The Relationship Between $\mathfrak{su}(2)$, $\mathfrak{so}(3)$ and SU(2), SO(3)

This note explains why the Lie algebras $\mathfrak{su}(2)$ and $\mathfrak{so}(3)$ have the same representation theory, while the corresponding Lie groups SU(2) and SO(3) do not. The central point is that (for topological reasons) exponentiating representations of the Lie algebra always yields representations of SU(2), but not always of SO(3).

1 Lie Algebra Equivalence: $\mathfrak{su}(2) \cong \mathfrak{so}(3)$

Although SU(2) and SO(3) are different groups, their Lie algebras are isomorphic:

$$\mathfrak{su}(2) \simeq \mathfrak{so}(3).$$

Both are real three-dimensional Lie algebras whose generators satisfy the same commutation relations:

$$[J_i, J_j] = \varepsilon_{ijk} J_k.$$

Consequently, any representation of $\mathfrak{su}(2)$ determines a representation of $\mathfrak{so}(3)$ and vice versa. Their irreducible representations are labelled by a non-negative half-integer j ("spin"), with dimension $2j+1$:

$$j = 0, \tfrac{1}{2}, 1, \tfrac{3}{2}, \ldots$$

2 Exponentiating Representations of the Lie Algebra: Examples

Let's examine what happens when we exponentiate the 1-, 2-, and 3-dimensional irreps of $\mathfrak{su}(2)$ and $\mathfrak{so}(3)$.

1D irrep ($j = 0$). In the trivial representation we have $J_x = J_y = J_z = 0$. Thus for any angle θ,

$$e^{i\theta J_z} = e^0 = 1.$$

In particular,

$$e^{i2\pi J_z} = 1,$$

so this clearly defines a representation of both SU(2) and SO(3) - the trivial rep.

2D irrep ($j = \frac{1}{2}$). We begin by defining the generators

$$J_x = \frac{1}{2}\sigma_x, \qquad J_y = \frac{1}{2}\sigma_y, \qquad J_z = \frac{1}{2}\sigma_z, \tag{1}$$

where the σ_i are the Pauli matrices. These satisfy

$$[\sigma_i, \sigma_j] = 2i\,\varepsilon_{ijk}\sigma_k,$$

and therefore the matrices in Eq. (1) obey

$$[J_i, J_j] = i\,\varepsilon_{ijk} J_k,$$

which are precisely the commutation relations of the Lie algebra $\mathfrak{su}(2)$. Thus they form a valid (and irreducible) representation of the algebra.

To understand what happens when we exponentiate these generators, it is helpful to recall the Bloch-sphere picture of a qubit. A pure state $|\psi\rangle$ corresponds to a point on the Bloch sphere via its Bloch vector $\boldsymbol{r}$, and a physical rotation of the Bloch sphere through angle θ about an axis $\hat{n}$ is implemented on the state by

$$U(\theta, \hat{n}) = \exp\left(-i\,\frac{\theta}{2}\,\hat{n}\cdot\boldsymbol{\sigma}\right).$$

The factor of 1/2 in the exponent means that the matrices $\frac{1}{2}\sigma_i$ generate all possible rotations of the Bloch sphere. In other words, by exponentiating the operators in Eq. (1) we recover exactly the unitary transformations that act on a qubit and rotate its Bloch vector. This is why exponentiating the 2D irrep "clearly" produces the standard fundamental representation of SU(2): it gives precisely the transformations that rotate qubit states in the way dictated by the Bloch-sphere geometry.

Now consider what happens under a full physical rotation by 2π. On the Bloch sphere,

$$R(2\pi) = I,$$

so every point on the sphere returns to itself. However,

$$U(2\pi) = \exp(-i\pi\,\hat{n}\cdot\boldsymbol{\sigma}) = -\,I.$$

Although $U(2\pi)$ multiplies the state vector by -1, the corresponding Bloch vector is unchanged, because the density matrix $|\psi\rangle\langle\psi|$ is invariant under the replacement $|\psi\rangle \mapsto -|\psi\rangle$.

A representation of SO(3) must send a full 2π rotation to $+I$, but exponentiating Eq. (1) sends it to $-I$. Hence this construction does *not* give a representation of SO(3), even though it is a perfectly good (and physically very important) representation of SU(2).

3D irrep ($j = 1$). In the spin–1 representation the generators are the 3×3 angular momentum matrices familiar from ordinary vector rotations in three dimensions. These matrices satisfy the same $\mathfrak{su}(2)$ commutation relations as in the spin $\frac{1}{2}$ case, but now they act on a three-dimensional real vector space, i.e., the space in which physical rotation matrices naturally live. For example, the generator of rotations about the z–axis is

$$J_z = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix},$$

which reflects the fact that the z–component is left unchanged while the x– and y–components mix under rotations in the usual way.

Exponentiating these generators reproduces the standard rotation matrices of three-dimensional Euclidean space. For instance, a rotation by angle θ about the z–axis is

$$e^{i\theta J_z} = \begin{pmatrix} e^{i\theta} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & e^{-i\theta} \end{pmatrix}.$$

If we restrict attention to real vectors, this is simply the familiar real rotation matrix acting on the (x, y)–plane with the z–axis fixed. Thus exponentiating the spin–1 generators yields exactly the rotations we use to move vectors around in physical three-dimensional space.

Evaluating this expression at $\theta = 2\pi$ gives

$$e^{i2\pi J_z} = \begin{pmatrix} e^{i2\pi} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & e^{-i2\pi} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} = I.$$

Thus a full 2π rotation returns every vector to itself, just as expected for an ordinary rotation in $\mathbb{R}^3$. In contrast to the spin–1/2 case, the matrices in this representation respond to a 2π rotation in exactly the way required of an element of SO(3).

Hence the spin–1 representation *does* give a representation of SO(3), since exponentiating its generators yields precisely the usual rotation matrices acting on vectors in three-dimensional space.

3 When is a representation of SU(2) also a representation of SO(3)?

Every representation of $\mathfrak{su}(2)$ gives, upon exponentiation, a representation of SU(2). But such a representation need not give a well-defined representation of SO(3). To understand why, we note (as seen in the examples above) that SO(3) is obtained from SU(2) by identifying pairs of elements that correspond to the same physical rotation. We now spell this out carefully.

Which elements of SU(2) represent the same rotation? In the examples above we saw that for the 2D irrep of $\mathfrak{su}(2)$ we had

$$e^{i2\pi J_z} = -I$$

whereas in the 1D and 3D representations,

$$e^{i2\pi J_z} = I\,.$$

This showed that in SU(2) a 2π rotation under the generator J_z is *not* necessarily the identity: it can be a distinct group element, $-I$.

However, from the point of view of physics, a 2π rotation brings every vector in $\mathbb{R}^3$ back to its starting orientation. Thus, in SO(3), a 2π rotation *is* the identity element.

This means that the two distinct matrices

$$I \qquad \text{and} \qquad -I$$

in SU(2) both correspond to the *same* physical rotation when we pass to SO(3). A fundamental fact from Lie group theory is that these are the *only* distinct elements of SU(2) that this happens to. In other words: apart from I and $-I$ (and so correspondingly U and $-U$), no other pairs of elements of SU(2) represent the same rotation in SO(3).

What does this mean for representations? Let R be a representation of SU(2). For R to define a representation of SO(3), it must assign the *same operator* to any two elements of SU(2) that correspond to the same rotation. Since I and $-I$ are the only such pair, the requirement is:

$$R(-I) = R(I) = I.$$

In other words, R must act *trivially* on $-I$.

What happens in the spin-j representation? In the spin-j representation, the J_z operator is diagonal with eigenvalues

$$j,\ j-1,\ \ldots,\ -j.$$

Thus

$$R(2\pi) = e^{i2\pi J_z} = \mathrm{diag}\Big(e^{i2\pi j},\ e^{i2\pi(j-1)},\ \ldots,\ e^{-i2\pi j}\Big)\,.$$

If j is an integer, then all exponents are integers and each factor is $e^{i2\pi m} = 1$. Hence $R(2\pi) = +I$. If j is a half-integer, then all exponents are half-integers and each factor is $e^{i2\pi(m+1/2)} = -1$. Hence $R(2\pi) = -I$. In both cases the result is a scalar matrix, which can be written compactly as

$$R(-I) = R(2\pi) = (-1)^{2j}\, I.$$

Thus,

$$R(-I) = \begin{cases} +I, & j \in \mathbb{Z}, \\ -I, & j \in \mathbb{Z} + \frac{1}{2}. \end{cases}$$

Connecting back to the earlier examples.

- In the **1D irrep** ($j = 0$), we found $R(2\pi) = +I$. This satisfies $R(-I) = I$, so this does give a representation of SO(3).
- In the **2D irrep** ($j = \frac{1}{2}$), we found $R(2\pi) = -I$. This violates the condition $R(-I) = I$, so the spin–$\frac{1}{2}$ representation does *not* give a representation of SO(3).
- In the **3D irrep** ($j = 1$), we found $R(2\pi) = +I$ again, and therefore this representation *does* give a representation of SO(3).

Conclusion. Only for integer values of j does the representation send $-I$ to the identity operator. These and only these representations of SU(2) give well-defined representations of SO(3).

SU(2) as a Double Cover of SO(3) The relationship between SU(2) and SO(3) is expressed by a map

$$\pi : \mathrm{SU}(2) \longrightarrow \mathrm{SO}(3).$$

This map is called a *covering map*. In this context, "covering" means the following concrete fact:

Every rotation in SO(3) comes from exactly two different matrices in SU(2).

More precisely, the matrices U and $-U$ in SU(2) always produce the same rotation when interpreted as acting on the Bloch sphere (or, equivalently, on vectors in $\mathbb{R}^3$). The covering map π therefore satisfies

$$\pi(U) = \pi(-U) \qquad \text{for every } U \in \mathrm{SU}(2).$$

These are the *only* pairs of distinct elements in SU(2) that map to the same rotation. In particular,

$$\pi(I) = \pi(-I),$$

so the identity element of SO(3) has two "preimages" in SU(2): the matrices I and $-I$.

Now consider a representation R of SU(2). For R to give a consistent representation of SO(3) —meaning that it depends only on the resulting physical rotation—it must assign the *same* operator to any two elements of SU(2) that represent the same rotation. Since the only such pair is U and $-U$, this requirement becomes

$$R(U) = R(-U) \qquad \text{for all } U \in \mathrm{SU}(2).$$

In particular, this requires

$$R(-I) = R(I) = I.$$

Using our earlier computation

$$R(-I) = (-1)^{2j}\, I,$$

we see that this happens precisely when j is an integer. Therefore, *a representation of SU(2) gives a well-defined representation of SO(3) exactly when the spin is an integer.*

4 Why does exponentiating representations of $\mathfrak{su}(2)$ always give a representation of SU(2) but exponentiating representations of $\mathfrak{so}(3)$ doesn't always give a representation of SO(3)?

The key idea is that SU(2) and SO(3) have different topology, and this difference affects which Lie algebra representations exponentiate to group representations. To explain this, we first explain what it means for a space to be *simply connected*.

4.1 Simply Connected Spaces

Formal definition. A topological space X is *simply connected* if

1. it is path–connected, and
2. every closed loop in X can be continuously deformed (contracted) to a single point while remaining inside X.

Intuition. A space is simply connected if it has no "holes" that a loop can get stuck around. On a sphere, for example, every loop can be pulled tight and made to shrink to a point. On a torus (a donut), however, a loop that winds around the hole cannot be contracted without leaving the torus. In this sense a torus is not simply connected, while a sphere is.

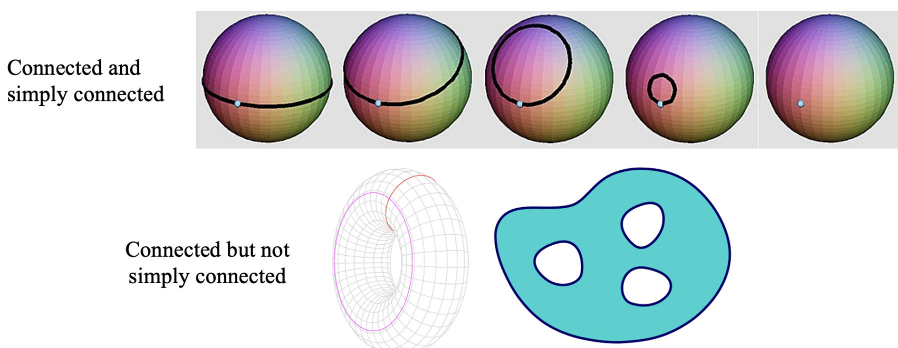

Figure 1: Some pictures from the wiki page on simply connected spaces.

Claim. A foundational theorem in Lie theory states:

> If a Lie group G is simply connected, then every representation of its Lie algebra $\mathfrak{g}$ exponentiates to a representation of G.

Thus, whether a Lie algebra representation integrates to the group depends not only on the algebra itself but also on the topology of G. Now SU(2) is simply connected but SO(3) is not. This is why every representation of the Lie algebra $\mathfrak{su}(2)$ exponentiates to a representation of SU(2), whereas representations of $\mathfrak{so}(3)$ do not necessarily exponentiate to representations of SO(3).

4.2 SU(2) is Simply Connected

SU(2) is simply connected because it is diffeomorphic to the three-sphere S^3, which is simply connected.

What is the 3-sphere? Just as the ordinary sphere S^2 is the set of points at unit distance from the origin in $\mathbb{R}^3$, the three-sphere S^3 is the set of points at unit distance from the origin in $\mathbb{R}^4$:

$$S^3 = \{(x_1, x_2, x_3, x_4) \in \mathbb{R}^4 : x_1^2 + x_2^2 + x_3^2 + x_4^2 = 1\}.$$

Although we cannot visualize it directly, S^3 behaves like a perfectly smooth "higher-dimensional sphere." In particular, it has no holes of any kind, and every loop drawn on it can be shrunk continuously to a point. In topology, this is exactly what it means for a space to be simply connected. (For a few more details on n-Spheres see Appendix A)

Why is SU(2) diffeomorphic to S^3? An element of SU(2) is a 2×2 complex unitary matrix with determinant 1:

$$U = \begin{pmatrix} a & b \\ -b^* & a^* \end{pmatrix}, \qquad |a|^2 + |b|^2 = 1.$$

Writing $a = x_1 + ix_2$ and $b = x_3 + ix_4$ turns this into the condition

$$x_1^2 + x_2^2 + x_3^2 + x_4^2 = 1.$$

Thus each element of SU(2) corresponds to a point on the 3-sphere within $\mathbb{R}^4$, and every point on the 3-sphere gives exactly one element of SU(2). This establishes a smooth one-to-one correspondence between SU(2) and S^3. Hence SU(2) has the same topology as a 3-sphere and is simply connected.

Consequence: Because SU(2) is simply connected, *every* representation of the Lie algebra $\mathfrak{su}(2)$ exponentiates to a representation of SU(2). This follows from the general theorem stated above.

4.3 SO(3) is Not Simply Connected

The situation for SO(3) is different. Although SO(3) also consists of rotation matrices, *its topology is not the same as that of a sphere.*

What is the topology of SO(3)? A key point in describing rotations is that a rotation is determined by an axis $\hat{n}$ and an angle θ. The axis is a point on the 2–sphere S^2, and the angle lies in $[0, \pi]$. However, there is an important ambiguity: the same physical rotation can be described in two different ways. Rotating by angle θ about axis $\hat{n}$ is physically identical to rotating by angle $-\theta$ about the opposite axis $-\hat{n}$. In particular, a rotation by π around $\hat{n}$ is exactly the same as a rotation by π around $-\hat{n}$. This means that opposite points on the axis sphere S^2 represent the same physical rotation when the angle is π.

At first this identification only appears to affect the "boundary" of the axis–angle description. But when we organize all rotations together into a single space—the group SO(3)—this ambiguity becomes global. To see why, imagine trying to assign to each rotation a unique point on some space. If we label a rotation by the pair $(\hat{n}, \theta)$, then the points $(\hat{n}, \pi)$ and $(-\hat{n}, \pi)$ must represent the *same* rotation. Thus the sphere of axes must have its antipodal points identified at the $\theta = \pi$ boundary.

Now, there is a well-known geometric construction that captures exactly this kind of identification: start with a 3–sphere S^3 (a three-dimensional surface in $\mathbb{R}^4$) and identify each point with its antipodal point. This produces the space

$$\mathbb{RP}^3 = S^3/\{\pm 1\}.$$

The quotient $S^3/\{\pm 1\}$ is precisely the space obtained by taking a sphere and identifying opposite points everywhere—not just on a boundary, but consistently throughout the whole space.

Thus the intuitive ambiguity that a rotation by π about $\hat{n}$ is the same as a rotation by π about $-\hat{n}$ leads naturally to the idea that the topology of the rotation group must involve identifying antipodal points. In fact, this construction produces exactly the correct topological space: SO(3) is homeomorphic to $\mathbb{RP}^3 = S^3/\{\pm 1\}$.

Now, because antipodal points are identified, the resulting space contains a "hidden hole": certain loops—such as the one obtained by continuously rotating through 0 to 2π—cannot be contracted to a point. This is the hallmark of a space that is not simply connected.

A particularly important example is the loop obtained by continuously rotating a rigid body through angles 0 to 2π about a fixed axis. Although the physical configuration of the body returns to its starting point after a full 2π rotation, the *path* traced out in SO(3) cannot be shrunk to a point.

The notorious plate trick. The following is the standard example to try and illustrate what a non-contractible loop in SO(3) looks like... Imagine holding a book, a plate, or any rigid object in your hand. Now rotate it smoothly about some fixed axis through angles 0 to 2π (as shown in this video). After the full 2π rotation, the object returns to the same physical orientation in space, so from the point of view of $\mathbb{R}^3$ nothing has changed.

However, if you pay attention to your arm while performing this rotation, you will notice that it becomes *twisted*. There is no way to untwist your arm back to its starting position without rotating the object further. In fact, only after a full 4π rotation can your arm return smoothly to its original configuration without moving the object again.

What this shows is that the path in SO(3) corresponding to a 2π rotation is not the same as the constant loop, even though the final orientation is identical. The path cannot be continuously deformed to a point, because doing so would require untwisting the arm without moving the object, which is impossible. A 4π rotation, on the other hand, *can* be contracted back to the trivial loop by gradually removing the twist.

The existence of such a loop that cannot be shrunk away is exactly what it means for SO(3) to be non-simply connected.

Punchline. Because SO(3) is *not* simply connected, a representation of the Lie algebra $\mathfrak{su}(2)$ does not automatically exponentiate to a representation of SO(3).

5 Summary

- The Lie algebras $\mathfrak{su}(2)$ and $\mathfrak{so}(3)$ are isomorphic, so their algebraic representation theory is the same.
- Every $\mathfrak{su}(2)$ representation exponentiates to SU(2) because SU(2) is simply connected.
- Only integer-spin representations exponentiate to SO(3) because SO(3) is not simply connected.
- SU(2) is the double cover of SO(3), and half-integer representations detect the nontrivial topology of SO(3).

A Appendix: Why a Sphere in $\mathbb{R}^4$ is Called a 3-Sphere

The notation for spheres follows the rule that an *n-sphere* is an n-dimensional surface living inside $\mathbb{R}^{n+1}$. Formally,

$$S^n = \{(x_1, \ldots, x_{n+1}) \in \mathbb{R}^{n+1} : x_1^2 + \cdots + x_{n+1}^2 = 1\}.$$

Although S^n sits inside $(n+1)$-dimensional space, the defining equation imposes one constraint on $(n+1)$ variables, leaving n degrees of freedom. Thus:

$$\dim S^n = n.$$

For example:

- $S^1 \subset \mathbb{R}^2$ is the unit circle (1-dimensional),
- $S^2 \subset \mathbb{R}^3$ is the usual sphere (2-dimensional),
- $S^3 \subset \mathbb{R}^4$ is a three-dimensional manifold.