AN INTRODUCTION TO LIE GROUPS

AARUSH KULKARNI

ABSTRACT. In basic terms, Lie groups are the language of continuous symmetry. The main idea is that these non-linear structures can be studied through a linear approximation called the Lie algebra, which is just the tangent space at the identity. We will begin by discussing the correspondence between a group and its algebra using the exponential map, (touching on the Baker-Campbell-Hausdorff formula and the famous covering homomorphism from $SU(2)$ to $SO(3)$). We conclude with a brief look at the bigger picture, such as the structure theory of Lie algebras and their role in representation theory.

CONTENTS

1. What Is Continuous Symmetry? 2
2. The Building Blocks of a Lie Group 2
3. Linearizing a Lie Group 3
4. Why Does It Take a 720° Turn to Get Back to Where You Started? 6
5. Structure and Representation Theory 7
6. The ‘Periodic Table’ of Simple Lie Algebras 8
7. Classifying Irreducible Representations 9
8. Homogeneous and Symmetric Spaces 10
Acknowledgments 11
References 11

1. What Is Continuous Symmetry?

While the symmetries of a crystal lattice are discrete (countable), the symmetries of physical laws under rotations or translations in spacetime, however, are continuous (uncountable). The role of the *Lie group* is to handle the continuous case that group theory could not. Lie theory was named after the Norwegian mathematician Sophus Lie. A Lie group is a smooth manifold that is also a group, with the two structures playing nicely together. In fact, Lie's original motivation grew from very concrete geometric problems involving line complexes and sphere mappings [2].

A Lie group is very powerful. We can use tools from calculus to kill problems in group theory, or use algebra to get a handle on geometry. Differential geometry teaches us how to think about curves and surfaces, which are manifolds of dimension 1 and 2. Lie theory is what appears when that geometric intuition is applied to higher-dimensional manifolds that have a group structure [4].

Our goal here is to provide an initial understanding of the subject. We will define the main objects, but the real aim is to understand the connection between a Lie group and its "linearization," the Lie algebra. This relationship is the bedrock of the entire theory.

2. The Building Blocks of a Lie Group

To properly define a Lie group, we need to merge ideas from two different worlds: abstract algebra and differential geometry. Let us briefly recall the key concepts from each.

Groups. An abstract group is a set with a binary operation satisfying a few simple rules. It is a sparse definition, but powerful enough to describe every kind of symmetry we know.

Definition 2.1. A *group* is a set G together with a binary operation $m : G \times G \to G$, where the image of (g, h) is written as gh, satisfying:

1. Associativity, $(gh)k = g(hk)$ for all $g, h, k \in G$.
2. An identity element, such that there exists an element $e \in G$ such that $eg = ge = g$ for all $g \in G$.
3. An inverse element, such that for each $g \in G$, there exists an element $g^{-1} \in G$ such that $gg^{-1} = g^{-1}g = e$.

Smooth manifolds. A smooth manifold is essentially a space that, when zoomed in far enough on any point, starts to look like ordinary Euclidean space. This local "flatness" is what allows us to do calculus on them. A good first picture to keep in mind is the surface of a sphere or a torus in $\mathbb{R}^3$.

Definition 2.2. A subset $S \subset \mathbb{R}^3$ is a *surface* if for every point $p \in S$, there is an open neighborhood W of p in $\mathbb{R}^3$ and a homeomorphism, which is a continuous bijection with a continuous inverse, from an open set $U \subset \mathbb{R}^2$ to $S \cap W$. Such a map is called a *surface patch* or chart.

More generally, an n-dimensional *topological manifold* is a Hausdorff, second-countable topological space that is locally homeomorphic to $\mathbb{R}^n$. To make it *smooth*, we require that the transition maps between any two overlapping charts are infinitely differentiable. This collection of compatible charts is called an *atlas*. Figure 1 illustrates this concept.

Definition 2.3. For a point p on a surface S, the *tangent space* T_pS is the plane of all possible velocity vectors of curves on S that pass through p. If a patch is given by $\sigma : U \to S$, the tangent space at $\sigma(u, v)$ is the two-dimensional vector space spanned by the partial derivatives $\{\sigma_u, \sigma_v\}$. This generalizes to an n-dimensional manifold, where T_pM is an n-dimensional real vector space.

Defining a lie group. We can now put these two ideas together.

Definition 2.4. A *Lie group* is a smooth manifold G that is also a group, such that the group operations of multiplication $m(g, h) = gh$ and inversion $i(g) = g^{-1}$ are smooth maps.

Remark 2.5. The smoothness requirement is the crucial glue that holds the definition together. It's not enough for a set to be a group and a manifold separately.

Example 2.6. By far the most common examples are *matrix Lie groups*. These are closed subgroups of the *general linear group* $GL_n(\mathbb{R})$, the group of all invertible $n \times n$ real matrices. $GL_n(\mathbb{R})$ is itself a manifold, since it's an open subset of the vector space $\mathbb{R}^{n^2}$, the space of all $n \times n$ matrices, carved out by the condition $\det(A) \neq 0$. Many familiar groups are defined by the matrices that leave some bilinear form $\beta(v, w) = v^T Bw$ invariant. A group G of this type is given by $G = \{g \in GL_n(\mathbb{R}) \mid g^T Bg = B\}$. The condition on the Lie algebra is the infinitesimal version of this, as shown in the following proposition.

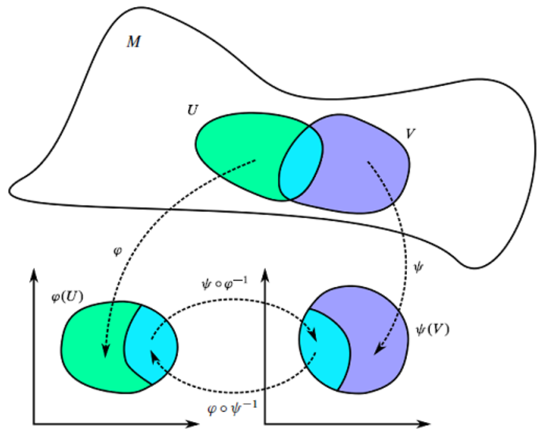

FIGURE 1. A smooth manifold M covered by charts mapping neighborhoods to Euclidean space. The transition maps between overlapping charts are required to be smooth.

Proposition 2.7. *The Lie algebra of* $G = \{g \in GL_n(\mathbb{R}) \mid g^T B g = B\}$ *is* $\mathfrak{g} = \{X \in M_n(\mathbb{R}) \mid X^T B + BX = 0\}$.

Proof. Let $X \in \mathfrak{g}$. By definition, the curve $\gamma(t) = \exp(tX)$ must lie in G for all $t \in \mathbb{R}$. This means $\exp(tX)^T B \exp(tX) = B$, and by differentiating at $t = 0$ we get 0. Using the product rule along with $\frac{d}{dt}\exp(tX)|_{t=0} = X$, we get

$$X^T \exp(0)^T B \exp(0) + \exp(0)^T BX \exp(0) = 0,$$

which simplifies to $X^T B + BX = 0$. For the other direction, assume $X^T B + BX = 0$. It is clear that $X^T B = -BX$, which implies $(X^T)^k B = (-1)^k BX^k$ for any integer $k \geq 0$. It follows that

$$\begin{aligned} \exp(tX)^T B = \exp(tX^T) B &= \left(\sum_{k=0}^{\infty} \frac{t^k (X^T)^k}{k!} \right) B \\ &= B \left(\sum_{k=0}^{\infty} \frac{t^k (-X)^k}{k!} \right) \\ &= B \exp(-tX). \end{aligned}$$

Thus, $\exp(tX)^T B \exp(tX) = B \exp(-tX)\exp(tX) = BI = B$. This shows $\exp(tX) \in G$ for all t, so $X \in \mathfrak{g}$, as desired. □

Here are some examples:

- If $B = I$, we get the *orthogonal group* $O(n)$, which leaves the dot product invariant. Its Lie algebra $\mathfrak{o}(n)$ is the space of skew-symmetric matrices, as $X^T + X = 0$. The subgroup $SO(n)$ of matrices with determinant 1 corresponds to rotations.
- If $B = \begin{pmatrix} 0 & I_k \\ -I_k & 0 \end{pmatrix}$, we get the *symplectic group* $Sp_{2k}(\mathbb{R})$.

The *special linear group* $SL_n(\mathbb{R})$ consists of matrices with determinant 1. Its Lie algebra $\mathfrak{sl}_n(\mathbb{R})$ is the space of traceless matrices.

Example 2.8. The simplest Lie groups are abelian. The real line $\mathbb{R}$ with addition, the circle group $S^1 = \{z \in \mathbb{C} \mid |z| = 1\}$ with complex multiplication, and the n-torus $T^n = S^1 \times \cdots \times S^1$ are all examples.

3. Linearizing a Lie Group

So, how can we possibly get a handle on such a complicated object as a Lie group? The central idea in all of Lie theory is *linearization*. We trade the non-linear group G for a related linear object, its Lie algebra $\mathfrak{g}$, which is much easier to work with.

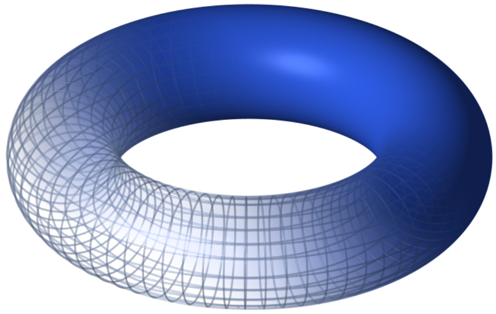

FIGURE 2. The 2-torus (T^2), a classic example of a compact, abelian Lie group.

The lie algebra as tangent space.

Definition 3.1. The *Lie algebra* $\mathfrak{g}$ of a Lie group G is the tangent space to G at the identity element e.

$$\mathfrak{g} := T_e G.$$

As a tangent space, $\mathfrak{g}$ is a vector space of the same dimension as G.To show the non-commutativity of the group, the algebra is equipped with a bilinear operation called the *Lie bracket* $[\cdot,\cdot] : \mathfrak{g} \times \mathfrak{g} \to \mathfrak{g}$. This bracket must be anti-symmetric, so $[X,Y] = -[Y,X]$, and satisfy the Jacobi identity, $[X,[Y,Z]] + [Y,[Z,X]] + [Z,[X,Y]] = 0$. For a matrix group $G \subseteq GL_n(\mathbb{R})$, the Lie algebra is the set of matrices

$$L(G) = \{X \in M_n(\mathbb{R}) \mid \exp(tX) \in G \text{ for all } t \in \mathbb{R}\}.$$

In this case, the Lie bracket is the familiar matrix commutator $[X,Y] = XY - YX$.

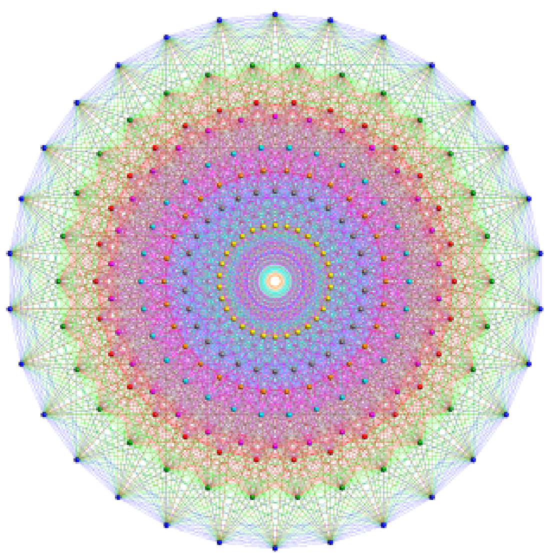

FIGURE 3. The Lie algebra $\mathfrak{g}$ is the tangent space to the Lie group G at the identity e. It is a flat, linear approximation of the group near that point.

The exponential map. To get from the algebra back to the group, we need a bridge: the exponential map. This map allows us to recover the local group structure entirely from the algebra.

Definition 3.2. The *exponential map*, $\exp : \mathfrak{g} \to G$, sends a vector $X \in \mathfrak{g}$ to the point on the group reached by following a path $\gamma_X(t)$ for one unit of time. The path γ_X is the unique one-parameter subgroup, which is a smooth homomorphism from $(\mathbb{R}, +)$ to G, whose starting velocity at $t = 0$ is X. We define $\exp(X) := \gamma_X(1)$.

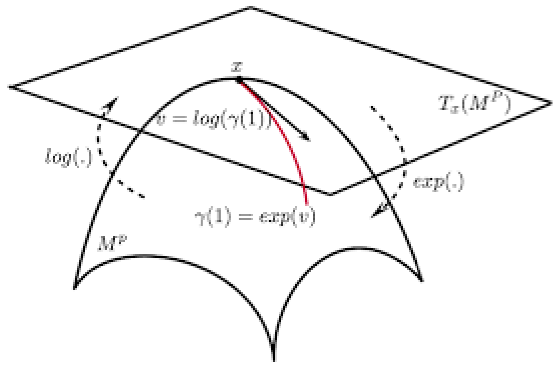

FIGURE 4. The exponential map takes a vector in the Lie algebra $\mathfrak{g}$ and maps it to a point on the group G by following the one-parameter subgroup it generates.

For matrix Lie groups, this is just the matrix exponential, $\exp(A) = \sum_{k=0}^{\infty} A^k/k!$. The map provides a local diffeomorphism from a neighborhood of $0 \in \mathfrak{g}$ to a neighborhood of $e \in G$. How the group's multiplication relates to the algebra's vector space structure is shown by the Baker-Campbell-Hausdorff formula.

> **The Baker-Campbell-Hausdorff Formula.** If X, Y are sufficiently small vectors in $\mathfrak{g}$, their product $\exp(X)\exp(Y)$ is equal to $\exp(Z)$ for some $Z \in \mathfrak{g}$. The formula gives Z as an infinite series of Lie brackets:
> $$Z = X + Y + \frac{1}{2}[X, Y] + \frac{1}{12}[X, [X, Y]] - \frac{1}{12}[Y, [X, Y]] + \dots .$$

The relationship between a Lie group homomorphism $\phi : G \to H$ and its corresponding Lie algebra homomorphism $d\phi_e : \mathfrak{g} \to \mathfrak{h}$ is combined with the exponential map. This is summarized by the following commutative diagram:

$$\begin{array}{ccc} \mathfrak{g} & \xrightarrow{\exp_G} & G \\ {\scriptstyle d\phi_e}\downarrow & & \downarrow{\scriptstyle \phi} \\ \mathfrak{h} & \xrightarrow[\exp_H]{} & H \end{array}$$

Remark 3.3. In practice, the full Baker-Campbell-Hausdorff formula is famously messy and almost never used for explicit calculations. The way that its actually applicable is because it guarantees that the group's multiplication is completely determined by the Lie bracket, at least near the identity. The first few terms are often sufficient for local approximations.

The adjoint representation. The Lie bracket isn't an ad-hoc definition; it comes naturally from the group structure itself. The group G acts on itself by conjugation: $c_g(h) = ghg^{-1}$. Taking the differential of this map at the identity gives a linear map on the Lie algebra, called the *adjoint representation* of the group.

$$\mathrm{Ad}(g) = (dc_g)_e : \mathfrak{g} \to \mathfrak{g}.$$

This gives a homomorphism $\mathrm{Ad} : G \to GL(\mathfrak{g})$, turning each group element g into an invertible linear map on the algebra. Differentiating *this* map at the identity gives a map from the Lie algebra of G to the Lie algebra of $GL(\mathfrak{g})$:

$$\mathrm{ad} := L(\mathrm{Ad}) : \mathfrak{g} \to \mathfrak{gl}(\mathfrak{g}) = \mathrm{End}(\mathfrak{g}).$$

As it turns out, this map 'ad' is the Lie bracket: $\mathrm{ad}(X)(Y) = [X, Y]$. The two are tied together by a beautiful formula:

$$\mathrm{Ad}(\exp(X)) = e^{\mathrm{ad}(X)} = \sum_{k=0}^{\infty} \frac{(\mathrm{ad}X)^k}{k!},$$

which cleanly connects group conjugation with the algebraic bracket structure [3]. This fundamental identity can be seen below:

$$\begin{array}{ccc} G & \xrightarrow{\mathrm{Ad}} & GL(\mathfrak{g}) \\ {\scriptstyle \exp}\uparrow & & \uparrow{\scriptstyle \exp} \\ \mathfrak{g} & \xrightarrow[\mathrm{ad}]{} & \mathfrak{gl}(\mathfrak{g}) \end{array}$$

4. Why Does It Take a 720° Turn to Get Back to Where You Started?

To make all this abstract theory concrete, there is no better example than the relationship between the special unitary group $SU(2)$ and the rotation group $SO(3)$. This connection is notably the basis for the quantum mechanical theory of electron spin.

Proposition 4.1. *The Lie algebras $\mathfrak{so}(3)$ and $\mathfrak{su}(2)$ are isomorphic.*

Proof. Note that $\mathfrak{so}(3)$ is the three-dimensional real vector space of 3×3 real skew-symmetric matrices, with a standard basis given by $E_1 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix}, E_2 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix}, E_3 = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$. Similarly, $\mathfrak{su}(2)$ has a basis $F_1 = \frac{1}{2}\begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}, F_2 = \frac{1}{2}\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, F_3 = \frac{1}{2}\begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}$. It is clear that for both sets of bases, $[X_i, X_j] = \epsilon_{ijk} X_k$ holds, where X is either E or F, thus the linear map sending F_i to E_i for each $i \in [\![1, 3]\!]$ is a Lie algebra isomorphism, as desired. □

Since their Lie algebras are the same, the groups $SU(2)$ and $SO(3)$ must look identical locally. Globally, however, they are different beasts.

Theorem 4.2. *There is a 2-to-1 surjective homomorphism $\phi : SU(2) \to SO(3)$. Topologically, $SU(2)$ is the simply connected universal cover of $SO(3)$.*

Proof. We define a map using the adjoint representation $\mathrm{Ad} : SU(2) \to \mathrm{Aut}(\mathfrak{su}(2))$. This action leaves the inner product on $\mathfrak{su}(2)$ given by $\langle X, Y \rangle := -2\mathrm{Tr}(XY)$ invariant, which guarantees that the image of Ad is a subgroup of $O(\mathfrak{su}(2)) \cong O(3)$. Because $SU(2)$ is connected, its image must lie in $SO(3)$, so we obtain a homomorphism $\phi : SU(2) \to SO(3)$, which is surjective. It follows that some $g \in SU(2)$ is in the kernel of ϕ if and only if it commutes with every element in $\mathfrak{su}(2)$, which implies that g lies in the center of $SU(2)$. The center of $SU(2)$ is $\{\pm I\}$, and thus, ϕ is a 2-to-1 surjective homomorphism. We can see this in the relationship below:

$$1 \longrightarrow \{\pm I\} \xrightarrow{i} SU(2) \xrightarrow{\phi} SO(3) \longrightarrow 1$$

Since $SU(2)$ is topologically the simply connected 3-sphere S^3, it is the universal covering group of $SO(3)$, and it follows that $SO(3) \cong SU(2)/\{\pm I\}$, as desired. □

Remark 4.3. This 2-to-1 covering has direct physical consequences. Objects that transform under $SO(3)$, such as vectors, come back to where they started after a 360° rotation. Objects that transform under $SU(2)$, such as spinors like electrons, pick up a minus sign after a 360° rotation and need a full 720° rotation to return to their original state. This strange "double covering" is one of the useful properties of the quantum mechanical description of spin.

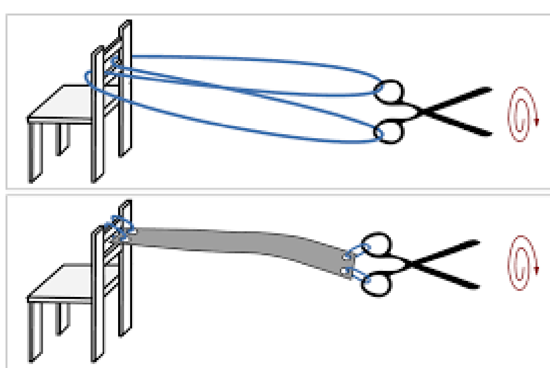

Figure 5. The "belt trick" is a nice physical analogy for the $SU(2)$ double cover of $SO(3)$.

5. Structure and Representation Theory

Lie theory (obviously) doesn't just end with the exponential map. The correspondence between Lie groups and Lie algebras is the starting point for a structure theory that lets us classify these objects and understand how they can act on other spaces.

Structure of lie algebras. We can analyze the structure of a Lie algebra $\mathfrak{g}$ by looking for its ideals, which are subspaces stable under the Lie bracket. One way to do this is with the derived series, defined by $D^0(\mathfrak{g}) = \mathfrak{g}$ and $D^{k+1}(\mathfrak{g}) = [D^k(\mathfrak{g}), D^k(\mathfrak{g})]$.

Definition 5.1. A Lie algebra $\mathfrak{g}$ is *solvable* if its derived series eventually hits $\{0\}$. A Lie algebra is *semisimple* if it has no non-zero solvable ideals. A non-abelian Lie algebra is *simple* if its only ideals are $\{0\}$ and itself.

Theorem 5.2 (Levi's Theorem [3])**.** *Any finite-dimensional real Lie algebra $\mathfrak{g}$ can be broken down as a semidirect product of its largest solvable ideal, called the radical and denoted $rad(\mathfrak{g})$, and a semisimple subalgebra $\mathfrak{s}$.*

$$\mathfrak{g} = rad(\mathfrak{g}) \rtimes \mathfrak{s}$$

This theorem important because it reduces the general study of Lie algebras to two more manageable cases: solvable and semisimple. One of the main tools for telling them apart is the Killing form.

> **The Killing Form and Cartan's Criteria**. The *Killing form* is a symmetric bilinear form on $\mathfrak{g}$ built from the adjoint representation:
> $$\kappa(X, Y) := \mathrm{tr}(\mathrm{ad}(X) \circ \mathrm{ad}(Y)).$$
> Cartan's criteria use the Killing form to sort algebras. A Lie algebra $\mathfrak{g}$ is solvable if and only if $\kappa(X, Y) = 0$ for all $X \in \mathfrak{g}, Y \in [\mathfrak{g}, \mathfrak{g}]$. A Lie algebra $\mathfrak{g}$ is semisimple if and only if its Killing form is non-degenerate. These are powerful computational tests. For example, one can compute the Killing form for $\mathfrak{sl}_n(\mathbb{R})$ to be $\kappa(X, Y) = 2n\mathrm{Tr}(XY)$, which is non-degenerate, proving it is semisimple. Semisimple Lie algebras can be broken down further into a direct sum of simple ones, which have been completely classified into the classical families of types A_n, B_n, C_n, D_n and five exceptional cases, E_6, E_7, E_8, F_4, G_2.

Proposition 5.3. *The Killing form is associative, meaning $\kappa([X, Y], Z) = \kappa(X, [Y, Z])$ for all $X, Y, Z \in \mathfrak{g}$.*

Proof. It suffices to expand the definition. Note that ad is a Lie algebra homomorphism, so $\mathrm{ad}([X, Y]) = [\mathrm{ad}(X), \mathrm{ad}(Y)]$. We have

$$\begin{aligned}
\kappa([X, Y], Z) &= \mathrm{tr}(\mathrm{ad}([X, Y]) \circ \mathrm{ad}(Z)) \\
&= \mathrm{tr}((\mathrm{ad}(X)\mathrm{ad}(Y) - \mathrm{ad}(Y)\mathrm{ad}(X)) \circ \mathrm{ad}(Z)) \\
&= \mathrm{tr}(\mathrm{ad}(X)\mathrm{ad}(Y)\mathrm{ad}(Z)) - \mathrm{tr}(\mathrm{ad}(Y)\mathrm{ad}(X)\mathrm{ad}(Z)).
\end{aligned}$$

Using $\mathrm{tr}(ABC) = \mathrm{tr}(BCA)$, on the second term, we see

$$\begin{aligned}
\kappa([X, Y], Z) &= \mathrm{tr}(\mathrm{ad}(X)\mathrm{ad}(Y)\mathrm{ad}(Z)) - \mathrm{tr}(\mathrm{ad}(X)\mathrm{ad}(Z)\mathrm{ad}(Y)) \\
&= \mathrm{tr}(\mathrm{ad}(X) \circ (\mathrm{ad}(Y)\mathrm{ad}(Z) - \mathrm{ad}(Z)\mathrm{ad}(Y))) \\
&= \mathrm{tr}(\mathrm{ad}(X) \circ [\mathrm{ad}(Y), \mathrm{ad}(Z)]) \\
&= \mathrm{tr}(\mathrm{ad}(X) \circ \mathrm{ad}([Y, Z])) = \kappa(X, [Y, Z]),
\end{aligned}$$

as desired. This property is also called the invariance of the Killing form. □

Representation theory. One of the main reasons to study Lie theory is for its applications in representation theory, which is the study of how these groups and algebras act on vector spaces. A representation of a Lie algebra $\mathfrak{g}$ on a vector space V is just a Lie algebra homomorphism $\rho : \mathfrak{g} \to \mathfrak{gl}(V)$, where $\mathfrak{gl}(V)$ is the algebra of endomorphisms of V with the commutator bracket.

Theorem 5.4 (Lie's Theorem [3])**.** *Let $\mathfrak{g}$ be a solvable Lie algebra over an algebraically closed field of characteristic zero. Then for any finite-dimensional representation $\rho : \mathfrak{g} \to \mathfrak{gl}(V)$, there is a non-zero vector $v \in V$ that is a simultaneous eigenvector for all operators $\rho(X)$ for $X \in \mathfrak{g}$.*

This theorem implies that any representation of a solvable Lie algebra can be put into an upper-triangular basis. For semisimple algebras, the situation is much nicer.

Theorem 5.5 (Weyl's Theorem on Complete Reducibility [3]). *Every finite-dimensional representation of a semisimple Lie algebra is completely reducible, meaning it is a direct sum of irreducible subrepresentations.*

Remark 5.6. Weyl's theorem is a really special property of semisimple algebras. For other Lie algebras, representations can be *indecomposable* without being irreducible. For instance, the representation of $\mathbb{R}$ by matrices $\begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix}$ fixes a 1D subspace but does not split into a direct sum of 1D representations.

This reduces the whole problem to classifying the irreducible representations, or "irreps." In physics, these irreps often correspond to fundamental particles. Combining systems, which corresponds to the tensor product of representations, means decomposing the resulting representation into its irreducible parts. For $GL(V)$, for instance, the irreps are the Schur functors $S_\lambda V$, indexed by partitions λ. Decomposing their tensor products is governed by combinatorial rules like the Littlewood-Richardson rule, which gives integer coefficients $c^\nu_{\lambda\mu}$ in the expansion [1]:

$$S_\lambda V \otimes S_\mu V \cong \bigoplus_\nu (S_\nu V)^{\oplus c^\nu_{\lambda\mu}}$$

For the standard representation V of $GL(V)$, where $\lambda = (1)$, the tensor product $V \otimes V$ splits into symmetric and anti-symmetric parts:

$$V \otimes V \cong \mathrm{Sym}^2(V) \oplus \Lambda^2(V),$$

corresponding to

$$S_{(1)}V \otimes S_{(1)}V \cong S_{(2)}V \oplus S_{(1,1)}V.$$

This kind of decomposition is a routine calculation in physics and math, from combining angular momenta in quantum mechanics to dealing with tensors in general relativity.

6. The 'Periodic Table' of Simple Lie Algebras

Weyl's Theorem tells us that representations of semisimple Lie algebras break down into irreducible building blocks. This shifts everything to classifying these building blocks: the simple Lie algebras themselves and their irreps. The complete classification of simple Lie algebras over $\mathbb{C}$ is one of the greatest achievements of modern mathematics, thanks mainly to Wilhelm Killing and Élie Cartan.

Root space decomposition. The trick to dissecting a semisimple Lie algebra $\mathfrak{g}$ is to find a maximal abelian subalgebra whose elements can all be diagonalized at the same time under the adjoint representation.

Definition 6.1. A *Cartan subalgebra* $\mathfrak{h}$ of $\mathfrak{g}$ is a nilpotent subalgebra that is its own normalizer. For semisimple Lie algebras, this is thankfully equivalent to being a maximal abelian subalgebra.

Since everything in $\mathfrak{h}$ commutes, the linear maps $\mathrm{ad}(H)$ for $H \in \mathfrak{h}$ form a commuting family of operators on $\mathfrak{g}$. Over $\mathbb{C}$, this means we can simultaneously diagonalize them, which leads to the *root space decomposition*.

Theorem 6.2 (Root Space Decomposition [3]). *Let $\mathfrak{g}$ be a semisimple Lie algebra over $\mathbb{C}$ with Cartan subalgebra $\mathfrak{h}$. Then $\mathfrak{g}$ splits into a direct sum of vector spaces:*

$$\mathfrak{g} = \mathfrak{h} \oplus \bigoplus_{\alpha \in \Phi} \mathfrak{g}_\alpha$$

where $\Phi \subset \mathfrak{h}^* \setminus \{0\}$ is a finite set of linear functionals called roots*, and $\mathfrak{g}_\alpha$ is the eigenspace*

$$\mathfrak{g}_\alpha = \{X \in \mathfrak{g} \mid [H, X] = \alpha(H)X \text{ for all } H \in \mathfrak{h}\}.$$

Miraculously, each root space $\mathfrak{g}_\alpha$ turns out to be one-dimensional.

The set of roots Φ forms a beautiful, highly symmetric geometric object in the dual space $\mathfrak{h}^*$, known as a *root system*. This object is the combinatorial skeleton of the Lie algebra.

Root systems and dynkin diagrams. The geometry of the root system tells you everything you need to know about the algebra. The Killing form gives us an inner product on $\mathfrak{h}^*$, turning it into a Euclidean space.

Definition 6.3. A *root system* is a finite set of non-zero vectors Φ in a Euclidean space V that satisfies:

(1) Φ spans V.
(2) If $\alpha \in \Phi$, the only multiples of α in Φ are $\pm\alpha$.
(3) For any $\alpha \in \Phi$, the reflection s_α across the hyperplane orthogonal to α maps the set Φ to itself.
(4) For any $\alpha, \beta \in \Phi$, the number $n_{\beta\alpha} = 2\frac{\langle \beta,\alpha \rangle}{\langle \alpha,\alpha \rangle}$ must be an integer.

Classifying simple Lie algebras now boils down to classifying irreducible root systems. These, in turn, can be encoded in simple graphs called *Dynkin diagrams.* To make one, you pick a basis of simple roots $\Delta \subset \Phi$. The diagram gets a node for each simple root. The number of edges between two nodes, say for α_i and α_j, is $n_{ij}n_{ji}$. If the roots have different lengths, an arrow points from the longer root to the shorter one.

> **The Classification Theorem**. Every simple complex Lie algebra corresponds to one of the following Dynkin diagrams. There are four infinite families corresponding to the classical algebras: A_n for $\mathfrak{sl}_{n+1}$, a simple chain of n nodes; B_n for $\mathfrak{so}_{2n+1}$, a chain with a double edge at the end; C_n for $\mathfrak{sp}_{2n}$, similar to B_n but with the arrow reversed; and D_n for $\mathfrak{so}_{2n}$, a chain with a fork at the end. In addition, there are exactly five exceptional algebras, denoted E_6, E_7, E_8, F_4, G_2. This classification is exhaustive; it is a complete "periodic table" for simple Lie algebras.

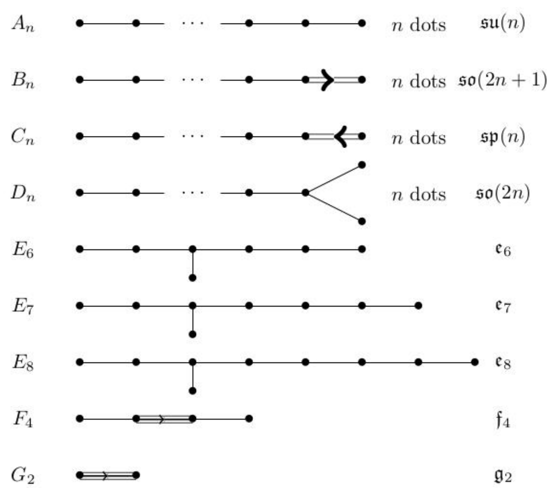

FIGURE 6. The complete classification of simple complex Lie algebras via their Dynkin diagrams, showing the four infinite classical families and five exceptional cases.

7. Classifying Irreducible Representations

Now that we understand the structure of semisimple algebras via root systems, we can describe their representations in a similar combinatorial way. The main idea is again to diagonalize the action of the Cartan subalgebra $\mathfrak{h}$ on the representation space V.

Weights and weight spaces. Just like the algebra itself, any finite-dimensional representation V splits into eigenspaces for the action of $\mathfrak{h}$.

Definition 7.1. Let (ρ, V) be a representation of a semisimple Lie algebra $\mathfrak{g}$. A functional $\lambda \in \mathfrak{h}^*$ is a *weight* of the representation if its *weight space*

$$V_\lambda = \{v \in V \mid \rho(H)v = \lambda(H)v \text{ for all } H \in \mathfrak{h}\}$$

is non-zero. The dimension of V_λ is called the *multiplicity* of the weight λ. The whole representation splits as a direct sum of its weight spaces: $V = \bigoplus_\lambda V_\lambda$.

The root vectors $X_\alpha \in \mathfrak{g}_\alpha$ act as ladder operators, moving vectors from one weight space to another: if $v \in V_\lambda$, then $X_\alpha \cdot v \in V_{\lambda+\alpha}$. This motivates the set of weights for any representation to form highly structured patterns, shifted copies of the root system.

The theorem of the highest weight. For an irreducible representation, this structure is extremely rigid. Once we choose a set of positive roots, there will always be a unique "highest" weight.

Theorem 7.2 (Theorem of the Highest Weight [1])**.** *For every finite-dimensional irreducible representation V of a semisimple Lie algebra $\mathfrak{g}$, there is a unique* highest weight λ*, such that all other weights are of the form $\lambda - \sum k_i\alpha_i$ for non-negative integers k_i and simple roots α_i. Furthermore: the highest weight space V_λ is one-dimensional; two irreducible representations are isomorphic if and only if they have the same highest weight; and for any "dominant integral weight" λ, which is a specific combinatorial condition, there exists a unique irreducible representation with that highest weight.*

This amazing theorem creates a nice bijection between irreducible representations and a set of simple combinatorial objects called dominant integral weights. It reduces the hard problem of classifying representations to a much easier combinatorial one.

Example 7.3. The Lie algebra $\mathfrak{sl}_3(\mathbb{C})$, of type A_2, has a 2D Cartan subalgebra, so its weights live on a plane. The standard representation $V \cong \mathbb{C}^3$ has three weights, forming a triangle. Its dual V^* has three weights forming an inverted triangle. The adjoint representation, which is the action of $\mathfrak{sl}_3$ on itself, has weights given by the six roots plus a zero weight of multiplicity two, forming a hexagon. If you tensor the standard representation with its dual, $V \otimes V^*$, the resulting 9D representation decomposes into the 8D adjoint representation and a 1D trivial one. This is exactly the decomposition of mesons into quarks and antiquarks in the "Eightfold Way" model of particle physics!

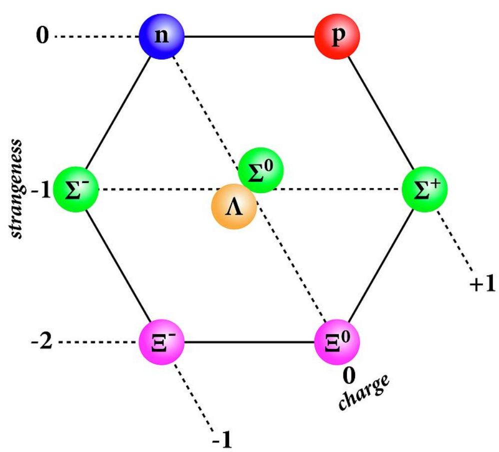

Figure 7. The baryon octet from the "Eightfold Way" particle physics model. This is a weight diagram for the 8-dimensional adjoint representation of the $SU(3)$ flavor symmetry group.

8. Homogeneous and Symmetric Spaces

The connection between algebra and geometry isn't just about the group manifold itself. Lie groups often act on other geometric spaces in a highly structured way.

Homogeneous spaces. A smooth manifold M is called a *homogeneous space* if a Lie group G acts on it transitively. This just means you can get from any point p to any other point q by acting with some group element $g \in G$. Roughly speaking, the space "looks the same" from every point.

Definition 8.1. If a Lie group G acts transitively on a manifold M, and H is the *stabilizer subgroup* of some point $p \in M$, meaning $H = \{g \in G \mid g \cdot p = p\}$, then H is a closed Lie subgroup of G, and the manifold M is diffeomorphic to the quotient space of left cosets G/H.

This gives the manifold M the structure of a fiber bundle over the base space G/H with fiber H:

$$H \xrightarrow{i} G \xrightarrow{\pi} G/H$$

Remark 8.2. The notation G/H formalizes the idea that a geometric object is defined by its symmetries. The sphere S^2 is "round" because the rotation group $SO(3)$ acts transitively on it; every point is the same as every other.

Example 8.3. Many of the most important spaces in geometry are homogeneous spaces. For instance, Euclidean space $\mathbb{R}^n$ is the quotient $E(n)/O(n)$, where $E(n)$ is the group of isometries of $\mathbb{R}^n$. The spaces of constant positive and negative curvature are also canonical examples.

Proposition 8.4. *The n-sphere S^n, the canonical space of constant positive curvature, is diffeomorphic to the homogeneous space $SO(n+1)/SO(n)$.*

Proof. Let the Lie group $G = SO(n+1)$ act on the manifold $M = S^n \subset \mathbb{R}^{n+1}$ by the standard matrix-vector multiplication. It is clear that this action is smooth. To show it is transitive, let $p, q \in S^n$ be any two points. Since both are unit vectors, there exists a rotation in $SO(n+1)$ that maps p to q. Thus, G acts transitively on M.

Note that S^n is diffeomorphic to the quotient G/H, where H is the stabilizer of a chosen point. Let us choose the north pole $p = e_1 = (1, 0, \dots, 0)^T$. The stabilizer is $H = \{g \in SO(n+1) \mid ge_1 = e_1\}$. If $ge_1 = e_1$, the first column of g must be e_1. Since $g \in O(n+1)$, its columns must be orthonormal, thus the first row of g to be $(1, 0, \dots, 0)$. Thus g must be of the block form

$$g = \begin{pmatrix} 1 & 0 \\ 0 & A \end{pmatrix},$$

where A is an $n \times n$ matrix. Also, from $g \in SO(n+1)$ it follows $g^T g = I$ and $\det(g) = 1$. The orthogonality condition implies $A^T A = I$, so $A \in O(n)$. The determinant condition implies $\det(g) = 1 \cdot \det(A) = 1$, so $A \in SO(n)$. It follows that the stabilizer H is isomorphic to $SO(n)$, and therefore $S^n \cong SO(n+1)/SO(n)$, as desired. □

Remark 8.5. Similarly, the hyperbolic space $\mathbb{H}^n$, the space of constant negative curvature, can be represented as the quotient $SO^+(n,1)/SO(n)$, where $SO^+(n,1)$ is the component of the identity of the Lorentz group leaving a quadratic form of signature $(n,1)$ invariant.

Symmetric spaces. An especially important class of homogeneous spaces are the symmetric spaces, introduced by Élie Cartan. A Riemannian manifold (M, g) is a *symmetric space* if, for every point $p \in M$, the geodesic symmetry s_p, which flips vectors in the tangent space by sending $\exp_p(v)$ to $\exp_p(-v)$, is a global isometry of the whole manifold.

Theorem 8.6 (Cartan)**.** *Every simply connected, complete Riemannian symmetric space is one of three types: Euclidean, which is a flat space like $\mathbb{R}^n$; compact, having positive curvature like S^n; or non-compact, having negative curvature like $\mathbb{H}^n$. Furthermore, every such space can be written as a homogeneous space G/H for a semisimple Lie group G.*

Acknowledgments

I want to thank Lucy Vuong for advice throughout the process of making this paper. I also want to thank Simon Rubenstein-Salzedo for running Euler Circle, which provided the opportunity for me to write this paper.

References

[1] W. Fulton and J. Harris, *Representation Theory: A First Course*, Graduate Texts in Mathematics, vol. 129, Springer-Verlag, New York, 1991.
[2] T. Hawkins, *The Birth of Lie's Theory of Groups*, The Mathematical Intelligencer, vol. 16, no. 2, 1994, pp. 6–17.
[3] J. Hilgert and K.-H. Neeb, *Structure and Geometry of Lie Groups*, Springer Monographs in Mathematics, Springer, 2012.
[4] S. Rubinstein-Salzedo, *Differential Geometry*, Euler Circle, 2021.

Euler Circle, Palo Alto, CA 94306