2. Lie Groups, Lie Algebras, and Their Representations

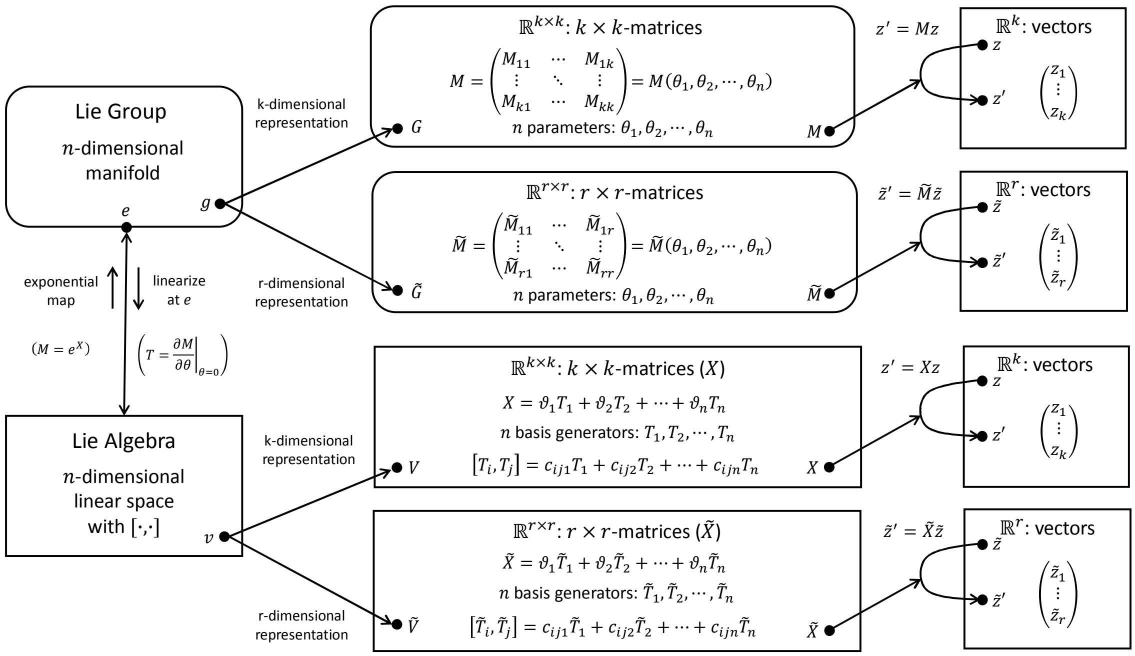

Before examining specific examples, we introduce the general ideas behind Lie groups, Lie algebras, and their representations. The above diagram helps us to get oriented. We start by giving a bird's eye view of the subject, then we go into some more details in Sections 2.1 to 2.4.

The basic idea behind a *Lie group* is to describe continuous symmetry transformations that depend on a couple of parameters. For example, all possible rotations of a 3D object parametrized by the rotation angles about the x, y, and z-axes form a Lie group.

The basic idea behind a *representation* is to map abstract transformations to concrete matrices. For example, the group of all possible rotations of a 3D object can be mapped to the set of 3×3 rotation matrices. These matrices, in turn, act on vectors, which are the elements of the *representation space*. We say that these vectors *furnish* the representation of the group. In our example, the representation space consists of 3-component column vectors.

Interestingly, Lie groups do not have just one, but many (usually infinitely many) representations! We will see many examples of that.

The basic idea behind a *Lie algebra* is to provide a simplified (linearized) description of the Lie group by considering only small transformations. Surprisingly, almost no information is lost by going from the Lie group to the Lie algebra! As with the Lie group, the abstract Lie algebra can be made concrete by mapping it to a set of matrices. These matrices act on the same representation space as the matrices of the corresponding group representation.

We can always go from a Lie-group representation to the corresponding Lie-algebra representation by differentiating (linearizing) it. Going from a Lie algebra back to the Lie group by integrating it is a bit more tricky because there can be more than one group for the same algebra. This backwards path is known as the *exponential map*.

The above diagram illustrates two different representations of the same Lie group: a k-dimensional representation in the upper branch and an r-dimensional one in the lower branch. The associated Lie algebra has two corresponding representations, which are shown below the group representations. We can identify the following items in the diagram:

- *The elements of the representation space*. In the upper representation, these are the k-dimensional column vectors z and in the lower representation, these are the r-dimensional column vectors $\tilde{z}$.
- *The elements of the Lie-group representation*. In the upper Lie-group representation, these are the $k \times k$ matrices M and in the lower representation, these are the $r \times r$ matrices $\widetilde{M}$. The matrices M and $\widetilde{M}$ act on the vectors in the representation space by matrix-vector multiplication: Mz and $\widetilde{M}\tilde{z}$. These matrices perform *transformations* on the vectors.
- *The elements of the Lie-algebra representation*. In the upper Lie-algebra representation, these are the $k \times k$ matrices X and in the lower representation, these are the $r \times r$ matrices $\tilde{X}$. The matrices X and $\tilde{X}$ act on the vectors in the representation space by matrix-vector multiplication: Xz and $\tilde{X}\tilde{z}$. These matrices are known as *generators* because they generate transformations.
- *The maps from the abstract Lie-group and Lie-algebra to the concrete representations*. These maps must be such that the matrix representations "do the same thing" as the abstract group and algebra. More technically, the maps must preserve the algebraic structure of the Lie group and Lie algebra. Such maps are known as *homomorphisms*.

The above diagram will serve as a template for the examples that follow. We will fill specific groups, matrices, and vectors into the boxes. We will also encounter representations that do not act on column vectors, but on objects such as spinors, tensors, functions, and fields. Nevertheless, we can still illustrate them using our diagram.

Next, we'll look at Lie groups, Lie-group representations, Lie algebras, and Lie-algebra representations one by one. Don't worry if not everything makes sense. We will go through many examples that illustrate and clarify the theory. The idea is to learn by exploring examples! You may want to skim over the subsequent theory and come back for more details after you have gone through some of the examples.

2.1 Lie Groups

A Lie group is a *set* with an infinite number of elements plus some structure that we'll come to in a moment. The elements in the set are *continuous*, that is, they can be specified by n real parameters, where n depends on the Lie group in question. For example, think of the set of all possible rotations of a 3D object. Any rotation can be specified by three real parameters ($n = 3$). These parameters could be the rotation angles about the x, y, and z axes, or they could be the orientation of the rotation axis in space (two parameters) and the rotation angle about this axis.

In addition to being a set, the Lie group has a *composition operation*: given element g and element h, we can form a new element $f = g \circ h$. A good example for such an operation is the composition of two 3D rotations: a first rotation h by some angle about some axis is followed by a second rotation g by another angle about another axis. From now on, we'll use the simpler notation $f = gh$ instead of $f = g \circ h$. The composition operation must satisfy the following four group axioms:

- *Closure*. The composition gh must again be a member of the group. Example: the composition of two 3D rotations is another 3D rotation.
- *Associativity*: $f(gh) = (fg)h$. Example: the grouping of 3D rotations doesn't matter.
- Existence of an *identity element* e such that $ge = eg = g$ for all g. Example: "no rotation" is the identity element of the group of rotations.
- Existence of an *inverse element* g^{-1} for all g such that $gg^{-1} = g^{-1}g = e$. Example: rotating counterclockwise is the inverse operation of rotating clockwise (by the same angle and about the same axis).

Do the elements of a Lie group commute? In general, they do not: $gh \neq hg$. Again, 3D rotation provides a good example: we get two different outcomes when we rotate about one axis and then rotate about another axis versus when we do the same in the reversed order. But if both rotations are about the *same axis*, then the order doesn't matter. So, while Lie groups in general are not commutative, some simple Lie groups are (e.g., 2D rotation or translation).

Continuous *symmetry transformations* together with the operation of composition always form a Lie group. Examples of such transformations are translations, rotations, Lorentz boosts, and gauge transformations. We can think of these as coordinate transformations due to a frame change (passive transformation) or as coordinate transformations due to an action on the physical object itself described in a fixed frame (active transformation).

Besides the discussed algebraic structure, Lie groups also possess a *topological structure*. Because the elements are smooth functions of n *continuous* parameters, each element in the group "lives" in a neighborhood of other elements. Globally, all the elements of the Lie group form a smooth n-dimesional *manifold*. You may visualize this manifold as a sphere and the group elements as the points that make up the sphere. However, this is just a crutch to help our intuition: there is no actual Lie group that has the topology of an ordinary sphere, S^2. Unfortunately, actual Lie-group topologies are either too trivial to be representative (like the circle, S^1) or too high-dimensional to be useful for intuition building (like the 3-sphere, S^3, which is a sphere in a 4-dimensional space).

2.2 Lie-Group Representations

The elements of a Lie group are abstract objects. To make them more concrete, we map them to (invertible) square matrices, which act as linear transformations on the vectors in the representation space. In the diagram, the abstract group element g is mapped to the concrete matrix G, similarly, m is mapped to M, and so on. The matrix M acts on the vector z in the representation space by matrix-vector multiplication: $z' = Mz$. We call a map from the abstract group elements to the concrete matrices a *representation* of the Lie group (sometimes the matrices themselves or even the objects in the representation space are also referred to as a representation of the group). To deserve this name, the matrices must "behave in the same way" as the abstract group elements when composed together. In other words, the map must preserve the algebraic structure of the group. A map of this kind is called a *group homomorphism*. For example, we can represent abstract 3D rotations by a parametrized 3×3 rotation matrix with the x, y, z rotation angles as parameters.

Surprisingly, Lie groups do not have just one, but many (usually infinitely many) representations! Different representations often act on representation spaces with different numbers of dimensions. For example, 3D rotations can be represented not only in the familiar 3-dimensional space, by what is known as the *defining* or *fundamental representation*, but also in a 5-dimensional space (or any other odd-dimensional space)! This 5-dimensional representation is *not* just rotating a 3-dimensional subspace, leaving the remaining two dimensions invariant. It is a truly 5-dimensional transformation mixing up all the dimensions. A 5-dimensional representation of 3D rotation is furnished, for example, by the quantum-mechanical wave function of an electron in the *d*-orbital of an atom. Rotating the atom about its nucleus changes the five coefficients describing the composition of the *d*-wave in terms of its five basis wave functions. The diagram illustrates the fact that the same abstract group can have multiple matrix representations with a k- and an r-dimensional representation.

It is important not to confuse the dimension n of the group (= the number of parameters) with the dimension k of its representation (= the number of basis vectors in the representation space). These are two entirely different things! For the defining representation of 3D rotations, they both happen to be three: three angles ($n = 3$) control all possible rotations in 3D space ($k = 3$). But that's just a coincidence! In 2D, we need only one angle to specify all possible 2D rotations ($n = 1$, $k = 2$) and in 4D, we need six angles to specify all possible 4D rotations ($n = 6$, $k = 4$). Thus, the Lie group of ordinary 3D-rotations, called SO(3), can be rather confusing when chosen as a first example. Later, when we are turning to specific groups, we will start with another group, namely SU(2), for which the dimension of the group and the dimension of the defining representation are different ($n = 3$, $k = 2$).

Are the representations of a Lie group linear? Yes and No. The action of the matrix on the vectors in the representation space is certainly *linear*: $Mz_1 + Mz_2 = M(z_1 + z_2)$. But the dependence of the matrix on the parameters is usually *nonlinear*: $M(\theta_1, \cdots, \theta_n) + M(\theta_1', \cdots, \theta_n') \neq M(\theta_1 + \theta_1', \cdots, \theta_n + \theta_n')$. To symbolize this nonlinear aspect, we use boxes with rounded corners for the Lie group and its representations in the diagram. Finally, we should mention that it is possible to generalize the idea of a representation to allow for a *nonlinear* action on the representation space, which is then known as a *realization.* But with many of the interesting transformations in physics being linear, especially in quantum mechanics, we focus here exclusively on representations.

2.3 Lie Algebras

To get from a Lie group to the corresponding Lie algebra, we *linearize* the Lie group at the identity element e and provide a multiplication-like operation known as the *Lie bracket*, $[v, w]$, which measures the degree to which the composition of two group elements (near e) fails to commute.

Linearization

If we picture the Lie group as a sphere (keeping in mind the caveat pointed out earlier), then the Lie algebra corresponds to the tangent plane touching the sphere at the identity element e. More generally, the Lie algebra corresponds to the n-dimensional tangent space touching the n-dimensional Lie-group manifold at the identity element. This tangent space is an n-dimensional linear space (= vector space), that is, all its elements (= vectors or points) can be expressed as linear combinations of n basis vectors. To symbolize this linearity, we use *rectangular* boxes for the Lie algebra and its representations in the diagram. Note that the dimension of the Lie algebra, n, must match that of the Lie-group manifold. For example, the tangent space to a 2-dimensional sphere is a 2-dimensional plane.

What do the elements of the Lie algebra represent? Group elements close to the identity e represent *small transformations* and because the origin of the Lie-algebra space touches the group manifold at e, Lie-algebra elements close to the origin can also be understood as small transformations. More accurately, a small Lie-algebra element represents the *displacement* resulting from a small transformation: the difference between the point in the representation space before and after the small transformation. Each displacement vector depends on the point in the representation space that the transformation acts on. Therefore, these displacement vectors form a *vector field* on the representation space. In summary, we can think of a small Lie-algebra element as

- a point in the tangent space to the group manifold near the origin,
- a small transformation (= a group element near the identity),
- or a vector field (of small displacements) on the representation space.

We will explain this again in more concrete terms when discussing matrix representations.

Normally when we linearize something, we lose a lot of information. For example, if we know only the first derivative of the function $f(x)$ at $x = 1$, there is no way we can reconstruct the entire function. Interestingly, this is not the case for Lie groups! Knowing the Lie algebra, we can reconstruct (almost) everything there is to know about the Lie group. Why? Because the group structure lets us build up (= compose) large transformations by combining many small ones. If we know how to rotate a 3D object by one degree about all three axes, we can rotate the object in any way we want (in steps of one degree). For this reason, the elements of a Lie algebra are also known as *generators*: they can be used to generate large transformations.

So, how do we get from the Lie-algebra elements (= generators) back to the Lie-group elements (= large transformations)? We need to perform some sort of integration. This results in the so-called *exponential map*. It is much easier to explain this map for a concrete matrix representation than for an abstract algebra and thus we defer this discussion until later.

The Lie algebra is a linear space and thus comes automatically equipped with the operations of addition and scalar multiplication. Given two Lie-algebra elements v and w, we can form the linear combination $\alpha v + \beta w$, where α and β are scalars. What do these operations represent? If v corresponds to the small transformation g (= element of the Lie group) and w corresponds to the small transformation h, then $v + w$ corresponds to the small transformation composed of g and h, that is, gh. The addition in the Lie algebra describes, to first order, the composition of two small transformations in the Lie group. (Whereas composition does not commute in general, it does commute to first order for small transformations and thus can be represented by addition.) Similarly, if v corresponds to the small transformation g and α is a whole number, then the product αv corresponds to the composition of α copies of the small transformation g, that is, g^{α}.

However, the linear structure of the tangent space is not enough to capture all of the information about the Lie group near the identity. For that, we need to add a multiplication-like operation, thus turning the vector space into an algebra. This new operation, the Lie bracket, measures the degree to which the composition of two Lie-group elements fails to commute.

Lie Bracket

How do we measure the degree to which the composition of two Lie-group elements g and h fails to commute? We compute $g^{-1}h^{-1}gh$, that is, we apply the two transformations sequentially and then undo them the "wrong way" around (note that $h^{-1}g^{-1}gh$ would always undo them perfectly: $h^{-1}g^{-1}gh = e$). If the two transformations commute, the result is the identity element e. For two *small* transformations (= group elements near the identity) that do *not* commute, the result is again the identity to *first order*, but there is a *second order* deviation. It is this second-order deviation that is measured by the Lie bracket $[v, w]$, where the Lie-algebra elements v and w correspond to the Lie-group elements g and h.

The Lie bracket of two elements, $[v, w]$, is again an element of the Lie algebra, that is, the Lie algebra is closed under the bracket. This follows from the fact that the corresponding transformation $g^{-1}h^{-1}gh$ is also again an element of the Lie group. The Lie bracket anticommutes, $[v, w] = -[w, v]$, reflecting the fact that the two small transformations $g^{-1}h^{-1}gh$ and $h^{-1}g^{-1}hg$ in the Lie group deviate from the identity in opposite ways. Finally, the Lie bracket is *not* associative $[u, [v, w]] \neq [[u, v], w]$, but this inequality can be turned into an equality by adding one more term:

$$\big[u, [v, w]\big] = \big[[u, v], w\big] + \big[v, [u, w]\big].$$

This formula is known as the *Jacobi identity* (easier to remember in its cyclic form: $[u, [v, w]] + [v, [w, u]] + [w, [u, v]] = 0$).

In summary, the Lie algebra has two operations: (i) linear combination, $\alpha v + \beta w$, which describes how two small transformations in the Lie group compose ($g^{\alpha}h^{\beta}$), and (ii) the Lie bracket, $[v, w]$, which describes how two small transformations in the Lie group fail to commute ($g^{-1}h^{-1}gh$). With these two operations, the Lie algebra captures almost everything there is to know about the Lie group. (Only some global features are left out, such as group components that are topologically separated from the identity component and the difference between a group and its covering group.)

2.4 Lie-Algebra Representations

The elements of a Lie algebra are abstract objects, just like those of the Lie group. To make them more concrete, we represent them by square matrices, which act linearly on the vectors in the representation space. In the diagram, the abstract algebra element v is mapped to the matrix V, similarly, x is mapped to X, and so on. The matrix X acts on the vector z in the representation space by matrix-vector multiplication: $z' = Xz$. We call a map from the abstract algebra elements to the concrete matrices a *representation* of the Lie algebra (sometimes the matrices themselves are also referred to as a representation of the algebra). For every k-dimensional representation of a Lie group, there is a corresponding k-dimensional representation of the Lie algebra. (Note: To avoid introducing too many symbols, we wrote in the diagram $z' = Xz$ for the algebra action and $z' = Mz$ for the group action, although the two z' are different.)

Linearization and Exponential Map

How do we find the Lie-algebra representation for a given Lie-group representation? We take the partial derivative of the group matrix $M(\theta_1, \cdots, \theta_n)$ with respect to each of the n parameters θ_i and evaluate the resulting matrices at $\theta_i = 0$, where we assumed that $\theta_i = 0$ parametrizes the identity matrix, $M(0, \cdots, 0) = I$:

$$T_i = \frac{\partial M(\theta_1, \cdots, \theta_n)}{\partial \theta_i}\bigg|_{\theta_1 = \cdots = \theta_n = 0}$$

The resulting n matrices, T_i, are the *basis generators* of the Lie algebra. A general element X of the Lie algebra, known as a *generator*, is a linear combination of these basis generators: $X = \vartheta_1 T_1 + \cdots + \vartheta_n T_n$. (Caution: the terms *generator* and *basis generator* are often not clearly distinguished in the literature.) The specific basis generators that we obtain from the above procedure depend on the parametrization of the group elements, but the linear space spanned by the basis generators is always the same. We may normalize the basis generators to suit a particular purpose (e.g., particular eigenvalues or structure constants). Note that a k-dimensional representation is given by $k \times k$ matrices and thus the linear space appears to have k^2 dimensions, but in truth, the Lie algebra is constrained to an n-dimensional subspace singled out by the n basis generators.

Given a small matrix εX in the Lie algebra, we can construct the corresponding small transformation $M = I + \varepsilon X$ in the Lie group, where I is the identity matrix and ε is a small number. To obtain the large transformation corresponding to the full matrix X we set $\varepsilon = 1/N$ and repeat the small transformation N times: $(I + \varepsilon X)^N$. In the limit of a large N, this leads to the transformation matrix

$$M = \lim_{N \to \infty} \left(I + \frac{X}{N}\right)^N,$$

which evaluates to the *matrix exponential*

$$M = e^X := I + X + \tfrac{1}{2}X^2 + \tfrac{1}{6}X^3 + \tfrac{1}{24}X^4 + \cdots.$$

Thus, for each matrix X in the Lie algebra, we can construct a large transformation e^X in the Lie group. We can generalize this result by considering all Lie-algebra elements in the linear subspace spanned by

X (elements in the "direction" of X). For each subspace $X\vartheta$ in the Lie algebra, we get a one-parameter set of transformation matrices $M(\vartheta) = e^{X\vartheta}$ in the Lie group, where ϑ is the parameter. This is the *exponential map* that takes us from the Lie algebra back to the Lie group. Let's check if this is consistent with what we said before: If we choose the parameter to be small, $\vartheta = \varepsilon$, and expand the exponential to first order, we find $M = e^{X\varepsilon} = I + \varepsilon X$, as expected. Moreover, taking the derivative of $e^{X\vartheta}$ with respect to ϑ and evaluating at $\vartheta = 0$ indeed yields the generator X.

Let's visualize the exponential map: picture the Lie group as a sphere (you know the caveat) with the south pole representing the identity element. Moreover, imagine a knob for adjusting the parameter ϑ. When the knob is set to zero, a light spot appears at the south pole of the sphere. This spot represents the identity transformation. When we turn the knob, the spot moves away from the south pole following a smooth arc. This trajectory represents the one-parameter set of transformations associated with the chosen generator (= Lie-algebra element). If we pick another (linearly independent) generator, we get another trajectory on the sphere and so on.

So far, we have recovered only a one-dimensional subset of the complete Lie group. How can we get the full n-parameter set of transformations? Two strategies come to mind: Linearly combine all n basis generators to obtain a general Lie-algebra element, $T_1\vartheta_1 + T_2\vartheta_2 \cdots + T_n\vartheta_n$, and then exponentiate it:

$$M(\vartheta_1, \vartheta_2, \cdots, \vartheta_n) = e^{T_1\vartheta_1 + T_2\vartheta_2 + \cdots + T_n\vartheta_n},$$

where T_i are the basis generators and ϑ_i are the coordinates of the general Lie-algebra element this basis. Alternatively, exponentiate each one of the n basis generators separately and multiply the resulting n one-parameter transformation matrices together (thus composing the transformations):

$$M(\vartheta_1', \vartheta_2', \cdots, \vartheta_n') = e^{T_1\vartheta_1'} \cdot e^{T_2\vartheta_2'} \cdot \ldots \cdot e^{T_n\vartheta_n'}.$$

At first glance, the two formulas may seem to give the same result, but in general they do *not*. They agree only if all the basis generators commute, $[T_i, T_j] = 0$, otherwise they are related by the Baker-Campbell-Hausdorff formula:

$$e^X e^Y = e^{X+Y+\frac{1}{2}[X,Y]+\frac{1}{12}[X,[X,Y]]-\frac{1}{12}[Y,[X,Y]]+\cdots}.$$

Moreover, we can combine the two strategies by exponentiating arbitrary (but complementary) subspaces of the algebra (e.g., $e^{T_1\vartheta_1''+T_2\vartheta_2''} \cdot \ldots \cdot e^{T_n\vartheta_n''}$). Finally, we can arbitrarily permute the exponential factors (e.g., $e^{T_2\vartheta_2'''} \cdot e^{T_1\vartheta_1'''} \cdot \ldots \cdot e^{T_n\vartheta_n'''}$). Thus, given a set of basis generators, there are many different ways to construct the parametrized transformation matrix. The difference between the resulting matrices is their parametrization (hence the distinction between ϑ, ϑ', etc.).

The elements of a Lie group can be parametrized in many different ways. We can think of this as putting different (curved) coordinate systems on the group manifold. If we start with a Lie-group representation parametrized in some way, derive the Lie algebra from it (as discussed earlier), and then recover the Lie group from some set of basis generators (as discussed above), we may not end up with the parametrization that we originally started with. For this reason, we distinguished between the original parametrization $M(\theta_1, \cdots, \theta_n)$ and the parametrization $M(\vartheta_1, \cdots, \vartheta_n)$ after exponentiation. It may even happen that after exponentiation, we end up with a representation of a slightly different Lie group, namely the *covering Lie-group*. We will see examples of this later.

Vector Fields and Integral Curves

It is instructive to look at the Lie-algebra elements from a somewhat different viewpoint. We know that the generator εX corresponds to the small transformation $I + \varepsilon X$. Acting with this transformation on the vector z in the representation space and writing the result as the original vector plus a small differential vector yields $z' = z + dz = (I + \varepsilon X)z$. After subtracting z from both sides, we obtain $dz = \varepsilon X z$. Thus, a particular generator εX maps a vector z to a small *displacement* dz located at z. In other words, each εX defines a *vector field* on the representation space. To visualize this field, imagine a little arrow at every point of the representation space. Each arrow indicates where the point at its root moves to under the small transformation. For example, if the transformation is a small 3D rotation, the vector field consists of little arrows swirling cylindrically around the axis of rotation; arrows closer to the axis are shorter, arrows farther form the axis are longer. For each axis orientation, that is, for each generator εX, there is a different vector field.

Trajectories through the representation space that "follow" the vectors of the field are known as *flow lines* or *integral curves.* We can write such a trajectory as a vector that depends on a parameter: $z(\vartheta)$. To make sure that the vectors of the field, $\varepsilon X z$, are tangent to this trajectory, we require

$$\frac{dz(\vartheta)}{d\vartheta} = Xz(\vartheta).$$

Given a starting point $z(0)$, the solution of this first-order differential equation is

$$z(\vartheta) = e^{X\vartheta} z(0).$$

Note that $M(\vartheta) = e^{X\vartheta}$, which acts on the starting point, is just the one-parameter set of transformation matrices that we discussed earlier! We can thus think of the exponential map as taking in a vector field and assembling it into integral curves.

What happens if we act with the transformation $e^{X\vartheta}$ (= Lie-group element) on the vector field associated with X (= Lie-algebra element)? Every arrow of the field gets moved from location z to $e^{X\vartheta}z$ and gets reoriented from Xz to $e^{X\vartheta}(Xz)$. What arrow was at the location $e^{X\vartheta}z$ *before* the transformation? The old arrow there was $X(e^{X\vartheta}z)$. But this is the *same* arrow as the one that is there now, namely $e^{X\vartheta}(Xz)$, because $e^{X\vartheta} = I + X\vartheta + \frac{1}{2}X^2\vartheta^2 + \cdots$ and X commute! So, every arrow in the field moves to a new location and changes its orientation, but the field *as a whole* stays the same: it remains *invariant*. For example, imagine rotating the vector field of rotation described above: it remains invariant. In summary, the vector field associated with X stays invariant under the transformations generated by X.

Eigenvectors and Eigenvalues of the Generators

Given a generator X, what can we say about the associated vector field? Much of the relevant information is contained in the *eigenvectors* and *eigenvalues* of the matrix X!

The eigenvectors specify the directions in which the flow of the vector field comes out of or goes into the origin (depending on the sign of the corresponding eigenvalue). The eigenvalues specify the magnitude of the flow in the direction of the corresponding eigenvectors. For example, the 2D generator $X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ has eigenvectors $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 1 \\ -1 \end{pmatrix}$ with eigenvalues 1 and -1, respectively, describing a vector field that flows out of the origin along the $y = x$ diagonal and into the origin along the $y = -x$ diagonal. Integrating this vector field yields the transformation $e^{X\vartheta} = \begin{pmatrix} \cosh\vartheta & \sinh\vartheta \\ \sinh\vartheta & \cosh\vartheta \end{pmatrix}$. For an illustration of these flows in the 3D case, see [RtR, Fig. 13.10]. If the sum of all the eigenvalues (= trace of X) is zero, then the in- and outgoing flows add up to zero, that is, there are no sources or sinks in the vector field: it is divergence free.

If the eigenvectors and/or eigenvalues are complex, the flow of the vector field does not come out of or go into the origin, instead, it goes *around* the origin. For example, the generator $X = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ has the complex eigenvectors $\begin{pmatrix} i \\ 1 \end{pmatrix}$ and $\begin{pmatrix} i \\ -1 \end{pmatrix}$ with eigenvalues i and $-i$, respectively, describing a circular flow around the origin. Integrating this vector field yields the transformation $e^{X\vartheta} = \begin{pmatrix} \cos\vartheta & -\sin\vartheta \\ \sin\vartheta & \cos\vartheta \end{pmatrix}$.

Lie Bracket, Commutator, and Structure Constants

What does the abstract Lie-bracket operation correspond to in a matrix representation? To find out, we determine the amount by which two small transformations, M and N, fail to commute, that is, we check by how much the matrix product $M^{-1}N^{-1}MN$ deviates from the identity. First, we express the two small transformations in terms of their generator matrices, X and Y, up to second order:

$$M = I + \varepsilon X + \tfrac{1}{2}\varepsilon^2 X^2, \quad N = I + \varepsilon Y + \tfrac{1}{2}\varepsilon^2 Y^2.$$

Then, we plug this into $M^{-1}N^{-1}MN$, expand, simplify, and ignore terms beyond second order, which leads (after some work) to

$$M^{-1}N^{-1}MN = I + \varepsilon^2(XY - YX) + \cdots$$

It is the second-order deviation from perfect commutation that defines the Lie bracket. Thus, the Lie bracket is represented by the matrix commutator: $[X, Y] = XY - YX$.

From the matrix expression $[X, Y] = XY - YX$, it is evident that the commutator (just like the abstract Lie bracket) anti-commutes: $[X, Y] = -[Y, X]$. Closure implies that the commutator of any two elements is again an element of the Lie algebra. Therefore, we can write the commutator of any two basis generators as a linear combination of the full set of basis generators:

$$[T_i, T_j] = \sum_{k=1}^{n} c_{ijk} T_k.$$

The coefficients c_{ijk} are known as the *structure constants* of the Lie algebra. Given the structure constants in a particular basis, we have everything we need to calculate any commutator in this basis. Lie algebras are considered to be isomorphic if they have the same structure constants in some basis. (For a cool visualization of structure constants, see http://visuallietheory.blogspot.com/2012/09/.)

Of course, there is a lot more to say about Lie groups, Lie algebras, and their representations, but this is all we need to get started. Now, let's look at some concrete examples!