LECTURE 14

The General Setup: Analyzing the Structure and Representations of an Arbitrary Semisimple Lie Algebra

This is the last of the four central lectures; in the body of it, §14.1, we extract from the examples of §11–13 the basic algorithm for analyzing a general semisimple Lie algebra and its representations. It is this algorithm that we will spend the remainder of Part III carrying out for the classical algebras, and the reader who finds the general setup confusing may wish to read this lecture in parallel with, for example, Lectures 15 and 16. In particular, §14.2 is less clearly motivated by what we have worked out so far; the reader may wish to skim it for now and defer a more thorough reading until after going through some more of the examples of Lectures 15–20.

§14.1. Analyzing Simple Lie Algebras in General

We said at the outset of Lecture 12 that once the analysis of the representations of $\mathfrak{sl}_3\mathbb{C}$ was understood, the analysis of the representations of any semisimple Lie algebra would be clear, at least in broad outline. Here we would like to indicate how that analysis will go in general, by providing an essentially algorithmic procedure for describing the representations of an arbitrary complex semisimple Lie algebra $\mathfrak{g}$. The process we give here is directly analogous, step for step, to that carried out in Lecture 12 for $\mathfrak{sl}_3\mathbb{C}$; the only difference is one change in the order of steps: having seen in the case of $\mathfrak{sl}_3\mathbb{C}$ the importance of the "distinguished" subalgebras $\mathfrak{s}_\alpha \cong \mathfrak{sl}_2\mathbb{C} \subset \mathfrak{g}$ and the corresponding distinguished elements $H_\alpha \in \mathfrak{s}_\alpha \subset \mathfrak{h}$, we will introduce them earlier here.

Step 0. *Verify that your Lie algebra is semisimple*; if not, none of the following will work (but see Remark 14.3). If your Lie algebra is not semisimple, pass as indicated in Lecture 9 to its semisimple part; a knowledge of the representations of this quotient algebra may not tell you everything about

W. Fulton et al., *Representation Theory*

the representations of the original, but it will at least tell you about the irreducible representations.

Step 1. *Find an abelian subalgebra* $\mathfrak{h} \subset \mathfrak{g}$ *acting diagonally.* This is of course the analogue of looking at the specific element H in $\mathfrak{sl}_2\mathbb{C}$ and the subalgebra $\mathfrak{h}$ of diagonal matrices in the case of $\mathfrak{sl}_3\mathbb{C}$; in general, to serve an analogous function it should be an abelian subalgebra that acts diagonally on one faithful (and hence, by Theorem 9.20, on any) representation of $\mathfrak{g}$. Moreover, in order that the restriction of a representation V of $\mathfrak{g}$ to $\mathfrak{h}$ carry the greatest possible information about V, $\mathfrak{h}$ should clearly be maximal among abelian, diagonalizable subalgebras; such a subalgebra is called a *Cartan subalgebra.*

It may seem that this step is somewhat less than algorithmic; in particular, while it is certainly possible to tell when a subalgebra of a given Lie algebra is abelian, and when it is diagonalizable, it is not clear how to tell whether it is maximal with respect to these properties. This defect will, however, be largely cleared up in the next step (see Remark 14.3).

Step 2. *Let* $\mathfrak{h}$ *act on* $\mathfrak{g}$ *by the adjoint representation, and decompose* $\mathfrak{g}$ *accordingly.* By the choice of $\mathfrak{h}$, its action on any representation of $\mathfrak{g}$ will be diagonalizable; applying this to the adjoint representation we arrive at a direct sum decomposition, called a *Cartan decomposition,*

$$\mathfrak{g} = \mathfrak{h} \oplus \left(\bigoplus \mathfrak{g}_\alpha\right), \tag{14.1}$$

where the action of $\mathfrak{h}$ preserves each $\mathfrak{g}_\alpha$ and acts on it by scalar multiplication by the linear functional $\alpha \in \mathfrak{h}^*$; that is, for any $H \in \mathfrak{h}$ and $X \in \mathfrak{g}_\alpha$ we will have

$$\operatorname{ad}(H)(X) = \alpha(H) \cdot X.$$

The second direct sum in the expression (14.1) is over a finite set of eigenvalues $\alpha \in \mathfrak{h}^*$; these eigenvalues—in the language of Lecture 12, the *weights of the adjoint representation*—are called the *roots* of the Lie algebra and the corresponding subspaces $\mathfrak{g}_\alpha$ are called the *root spaces.* Of course, $\mathfrak{h}$ itself is just the eigenspace for the action of $\mathfrak{h}$ corresponding to the eigenvalue 0 (see Remark 14.3 below); so that in some contexts—such as the following paragraph, for example—it will be convenient to adopt the convention that $\mathfrak{g}_0 = \mathfrak{h}$; but we do not usually count $0 \in \mathfrak{h}^*$ as a root. The set of all roots is usually denoted $R \subset \mathfrak{h}^*$.

As in the previous cases, we can picture the structure of the Lie algebra in terms of the diagram of its roots: by the fundamental calculation of §11.1 and Lecture 12 (which we will not reproduce here for the fourth time) we see that the adjoint action of $\mathfrak{g}_\alpha$ carries the eigenspace $\mathfrak{g}_\beta$ into another eigenspace $\mathfrak{g}_{\alpha+\beta}$.

There are a couple of things we can anticipate about how the configuration of roots (and the corresponding root spaces) will look. We will simply state them here as

Facts 14.2

(i) *each root space* $\mathfrak{g}_\alpha$ *will be one dimensional.*
(ii) R *will generate a lattice* $\Lambda_R \subset \mathfrak{h}^*$ *of rank equal to the dimension of* $\mathfrak{h}$.

(iii) *R is symmetric about the origin, i.e., if $\alpha \in R$ is a root, then $-\alpha \in R$ is a root as well.*

These facts will all be proved in general in due course; for the time being, they are just things we will observe as we do the analysis of each simple Lie algebra in turn. We mention them here simply because some of what follows will make sense only given these facts. Note in particular that by (ii), the roots all lie in (and span) a real subspace of $\mathfrak{h}^*$; all our pictures clearly will be of this real subspace.

Remark 14.3. If indeed 0 does appear as an eigenvalue of the action of $\mathfrak{h}$ on $\mathfrak{g}/\mathfrak{h}$, then we may conclude from this that $\mathfrak{h}$ was not maximal to begin with: by the above, anything in the 0-eigenspace of the action of $\mathfrak{h}$ commutes with $\mathfrak{h}$ and (given the fact that the $\mathfrak{g}_\alpha$ are one dimensional) acts diagonally on $\mathfrak{g}$, so that if it not already in $\mathfrak{h}$, then $\mathfrak{h}$ could be enlarged while still retaining the properties of being abelian and diagonalizable. Similarly, the assertion in (ii) that the roots span $\mathfrak{h}^*$ follows from the fact that an element of $\mathfrak{h}$ in the annihilator of all of them would be in the center of $\mathfrak{g}$.

From what we have done so far, we get our first picture of the structure of an arbitrary irreducible finite-dimensional representation V of $\mathfrak{g}$. Specifically, V will admit a direct sum decomposition

$$V = \bigoplus V_\alpha, \tag{14.4}$$

where the direct sum runs over a finite set of $\alpha \in \mathfrak{h}^*$ and $\mathfrak{h}$ acts diagonally on each V_α by multiplication by the eigenvalue α, i.e., for any $H \in \mathfrak{h}$ and $v \in V_\alpha$ we will have

$$H(v) = \alpha(H) \cdot v.$$

The eigenvalues $\alpha \in \mathfrak{h}^*$ that appear in this direct sum decomposition are called the *weights* of V; the V_α themselves are called *weight spaces*; and the dimension of a weight space V_α will be called the *multiplicity* of the weight α in V. We will often represent V by drawing a picture of the set of its weights and thinking of each dot as representing a subspace; this picture (often with some annotation to denote the multiplicity of each weight) is called the *weight diagram* of V.

The action of the rest of the Lie algebra on V can be described in these terms: for any root β, we have

$$\mathfrak{g}_\beta : V_\alpha \to V_{\alpha+\beta},$$

so we can think of the action of $\mathfrak{g}_\beta$ on V as a translation in the weight diagram, shifting each of the dots over by β and mapping the weight spaces correspondingly.

Observe next that all the weights of an irreducible representation are congruent to one another modulo the root lattice Λ_R: otherwise, for any weight α of V the subspace

$$V' = \bigoplus_{\beta \in \Lambda_R} V_{\alpha+\beta}$$

would be a proper subrepresentation of V. In particular, in view of Fact 14.2(ii), this means that the weights all lie in a translate of the real subspace spanned by the roots, so that it is not so unreasonable to draw a picture of them.

Step 3. *Find the distinguished subalgebras* $\mathfrak{s}_\alpha \cong \mathfrak{sl}_2\mathbb{C} \subset \mathfrak{g}$. As we saw in the example of $\mathfrak{sl}_3\mathbb{C}$, a crucial ingredient in the analysis of an arbitrary irreducible finite-dimensional representation is the restriction of the representation to certain special copies of the algebra $\mathfrak{sl}_2\mathbb{C}$ contained in $\mathfrak{g}$, and the application of what we know from Lecture 11 about such representations. To generalize this to our arbitrary Lie algebra $\mathfrak{g}$, let $\mathfrak{g}_\alpha \subset \mathfrak{g}$ be a root space, one dimensional by (i) of Fact 14.2. Then by (iii) of Fact 14.2, there is another root space $\mathfrak{g}_{-\alpha} \subset \mathfrak{g}$; and their commutator $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$ must be a subspace of $\mathfrak{g}_0 = \mathfrak{h}$, of dimension at most one. The adjoint action of the commutator $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$ thus carries each of $\mathfrak{g}_\alpha$ and $\mathfrak{g}_{-\alpha}$ into itself; so that *the direct sum*

$$\mathfrak{s}_\alpha = \mathfrak{g}_\alpha \oplus \mathfrak{g}_{-\alpha} \oplus [\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}] \tag{14.5}$$

is a subalgebra of $\mathfrak{g}$. The structure of $\mathfrak{s}_\alpha$ is not hard to describe, given two further facts that we will state here, verify in cases, and prove in general in Appendix D.

Facts 14.6.

(i) $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}] \neq 0$; *and*
(ii) $[[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}], \mathfrak{g}_\alpha] \neq 0$.

Given these, it follows that *the subalgebra* $\mathfrak{s}_\alpha$ *is isomorphic* to $\mathfrak{sl}_2\mathbb{C}$. In particular, we can pick a basis $X_\alpha \in \mathfrak{g}_\alpha$, $Y_\alpha \in \mathfrak{g}_{-\alpha}$, and $H_\alpha \in [\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$ satisfying the standard commutation relations (9.1) for $\mathfrak{sl}_2\mathbb{C}$; X_α and Y_α are not determined by this, but H_α is, being the unique element of $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$ having eigenvalues 2 and -2 on $\mathfrak{g}_\alpha$ and $\mathfrak{g}_{-\alpha}$, respectively [i.e., H_α is uniquely characterized by the requirements that $H_\alpha \in [\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$ and $\alpha(H_\alpha) = 2$.]

Step 4. *Use the integrality of the eigenvalues of the* H_α. The distinguished elements $H_\alpha \in \mathfrak{h}$ found above are important first of all because, by the analysis of the representations of $\mathfrak{sl}_2\mathbb{C}$ carried out in Lecture 9, in any representation of $\mathfrak{s}_\alpha$—and hence in any representation of $\mathfrak{g}$—*all eigenvalues of the action of* H_α *must be integers.* Thus, every eigenvalue $\beta \in \mathfrak{h}^*$ of every representation of $\mathfrak{g}$ must assume integer values on all the H_α. We correspondingly let Λ_W be the set of linear functionals $\beta \in \mathfrak{h}^*$ that are integer valued on all the H_α; Λ_W will be a lattice, called the *weight lattice* of $\mathfrak{g}$, with the property that

all weights of all representations of $\mathfrak{g}$ *will lie in* Λ_W.

Note, in particular, that $R \subset \Lambda_W$ and hence $\Lambda_R \subset \Lambda_W$; in fact, the root lattice will in general be a sublattice of finite index in the weight lattice.

Step 5. *Use the symmetry of the eigenvalues of the* H_α. The integrality of the eigenvalues of the H_α under any representation is only half the story; it is also true that they are symmetric about the origin in $\mathbb{Z}$. To express this, for any α we introduce the involution W_α on the vector space $\mathfrak{h}^*$ with $+1$-eigenspace the hyperplane

$$\Omega_\alpha = \{\beta \in \mathfrak{h}^* \colon \langle H_\alpha, \beta \rangle = 0\} \tag{14.7}$$

and minus 1 eigenspace the line spanned by α itself.[1] In English, W_α is the reflection in the plane Ω_α with axis the line spanned by α:

$$W_\alpha(\beta) = \beta - \frac{2\beta(H_\alpha)}{\alpha(H_\alpha)}\alpha = \beta - \beta(H_\alpha)\alpha. \tag{14.8}$$

Let $\mathfrak{W}$ be the group generated by these involutions; $\mathfrak{W}$ is called the *Weyl group* of the Lie algebra $\mathfrak{g}$.

Now suppose that V is any representation of $\mathfrak{g}$, with eigenspace decomposition $V = \bigoplus V_\beta$. The weights β appearing in this decomposition can then be broken up into equivalence classes mod α, and the direct sum

$$V_{[\beta]} = \bigoplus_{n \in \mathbb{Z}} V_{\beta + n\alpha} \tag{14.9}$$

of the eigenspaces in a given equivalence class will be a subrepresentation of V for $\mathfrak{s}_\alpha$. It follows then that the set of weights of V congruent to any given β mod α will be invariant under the involution W_α; in particular,

The set of weights of any representation of $\mathfrak{g}$ *is invariant under the Weyl group.*

To make this more explicit, the string of weights that correspond to nonzero summands in (14.9) are, possibly after replacing β by a translate by a multiple of α:

$$\beta, \beta + \alpha, \beta + 2\alpha, \ldots, \beta + m\alpha, \quad \text{with } m = -\beta(H_\alpha). \tag{14.10}$$

(Note that by our analysis of $\mathfrak{sl}_2\mathbb{C}$ this must be an uninterrupted string.) Indeed if we choose β and $m \geq 0$ so that (14.10) is the string corresponding to nonzero summands in (14.9), then the string of integers

$$\beta(H_\alpha), (\beta + \alpha)(H_\alpha) = \beta(H_\alpha) + 2, \ldots, (\beta + m\alpha)(H_\alpha) = \beta(H_\alpha) + 2m$$

must be symmetric about zero, so $\beta(H_\alpha) = -m$. In particular,

$$W_\alpha(\beta + k\alpha) = \beta + (-\beta(H_\alpha) - k)\alpha = \beta + (m - k)\alpha.$$

Note also that by the same analysis the multiplicities of the weights are invariant under the Weyl group.

We should mention one other fact about the Weyl group, whose proof we also postpone:

[1] Note that by the nondegeneracy assertion (ii) of Fact 14.6, the line $\mathbb{C}\cdot\alpha$ does not lie in the hyperplane Ω_α. Recall that $\langle\ ,\ \rangle$ is the pairing between $\mathfrak{h}$ and $\mathfrak{h}^*$, so $\langle H_\alpha, \beta\rangle = \beta(H_\alpha)$.

Fact 14.11. *Every element of the Weyl group is induced by an automorphism of the Lie algebra* $\mathfrak{g}$ *carrying* $\mathfrak{h}$ *to itself.*

We can even say what automorphism of $\mathfrak{g}$ does the trick: to get the involution W_α, take the adjoint action of the exponential $\exp(\pi i U_\alpha) \in G$, where G is any group with Lie algebra $\mathfrak{g}$ and U_α is a suitable element of the direct sum of the root spaces $\mathfrak{g}_\alpha$ and $\mathfrak{g}_{-\alpha}$. To prove that $\mathrm{Ad}(\exp(\pi i U_\alpha))$ actually does this requires more knowledge of $\mathfrak{g}$ than we currently possess; but it would be an excellent exercise to verify this assertion directly in each of the cases studied below. (For the general case see (23.20) and (26.15).)

Step 6. *Draw the picture* (optional). While there is no logical need to do so at this point, it will be much easier to think about what is going on in $\mathfrak{h}^*$ if we introduce the appropriate inner product, called the *Killing form*, on $\mathfrak{g}$ (hence by restriction on $\mathfrak{h}$, and hence on $\mathfrak{h}^*$). Since the introduction of the Killing form is, logically, a digression, we will defer until later in this lecture a discussion of its various definitions and properties. It will suffice for now to mention the characteristic property of the induced inner product on $\mathfrak{h}^*$: up to scalars it is the unique inner product on $\mathfrak{h}^*$ preserved by the Weyl group, i.e., in terms of which the Weyl group acts as a group of orthogonal transformations. Equivalently, it is the unique inner product (up to scalars) such that the line spanned by each root $\alpha \in \mathfrak{h}^*$ is actually perpendicular to the plane Ω_α (so that the involution W_α is just a reflection in that hyperplane). Indeed, in practice this is most often how we will compute it. In terms of the Killing form, then, we can say that *the Weyl group is just the group generated by the reflections in the hyperplanes perpendicular to the roots of the Lie algebra.*

Step 7. *Choose a direction in* $\mathfrak{h}^*$. By this we mean a real linear functional l on the lattice Λ_R irrational with respect to this lattice. This gives us a decomposition of the set

$$R = R^+ \cup R^-, \tag{14.12}$$

where $R^+ = \{\alpha: l(\alpha) > 0\}$ (the $\alpha \in R^+$ are called the *positive* roots, those in R^- *negative*); this decomposition is called an *ordering of the roots*. For most purposes, the only aspect of l that matters is the associated ordering of the roots.

The point of choosing a direction—and thereby an ordering of the roots $R = R^+ \cup R^-$—is, of course, to mimic the notion of highest weight vector that was so crucial in the cases of $\mathfrak{sl}_2\mathbb{C}$ and $\mathfrak{sl}_3\mathbb{C}$. Specifically, we make the

Definition. Let V be any representation of $\mathfrak{g}$. A nonzero vector $v \in V$ that is both an eigenvector for the action of $\mathfrak{h}$ and in the kernel of $\mathfrak{g}_\alpha$ for all $\alpha \in R^+$ is called a *highest weight vector* of V.

Just as in the previous cases, we then have

Proposition 14.13. *For any semisimple complex Lie algebra* $\mathfrak{g}$,

(i) *every finite-dimensional representation* V *of* $\mathfrak{g}$ *possesses a highest weight vector;*

(ii) *the subspace W of V generated by the images of a highest weight vector v under successive applications of root spaces $\mathfrak{g}_\beta$ for $\beta \in R^-$ is an irreducible subrepresentation;*

(iii) *an irreducible representation possesses a unique highest weight vector up to scalars.*

PROOF. Part (i) is immediate: we just take α to be the weight appearing in V for which the value $l(\alpha)$ is maximal and choose v any nonzero vector in the weight space V_α. Since $V_{\alpha+\beta} = (0)$ for all $\beta \in R^+$, such a vector v will necessarily be in the kernel of all root spaces $\mathfrak{g}_\beta$ corresponding to positive roots β.

Part (ii) may be proved by the same argument as in the two cases we have already discussed: we let W_n be the subspace spanned by all $w_n \cdot v$ where w_n is a word of length at most n in elements of $\mathfrak{g}_\beta$ for negative β. We then claim that for any X in any positive root space, $X \cdot W_n \subset W_n$. To see this, write a generator of W_n in the form $Y \cdot w$, $w \in W_{n-1}$, and use the commutation relation $X \cdot Y \cdot w = Y \cdot X \cdot w + [X, Y] \cdot w$; the claim follows by induction, since $[X, Y]$ is always in $\mathfrak{h}$. The subspace $W \subset V$ which is a union of all the W_n's is thus a subrepresentation; to see that it is irreducible; note that if we write $W = W' \oplus W''$, then either W' or W'' will have to contain the one-dimensional weight space W_α, and so will have to equal W.

The uniqueness of the highest weight vector of an irreducible representation follows immediately: if $v \in V_\alpha$ and $w \in V_\beta$ were two such, not scalar multiples of each other, we would have $l(\alpha) > l(\beta)$ and vice versa. □

Exercise 14.14. Show that in (ii) one need only apply those $\mathfrak{g}_\beta$ for which $\mathfrak{g}_\beta \cdot v \neq 0$. (Note: with W_n defined using only these $\mathfrak{g}_\beta$, and X in any root space, the same inductive argument shows that $X \cdot W_n \subset W_{n+1}$. On the other hand, if one uses all $\mathfrak{g}_\beta$ with β negative and primitive, as in Observation 14.16, then $X \cdot W_n \subset W_{n-1}$. One cannot combine these, however: V may *not* be generated by successively applying those $\mathfrak{g}_\beta$ with β negative, primitive, and $\mathfrak{g}_\beta \cdot v \neq 0$, e.g., the standard representation of $\mathfrak{sl}_3\mathbb{C}$.)

The weight α of the highest weight vector of an irreducible representation will be called, not unreasonably, the *highest weight* of that representation; the term *dominant weight* is also common.

We can refine part (ii) of this proposition slightly in another direction; this is not crucial but will be useful later on in estimating multiplicities of various representations. This refinement is based on

Exercise 14.15*. (a) Let $\alpha_1, \ldots, \alpha_k$ be roots of a semisimple Lie algebra $\mathfrak{g}$ and $\mathfrak{g}_{\alpha_i} \subset \mathfrak{g}$ the corresponding root spaces. Show that the subalgebra of $\mathfrak{g}$ generated by the Cartan subalgebra $\mathfrak{h}$ together with the $\mathfrak{g}_{\alpha_i}$ is exactly the direct sum $\mathfrak{h} \oplus (\bigoplus \mathfrak{g}_\alpha)$, where the direct sum is over the intersection of the set R of roots of $\mathfrak{g}$ with the semigroup $\mathbb{N}\{\alpha_1, \ldots, \alpha_k\} \subset \mathfrak{h}$ generated by the α_i.

(b) Similarly, let $\alpha_1, \ldots, \alpha_k$ be negative roots of a semisimple Lie algebra $\mathfrak{g}$ and $\mathfrak{g}_{\alpha_i} \subset \mathfrak{g}$ the corresponding root spaces. Show that the subalgebra of $\mathfrak{g}$ generated by the $\mathfrak{g}_{\alpha_i}$ is exactly the direct sum $\bigoplus \mathfrak{g}_\alpha$, where the direct sum is over the intersection of the set R of roots of $\mathfrak{g}$ with the semigroup $\mathbb{N}\{\alpha_1, \ldots, \alpha_k\} \subset \mathfrak{h}$ generated by the α_i.

(Note that by the description of the adjoint action of a Lie algebra on itself we have an obvious inclusion; the problem here is to show—given the facts above—that if $\alpha + \beta \in R$, then $[\mathfrak{g}_\alpha, \mathfrak{g}_\beta] \neq 0$.)

From this exercise, it is clear that generating a subrepresentation W of a given representation V by successive applications of root spaces $\mathfrak{g}_\beta$ for $\beta \in R^-$ to a highest weight vector v is inefficient; we need only apply the root spaces $\mathfrak{g}_\beta$ corresponding to a set of roots β generating R^- as a semigroup. We accordingly introduce another piece of terminology: we say that a positive (resp., negative) root $\alpha \in R$ is *primitive* or *simple* if it cannot be expressed as a sum of two positive (resp. negative) roots. (Note that, since there are only finitely many roots, every positive root can be written as a sum of primitive positive roots.) We then have

Observation 14.16. *Any irreducible representation V is generated by the images of its highest weight vector v under successive applications of root spaces $\mathfrak{g}_\beta$ where β ranges over the primitive negative roots.*

We have already seen one example of this in the case of $\mathfrak{sl}_3\mathbb{C}$, where we observed (in the proof of Claim 12.10 and in the analysis of $\mathrm{Sym}^2 V \otimes V^*$ in Lecture 13) that any irreducible representation was generated by applying the two elements $E_{2,1} \in \mathfrak{g}_{L_2 - L_1}$ and $E_{3,2} \in \mathfrak{g}_{L_3 - L_2}$ to a highest weight vector.

To return to our description of the weights of an irreducible representation V, we observe next that in fact *every vertex of the convex hull of the weights of V must be conjugate to α under the Weyl group.* To see this, note that by the above the set of weights is contained in the cone $\alpha + C_\alpha^-$, where C_α^- is the positive real cone spanned by the roots $\beta \in R^-$ such that $\mathfrak{g}_\beta(v) \neq 0$—that is, such that $\alpha(H_\beta) \neq 0$. Conversely, the weights of V will contain the string of weights

$$\alpha, \alpha + \beta, \alpha + 2\beta, \ldots, \alpha + (-\alpha(H_\beta))\beta \tag{14.17}$$

for any $\beta \in R^-$. Thus, any vertex of the convex hull of the set of weights of V adjacent to α must be of the form

$$\alpha - \alpha(H_\beta)\beta = W_\beta(\alpha)$$

for some β; applying the same analysis to each successive vertex gives the statement.

From the above, we deduce that the set of weights of V will lie in the convex hull of the images of α under the Weyl group. Since, moreover, we know that the intersection of this set with any set of weights of the form $\{\beta + n\gamma\}$ will be a connected string, it follows that *the set of weights of V will be exactly the weights that are congruent to α modulo the root lattice Λ_R and that lie in the convex hull of the images of α under the Weyl group.*

One more bit of terminology, and then we are done. By what we have seen (cf. (14.17)), the highest weight of any representation of V will be a weight α satisfying $\alpha(H_\gamma) \geq 0$ for every $\gamma \in R^+$. The locus $\mathscr{W}$, in the real span of the roots, of points satisfying these inequalities—in terms of the Killing form, making an acute or right angle with each of the positive roots—is called the (closed) *Weyl chamber* associated to the ordering of the roots. A Weyl chamber could also be described as the closure of a connected component of the complement of the union of the hyperplanes Ω_α. The Weyl group acts simply transitively on the set of Weyl chambers and likewise on the set of orderings of the roots. As usual, these statements will be easy to see in the cases we study, while the abstract proofs are postponed (to Appendix D).

Step 8. *Classify the irreducible, finite-dimensional representations of* $\mathfrak{g}$. Where all the above is leading should be pretty clear; it is expressed in the fundamental existence and uniqueness theorem:

Theorem 14.18. *For any α in the intersection of the Weyl chamber $\mathscr{W}$ associated to the ordering of the roots with the weight lattice Λ_W, there exists a unique irreducible, finite-dimensional representation Γ_α of $\mathfrak{g}$ with highest weight α; this gives a bijection between $\mathscr{W} \cap \Lambda_W$ and the set of irreducible representations of $\mathfrak{g}$. The weights of Γ_α will consist of those elements of the weight lattice congruent to α modulo the root lattice Λ_R and lying in the convex hull of the set of points in $\mathfrak{h}^*$ conjugate to α under the Weyl group.*

HALF-PROOF. We will give here just the proof of uniqueness, which is easy. The existence part we will demonstrate explicitly in each example in turn; and later on we will sketch some of the constructions that can be made in general.

The uniqueness part is exactly the same as for $\mathfrak{sl}_3\mathbb{C}$. If V and W are two irreducible, finite-dimensional representations of $\mathfrak{g}$ with highest weight vectors v and w, respectively, both having weight α, then the vector $(v, w) \in V \oplus W$ will again be a highest weight vector of weight α in that representation. Let $U \subset V \oplus W$ be the subrepresentation generated by (v, w); since U will again be irreducible the projection maps $\pi_1: U \to V$ and $\pi_2: U \to W$, being nonzero, will have to be isomorphisms. □

Another fact which we will see as we go along—and eventually prove in general—is that there are always *fundamental weights* $\omega_1, \ldots, \omega_n$ with the property that any dominant weight can be expressed uniquely as a non-negative integral linear combination of them. They can be characterized geometrically as the first weights met along the edges of the Weyl chamber, or algebraically as those elements ω_i in $\mathfrak{h}^*$ such that $\omega_i(H_{\alpha_j}) = \delta_{i,j}$, where $\alpha_1, \ldots, \alpha_n$ are the simple roots (in some order). When we have found them, we often write $\Gamma_{a_1,\ldots,a_n}$ for the irreducible representation with highest weight $a_1\omega_1 + \cdots + a_n\omega_n$; i.e.,

$$\Gamma_{a_1,\ldots,a_n} = \Gamma_{a_1\omega_1 + \cdots + a_n\omega_n}.$$

As with most of the material in this section, general proofs will be found in Lecture 21 and Appendix D.

One basic point we want to repeat here (and that we hope to demonstrate in succeeding lectures) is this: that actually carrying out this process in practice is completely elementary and straightforward. Any mathematician, stranded on a desert island with only these ideas and the definition of a particular Lie algebra $\mathfrak{g}$ such as $\mathfrak{sl}_n\mathbb{C}$, $\mathfrak{so}_n\mathbb{C}$, or $\mathfrak{sp}_{2n}\mathbb{C}$, would in short order have a complete description of all the objects defined above in the case of $\mathfrak{g}$. We should say as well, however, that at the conclusion of this procedure we are left without one vital piece of information about the representations of $\mathfrak{g}$, without which we will be unable to analyze completely, for example, tensor products of known representations; this is, of course, a description of the multiplicities of the basic representations Γ_α. As we said, we will, in fact, describe and prove such a formula (the Weyl character formula); but it is of a much less straightforward character (our hypothetical shipwrecked mathematician would have to have what could only be described as a pretty good day to come up with the idea) and will be left until later. For now, we will conclude this lecture with the promised introduction to the Killing form.

§14.2. About the Killing Form

As we said, the Killing form is an inner product (symmetric bilinear form) on the Lie algebra $\mathfrak{g}$; abusing our notation, we will denote by B both the Killing form and the induced inner products on $\mathfrak{h}$ and $\mathfrak{h}^*$. B can be defined in several ways; the most common is by associating to a pair of elements $X, Y \in \mathfrak{g}$ the trace of the composition of their adjoint actions on $\mathfrak{g}$, i.e.,

$$B(X, Y) = \mathrm{Tr}(\mathrm{ad}(X) \circ \mathrm{ad}(Y)\colon \mathfrak{g} \to \mathfrak{g}). \tag{14.19}$$

As we will see, the Killing form may be computed in practice either from this definition, or (up to scalars) by using its invariance under the group of automorphisms of $\mathfrak{g}$. We remark that this definition is not as opaque as it may seem at first. For one thing, the description of the adjoint action of the root space $\mathfrak{g}_\alpha$ as a "translation" of the root diagram—that is, carrying each root space $\mathfrak{g}_\beta$ into $\mathfrak{g}_{\alpha+\beta}$—tells us immediately that $\mathfrak{g}_\alpha$ is perpendicular to $\mathfrak{g}_\beta$ for all β other than $-\alpha$; in other words, the decomposition

$$\mathfrak{g} = \mathfrak{h} \oplus \left(\bigoplus_{\alpha \in R^+} (\mathfrak{g}_\alpha \oplus \mathfrak{g}_{-\alpha})\right) \tag{14.20}$$

is orthogonal. As for the restriction of B to $\mathfrak{h}$, this is more subtle, but it is not hard to write down: if X, Y are in $\mathfrak{h}$, and Z_α generates $\mathfrak{g}_\alpha$, then $\mathrm{ad}(X) \circ \mathrm{ad}(Y)(Z_\alpha) = \alpha(X)\alpha(Y)Z_\alpha$, so $B(X, Y) = \sum \alpha(X)\alpha(Y)$, the sum over the roots; viewing $B|_\mathfrak{h}$ as an element of the symmetric square $\mathrm{Sym}^2(\mathfrak{h}^*)$, we have

$$B|_\mathfrak{h} = \frac{1}{2} \sum_{\alpha \in R} \alpha^2. \tag{14.21}$$

A key fact following from this—one that, if nothing else, makes picturing $\mathfrak{h}^*$ with the inner product B involve less eyestrain—is

(14.22) *B is positive definite on the real subspace of $\mathfrak{h}$ spanned by the vectors $\{H_\alpha: \alpha \in R\}$.*

Indeed, all roots take on real values on this space (since all $\alpha(H_\beta) \in \mathbb{Z} \subset \mathbb{R}$), so for H in this real subspace of $\mathfrak{h}$, $B(H, H)$ is non-negative, and is zero only when all $\alpha(H) = 0$, which implies $H = 0$, since the roots span $\mathfrak{h}^*$.

To see that the Killing form is nondegenerate on all of $\mathfrak{g}$, we need the useful identity:

$$B([X, Y], Z) = B(X, [Y, Z]) \tag{14.23}$$

for all X, Y, Z in $\mathfrak{g}$. This follows from the identity

$$\mathrm{Trace}((\bar{X}\bar{Y} - \bar{Y}\bar{X})\bar{Z}) = \mathrm{Trace}(\bar{X}(\bar{Y}\bar{Z} - \bar{Z}\bar{Y}))$$

for any endomorphisms $\bar{X}$, $\bar{Y}$, $\bar{Z}$ of a vector space. And this, in turn, follows from

$$\mathrm{Trace}(\bar{Y}\bar{X}\bar{Z} - \bar{X}\bar{Z}\bar{Y}) = \mathrm{Trace}([\bar{Y}, \bar{X}\bar{Z}]) = 0.$$

An immediate consequence of (14.23) is that if $\mathfrak{a}$ is any ideal in a Lie algebra $\mathfrak{g}$, then its orthogonal complement $\mathfrak{a}^\perp$ with respect to B is also an ideal. In particular, if $\mathfrak{g}$ is simple, the kernel of B is zero (note that the kernel cannot be $\mathfrak{g}$ since it does not contain $\mathfrak{h}$). Since the Killing form of a direct sum is the sum of the Killing forms of the factors, it follows that *the Killing form is nondegenerate on a semisimple Lie algebra* $\mathfrak{g}$.

One of the reasons the Killing form helps to picture $\mathfrak{h}^*$ is the fact mentioned above:

Proposition 14.24. *With respect to B, the line spanned by each root α is perpendicular to the hyperplane Ω_α.*

As we observed, this is equivalent to saying that the involutions W_α above are simply reflections in hyperplanes, and in turn to saying that the whole Weyl group is orthogonal. Note also that Proposition 14.24 thereby follows immediately from the Fact 14.11: from the definition of B above, it is clearly invariant under any automorphism of $\mathfrak{g}$. Nevertheless, we would prefer not to rely on this fact; and anyway giving a direct proof of the proposition is not hard, in terms of the picture we have of the adjoint action of $\mathfrak{g}$ on itself. To prove the assertion $\alpha \perp \Omega_\alpha$, it suffices to prove the dual assertion that $H \perp H_\alpha$ for all H in the annihilator of α. But now by construction H_α is the commutator $[X_\alpha, Y_\alpha]$ of an element $X_\alpha \in \mathfrak{g}_\alpha$ and an element $Y_\alpha \in \mathfrak{g}_{-\alpha}$. Using (14.23) we have for any H in $\mathfrak{h}$,

$$\begin{aligned} B(H_\alpha, H) &= B([X_\alpha, Y_\alpha], H) = B(X_\alpha, [Y_\alpha, H]) \\ &= B(X_\alpha, \alpha(H)Y_\alpha) = \alpha(H)B(X_\alpha, Y_\alpha), \end{aligned} \tag{14.25}$$

which vanishes since $\alpha(H) = 0$.

Note that as a consequence of this, we can characterize the Weyl chamber associated to an ordering of the roots as exactly those vectors in the real span of the roots forming an acute angle with all the positive roots (or, equivalently, with all the primitive ones); the Weyl chamber is thus the cone whose faces lie in the hyperplanes perpendicular to the primitive positive roots.

Equation (14.25) leads to a formula for the isomorphism of $\mathfrak{h}$ with $\mathfrak{h}^*$ determined by the Killing form. First note that for $H = H_\alpha$ it gives

$$B(H_\alpha, H_\alpha) = 2B(X_\alpha, Y_\alpha) \neq 0,$$

for if $B(X_\alpha, Y_\alpha)$ were zero we would have $B(H_\alpha, H) = 0$ for all H, contradicting the nondegeneracy of B on $\mathfrak{h}$. The element T_α of $\mathfrak{h}$ which corresponds to $\alpha \in \mathfrak{h}^*$ by the Killing form is by definition the element of $\mathfrak{h}$ that satisfies the condition

$$B(T_\alpha, H) = \alpha(H) \quad \text{for all } H \in \mathfrak{h}. \tag{14.26}$$

Looking at (14.25), we see that $T_\alpha = H_\alpha/B(X_\alpha, Y_\alpha) = 2H_\alpha/B(H_\alpha, H_\alpha)$. This proves

Corollary 14.27. *The isomorphism of $\mathfrak{h}^*$ and $\mathfrak{h}$ determined by the Killing form B carries α to $T_\alpha = (2/B(H_\alpha, H_\alpha))\cdot H_\alpha$.*

The Killing form on $\mathfrak{h}^*$ is defined by $B(\alpha, \beta) = B(T_\alpha, T_\beta)$.

Exercise 14.28. Show that the inverse isomorphism from $\mathfrak{h}$ to $\mathfrak{h}^*$ takes H_α to $(2/B(\alpha, \alpha))\cdot\alpha$.

The orthogonality of W_α can be expressed by the formula

$$W_\alpha(\beta) = \beta - \frac{2B(\beta, \alpha)}{B(\alpha, \alpha)}\alpha.$$

Comparing with (14.8) this says:

Corollary 14.29. *If α and β are roots, then*

$$2B(\beta, \alpha)/B(\alpha, \alpha) = \beta(H_\alpha)$$

is an integer.

By the above identification of $\mathfrak{h}$ with $\mathfrak{h}^*$, (14.22) translates to

Corollary 14.30. *The Killing form B is positive definite on the real vector space spanned by the root lattice Λ_R.*

Note that it follows immediately from (14.22) that the Weyl group $\mathfrak{W}$ is finite, being simultaneously discrete ($\mathfrak{W}$ preserves the set R of roots of $\mathfrak{g}$ and hence the lattice Λ_R; it follows that $\mathfrak{W}$ can be realized as a subgroup of $\mathrm{GL}_n\mathbb{Z}$)

and compact ($\mathfrak{W}$ preserves the Killing form, and hence is a subgroup of the orthogonal group $O_n\mathbb{R}$.) Alternatively, $\mathfrak{W}$ is a subgroup of the permutation group of the set of roots.

As we observed, the Killing form on $\mathfrak{h}^*$ is preserved by the Weyl group. In fact, in case $\mathfrak{g}$ is simple, the Killing form is, up to scalars, the unique inner product preserved by the Weyl group. This will follow from

Proposition 14.31. *The space $\mathfrak{h}^*$ is an irreducible representation of the Weyl group $\mathfrak{W}$.*

PROOF. Suppose that $\mathfrak{z} \subset \mathfrak{h}^*$ were preserved by the action of $\mathfrak{W}$. This means that every root $\alpha \in \mathfrak{h}^*$ of $\mathfrak{g}$ will either lie in the subspace $\mathfrak{z}$ or be perpendicular to it, i.e., for every $\alpha \in \mathfrak{z}$ and $\beta \notin \mathfrak{z}$ we will have $\beta(H_\alpha) = 0$. We claim then that *the subspace $\mathfrak{g}'$ of $\mathfrak{g}$ spanned by the subalgebras $\{\mathfrak{s}_\alpha\}_{\alpha \in \mathfrak{z}}$ will be an ideal in $\mathfrak{g}$.* Clearly it will be a subalgebra; the space spanned by the distinguished subalgebras $\mathfrak{s}_\alpha$ corresponding to the set of roots lying in any subspace of $\mathfrak{h}^*$ will be. To see that it is in fact an ideal, let $Y \in \mathfrak{g}_\beta$ be an element of a root space. Then for any $\alpha \in \mathfrak{z}$, we have

$$[Y, Z] \in \mathfrak{g}_{\alpha+\beta} = 0$$

since $\alpha + \beta$ is neither in $\mathfrak{z}$ nor perpendicular to it, and so cannot be a root; and

$$[Y, H_\alpha] = -[H_\alpha, Y] = \beta(H_\alpha) \cdot Y = 0.$$

Thus, $\mathrm{ad}(Y)$ kills $\mathfrak{g}'$; since, of course, all of H itself will preserve $\mathfrak{g}'$, it follows that $\mathfrak{g}'$ is an ideal. Thus, either all the roots lie in $\mathfrak{z}$ and so $\mathfrak{z} = \mathfrak{h}^*$, or all roots are perpendicular to $\mathfrak{z}$ and correspondingly $\mathfrak{z} = (0)$. □

Note that given Fact 14.11, we can also express the last statement by saying that (in case $\mathfrak{g}$ is simple) the Killing form on $\mathfrak{h}$ is the unique form preserved by every automorphism of the Lie algebra $\mathfrak{g}$ carrying $\mathfrak{h}$ to itself. As we will see, in practice this is most often how we will first describe the Killing form.

Exercise 14.32. Find the Killing form on the Lie algebras $\mathfrak{sl}_2\mathbb{C}$ and $\mathfrak{sl}_3\mathbb{C}$ by explicit computation, and verify the statements made above in these cases.

Exercise 14.33*. If a semisimple Lie algebra is a direct sum of simple subalgebras, then its Killing form is the orthogonal sum of the Killing forms of the factors. Show that, conversely, if the roots of a semisimple Lie algebra lie in a collection of mutually perpendicular subspaces, then the Lie algebra decomposes accordingly.

Exercise 14.34*. Suppose $\mathfrak{g}$ is a Lie algebra that has an abelian subalgebra $\mathfrak{h}$ such that $\mathfrak{g}$ has a decomposition (14.1), satisfying the conditions of Facts 14.2 and 14.6. Show that $\mathfrak{g}$ is semisimple, and $\mathfrak{h}$ is a Cartan subalgebra.

The preceding exercise can be used instead of Weyl's unitary trick or any abstract theory to verify that the algebras we meet in the next few lectures are all semisimple. It is tempting to call such a Lie algebra "visibly semisimple."

The discussion of the geometry of the roots of a semisimple Lie algebra will be continued in Lecture 21 and completed in Appendix D. The Killing form becomes particularly useful in the general theory; for example, solvability and semisimplicity can both be characterized by properties of the Killing form (see Appendix C).

Exercise 14.35*. Show that $\mathfrak{b} = \mathfrak{h} \oplus \bigoplus_{\alpha > 0} \mathfrak{g}_\alpha$ is a maximal solvable subalgebra of $\mathfrak{g}$; $\mathfrak{b}$ is called a *Borel subalgebra.* Show that $\bigoplus_{\alpha > 0} \mathfrak{g}_\alpha$ is a maximal nilpotent subalgebra of $\mathfrak{g}$. These will be discussed in Lecture 25.

Exercise 14.36*. Show that the Killing form on the Lie algebra $\mathfrak{gl}_m$ is given by the formula

$$B(X, Y) = 2m \operatorname{Tr}(X \circ Y) - 2 \operatorname{Tr}(X) \operatorname{Tr}(Y).$$

Find similar formulas for $\mathfrak{sl}_m$, $\mathfrak{so}_m$, and $\mathfrak{sp}_m$, showing in each case that $B(X, Y)$ is a constant multiple of $\operatorname{Tr}(X \circ Y)$.

Exercise 14.37. If G is a real Lie group, the Killing form on its Lie algebra $\mathfrak{g} = T_e G$ may not be positive definite. When it is, it determines, by left translation, a Riemannian metric on G. Show that the Killing form is positive definite for $G = \mathrm{SO}_n\mathbb{R}$, but not for $\mathrm{SL}_n\mathbb{R}$.