces. If $V \cong W$, then we have a short exact sequence $0 \rightarrow V \rightarrow W \rightarrow 0 \rightarrow 0$ which tells us that $[W] = [V] + [0]$. Applying this to $V = W = 0$, we conclude that $[0] = 0$ and in particular, $[V] = [W]$, so isomorphic objects have the same class (this observation was not special to vector spaces). Next, everything is isomorphic to \mathbf{C}^n for some n , and we have short exact sequences $0 \rightarrow \mathbf{C}^a \rightarrow \mathbf{C}^n \rightarrow \mathbf{C}^{n-a} \rightarrow 0$ for any $a \leq n$, so $[\mathbf{C}^n] = n[\mathbf{C}^1]$ for all $n \geq 0$.

Next, $K(\text{Vec})$ is nonzero: there is a surjective homomorphism $\dim: K(\text{Vec}) \rightarrow \mathbf{Z}$ given by

$$\dim\left(\sum_V c_V [V]\right) = \sum_V c_V \dim(V).$$

since $\dim(B) = \dim(A) + \dim(C)$ for any short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$. We conclude from this that $K(\text{Vec}) \cong \mathbf{Z}$. \square

Note that $\text{ch}_B = \text{ch}_A + \text{ch}_C$ if there is a short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$, which tells us that there is a well-defined linear map

$$\text{ch}: K(\mathcal{O}) \rightarrow \mathcal{X}.$$

Since we know that \mathcal{O} is artinian, we have

$$[M] = \sum_{\lambda \in \mathfrak{h}^*} [M : L(\lambda)] [L(\lambda)].$$

(the point is that the sum is finite and $[M : L(\lambda)]$ is well-defined). So the $[L(\lambda)]$ span $K(\mathcal{O})$. They're also linearly independent (use ch; I'll leave it as an exercise).

This is a nice organizational tool. One obvious question is how to compute the numbers $[M(\lambda) : L(\mu)]$. We know that this is 0 unless $\mu \in W \bullet \lambda$ and $\mu \leq \lambda$. Furthermore, $[M(\lambda) : L(\lambda)] = 1$. This tells us that the change of basis matrix between $\{M(w \bullet \lambda)\}_{w \in W}$ and $\{L(w \bullet \lambda)\}_{w \in W}$ is upper-triangular with 1's on the diagonal if we pick an ordering of $w \bullet \lambda$ that extends the partial order \leq , and hence is invertible. So we have integers $b(\lambda, \mu)$ such that

$$[L(\lambda)] = \sum_{\substack{\mu \in W \bullet \lambda \\ \mu \leq \lambda}} b(\lambda, \mu) [M(\mu)].$$

It would also be nice to understand how to compute these numbers since this tells us how to compute the character of $L(\lambda)$:

$$\text{ch}_{L(\lambda)} = \sum_{\mu} \frac{b(\lambda, \mu) e^{\mu}}{\prod_{\alpha \in \Phi^-} (1 - e^{\alpha})}.$$

Remark 5.5.2. $L(\lambda)$ is finite-dimensional if and only if $\lambda \in \Lambda \cap \Gamma$, i.e., $\langle \lambda, \alpha \rangle \in \mathbf{Z}_{\geq 0}$ for all positive roots α . In that case, the *Weyl character formula* gives the coefficients $b(\lambda, \mu)$ explicitly. In fact, they are always ± 1 ; more precisely, if $\mu = w \bullet \lambda$ (the condition on λ guarantees that there is a single w that works for μ), then $b(\lambda, \mu) = (-1)^{\ell(w)}$. The BGG resolution, to be discussed later, will give a different way to understand this identity. \square

6. BLOCK DECOMPOSITION

6.1. **Yoneda Ext.** Category \mathcal{O} is quite large, but it turns out to decompose into much smaller, more manageable pieces, called blocks. To discuss this notion, we review some basics on extensions (in the sense of homological algebra) from the Yoneda perspective.

First, fix modules A, C . Consider a short exact sequence E of the form

$$E : 0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$$

(in short, this means f is injective, g is surjective, and $\ker g = \text{image } f$). Given another extension E'

$$E' : 0 \rightarrow A \xrightarrow{f'} B' \xrightarrow{g'} C \rightarrow 0,$$

we say that E and E' are **Yoneda equivalent** if there is a homomorphism $\varphi : B \rightarrow B'$ such that $f' = \varphi \circ f$ and $g = g' \circ \varphi$. Pictorially, we are asking for this diagram to commute:

$$\begin{array}{ccccccc} 0 & \longrightarrow & A & \xrightarrow{f} & B & \xrightarrow{g} & C \longrightarrow 0 \\ & & \text{id} \downarrow & & \downarrow \varphi & & \downarrow \text{id} \\ 0 & \longrightarrow & A & \xrightarrow{f'} & B' & \xrightarrow{g'} & C \longrightarrow 0 \end{array}$$

Here are a few key facts (without proof, this can take a long time to do everything properly, so you'll have to find a reference on homological algebra if you want the details):

- This is an equivalence relation; let $\text{Ext}_{\mathcal{O}}^1(C, A)$ denote the set of equivalence classes.
- $\text{Ext}_{\mathcal{O}}^1(C, A)$ has the structure of a complex vector space, where the 0 vector is represented by the split sequence

$$0 \rightarrow A \rightarrow A \oplus C \rightarrow C \rightarrow 0$$

with the first map being $x \mapsto (x, 0)$ and the second map being $(x, y) \mapsto y$.

- For fixed M , the assignment $A \mapsto \text{Ext}_{\mathcal{O}}^1(M, A)$ is a functor from \mathcal{O} to the category of complex vector spaces; and the assignment $A \mapsto \text{Ext}_{\mathcal{O}}^1(A, M)$ is a contravariant functor.
- For fixed M , and any short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C$, we have exact sequences

$$0 \rightarrow \text{Hom}_{\mathcal{O}}(M, A) \rightarrow \text{Hom}_{\mathcal{O}}(M, B) \rightarrow \text{Hom}_{\mathcal{O}}(M, C) \rightarrow \text{Ext}_{\mathcal{O}}^1(M, A) \rightarrow \text{Ext}_{\mathcal{O}}^1(M, B) \rightarrow \text{Ext}_{\mathcal{O}}^1(M, C),$$

$$0 \rightarrow \text{Hom}_{\mathcal{O}}(C, M) \rightarrow \text{Hom}_{\mathcal{O}}(B, M) \rightarrow \text{Hom}_{\mathcal{O}}(A, M) \rightarrow \text{Ext}_{\mathcal{O}}^1(C, M) \rightarrow \text{Ext}_{\mathcal{O}}^1(B, M) \rightarrow \text{Ext}_{\mathcal{O}}^1(A, M).$$

- We have $\text{Ext}_{\mathcal{O}}^1(C, A \oplus A') \cong \text{Ext}_{\mathcal{O}}^1(C, A) \oplus \text{Ext}_{\mathcal{O}}^1(C, A')$ and similarly in the other argument.

If we know that $\text{Ext}_{\mathcal{O}}^1(C, A) = 0$, then for every short exact sequence $0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$, it must be equivalent to the split sequence. This implies that there exists $C \xrightarrow{s} B$ such that $gs = \text{id}_B$, and in particular, $B \cong A \oplus C$.

Remark 6.1.1. If you have seen derived functors, then $\text{Ext}_{\mathcal{O}}^1(A, -)$ can also be defined as the first derived functor of $\text{Hom}_{\mathcal{O}}(A, -)$ and similarly for the other argument. It takes a bit of work to prove that these two definitions agree, so we won't cover it here. \square

6.2. Blocks. Blocks are certain subcategories of \mathcal{O} which will decompose it. Given two irreducible representations L and L' , we say that they belong to the same block if $\text{Ext}_{\mathcal{O}}^1(L, L') \neq 0$, i.e., there exists a non-split exact sequence $0 \rightarrow L' \rightarrow M \rightarrow L \rightarrow 0$ or if $\text{Ext}_{\mathcal{O}}^1(L', L) \neq 0$. More generally, two irreducible representations L and L' are in the same block if there exists a sequence $L = L_1, L_2, \dots, L_n = L'$ such that L_i and L_{i+1} are in the same block for all $i = 1, \dots, n-1$. We'll denote the set of blocks by \mathcal{B} . A general $M \in \mathcal{O}$ belongs to a block b if all of its composition factors belong to b .

Given a block $b \in \mathcal{B}$ and $M \in \mathcal{O}$, let $[M]_b$ be the sum of $[M : L(\lambda)][L(\lambda)]$ ranging over all $L(\lambda) \in b$. Then we get $[M] = \sum_{b \in \mathcal{B}} [M]_b$; we'd like to lift this decomposition to the level of modules.

Given a general object $M \in \mathcal{O}$, and a block $b \in \mathcal{B}$, let M_b denote the sum of all submodules M' of M such that all composition factors of M' belong to b . Since M is artinian, $M_b \neq 0$ for only finitely many b . The following proof was adapted from [J, §II.7.1].

Proposition 6.2.1. *We have*

$$M = \bigoplus_{b \in \mathcal{B}} M_b.$$

In particular, given any other $M' \in \mathcal{O}$, we have

$$\text{Hom}_{\mathcal{O}}(M, M') \cong \bigoplus_{b \in \mathcal{B}} \text{Hom}_{\mathcal{O}}(M_b, M'_b),$$

and if M is indecomposable, then all of its composition factors belong to the same block.

Proof. Let $N = \sum_{b \in \mathcal{B}} M_b$. For any b , we have $M_b \cap (\sum_{b' \neq b} M_{b'}) = 0$ since they don't share any composition factors and hence can't have any common submodules. Hence the sum defining N is direct. We need to show that $N = M$.

Suppose not. Then $M/N \neq 0$, so it has a simple submodule L (since it is artinian) and so its preimage in M gives a submodule N' that contains N such that $N'/N \cong L$. Let b be the

block that contains L . The sequence $0 \rightarrow N \rightarrow N' \rightarrow L \rightarrow 0$ is non-split: otherwise there is a preimage of L that belongs to M_b and contradicts that $N' \neq N$. So

$$0 \neq \text{Ext}_{\mathcal{O}}^1(L, N) \cong \bigoplus_{b' \in \mathcal{B}} \text{Ext}_{\mathcal{O}}^1(L, M_{b'})$$

so there exists some composition factor L' of some $M_{b'}$ such that $\text{Ext}_{\mathcal{O}}^1(L, L') \neq 0$. This implies that $b = b'$ by definition. Now set $N'' = \bigoplus_{b' \neq b} M_{b'}$. Since all composition factors of N'/N'' belong to b , the exact sequence $0 \rightarrow N'' \rightarrow N' \rightarrow N'/N'' \rightarrow 0$ must be split. But then by definition, this means that there is a preimage of N'/N'' in N' that is contained in M_b , which implies that $N' \subseteq N$, a contradiction.

For the statement about Hom, it follows from the fact that Hom commutes with finite direct sums in both arguments, and there are no nonzero maps $M_b \rightarrow M_{b'}$ if $b \neq b'$ by definition. \square

We can think of the assignment $M \mapsto M_b$ as a functor $\mathcal{O} \mapsto b$, where we think of b as being the full subcategory on objects belonging to b . The above discussion implies that this functor is exact (i.e., preserves exact sequences). We can also think of the above result as giving a decomposition

$$\mathcal{O} \simeq \bigoplus_{b \in \mathcal{B}} b.$$

Our next goal is to determine what these blocks are with a combinatorial description.

6.3. Subcategories \mathcal{O}_χ . First, a reminder on generalized eigenvectors. Let X be a linear operator on a space V and let λ be a scalar. Then $v \in V$ is a generalized eigenvector with eigenvalue λ if there exists n such that $(X - \lambda)^n v = 0$ and the set of all generalized eigenvectors of eigenvalue λ is the generalized eigenspace. If $\dim V < \infty$, then we can rephrase Jordan canonical form as saying that V is a direct sum of its generalized eigenspaces.

If X, Y commute, then the generalized eigenspaces of X are invariant under Y , and vice versa.

Now let χ be a central character (an algebra homomorphism $Z(\mathfrak{g}) \rightarrow \mathbf{C}$). Given $M \in \mathcal{O}$, define M^χ to be the set of $v \in M$ such that v is a generalized eigenvector with eigenvalue $\chi(x)$ for all $x \in Z(\mathfrak{g})$. Each M^χ is a subrepresentation of M (check).

Proposition 6.3.1. *If $M \in \mathcal{O}$, then $M = \bigoplus_\chi M^\chi$.*

Proof. Given a weight μ , M_μ is finite-dimensional. Furthermore, M_μ is a $Z(\mathfrak{g})$ -submodule of M . Hence the intersections $M_\mu \cap M^\chi$ are generalized eigenspaces for the action of $Z(\mathfrak{g})$ on M_μ , so M_μ is the direct sum of them. Since M is a sum of its weight spaces, this gives the result. \square

Define \mathcal{O}_χ to be the full subcategory on objects M such that $M = M^\chi$.

Proposition 6.3.2. *If $M \in \mathcal{O}_\chi$ and $M' \in \mathcal{O}_{\chi'}$ with $\chi \neq \chi'$, then $\text{Ext}_{\mathcal{O}}^1(M, M') = 0$.*

Proof. Consider an extension $0 \rightarrow M' \rightarrow N \rightarrow M \rightarrow 0$. But then $M' = N^{\chi'}$ and N^χ is a submodule of N which maps isomorphically to M . We can use an inverse of this isomorphism to split the sequence. \square

In particular, every block belongs to some \mathcal{O}_χ and we have a decomposition

$$\mathcal{O} \simeq \bigoplus_x \mathcal{O}_\chi.$$

Recall that we already discussed that every central character is of the form χ_λ for some weight λ .

Proposition 6.3.3. *If $\lambda \in \Lambda$ is an integral weight, then $\mathcal{O}_{\chi_\lambda}$ is a block.*

Proof. The irreducible representations of $\mathcal{O}_{\chi_\lambda}$ are all of the form $L(w \bullet \lambda)$ for $w \in W$. Since W is generated by simple reflections, it will suffice to show that $L(\lambda)$ and $L(\mu)$ are in the same block where $\mu = s_\alpha \bullet \lambda$ for some simple root α . Let $n = \langle \lambda, \alpha \rangle$. Without loss of generality, we can assume that $n \geq 0$ (if $n = -1$, then $\lambda = \mu$ and if $n \leq -2$ we may reverse the roles of λ and μ).

By Proposition 5.2.1, we know that there is a nonzero homomorphism $f: M(\mu) \rightarrow M(\lambda)$ (whose image lies in the maximal submodule $N(\lambda)$). Let $N = f(N(\mu))$. Then $M(\lambda)/N$ is a highest weight module containing a submodule isomorphic to $L(\mu)$ (the image of $M(\mu)$ under f) and a quotient module isomorphic to $L(\lambda)$ (since its highest weight is λ). Highest weight modules are indecomposable (since they are cyclic modules) so we get the result since we already saw that indecomposable modules belong to a single block. \square

For the central weight χ_0 , the corresponding block \mathcal{O}_{χ_0} is called the **principal block**. Recall this comes from the trivial representation, so $\chi_0(x) = 1$ for all $x \in Z(\mathfrak{g})$.

When λ is not integral, $\mathcal{O}_{\chi_\lambda}$ is not going to be a single block in general.

Example 6.3.4. Consider $\mathfrak{g} = \mathfrak{sl}_2$. In previous examples, we identified Λ with \mathbf{Z} , and \mathfrak{h}^* with \mathbf{C} ; then $\rho = 1$ with this identification.

The two irreducible representations of $\mathcal{O}_{\chi_\lambda}$ are $L(\lambda)$ and $L(-\lambda - 2)$.

We saw that if $\alpha \notin \mathbf{Z}_{\geq 0}$, then the Verma module $M(\alpha)$ is irreducible, so $M(\alpha) = L(\alpha)$. In particular, if $\lambda \notin \mathbf{Z}$ then $\mathcal{O}_{\chi_\lambda}$ has two irreducible representations which are both Verma modules.

However, $M(\lambda)$ and $M(-\lambda - 2)$ are not in the same block: suppose N is an extension of these two modules (in either order). Starting with any weight vector, applying elements of $U(\mathfrak{sl}_2)$ can only change its weight by integer multiples of 2 (since the weight of X is 2 and the weight of Y is -2), and so the sum of weight spaces for either $M(\lambda)$ or $M(-\lambda - 2)$ in N must both be subrepresentations, so N is the direct sum of the two.

So we see a few different kinds of behavior here:

- If $\lambda \in \mathbf{Z} \setminus \{-1\}$, then $\mathcal{O}_{\chi_\lambda}$ is a single block and has two irreducible representations.
- If $\lambda = -1$, then $\mathcal{O}_{\chi_{-1}}$ is still a single block, but only has one irreducible representation.
- If $\lambda \notin \mathbf{Z}$, then $\mathcal{O}_{\chi_\lambda}$ splits into two blocks, each having one irreducible representation.

\square

6.4. Dominant and antidominant weights. In the previous example, we used the fact that in a nontrivial extension of irreducible representations of \mathfrak{sl}_2 , the weights should always differ by integer multiples of 2, or said another way, should all belong the same coset of $2\mathbf{Z}$ in \mathbf{C} . For general \mathfrak{g} , we can replace $2\mathbf{Z}$ with the root lattice Λ_r and formalize it as follows.

Proposition 6.4.1. *Let $V \in \mathcal{O}$ and for each coset $[\alpha]$ of \mathfrak{h}^*/Λ_r , define $V_{[\alpha]} = \bigoplus_{\lambda \in [\alpha]} V_\lambda$. Then each $V_{[\alpha]}$ is a subrepresentation of V and $V = \bigoplus_{[\alpha] \in \mathfrak{h}^*/\Lambda_r} V_{[\alpha]}$.*

This motivates the following definition. For $\lambda \in \mathfrak{h}^*$, define $[\lambda]$ to be its coset in \mathfrak{h}^*/Λ_r , and define (recall that W is the Weyl group)

$$W_{[\lambda]} = \{w \in W \mid w \bullet \lambda - \lambda \in \Lambda_r\}.$$

As the notation suggests, W acts on \mathfrak{h}^*/Λ_r via the dotted action and $W_{[\lambda]}$ is the stabilizer of $[\lambda]$ under this action.

In particular, $\mathcal{O}_{\chi_\lambda}$ decomposes further as a direct sum of subcategories indexed by the orbits of $W_{[\lambda]}$ acting on $W \bullet \lambda$ (via the dotted action). These are in fact blocks, but we probably won't get to that proof, see [H3, §4.9].

Recall for $\lambda, \mu \in \mathfrak{h}^*$, that we defined $\langle \lambda, \mu \rangle = 2(\lambda, \mu)/(\mu, \mu)$.

Remark 6.4.2. In [H3, §3.4], it is proven that the collection of vectors

$$\Phi_{[\lambda]} = \{\alpha \in \Phi \mid \langle \lambda, \alpha \rangle \in \mathbf{Z}\},$$

is a root system inside of its linear span, and that $W_{[\lambda]}$ is its Weyl group. \square

Definition 6.4.3. $\lambda \in \mathfrak{h}^*$ is **antidominant** if $\langle \lambda + \rho, \alpha \rangle \notin \mathbf{Z}_{>0}$ for all positive roots $\alpha \in \Phi^+$. Analogously, λ is **dominant** if $\langle \lambda + \rho, \alpha \rangle \notin \mathbf{Z}_{<0}$ for all $\alpha \in \Phi^+$. \square

As Humphreys warns, this definition of dominant is a relaxation of the usual definition of dominance (which also requires that $\lambda \in \Lambda$).

Recall that $\lambda < \mu$ means that $\mu - \lambda$ can be written as a sum of positive roots with positive integer coefficients.

Proposition 6.4.4. *Pick $\lambda \in \mathfrak{h}^*$. An element of the dotted orbit $W_{[\lambda]} \bullet \lambda$ is antidominant if and only if it is minimal with respect to the partial order $<$, and there is a unique such element.*

Similarly, an element is dominant if and only if it is maximal with respect to $<$, and there is a unique such element.

I'll just prove the easier implication; for the rest, see [H3, §3.5].

First, if μ is a minimal element of $W_{[\lambda]} \bullet \lambda$, then it must be antidominant. If not, then pick $\alpha \in \Phi^+$ such that $\langle \mu + \rho, \alpha \rangle \in \mathbf{Z}_{>0}$. But then

$$s_\alpha \bullet \mu = s_\alpha(\mu + \rho) - \rho = (\mu + \rho) - \langle \mu + \rho, \alpha \rangle \alpha - \rho = \mu - \langle \mu + \rho, \alpha \rangle \alpha,$$

so $s_\alpha \bullet \mu < \mu$ and $s_\alpha \in W_{[\mu]} = W_{[\lambda]}$, which contradicts minimality.

In particular, there are two canonical choices for a representative of the coset $[\lambda]$ in \mathfrak{h}^*/Λ_r : either the antidominant or dominant weight. They will play special roles. For example, here's an easy property:

Proposition 6.4.5. *If λ is antidominant, then $M(\lambda)$ is irreducible.*

6.5. Duality functor. First, there is an anti-involution τ of \mathfrak{g} (i.e., $\tau([x, y]) = [\tau(y), \tau(x)]$ and $\tau^2 = 1$) such that $\tau(h) = h$ for all $h \in \mathfrak{h}$ called the **transpose map**. In particular, for each root α , we have $\tau(\mathfrak{g}_\alpha) = \mathfrak{g}_{-\alpha}$.

Example 6.5.1. For \mathfrak{sl}_n , we can take $\tau(x) = x^T$, the usual transpose operation. Actually this works for all of the classical cases if we use the particular choices of bilinear forms from §4.3. \square

Given a representation M , we will denote by $M^{*,\tau}$ the dual space M^* with the action by \mathfrak{g} defined by

$$(xf)(v) = f(\tau(x)v)$$

for $f \in M^*$, $x \in \mathfrak{g}$, and $v \in M$.

Temporarily denote by \mathcal{C} the full subcategory of $U(\mathfrak{g})$ -modules which admit a weight decomposition with finite-dimensional weight spaces.

If $M \in \mathcal{C}$, then for every weight λ , we can identify M_λ^* with the subspace of $f \in M^*$ such that $f(M_\mu) = 0$ for all $\mu \neq \lambda$. We define

$$M^\vee = \bigoplus_{\lambda \in \mathfrak{h}^*} M_\lambda^* \subseteq M^{*,\tau}.$$

Then M^\vee is a subrepresentation, so in particular, is a representation. Since τ fixes \mathfrak{h} , we have $M_\lambda^\vee = M_\lambda^*$, so $M^\vee \in \mathcal{C}$.

If $\varphi: M \rightarrow N$ is a homomorphism, then it preserves weight spaces, and so we get maps $\varphi_\lambda^*: M_\lambda^* \rightarrow N_\lambda^*$; taking the sum gives a dual homomorphism $\varphi^\vee: N^\vee \rightarrow M^\vee$.

We start with some basic observations:

- Proposition 6.5.2.** (1) If $M \in \mathcal{C}$, then we have a natural isomorphism $M \cong (M^\vee)^\vee$.
 In particular, \vee gives an equivalence $\mathcal{C} \simeq \mathcal{C}^{\text{op}}$.
 (2) Given a short exact sequence $0 \rightarrow M_1 \rightarrow M_2 \rightarrow M_3 \rightarrow 0$, duality induces a short exact sequence $0 \rightarrow M_3^\vee \rightarrow M_2^\vee \rightarrow M_1^\vee \rightarrow 0$.
 (3) For every weight λ , we have $L(\lambda)^\vee \cong L(\lambda)$.
 (4) If $M \in \mathcal{C}$, then $\text{ch}_M = \text{ch}_{M^\vee}$.

Proof. (1) We're using here that each weight space is finite-dimensional, and hence there is a natural isomorphism $\psi: M_\lambda \rightarrow (M_\lambda^*)^*$ where for $v \in M_\lambda$ and $f \in M_\lambda^*$, $\psi(v)(f)$ is defined to be $f(v)$. This is compatible with the representation structure we've defined since $\tau^2 = 1$.

(2) Not much to say, we can check this on weight spaces one at a time.

(3) By (2), if $L(\lambda)^\vee$ had a nonzero proper submodule, then this would give a nonzero proper quotient of $(L(\lambda)^\vee)^\vee \cong L(\lambda)$ which contradicts that it is irreducible. So $L(\lambda)^\vee$ is also irreducible and a highest weight representation of highest weight λ , so the two are isomorphic by uniqueness.

(4) Clear from construction. □

Proposition 6.5.3. If $M \in \mathcal{O}$, then $M^\vee \in \mathcal{O}$. In particular, \vee gives an equivalence $\mathcal{O} \simeq \mathcal{O}^{\text{op}}$.

Proof. We know that M is artinian, and so has a composition series

$$0 = M_0 \subseteq M_1 \subseteq \cdots \subseteq M_n = M.$$

Define $N_i = \ker(M^\vee \rightarrow M_i^\vee) \cong (M/M_i)^\vee$. Then we get

$$0 = N_n \subseteq N_{n-1} \subseteq \cdots \subseteq N_0 = M^\vee.$$

Also, $N_i/N_{i+1} \cong (M_{i+1}/M_i)^\vee$. From the previous result, these are irreducible, so we have a composition series for M^\vee . In particular, M^\vee is a finitely generated $U(\mathfrak{g})$ -module.

By construction, M^\vee is \mathfrak{h} -semisimple.

Finally, we need to check that M^\vee is locally \mathfrak{n} -finite. First recall that since $M \in \mathcal{O}$, there are finitely many weights $\lambda_1, \dots, \lambda_r$ so that if $M_\beta \neq 0$, then $\beta \leq \lambda_i$ for some i . Now pick a weight vector $v \in M_\beta$. Then $U(\mathfrak{n}) \cdot v$ is contained in the sum of weight spaces M_γ such

that $\gamma \geq \beta$. Also, the intervals $\{\gamma \mid \beta \leq \gamma \leq \lambda_i\}$ are finite (write $\lambda_i - \beta = \sum_j c_j \alpha_j$ as a nonnegative integer sum of simple roots; any γ is of the form $\beta + \sum_j c'_j \alpha_j$ with $c'_j \leq c_j$), and so $U(\mathfrak{n}) \cdot v$ is contained in finitely many weight spaces, and each one is finite dimensional. Finally, everything is a finite sum of weight vectors, so we're done. \square

Now that we know that \vee preserves objects of \mathcal{O} , we get some more consequences:

Corollary 6.5.4. (1) For $M, N \in \mathcal{O}$, we have

$$\mathrm{Ext}_{\mathcal{O}}^1(M, N) \cong \mathrm{Ext}_{\mathcal{O}}^1(N^\vee, M^\vee).$$

In particular, for any pairs of weights λ, μ , we have

$$\mathrm{Ext}_{\mathcal{O}}^1(L(\lambda), L(\mu)) \cong \mathrm{Ext}_{\mathcal{O}}^1(L(\mu), L(\lambda)).$$

(2) We have $(M \oplus N)^\vee \cong M^\vee \oplus N^\vee$. In particular, if $M \in \mathcal{O}$ is indecomposable, then so is M^\vee .

(3) If $M \in \mathcal{O}_\chi$, then $M^\vee \in \mathcal{O}_\chi$. In particular, \vee gives an equivalence $\mathcal{O}_\chi \simeq \mathcal{O}_\chi^{\mathrm{op}}$.

Proof. (1) can be proven using the fact that \vee preserves short exact sequences.

(2) The isomorphism is straightforward; if M^\vee decomposes as $P \oplus Q$, then we get $M \cong P^\vee \oplus Q^\vee$, which means that M is decomposable.

(3) We recall some facts. We have the Harish-Chandra homomorphism $\xi: Z(\mathfrak{g}) \rightarrow U(\mathfrak{h})$ where $Z(\mathfrak{g})$ is the center of $U(\mathfrak{g})$ which is defined by restricting the projection map $U(\mathfrak{g}) \rightarrow U(\mathfrak{h})$ to $Z(\mathfrak{g})$ (specifically, we pick a basis compatible with $\mathfrak{g} = \mathfrak{n}^- \oplus \mathfrak{h} \oplus \mathfrak{n}$ and order it so that $\mathfrak{n}^- < \mathfrak{h} < \mathfrak{n}$ and send all PBW basis elements that use \mathfrak{n}^- or \mathfrak{n} to 0).

Next, τ fixes \mathfrak{h} pointwise and $U(\mathfrak{h})$ is commutative; also it swaps \mathfrak{n} and \mathfrak{n}^- , and so $\xi(\tau(z)) = \xi(z)$ for all $z \in Z(\mathfrak{g})$. In particular, τ fixes $Z(\mathfrak{g})$ pointwise since ξ is injective. Hence, if $z \in Z(\mathfrak{g})$ acts on M by the scalar $\chi(z)$, then it also acts on M^\vee by the same scalar. \square

The dual Verma modules $M(\lambda)^\vee$ gives us a new class of modules to consider. Some more properties:

Proposition 6.5.5. (1) $M(\lambda)^\vee$ has a unique irreducible subrepresentation which is $L(\lambda)$.

If $L(\mu)$ is a composition factor of $M(\lambda)^\vee$, then $\mu \leq \lambda$.

(2) We have

$$\dim \mathrm{Hom}_{\mathcal{O}}(M(\mu), M(\lambda)^\vee) = \delta_{\lambda, \mu}.$$

Every nonzero homomorphism $M(\lambda) \rightarrow M(\lambda)^\vee$ has image equal to $L(\lambda)$.

(3) We have $\mathrm{Ext}_{\mathcal{O}}^1(M(\mu), M(\lambda)^\vee) = 0$ for all λ, μ .

Proof. (1) This follows by applying \vee to the fact that $M(\lambda)$ has a unique irreducible quotient representation which is $L(\lambda)$. Also, the composition factors of $M(\lambda)$ and $M(\lambda)^\vee$ are the same, so the second part follows from previous results.

(2) Suppose there is a nonzero homomorphism $M(\mu) \rightarrow M(\lambda)^\vee$. Its image is artinian, so it has an irreducible submodule, which must be isomorphic to $L(\lambda)$ by (1). This implies that $\lambda \leq \mu$ since then $L(\lambda)$ is a composition factor of $M(\mu)$. But we can dualize to get a nonzero homomorphism $M(\lambda) \rightarrow M(\mu)^\vee$ so we also conclude that $\mu \leq \lambda$, and hence $\lambda = \mu$.

Next, since $M(\lambda)$ is generated by $M(\lambda)_\lambda$, every homomorphism $M(\lambda) \rightarrow M(\lambda)^\vee$ is determined by the image of this subspace. But $\dim M(\lambda)_\lambda^\vee = 1$ so the space of maps is at most 1-dimensional. We do actually get a nonzero map though, since we can take the composition $M(\lambda) \rightarrow L(\lambda) \rightarrow M(\lambda)^\vee$ (the rest are scalar multiples of this map).

(3) Since $\text{Ext}_{\mathcal{O}}^1(M(\mu), M(\lambda)^\vee) \cong \text{Ext}_{\mathcal{O}}^1(M(\lambda), M(\mu)^\vee)$, we may assume without loss of generality that $\lambda \not\preceq \mu$, i.e., $\lambda \leq \mu$ or they are incomparable. This means that μ is a maximal element of the set of weights appearing in either $M(\mu)$ or $M(\lambda)^\vee$. Suppose we have an extension

$$0 \rightarrow M(\lambda)^\vee \rightarrow M \rightarrow M(\mu) \rightarrow 0.$$

Let v be a highest weight vector in $M(\mu)_\mu$ and let $v' \in M_\mu$ be any preimage. Then v' must also be a highest weight vector by maximality of μ , so there is a map $M(\mu) \rightarrow M$ taking v to v' ; this splits the above extension. \square

7. PROJECTIVES AND INJECTIVES

7.1. Projective objects. Recall that a $U(\mathfrak{g})$ -module P is **projective** if the functor $\text{Hom}_{\mathfrak{g}}(P, -)$ is exact. Alternatively, given any surjection $\pi: M \rightarrow N$ and arbitrary homomorphism $\varphi: P \rightarrow N$, there exists a homomorphism $\psi: P \rightarrow M$ (the “lift”) such that $\varphi = \pi \circ \psi$. Free modules are projective, so we can write any representation as a quotient of a projective module; we say that “the category of $U(\mathfrak{g})$ -modules has **enough projectives**”. However, free $U(\mathfrak{g})$ -modules do not belong to \mathcal{O} since they fail to be locally \mathfrak{n} -finite.

We would like to prove that nonetheless, \mathcal{O} has enough projectives. There is a subtle distinction though: if $M \in \mathcal{O}$ is a projective object of \mathcal{O} , it may not be projective as a $U(\mathfrak{g})$ -module because we only require that the lifting property hold for surjections between modules that belong to \mathcal{O} . To emphasize this, we will usually say “projective object of \mathcal{O} ”.

Proposition 7.1.1. *If λ is dominant, i.e., maximal in its $W_{[\lambda]}$ -orbit, then $M(\lambda)$ is a projective object of \mathcal{O} .*

Proof. To check that $M(\lambda)$ is projective, it suffices to only consider modules in the same block as $M(\lambda)$. Let $\pi: U \rightarrow V$ be a surjection in this block and suppose we’re given a map $\varphi: M(\lambda) \rightarrow V$. Let $m_+ \in M(\lambda)_\lambda$ be a highest weight vector and set $v_+ = \varphi(m_+)$. Finally, let $u_+ \in U_\lambda$ be any preimage of v_+ under π . Let β be a maximal element in $\{\mu \mid \mu \geq \lambda, U_\mu \neq 0\}$. Then every vector of U_β is a highest weight vector. However, since U is in the same block as $M(\lambda)$, β is in the $W_{[\lambda]}$ -orbit of λ , and so $\beta = \lambda$ by assumption. Hence there is a well-defined map $\psi: M(\lambda) \rightarrow U$ given by $\psi(m_+) = u_+$, and so $M(\lambda)$ is projective. \square

Proposition 7.1.2. *If $P \in \mathcal{O}$ is a projective object and L is a finite-dimensional representation, then $P \otimes_{\mathbf{C}} L$ is projective.*

Proof. By hom-tensor adjunction, for any $M \in \mathcal{O}$, we have a natural isomorphism

$$\text{Hom}_{\mathcal{O}}(P \otimes_{\mathbf{C}} L, M) = \text{Hom}_{\mathcal{O}}(P, \text{Hom}_{\mathbf{C}}(L, M)),$$

here $\text{Hom}_{\mathbf{C}}(L, M)$ is isomorphic to $L^* \otimes M$ as a representation (usual dual $*$, not \vee). In particular, tensoring with L^* (over \mathbf{C}) is exact, and the functor $\text{Hom}_{\mathcal{O}}(P, -)$ is exact since P is projective, so we conclude that $P \otimes_{\mathbf{C}} L$ is projective. \square

Before continuing, we should examine the structure of tensor products $M(\lambda) \otimes L$ where L is finite-dimensional. In order to do that, let me recall a general associativity property of tensor products. Let R, S be rings, let M_1 be a right R -module, M_2 be an (R, S) -bimodule (i.e., left R -module and right S -module such that $(rm)s = r(ms)$ for all $r \in R, m \in M_2$ and $s \in S$) and M_3 a left S -module. Then we have a natural isomorphism

$$(M_1 \otimes_R M_2) \otimes_S M_3 \cong M_1 \otimes_R (M_2 \otimes_S M_3).$$

Implicit here is that $M_1 \otimes_R M_2$ is a right S -module using the right S -module structure on M_2 , and similarly, $M_2 \otimes_S M_3$ is a left R -module. This gives the following:

Proposition 7.1.3. *If L is a finite-dimensional representation of \mathfrak{g} , then we have an isomorphism of \mathfrak{g} -representations*

$$M(\lambda) \otimes L \cong \text{Ind}_{\mathfrak{b}}^{\mathfrak{g}}(\mathbf{C}_\lambda \otimes L)$$

where on the right hand side, we restricted the \mathfrak{g} -action on L to \mathfrak{b} .

Proof. We are using the previous isomorphism with $R = U(\mathfrak{b})$ and $S = \mathbf{C}$, $M_1 = U(\mathfrak{g})$, $M_2 = \mathbf{C}_\lambda$, and $M_3 = L$. Since the $U(\mathfrak{g})$ -module structure on both sides is compatible by naturality. \square

First consider the structure of L as a \mathfrak{b} -representation. Let v_1, \dots, v_N be a weight basis for L with respective weights μ_1, \dots, μ_N . We can assume that they are ordered so that if $\mu_i \leq \mu_j$, then $i \leq j$. Letting z denote $1 \in \mathbf{C}_\lambda$, we have that $z \otimes v_1, \dots, z \otimes v_N$ is a weight basis for $\mathbf{C}_\lambda \otimes L$ with weights $\lambda + \mu_1, \dots, \lambda + \mu_N$. Then the \mathfrak{b} -subrepresentation generated by any $z \otimes v_k$ is contained in the span of $\{z \otimes v_k, z \otimes v_{k+1}, \dots, z \otimes v_N\}$.

We can phrase this as follows: let V_k be the span of $z \otimes v_k, \dots, z \otimes v_N$ in $\mathbf{C}_\lambda \otimes L$. Then we have a filtration by \mathfrak{b} -submodules:

$$V_N \subset V_{N-1} \subset \dots \subset V_1 = \mathbf{C}_\lambda \otimes L.$$

Finally, note that induction is exact, so if we set $V'_k = \text{Ind}_{\mathfrak{b}}^{\mathfrak{g}} V_k$, then $\text{Ind}_{\mathfrak{b}}^{\mathfrak{g}}(\mathbf{C}_\lambda \otimes L)$ has a filtration

$$V'_N \subset V'_{N-1} \subset \dots \subset V'_1 = \text{Ind}_{\mathfrak{b}}^{\mathfrak{g}}(\mathbf{C}_\lambda \otimes L),$$

such that (setting $V'_{N+1} = 0$)

$$V'_k/V'_{k+1} \cong M(\lambda + \mu_k).$$

We can summarize this as follows:

Proposition 7.1.4. *Let μ_1, \dots, μ_N be the weights of L , listed with multiplicities. Then $M(\lambda) \otimes L$ has a filtration whose successive quotients are the Verma modules $M(\lambda + \mu_k)$ for $k = 1, \dots, N$. Furthermore, if ν is a maximal weight amongst μ_1, \dots, μ_N , then $M(\lambda) \otimes L$ contains a submodule isomorphic to $M(\lambda + \nu)$, and similarly, if η is a minimal weight, then $M(\lambda) \otimes L$ has a quotient module isomorphic to $M(\lambda + \eta)$.*

For future reference, we make a definition:

Definition 7.1.5. A filtration of $M \in \mathcal{O}$ whose successive quotients are Verma modules is called a **standard filtration** (or **Verma flag**). \square

In general, M might not have a standard filtration.

First, let me comment that there is a unique element $w_0 \in W$ such that $w_0(\Phi^+) = \Phi^-$. This has maximal possible length ($\ell(w_0) = |\Phi^+|$). In particular,

$$w_0(\rho) = -\rho$$

(recall that ρ is $\frac{1}{2}$ times the sum of all positive roots). One last thing: for each simple root α , we have $s_\alpha(\rho) = \rho - \alpha$ (since α permutes $\Phi^+ \setminus \{\alpha\}$) and hence $\langle \rho, \alpha \rangle = 1$.

When $L = L(\mu)$, μ is the (unique) maximal weight that appears, and $w_0(\mu)$ is the (unique) minimal weight.

Now we have everything to prove our goal.

Proposition 7.1.6. *Category \mathcal{O} has enough projectives.*

Proof. First we show that given a weight λ , there is a projective object P such that $L(\lambda)$ is a quotient of P . If λ is dominant, then we saw that we can take $P = M(\lambda)$. Recall that dominant means that $\langle \lambda + \rho, \alpha \rangle$ is not a negative integer for all positive roots α . If not, then since $\langle \rho, \alpha \rangle = 1$ for each simple root, there exists an integer $n \gg 0$ such that $\lambda + n\rho$ is dominant.

Set $\mu = \lambda + n\rho$. Then $M(\mu)$ is projective. The lowest weight of $L(n\rho)$ is $w_0(n\rho) = -n\rho$, and so $M(\mu) \otimes L(n\rho)$ is a projective object which has a quotient isomorphic to $M(\lambda)$, and hence we have a surjection $M(\mu) \otimes L(n\rho) \rightarrow L(\lambda)$.

Finally, we prove that every $M \in \mathcal{O}$ is a quotient of a projective object by induction on the length of its composition series. The base case has just been handled. Otherwise, we can find a simple submodule $L(\lambda)$ of M ; let N be the quotient module. Then N has shorter length, so there is a surjection $P \rightarrow N$ where P is projective. In particular, we can lift this to a map $P \rightarrow M$. If this is surjective, we're done. Otherwise, the image is a submodule M' such that $M' \cap L(\lambda) = 0$, which means we have a decomposition $M \cong M' \oplus L(\lambda)$. But we can find a projective object P' that surjects onto $L(\lambda)$ and hence we can write M as a quotient of $P \oplus P'$, which is again projective. \square

7.2. Projective covers. Now we know that every $M \in \mathcal{O}$ is the quotient of at least one projective object. We now want to show there is a “smallest” choice of projective object, which we will call its projective cover. A lot of this will result from some general algebra facts, but since they may not be familiar, we'll cover them here, and try to minimize technical background.

Let's start with irreducible representations. Our first claim is that there is a natural correspondence between irreducible representations and *indecomposable* projective objects.

If P is projective and $P \cong P_1 \oplus P_2$, then both P_1 and P_2 are also projective (I'll leave the check to you if you haven't seen this before). An object is **indecomposable** if it is not isomorphic to a direct sum of two nonzero modules. Otherwise, we can keep decomposing it into smaller pieces; since every object is artinian, this process always terminates, so every object is isomorphic to a finite direct sum of indecomposable modules².

Suppose that M is irreducible and P is a projective object surjecting onto M . If $P \cong P_1 \oplus P_2$, then the images of P_1 and P_2 are submodules of M that span it, and in particular, either P_1 or P_2 surjects onto M (possibly both). So we can always write irreducible representations as a quotient of an indecomposable projective. Furthermore, if P is projective, then it has at least one irreducible quotient since it has a composition series.

Now we need to refine this discussion. Before proceeding, we begin with Fitting's lemma. Note that this will equally apply to any module over a ring which is simultaneously noetherian and artinian.

Proposition 7.2.1 (Fitting's lemma). *Pick $M \in \mathcal{O}$ and a \mathfrak{g} -linear map $f: V \rightarrow V$. Then for all $n \gg 0$, we have $M = \ker(f^n) \oplus \text{image}(f^n)$. In particular, if M is indecomposable, then either f is an isomorphism or f is nilpotent.*

Proof. Since M is noetherian, the increasing chain of submodules $\ker(f) \subseteq \ker(f^2) \subseteq \dots$ must stabilize. Similarly, since M is artinian, the decreasing chain of submodules

²There is a uniqueness property here, generally falling under the name Krull–Schmidt theorem; we might not need to get into this.

$\text{image}(f) \supseteq \text{image}(f^2) \supseteq \dots$ also stabilizes. Pick m so that $\ker(f^n) = \ker(f^m)$ and $\text{image}(f^n) = \text{image}(f^m)$ for all $n \geq m$.

If $n \geq m$, we claim that $\ker(f^n) \cap \text{image}(f^n) = 0$. Pick x in this intersection. Then there exists y such that $x = f^n(y)$, which means that $f^{2n}(y) = f^n(x) = 0$. But $\ker(f^{2n}) = \ker(f^n)$, and so $f^n(y) = 0$; and hence $x = 0$. Now we claim that $\ker(f^n) + \text{image}(f^n) = M$. Pick arbitrary x . Since $\text{image}(f^n) = \text{image}(f^{2n})$, there exists z such that $f^n(x) = f^{2n}(z)$. So $x - f^n(z) \in \ker(f^n)$ and we have $x = (x - f^n(z)) + f^n(z)$. In particular, $M = \ker(f^n) \oplus \text{image}(f^n)$.

Now suppose that M is indecomposable. Then either $\ker(f^n) = 0$ or $\text{image}(f^n) = 0$. In the first case, f^n is injective, which means that f must also be injective. Furthermore, $M = \text{image}(f^n)$ implies that f^n is surjective, so in particular f is surjective. So f is an isomorphism. In the second case, $f^n = 0$, so by definition f is nilpotent. \square

Proposition 7.2.2. *Let $M \in \mathcal{O}$ be irreducible (and nonzero), and suppose we are given surjections $\pi: P \rightarrow M$ and $\pi': P' \rightarrow M$ where P, P' are indecomposable projective objects. Then $P \cong P'$.*

Proof. Since P is projective, there is a map $g: P \rightarrow P'$ such that $\pi = \pi'g$. Similarly, there is a map $g': P' \rightarrow P$ such that $\pi' = \pi g'$.

$$\begin{array}{ccc} P & & P' \\ \downarrow g & \searrow \pi & \downarrow g' \\ P' & \xrightarrow{\pi'} & M \end{array} \quad \begin{array}{ccc} P' & & P \\ \downarrow g' & \searrow \pi' & \downarrow \pi \\ P & \xrightarrow{\pi} & M \end{array}$$

Define $f = g'g: P \rightarrow P$. By Fitting's lemma, either f is an isomorphism or f is nilpotent. In the second case, suppose $f^n = 0$. Note that $\pi = \pi'g = \pi g'g = \pi f$. So then $0 = \pi f^n = \pi f^{n-1} = \dots = \pi$, which contradicts that π is surjective. Hence f is an isomorphism. If we define $f' = gg'$, then in the same way we conclude that f' is an isomorphism. This means that g and g' are isomorphisms. \square

Lemma 7.2.3. *Let $P \in \mathcal{O}$ be an indecomposable projective object and suppose we have a surjection $\pi: P \rightarrow M$ with M nonzero and irreducible. If Q is any proper submodule of P , then $\pi(Q) = 0$.*

Proof. Let $\psi: Q \rightarrow M$ be the restriction of π to Q . Since M is irreducible, ψ is either surjective or 0. Assume ψ is surjective, then since P is projective, there exists a map $g: P \rightarrow Q$ such that $\psi g = \pi$. Let $i: Q \rightarrow P$ be the inclusion map.

$$\begin{array}{ccc} P & & \\ \uparrow g & \nearrow \pi & \\ Q & \xrightarrow{\psi} & M \end{array}$$

Then $f = ig$ is a map from P to itself. So by Fitting's lemma, either f is an isomorphism or nilpotent. The first case implies that i is surjective, i.e., $Q = P$, which contradicts the assumption that Q is a proper submodule. So f is nilpotent, say $f^n = 0$. But note that $\psi = \pi i$ by definition, so

$$\pi = \psi g = \pi i g = \pi f$$

But then repeating that, we conclude that $\pi = \pi f^n = 0$, which contradicts that $M \neq 0$. So we conclude that $\psi = 0$. \square

Proposition 7.2.4. *Let $P \in \mathcal{O}$ be an indecomposable projective object and let M, M' be nonzero irreducible representations together with surjections $\pi: P \rightarrow M$ and $\pi': P \rightarrow M'$. Then $M \cong M'$.*

Proof. Apply the previous lemma to π with $Q = \ker \pi'$. Then $\pi(Q) = 0$ implies that $\ker \pi' \subseteq \ker \pi$. A symmetric argument implies that $\ker \pi \subseteq \ker \pi'$. In particular, $\ker \pi = \ker \pi'$ implies that $M \cong M'$. \square

In conclusion, the above discussion gives us a bijection between irreducible objects and indecomposable projective objects.

Definition 7.2.5. For each weight λ , we let $P(\lambda)$ denote the unique (up to isomorphism) indecomposable projective object that has $L(\lambda)$ as a quotient module. \square

When λ is dominant, we have $P(\lambda) = M(\lambda)$.

Now to extend this discussion to general objects of \mathcal{O} . Given $M \in \mathcal{O}$, and a surjection $\pi: P \rightarrow M$ where P is projective, we say that π is a **projective cover** (or by abuse of notation, that P is a projective cover) if, for every proper submodule $N \subset P$, we have $\pi(N) \neq M$.

Proposition 7.2.6. *Every object of \mathcal{O} has a projective cover, and they are unique up to isomorphism.*

Proof. We will prove this by induction on the length of a composition series. For the base case, when our module is irreducible, it follows from Lemma 7.2.3 that if $M = L(\lambda)$, then $P(\lambda) \rightarrow L(\lambda)$ is indeed a projective cover.

Otherwise, we mimic the proof that \mathcal{O} has enough projectives. Given $M \in \mathcal{O}$, we have a short exact sequence

$$0 \rightarrow L(\lambda) \rightarrow M \rightarrow N \rightarrow 0$$

for some λ . Since N has shorter length, it has a projective cover P , and the map $P \rightarrow N$ can be lifted to a map $P \rightarrow M$. If this is surjective, then this is also a projective cover. If not, then we have $M \cong L(\lambda) \oplus N$ and $P(\lambda) \oplus P$ is a projective cover of M (check). The uniqueness is left as an exercise. \square

Proposition 7.2.7. (1) *Let P be a projective object. If $P \cong \bigoplus_{\lambda} P(\lambda)^{\oplus m_{\lambda}}$, then $m_{\lambda} = \dim \operatorname{Hom}_{\mathcal{O}}(P, L(\lambda))$.*

(2) *For any $M \in \mathcal{O}$, $\dim \operatorname{Hom}_{\mathcal{O}}(P(\lambda), M) = [M : L(\lambda)]$.*

Proof. (1) clear from previous discussion

(2) Let $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ be a short exact sequence in \mathcal{O} . Since $P(\lambda)$ is projective in \mathcal{O} , the functor $\operatorname{Hom}_{\mathcal{O}}(P(\lambda), -)$ is exact and so we get a short exact sequence

$$0 \rightarrow \operatorname{Hom}_{\mathcal{O}}(P(\lambda), A) \rightarrow \operatorname{Hom}_{\mathcal{O}}(P(\lambda), B) \rightarrow \operatorname{Hom}_{\mathcal{O}}(P(\lambda), C) \rightarrow 0.$$

In particular, $[M] \mapsto \dim \operatorname{Hom}_{\mathcal{O}}(P(\lambda), M)$ gives a well-defined homomorphism $f_{\lambda}: K(\mathcal{O}) \rightarrow \mathbf{Z}$. Note that $f_{\lambda}([L(\mu)]) = \delta_{\lambda, \mu}$. On the other hand, $[M] \mapsto [M : L(\lambda)]$ also defines a well-defined homomorphism taking the same values on $[L(\mu)]$, so they must agree since $\{[L(\mu)]\}$ spans $K(\mathcal{O})$. \square

7.3. Injective objects. Recall the definition of injective object: $I \in \mathcal{O}$ is **injective** if the functor $\text{Hom}_{\mathcal{O}}(-, I)$ is exact, i.e., given any injective homomorphism $i: M \rightarrow N$ of objects in \mathcal{O} together with a homomorphism $\varphi: M \rightarrow I$, there exists $\psi: N \rightarrow I$ such that $\psi \circ i = \varphi$. Note we can think of M as a submodule of N , so we're saying that homomorphisms from submodules can always be extended to the whole module. As with projective objects, there is a distinction between being an injective $U(\mathfrak{g})$ -module and being an injective object of \mathcal{O} , but we won't have a use for the former notion.

As before, we'll say that \mathcal{O} has **enough injectives** if, for all $M \in \mathcal{O}$, there exists an injective object I and an injective homomorphism $i: M \rightarrow I$. There is a pleasant shortcut here: we already know some facts about projective objects, and the definitions for injective and projective are the same if we reverse the direction of all arrows. We also have a duality \vee on \mathcal{O} , so we can translate statements freely.

So we just state them:

Proposition 7.3.1. (1) \mathcal{O} has enough injectives.

(2) For each weight λ , there is a unique (up to isomorphism) indecomposable injective object, which we will denote by $I(\lambda)$, that has $L(\lambda)$ as a submodule.

We can take $I(\lambda) = P(\lambda)^\vee$.

The dual notion of a projective cover is an injective envelope: given an injective object I and $M \in \mathcal{O}$ together with an injection $i: M \rightarrow I$, we call I an **injective envelope** if, for all nonzero submodules N of I , we have $M \cap N \neq 0$.

Proposition 7.3.2. Every object of \mathcal{O} has an injective envelope and it is unique up to isomorphism.

7.4. Standard filtrations. Recall that we define a standard filtration of $M \in \mathcal{O}$ to be a filtration whose successive quotients are Verma modules. We saw that the tensor product of a Verma module with a finite-dimensional module has a standard filtration.

Proposition 7.4.1. Suppose M has a standard filtration.

- (1) If λ is a maximal weight of M , then $\text{Hom}_{\mathcal{O}}(M(\lambda), M) \neq 0$, every nonzero homomorphism $\varphi: M(\lambda) \rightarrow M$ is injective, and $\text{coker } \varphi$ has a standard filtration.
 (2) If $M \cong M' \oplus M''$, then both M' and M'' have standard filtrations.

Proof. (1) Pick nonzero $v \in M_\lambda$. Then v is a highest weight vector and hence gives a nonzero map $\varphi: M(\lambda) \rightarrow M$. Let $M^0 \subset M^1 \subset \dots$ be a standard filtration of M , and pick i minimal so that $\text{image } \varphi \subseteq M^i$. In particular, the map $M(\lambda) \rightarrow M^i/M^{i-1}$ is nonzero, but on the other hand, $M^i/M^{i-1} \cong M(\mu)$ for some weight μ . Then $\lambda \leq \mu$, and so maximality of λ forces $\lambda = \mu$. But every nonzero endomorphism of $M(\lambda)$ is an isomorphism, so we conclude that φ is injective.

This also implies that $\text{image } \varphi \cap M^{i-1} = 0$ and we have a short exact sequence

$$0 \rightarrow M^{i-1} \rightarrow \text{coker } \varphi \rightarrow M/M^i \rightarrow 0.$$

By definition, both M^{i-1} and M/M^i have standard filtrations, so we can combine them to get one on $\text{coker } \varphi$ as well.

(2) We prove this by induction on the length of a standard filtration on M . The base case is that M is a Verma module; hence M is indecomposable, so there is nothing to say. Otherwise, let λ be a maximal weight of M . Then either $M'_\lambda \neq 0$ or $M''_\lambda \neq 0$; without loss of generality, suppose $M'_\lambda \neq 0$. Pick nonzero $v \in M'_\lambda$, which induces a map $\varphi: M(\lambda) \rightarrow M'$.

The composition $M(\lambda) \rightarrow M' \rightarrow M$ is injective by (1) and hence φ is also injective. From (1), we know that $M/M(\lambda) = M'/M(\lambda) \oplus M''$ has a standard filtration; by induction $M'/M(\lambda)$ and M'' both have standard filtrations, so we're done. \square

If M has a standard filtration, then its restriction to $U(\mathfrak{n}^-)$ is free. In particular, the number of quotients that are isomorphic to $M(\lambda)$ is well-defined since this is the dimension of the space of free generators of weight λ , and we denote this number by $(M : M(\lambda))$.

As we discussed before, the classes of simple objects $[L(\lambda)]$ form a basis for $K(\mathcal{O})$ as we vary over all weights of λ . We also know that

$$[M(\lambda)] = [L(\lambda)] + \sum_{\mu < \lambda} c_{\mu}^{\lambda} [L(\mu)]$$

for some set of non-negative integers c_{μ}^{λ} , which implies that the $[M(\lambda)]$ also forms a basis. If M has a standard filtration, then

$$[M] = \sum_{\lambda} (M : M(\lambda)) [M(\lambda)].$$

Proposition 7.4.2. *If M has a standard filtration, then*

$$(M : M(\lambda)) = \dim \operatorname{Hom}_{\mathcal{O}}(M, M(\lambda)^{\vee}).$$

Proof. We prove this by induction on the length of a standard filtration. For the base case, with $M = M(\mu)$, we have

$$(M(\mu) : M(\lambda)) = \delta_{\lambda, \mu} = \dim \operatorname{Hom}_{\mathcal{O}}(M(\mu), M(\lambda)^{\vee})$$

where the second equality comes from Proposition 6.5.5.

In general, we have a short exact sequence of the form

$$0 \rightarrow N \rightarrow M \rightarrow M(\mu) \rightarrow 0$$

where N has a shorter standard filtration than M . Since $[M] = [M(\mu)] + [N]$ in $K(\mathcal{O})$, we have

$$\begin{aligned} (M : M(\lambda)) &= (M(\mu) : M(\lambda)) + (N : M(\lambda)) \\ &= \dim \operatorname{Hom}_{\mathcal{O}}(M(\mu), M(\lambda)^{\vee}) + \dim \operatorname{Hom}_{\mathcal{O}}(N, M(\lambda)^{\vee}). \end{aligned}$$

On the other hand, we also have an exact sequence

$$0 \rightarrow \operatorname{Hom}_{\mathcal{O}}(M(\mu), M(\lambda)^{\vee}) \rightarrow \operatorname{Hom}_{\mathcal{O}}(M, M(\lambda)^{\vee}) \rightarrow \operatorname{Hom}_{\mathcal{O}}(N, M(\lambda)^{\vee}) \rightarrow \operatorname{Ext}_{\mathcal{O}}^1(M(\mu), M(\lambda)^{\vee})$$

The last term is 0 by Proposition 6.5.5, so we conclude that

$$\dim \operatorname{Hom}_{\mathcal{O}}(M, M(\lambda)^{\vee}) = \dim \operatorname{Hom}_{\mathcal{O}}(M(\mu), M(\lambda)^{\vee}) + \dim \operatorname{Hom}_{\mathcal{O}}(N, M(\lambda)^{\vee}). \quad \square$$

Theorem 7.4.3. (1) *If $P \in \mathcal{O}$ is projective, then P has a standard filtration.*

(2) *(BGG reciprocity) Given weights λ, μ , we have*

$$(P(\lambda) : M(\mu)) = [M(\mu) : L(\lambda)].$$

(3) *If $(P(\lambda) : M(\mu)) \neq 0$, then $\mu \geq \lambda$ and μ and λ are in the same dotted W -orbit. Furthermore, $(P(\lambda) : M(\lambda)) = 1$.*

(4) *In particular, $\{[P(\lambda)]\}$ is a basis for $K(\mathcal{O})$, and if P and P' are projectives with the same character, then $P \cong P'$.*

Proof. (1) We can write P as a direct sum of $P(\lambda)$; it suffices to prove the result for $P(\lambda)$. Pick n so that $\nu = \lambda + n\rho$ is dominant. As we saw before, $M(\nu) \otimes L(n\rho)$ is projective and has $L(\lambda)$ as a quotient. In particular, $P(\lambda)$ is a direct summand of $M(\nu) \otimes L(n\rho)$; we also know that $M(\nu) \otimes L(n\rho)$ has a standard filtration (Proposition 7.1.4), so by the previous result, $P(\lambda)$ does too.

(2) We have

$$(P(\lambda) : M(\mu)) = \dim \operatorname{Hom}_{\mathcal{O}}(P(\lambda), M(\mu)^\vee) = [M(\mu)^\vee : L(\lambda)] = [M(\mu) : L(\lambda)]$$

where the last equality follows since $M(\mu)$ and $M(\mu)^\vee$ have the same character.

(3) If the above quantity is nonzero, then $\lambda \leq \mu$ and λ and μ belong to the same dotted W -orbit; it is 1 when $\lambda = \mu$.

(3) The previous result tells us that we have

$$[P(\lambda)] = [M(\lambda)] + \sum_{\mu > \lambda} d_\mu^\lambda [M(\mu)]$$

for some set of coefficients d_μ^λ . Furthermore, we can restrict to weights belonging to the same dotted W -action (in particular a finite set). Since the $[M(\lambda)]$ form a basis, the result follows since we have a unitriangular change of bases. The last statement follows since every projective P is isomorphic to a direct sum of the $P(\lambda)$. \square

Example 7.4.4. Consider the case of \mathfrak{sl}_2 . Our positive root is $\alpha = 2$ and $\rho = 1$, so $\langle \lambda + \rho, \alpha \rangle = 2\lambda + 2$. So dominance means $\lambda \notin \{-2, -3, \dots\}$.

Recall that if $\lambda = -1$ or $\lambda \notin \mathbf{Z}$, then the block containing $L(\lambda)$ just contains one irreducible, and it is dominant and $L(\lambda) = M(\lambda) = P(\lambda)$.

Otherwise, if $\lambda \in \mathbf{Z}_{\geq 0}$, we have a block consisting of $L(\lambda)$ and $L(\mu)$ where $\mu = -\lambda - 2$. By the above comment, $P(\lambda) = M(\lambda)$. What is the structure of $P(\mu)$? We know that it has a standard filtration and from BGG reciprocity, we get

$$(P(\mu) : M(\mu)) = 1, \quad (P(\mu) : M(\lambda)) = [M(\lambda) : L(\mu)] = 1$$

since $[M(\lambda)] = [L(\lambda)] + [L(\mu)]$. Since $P(\mu)$ is indecomposable and has $M(\mu)$ as a quotient, it must be an extension of the form

$$0 \rightarrow M(\lambda) \rightarrow P(\mu) \rightarrow M(\mu) \rightarrow 0. \quad \square$$

Remark 7.4.5. Fix a central character $\chi = \chi_\lambda$ and let $P = \bigoplus_{w \in W} P(w \bullet \lambda)$. Then P is known as a *projective generator* for \mathcal{O}_χ since every module is a quotient of $P^{\oplus n}$ for some n . Also, $A = \operatorname{Hom}_{\mathcal{O}}(P, P)$ is a finite-dimensional algebra, and one can show that \mathcal{O}_χ is equivalent to the category of finite-dimensional right A -modules. There are some advantages to this perspective since the representation theory of finite-dimensional algebras is fairly well-developed, especially if A can be computed explicitly. \square

8. BGG RESOLUTION

For $M \in \mathcal{O}$, the linear combination in $K(\mathcal{O})$ given by $[M] = \sum_\lambda c_\lambda [L(\lambda)]$ has a concrete meaning: c_λ is the multiplicity of $L(\lambda)$ in any composition series of M . This makes sense since the coefficients c_λ must be a non-negative integer. However, we've already seen that when we invert linear combinations, we get negative coefficients, such as writing $[L(\lambda)]$ in terms of $[M(\mu)]$. Negative coefficients can also be interpreted naturally in terms of *Euler characteristics*.

8.1. General discussion of resolutions. Generally speaking, given a module M , we can informally call any chain complex

$$\mathbf{F}_\bullet : \cdots \rightarrow \mathbf{F}_i \rightarrow \mathbf{F}_{i-1} \rightarrow \cdots \rightarrow \mathbf{F}_0$$

a “(left) resolution” of M if $H_0(\mathbf{F}_\bullet) = M$ and $H_i(\mathbf{F}_\bullet) = 0$ for $i > 0$. Usually, we want to require that the \mathbf{F}_i satisfy some special property. A common example is requiring that the \mathbf{F}_i are free modules, in which case this is called a **free resolution**. Another common example is requiring that the \mathbf{F}_i are projective, in which case this is called a **projective resolution**.

Heuristically, we can think of a resolution as an approximation of a general module M by modules of a special class. More practically, projective resolutions can be used to compute right derived functors, such as Tor .

There is yet another angle. If the resolution has finitely many terms, then in the Grothendieck group, we get an equation

$$[M] = \sum_i (-1)^i [\mathbf{F}_i].$$

If we’re talking about \mathcal{O} , then this implies an equation for formal characters

$$\text{ch}_M = \sum_i (-1)^i \text{ch}_{\mathbf{F}_i}.$$

In category \mathcal{O} we don’t have free modules, and projective modules have complicated characters (as we saw from BGG reciprocity). We’ll be interested in the case when the \mathbf{F}_i have standard filtrations. I’m not sure if this is typical terminology, but we’ll call that a **standard resolution**. This exists when M is finite-dimensional and we’ll be able to recover the well-known Weyl character formula. This is perhaps a strange way to deduce it, but it’ll give us an excuse to discuss some other pieces of algebra that we haven’t yet seen.

8.2. (Relative) Chevalley–Eilenberg complex. For the purposes of this section, let \mathfrak{g} be any complex Lie algebra (other fields are fine too). For $k = 0, \dots, \dim \mathfrak{g}$ (though it’s not necessary to assume that $\dim \mathfrak{g}$ is finite), we define free $U(\mathfrak{g})$ -modules

$$\mathbf{K}_k = U(\mathfrak{g}) \otimes_{\mathbb{C}} \bigwedge^k(\mathfrak{g})$$

and $U(\mathfrak{g})$ -linear maps $\varphi_k, \psi_k : \mathbf{K}_k \rightarrow \mathbf{K}_{k-1}$ for each $k \geq 1$ ($f \in U(\mathfrak{g})$, $x_i \in \mathfrak{g}$):

$$\begin{aligned} \varphi_k(f \otimes x_1 \wedge \cdots \wedge x_k) &= \sum_{i=1}^k (-1)^{i-1} f x_i \otimes x_1 \wedge \cdots \widehat{x}_i \cdots \wedge x_k, \\ \psi_k(f \otimes x_1 \wedge \cdots \wedge x_k) &= \sum_{1 \leq i < j \leq k} (-1)^{i+j} f \otimes [x_i, x_j] \wedge x_1 \wedge \cdots \widehat{x}_i \cdots \widehat{x}_j \cdots \wedge x_k, \end{aligned}$$

where a hat over an item means that it has been removed.

Now set $\partial_k = \varphi_k + \psi_k$. It can be shown (though we will not do it) that $\partial_{k-1} \partial_k = 0$ for all $k \geq 2$ so that we get a chain complex

$$\mathbf{K}_\bullet : \cdots \rightarrow \mathbf{K}_k \xrightarrow{\partial_k} \mathbf{K}_{k-1} \xrightarrow{\partial_{k-1}} \mathbf{K}_{k-2} \rightarrow \cdots \rightarrow \mathbf{K}_0,$$

called the **Chevalley–Eilenberg complex** of \mathfrak{g} (sometimes also the **Koszul complex**).

For concreteness, here is ∂_k for $k = 1, 2$:

$$\partial_1(f \otimes x) = fx, \quad \partial_2(f \otimes x_1 \wedge x_2) = fx_1 \otimes x_2 - fx_2 \otimes x_1 - f \otimes [x_1, x_2].$$

It follows immediately from the description of ∂_1 that $H_0(\mathbf{K}_\bullet) \cong \mathbf{C}$.

Theorem 8.2.1 (Koszul). *For $i \geq 1$, $H_i(\mathbf{K}_\bullet) = 0$.*

Proof. Just a sketch.

Pick an ordered basis $\{e_i\}_{i \in I}$ for \mathfrak{g} and consider the corresponding PBW basis for $U(\mathfrak{g})$. For each p and k , let $F_p \mathbf{K}_k$ be the subspace of \mathbf{K}_k spanned by the elements $e_{i_1} \cdots e_{i_n} \otimes e_{j_1} \wedge \cdots \wedge e_{j_k}$ where $i_1 \leq \cdots \leq i_n$, and $j_1 < \cdots < j_k$, and $n + k \leq p$. From the definition of ∂ , it follows that $\partial_k(F_p \mathbf{K}_k) \subseteq F_p \mathbf{K}_{k-1}$. In particular, for each p , we get a subcomplex $F_p \mathbf{K}_\bullet$. General homological algebra gives an upper bound

$$\dim H_i(\mathbf{K}_\bullet) \leq \sum_p \dim H_i(F_p \mathbf{K}_\bullet / F_{p-1} \mathbf{K}_\bullet).$$

The advantage of the quotient complexes on the right is that $\psi_k(F_p \mathbf{K}_k) \subseteq F_{p-1} \mathbf{K}_{k-1}$, and hence $\bigoplus_p F_p \mathbf{K}_\bullet / F_{p-1} \mathbf{K}_\bullet$ is isomorphic to the Koszul complex on \mathfrak{g} where \mathfrak{g} is now thought of as an *abelian* Lie algebra. In this case, there are direct methods to show that $H_i = 0$ for $i \geq 1$, and I will outline some in the exercises. \square

Remark 8.2.2. The Chevalley–Eilenberg complex is a free resolution of the trivial representation \mathbf{C} . But, when \mathfrak{g} is semisimple, free $U(\mathfrak{g})$ -modules do not belong to \mathcal{O} , so it will be better to construct something “smaller”. \square

We now want to generalize the above construction to a “relative” setting. Let $\mathfrak{b} \subseteq \mathfrak{g}$ be a Lie subalgebra. Then \mathfrak{b} acts on $\mathfrak{g}/\mathfrak{b}$ and hence also on its exterior powers. Now define

$$\mathbf{K}(\mathfrak{g}, \mathfrak{b})_k := U(\mathfrak{g}) \otimes_{U(\mathfrak{b})} \bigwedge^k (\mathfrak{g}/\mathfrak{b}).$$

We can generalize the above formulas to define $U(\mathfrak{g})$ -linear maps $\varphi_k, \psi_k: \mathbf{K}(\mathfrak{g}, \mathfrak{b})_k \rightarrow \mathbf{K}(\mathfrak{g}, \mathfrak{b})_{k-1}$ (for $x \in \mathfrak{g}$, let \bar{x} denote its coset in $\mathfrak{g}/\mathfrak{b}$):

$$\begin{aligned} \varphi_k(f \otimes \bar{x}_1 \wedge \cdots \wedge \bar{x}_k) &= \sum_{i=1}^k (-1)^{i-1} f x_i \otimes \bar{x}_1 \wedge \cdots \wedge \widehat{x}_i \cdots \wedge \bar{x}_k, \\ \psi_k(f \otimes \bar{x}_1 \wedge \cdots \wedge \bar{x}_k) &= \sum_{1 \leq i < j \leq k} (-1)^{i+j} f \otimes \overline{[x_i, x_j]} \wedge \bar{x}_1 \wedge \cdots \wedge \widehat{x}_i \cdots \wedge \widehat{x}_j \cdots \wedge \bar{x}_k, \end{aligned}$$

and define $\partial_k = \varphi_k + \psi_k$. It needs to be shown that these maps are well-defined (i.e., independent of choices of coset representatives) but we will omit this somewhat tedious, but straightforward, task. As before, one can check that $\partial_k \partial_{k-1} = 0$ so that we get a chain complex $\mathbf{K}(\mathfrak{g}, \mathfrak{b})_\bullet$ (the **relative Chevalley–Eilenberg complex**). Using a variation of the above argument, it can also be shown that its homology vanishes in positive degrees (and is \mathbf{C} in degree 0).

8.3. Weights in the relative Chevalley–Eilenberg complex. As the notation suggests, we will apply the previous construction in the case that \mathfrak{g} is a semisimple complex Lie algebra and \mathfrak{b} is a Borel subalgebra. In that case we have

$$\mathbf{K}(\mathfrak{g}, \mathfrak{b})_k = \text{Ind}_{\mathfrak{b}}^{\mathfrak{g}} \left(\bigwedge^k (\mathfrak{g}/\mathfrak{b}) \right).$$

We recall an earlier discussion. Let μ_1, \dots, μ_N be the weights of $\bigwedge^k(\mathfrak{g}/\mathfrak{b})$ (repeated to account for multiplicities). Then $\mathbf{K}(\mathfrak{g}, \mathfrak{b})_k$ has a standard filtration whose quotient modules are $M(\mu_1), \dots, M(\mu_N)$.

The μ_i have an explicit description. First, the weights of $\mathfrak{g}/\mathfrak{b}$ are the negative roots Φ^- (each appearing with multiplicity 1). In particular, the weights of $\bigwedge^k(\mathfrak{g}/\mathfrak{b})$ is the multiset $\alpha_1 + \dots + \alpha_k$ where $\alpha_i \in \Phi^-$ and the α_i are all distinct.

Example 8.3.1. If $\mathfrak{g} = \mathfrak{sl}_3$, then $\Phi^- = \{(-1, 1, 0), (-1, 0, 1), (0, -1, 1)\}$. So the weights for $k = 0, 1, 2, 3$ are

k	
0	(0, 0, 0)
1	(-1, 1, 0), (-1, 0, 1), (0, -1, 1)
2	(-2, 1, 1), (-1, 0, 1), (-1, -1, 2)
3	(-2, 0, 2)

□

Let's examine the Euler characteristic of this complex, or more precisely its value in $K(\mathcal{O})$. By general principles, we have (setting $N = \dim(\mathfrak{g}/\mathfrak{b})$):

$$\sum_{k=0}^N (-1)^k [H_k(\mathbf{K}(\mathfrak{g}, \mathfrak{b})_\bullet)] = \sum_{k=0}^N (-1)^k [\text{Ind}_{\mathfrak{b}}^{\mathfrak{g}}(\bigwedge^k(\mathfrak{g}/\mathfrak{b}))].$$

The left hand side is just $[\mathbf{C}]$, where \mathbf{C} is the trivial representation by what we said above. Furthermore, each term on the right can be expanded into Verma modules using the filtration above, so we get

$$[\mathbf{C}] = \sum_{S \subseteq \Phi^-} (-1)^{|S|} [M(\sum_{\alpha \in S} \alpha)].$$

Taking the formal character of both sides gives (with the identification $e^0 = 1$)

$$1 = \sum_{S \subseteq \Phi^-} (-1)^{|S|} \frac{\prod_{\alpha \in S} e^\alpha}{\prod_{\alpha \in \Phi^-} (1 - e^\alpha)}$$

which, by clearing denominators, simply gives us the (tautological) identity

$$\prod_{\alpha \in \Phi^-} (1 - e^\alpha) = \sum_{S \subseteq \Phi^-} (-1)^{|S|} \prod_{\alpha \in S} e^\alpha.$$

Example 8.3.2. In the previous \mathfrak{sl}_3 example, the weight $(-1, 0, 1)$ appears in degrees $k = 1$ and $k = 2$, so that in the alternating sum above, the terms cancel. So rather than write it as a sum of 8 terms, we just need 6. The remaining 6 terms are all actually of the form $w \bullet 0$ where w ranges over all permutations. Recall that $\rho = (1, 0, -1)$. Here's a table where $w \in \mathfrak{S}_3$ is written as $w(1)w(2)w(3)$:

w	$\ell(w)$	$w \bullet 0$
123	0	(0, 0, 0)
213	1	(-1, 1, 0)
132	1	(0, -1, 1)
231	2	(-2, 1, 1)
312	2	(-1, -1, 2)
321	3	(-2, 0, 2)

Here I've also listed the length $\ell(w)$ for each permutation w to show that this matches up with the homological degree k listed in the previous table. \square

This suggests that we try to “trim” the complex to get a more interesting identity (and also a smaller complex). Concretely, for any central character χ , we can project onto the subcategory \mathcal{O}_χ . Since this is an exact functor which either preserves a Verma module or sends it to 0, the result will be another chain complex whose terms have standard filtrations. Exactness also says that projection commutes with taking homology; so if we project to the principal block \mathcal{O}_{χ_0} , the 0th homology of our complex will still be \mathbf{C} . This is what we'll be interested in doing.

More explicitly, $M(\mu)$ will go to 0 if and only if $\mu \notin W \bullet 0$, so we want to determine which weights of $\bigwedge^\bullet(\mathfrak{g}/\mathfrak{b})$ belong to $W \bullet 0$.

First, let's make an important observation. For $w \in W$, define $\Pi_w = \Phi^- \cap w(\Phi^+)$, and for any subset $S \subseteq \Phi^-$, define $\bar{S} = \sum_{\alpha \in S} \alpha$.

Lemma 8.3.3. *For $w \in W$, we have $\bar{\Pi}_w = w \bullet 0$.*

Proof. We have

$$w \bullet 0 = \frac{1}{2} \sum_{\alpha \in \Phi^+} w(\alpha) - \frac{1}{2} \sum_{\alpha \in \Phi^+} \alpha.$$

Hence for $\alpha \in \Phi^+$, if $w(\alpha) \in \Phi^+$, then it gets cancelled when combining the above sums, while if $w(\alpha) \in \Phi^-$, we get a contribution of $w(\alpha)$. \square

Next, the subsets are all distinct: if $\Pi_w = \Pi_u$, then $w^{-1}u$ preserves all positive roots (hence has length 0) and must be the identity. We want to show something stronger: if $\bar{\Pi} = \bar{\Pi}_w$, then $\Pi = \Pi_w$. First, we develop a preliminary step.

Recall that the simple reflections are denoted σ_α and that $\ell(w)$ is the least r such that w can be written as a product of r simple reflections. Recall also that

$$\ell(w) = |\Phi^- \cap w(\Phi^+)| = |\Pi_w|.$$

Lemma 8.3.4. *Let α be a simple root and set $w' = \sigma_\alpha w$. Then $\ell(w) = \ell(w') \pm 1$ and*

- $\ell(w) = \ell(w') + 1$ if and only if $-\alpha \notin \Pi_{w'}$ if and only if $\Pi_w = \sigma_\alpha(\Pi_{w'}) \cup \{-\alpha\}$, and
- $\ell(w) = \ell(w') - 1$ if and only if $-\alpha \in \Pi_{w'}$ if and only if $\Pi_w = \sigma_\alpha(\Pi_{w'}) \setminus \{-\alpha\}$.

Proof. First, recall that $\sigma_\alpha(\Phi^- \setminus \{-\alpha\}) = \Phi^- \setminus \{-\alpha\}$ and $\sigma_\alpha(-\alpha) = \alpha$.

This implies that

$$\Pi_w = \begin{cases} \sigma_\alpha(\Pi_{w'}) \cup \{-\alpha\} & \text{if } -\alpha \notin \Pi_{w'} \\ \sigma_\alpha(\Pi_{w'}) \setminus \{-\alpha\} & \text{if } -\alpha \in \Pi_{w'} \end{cases}.$$

In the first case, we see that $\ell(w) = \ell(w') + 1$ and in the second case we have $\ell(w) = \ell(w') - 1$. \square

In particular, there exists a simple root α such that $\ell(\sigma_\alpha w) < \ell(w)$ (pick a minimal expression for w as a product of simple reflections, and take σ_α to be the leftmost term).

Proposition 8.3.5. *Given $w \in W$, if $\Pi \subseteq \Pi^-$ and $\bar{\Pi} = \bar{\Pi}_w$, then $\Pi = \Pi_w$.*

Proof. We prove the above statement by induction on $\ell(w)$. If $\ell(w) = 0$, then $\Pi_w = \emptyset$, and the statement follows from the fact that all elements of Π^- are non-positive integer combinations of the simple roots (which are linearly independent).

In general, suppose that $\bar{\Pi} = \bar{\Pi}_w$ and pick a simple root α such that $\ell(w') = \ell(w) - 1$ where $w' = \sigma_\alpha w$. We claim that $-\alpha \in \Pi$. If not, then $\sigma_\alpha(\Pi) \subseteq \Phi^-$, and we have

$$\begin{aligned} \overline{\sigma_\alpha \Pi \cup \{-\alpha\}} &= \sigma_\alpha \bar{\Pi} - \alpha \\ &= \sigma_\alpha \bar{\Pi}_w - \alpha \\ &= \sigma_\alpha w(\rho) - \sigma_\alpha(\rho) - \alpha \\ &= w' \bullet 0 = \bar{\Pi}_{w'}. \end{aligned}$$

Since $\ell(w') < \ell(w)$, we use induction to conclude that $\Pi_{w'} = \sigma_\alpha \Pi \cup \{-\alpha\}$, but the previous result says that $-\alpha \notin \Pi_{w'}$, so we get a contradiction.

In particular, $-\alpha \in \Pi$. Set $\Pi' = \sigma_\alpha(\Pi \setminus \{-\alpha\})$, so that

$$\bar{\Pi}' = \sigma_\alpha(\bar{\Pi} + \alpha) = \sigma_\alpha(\bar{\Pi}_w) - \alpha = w' \bullet 0 = \bar{\Pi}_{w'}.$$

Again, by induction, we conclude that $\Pi' = \Pi_{w'}$ and hence $\Pi = \Pi_w$ by the previous lemma. \square

The upshot is that if we project to the principal block, then we get a complex \mathbf{F}_\bullet such that \mathbf{F}_i has a standard filtration by Verma modules $M(w \bullet 0)$ where $\ell(w) = i$ since the only subsets $\Pi \subset \Phi^-$ that contribute are those of the form $\bar{\Pi}_w$. Its homology is still \mathbf{C} in degree 0 and 0 in positive degrees. This gives us the following identity:

Theorem 8.3.6 (Weyl denominator formula).

$$\prod_{\alpha \in \Phi^-} (1 - e^\alpha) = \sum_{w \in W} (-1)^{\ell(w)} e^{w \bullet 0}.$$

Proof. Taking the Euler characteristic of \mathbf{F}_\bullet (its class in $K(\mathcal{O})$), we get

$$[\mathbf{C}] = \sum_{k=0}^N (-1)^k [\mathbf{F}_k].$$

But also, we have $[\mathbf{F}_k] = \sum_{\substack{w \in W \\ \ell(w)=k}} [M(w \bullet 0)]$. The result follows by taking formal characters and multiplying both sides by $\prod_{\alpha \in \Phi^-} (1 - e^\alpha)$. \square

Example 8.3.7. Consider the case $\mathfrak{g} = \mathfrak{sl}_n$. In that case, we identified \mathfrak{h}^* with the space $\{(a_1, \dots, a_n) \in \mathbf{C}^n \mid a_1 + \dots + a_n = 0\}$. Let's think of $e^{(a_1, \dots, a_n)}$ as a Laurent monomial $x_1^{a_1} \cdots x_n^{a_n}$, where x_1, \dots, x_n are some auxiliary variables. Then the product $\prod_{\alpha \in \Phi^-} (1 - e^\alpha)$ becomes the product

$$\prod_{1 \leq i < j \leq n} \left(1 - \frac{x_j}{x_i}\right).$$

Also, W is the n th symmetric group \mathfrak{S}_n and $(-1)^{\ell(w)}$ is the usual sign $\text{sgn}(w)$ (1 for an even permutation and -1 for an odd permutation). The sum $\sum_{w \in \mathfrak{S}_n} \text{sgn}(w) e^{w \bullet 0}$ can be interpreted as a determinant. More specifically, note that the i th component of $w \bullet 0$ is $i - w^{-1}(i)$ so that

$$e^{w \bullet 0} = x_1^{1-w^{-1}(1)} \cdots x_n^{n-w^{-1}(n)}.$$

Hence the sum can be written as (and simplified using that $\text{sgn}(w) = \text{sgn}(w^{-1})$)

$$\sum_{w \in \mathfrak{S}_n} \text{sgn}(w) x_1^{1-w^{-1}(1)} \cdots x_n^{n-w^{-1}(n)} = \sum_{w \in \mathfrak{S}_n} \text{sgn}(w) x_1^{1-w(1)} \cdots x_n^{n-w(n)}.$$

Using that the determinant of a matrix is a sum over permutations, we can write this more compactly as

$$\det(x_i^{i-j})_{i,j=1,\dots,n}.$$

For example, when $n = 3$, we get the matrix $\begin{pmatrix} 1 & x_1^{-1} & x_1^{-2} \\ x_2 & 1 & x_2^{-1} \\ x_3^2 & x_3 & 1 \end{pmatrix}$. In particular, we get the identity

$$\prod_{1 \leq i < j \leq n} \left(1 - \frac{x_j}{x_i}\right) = \det(x_i^{i-j}).$$

This is actually just one form of the Vandermonde identity. To see that, multiply both sides by $x_1^{n-1}x_2^{n-2}\cdots x_{n-1}$ to get

$$\prod_{1 \leq i < j \leq n} (x_i - x_j) = \det(x_i^{n-j}).$$

The Weyl denominator formula for the special orthogonal and symplectic Lie algebras also have interpretations like this, but I'll leave it to you to work out.

One can also consider certain infinite-dimensional analogues of semisimple Lie algebras (Kac–Moody algebras) which also have their own version of the Weyl denominator formula. Other interesting identities can be realized this way, such as the Jacobi triple product formula. \square

What about other finite-dimensional $L(\lambda)$? There's an obvious thing we can try to do. The tensor product $\mathbf{F}_\bullet \otimes L(\lambda)$ is a resolution of $L(\lambda)$ and we know that each $\mathbf{F}_i \otimes L(\lambda)$ has a standard filtration. However, just like the relative Chevalley–Eilenberg complex $\mathbf{K}(\mathfrak{g}, \mathfrak{b})_\bullet$, it will have a lot of terms that cancel when we consider the Euler characteristic. Again, we can try to project to the block $\mathcal{O}_{\chi_\lambda}$. This leads to the general notion of translation functors, which we discuss (unfortunately without much detail) in a later section.

But for now, let's finish this case. This requires one fact: if $L(\lambda)$ is finite-dimensional, then its set of weights is invariant under W .

First, we know that each term $\mathbf{F}_i \otimes L(\lambda)$ has a standard filtration whose quotients are of the form $M(w \bullet 0 + \mu)$ where $\ell(w) = i$ and μ is a weight of $L(\lambda)$ (counted with multiplicity). We need to determine when this is of the form $w' \bullet \lambda$.

Proposition 8.3.8. *If μ is a weight of $L(\lambda)$, then $w \bullet 0 + \mu \in W \bullet \lambda$ if and only if $\mu = w\lambda$. In that case, we have $w \bullet 0 + \mu = w \bullet \lambda$.*

Proof. Suppose that $w \bullet 0 + \mu = u \bullet \lambda$ for some $u \in W$. Then

$$w\rho - \rho + \mu = u\lambda + u\rho - \rho,$$

which we can rewrite as

$$\lambda + \rho = u^{-1}\lambda + u^{-1}w\rho.$$

Since $L(\lambda)$ is finite-dimensional, $u^{-1}\lambda$ is a weight of it, and so $u^{-1}\lambda \leq \lambda$, i.e., $u^{-1}\lambda = \lambda - v$ where v is a nonnegative sum of simple roots; similarly, $L(\rho)$ is finite-dimensional, and so $u^{-1}w\rho = \rho - v'$ where v' is a nonnegative sum of simple roots. But then the equation above forces $v + v' = 0$, i.e., $v = v' = 0$, and so $u^{-1}w\rho = \rho$. Finally, the stabilizer of W on ρ is trivial, and so $u = w$.

In that case, we have $w \bullet 0 + \mu = w\rho - \rho + w\lambda = w \bullet \lambda$. \square

In conclusion, let \mathbf{F}_i^λ be the projection of $\mathbf{F}_i \otimes L(\lambda)$ to $\mathcal{O}_{\chi_\lambda}$. Then \mathbf{F}_i^λ has a standard filtration whose quotients are exactly of the form $M(w \bullet \lambda)$ for $\ell(w) = i$.

Theorem 8.3.9 (Weyl character formula). *If $L(\lambda)$ is finite-dimensional, i.e., λ is an integral dominant weight, then*

$$\mathrm{ch}_{L(\lambda)} = \frac{\sum_{w \in W} (-1)^{\ell(w)} e^{w \bullet \lambda}}{\sum_{w \in W} (-1)^{\ell(w)} e^{w \bullet 0}}.$$

Proof. The above discussion gives the following formula in $K(\mathcal{O})$:

$$[L(\lambda)] = \sum_{w \in W} (-1)^{\ell(w)} [M(w \bullet \lambda)].$$

So we can take formal characters of both sides and use the formulas we've already established (we're using the Weyl denominator formula). \square

Example 8.3.10. Going back to \mathfrak{sl}_n , we previously interpreted the denominator as a determinant, and we can do similarly for the numerator. Using the notation from before, we have

$$e^{w \bullet \lambda} = x_1^{\lambda_{w^{-1}(1)} + 1 - w^{-1}(1)} \cdots x_n^{\lambda_{w^{-1}(n)} + 1 - w^{-1}(n)}$$

and so

$$\sum_{w \in W} \mathrm{sgn}(w) e^{w \bullet \lambda} = \det(x_i^{\lambda_j + i - j})_{i,j=1,\dots,n}.$$

This gives the following ratio of determinants:

$$\mathrm{ch}_{L(\lambda)} = \frac{\det(x_i^{\lambda_j + i - j})_{i,j=1,\dots,n}}{\det(x_i^{i-j})_{i,j=1,\dots,n}} = \frac{\det(x_i^{\lambda_j + n - j})_{i,j=1,\dots,n}}{\det(x_i^{n-j})_{i,j=1,\dots,n}}.$$

For the second equality, we multiplied both determinants by $x_1^{n-1} x_2^{n-2} \cdots x_{n-1}$. The last expression is the classical formula for Schur polynomials $s_\lambda(x_1, \dots, x_n)$.

For the special orthogonal and symplectic Lie algebras, you can similarly turn the Weyl character formula into a ratio of determinants. \square

8.4. Bott's theorem. Recall that the first complex we considered, the Chevalley–Eilenberg complex \mathbf{K}_\bullet , gives a resolution of \mathbf{C} consisting of free $U(\mathfrak{g})$ -modules. This can be used to compute various homological invariants such as Tor and Ext . We won't consider that, rather let's focus on the complexes $\mathbf{F}_\bullet^\lambda$ we just constructed which are a resolution of $L(\lambda)$ consisting of modules which have standard filtrations.

Another way to say that is that when restricted to $U(\mathfrak{n}^-)$, each \mathbf{F}_i^λ is a free $U(\mathfrak{n}^-)$ -module. Hence we can use it to compute Tor and Ext (but over the smaller ring $U(\mathfrak{n}^-)$). There's one particularly important calculation that falls into this scenario.

We've only discussed Ext^1 , and now we'll consider Ext^n for higher n . To understand this properly, we'd need to discuss derived functors. Let me skip that and just state the one fact we'll need. For a ring R and (left) R -modules M, N , to compute $\mathrm{Ext}_R^n(M, N)$, we do the following:

- Construct an R -free resolution \mathbf{F}_\bullet of M (a projective resolution is good enough).

- Apply the functor $\text{Hom}_R(-, N)$ to \mathbf{F}_\bullet to get a complex

$$\text{Hom}_R(\mathbf{F}_0, N) \xrightarrow{d^0} \text{Hom}_R(\mathbf{F}_1, N) \xrightarrow{d^1} \text{Hom}_R(\mathbf{F}_2, N) \xrightarrow{d^2} \dots$$

- $\text{Ext}_R^n(M, N)$ is the n th homology counting from the left, i.e., $\ker(d^n)/\text{image}(d^{n-1})$.

Notably, the final answer will not depend on the particular resolution \mathbf{F}_\bullet .

We will be interested in the special case where $R = U(\mathfrak{n}^-)$, $N = \mathbf{C}$ is the trivial representation, and $M = L(\lambda)$, where λ is integral and dominant, and the action is restricted from the action of $U(\mathfrak{g})$.

In that case $\mathbf{F}_\bullet^\lambda$ is the free resolution that we will use. In that case, note that for any weight μ , we have $\dim_{\mathbf{C}} \text{Hom}_{U(\mathfrak{n}^-)}(M(\mu), \mathbf{C}) = 1$ since a map only depends on the image of the highest weight vector. It will be helpful to keep track of the \mathfrak{h} -action on everything, and in that case, we can say that

$$\text{Hom}_{U(\mathfrak{n}^-)}(M(\mu), \mathbf{C}) = \mathbf{C}_{-\mu}.$$

In particular, our complex $\text{Hom}_{U(\mathfrak{n}^-)}(\mathbf{F}_\bullet^\lambda, \mathbf{C})$ gives us a complex of \mathfrak{h} -modules

$$\mathbf{C}_{-\lambda} \rightarrow \bigoplus_{\substack{w \in W \\ \ell(w)=1}} \mathbf{C}_{-w \bullet \lambda} \rightarrow \bigoplus_{\substack{w \in W \\ \ell(w)=2}} \mathbf{C}_{-w \bullet \lambda} \rightarrow \dots$$

Since all of the weights $-w \bullet \lambda$ are distinct, the maps are necessarily 0 (since they have to be \mathfrak{h} -linear).

We conclude that, as \mathfrak{h} -modules, we have

$$\text{Ext}_{U(\mathfrak{n}^-)}^k(L(\lambda), \mathbf{C}) = \bigoplus_{\substack{w \in W \\ \ell(w)=k}} \mathbf{C}_{-w \bullet \lambda}.$$

There are a few small things we can say. First, if M and N are finite-dimensional representations of \mathfrak{a} , where \mathfrak{a} is an arbitrary finite-dimensional Lie algebra, then we have isomorphisms

$$\text{Ext}_{U(\mathfrak{a})}^k(M, N) \cong \text{Ext}_{U(\mathfrak{a})}^k(N^*, M^*).$$

Second, the ext group $\text{Ext}_{U(\mathfrak{a})}^k(\mathbf{C}, N)$ is also called the k th Lie algebra cohomology group with coefficients in N , and is denoted by $H^k(\mathfrak{a}; N)$.

With this new terminology, let us summarize.

Theorem 8.4.1 (Bott). *As \mathfrak{h} -representations, we have*

$$H^k(\mathfrak{n}; L(\lambda)) = \bigoplus_{\substack{w \in W \\ \ell(w)=k}} \mathbf{C}_{w \bullet \lambda}.$$

Proof. We've done all of the work and just need to do some change of notation, so to speak. First, from above, we have

$$H^k(\mathfrak{n}^-; L(\lambda)^*) = \text{Ext}_{U(\mathfrak{n}^-)}^k(\mathbf{C}, L(\lambda)^*) = \text{Ext}_{U(\mathfrak{n}^-)}^k(L(\lambda), \mathbf{C}) = \bigoplus_{\substack{w \in W \\ \ell(w)=k}} \mathbf{C}_{-w \bullet \lambda}.$$

Next, as was used before, there is a unique longest element $w_0 \in W$ and $L(\lambda)^* \cong L(-w_0\lambda)$. Then the above becomes

$$H^k(\mathfrak{n}^-; L(\lambda)) = \bigoplus_{\substack{w \in W \\ \ell(w)=k}} \mathbf{C}_{-w \bullet (-w_0\lambda)}.$$

Finally, we apply one more trick: \mathfrak{n} is the span of the positive roots with respect to our fixed choice of simple roots Δ . However, $-\Delta$ is another choice of simple roots, and \mathfrak{n} is the span of the negative roots with this new system. Furthermore, with this new choice, ρ is replaced with $-\rho$ and the highest weight of $L(\lambda)$ gets replaced with its lowest weight, which is $w_0\lambda$. Hence, with respect to $-\Delta$, the subscript $-w \bullet (-w_0\lambda)$ becomes $-(w(-\lambda - \rho) + \rho) = w \bullet \lambda$ and the isomorphism above becomes

$$H^k(\mathfrak{n}; L(\lambda)) = \bigoplus_{\substack{w \in W \\ \ell(w)=k}} \mathbf{C}_{w \bullet \lambda}. \quad \square$$

Remark 8.4.2. There is a geometric version of this theorem which is better known. In that case, each weight μ corresponds to a line bundle $\mathcal{L}(\mu)$ on the flag variety X of \mathfrak{g} . Bott's theorem in that context says that its (coherent) sheaf cohomology is nonzero if and only if $\mu = w \bullet \lambda$ for some dominant integral weight λ , and in that case, it is concentrated in degree $\ell(w)$ and gives the representation $L(\lambda)$. The case when μ is already dominant and integral is also known as the Borel–Weil theorem and allows one to construct $L(\lambda)$ geometrically as the space of sections of the line bundle $\mathcal{L}(\lambda)$. I don't plan to discuss the connection between the two or elaborate any further. \square

8.5. Complements on the BGG complex. We were able to get two applications from the complexes \mathbf{F}_k^λ just knowing that the terms have standard filtrations (and which Verma modules show up). We can say much more actually, and we'll do so without proving the properties below.

First, the terms \mathbf{F}_k^λ are actually isomorphic to a direct sum of Verma modules:

$$\mathbf{F}_k^\lambda \cong \bigoplus_{\substack{w \in W \\ \ell(w)=k}} M(w \bullet \lambda).$$

This can be proven by showing that $\text{Ext}_{\mathcal{O}}^1(M(w \bullet \lambda), M(u \bullet \lambda)) = 0$ when $\ell(w) = \ell(u)$. In fact, a more refined statement can be made if we make use of the Bruhat order on W , which we define now.

Let $S = \{\sigma_1, \dots, \sigma_n\}$ be the simple reflections of W and let $T = \bigcup_{w \in W} wSw^{-1}$ be the set of elements conjugate to a simple reflection. Given $t \in T$ and $w, w' \in W$, we write $w' \xrightarrow{t} w$ if $w = tw'$ and $\ell(w) > \ell(w')$; we also write $w' \rightarrow w$ if there exists t such that $w' \xrightarrow{t} w$. Then we define $w \geq u$ if there is a sequence w_0, \dots, w_n such that $u = w_0 \rightarrow w_1 \rightarrow \dots \rightarrow w_n = w$. This is a very rich combinatorial structure, but we'll only touch on some aspects.

Now if λ is integral and dominant, then the Bruhat order can be related to the ordering on weights: $u \bullet \lambda < w \bullet \lambda$ if and only if $u > w$. This has an alternative formulation: if we instead assume that μ is integral and *antidominant*, then $u \bullet \mu > w \bullet \mu$ if and only if $u > w$. We will skip the verification, but see [H3, §5.2] (which phrases it the second way).

Proposition 8.5.1. *If λ is dominant and integral and $\text{Ext}_{\mathcal{O}}^1(M(w \bullet \lambda), M(u \bullet \lambda)) \neq 0$, then $u < w$ in Bruhat order (and hence $\ell(u) < \ell(w)$).*

See [H3, §6.5] for a statement about Ext groups for general weights.

In particular, the map $\mathbf{F}_k^\lambda \rightarrow \mathbf{F}_{k-1}^\lambda$ can be written as a sum of maps of the form $M(u \bullet \lambda) \rightarrow M(w \bullet \lambda)$ where $\ell(u) = k$ and $\ell(w) = k - 1$. If such a map is nonzero, then $u \bullet \lambda < w \bullet \lambda$ and hence $u > w$.

To go further, we quote some additional results.

Proposition 8.5.2. *Given weights μ, ν , we have $\dim \text{Hom}_{\mathcal{O}}(M(\mu), M(\nu)) \leq 1$. If equal to 1, then any nonzero homomorphism $M(\mu) \rightarrow M(\nu)$ is injective.*

Furthermore, for λ integral and dominant, we have $\dim \text{Hom}_{\mathcal{O}}(M(u \bullet \lambda), M(w \bullet \lambda)) = 1$ if and only if $u > w$.

This tells us the following: taking w to be the identity in the previous statement, $M(\lambda)$ contains a unique submodule which is isomorphic to $M(u \bullet \lambda)$ for each $u \in W$. So we can fix, once and for all, an identification for each u . Concretely, the space of highest weight vectors in $M(\lambda)$ of weight $u \bullet \lambda$ is 1-dimensional, and we are choosing a particular nonzero vector v_u in each. That means that for $u > w$, there is a distinguished map

$$i_{u,w}: M(u \bullet \lambda) \rightarrow M(w \bullet \lambda)$$


such that the image of the generator of $M(u \bullet \lambda)$ under $i_{u,w}$ gets sent to v_u when we embed $M(w \bullet \lambda)$ into $M(\lambda)$. Every other map is a scalar times $i_{u,w}$, so we can just identify them with this scalar. Note that everything depends on the choices of v_u , but once we make these choices, the maps $i_{u,w}$ are determined.

Now suppose we've made these choices and isomorphisms. Then in $\mathbf{F}_{\bullet}^{\lambda}$, we can identify each map $M(u \bullet \lambda) \rightarrow M(w \bullet \lambda)$ with a scalar $e(u, w)$. The condition that $\mathbf{F}_{\bullet}^{\lambda}$ is a complex translates to the fact that whenever $\ell(u) = \ell(w) + 2$, we have

$$\sum_{u < v < w} e(u, v)e(v, w) = 0.$$

Actually, the structure of these sums is not so complicated. There are exactly two cases:

- u and w are incomparable in Bruhat order, so that the sum is empty (and hence 0).
- $u < w$. In that case, there are exactly two choices of v , say v_1, v_2 , in that sum. This is referred to as a “square”:



We won't go any further, but it is known that if $\ell(u) = \ell(w) + 1$ and $u > w$, then $e(u, w) \neq 0$. Furthermore, it's possible to rescale all of the values so that we always have $e(u, w) \in \{-1, 0, 1\}$.

8.6. Translation functors. Given a weight λ , let $\text{pr}_{\lambda}: \mathcal{O} \rightarrow \mathcal{O}_{\chi_{\lambda}}$ denote the projection onto the subcategory $\mathcal{O}_{\chi_{\lambda}}$ as discussed before.

Now suppose that λ, μ are weights such that the difference $\mu - \lambda$ is integral. Then there is a unique dominant integral weight ν in the usual (not dotted) W -orbit of $\mu - \lambda$ and $L(\nu)$ is finite-dimensional. The **translation functor** T_{λ}^{μ} is defined to be the operation

$$M \mapsto \text{pr}_{\mu}(L(\nu) \otimes \text{pr}_{\lambda}(M)).$$

Since projection and tensoring (remember it's over \mathbf{C}) are exact, the translation functor is also exact. This is technically a functor $\mathcal{O} \rightarrow \mathcal{O}_{\chi_{\mu}}$, but we can also think of it as a functor $\mathcal{O}_{\chi_{\lambda}} \rightarrow \mathcal{O}_{\chi_{\mu}}$.

This topic is discussed in-depth in [H3, §7]. We've already seen this used above: starting with the Chevalley–Eilenberg complex \mathbf{K}_\bullet , we applied T_0^λ to get the BGG complex $\mathbf{F}_\bullet^\lambda$.

Here I'll just state a special case of the theorem in [H3, §7.8]. To be consistent, I'll use antidominant weights.

Theorem 8.6.1. *Let λ, μ be antidominant and integral weights. Then T_λ^μ gives an equivalence between $\mathcal{O}_{\chi_\lambda}$ and \mathcal{O}_{χ_μ} (with inverse T_μ^λ). For any $w \in W$, we have*

$$T_\lambda^\mu(M(w \bullet \lambda)) \cong M(w \bullet \mu), \quad T_\lambda^\mu(L(w \bullet \lambda)) \cong L(w \bullet \mu).$$

A more general theorem is stated there for not necessarily integral weights, but involves some more notation. The point here is that if one is focused on integral weights, then it's really enough to study the principal block \mathcal{O}_{χ_0} . In general, there are many results that state that a block is equivalent to a principal block (but perhaps in a semisimple Lie algebra of smaller rank).

9. KAZHDAN–LUSZTIG POLYNOMIALS

This section uses material from [H1, §7]. This is also an abridged form of my notes from Math 264C (Spring 2021).

I want to revisit the issue of expanding the class of a Verma module in $K(\mathcal{O})$ in terms of classes of irreducible modules. The solution relies on a special class of polynomials called the Kazhdan–Lusztig polynomials. There is a lot of ground to cover there; what I'd like to do is give some indication of how they are defined (this is already complicated).

9.1. Coxeter groups. The Kazhdan–Lusztig polynomials for \mathfrak{g} actually only depend on its Weyl group W (in the simple case, generally W determines \mathfrak{g} except that \mathfrak{so}_{2n+1} and \mathfrak{sp}_{2n} have isomorphic Weyl groups). In fact, they can be defined for a much larger class of groups, so I want to take a moment to explain that context.

First, let W be the Weyl group of \mathfrak{g} and let S be the set of simple reflections. Then each $s \in S$ is an involution, i.e., $s^2 = 1$. Furthermore, since W is finite, for each $s, s' \in S$, the product ss' has finite order, call it $m(s, s')$. A nontrivial fact is that this information is enough to describe the group. More precisely, W has the following presentation:

$$W = \langle s \in S \mid s^2 = 1, \quad (ss')^{m(s, s')} = 1 \rangle.$$

This is essentially the definition of a Coxeter group.

Definition 9.1.1. Let S be a finite set and $m: S \times S \rightarrow \mathbf{Z}_{>0} \cup \{\infty\}$ be a symmetric function such that $m(s, s) = 1$ and $m(s, s') \geq 2$ if $s \neq s'$. The associated **Coxeter group** is denoted W and is defined as the group with generators $s \in S$ with the relations $(ss')^{m(s, s')} = 1$ whenever $m(s, s') < \infty$. This information is usually given as (W, S) . \square

Since $m(s, s) = 1$, the relation above implies that each $s \in S$ has order 2. If $m(s, s') = 2$, the relation says that s and s' commute with each other.

Example 9.1.2. For the symmetric group, we can take $S = \{s_1, \dots, s_{n-1}\}$ with $m(s_i, s_j) = 2$ if $|i - j| > 1$ and $m(s_i, s_j) = 3$ if $|i - j| = 1$.

The dihedral group of order $2r$ is a Coxeter group with $S = \{s, t\}$ and $m(s, t) = r$. This is not a Weyl group if $r = 5$ or $r \geq 7$. \square

Remark 9.1.3. Given two Coxeter groups (W, S) and (W', S') , the product $W \times W'$ is the Coxeter group associated with $S \amalg S'$ where $m(s, s') = 2$ whenever $s \in S$ and $s' \in S'$. We will call this **reducible**; and (W, S) is **irreducible** if it is not of this form.

The Weyl groups of simple Lie algebras all give finite irreducible Coxeter groups, but they don't exhaust all of them. Besides the dihedral groups, there are two more examples: one is the symmetry group of the icosahedron (a 3-dimensional polytope), and the other is the symmetry group of a 4-dimensional polytope known as the "600-cell".

What is interesting is that the constructions below will have significance for Lie algebras, but still make sense even when W is not a Weyl group. \square

Remark 9.1.4. Being finite is a rather special condition on the function m , so we have many infinite Coxeter groups. Some of them can also be interpreted as "Weyl groups" of certain infinite-dimensional Lie algebras known as Kac–Moody algebras. \square

Finally, the definition of Bruhat order that we previously gave extends to the general setting of Coxeter groups.

9.2. Hecke algebras. Let (W, S) be a Coxeter group.

Let A be a commutative ring. Let \mathcal{E} be the free A -module with basis $\{T_w \mid w \in W\}$.

Theorem 9.2.1. *Let $a, b \in A$. There is a unique associative algebra structure on \mathcal{E} such that for all $s \in S$ and $w \in W$:*

$$T_s T_w = \begin{cases} T_{sw} & \text{if } \ell(sw) > \ell(w) \\ aT_w + bT_{sw} & \text{if } \ell(sw) < \ell(w) \end{cases}.$$

This algebra will be denoted $\mathcal{E}_A(a, b)$.

Uniqueness is clear: the relations say that $T_w = T_{s_1} \cdots T_{s_r}$ whenever $s_1 \cdots s_r$ is a reduced expression for w (and T_1 is the multiplicative identity). Hence, the product $T_v T_w$ for any $v, w \in W$ can be deduced from the relations above.

Existence is more subtle, and we will omit it.

Let $A = \mathbf{Z}[q^{\pm 1/2}]$ (this is the ring of Laurent polynomials in q with a square root of q adjoined). The **Hecke algebra** of (W, S) is $\mathcal{H} = \mathcal{E}_A(q - 1, q)$.

There is a different approach to constructing \mathcal{H} : we can instead define it as being generated by T_s for simple reflections and give their relations; the difficulty then is shifted to proving that \mathcal{H} is a free A -module with basis T_w for $w \in W$.

Remark 9.2.2. If we specialize $q = 1$, then \mathcal{H} becomes the group algebra of W . \square

9.3. R-polynomials. For $s \in S$, we have $T_s^2 = (q - 1)T_s + q$ in \mathcal{H} , which we rewrite as $T_s(T_s + 1 - q) = q$. Hence T_s is invertible with

$$T_s^{-1} = q^{-1}(T_s - (q - 1)).$$

This implies that T_w is invertible in general.

Theorem 9.3.1. *For $x \leq w$, there exist polynomials $R_{x,w}(q)$ of degree $\ell(w) - \ell(x)$ such that $R_{w,w}(q) = 1$ and*

$$T_w^{-1} = (-q)^{-\ell(w)} \sum_{x \leq w} (-1)^{\ell(x)} R_{x,w}(q) T_x.$$

Furthermore, these polynomials are nonzero.

We call the $R_{x,w}(q)$ the **R-polynomials**. They satisfy the following recursion:

Proposition 9.3.2. *With the convention that $R_{a,b}(q) = 0$ if $a \not\leq b$, pick $s \in S$ so that $w > sw$. Then for $x \leq w$, we have*

$$R_{x,w}(q) = \begin{cases} R_{sx,sw}(q) & \text{if } x > sx \\ qR_{sx,sw}(q) + (q-1)R_{x,sw}(q) & \text{if } x < sx \end{cases}.$$

For our purposes, these are an intermediate tool towards defining the Kazhdan–Lusztig polynomials (though we don't prove anything so we could have skipped this discussion). They do have some significance in their own right, but we won't discuss it.

9.4. Kazhdan–Lusztig polynomials. Define $\iota: \mathcal{H} \rightarrow \mathcal{H}$ on $\mathbf{Z}[q^{\pm 1/2}]$ by $\iota(q^{1/2}) = q^{-1/2}$ and extend it to \mathcal{H} on basis elements by $\iota(T_w) = T_{w^{-1}}$.

Lemma 9.4.1. *ι is a ring homomorphism, so in particular, $\iota^2 = 1$.*

Theorem 9.4.2. *For each $w \in W$, there exists a unique $C_w \in \mathcal{H}$ such that*

- (1) $\iota(C_w) = C_w$,
- (2) *there exist polynomials $P_{x,w}(q) \in \mathbf{Z}[q]$ for $x \leq w$ such that $P_{w,w}(q) = 1$ and $\deg P_{x,w}(q) \leq \frac{1}{2}(\ell(w) - \ell(x) - 1)$ for $x < w$, and*

$$C_w = (-q^{1/2})^{\ell(w)} \sum_{x \leq w} (-q)^{-\ell(x)} P_{x,w}(q^{-1}) T_x.$$

The $P_{x,w}(q)$ are the **Kazhdan–Lusztig polynomials**. There's a long road towards the next result, but I'm just going to state it. The result is attributed to Beilinson–Bernstein and also Brylinski–Kashiwara (see [H3, §8.8] for references).

Below, we focus on the principal block \mathcal{O}_{χ_0} ; the dominant weight there is 0 and the antidominant weight is -2ρ . For each $w \in W$, we let $M_w = M(w \bullet (-2\rho))$ and $L_w = L(w \bullet (-2\rho))$ so that we're shifting the focus to the antidominant weight. As usual, let $w_0 \in W$ be the unique element of longest length.

Theorem 9.4.3. *In $\mathbf{K}(\mathcal{O})$, we have*

$$\begin{aligned} [M_w] &= \sum_{x \leq w} P_{w_0 w, w_0 x}(1) [L_x], \\ [L_w] &= \sum_{x \leq w} (-1)^{\ell(w) - \ell(x)} P_{x,w}(1) [M_x]. \end{aligned}$$

From our discussion of translation functors, this equally applies to Verma modules and simples that belong to the other blocks $\mathcal{O}_{\chi_\lambda}$ where λ is integral and dominant. There is a lot known beyond that, but I won't discuss it.

These formulas only make use the specialization of the Kazhdan–Lusztig polynomials at 1, i.e., the sum of their coefficients. It turns out that these coefficients are always nonnegative, and it is natural to ask if they individually have any meaning.

One interpretation is in terms of the socle filtration of M_w , which we define now. Given any module M over a ring R , define its **socle**, denoted $\text{Soc}(M)$, to be the sum of all of its simple submodules. Note that this sum is necessarily a direct sum. The **socle filtration** is defined inductively by $\text{Soc}^0(M) = 0$ and, for $k > 0$, $\text{Soc}^k(M)$ is the preimage of $\text{Soc}(M/\text{Soc}^{k-1}(M))$ under the quotient map $M \rightarrow M/\text{Soc}^{k-1}(M)$. If M is noetherian, then this is a finite

filtration, and if M is artinian, the final term is M (if a nonzero module is artinian, then its socle is always nonzero).

Finally, for $k > 0$, we define $\text{Soc}_k(M) = \text{Soc}^k(M)/\text{Soc}^{k-1}(M)$. If $M \in \mathcal{O}$, we get an identity in $K(\mathcal{O})$:

$$[M] = \sum_k [\text{Soc}_k(M)].$$

Theorem 9.4.4. *If $x < w$, we have*

$$P_{w_0w, w_0x}(q) = \sum_k [\text{Soc}_{\ell(x)+1+2k}(M_w) : L_x] q^k.$$

Remark 9.4.5. • See [H3, Chapter 8] for a connection between Kazhdan–Lusztig polynomials and Schubert varieties. In fact, this was one of the early proofs that they have non-negative coefficients in the case when the Coxeter group is the Weyl group of a semisimple Lie algebra (or more generally, a Kac–Moody algebra). The general situation was resolved in [EW].

- It is an open problem (interval conjecture) to determine if $P_{x,y}(q)$ is a combinatorial invariant of the interval $[x, y]$ in Bruhat order. That is, $[x, y]$ carries a poset structure and if this is isomorphic to $[x', y']$ for some other elements in a Bruhat order, does this force $P_{x,y}(q) = P_{x',y'}(q)$?
- Kazhdan–Lusztig polynomials arise in many other parts of representation theory. For instance, they play a role in determining the characters of finite-dimensional representations of the superalgebra versions of semisimple Lie algebras. \square

10. ABSTRACT HIGHEST WEIGHT CATEGORIES

We'll follow the exposition in [CPS].

10.1. Highest weight categories. Let \mathbf{k} be a field and let \mathcal{C} be an abelian \mathbf{k} -linear category. Rather than define this, let me just point out the main examples of interest. If A is a \mathbf{k} -algebra, then the category of all left A -modules, denoted Mod_A is an abelian \mathbf{k} -linear category, as is the category of all right A -modules, denoted ${}_A\text{Mod}$. Any full subcategory which is closed under taking kernels and cokernels is also an example. For instance, if A is (left-)noetherian, then the full subcategory of finitely generated (left-)modules forms an abelian subcategory, denoted Mod_A^{fg} . So for instance, \mathcal{O} is an abelian \mathbf{C} -linear category. We will just use module categories (and their subcategories) as our examples; one can do things more generally, but I want to keep technicalities to a minimum.

A module M is **locally artinian** if it is the union of its finite length submodules; alternatively, if for every $x \in M$, the submodule generated by x has finite length. If M is locally artinian and S is simple, then S is defined to be a composition factor of M if it is a composition factor of some finite length submodule of M , and $[M : S]$ is defined to be the supremum of $[N : S]$ where N ranges over all finite length submodules of M .

We'll say that \mathcal{C} is locally artinian if all of its objects are locally artinian.

Given a partially ordered set (**poset** for short) (Λ, \leq) and $x, y \in \Lambda$, define the **interval** $[x, y]$ to be

$$[x, y] = \{z \in \Lambda \mid x \leq z, z \leq y\}.$$

Then Λ is called **interval-finite** if $[x, y]$ is a finite set for all x, y . We have seen an example: $\Lambda = \mathfrak{h}^*$ where \mathfrak{h} is a Cartan subalgebra in \mathfrak{g} , and $x \leq y$ is the usual order, i.e., $y - x$ is a $\mathbf{Z}_{\geq 0}$ -linear combination of positive roots.

Definition 10.1.1. Let \mathcal{C} be a locally artinian \mathbf{k} -linear category. Let Λ be the set of isomorphism classes of simple modules in \mathcal{C} and suppose for each $\lambda \in \Lambda$, we have a representative $S(\lambda)$ for this isomorphism class. We call \mathcal{C} a **highest weight category** if:

- (1) There is a partial ordering \leq on Λ that makes it interval-finite.
- (2) There is a collection of objects $\{A(\lambda) \mid \lambda \in \Lambda\}$ of \mathcal{C} such that:
 - (a) For all λ , we have $S(\lambda) \subset A(\lambda)$, and the composition factors $S(\mu)$ of $A(\lambda)/S(\lambda)$ satisfy $\mu < \lambda$.
 - (b) For all $\lambda, \mu \in \Lambda$, both $\dim_{\mathbf{k}} \text{Hom}_{\mathcal{C}}(A(\lambda), A(\mu))$ and $[A(\lambda) : S(\mu)]$ are finite.
- (3) Every $S(\lambda)$ has an injective envelope $I(\lambda)$ in \mathcal{C} . Furthermore, $I(\lambda)$ has a filtration (possibly infinite) $F_1 \subset F_2 \subset \cdots$ such that:
 - (a) $F_1 \cong A(\lambda)$.
 - (b) For $n \geq 1$, we have $F_{n+1}/F_n \cong A(\mu(n))$ for some $\mu(n) > \lambda$.
 - (c) For all $\mu \in \Lambda$, there are only finitely many n such that $\mu = \mu(n)$.
 - (d) The filtration is exhaustive, i.e., $\bigcup_n F_n = I(\lambda)$. □

The objects $A(\lambda)$ are called **co-standard modules** and any filtration in part (3) is called a **good filtration**.

Remark 10.1.2. If we also assume that each module is also finitely generated, then the locally artinian condition forces each object of \mathcal{C} to be finite length. Then the finiteness conditions (2b), (3c), and (3d) above become redundant. □

Example 10.1.3. Let $\mathcal{C} = \mathcal{O}$ for a semisimple complex Lie algebra \mathfrak{g} . Then we can take $\Lambda = \mathfrak{h}^*$ with the usual partial order. We've already seen that objects of \mathcal{O} are artinian. The objects $A(\lambda)$ can be taken to be the dual Verma modules $M(\lambda)^\vee$ and the properties about composition factors and injective envelopes follow by applying \vee to our obtained results on projective covers. Hence \mathcal{O} is an example of a highest weight category (in fact, it is one of the motivating examples). □

Example 10.1.4. Let \mathbf{k} be an algebraically closed field and let $G = \mathbf{GL}_n(\mathbf{k})$ be the group of $n \times n$ invertible matrices with entries in \mathbf{k} . A representation of G on a \mathbf{k} -vector space V is a homomorphism $\varphi: G \rightarrow \mathbf{GL}(V)$, the group of invertible operators on V . This representation is called **polynomial** if there is a choice of basis for V so that the coordinate functions of φ are polynomial functions in the n^2 coordinates of G .

The category of finite-dimensional polynomial representations is a highest weight category. If the characteristic of \mathbf{k} is 0, this is a boring statement since all representations are direct sums of irreducible representations, but this breaks down in positive characteristic.

Without going into too much detail, the irreducibles are indexed by integer partitions with at most n parts, that is, $\lambda = (\lambda_1, \dots, \lambda_n) \in \mathbf{Z}_{\geq 0}^n$ such that $\lambda_1 \geq \dots \geq \lambda_n \geq 0$. Let $S^d = \text{Sym}^d(\mathbf{k}^n)$ denote the d th symmetric power of the standard representation \mathbf{k}^n (i.e., the space of column vectors). Then for each λ , define $S^\lambda = S^{\lambda_1} \otimes \cdots \otimes S^{\lambda_n}$. This is always an injective module independent of characteristic. It contains a certain submodule called the **Schur module** \mathbf{S}_λ . This is irreducible in characteristic 0, but otherwise will play the role of the co-standard object.

The partial ordering on partitions is the dominance order, which is defined as follows: first define $|\lambda| = \lambda_1 + \cdots + \lambda_n$. If $|\lambda| \neq |\mu|$, we do not compare the elements, but otherwise, $\lambda \leq \mu$ means that $\lambda_1 + \cdots + \lambda_i \leq \mu_1 + \cdots + \mu_i$ for all $i = 1, \dots, n$. This is closely related to the usual order on weights for the Lie algebra \mathfrak{sl}_n .

The axioms for a highest weight category follow from a few non-trivial properties which we outline. First, a vector $v \in V$ of a polynomial representation V is a weight vector of weight $(\alpha_1, \dots, \alpha_n)$ if

$$\text{diag}(x_1, \dots, x_n)v = x_1^{\alpha_1} \cdots x_n^{\alpha_n}v$$

for all $x_1, \dots, x_n \in \mathbf{k}^\times$ and diag means the diagonal matrix with those entries. Then V has a basis of weight vectors, and we define its character $\text{ch}_V(x_1, \dots, x_n) = \sum_{\alpha} x^\alpha$, the sum over the weights of these basis vectors.

The character of the irreducible representation $L(\lambda)$ is of the form $x^\lambda + \sum_{\alpha \leq \lambda} c_{\lambda, \alpha} x^\alpha$ for some non-negative integers $c_{\lambda, \alpha}$.

The character of \mathbf{S}_λ is the Schur polynomial $s_\lambda(x_1, \dots, x_n)$. The coefficient of x^α is something combinatorial: it is the number of semistandard Young tableaux of shape λ and using α_i many i 's for all i . This is nonzero if and only if $\alpha \leq \lambda$.

Second, and more subtle, is that S^λ has a good filtration. An easier statement is that its character is the sum of Schur polynomials s_μ such that $\mu \geq \lambda$. More precisely, the coefficient of s_μ in ch_{S^λ} is the number of semistandard Young tableaux of shape μ using λ_i many i 's (note the reversal of the roles). This follows from a combinatorial result known as Pieri's rule. \square

Proposition 10.1.5. *Let \mathcal{C} be a highest weight category.*

- (1) *If $\text{Hom}_{\mathcal{C}}(S(\mu), A(\lambda)) \neq 0$, then $\mu = \lambda$.*
- (2) *If $\text{Ext}_{\mathcal{C}}^1(S(\mu), A(\lambda)) \neq 0$ or $\text{Ext}_{\mathcal{C}}^1(A(\mu), A(\lambda)) \neq 0$, then $\mu > \lambda$.*
- (3) *$I(\lambda)$ has a good filtration F_\bullet such that if we define $F_{n+1}/F_n \cong A(\mu(n))$, then for all $i, j > 0$, $\mu(i) < \mu(j)$ implies that $i < j$.*

Proof. (1) If $\text{Hom}_{\mathcal{C}}(S(\mu), A(\lambda)) \neq 0$, then a nonzero map $S(\mu) \rightarrow A(\lambda)$ is injective and hence gives an inclusion $S(\mu) \subset I(\lambda)$. However, $S(\lambda)$ is an essential submodule of $I(\lambda)$, which implies that $S(\mu) \cap S(\lambda) \neq 0$, and hence we must have $\mu = \lambda$.

(2) First suppose that $\text{Ext}_{\mathcal{C}}^1(S(\mu), A(\lambda)) \neq 0$. Then consider the short exact sequence

$$0 \rightarrow A(\lambda) \rightarrow I(\lambda) \rightarrow I(\lambda)/A(\lambda) \rightarrow 0.$$

Since $I(\lambda)$ is injective, $\text{Ext}_{\mathcal{C}}^1(S(\mu), I(\lambda)) = 0$, and so $\text{Hom}_{\mathcal{C}}(S(\mu), I(\lambda)/A(\lambda)) \neq 0$ since it has $\text{Ext}_{\mathcal{C}}^1(S(\mu), A(\lambda))$ as a quotient (consider the long exact sequence from applying $\text{Hom}_{\mathcal{C}}(S(\mu), -)$). In particular, if F_\bullet is a good filtration of $I(\lambda)$, there exists j such that $\text{Hom}_{\mathcal{C}}(S(\mu), F_j/F_1) \neq 0$. In particular, $\text{Hom}_{\mathcal{C}}(S(\mu), A(\mu(n))) \neq 0$ for some $n \leq j$, which means that $\mu(n) = \mu$ by (1). But also $\mu > \lambda$ by axiom (3b).

Next, suppose that $\text{Ext}_{\mathcal{C}}^1(A(\mu), A(\lambda)) \neq 0$. Then there exists $S(\nu)$ in the composition series of $A(\mu)$ such that $\text{Ext}_{\mathcal{C}}^1(S(\nu), A(\lambda)) \neq 0$. But then $\nu \leq \mu$ by axiom (2) and $\nu > \lambda$ by what we just proved, so we conclude that $\mu > \lambda$.

(3) First, pick any good filtration of $I(\lambda)$. Suppose there is an n so that

$$\text{Ext}_{\mathcal{C}}^1(A(\mu(n+1)), A(\mu(n))) = 0.$$

Then the short exact sequence

$$0 \rightarrow F_{n+1}/F_n \rightarrow F_{n+2}/F_n \rightarrow F_{n+2}/F_{n+1} \rightarrow 0$$

must split, which means there exists a submodule $F'_{n+1} \subset F_{n+2}$ that contains F_n such that

$$F'_{n+1}/F_n \cong A(\mu(n+1)), \quad F_{n+2}/F'_{n+1} \cong A(\mu(n)),$$

i.e., we can “swap” the order of $A(\mu(n))$ and $A(\mu(n+1))$ in this good filtration.

To finish, extend the partial ordering on the set $\{\mu(n)\}$ to a total ordering \prec (by the interval-finite property, we can do this so that there are no infinite decreasing sequences). Then if $\mu(n+1) \prec \mu(n)$, we know that $\mu(n) \not\prec \mu(n+1)$, and so we can swap these terms in the filtration. In particular, keep swapping terms until we have $\mu(1) \preceq \mu(2) \preceq \dots$. Since there are no infinite decreasing sequences, any given term will only be swapped finitely many times, and so the n th term of the resulting filtration is well-defined for all n . \square

10.2. Path algebras. In representation theory literature, a **quiver** Q is a directed graph. Formally, it consists of a finite vertex set V and for each pair $(v, w) \in V \times V$, we have a positive integer $a(v, w)$, which represents the number of arrows pointing from v to w , and will usually denote them by $a: v \rightarrow w$. Pictorially, here is an example with 3 vertices (and we have named the arrows for reference’s sake):

$$v_1 \begin{array}{c} \xrightarrow{a_1} \\ \xrightarrow{a_2} \end{array} v_2 \xrightarrow{a_3} v_3$$

A **representation** M of Q (over a field \mathbf{k}) consists of two pieces of data:

- For each $x \in V$, we have a \mathbf{k} -vector space M_x ,
- For each arrow $a: x \rightarrow y$, we have a linear map $M_a: M_x \rightarrow M_y$.

Notably, no compatibilities between these arrows are required. A homomorphism f between representations M and N consists of a linear map $f_x: M_x \rightarrow N_x$ for all $x \in V$ such that for all arrows $a: x \rightarrow y$, we have $f_y M_a = N_a f_x$.

The **path algebra** of Q , denoted $\mathbf{k}[Q]$ is the \mathbf{k} -vector space whose basis is the set of all directed paths in Q (including paths of length 0, which correspond to the choice of a vertex). In the example above, $\mathbf{k}[Q]$ has basis

$$\{v_1, v_2, v_3, a_1, a_2, a_3, a_1 a_3, a_2 a_3\}.$$

Multiplication of two paths is defined by concatenating paths when the endpoints match up, and 0 otherwise. In our example, $a_1 \cdot a_2 = 0$, $a_1 \cdot a_3 = a_1 a_3$, $v_1 \cdot a_1 = a_1$, $v_2 \cdot a_1 = 0$, etc. In particular, each vertex is an idempotent element in $\mathbf{k}[Q]$ and the multiplicative identity is the sum of the vertices.

Proposition 10.2.1. *The category of representations of Q over \mathbf{k} is equivalent to the category of right $\mathbf{k}[Q]$ -modules.*

Proof. Here’s just a sketch.

If M is a representation of Q , then $\bigoplus_{x \in V} M_x$ will be our $\mathbf{k}[Q]$ -module. If $a: x \rightarrow y$ is an arrow, then $ma = 0$ if $m \in M_z$ for $z \neq x$ and otherwise $ma = M_a(m) \in M_y$. Also, for each $x \in V$, the idempotent x acts on M_x by the identity and acts on M_z for $z \neq x$ by 0. The idempotents and arrows generate $\mathbf{k}[Q]$, so it remains to check that this extends to a well-defined module structure; it is a right module because of our convention on how paths are written.

Conversely, if M is a right $\mathbf{k}[Q]$ -module, then we get a representation by setting $M_x = Mx$ and defining M_a to be action of $a \in \mathbf{k}[Q]$. \square

An **acyclic quiver** is a quiver which does not have any directed cycles (or loops, i.e., $a(v, v) = 0$ for all v). This is equivalent to $\mathbf{k}[Q]$ being finite-dimensional.

In that case, there is one simple module $S(x)$ for each $x \in V$. This corresponds to the representation defined by

$$S(x)_y = \begin{cases} \mathbf{k} & \text{if } x = y \\ 0 & \text{if } x \neq y \end{cases}.$$

The radical of $\mathbf{k}[Q]$ is the span of all paths of positive length.

Since each vertex $x \in V$ is an idempotent, we get a projective module $P(x) = x\mathbf{k}[Q]$. This is the span of all paths that begin at x . In fact, this is the projective cover of $S(x)$. Dually, $I(x) = (\mathbf{k}[Q]x)^*$ is the injective envelope of $S(x)$ and is the dual space of the span of all paths that end at x .

This already tells us how to define a highest weight structure. Let $\Lambda = V$ and define $x \leq y$ if there exists a path starting at x and ending at y (since there are no directed cycles, this is indeed a partial ordering). Then we can take $A(x) = I(x)$, and we're done.

Remark 10.2.2. If A is a finite-dimensional algebra, then there exists a quiver Q and a 2-sided ideal $I \subset \mathbf{k}[Q]$ such that Mod_A is equivalent to $\text{Mod}_{\mathbf{k}[Q]/I}$. Hence, the study of quotients of path algebras can be seen as encompassing the study of general finite-dimensional algebras. However, the translation may not be so easy to understand in practice. \square

10.3. Quasi-hereditary algebras. A natural question is, given a \mathbf{k} -algebra A , is $\mathcal{C} = \text{Mod}_A$ (or ${}_A\text{Mod}$) a highest weight category? Cline, Parshall, and Scott characterized the finite-dimensional algebras with this property. We will need some definitions from noncommutative algebra.

Below, A will be a finite-dimensional algebra over \mathbf{k} . Given a left A -module M , its **annihilator** is defined to be

$$\text{Ann}(M) = \{a \in A \mid am = 0 \text{ for all } m \in M\}.$$

We can define the same for right modules.

Proposition 10.3.1. *The following sets are the same:*

- (1) *The intersection of $\text{Ann}(M)$ as M ranges over all simple left A -modules.*
- (2) *The intersection of $\text{Ann}(M)$ as M ranges over all simple right A -modules.*
- (3) *The intersection of all maximal left ideals of A .*
- (4) *The intersection of all maximal right ideals of A .*

The ideal characterized by the previous proposition is denoted $\text{rad}(A)$ and is called the **Jacobson radical** of A .

Generalizing (3), given any left A -module M , the **radical** of M , denoted $\text{rad}(M)$, is defined to be the intersection of all maximal proper submodules of M . If there are no maximal proper submodules, define $\text{rad}(M) = M$ (if M is finitely generated, maximal proper submodules do indeed exist). We can also make a similar definition for right A -modules. When $M = A$, this agrees with the previous definition.

Proposition 10.3.2. (1) *If M is a finite length module, then the quotient $M/\text{rad}(M)$ is semisimple, i.e., a direct sum of simple modules.*
 (2) *If A is finite-dimensional (the case we'll consider), then for any A -module M , we have $\text{rad}(M) = \text{rad}(A)M$.*

I'll skip the proof.

Definition 10.3.3. A nonzero 2-sided ideal J is **heredity** if

- (1) $J^2 = J$,
- (2) $J \operatorname{rad}(A)J = 0$, and
- (3) J is both a projective left A -module and a projective right A -module. □

The conditions might be difficult to parse, so let's see what they say about simple modules (how we'll use them).

Proposition 10.3.4. *Suppose that J is heredity.*

- (1) For any left A/J -module M , we have $\operatorname{Hom}_A(J, M) = 0$.
- (2) Every simple left A -module is a composition factor of exactly one of $J/\operatorname{rad}(J)$ or A/J .

Proof. (1) To see this, suppose there is a nonzero homomorphism $J \rightarrow M$. We can write M as a quotient of free A/J -module, and since J is projective, this gives us a nonzero homomorphism $J \rightarrow A/J$, i.e., a homomorphism $J \rightarrow A$ which is not contained in J . However, the image of J^2 under this map must be contained in J , but $J = J^2$, so we have a contradiction.

(2) Let S be a simple left A -module. If S appears in the composition series for A/J , then J annihilates S , i.e., S is an A/J -module. Then (1) implies that $\operatorname{Hom}_A(J, S) = 0$ and hence S is cannot be a summand of $J/\operatorname{rad}(J)$. In particular, the composition factors of $J/\operatorname{rad}(J)$ and A/J do not overlap.

On the other hand, every simple A -module is a composition factor of A , and hence must be a composition factor of $\operatorname{rad}(J)$, $J/\operatorname{rad}(J)$, or A/J . If S is a composition factor of $\operatorname{rad}(J)$, then since $J \operatorname{rad}(J) = J \operatorname{rad}(A)J = 0$, we have that S is an A/J -module, and hence is a composition factor of A/J . □

Finally, a finite-dimensional algebra A is **quasi-hereditary** if there exists a chain of 2-sided ideals

$$0 = J_0 \subset J_1 \subset \cdots \subset J_m = A$$

such that J_i/J_{i-1} is a heredity ideal in A/J_{i-1} for all $i = 1, \dots, m$. Such a chain is called a **heredity chain**.

Remark 10.3.5. If we prefer a recursive definition, we can replace this condition and instead declare that A is quasi-hereditary if either:

- A is semisimple, or
- A has a heredity ideal J such that A/J is quasi-hereditary. □

I couldn't think of any good simple examples to illustrate this definition. Most of the literature treats these algebras using different perspectives beyond the scope of what I have time to talk about. However, there is a rather nice fact which illustrates the point of these algebras.

Theorem 10.3.6. *Let A be a finite-dimensional \mathbf{k} -algebra. Then A is quasi-hereditary if and only if ${}_A \operatorname{Mod}$ is a highest weight category.*

Let's first tackle one implication.

Proposition 10.3.7. *If A is a quasi-hereditary \mathbf{k} -algebra, then ${}_A \operatorname{Mod}$ and Mod_A are both highest weight categories.*

Before we begin the proof, let me make a few remarks. First, the opposite algebra A^{op} is defined to have the same vector space as A , but with multiplication reversed: xy in A^{op} is defined to be the value of yx in A . Then we have an equivalence ${}_A \text{Mod}^{\text{fg}} \simeq \text{Mod}_{A^{\text{op}}}^{\text{fg}}$: if M is a right A -module, then the \mathbf{k} -linear dual M^* is a left A^{op} -module: for $x \in A^{\text{op}}$ and $m \in M$, xm is defined to be whatever mx originally was (and going the other way is similar; double dual is an equivalence for finite-dimensional spaces).

Note that the 2-sided ideals of A are the same as the 2-sided ideals of A^{op} . In particular, if A is quasi-hereditary, then so is A^{op} .

Finally, we'll focus on ${}_A \text{Mod}$ since the proof for Mod_A is similar but with all instances of left and right below swapped.

Proof. Let $0 = J_0 \subset \cdots \subset J_m = A$ be a heredity chain for A . We will view it as a heredity chain for A^{op} . For $i = 1, \dots, m$, let $M_i = J_i/J_{i-1}$. Then $M_i/\text{rad}(M_i)$ is a direct sum of simple left A^{op} -modules; let Λ_i denote the isomorphism classes of the modules that appear and define $\Lambda = \bigcup_i \Lambda_i$. The previous result implies that the Λ_i are pairwise disjoint and that Λ exhausts all isomorphism classes of simple left A^{op} -modules. We define a partial ordering on Λ by declaring that $\lambda < \mu$ if and only if $\lambda \in \Lambda_i$ and $\mu \in \Lambda_j$ such that $i > j$ (in other words, $\Lambda_1 > \Lambda_2 > \cdots > \Lambda_m$).

For $\lambda \in \Lambda_i$, let $T(\lambda)$ be the corresponding simple left A^{op} -module. In fact, it is a left A^{op}/J_{i-1} -module. Let $P'(\lambda)$ be its projective cover as a module over A^{op}/J_{i-1} ³. Then $P'(\lambda)$ is a direct summand of J_i/J_{i-1} and every composition factor of $P'(\lambda)$ is a composition factor of A^{op}/J_{i-1} and hence belongs to $\bigcup_{j \geq i} \Lambda_j$. Furthermore,

$$\text{rad}(P'(\lambda)) = \ker(P'(\lambda) \rightarrow T(\lambda))$$

and hence every composition factor of $\text{rad}(P'(\lambda)) \subset \text{rad}(J_i/J_{i-1})$ is in fact a composition factor of A^{op}/J_i by the proof of the previous result.

Now define $S(\lambda) = T(\lambda)^*$ and $A(\lambda) = P'(\lambda)^*$. Then $A(\lambda)$ is a right A -module containing $S(\lambda)$ such that the composition factors $S(\mu)$ of $A(\lambda)/S(\lambda) \cong \text{rad}(P'(\lambda))^*$ satisfy $\mu < \lambda$.

Next, let $P(\lambda)$ be the projective cover of $T(\lambda)$ as an A^{op} -module. Then $P(\lambda)$ is a direct summand of A^{op} , and hence, for all i , $J_i P(\lambda)/J_{i-1} P(\lambda)$ is a direct summand of J_i/J_{i-1} . Since the latter is a projective A^{op}/J_{i-1} -module, we conclude that $J_i P(\lambda)/J_{i-1} P(\lambda)$ is isomorphic to a direct sum of $P'(\mu)$ for $\mu \in \Lambda_i$. Also, if $\lambda \in \Lambda_k$, then $P(\lambda)/J_{k-1}$ is a direct summand of J_k/J_{k-1} , and since $J_k^2/J_{k-1} = J_k/J_{k-1}$, we conclude that $J_k P(\lambda) = P(\lambda)$ and that $P(\lambda)/J_{k-1} P(\lambda) \cong P'(\lambda)$.

Finally, $I(\lambda) = P(\lambda)^*$ is an injective envelope of $S(\lambda)$. Starting with the filtration

$$0 \subset J_1 P(\lambda) \subset J_2 P(\lambda) \subset \cdots \subset J_k P(\lambda) = P(\lambda),$$

we get a filtration $F'_1 \subset \cdots \subset F'_k$ on $I(\lambda)$ by taking $F'_{k-i} = (P(\lambda)/J_i P(\lambda))^*$. Then

$$F'_1 = (P(\lambda)/J_{k-1} P(\lambda))^* \cong A(\lambda)$$

and the remaining quotients are direct sums of $A(\mu)$ where $\mu > \lambda$. So we can refine this filtration to get the kind required by axiom (3) in the definition of highest weight category. \square

³We discussed projective covers in the context of category \mathcal{O} , but as was remarked there, the methods apply equally well to arbitrary modules which are artinian and noetherian, e.g., finitely generated modules over a finite-dimensional algebra.

Now we'd like to prove that if Mod_A is a highest weight category, then A must be quasi-hereditary. The definition is recursive in nature, so we'll need some information about subcategories of highest weight categories.

A subset $\Gamma \subset \Lambda$ is a **lower ideal** if $x \in \Gamma$ and $y \leq x$ implies that $y \in \Gamma$. It is **finitely generated** there is a finite list of elements x_1, \dots, x_n such that every $y \in \Gamma$ satisfies $y \leq x_i$ for some i . We define $\mathcal{C}[\Gamma]$ to be the full subcategory of \mathcal{C} consisting of modules M such that every composition factor S of M is indexed by an element in Γ .

Proposition 10.3.8. *If Γ is a finitely generated lower ideal of Λ , then $\mathcal{C}[\Gamma]$ is a highest weight category and each injective envelope has a good filtration of finite length.*

Proof. We take Γ as our partially ordered set indexing simple objects with the induced partial ordering. If $\lambda \in \Gamma$, then $A(\lambda)$ belongs to $\mathcal{C}[\Gamma]$ since Γ is a lower ideal, so we take these objects to satisfy axiom (2). Next, let $I(\lambda)$ be the injective envelope of $S(\lambda)$ in \mathcal{C} , and define $I(\lambda)_\Gamma$ to be the maximal submodule of $I(\lambda)$ that belongs to $\mathcal{C}[\Gamma]$.

We claim that $I(\lambda)_\Gamma$ is the injective envelope of $S(\lambda)$ in $\mathcal{C}[\Gamma]$. First, if $M \rightarrow N$ is an injection of modules in $\mathcal{C}[\Gamma]$, and $M \rightarrow I(\lambda)_\Gamma$ is any homomorphism, then it can be extended to a homomorphism $N \rightarrow I(\lambda)$. However, by definition of $I(\lambda)_\Gamma$, its image lies in $I(\lambda)_\Gamma$, so that $I(\lambda)_\Gamma$ is injective in the subcategory $\mathcal{C}[\Gamma]$. Second, the only simple submodule of $I(\lambda)_\Gamma$ is $S(\lambda)$ by definition, and so $S(\lambda) \subset I(\lambda)_\Gamma$ is an essential submodule, and the claim is proven.

Finally, let F_\bullet be a good filtration of $I(\lambda)$. Since Λ is interval-finite and Γ is finitely generated, the number $\mu \in \Gamma$ such that $A(\mu)$ is a quotient of the form F_i/F_{i-1} is finite. By Proposition 10.1.5, we can pick a new good filtration F' such that there exists n such that if $F'_{i+1}/F'_i \cong A(\mu(i))$, then $\mu(i) \in \Gamma$ if and only if $i \leq n$. Then $F'_1 \subset \dots \subset F'_{n+1} = I(\lambda)_\Gamma$ is a good filtration for $I(\lambda)_\Gamma$ in $\mathcal{C}[\Gamma]$. \square

Now we come to the last part of the theorem. For two modules M, N , define the **trace** $\text{Tr}_M(N)$ to be the submodule of N generated by all images of homomorphisms $M \rightarrow N$.

Proposition 10.3.9. *If A is finite-dimensional and ${}_A \text{Mod}$ is a highest weight category, then A is quasi-hereditary.*

Proof. Let Λ be the poset associated some highest weight structure on ${}_A \text{Mod}$. Since A is finite-dimensional, it has finitely many simple modules, so the poset Λ is finite. If $|\Lambda| = 1$, then A is a semisimple algebra, so there is nothing to show.

Otherwise, let $\lambda \in \Lambda$ be any maximal element. Then $I(\lambda) = A(\lambda)$ is an injective right A -module. Its dual $P = I(\lambda)^*$ is the projective cover of the simple left A^{op} -module $S(\lambda)^*$. In particular, $P(\lambda)$ is a summand of A^{op} , so we have $P = A^{\text{op}}e$ for some idempotent e . Define $J = A^{\text{op}}eA^{\text{op}} = AeA$, which is also the trace $\text{Tr}_P(A^{\text{op}})$.

Since $\text{Hom}_{A^{\text{op}}}(P, -)$ is exact and any nonzero map $P \rightarrow P$ is an isomorphism, we conclude that for any indecomposable projective $P(\mu) = I(\mu)^*$, the trace $\text{Tr}_P(P(\mu))$ is isomorphic to $P^{\oplus r_\mu}$, where r_μ is the number of times that $A(\mu)$ appears in a good filtration for $I(\mu)$. In particular, since trace distributes over direct sums and A^{op} is a direct sum of the $P(\mu)$, we conclude that J is a direct sum of copies of P , and hence is projective.

Next, $J^2 = (AeA)(AeA) = AeA = J$ (since $e^2 = e$).

Finally, since all simples in the composition series of $\text{rad } P = (P/S(\lambda))^*$ are different from $S(\lambda)^*$, we have $\text{Hom}_{A^{\text{op}}}(P, \text{rad } P) = 0$. Finally, $\text{rad } P = (\text{rad } A^{\text{op}})A^{\text{op}}e$, so we conclude that $e(\text{rad } A^{\text{op}})A^{\text{op}}e = 0$. This implies that $J(\text{rad } A^{\text{op}})J = 0$. In conclusion, J is a heredity ideal of A^{op} , and hence also of A .

To finish, note that if $\mu \neq \lambda$, then $eS(\mu)^* = 0$, and so J annihilates any simple other than $S(\lambda)$. In particular, if we set $\Gamma = \Lambda \setminus \{\lambda\}$, then every module in $\mathcal{C}[\Gamma]$ is annihilated by a power of J , and hence by J itself since $J^2 = J$. We conclude that ${}_{A/J} \text{Mod}$ can be identified with $\mathcal{C}[\Gamma]$ which is a highest weight category, and so we are done. \square

Finally, let's conclude with one more related statement.

Proposition 10.3.10. *Let \mathcal{C} be a highest weight category such that Λ finite. Let $I = \bigoplus_{\lambda \in \Lambda} I(\lambda)$ and let $A = \text{End}_{\mathcal{C}}(I)$. Then A is quasi-hereditary and ${}_A \text{Mod} \simeq \mathcal{C}$.*

We mentioned the dual version of this for projective covers earlier.

REFERENCES

- [BB] Anders Björner, Francesco Brenti, *Combinatorics of Coxeter Groups*
- [CPS] E. Cline, B. Parshall, L. Scott, Finite dimensional algebras and highest weight categories, *J. reine angew. Math.* **391** (1988), 85–99.
- [EW] Ben Elias, Geordie Williamson, The Hodge theory of Soergel bimodules, *Ann. of Math. (2)* **180** (2014), no. 3, 1089–1136, [arXiv:1212.0791v2](https://arxiv.org/abs/1212.0791v2).
- [H1] James E. Humphreys, *Reflection Groups and Coxeter Groups*
- [H2] James E. Humphreys, *Introduction to Lie Algebras and Representation Theory*
- [H3] James E. Humphreys, *Representations of Semisimple Lie Algebras in the BGG category \mathcal{O}*
- [J] Jens Carsten Jantzen, *Representations of Algebraic Groups*