Lie Groups and Quantum Mechanics

Michael Weiss

1 Introduction

These notes attempt to develop some intuition about Lie groups, Lie algebras, spin in quantum mechanics, and a network of related ideas. The level is rather elementary— linear algebra, a little topology, a little physics. I don't see any point in copying proofs or formal definitions that can be had from a shelfful of standard texts. I focus on a couple of concrete examples, at the expense of precision, generality, and elegance. See the first paragraph on Lie groups to get the flavor of my "definitions". I state many facts without proof. Verification may involve anything from routine calculation to a deep theorem. Phrases like "Fact:" or "it turns out that" give warning that an assertion is not meant to be obvious.

A quote from the Russian mathematician V. I. Arnol'd:

> It is almost impossible for me to read contemporary mathematicians who, instead of saying "Petya washed his hands," write simply: "There is a $t_1 < 0$ such that the image of t_1 under the natural mapping $t_1 \mapsto \text{Petya}(t_1)$ belongs to the set of dirty hands, and a t_2, $t_1 < t_2 \leq 0$, such that the image of t_2 under the above-mentioned mapping belongs to the complement of the set defined in the preceding sentence."

A taste of things to come: consider the spin of an electron. One would like to visualize the electron as a little spinning ball. This is not right, yet not totally wrong. A spinning ball spins about an axis, and the angular velocity vector points along this axis. You can imagine changing the axis by rotating the space containing the ball. [1] Analogously, the quantum spin state of an electron has an associated axis, which can be changed by rotating the ambient space.

The classical concepts of rotation and angular velocity are associated with $SO(3)$, the group of rotations in 3-space. $SO(3)$ is an example of a Lie group.

[1] If you want to be really concrete, imagine a spinning gyroscope fitting snugly in a box. Rotate the box.

Another Lie group, $SU(2)$, plays a key role in the theory of electron spin. Now $SO(3)$ and $SU(2)$ are not isomorphic, but they are "locally isomorphic", meaning that as long as we consider only small rotations, we can't detect any difference. However, a rotation of $360°$ corresponds to a element of $SU(2)$ that is not the identity. Technically, $SU(2)$ is a double cover of $SO(3)$.

Associated with every Lie group is something called its Lie algebra. The Lie algebra is a vector space, but it has additional structure: a binary operation called the Lie bracket. For the rotation group, the elements of the corresponding Lie algebra can be thought of as angular velocities. Indeed, angular velocities are usually pictured as vectors in elementary physics (right hand rule of thumb). The Lie bracket for this example turns out to be the familiar cross-product from vector algebra. (Unfortunately, I won't get round to discussing the Lie bracket.)

The Lie algebras of $SO(3)$ and $SU(2)$ are isomorphic. This is the chief technical justification for the "electron = spinning ball" analogy. The non-isomorphism of $SU(2)$ and $SO(3)$ has subtle consequences. I can't resist mentioning them, though these notes contain few further details. Electrons are fermions, a term in quantum mechanics which implies (among other things) that the Pauli exclusion principle applies to them. Photons on the other hand are bosons, and do not obey the exclusion principle. This is intimately related to the difference between the groups $SU(2)$ and $SO(3)$. Electrons have spin $\frac{1}{2}$, and photons have spin 1. In general, particles with half-odd-integer spin are fermions, and particles with integer spin are bosons.

The deeper study of the electron involves the Dirac equation, which arose out of Dirac's attempt to marry special relativity and quantum mechanics. The relevant Lie group here is the group of all proper Lorentz transformations.

A Rough Road-map. The basic plan of attack: show how elements of $SU(2)$ correspond to rotations; then apply this to the spin of the electron.

I start with the most basic concepts of Lie group and Lie algebra theory. $SO(3)$ is the ideal illustrative example: readily pictured, yet complicated enough to be interesting. The main goal is the double covering result. I do not take the most direct path to this goal, attempting to make it appear "naturally". Once we do have it, the urge to explore some of the related topology is irresistible.

Next comes physics. Usually introductory quantum mechanics starts off with things like wave/particle duality, the Heisenberg uncertainty principle, and so forth. Technically these are associated with the Hilbert space of complex-valued L^2 functions on $\mathbf{R}^3$— not the simplest Hilbert space to start with. If one ignores these issues and concentrates solely on spin, the relevant Hilbert space is $\mathbf{C}^2$. (Feynman's *Lectures on Physics*, volume III, was the first textbook to take this approach in its first few chapters.) $SU(2)$ makes its entrance as a symmetry group on $\mathbf{C}^2$. I conclude with a few hand-waves on some loose ends.

2 Lie Groups

A Lie matrix group is a continuous subgroup of the group of all non-singular $n \times n$ matrices over a field K, where K is either $\mathbf{R}$or $\mathbf{C}$. "Continuous" really is a shorthand for saying that the Lie group is a manifold. The rough idea is that the components of a matrix in the group can vary smoothly; thus, concepts like "differentiable function $f : \mathbf{R} \to G$" should make sense. I'll just say "Lie group" for Lie matrix group, though many mathematicians would groan at this.

Example: $O(n)$ is the group of all orthogonal $n \times n$ matrices, i.e. all matrices A with real components such that $A^t A = \mathbf{1}$. This is just the group of all isometries of $\mathbf{R}^n$ which leave the origin fixed. (Standard calculation: let x be a column vector. Then $(Ax)^t(Ax) = x^t x$, i.e., the norm of x equals the norm of Ax.) Note also that the equations $A^t = A^{-1}$ and $AA^t = \mathbf{1}$ follow from $A^t A = \mathbf{1}$.

If $A^t A = \mathbf{1}$, then we have immediately $\det(A)^2 = 1$, i.e., $\det(A) = \pm 1$. $SO(n)$ is the subgroup of all matrices in $O(n)$ with determinant 1. Fact: $SO(n)$ is connected, and is in fact the connectedness component of $\mathbf{1}$in $O(n)$. I will focus initially on $SO(2)$ and $SO(3)$. These are, colloquially, the groups of rotations in 2-space and 3-space. $O(n)$ is the group of reflections and rotations.

Digression: the well-known puzzle, "why do mirrors reverse left and right, but not up and down?" is resolved mathematically by pointing out that a mirror perpendicular to the y-axis performs the reflection:

$$(x, y, z) \mapsto (x, -y, z)$$

i.e., corresponds to the following matrix in $O(3)$:

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

For some psychological reason, people tend to think of this as the composition of a 180° rotation about the z-axis followed by a reflection in a plane perpendicular to the x-axis: $(x, y, z) \mapsto (-x, -y, z) \mapsto (x, -y, z)$. This makes it seem that the mirror is treating the x and z axes differently (left/right vs. updown), though it really isn't. End of digression.

Example: $SO(2)$, rotations in 2-space. Since $\det(A) = 1$, it is easy to write down the components of A^{-1}. Equating these to A^t, we see that A has the form:

$$\begin{bmatrix} a & -b \\ b & a \end{bmatrix}$$

with the constraint that $a^2 + b^2 = 1$. We can set up a one-one correspondence between this matrix and the complex number $a + ib$ on the unit circle. This is a Lie group isomorphism between $SO(2)$ and the unit circle. We can of course find an angle θ for which $a = \cos\theta$ and $b = \sin\theta$.

Elements of $SO(2)$ have real components, but it is enlightening to consider $SO(2)$ as a subgroup of the group of all non-singular complex 2×2 matrices. Fact: any matrix in $SO(2)$ is similar to a matrix of the form

$$\begin{bmatrix} a+ib & 0 \\ 0 & a-ib \end{bmatrix} = \begin{bmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta} \end{bmatrix}$$

Of course, the new basis vectors have complex components.

Example: $SO(3)$, rotations in 3-space. Fact: any element of $SO(3)$ leaves a line through the origin fixed. This seems obvious to anyone who has rolled a ball, but it is not totally trivial to prove. I will outline two arguments, both instructive.

First, an approach Euclid might have followed. Fact: any isometry of 3-space that leaves the origin fixed is the composition of at most three reflections. For consider the initial and final positions of the x, y, and z axes. One reflection moves the x-axis to its final position, pointing in the correct direction; a second reflection takes care of the y-axis; a third may be needed to reverse the direction of the z-axis. We see that an orientation-preserving isometry is the composition of two reflections. The intersection of the planes of these two reflections gives a line that is left fixed.

Next, a linear algebra approach. We want to show that $A \in SO(3)$ leaves a vector fixed ($Av = v$), or in other words that 1 is an eigenvalue. Note first that the characteristic polynomial for A is a cubic with real coefficients. Hence it has at least one real root. Furthermore, $A^t A = \mathbf{1}$ implies that $A^* A = \mathbf{1}$ trivially, since all components of A are real. By the spectral theorem, A can be diagonalized (working over $\mathbf{C}$) and its eigenvalues are all on the unit circle. Since the characteristic polynomial has real coefficients, non-real roots appear in conjugate pairs. It follows that A is similar to a matrix of the form $\mathrm{diag}(\pm 1, \pm 1, \pm 1)$ or $\mathrm{diag}(\pm 1, e^{i\theta}, e^{-i\theta})$. Since $\det(A) = 1$, and the determinant is the product of the roots, at least one eigenvalue must be 1.

The first argument can also be cast in linear algebra form. This leads to Householder transformations.

3 Lie Algebras

Let G be a Lie group. Let $x(t)$ be a smooth curve in G passing through the unit element $\mathbf{1}$of G, i.e., a smooth mapping from a neighborhood of 0 on the real line into G with $x(0) = \mathbf{1}$. $T(G)$, the tangent space of G at $\mathbf{1}$, consists of all matrices of the form $\left.\frac{dx(t)}{dt}\right|_{t=0}$, or just $x'(0)$ in a less clumsy notation.

$T(G)$ is the Lie algebra of G. I will show in a moment that $T(G)$ is a vector space over $\mathbf{R}$, and I really should (but I won't) define a binary operation $[x, y]$ (the Lie bracket) on $T(G) \times T(G)$.

Proof that $T(G)$ is a vector space over $\mathbf{R}$: if $x(t)$ is a smooth curve and $x(0) = \mathbf{1}$, then set $y(t) = x(kt)$, $k \in \mathbf{R}$. This is also a smooth curve and $y'(0) = kx'(0)$. So $T(G)$ is closed under multiplication by elements of $\mathbf{R}$. (Note this argument fails for complex k.) Similarly, differentiating $z(t) = x(t)y(t)$ (with $x(0) = y(0) = \mathbf{1}$, as usual) shows that $T(G)$ is closed under addition.

Historically, the Lie algebra arose from considering elements of G "infinitesimally close to the identity". Suppose $\epsilon \in \mathbf{R}$ is very small, or (pardon the expression), "infinitesimally small". Then $x'(0)$ is approximately $\frac{x(\epsilon)-x(0)}{\epsilon}$, or (remembering $x(0) = \mathbf{1}$)

$$x(\epsilon) \approx \mathbf{1} + \epsilon x'(0)$$

Historically, $x(\epsilon)$ is a so-called infinitesimal generator of G.

Robinson has shown how this classical approach can be made rigorous, using non-standard analysis. Even without this, the classical notions provide a lot of insight. For example, let n be an "infinite" integer. Then if $t \in \mathbf{R}$ is an ordinary real number (not "infinitesimal"), we can let $\epsilon = t/n$ and so

$$x(\epsilon)^n \approx \left(\mathbf{1} + \frac{tx'(0)}{n}\right)^n \approx e^{tx'(0)}$$

Assume that the left hand side is an ordinary "finite" element of G. Write v for $x'(0)$, an arbitrary element of the Lie algebra $T(G)$. This suggests there should be a map $(t, v) \mapsto e^{tv}$ from $\mathbf{R} \times T(G)$ into G.

In fact, the following is true: for any Lie group G with Lie algebra $T(G)$, we have a mapping exp from $T(G)$ into G such that $\exp(\mathbf{0}) = \mathbf{1}$, and $\exp((t_1 + t_2)v) = \exp(t_1 v)\exp(t_2 v)$, for any $t_1, t_2 \in \mathbf{R}$ and $v \in T(G)$.

It also turns out that the Lie algebra structure determines the Lie group structure "locally": if the Lie algebras of two Lie groups are isomorphic, then the Lie groups are locally isomorphic. Here, the Lie algebra structure includes the bracket operation, and of course one has to define local isomorphism.

Now for our standard example, $SO(n)$. Notation: the Lie algebra of $SO(n)$ is $so(n)$. If you differentiate the condition $A^tA = \mathbf{1}$ and plug in $A(0) = \mathbf{1}$, you will conclude that all elements of $so(n)$ are anti-symmetric. Fact: the converse is true.

Example: $SO(2)$, rotations in 2-space. All elements of $so(2)$ have the form

$$\begin{bmatrix} 0 & -c \\ c & 0 \end{bmatrix}$$

An element of $so(2)$ can be thought of as an angular speed.

In the earlier discussion of $SO(2)$, I set up a one-one correspondence

$$\begin{bmatrix} a & -b \\ b & a \end{bmatrix} \leftrightarrow a + ib$$

between matrices and complex numbers. (We had the restriction $a^2 + b^2 = 1$, but we can drop this and still get an isomorphism between matrices of this form and $\mathbf{C}$.) Then the displayed element of $so(2)$ corresponds to the purely imaginary number ic, and we have a map $(t, ic) \mapsto e^{ict}$ mapping $\mathbf{R} \times so(2)$ onto the unit circle, which in turn is isomorphic to $SO(2)$.

$SO(3)$, rotations in 3-space. Elements of $so(3)$ have the form:

$$\begin{bmatrix} 0 & -c & b \\ c & 0 & -a \\ -b & a & 0 \end{bmatrix} \leftrightarrow (a, b, c) \in \mathbf{R}^3$$

(We'll see the reason for the peculiar choice of signs and arrangement of a, b, and c shortly.)

Fact: the vector (a, b, c) is the angular velocity vector for the above element of $so(3)$. What does this mean? Well first, let $v_0 \in \mathbf{R}^3$ be some arbitrary vector; if $A(t)$ is a curve in $SO(3)$, and we set $v(t) = A(t)v_0$, then $v(t)$ is a rotating vector, whose tip traces out the trajectory of a moving point. The velocity of this point at $t = 0$ is $A'(0)v_0$. It turns out that $A'(0)v_0$ equals the cross-product $(a, b, c) \times v_0$, which characterizes the angular velocity vector. The next few paragraphs demonstrate this equality less tediously than by direct calculation.

Let

$$\hat{\mathbf{x}} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix}, \quad \hat{\mathbf{y}} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{bmatrix}, \quad \hat{\mathbf{z}} = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

so the general element of $so(3)$ can be written $a\hat{\mathbf{x}} + b\hat{\mathbf{y}} + c\hat{\mathbf{z}}$. And $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$, and $\hat{\mathbf{z}}$are simply the elements of $so(3)$ corresponding to unit speed uniform rotations about the x, y, and z axes, respectively— as can be seen by considering their effects on the standard orthonormal basis.

This verifies the equation $A'(0)v_0 = (a, b, c) \times v_0$ for the special cases of $A'(0) = \hat{\mathbf{x}}, \hat{\mathbf{y}}$, and $\hat{\mathbf{z}}$. The general case now follows by linearity.

The Adjoint Representation. Elements of $SO(3)$ act on $\mathbf{R}^3$, or equivalently on orthonormal bases of $\mathbf{R}^3$ (frames).[2] But $SO(3)$ can also be regarded as a set of transformations on the vector space $so(3)$, as we will see in a moment. Intuitively, the triple ($\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$, $\hat{\mathbf{z}}$) takes the place of the standard frame for $\mathbf{R}^3$.

[2] In other words, a element of $SO(3)$ determines a mapping of $\mathbf{R}^3$ to $\mathbf{R}^3$. As it happens, the action is faithful, i.e., the mapping determines the element of $SO(3)$.

General fact: for any Lie group G, there is a homomorphism (also known as a representation) of G into the group of non-singular linear transformations on the vector space $T(G)$, with kernel $Z(G)$, the center of G.

Here's how it goes. For any group G, we have the group of inner automorphism $\mathrm{Inn}(G)$ and a homomorphism $G \to \mathrm{Inn}(G)$ defined by $g \mapsto \iota_g$, where $\iota_g(h) = ghg^{-1}$. The kernel is $Z(G)$. The automorphism ι_g is furthermore determined completely by its effects on any set of generators for G.

Now take G to be a Lie group. Let's consider the effect of ι_g on an "infinitesimal" generator $\mathbf{1} + \epsilon h$, where $h \in T(G)$.

$$g(\mathbf{1} + \epsilon h)g^{-1} = \mathbf{1} + \epsilon ghg^{-1}$$

Or in terms of derivatives, if $x(t)$ is our prototypical smooth curve through $\mathbf{1}$, then the derivative of $gx(t)g^{-1}$ at $t = 0$ is $gx'(0)g^{-1}$. So the vector space $T(G)$ is closed under the map $h \mapsto ghg^{-1}$. (Remember that both G and $T(G)$ are sets of matrices.) $Z(G)$ is clearly contained in the kernel, and in fact: $Z(G)$ is the kernel.

For $SO(3)$, this is rather intuitive. Suppose $h \in SO(3)$ is a rotation about the axis determined by vector $v \in \mathbf{R}^3$. Then ghg^{-1} is a rotation about the axis gv: $(ghg^{-1})(gv) = ghv = gv$. If we think of h as an infinitesimal rotation, then we see that the action of $SO(3)$ on $so(3)$ given by ι_g looks just like the action of $SO(3)$ on $\mathbf{R}^3$.

Only in three dimensions do things work out so neatly. $SO(2)$ is abelian, and the adjoint representation for abelian Lie groups is boring— the trivial homomorphism. And the vector space $so(4)$ has dimension 6, so the adjoint representation gives an imbedding of $SO(4)$ in the group of non-singular 6×6 matrices.

Unitary Matrices: $SU(n)$. Now for a different example. $U(n)$ is the group of unitary $n \times n$ matrices, i.e., complex matrices satifying $A^*A = \mathbf{1}$. An easy computation shows that $|\det(A)| = 1$. $SU(n)$ is the subgroup for which the determinant is 1 (unimodular matrices). Unlike the situation with $O(n)$ and $SO(n)$, the dimensions of $U(n)$ and $SU(n)$ (as manifolds) differ by 1.

The Lie algebras of $U(n)$ and $SU(n)$ are denoted $u(n)$ and $su(n)$, respectively. Differentiating $A^*A = \mathbf{1}$ we conclude that $u(n)$ consists of anti-Hermitian matrices: $B^* = -B$. Note that B is anti-Hermitian if and only if iB is Hermitian.

Fact: if $A(0) = \mathbf{1}$, then $\left.\frac{d \det A(t)}{dt}\right|_{t=0} = \mathrm{tr}\ A'(0)$ (where tr = trace). (Expanding by minors does the trick.) This makes one half of the following fact obvious: the Lie algebra for the Lie group of unimodular matrices consists of all the traceless matrices.

For the special case $SU(2)$, things work out very nicely. Since $\det(A) = 1$, one can write down the components for A^{-1} easily, and equating them to A^*, one concludes that $SU(2)$ consists of all matrices of the form:

$$\begin{bmatrix} a+id & c+ib \\ -c+ib & a-id \end{bmatrix} = a\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + b\begin{bmatrix} 0 & i \\ i & 0 \end{bmatrix} + c\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} + d\begin{bmatrix} i & 0 \\ 0 & -i \end{bmatrix}$$

$$= a\mathbf{1} + b\mathbf{i} + c\mathbf{j} + d\mathbf{k}, \quad a^2+b^2+c^2+d^2 = 1$$

defining **i**,**j**,**k**as the given 2×2 matrices in $SU(2)$. Exercise: these four elements satisfy the multiplication table of the quaternions, so $SU(2)$ is isomorphic to the group of quaternions of norm 1. (The somewhat peculiar arrangement of a, b, c, d in the displayed element of $SU(2)$ is dictated by convention.)

Next, an arbitary anti-Hermitian matrix looks like:

$$\begin{bmatrix} i(a+d) & c+ib \\ -c+ib & i(a-d) \end{bmatrix} = a\begin{bmatrix} i & 0 \\ 0 & i \end{bmatrix} + b\begin{bmatrix} 0 & i \\ i & 0 \end{bmatrix} + c\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} + d\begin{bmatrix} i & 0 \\ 0 & -i \end{bmatrix}$$

$$= ai\mathbf{1} + b\mathbf{i} + c\mathbf{j} + d\mathbf{k}$$

This is traceless if and only if $a = 0$. So we have a canonical 1–1 correspondence between $su(2)$ and $\mathbf{R}^3$, and so also with $so(3)$: $b\mathbf{i} + c\mathbf{j} + d\mathbf{k} \leftrightarrow (b, c, d) \leftrightarrow b\hat{\mathbf{x}} + c\hat{\mathbf{y}} + d\hat{\mathbf{z}}$.

It turns out that this correspondence is a Lie algebra isomorphism. $SU(2)$ and $SO(3)$ are locally isomorphic, but not isomorphic— as we will see next.

$SU(2)$ acts on $su(2)$ via the adjoint representation. But we have a 1–1 correspondence between $su(2)$ and $\mathbf{R}^3$, so we have a representation of $SU(2)$ in the group of real 3×3 matrices. Let A be an element of $SU(2)$ and $v = b\mathbf{i}+c\mathbf{j}+d\mathbf{k}$ be an element of $su(2)$. Note that $\det v = b^2 + c^2 + d^2$. Since the map $v \mapsto AvA^{-1}$ preserves determinants, it preserves norms when considered as acting on $\mathbf{R}^3$. So the adjoint representation maps $SU(2)$ into $O(3)$. Fact: it maps $SU(2)$ onto $SO(3)$.

Incidentally, you can see directly that $v \mapsto AvA^{-1}$ preserves anti-Hermiticity by writing it $v \mapsto AvA^*$.

The kernel of the adjoint representation for $SU(2)$ is its center, which clearly contains $\pm\mathbf{1}$— and in fact, consists of just those two elements. So we have a 2–1 mapping $SU(2) \to SO(3)$. Our double cover! I'll look at the topology of this in a moment.

Physicists prefer to work with the Pauli spin matrices instead of the quaternions. The Pauli matrices are just the Hermitian counterparts to **i**, **j**, and **k**:

$$\mathbf{i} = i\sigma_x, \quad \mathbf{j} = i\sigma_y, \quad \mathbf{k} = i\sigma_z$$

They form a basis (with **1**)for the vector space of Hermitian 2×2 matrices:

$$\begin{bmatrix} a+d & b-ic \\ b+ic & a-d \end{bmatrix} = a\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + b\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} + c\begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix} + d\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

$$= a\mathbf{1} + b\sigma_x + c\sigma_y + d\sigma_z$$

$SU(2)$ acts on the space of traceless Hermitian 2×2 matrices in the same way as on $su(2)$: $h \mapsto ghg^{-1}$.

Picturing the correspondences. We have an embarrassment of riches: so many correspondences that it is easy to get confused. This paragraph tries to fit things into a coherent framework, to make them easier to remember.

We have two fundamental concepts that are easy to visualize:

$$\boxed{\text{Rotations}} \overset{\text{acts on}}{\longrightarrow} \boxed{\begin{array}{l}\text{vectors in 3-space} \equiv \\ \text{angular velocities}\end{array}}$$

The rotation group acts on the space of vectors. For any representation of the rotation group and any representation of the vector space, we would like to have an intuitive grasp of the action. (Since the action preserves distances, one can also consider the action of the group just on the set of vectors of norm 1, i.e., on the sphere.)

Possible representations for the rotation group: $SO(3)$, $SU(2)$, quaternions of norm 1. Possible representations for the vector space: $\mathbf{R}^3$, $so(3)$, $su(2)$, and the space of traceless Hermitian matrices. For $SO(3)$ with $\mathbf{R}^3$, the action is matrix multiplication on the left: $v \mapsto Av$. For $SO(3)$ with $so(3)$, or $SU(2)$ with $su(2)$ or the traceless Hermitian matrices, the action is conjugation: $v \mapsto AvA^{-1} = AvA^*$.

The actions with $SO(3)$ are all faithful. The actions with $SU(2)$ are all two-to-one— A and $-A$ determine the same action (i.e., rotation).

Let's look at the $SU(2)$ actions in more detail. The space of traceless Hermitian matrices consists of all matrices of the form $x\sigma_x + y\sigma_y + z\sigma_z$, $x, y, z \in \mathbf{R}$. This is in one-one correspondence with $su(2)$:

$$x\sigma_x + y\sigma_y + z\sigma_z \leftrightarrow i(x\sigma_x + y\sigma_y + z\sigma_z) = x\mathbf{i} + y\mathbf{j} + z\mathbf{k}$$

So if we understand one $SU(2)$ action, we understand the other. I'll use Pauli matrices from now on.

An arbitary element A of $SU(2)$ looks like

$$A = \begin{bmatrix} a+id & c+ib \\ -c+ib & a-id \end{bmatrix} = a\mathbf{1} + bi\sigma_x + ci\sigma_y + di\sigma_z, \quad a^2+b^2+c^2+d^2 = 1$$

and we see that $A^* = a\mathbf{1} - bi\sigma_x - ci\sigma_y - di\sigma_z$. So the result of acting with A on v can be computed simply by working out the product $(a\mathbf{1} + bi\sigma_x + ci\sigma_y + di\sigma_z)(x\sigma_x + y\sigma_y + z\sigma_z)(a\mathbf{1} - bi\sigma_x - ci\sigma_y - di\sigma_z)$. For this we need the multiplication table of the σ's. This is simply:

$$\begin{array}{ccccccc}\sigma_x^2 & = & \sigma_y^2 & = & \sigma_z^2 & = & \mathbf{1} \\ \sigma_x\sigma_y & = & -\sigma_y\sigma_x & = & i\sigma_z & & \\ \sigma_y\sigma_z & = & -\sigma_z\sigma_y & = & i\sigma_x & & \\ \sigma_z\sigma_x & = & -\sigma_x\sigma_y & = & i\sigma_y & & \end{array}$$

The easiest way to check this is to work out the action of the σ's on $(p,q) \in \mathbf{C}^2$:

$$\sigma_x \begin{bmatrix} p \\ q \end{bmatrix} = \begin{bmatrix} q \\ p \end{bmatrix}$$

$$\sigma_y \begin{bmatrix} p \\ q \end{bmatrix} = \begin{bmatrix} -iq \\ ip \end{bmatrix}$$

$$\sigma_z \begin{bmatrix} p \\ q \end{bmatrix} = \begin{bmatrix} p \\ -q \end{bmatrix}$$

We have here non-commuting Hermitian matrices whose product is not Hermitian— in fact is anti-Hermitian. Now, one has an analogy between matrices and complex numbers, under which “Hermitian” goes with “real”, “anti-Hermitian” goes with “purely imaginary”, and “unitary” goes with “on the unit circle”. The σ matrices provide a striking example of the analogy breaking down. Non-commutativity is the culprit— for of course the product of commuting Hermitian matrices is Hermitian.

Example: what rotation does $i\sigma_z$ represent? Its action on the unit x vector, σ_x, is just $\sigma_x \mapsto i\sigma_z\sigma_x(-i)\sigma_z = i\sigma_y\sigma_z = -\sigma_x$. Similar calculations show that y goes to $-$y and z stays put, so we have a 180° rotation about the z-axis. Similarly for $i\sigma_x$ and $i\sigma_y$.

The exponential map provides a mapping from $su(2)$ into $SU(2)$:

$$ib\sigma_x + ic\sigma_y + id\sigma_z \mapsto \exp(ib\sigma_x + ic\sigma_y + id\sigma_z)$$

(Warning: the exponential of a sum is *not* in general the product of the exponentials, because of non-commutativity.) For the rotation group (as we've seen) this says simply that an angular velocity determines a rotation— e.g., by the prescription “rotate at the given angular velocity for one time unit”. The basis of “angular velocities” in $su(2)$ is $(i\sigma_z, i\sigma_y, i\sigma_z)$. Let us consider rotations about the z-axis.

$$e^{ib\sigma_z} = \begin{bmatrix} e^{ib} & 0 \\ 0 & e^{-ib} \end{bmatrix} = \cos b\,\mathbf{1} + i\sin b\,\sigma_z$$

(since $i\sigma_z$ acts separately on each coordinate.) Perhaps the clearest way to exhibit the action of this rotation on $v = x\sigma_x + y\sigma_y + z\sigma_z$ is to work entirely in matrix form:

$$\begin{bmatrix} e^{ib} & 0 \\ 0 & e^{-ib} \end{bmatrix} \begin{bmatrix} z & x-iy \\ x+iy & -z \end{bmatrix} \begin{bmatrix} e^{-ib} & 0 \\ 0 & e^{ib} \end{bmatrix} = \begin{bmatrix} z & e^{2ib}(x-iy) \\ e^{-2ib}(x+iy) & -z \end{bmatrix}$$

i.e., a rotation about the z-axis of $-2b$ radians. Exercise: the same sort of thing holds for $i\sigma_y$ and $i\sigma_z$.

So one can directly picture the action of $SU(2)$ on vectors in 3-space. The "double angles" $2b$, etc., stem from the two multiplications in the action: $v \mapsto AvA^*$. And the double angles in turn are the reason the map from $SU(2)$ to $SO(3)$ is two-to-one.

$SU(2)$ acts on $\mathbf{C}^2$ via left multiplication: $v \mapsto Av$, where v is a column vector. Can one picture v as some kind of geometric object in 3-space? Yes indeed! An object known as a spin vector embodies v geometrically. But I won't get to them.

Topology. Fact: $SU(n)$ is simply connected. So $SU(2)$ is the universal covering space of $SO(3)$. The covering map is 2–1. It follows from standard results in topology that the fundamental group of $SO(3)$ is $\mathbf{Z}_2$.

Recall that elements of $SU(2)$ all take the form $a\mathbf{1} + b\mathbf{i} + c\mathbf{j} + d\mathbf{k}$ with $a^2 + b^2 + c^2 + d^2 = 1$. Therefore $SU(2)$ is topologically S^3, the 3-dimensional hypersphere. (This is enough to show that $SU(2)$ is simply connected.) Because the kernel is $\{\pm\mathbf{1}\}$ (verify!), antipodal points are identified on mapping into $SO(3)$, so $SO(3)$ is homeomorphic to real projective 3-space.

This can be seen another way. A rotation can be specified by a vector along the axis of the rotation, with magnitude giving the angle of the rotation. This serves to identify elements of $SO(3)$ with points inside or on a ball of radius π. However, antipodal points on the surface of the ball represent the same rotation. The resulting space (a three-dimensional ball with antipodal points on the surface identified) is well-known to be homeomorphic to real projective 3-space. (If you think about this argument for a bit, you should see an implicit use of the exponential mapping from $so(3)$ into $SO(3)$.)

A loop in the topological space $SO(3)$ can be visualized as a continous "trajectory" of rotations: we take a rigid object and turn it around in some fashion, finally restoring it to its original orientation. The following fact can be deduced from this: if a solid object is connected by threads to a surrounding room, and the object is turned through 720°, then the threads can be untangled without turning the object any more. However, if the object is turned through 360°, then the threads cannot be untangled. (The two-thread version of this is known as "Dirac's string trick".) In this sense, a continuous turn through 360° is not the same as no turn at all (but a 720° turn is.)

4 Quantum Mechanics: Two-state Systems

The framework of quantum mechanics rests on three pillars: the Hilbert space of quantum states; the Hermitian operators, also called observables; and the unitary evolution operators. I start by trying to attach some pictures to these abstractions.

The simplest classical system consists of a single point particle coasting along in space (perhaps subject to a force field). To "quantize" this, you'll need the Hilbert space of complex-valued L^2 functions on $\mathbf{R}^3$, and you'll encounter unbounded operators on this space. So goes the tale of history: Heisenberg, Schrödinger, Dirac and company cut their milk teeth on this problem.

I will take an ahistorical but mathematically gentler approach. The Hilbert space for a two-state quantum system is $\mathbf{C}^2$, and the operators can all be represented as 2×2 complex matrices. The spin of an electron provides a physical example. That is, if we simply ignore position and momentum ("mod them out", so to speak), we have a physical picture that can be modelled by this (relatively) simple framework. (As noted before, Feynman's *Lectures*, volume III, starts off like this.)

Quantum states. Generally, the quantum state of a physical system is specified by a non-zero vector in a Hilbert space over the complex numbers (call it H), with the understanding that x and cx specify the same quantum state, where x is in H and c is a non-zero complex number. In other words, the set of quantum states is $(H - \{0\})/\sim$, where $\sim$ is the equivalence relation defined by $x \sim y$ if and only if $x = cy$ for some $c \neq 0$. Sometimes I will abuse language, and say "the state vector of the system is $v \in H$" instead of saying "the state of the system is specified by $v \in H$."

The state of our spinning electron is therefore specified by giving a ratio $a : b$ of two complex numbers (since $H = \mathbf{C}^2$), not both zero. Mathematically, this is just the complex projective line. The correspondence between points on the complex projective line and points on the Riemann sphere is well-known; I will make use of it in a moment. [3]

Physically, we want to picture the electron as a little ball, spinning about some axis. It has angular momentum. In more detail, the angular momentum vector

[3] A quick review: write $(a : b)$ for the equivalence class of (a, b). We will associate either a complex number w or else ∞ with each class $(a : b)$. If $b \neq 0$, then $(a : b) = (w : 1)$, where $w = a/b$. All pairs of the form $(a, 0)$ belong to the class $(1 : 0)$. So we associate a/b with $(a : b)$ if $b \neq 0$, and ∞ with $(1 : 0)$. Mapping the complex plane plus ∞ to the Riemann sphere via the usual stereographic projection completes the trick. Some sample points to bear in mind: the north pole is $(1 : 0)$; the south pole is $(0 : 1)$; points on the equator have the form $(e^{i\theta} : 1)$. (For the purist, the special treatment of ∞ rankles. It is not singled out on either the complex projective line or on the Riemann sphere. Later I will show how to set up the correspondence without this blemish.)

is a vector pointing along the axis of rotation, with length proportional to the speed of rotation. This defines the vector up to sign. The sign ambiguity is resolved, conventionally, by the right hand rule of thumb: if you curl the fingers of your right hand in the direction of rotation, the thumb points in the direction of the angular momentum vector.

Classically, the electron could be spinning at any speed, so its angular momentum could have any magnitude. In quantum mechanics, the angular momentum is quantized: its magnitude (measured along the axis of rotation) must be $\hbar/2$ for any spin $\frac{1}{2}$ particle. (Note: $\hbar = h/2\pi$, where h of course is Planck's constant.)

The quantization of angular momentum, although inexplicable classically, is easy enough to picture: we just stipulate that our spinning ball must be spinning at a particular rate. So we should be able to specify the spin state of the electron just by giving a direction in 3-space, or equivalently, by picking a point on the sphere S^2. But I just noted that the set of states is "isomorphic" to the Riemann sphere. So everything fits.

Hermitian observables. So much for quantum states. Next I look at measurement. Generally, a measurement of something in quantum mechanics (spin, energy, momentum ...) is associated with a Hermitian operator, say A, on the Hilbert space of states. In finite dimensional cases, possible results of the measurement must be eigenvalues of A. If v is an eigenvector of A with eigenvalue λ, then measuring a system with state vector v will always give you the result λ.

What happens if the state vector of the system is not an eigenvector? Here we skate perilously close to philosophy! To make things simple[4], suppose $(v_1, ..., v_n)$ is an orthonormal basis for H consisting of eigenvectors for A, say $Av_i = \lambda_i v_i$. Suppose the system has state vector $w = a_1 v_1 + \ldots + a_n v_n$, and assume w is normalized (i.e., $|w| = 1$). The so-called "collapse" interpretation of quantum mechanics says: (i) The measurement process forces the system to jump randomly to a state with state vector v_1, or v_2, or $\ldots$ or v_n. (ii) The probability that the system jumps to a state with state vector v_i is $|a_i|^2 = |\langle w, v_i \rangle|^2$. (iii) If the system ends up with state vector v_i, then the measurment yields result λ_i. (Note that you'll get real-valued measurement results because A is Hermitian.)

Nearly everyone agrees that the collapse interpretation will correctly predict the results of experiments. Whether it is what's "really going on" is the subject of endless debates. What we have here are rules for calculating probabilities. At least four different philosophies have been draped around the rules.

How does measurement look for the electron? Say we want to measure the

[4]In the infinite dimensional case, you have to use the spectral decomposition of A instead of an orthonormal basis of eigenvectors; another reason why spin is simpler than position.

component of spin along the z-axis. Since the electron is charged, it acts like a little magnet, with north and south poles along the axis of rotation. (Circulating charge causes a magnetic field. Think of an electromagnet— a coil of wire with an electric current flowing around in it.) A physicist would say that the electron has a magnetic moment.

We can use the magnetic moment to measure the spin. Stern and Gerlach got a Nobel prize for doing just that. They sent a beam of electrically neutral silver atoms through a magnetic field. It turns out that the magnetic moments of the electrons in a silver atom cancel out in pairs except for one electron, so we can pretend (so far as the spin is concerned) that we're looking at a beam of electrons passing through a magnetic field.

The magnetic field was designed to produce a force on the electrons. An electron with spin pointing up would look like an $\boxed{\substack{\text{N}\\ \text{S}}}$ magnet, and would experience an upward force; an electron with spin pointing down would look like an $\boxed{\substack{\text{S}\\ \text{N}}}$ magnet, and would experience a downward force. Classically, you would expect an electron with spin at angle α to the vertical to experience an upward force proportional to $\cos\alpha$.[5] So the electron beam should be spread out into a vertical smear, according to classical mechanics.

In fact, the beam splits into two beams, one up, one down. In other words, if we measure the component of the spin along the vertical axis, we always find that the spin is entirely up or entirely down. *This is the most basic sense in which the "spinning ball" analogy is wrong.* The same two-valued behavior holds for *any* measurement axis.

Classically this is inexplicable. How can the electron have spin up and spin sideways at the same time? Answer: it doesn't. After you've measured the spin along the z-axis, the electron has vertical spin (say spin up). If you take your vertically spinning electron and measure its spin along the x-axis, you have a 50–50 chance at getting spin left or spin right. If you now repeat the spin measurement along the z-axis, you have a 50–50 chance of getting spin up or spin down. The x-axis measurement has destroyed the information obtained from the z-axis measurement.

Let A be the Hermitian operator corresponding to "measure the spin along the z-axis". The eigenvalues (i.e., possible results) will be 1 and -1, if we choose our units right. Pick a basis of two eigenvectors; then the matrix for A in this basis is just $\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$, i.e., the Pauli matrix σ_z. (Common notation for the eigenvectors is $|\text{up}\rangle$ and $|\text{down}\rangle$, although $|\text{dead}\rangle$ and $|\text{alive}\rangle$ are popular for that other famous two-state system, Schrödinger's Cat.)

[5]Why wouldn't the electron simply snap into alignment with the magnetic field? Answer: the spinning electron would act like a gyroscope, and precess in response to the torque exerted by the field. Thus it would maintain its angle of inclination to the field.

If σ_z is here, can σ_x and σ_y be far behind? In fact, these are the matrices for measuring the x-component (respectively y-component) of spin, provided we continue to use the same basis $|\text{up}\rangle$ and $|\text{down}\rangle$.

It turns out that $(1:0)$ represents spin up along the z-axis, $(1:1)$ represents spin along the x-axis, and $(i:1)$ represents spin along the y-axis. In a different notation, $|\text{up}\rangle + |\text{down}\rangle$ and $i|\text{up}\rangle + |\text{down}\rangle$ are the state vectors for these two spin directions. The x-axis and y-axis state vectors are not eigenvectors of σ_z. The rules for calculating probabilities (clothed in any philosophy you like) yield the 50–50 chances mentioned earlier.

As an exercise, you may like to chew on these remarks: if two measurements can be done simultaneously, then the associated Hermitian operators must have the same set of eigenvectors, and so the operators must commute. But the σ matrices don't commute. This accounts mathematically for the non-intuitive (or at least non-classical) results of the Stern-Gerlach experiment. The Heisenberg uncertainty principle stems from the same sort of considerations.

Now a general comment. Any linear operator on a Hilbert space H induces a mapping on the space of states, since if x and cx are two state vectors for the same state, then Ax and $Acx = cAx$ will represent the same state. Can I dispense with the Hilbert space entirely and just work with the space of states and the induced mappings? The answer is yes, but it would be inconvenient. If A is an observable with eigenvector v, say $Av = \lambda v$, then the eigenvalue λ has physical significance. But when we look at the action of A on the space of states, all we notice (at first) is that A leaves the state corresponding to v fixed.

Nonetheless, the results of measurement are encoded in the action of A on the space of states. A and cA, $c \neq 0$, induce the same mapping on the states, and the converse holds for the cases of interest to us (if A and B induce the same state mapping, then $A = cB$ for some scalar $c \neq 0$). This scalar c will be real for Hermitian A and B. If $c \neq 1$, then A and B really represent the same measurement, but expressed in different units (e.g., foot-pounds vs. ergs.)

Example: suppose $Av = \lambda v$ and $Aw = \mu w$, and $Bv = \lambda' v$, $Bw = \mu w$, with $\lambda \neq \lambda'$. A and B do the same thing to the quantum states determined by v and w— namely, the states are left fixed. However, A and B send the state determined by $v + w$ to different states.

Unitary evolution operators. Suppose I change the state of a quantum system "smoothly". For example, I could move the system through space, or I could "move it through time" (i.e., just wait— hence the term, "evolution operator"), or I could (surprise!) rotate it.

It happens quite generally in quantum mechanics that all such state changes are induced by unitary operators. Proving this clearly would require some physical assumptions, and I won't go into this at all. For our spinning electron, it turns out you can even assume a little more: the change of state caused by rotating the electron is induced by an operator in $SU(2)$. The operator in question is called a rotation operator.

How should we visualize the action of a rotation operator R on a state vector v? We saw how to picture R as a rotation in 3-space by looking at its effects on traceless Hermitian matrices: $A \mapsto RAR^*$, where $A = x\sigma_x + y\sigma_y + z\sigma_z$. How can we hook up the action of R on state vectors with the action of R on traceless Hermitian matrices? It seems we need a correspondence between states (say $(a:b)$) and matrices of the form $x\sigma_x + y\sigma_y + z\sigma_z$. We won't get quite this, but we'll get something just as good.

The trick is to set up a correspondence between states and yet another kind of matrix: a projection matrix. You probably noticed that the matrix σ_z does a pretty good job specifying the state "spin up along the z-axis". As it turns out, σ_z is not a projection matrix, but it corresponds in a natural fashion to $\frac{1}{2}\mathbf{1} + \frac{1}{2}\sigma_z$), which is.

Here's how it goes for an arbitrary state vector $v = a|\text{up}\rangle + b|\text{down}\rangle$. Suppose v is normalized, so $|a|^2 + |b|^2 = 1$. The projection matrix for v is given by taking the product of the column vector v with the row vector v^*:

$$vv^* = \begin{bmatrix} a \\ b \end{bmatrix} [a^*, b^*] = \begin{bmatrix} aa^* & ab^* \\ a^*b & bb^* \end{bmatrix}$$

(Standard physicists' notation is $|v\rangle$ for v and $\langle v|$ for v^*. The product is $|v\rangle\langle v|$. The norm is $\langle v|v\rangle$. Mathematicians prefer to talk about a vector space and its dual instead of column vectors and row vectors, but these notes prefer concreteness to elegance.)

Now, vv^* is a Hermitian matrix with determinant 0 (check!). It must therefore take the form:

$$\begin{bmatrix} t+z & x-iy \\ x+iy & t-z \end{bmatrix} = t\mathbf{1} + x\sigma_x + y\sigma_y + z\sigma_z, \quad t^2 - x^2 - y^2 - z^2 = 0$$

(If the appearence of $t^2 - x^2 - y^2 - z^2$ makes you think "Special Relativity!", you're on the right track, but I won't get into that.) However, the trace is not 0, but $aa^* + bb^* = 2t$. Since I took v to be normalized ($|a|^2 + |b|^2 = 1$), it follows that $t = \frac{1}{2}$.

So we have a mapping from state vectors to Hermitian matrices of the form $\frac{1}{2}\mathbf{1} + x\sigma_x + y\sigma_y + z\sigma_z$ with $x^2 + y^2 + z^2 = \frac{1}{4}$. And the latter are in an obvious one-one correspondence with points on a sphere of radius one-half.

The mapping $v \mapsto vv^*$ (restricted to normalized vectors) actually establishes a one-one correspondence between *states* and our special class of Hermitian matrices. For let c be a complex number of norm 1; then $(cv)(cv)^* = vv^*$.

Why do I call vv^* a projection matrix? Answer: by analogy with projections in ordinary real vector spaces, say $\mathbf{R}^3$. If $\boldsymbol{v}$ is a vector of norm 1, and $\boldsymbol{w}$ is an arbitrary vector, then the projection of $\boldsymbol{w}$ "along the vector $\boldsymbol{v}$" (i.e., in the subspace spanned by $\boldsymbol{v}$) is $(\boldsymbol{v} \cdot \boldsymbol{w})\boldsymbol{v}$. Analogous to this, we define $\mathrm{proj}_v(w) = \langle v, w\rangle v$, using the notation $\langle v, w\rangle$ for the inner product. In the "row vector, column vector" notation, this is $(v^*w)v = v(v^*w) = vv^*w$. In physicists' notation, this is $|v\rangle\langle v|w\rangle$.

We have acquired a new way of picturing the action of $SU(2)$ on 3-space. The formula $v \mapsto Avv^*A^*$ captures it succinctly. The mapping $v \mapsto vv^*$ sets up a one-one correspondence between the states $(a : b)$ (i.e., the complex projective line) and points on a sphere in 3-space. In fact this is just the Riemann sphere mapping!

So the quantum states for the spin of an electron can be pictured as points on a sphere. Elements of $SU(2)$ correspond to the change in state induced by rotating the electron, and this action of $SU(2)$ can be pictured as a rotation of the sphere. The naive pictures match up with the $SU(2)$ formalism flawlessly. The element $-\mathbf{1}$ of $SU(2)$ induces the identity mapping on the space of states, since v and $-v$ represent the same quantum state.

A simple computation illustrates how everything meshes. The rotation operator for a clockwise 90° rotation about the y-axis is $\frac{1}{\sqrt{2}}(\mathbf{1}+i\sigma_y)$. Indeed, if you work out $(\mathbf{1}+i\sigma_y)\sigma_x(\mathbf{1}-i\sigma_y)$, you get $2\sigma_z$, and likewise $(\mathbf{1}+i\sigma_y)\sigma_z(\mathbf{1}-i\sigma_y) = -2\sigma_x$. The x-axis maps to the z-axis, and the z-axis maps to minus the x-axis.

The example of electron spin illustrates two features of quantum mechanics very clearly.

- **Probabalistic character:** The quantum state does not uniquely determine the result of experiment. Hence Einstein's famous complaint, "I shall never believe that God plays dice with the universe." (Perhaps he plays cards with the physicists?)

- **Correspondence with classical physics:** Classical physics emerges from quantum mechanics by taking averages.

Some more technical features, also embodied in this example, and typical of quantum mechanics:

- **Need for complex numbers:** The neat correspondence with the classical spinning ball picture wouldn't work if we did everything over $\mathbf{R}$.

- **Non-commuting observables:** You cannot simultaneously measure the x and z components of spin (for example), because σ_x and σ_z do not commute.

- **Symmetry groups and observables:** The rotation symmetry group gives rise indirectly to the σ matrices, and ultimately to the notion of angular momentum. The mathematical basis is the Lie groupLie algebra correspondence.

Had I started with the first historical example, the single spinless particle coasting in space, I would be illustrating the same morals with different actors:

- **Need for complex numbers:** The appropriate Hilbert space is the space of complex-valued L^2 functions on 3-space.
- **Non-commuting observables:** The momentum and position operators do not commute, and you cannot simultaneously measure position and momentum.
- **Symmetry groups and observables:** The group of translations in 3-space gives rise to Lie group acting on the L^2 Hilbert space; the momentum operator emerges from the corresponding Lie algebra.

5 Loose Ends

There are so many loose ends that it seems pointless to try to tie them all up. I will finish off with a few observations, meant more to tantalize than enlighten.

The Lie bracket. Can one reconstruct the Lie group from the Lie algebra? In general, the multiplication table of a group is determined if you know the multiplication table for its generators; why not try this with the "infinitesimal" generators? If you try this approach, you will find you need to know the commutators of infinitesimal elements, like $x(\epsilon)y(\epsilon)x(\epsilon)^{-1}y(\epsilon)^{-1}$.

My "definition" of the Lie algebra involved approximating infinitesimal generators by Taylor expansions out to the first order. In other words, I used only first order derivatives. But to the first order, the commutators are zero!

Say we approximate an "infinitesimal" element of the Lie group out to the second order:

$$x(\epsilon) \approx \mathbf{1} + \epsilon x'(0) + \frac{\epsilon^2}{2}x''(0)$$

If you work out the commutator, you will find expressions of the form $vw - wv$ appearing, where v and w belong to the Lie algebra. And one can verify that $vw - wv$ belongs in fact to the Lie algebra, as I've defined it, although vw and wv in general don't.

Remarkably, knowledge of these second order terms completely specifies the structure of the Lie group near the identity. That is, if the Lie algebras are isomorphic, then the Lie groups are locally isomorphic. Third and higher-order terms are not needed.

Special relativity, spinors, and the Dirac equation. The crucial Lie group for special relativity is the Poincaré group: all transformations of Minkowski 4-space (spacetime) that preserve the Minkowski pseudo-metric. The Lorentz group is the subgroup that leaves the origin fixed, and the proper Lorentz group is the subgroup of orientation preserving Lorentz transformations. The proper Lorentz group in turn contains the rotation group of 3-space, $SO(3)$.

Just as $SU(2)$ is the double cover of $SO(3)$, so $SL(2)$ is the double cover of the proper Lorentz group, where $SL(2)$ is the group of unimodular 2×2 complex matrices.

Say A is in $SL(2)$ and v is in $\mathbf{C}^2$. It turns out to be important to pry the mapping $vv^* \mapsto Avv^*A^*$ apart into $v \mapsto Av$ and $v^* \mapsto v^*A^*$. The vector v can be pictured as a geometric object consisting of a vector in space (rooted at the orgin) with an attached "flag", i.e., a half-plane whose "edge" contains the vector. Moreover, if the flag is rotated through 360°, v turns into $-v$! (Recall the earlier remarks on untangling threads.) Such an object is called a spin vector. And just as one can create tensors out of the raw material of vectors, so one creates spinors out of spin vectors.

Dirac invented spinors in the course of inventing (or discovering) the Dirac equation, the correct relativistic wave equation for the electron. As it happens, the σ matrices are not enough to carry the load; Dirac had to go up to 4×4 matrices (called the Dirac matrices). The σ matrices are imbedded in the Dirac matrices.

I won't repeat the story of how Dirac discovered antiparticles. Nor the story of how he rediscovered knitting and purling (see Gamow's *Thirty Years that Shook Physics*.)

Spin and Statistics. The spin-statistics theorem of quantum field theory says that particles with half-odd-integer spin (like the electron) must be fermions, while particles with integer spin (like the photon) must be bosons. Fermions obey Fermi-Dirac statistics, and hence obey the Pauli exclusion principle. Bosons obey Bose-Einstein statistics.

The difference in statistics stems from the properties of the exchange operator. This is a unitary operator, say P_{ex}, which represents the effect of exchanging two identical fermions, or two identical bosons. For the fermion case, one has at a critical point in the calculations

$$P_{\mathrm{ex}}v = -v$$

and for the boson case,

$$P_{\text{ex}} v = v$$

The minus sign for fermions ultimately derives from the double covering of $SO(3)$ via $SU(2)$. Spinors also get into the act. Since I don't fully understand the story myself, this seems like a good place to stop.