Representation Theory

Elisha Peterson

February 19, 2004

Contents

1 Getting Oriented

Using mathematical formalism, a group G may be "represented" by transformations on a vector space:

2 Representations

2.1 The Basics

A representation of a group G is a vector space V for which one can view each $g \in G$ as a transformation $g : V \to V$.

> **Representation:** a pair (π, V), where V is a finite dimensional vector space over $\mathbb{C}$, and π is a continuous homomorphism $\pi : G \to GL(V)$ into the Lie group $GL(V)$ of invertible linear transformations.

We will write g for an element of G and π_g for its corresponding transformation. The **dimension** of the representation is the dimension of the vector space V. Simple examples of representations include:

- The **trivial representation** with $\pi_g = I$ or $g(v) = v$ for all $v \in V$.
- The **standard representation** for $G \subset GL(n, \mathbb{C})$ with $g(v) = g \cdot v$ for all $v \in \mathbb{C}^n$.
- A *group action* $G \times X \to X$ provides a representation with $V = \{f : X \to \mathbb{C}\} = X^*$ and $(\pi_g f)(x) = f(g^{-1}x)$. Thus, g takes the map $x \mapsto f(x)$ to the map $x \mapsto f(g^{-1}x)$.

2.2 Reducible and Unitary Representations

Given an **invariant subspace** U of a representation (π, V), meaning $\pi_g u \in U$ for all $u \in U$, one may restrict (π, V) to the **subspace representation** (π', U) with $\pi'_g = \pi_g|_U$. The existence of invariant subspaces allows one to break up a representation into smaller pieces. Those without invariant subspaces are:

> **Irreducible representation:** a (π, V) with no nontrivial invariant subspaces.

As a simple example, a 1-dimensional representation must be irreducible.

Irreducible representations form the building blocks for all others, just as primes do for integers. And as we ask how to decompose a given integer into primes, it is natural to ask when and how we can write $V = V_1 \oplus \cdots \oplus V_n$ where the V_i are all irreducible invariant subspaces. Here are some things to consider:

- By extending a basis for an invariant subspace U to a basis for V, we can write all π_g in block-upper diagonal form $[**; 0*]$.
- If we can write $V = U_1 \oplus U_2$, we can similarly write π_g in block-diagonal form $[*0; 0*]$.
- Not every invariant subspace provides such a decomposition. Letting $G = (\mathbb{R}, +) = \{(1a; 01) : a \in \mathbb{R}\}$ with the standard representation, we have invariant subspace $U = \{(x; 0)\}$ but its complement $\{(0; y)\}$ clearly not.

It turns out that a representation (π, V) is **completely reducible**, meaning it can be broken down into irreducible invariant subspace representations, when it finite-dimensional and invariant under a certain inner product:

> **Unitary representation:** a representation (π, V) for which $\langle \pi_g v, \pi_g w \rangle = \langle v, w \rangle$ for some positive-definite Hermitian inner product $\langle , \rangle$. To be Hermitian, it must satisfy $\langle v, w \rangle = \sum v_i \bar{w}_i$ and $\langle v, w \rangle = \overline{\langle w, v \rangle}$. This is equivalent to requiring the matrices π_g to be unitary $\pi_g^{-1} = \pi_g^*$ or $\langle \pi_g v, w \rangle = \langle v, \overline{\pi_g^T} w \rangle$.

The finite-dimensional unitary representations are completely reducible for the following simple fact:

- if U is an invariant subspace, then $U^\perp = \{v : \langle u, v \rangle \quad \forall \quad u \in U\}$ is also invariant.

Thus, an algorithm for finding invariant subspaces provides an algorithm for finding the irreducible ones, but *only in the finite-dimensional case.*

2.3 Some Completely Reducible Representations

In order to show that a representation is completely reducible, it suffices to show that it is unitary. In the case of a finite group G, this is easy: take any positive-definite Hermitian inner product $(\, , \,)$ on V and *average* it over the group by defining:

$$\langle v, w \rangle = \frac{1}{|G|} \sum_{g \in G} (\pi_g v, \pi_g w).$$

It is simple to show that $\langle \, , \, \rangle$ is again positive-definite and Hermitian and it is obvious from the definition that it is invariant under $\pi(x)$ for $x \in G$.

This averaging trick also works for any compact group G, although we must now integrate. Again, assuming the existence of a non-degenerate Hermitian inner product $(\, , \,)$ on V, we define

$$\langle v, w \rangle = \int_G (\pi_g v, \pi_g w) d\mu(g).$$

The catch is that we must the measure μ must satisfy certain conditions for this new inner product to be π-invariant. The appropriate notion is:

> **(Right) Haar measure:** a measure μ on a locally compact topological group G defined on the Borel subsets $\mathcal{B}$ and satisfying (1) finiteness $\mu(K) < \infty$ for K compact, (2) non-degeneracy $\mu(U) > 0$ for U non-empty and open, and (3) translation invariance $\mu(Eg) = \mu(E)$ for $E \in \mathcal{B}$ and $g \in G$.

For such a measure, we have $\int_G f(x) d\mu(x) = \int_G f(xg) d\mu(x)$, making the above inner product invariant under $\pi(x)$.

Actually, a Haar measure is unique up to a constant factor, usually chosen so that $\mu(G) = 1$. In the finite case, we have $\mu(E) = \frac{|E|}{|G|}$, and this process gives the same inner product as above.

2.4 Dual and Tensor Representations

An arbitrary representation (π, V) of G gives rise to a **dual representation** $(\check{\pi}, V^*)$ with V^* the dual vector space to V and $\check{\pi}$ the **contragradient** defined by $\langle \check{\pi}_g y, x \rangle = \langle y, \pi_{g^{-1}}) x \rangle$. Note that $y \in V^*$ and $x \in V$, so this inner product is given by $\langle y, x \rangle = y(x)$. In terms of matrices (using a given basis and the dual basis), we have $\check{\pi}_g = [\pi_g^{-1}]^T$. Also, an invariant subspace $U \subset V$ corresponds to an invariant subspace $U^\perp \subset V^*$ where $U^\perp = \{y \in V^* : \langle y, x \rangle = 0 \quad \forall \quad x \in U\}$.

Two representations (π, V) of G and (ρ, W) of H give a **tensor representation** $(\pi \otimes \rho, V \otimes W)$ of $G \times H$ defined in the obvious manner by $(\pi \otimes \rho)_{(g,h)}(v \otimes w) = \pi_g v \otimes \pi_h w$. In the case that $G = H$, we also have a representation on G with $(\pi \otimes \rho)_g(v \otimes w) = \pi_g v \otimes \pi_g w$. A major question of interest in representation theory is:

- Given (irreducible) representations V and W, how does one decompose $V \otimes W$ in terms of irreducible invariant subspaces?

2.5 Equivalence of Representations

A map between representations (π, V) and (ρ, W) must respect the structure of G, and thus is called a G-**map**, meaning a linear map $A : V \to W$ such that $A(\pi_g v) = \rho_g(Av)$. These maps form the group $\mathrm{Hom}_G(V, W)$. This also provides the notion of equivalence:

> **Equivalent Representations:** those for which there exists a G-map $A : V \to W$ which is also a vector space isomorphism (a G-**isomorphism**).

In this case, the matrix of π_g with respect to basis $\{v_i\}$ and that of ρ_g with respect to basis $\{Av_i\}$ coincide. This brings up another major problem in representation theory:

- Find all irreducible representations of a given group G up to equivalence.

A first step in solving this problem is the following intuitive theorem:

> **Schur's Lemma:** Let a group G and a G-map $A \in \mathrm{Hom}_G(V, W)$ between irreducible representations (π, V) and (ρ, W) be given. We have:
>
> 1. If $\pi = \rho$, then $A = \gamma I$ for some $\gamma \in \mathbb{C}$;
> 2. If $\pi \approx \rho$ (just an equivalence), then $\dim_{\mathbb{C}}[\mathrm{Hom}_G(V, W)] = 1$;
> 3. Otherwise, $\pi \not\approx \rho$, and all such A vanish, so that $\mathrm{Hom}_G(V, W) = \{0\}$.

The proof relies on the fact that G-maps preserve invariant subspaces, such as $\ker A$, $\mathrm{im}\ A$, and λ-eigenspaces. This lemma is a simple but powerful tool in representation theory and will be widely used. It follows, for example, that all irreducible representations of abelian groups are 1-dimensional, since all 1-dimensional subspaces are invariant.

Following abelian groups, the next simplest case is that of finite groups. Representations in this case are also easily classified, and are a good model for the general theory. Before we turn to such taxonomical questions, however, we will look at classifying only *pieces* of representations... specifically *class functions...* in the Peter-Weyl Theorem.

3 The Peter-Weyl Theorem

The Peter-Weyl Theorem is really just a generalization of Fourier series:

- Fourier: any function $S^1 \to \mathbb{C}$ from the circle to the complex numbers can be approximated by a finite sum of exponentials.
- Peter-Weyl Theorem: any **class function** on a compact group G (those with $f(gxg^{-1}) = f(x)$) can be approximated by a sum of **characters** (a character is the trace of a representation).

There are actually many ways to write the Peter-Weyl Theorem, and the classical statement says nothing about class functions and characters. Rather, this result will arise as a corollary.

3.1 Fourier Series

We can reinterpret Fourier series in the language of representation theory. First, our group is $G = \mathbb{R}/\mathbb{Z} = S^1$, which is abelian. Hence, all representations are 1-dimensional and of the form $e_m : G \to \mathbb{C}^\times$ with $e_m(x) = e^{imx}$ for $m \in \mathbb{Z}$.

We can identify the representation e_m with the vector space $V_m = \mathbb{C}e_m$, which is invariant under translation: $[R(g)e_m](x) = e_m(x+g) = e_m(x)e_m(g) = \lambda_g e_m(x) \in \mathbb{C}e_m$. Then, Fourier theory states that

$$L^2(G) = \bigoplus \sum_m \mathbb{C}e_m,$$

where $L^2(G)$ is the space of square-integrable functions on S^1 with respect to the standard Haar measure $\int_G f(g)dg = \frac{1}{2\pi}\int_0^{2\pi} f(x)dx$. Moreover, $\{e_m\}$ is a *complete orthonormal basis* for $L^2(G)$.

As we have said, this example is a prototype for the Peter-Weyl Theorem, the stated form of which will mimic the $L^2(G) = \bigoplus \sum_m \mathbb{C}e_m$ decomposition.

3.2 Preliminary Results

In general, we will assume that G is a compact group. This compactness implies that left and right Haar measures on G are equivalent; thus, all Haar measures are invariant under both left $\mu(gE) = \mu(E)$ and right $\mu(Eg) = \mu(E)$ translations.

We now give a slew of results which lead up to the Peter-Weyl Theorem, and can be thought of as extensions of Schur's Lemma:

- (Lemma) For irreducible representations (π, V) and (ρ, W) of a compact group G, and $A \in \mathrm{Hom}_G(V, W)$ a G-map, set $A^0 = \int_G \rho(g^{-1}A\pi(g)dg$ for a normalized Haar measure. Then, $A^0 = 0$ if $\pi \not\approx \rho$, but $A^0 = (\frac{\mathrm{tr}(A)}{\deg \pi}) \cdot I$ if $\pi = \rho$.

  The map A^0 is a sort of average of A between the two representations. The proof is straightforward: one must show that A^0 is a G-map, so Schur's Lemma can be applied. Then one must show how the lemma's claims for the G-map A extend to A^0.

- (Corollary) If $v \in V, v^* \in V^*, w \in W, andw^* \in W^*$, then $\int_G \langle \pi(g)v, v^*\rangle\langle \rho(g^{-1})w, w^*\rangle dg$ equals 0 if $\pi \not\approx \rho$ but $\frac{1}{\deg \pi}\langle v, w^*\rangle\langle w, v^*\rangle$ if $\pi = \rho$.

  This result follows from the previous lemma using the map $A(x) = \langle x, v^*\rangle w$. A slight change in language for unitary representations gives the next result.

- (Corollary) Likewise, for unitary representations and Hermitian inner product $(\cdot, \cdot)$, we may take $x_i \in V, y_i \in W$ and obtain $\int_G (\pi(g)x_1, x_2)\overline{(\rho(g)y_1, y_2)}dg$ equals 0 if $\pi \not\approx \rho$ but $\frac{1}{\deg \pi}(x_1, y_1)\overline{(x_2, y_2)}$ if $\pi = \rho$.

- (Schur Orthogonality Relation) Here we restrict to matrix elements: if $\pi_{ij}(g)$ and $\rho_{kl}(g)$ are matrix elements of π and ρ with respect to orthonormal bases of V and W, then $\int_G \pi_{ij}(g)\overline{\rho_{kl}(g)}dg$ equals 0 if $\pi \not\approx \rho$ or $\frac{1}{\deg \pi}\delta_{ik}\delta_{jl}$ if $\pi = \rho$.

  This follows from the above by noting that for basis elements e_i and e_j, the matrix element is given by $\pi_{ij}(g) = (\pi(g)e_j, e_i)$.

3.3 The Theorem

We are finally ready for the main result:

> **Peter-Weyl Theorem:** Let G be a compact group, and $\{(\pi^\lambda, V^\lambda)\}_{\lambda\in\Lambda}$ a complete set of inequivalent unitary representations of G. Given an orthonormal basis $\{e_i^\lambda\}$ for V^λ, define $\pi_{ij}^\lambda(g) = (\pi^\lambda(g)e_j^\lambda, e_i^\lambda)$. Then, $\{f_{ij}^\lambda\}$ where $f_{ij}^\lambda = \sqrt{\deg \pi^\lambda}\pi_{ij}^\lambda$ is a *complete* orthonormal set in $L^2(G)$, and Λ must be countable.

For the proof, first note that $\{f_{ij}^\lambda\}$ is clearly orthonormal by the Schur Orthogonality Relation given above. It remains to show completeness and countability. In a *separable* Hilbert space, every orthonormal set is countable, and $L^2(G)$ is separable since the set of rational polynomials in $L^2(G)$ forms a countable dense subset.

Actually, all polynomials are matrix coefficients: if $p(x)$ is a polynomial of degree d, then (R, V_d) where R is right translation $R_g(q)(x) = q(xg)$ and V_d is the set of polynomials of degree $\leq d$ is a representation. Moreover, selecting e^* such that $\langle e^*, q\rangle = q(I)$ gives $p(x) = \langle R_g(q), e^*\rangle$ meaning p is a matrix coefficient. So the matrix coefficients are dense in $L^2(G)$ since the polynomials are, and this verifies completeness.

3.4 Consequences for Characters

The Peter-Weyl Theorem allows us to decompose $L^2(G)$ into subspaces invariant under G. Specifically, we can write

$$L^2(G) = \bigoplus \sum_{\lambda \in \Lambda} M(V^\lambda),$$

where $M(V^\lambda)$ is spanned by the matrix coefficients of the representation (π^λ, V^λ). Note that $M(V^\lambda)$ is invariant under the action of G by left/right translation.

The Theorem also has consequences for characters and class functions. Characters are functions in G given by the trace of a representation, while a class function is any function invariant under conjugation:

Character: given a finite-dimensional representation (π, V), it is the function $\chi_\pi : G \to G$ given by $g \mapsto \mathrm{tr}(\pi_g)$.

Class Function: a function $f : G \to \mathbb{C}$ such that $f(gxg^{-1}) = f(x)$.

Of course, all characters are class functions since $\mathrm{tr}(ABA^{-1}) = \mathrm{tr}(B)$. Some properties of characters follow:

- Equivalent representations have the same character: $\pi \approx \rho$ implies $\chi_\pi = \chi_\rho$. Actually, the reverse implication also holds.
- For direct sum and tensor product representations: $\chi_{\pi\oplus\rho} = \chi_\pi + \chi_\rho$ and $\chi_{\pi\otimes\rho} = \chi_\pi \chi_\rho$.
- For the dual representation: $\chi_{\tilde{\pi}}(g) = \chi_\pi(g^{-1})$; for unitary representations, $\chi_{\tilde{\pi}}(g) = \overline{\chi_\pi(g)} = \chi_\pi(g^{-1})$.
- (Schur Orthogonality for Characters) Given irreducible representations (π, V) and (ρ, W) of G, the characters satisfy $(\chi_\pi, \chi_\rho) = \delta_{\pi\rho}$, 1 if $\pi \approx \rho$ and 0 if $\pi \not\approx \rho$.

4 Representations of Finite Groups

Representation theory for finite groups contains a surprisingly large amount of the theory used in the more general case. We have already seen that all finite representations are unitary using the averaging technique

$$\langle v, w \rangle = \frac{1}{|G|} \sum_{g \in G} (\pi_g v, \pi_g w),$$

a fact which will bring us to several more formula allowing us to completely classify the irreducible representations of a finite group.

5 Computational Examples

We now turn to representations of non-finite groups, and illustrate a few cases which will be especially demonstrative of the general classification theory.

5.1 Representations of $sl_2\mathbb{C}$

The representations of $sl_2\mathbb{C}$ are a very important case because they indicate the fundamental structure behind every representation of a Lie algebra. The eigenvalues of any irreducible representation form a neat structure. In the case of $sl_2\mathbb{C}$, these eigenvalues are colinear and occur on the nice integer lattice, while in the more general case they also have a very nice structure on a (higher-dimensional) lattice.

The best part about the irreducible representations of $sl_2\mathbb{C}$ is that there is one, and only one, for each integer $n \geq 0$, and we can describe how to write down every one of them. Our basic strategy involves three steps:

- Find a suitable abelian subspace h of the Lie algebra. Break the Lie algebra into pieces corresponding to this h. Check that all these pieces preserve the given decomposition.
- Decompose V in terms of the eigenspaces of h. Check that this decomposition is preserved by the other pieces of the Lie algebra.
- Determine the possible configuration of the eigenspaces by examining their eigenvalues.

Actually, these are the three steps used to write down the representations of *any* Lie algebra, although it certainly becomes more complicated. This is especially true since W will be cyclic for the case at hand, but not in general.

Step 1:

For our first task, we write down the following basis of $sl_2\mathbb{C}$:

$$H = (1, 0; 0, -1), X = (0, 1; 0, 0), Y = (0, 0; 1, 0).$$

Our abelian subspace h consists of the diagonal elements and is generated by H. Note that we have $[H, X] = 2X, [H, Y] = -2Y$, and $[X, Y] = H$.

Step 2:

The action of H on our representation V is diagonalizable (a Jordan decomposition argument), so we may decompose $V = \oplus V_\alpha$ where $H(v) = \alpha \cdot v$ for all $v \in V_\alpha$.

Step 3:

Now, the question is how the rest of $sl_2\mathbb{C}$ acts on V_α; specifically, we must consider X and Y. We have

$$H(X(v)) = X(H(v)) + [H, X](v) = X(\alpha \cdot v) + 2X(v) = (\alpha + 2)X(v),$$

and so $X : V_\alpha \to V_{\alpha+2}$. Likewise, for Y we have $Y : V_\alpha \to V_{\alpha-2}$.

Finite-dimensionality and irreducibility combine to give the existence of a vector v such that $\{v, Y(v), Y^2(v), \ldots\}$ spans V. Here, we let $v \in V_\beta$, the largest nonzero eigenspace. This is a finite-dimensional vector space, and invariant under the actions of X, Y, and H, hence under all of $sl_2\mathbb{C}$. Therefore, it must be all of V.

This means that all eigenspaces are 1-dimensional, and that V is the direct sum of $V_\beta, V_{\beta-2}, \ldots, V_{\beta-2k}$. Now, if $w \in V_\beta$, then $Y^{k+1}(w) = 0$ but $Y^k(w) \neq 0$, and so $X(Y^{k+1}(v)) = 0 = k(\beta - k + 1)Y^k(v)$. Thus, $\beta = k - 1$ is an integer, and an irreducible representation $V^{(n)}$ looks like:

$$V^{(n)} = V_{-n} \oplus V_{-n+2} \oplus \cdots \oplus V_{n-2} \oplus V_n.$$

We can actually show that $V^{(n)}$ is isomorphic to the representation $\mathrm{Sym}^n V$ by analyzing eigenvalues. It is clear that a representation is uniquely defined by its set of eigenvalues. This fact allows us to solve the *Clebsch-Gordon Problem*: write down formulae for various products of representations in terms of the irreducibles. In particular, since the eigenvalues of $W \otimes W'$ are just the sum of those in W and those in W', we have:

$$\mathrm{Sym}^a V \otimes \mathrm{Sym}^b V = \mathrm{Sym}^{a-b} V \oplus \cdots \mathrm{Sym}^{a+b-2} V \oplus \mathrm{Sym}^{a+b} V.$$

We also have $\mathrm{Sym}^n(\mathrm{Sym}^2 V) = \bigoplus_{a=0}^{[n/2]} \mathrm{Sym}^{2n-4a} V$. One can also obtain these formulae through algebraic geometry...

5.2 Representations of $sl_3\mathbb{C}$

As might be expected, raising the dimension greatly complicates the problem of writing down the irreducible representations. It is still quite possible, however, and still involves the same three basic steps. The approach must be generalized somewhat, however, and also streamlined. We will introduce a few new terms here which occur in the general theory.

Step 1:

The first difficulty to overcome is that the abelian subspace h from above is no longer generated by a single element. Thus, we must redefine our notions of eigenvector, eigenvalue, and eigenspace to accommodate several generators. Given an abelian subspace h, an *eigenvector* will satisfy the equation $H(v) = \alpha_H \cdot v$ for all $H \in \mathsf{h}$. Of course, now α_H depends on H, and thus can be viewed as a element $\alpha \in \mathsf{h}^*$; so the *eigenvalues* are now elements of the dual space rather than scalars.

Remarkably, for a suitable choice of h in a general Lie algebra, the action of h on the Lie algebra g and all its representations is *simultaneously diagonalizable*. What this means is that we can decompose $sl_3\mathbb{C} = \mathsf{h} \oplus (\oplus \mathsf{g}_\alpha)$, where g_α is an eigenspace for the *adjoint action* of h: whenever $H \in \mathsf{h}$, $Y \in \mathsf{g}_\alpha$ we have $\mathsf{ad}(H)(Y) = [H, Y] = \alpha_H \cdot Y$.

Our choice of h will (again) be the set of diagonal matrices. Finding the g_α is not difficult; the only elements M with $[D, M] = \lambda M$ for all diagonal matrices D are those only one nonzero entry. Hence, the subspaces g_α are generated by the matrix E_{ij} whose only nonzero entry is $e_{ij} = 1$. Letting $D = (a_1, 0, 0; 0, a_2, 0; 0, 0, a_3)$, the eigenvectors are elements of the space h^* generated by $L_i : D \mapsto a_i$. Since $[D, E_{ij}] = (a_i - a_j)E_{ij}$, we see that E_{ij} generates the space $\mathsf{g}_{L_i - L_j}$. In particular, we have:

$$sl_3\mathbb{C} = \mathsf{h} \oplus \mathsf{g}_{L_1-L_2} \oplus \cdots \oplus \mathsf{g}_{L_3-L_2}.$$

Clearly, h preserves this eigenspace decomposition, but it is also true that each g_α preserves it. In fact, by considering elements of the given subspaces, one can show that $\mathrm{ad}(\mathsf{g}_\alpha) : \mathsf{g}_\beta \to \mathsf{g}_{\alpha+\beta}$.

Step 2:

Now, we consider an arbitrary representation V of $sl_3\mathbb{C}$, and break it up into eigenspaces $V = \oplus V_\alpha$, so that whenever $v \in V_\alpha$ and $H \in h$ we have $Hv = \alpha_H \cdot v$. When $X \in \mathsf{g}_\beta$ is outside h, we have $H(X(v)) = (\alpha_H + \beta_H) \cdot X(v)$, and so the action of g_β on the representation carries V_α to $V_{\alpha+\beta}$. In particular, we see that the eigenvalues of the representation are all in some translate of the lattice Λ_R in h^* generated by $\{L_i - L_j\}$.

Now for some terminology. The eigenvalues $L_i - L_j$ of the adjoint action on $sl_3\mathbb{C}$ are called **roots**, with the corresponding $\mathsf{g}_{L_i-L_j}$ **root spaces**. The lattice formed by the roots is called the **weight lattice**, and denote Λ_R. For an arbitrary representation, the eigenvalues $\alpha \in \mathsf{h}^*$ are called **weights**, and the V_α **weight spaces**. So, the roots determine how we decompose the Lie algebra, while the weights determine how we decompose the representation.

Step 3:

To finish the proof requires a generalization of an 'extremal' vector. In one dimension, the notion is trivial, but here we must choose a linear functional which is positive on half the roots and negative on the other half. The mechanism used to define an extreme vector is not so important, however. What is important is that there is a vector $v \in V_\alpha$ which is killed by each of E_{12}, E_{13}, and E_{23}. This corresponds to the case for $sl_2(\mathbb{C})$ when v was the extremal vector killed by X.

Given this, we see as before that V is generated by the images of v under the opposite elements E_{21}, E_{31}, and E_{32}. This restricts us to a sort of 'corner region' of h^*. There are a few more key observations to make:

- The eigenspaces on the boundaries of the region correspond to irreducible representations on $sl_2\mathbb{C}$.
- There is a symmetry about the three lines $\langle [E_{ij}, E_{ji}], L\rangle = 0$.

This implies that the weights (eigenvalues) of the representations are contained in either a hexagon or a triangle (with $\mathbb{Z}_3$ symmetry). It is clear that the interior must be filled out, and the interior weights must have multiplicity 1. This completes the $sl_3\mathbb{C}$ case.

6 The General Case

In order to classify all (irreducible/finite-dimensional) representations for an arbitrary (semisimple) Lie algebra, one must take steps analogous to those above for $sl_(2\mathbb{C})$ and $sl_3(\mathbb{C})$. The theory is extremely practical, and provides a mechanical calculation of possibilities. We will go through the steps in detail here.

6.1 The Cartan Decomposition (Step 1)

We need to find a suitable abelian subspace, analogous to the set of diagonal matrices for $sl_n(\mathbb{C})$. Our choice is restricted to maximal abelian subalgebras which act diagonally on the adjoint representation. This is called a **Cartan subalgebra**, and is a well-defined piece of every semisimple Lie algebra.

The Cartan subalgebra h is used to define the **Cartan decomposition** of the Lie algebra: $\mathsf{g} = \mathsf{h} \oplus (\oplus \mathsf{g}_\alpha)$. There will be a finite number of eigenvalues α called **roots**, with corresponding subspaces g_α called **root spaces**. As we saw before, the roots comprise a subset $R \subset \mathsf{h}^*$ of the dual space and their configuration therein is a useful depiction of the structure of g. Note that the adjoint action of g_α takes g_β to $\mathsf{g}_{\alpha+\beta}$, so that the decomposition is preserved by the action of h.

We note a few facts regarding this decomposition. First, each root space g_α is one-dimensional. Second, R generates a lattice $\Lambda_R \subset \mathsf{h}^*$ with rank equal to $\dim(\mathsf{h})$, meaning the roots span h^*. Last, α is a root iff $-\alpha$ is also a root.

6.2 Decomposing the Representation (Step 2)

We now turn our attention to an arbitrary representation V, and note that the action of h on V gives the decomposition $V = \oplus V_\alpha$. The eigenvalues are now called **weights** and the V_α are called **weight spaces**. The dimension of V_α is called the **multiplicity of the weight**. A simple diagram in h^* suffices to depict the weights and their dimensions.

Note that the weights are all congruent modulo the root lattice, meaning they differ by linear combinations of the roots of the Lie algebra. This is related to the fact that a piece g_β of the Cartan decomposition takes V_α into $V_{\alpha+\beta}$.

6.3 Analyzing the Configuration of Eigenvalues/Weights (Step 3)

This third step is by far the most detailed, although it is really not much harder than the previous two. It basically amounts to finding an extremal piece of the representation (which turns out to be an $sl_2(\mathbb{C})$ representation, and using it to fill in the eigenvalues for the rest of the representation.

First, note that for any root space g_α, we have a corresponding subalgebra $s_\alpha = \mathsf{g}_\alpha \oplus \mathsf{g}_{-\alpha} \oplus [\mathsf{g}_\alpha, \mathsf{g}_{-\alpha}]$ which is, in fact, isomorphic to $sl_2(\mathbb{C})$. To see this, one can choose elements $X_\alpha \in \mathsf{g}_\alpha$, $Y_\alpha \in \mathsf{g}_{-\alpha}$, and $H_\alpha \in [\mathsf{g}_\alpha, \mathsf{g}_{-\alpha}]$ which will act exactly like $X, Y, H \in sl_2(\mathbb{C})$. Thus, when we restrict to s_α we obtain a representation of $sl_2(\mathbb{C})$. In particular, all eigenvalues $\beta \in \mathsf{h}^*$ of the representation must be $\mathbb{Z}$-valued on all elements H_α.

We can say more than that about the eigenvalues of H_α, however: they are symmetric about the hyperplane $\{\langle H_\alpha, \beta\rangle = 0\}$. Letting $W_\alpha(\beta) = \beta - \frac{2\beta(H_\alpha)}{\alpha(H_\alpha)}\alpha = \beta - \beta(H_\alpha)\alpha$ be the reflection in this hyperplane, we see that α-equivalence classes of V such as $V_{[\beta]} = \oplus V_{\beta+n\alpha}$ are all invariant under the action of s_α, and so fixed by the W_α involution. Moreover, from the structure of $sl_2(\mathbb{C})$ representations, one sees that each equivalence class can be written $V_{[\beta]} = V_{\beta-m} \oplus V_{\beta-m+2} \oplus \cdots \oplus V_{\beta+m}$.

The set of such hyperplane reflections W_α forms a group called the **Weyl Group**. In general, the set of weights, and the multiplicities of the weights, of any representation of g is invariant under the action of the Weyl group. One can see this by noting that it is true for each equivalence class $V_{[\beta]}$. Here is another fact: every element of the Weyl group is induced by an automorphism of g which fixes h.

7 The Road Ahead