A Visual Invitation to Cartan Geometries

Jacob W. Erickson

Email address: jwericks@umd.edu

To Symmetry,
for everything.

Contents

Preface

We exist within a geometric universe, full of geometric information. In order to live within this universe, we must have some way of processing geometric information from the world around us. So, how do we do it? How do we *experience* geometry?

As it turns out, pondering this question leads very naturally to the idea of Cartan geometries.

Unfortunately, Cartan geometries have a rather unfair reputation for being intimidating and geometrically unintuitive. It is easy to see where this reputation comes from: at first glance, Cartan geometries look like cognitively impenetrable contrivances, requiring a weird Lie algebra-valued one-form on a principal bundle to satisfy a bunch of technical-looking, arcane conditions. Over the course of this book, we intend to convince the reader that, contrary to this first impression, Cartan geometries actually encode something *profoundly* intuitive.

Our singular goal for this book is to give meaningful geometric intuition for the various major ideas in the modern study of Cartan geometries. We hope that, in doing so, we can make the topic a bit more accessible to newcomers and deepen the understanding of those who already have some knowledge of the subject.

Of course, goals, hopes, and dreams are nice, but we need to be concrete about this. Reading a math book, even a small one like this, is a significant time commitment, and we need to be clear about what we are doing so that we do not waste the reader's valuable time. So, why read this book in particular? What are we doing differently?

The first notable difference is our approach to principal bundles. To be blunt, doing differential geometry on principal bundles is really intimidating to beginners. Even *I* remember thinking of principal bundles as a "necessary evil" when I first started learning differential geometry. The reasons for why we *want* to work on principal bundles are often poorly presented, leaving students to simply accept that we do things a particular way. This must change. In this book, we will make considerable effort to clarify why principal bundles should be our friends: they are the ideal way of placing ourselves inside of the geometry.

The second major departure from other works on the topic is that we actually know how to explain the intuition behind the machinery. At time of writing, I do not think most Cartan geometers realize the true depth of intuition behind Cartan geometries. This is not to say that they do not utilize this intuition, in the same way that one does not need to know the function of the eyelid in order to blink. Indeed, this is, in some sense, the primary issue with trying to explain the subject: the machinery is conveying something so deeply instinctive to how we experience geometry that it is difficult to put into words. Once one fully "gets" the idea, though, this instinct makes the machinery very easy to use, even if one does not know how to explain why it is easy.

Somewhat orthogonal to our main goal for the book, our third significant distinction from previous literature is how fully we commit to the analogy between Cartan geometries and Lie groups. This is mostly just a practice we want to make standard in future works. Several papers on the topic fumble and fuddle around with complicated additional machinery that obfuscates the points of their constructions; often, such machinery can be vastly simplified or discarded entirely by using the concept of development to translate the ideas into Lie-theoretic ones.

Of course, the final—and very prominent—difference is our focus on pictures and intuition. We are not kidding around when we include the word "visual" in the title: a solid chunk of the book is filled with illustrations of geometric ideas.

To be clear, this will not be a particularly comprehensive introduction, since we do not think that is what the subject needs right now. Between the valiant ambition of Sharpe's flawed masterpiece [**5**] and the fastidious ministrations of Čap and Slovák's encyclopedic juggernaut [**2**], the existing literature already covers most of the fundamental technical details for Cartan geometries quite thoroughly. Instead, this will be more of an *invitation* to the subject: we will show the reader how to experience these ideas for themselves, with the hope that this will prepare and encourage them to read more challenging works as well.

That being said, we will be assuming a nontrivial amount of background. Specifically, we will need the reader to have at least some proficiency with basic differential geometry. The preliminaries for working on manifolds are covered so well and in so many places that rehashing them here would be a disservice to them. For those seeking references to such material, we can confidently recommend the first volume [**6**] of Spivak's well-known introduction to the subject, as well as [other works that are also good].

As a corollary of this, we will require the reader to know the basics of Lie theory. Most introductory courses to the topic should be sufficient. We will, for example, be writing under the assumption that the reader already knows what the orthogonal group is and understands what semidirect products are. However, we have written Chapter 1 with the understanding that, at time of writing, most introductory courses on Lie theory do not actually teach how to visualize Lie groups and their basic operations, so such gaps will be filled by the text.

[Acknowledgements] [Be sure to mention that this is based on the lecture notes we made for "Parabolic Geometries for People that Like Pictures"]

Jacob W. Erickson

Progress notes

Because I want to reference some of the techniques I use for building intuition with Cartan geometries in my PhD dissertation, I'm putting this online much earlier than I'd like; there's a *lot* of editing and writing left to do before this even resembles the form I want it to eventually take. Still, I think the hundred pages I have included so far, and the pictures drawn within them, will be quite useful to those seeking to understand the intuition behind the subject.

Aside from general editing, here is a quick run-down of what still needs to be done:

Chapter 1
- Finish the section on the adjoint representation and the bracket

Chapter 2
- Completely redo this chapter. Sloppy!

Chapter 3
- Redo the sections where we go over the examples
- Especially focus on expanding hyperbolic geometry, since we need that for doing representation theory later
- Write the section on comparing model geometries (extending Euclidean geometry to affine geometry, etc.)

Chapter 4
- This chapter is based on the lecture I usually give when I try to explain Cartan geometries to newcomers, where the point is to trick the audience into realizing that they actually already know how to use Cartan geometries. I'm not convinced it translates well into words on a page, so I might try fiddling with the wording. In particular, I'm not yet happy with how the beginning of the chapter connects to the rest of it.
- Write (and also draw the pictures for) the section on Klein geometries and curvature

Chapter 5
- Finish writing it, obviously.
- A lot of the ideas will come from my holonomy paper, if you want the ideas but not the pictures

Chapter 6
- Write the chapter
- Consider the placement of the chapter and its contents; the point of it is to give the reader a chance to practice with some of the ideas in a more concrete setting, so it should come after introducing enough of the elementary tools but before doing anything too intense with the structure theory.
- In particular, we want this to come *before* the chapter on non-compact Riemannian symmetric spaces.

Chapter 7
- Write the chapter
- Consider tone?

Chapter 8
- Rewrite the lecture so that it works as a chapter of the book
- Find a way to explain that "it's just a trace form" is hand-waving and not intuition, but without sounding like a jerk to people that use this handwaving as if it were intuition.

Chapter 9
- Rewrite the lecture so that it works as a chapter of the book
- Expand upon the maximal compact subgroup stuff?
- Consider whether to move the introduction to parabolic subgroups to the next chapter

Chapter 10
- Rewrite the lecture so that it works as a chapter of the book
- Consider whether restricting this chapter to just being about structure theory would be better

Chapter 11
- Rewrite the lecture so that it works as a chapter of the book
- Add in the stuff we figured out about the line bundle and how to explain unipotent tilting formally?
- Consider whether talking about curvature here is a good idea? Maybe if the chapter on Riemannian geometry does a good enough job with explaining torsion?

Chapter 12
- Rewrite the lecture so that it works as a chapter of the book
- The pictures *really* need to be redrawn...
- Consider the curved case?

Chapter 13
- Rewrite the lecture so that it works as a chapter of the book
- Consider the curved case?
- Consider moving the description of Tanaka prolongation to the next chapter

Chapter 14
- Write the chapter

Appendix A
- Write the appendix

Chapter 1

Visualizing the Fundamental Tools of Lie Theory

In essence, the idea behind Cartan geometries is to make a principal bundle over a manifold resemble a particular Lie group. Much of their immense potency comes from allowing us to "do Lie theory" on geometric objects where there is no inherent geometric symmetry. Of course, this should suggest to us that, if we want to talk about geometric intuition for these things, then we really need to *start* with Lie groups.

Therefore, in this first chapter, we will answer some basic questions regarding visualization of Lie groups, including:

- What does left-translation look like?
- What does right-translation look like?
- What does conjugation look like?
- What is the Maurer-Cartan form, and why is it so amazing?

We will also introduce—and hopefully justify—how we will use the terms *transformation* and *motion*, as well as the somewhat technical distinction we make between them. To do all of this, we will focus primarily on the main example we will use throughout the beginning of the book: the Lie group of Euclidean isometries.

By the end of this chapter, the reader should start to have an intuitive grasp of what it is like to move around inside of a Lie group, and in the next chapter, we will practice using this to do geometry.

1.1. Picturing the group of Euclidean isometries

To start, we give a way of placing ourselves inside of Euclidean geometry: orthonormal frames.

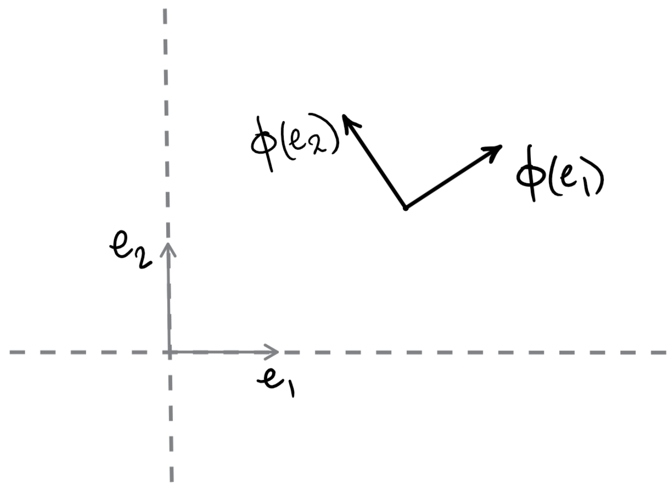

Figure 1. We can depict an orthonormal frame ϕ on $\mathbb{R}^2$ by the pair of tangent vectors $(\phi(e_1), \phi(e_2))$

Consider the plane $\mathbb{R}^2$ equipped with the usual Euclidean structure. An *orthonormal frame* over $u \in \mathbb{R}^2$ is just a linear isometry ϕ from the tangent space $T_0\mathbb{R}^2$ at 0, which we will identify with $\mathbb{R}^2$ itself, to the tangent space $T_u\mathbb{R}^2$ at u. Fixing a pair e_1 and e_2 of orthonormal vectors in $\mathbb{R}^2 \approx T_0\mathbb{R}^2$, we can uniquely determine an orthonormal frame ϕ by its values on e_1 and e_2, since linear maps are uniquely determined by their values on a basis. In particular, we can pictorially depict an orthonormal frame ϕ on the plane as the pair of tangent vectors $(\phi(e_1), \phi(e_2))$, as we have in Figure 1.

While using pairs of vectors is useful for drawing simple pictures, there is a different way of visualizing orthonormal frames that will be useful in far more general settings. Imagine we are walking around on $\mathbb{R}^2$. Look directly in front of us; along this tangent direction, there is a unique unit vector that corresponds to moving "forward" with unit speed. Similarly, to our left, perpendicular to the forward direction, there is a unique unit vector corresponding to leftward motion with unit speed. With this information, we can identify our *configuration* on the plane with the unique orthonormal frame ϕ such that $\phi(e_1)$ is the unit forward vector and $\phi(e_2)$ is the unit leftward vector.[1] In other words, orthonormal frames allow us to place ourselves inside of Euclidean geometry, as illustrated in Figure 2.

[1] Of course, the choice to use forward and left is arbitrary, and we could just as easily have chosen something else as long as we remained consistent.

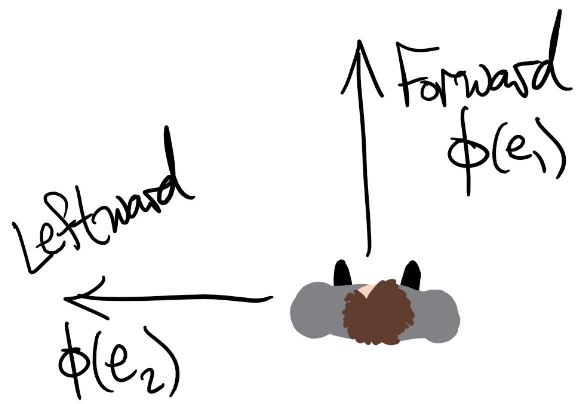

Figure 2. Each orthonormal frame corresponds to a unique configuration for ourselves as pedestrians on the Euclidean plane

Now, let us consider the Lie group $\mathrm{I}(2)$ of Euclidean isometries of $\mathbb{R}^2$ under composition.

To each $u \in \mathbb{R}^2$, there is a unique isometry $\tau_u \in \mathrm{I}(2)$ given by $v \mapsto u+v$, called *translation*[2] by u. In particular, each isometry $\phi \in \mathrm{I}(2)$ uniquely decomposes as a composition

$$\phi = \tau_{\phi(0)} \circ \big(\tau_{\phi(0)}^{-1} \circ \phi\big),$$

where $\tau_{\phi(0)}$ is a translation and $\tau_{\phi(0)}^{-1} \circ \phi = \tau_{-\phi(0)} \circ \phi$ is an isometry that fixes 0. Since isometries preserve lines in the plane, an isometry that fixes 0 must be linear, hence the subgroup of isometries that fix 0 is precisely the orthogonal group $\mathrm{O}(2)$ of linear isometries of $\mathbb{R}^2$. In other words, every $\phi \in \mathrm{I}(2)$ can be written uniquely as a composition $\tau_u \circ A$ for some $u \in \mathbb{R}^2$ and $A \in \mathrm{O}(2)$.

Given two isometries $\tau_u \circ A$ and $\tau_v \circ B$, we can compute their composition: for $x \in \mathbb{R}^2$,

$$\begin{aligned}(\tau_u \circ A) \circ (\tau_v \circ B)(x) &= (\tau_u \circ A)(v + B(x)) = u + A(v + B(x)) \\ &= (u + A(v)) + AB(x) = (\tau_{u+A(v)} \circ AB)(x),\end{aligned}$$

so $(\tau_u \circ A) \circ (\tau_v \circ B) = \tau_{u+A(v)} \circ AB$. In particular, we may consider the Lie group $\mathrm{I}(2)$ as the semidirect product $\mathbb{R}^2 \rtimes \mathrm{O}(2)$, with group operation given by

$$(u, A)(v, B) = \big(u + A(v), AB\big).$$

[2]These correspond to both left-translations and right-translations on $\mathbb{R}^2$, viewed as a Lie group.

Elements of $\mathrm{O}(2)$ are, by definition, linear isometries from $\mathbb{R}^2 \approx T_0\mathbb{R}^2$ to itself, so $\mathrm{O}(2)$ can be viewed as the space of orthonormal frames over 0. By adding in translations, this perspective then allows us to identify $\mathrm{I}(2)$ with the space of *all* orthonormal frames over $\mathbb{R}^2$, which we would usually call the *orthonormal frame bundle* over $\mathbb{R}^2$. Specifically, for each isometry $\phi \in \mathrm{I}(2)$, the pushforward $\phi_* : \mathbb{R}^2 \approx T_0\mathbb{R}^2 \to T_{\phi(0)}\mathbb{R}^2$ gives a linear isometry from the tangent space at 0 to the tangent space at $\phi(0)$, hence ϕ_* is an orthonormal frame at $\phi(0)$.

In review, we identify each element $\phi \in \mathrm{I}(2)$ of the isometry group $\mathrm{I}(2)$ with the orthonormal frame ϕ_* determined by its pushforward at 0. Every $\phi \in \mathrm{I}(2)$ uniquely decomposes as a composition of the form $\tau_u \circ A$ for some $u \in \mathbb{R}^2$ and $A \in \mathrm{O}(2)$. Pictorially, we can depict the orthonormal frame $\tau_{u*} \circ A : \mathbb{R}^2 \approx T_0\mathbb{R}^2 \to T_u\mathbb{R}^2$ corresponding to $\tau_u \circ A$ as the pair of tangent vectors $(\tau_{u*}(A(e_1)), \tau_{u*}(A(e_2)))$ at $\phi(0) = u$. More importantly, however, we can identify $\tau_u \circ A$ with the configuration of ourselves on the plane such that $\tau_{u*}(A(e_1))$ is the unit forward vector and $\tau_{u*}(A(e_2))$ is the unit leftward vector.

1.2. Transformation and Motion

Throughout, we will use $\mathrm{L}_a : g \mapsto ag$ to denote left-translation by a and $\mathrm{R}_a : g \mapsto ga$ to denote right-translation by a. What do left-translation and right-translation look like? In particular, how are they different? Since we have some intuition for what elements of $\mathrm{I}(2)$ look like, it gives a good place to investigate these questions.

To start, let us look at how left-translation by τ_{e_1} behaves. For an arbitrary element $\tau_u \circ A \in \mathrm{I}(2)$, we have

$$\mathrm{L}_{\tau_{e_1}}(\tau_u \circ A) = \tau_{e_1} \circ (\tau_u \circ A) = (\tau_{e_1} \circ \tau_u) \circ A = \tau_{e_1+u} \circ A.$$

In other words, it behaves basically the same as it does as a transformation of $\mathbb{R}^2$, shifting every orthonormal frame uniformly by e_1. This has been illustrated in Figure 3, and the corresponding situation for rotations has similarly been illustrated in Figure 4.

To see where this behavior comes from, note that $\mathrm{I}(2)$ acts transitively on $\mathbb{R}^2$, so since $\mathrm{O}(2)$ is the stabilizer of 0 in $\mathrm{I}(2)$, we can think of $\mathbb{R}^2$ as the homogeneous space $\mathrm{I}(2)/\mathrm{O}(2)$. In particular, we have a natural quotient map

$$q_{\mathrm{O}(2)} : \mathrm{I}(2) \to \mathbb{R}^2 \cong \mathrm{I}(2)/\mathrm{O}(2)$$

given by $\phi \mapsto \phi(0)$, or equivalently, by $\tau_u \circ A \mapsto u$.

This map $q_{\mathrm{O}(2)}$ lets us think of the space $\mathrm{I}(2)$ of orthonormal frames of $\mathbb{R}^2$ as a bundle over $\mathbb{R}^2$. In terms of orthonormal frames, $q_{\mathrm{O}(2)}$ just takes

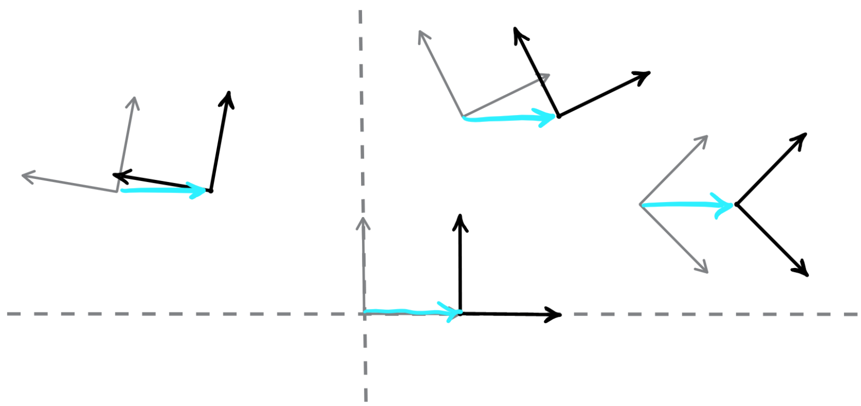

Figure 3. Left-translating by τ_{e_1} shifts all orthonormal frames uniformly by e_1, as if we were applying it as a transformation to the plane and the orthonormal frames were thought of as being inside of the plane

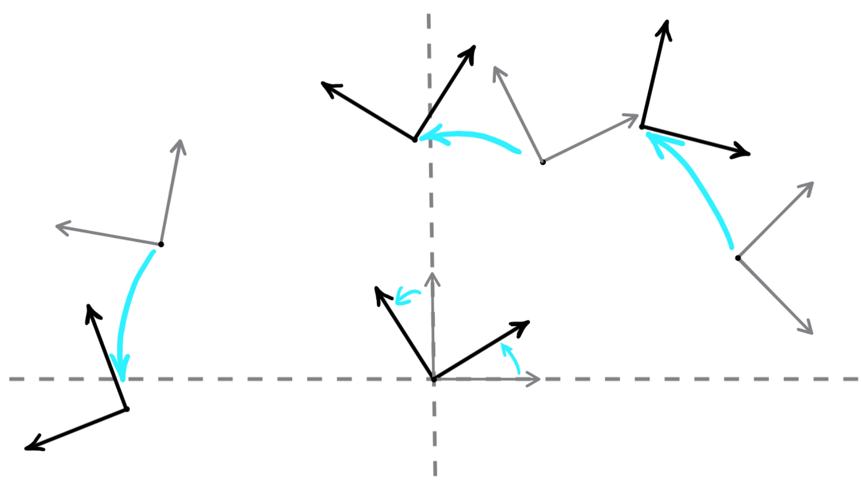

Figure 4. Left-translating by the rotation given by the linear isometry $\mathrm{rot}(\theta) := \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} \in \mathrm{O}(2)$ uniformly rotates all orthonormal frames by θ around 0

orthonormal frames over $u \in \mathbb{R}^2$ and maps them all to u. Equivalently, thinking as a pedestrian on the Euclidean plane, $q_{\mathrm{O}(2)}$ takes our precise configuration on the Euclidean plane and maps it to the point of $\mathbb{R}^2$ at which we are positioned.

For $\phi, \psi \in \mathrm{I}(2)$, we have

$$q_{\mathrm{O}(2)}(\mathrm{L}_\phi(\psi)) = q_{\mathrm{O}(2)}(\phi \circ \psi) = \phi \circ \psi(0) = \phi(q_{\mathrm{O}(2)}(\psi)),$$

so under the quotient map $q_{\mathrm{O}(2)}$, left-translation by ϕ in $\mathrm{I}(2)$ corresponds to just applying ϕ as a *transformation*.

What does this mean for right-translation? For $\phi, \psi \in \mathrm{I}(2)$, we have

$$q_{\mathrm{O}(2)}(\mathrm{R}_\phi(\psi)) = q_{\mathrm{O}(2)}(\psi \circ \phi) = \psi(\phi(0)).$$

The key here is to notice that, because $\mathrm{I}(2)$ acts on $\mathbb{R}^2 \cong \mathrm{I}(2)/\mathrm{O}(2)$ from the left, ϕ gets to act *before* ψ when we apply the transformation $\psi \circ \phi = \mathrm{R}_\phi(\psi)$. This means that right-translation by ϕ moves each orthonormal frame as if ϕ is acting on the orthonormal frame at the identity.

In an attempt to clarify what this means, let us see what right-translation by τ_{e_1} does. For an arbitrary $\tau_u \circ A \in \mathrm{I}(2)$, we have

$$\mathrm{R}_{\tau_{e_1}}(\tau_u \circ A) = (\tau_u \circ A) \circ \tau_{e_1} = \tau_{u+A(e_1)} \circ A.$$

Under the appropriate identifications between vectors and translations, this amounts to right-translation by τ_{e_1} shifting each orthonormal frame ϕ_* by the vector $\phi_*(e_1)$ to which that orthonormal frame maps e_1. We have illustrated this in Figure 5.

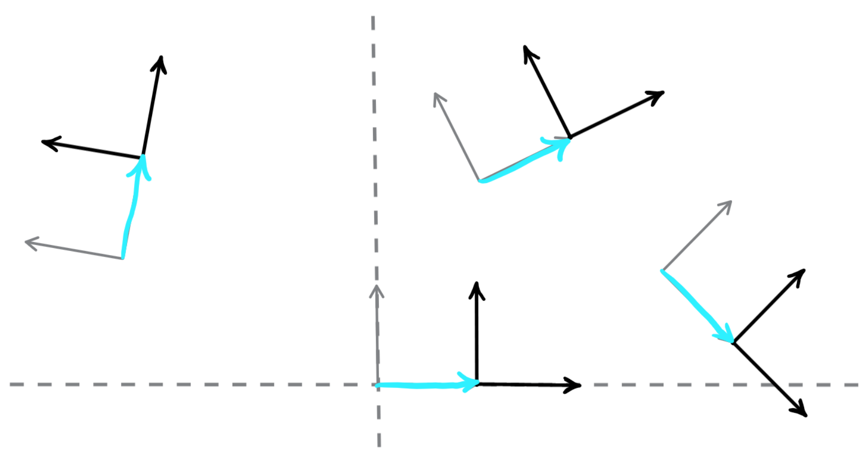

Figure 5. Right-translating by τ_{e_1} shifts each orthonormal frame by the vector to which that orthonormal frame maps e_1

Under our identification between orthonormal frames and configurations for ourselves as pedestrians on the Euclidean plane, this gives right-translation by τ_{e_1} a very simple description: **it corresponds to walking forward by one unit**. We have shown this in Figure 6.

More generally, we can think of right-translation by an arbitrary $\phi \in \mathrm{I}(2)$ as applying ϕ from the perspective of our configuration. For example, we see in Figure 7 that right-translation by the rotation $\mathrm{rot}(\theta) := \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}$ just turns us on the spot by θ, rather than necessarily rotating us around the origin at 0.

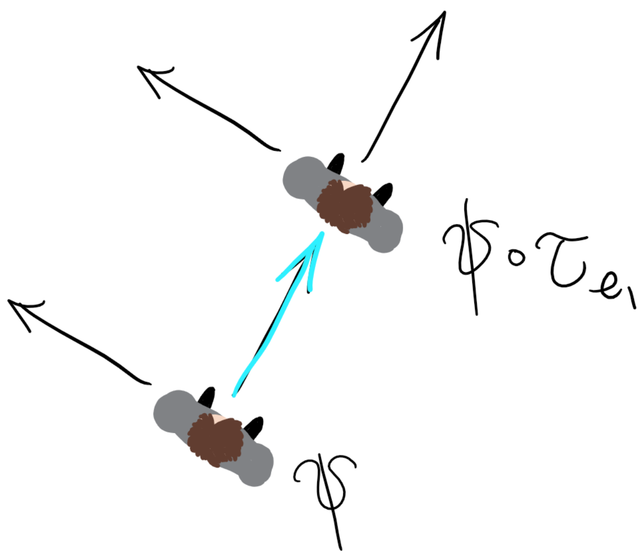

Figure 6. Right-translating by τ_{e_1} corresponds to walking forward one unit

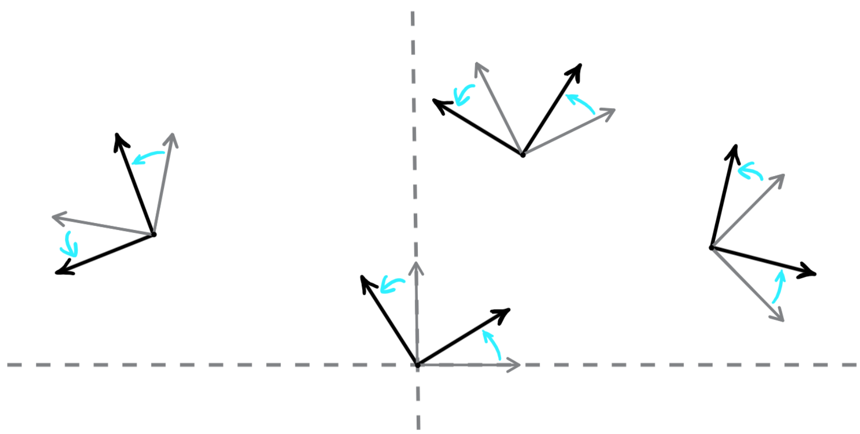

Figure 7. Right-translating by $\mathrm{rot}(\theta)$ turns each orthonormal frame on the spot by θ

Let us attempt to summarize this intuition in just a few words. Under left-translation, elements act as *transformations*, effecting each orthonormal frame uniformly according to the usual action on the homogeneous space $\mathrm{I}(2)/\,\mathrm{O}(2)$. Under right-translation, on the other hand, elements act as *motions*, moving orthonormal frames according to their own perspectives. More evocatively, left-translating by a rotation is like rotating the whole Earth (with us on it) and right-translating by a rotation is like turning around.

1.3. Conjugation

Now that we have some idea of what left-translation and right-translation look like, we are naturally led to ask: what does conjugation look like? There are several ways to approach this question, and as we will throughout these notes, we encourage the reader to wander off the path we are following and explore when they feel motivated to do so. However, the author has found one interpretation in particular that is consistently useful and easy to see.

For $g, h \in \mathrm{I}(2)$, observe that

$$\mathrm{R}_h(g) = gh = (ghg^{-1})g = \mathrm{L}_{ghg^{-1}}(g).$$

On the left-hand side, we have g right-translated by h, which we can interpret as moving by h from the perspective of g. On the right-hand side, we have g left-translated by ghg^{-1}, which we can interpret as just applying the transformation ghg^{-1} to g. In other words, ghg^{-1} is the element we can apply to g as a transformation to reproduce the motion given by h.

Let us give some examples. For $u \in \mathbb{R}^2$, consider the translation τ_u and the rotation $\mathrm{rot}(\theta)$ of angle θ around 0. If we right-translate τ_u by $\mathrm{rot}(\theta)$, then this corresponds to turning on the spot by θ around u. Thus, the conjugate $\tau_u \mathrm{rot}(\theta)\tau_u^{-1}$ is the *transformation* that does this to τ_u, namely rotation of the whole plane by θ around the point u. We have illustrated this in Figure 8.

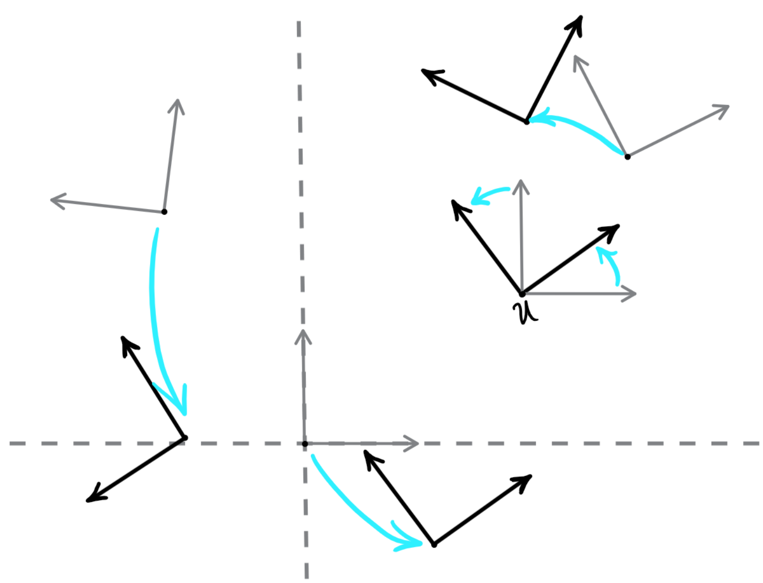

Figure 8. $\tau_u \mathrm{rot}(\theta)\tau_u^{-1}$ rotates the plane by θ around u

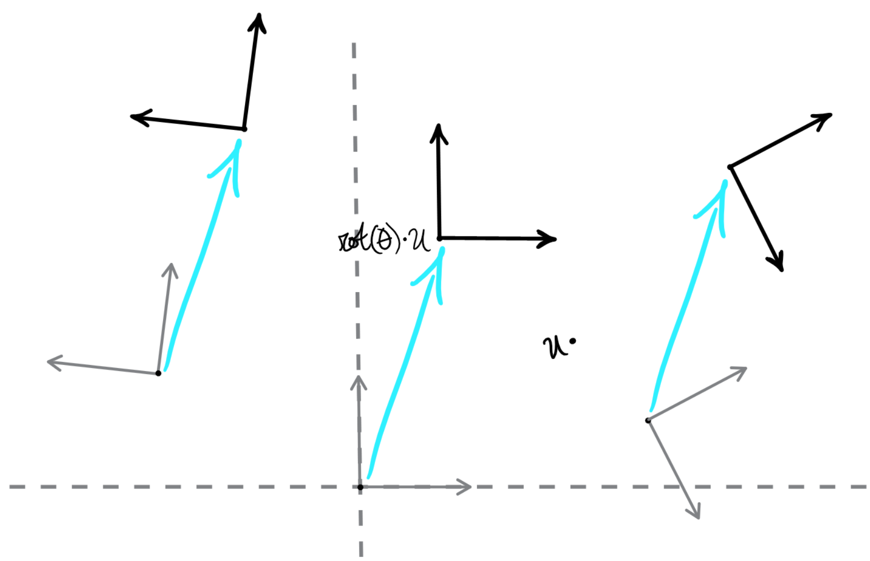

Figure 9. $\mathrm{rot}(\theta)\tau_u\mathrm{rot}(\theta)^{-1}$ shifts the plane by the vector to which the orthonormal frame corresponding to $\mathrm{rot}(\theta)$ sends u

Similarly, if we right-translate $\mathrm{rot}(\theta)$ by τ_u, then this corresponds to moving by the vector u according to the perspective of the orthonormal frame at $\mathrm{rot}(\theta)$. Therefore, the conjugate $\mathrm{rot}(\theta)\tau_u\mathrm{rot}(\theta)$ is the *transformation* that does this to $\mathrm{rot}(\theta)$, namely translation by $\mathrm{rot}(\theta) \cdot u$.

Using the terminology of transformation and motion, a pithy way of summarizing this interpretation is that conjugation by g converts motions to transformations that look like those motions at g.

1.4. The Maurer-Cartan Form

Before talking about the Maurer-Cartan form, let us first examine the structure of the Lie algebra $\mathfrak{i}(2)$ of $\mathrm{I}(2)$. Using the decomposition of $\mathrm{I}(2)$ as the semidirect product $\mathbb{R}^2 \rtimes \mathrm{O}(2)$, we can decompose $\mathfrak{i}(2)$ as the semidirect sum $\mathbb{R}^2 \ni \mathfrak{o}(2)$.

This decomposition has a fairly nice interpretation: it tells us that every element of $\mathfrak{i}(2)$ can be written as the sum of a (translational) velocity and an angular velocity. Characterizing elements of the Lie algebra as generators for one-parameter subgroups, we can think of the Lie algebra as the space of *instantaneous motions* through the identity. The velocities in $\mathbb{R}^2 < \mathfrak{i}(2)$ determine the obvious one-parameter subgroups: for each $v \in \mathbb{R}^2 \approx T_0\mathbb{R}^2$, $\exp(tv) = \tau_{tv}$. Similarly, the angular velocities $t\left[\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right] \in \mathfrak{o}(2) < \mathfrak{i}(2)$ map to $\mathrm{rot}(t)$ under the exponential map.

By definition, the Lie algebra $\mathfrak{i}(2)$ of I(2) is the tangent space of I(2) at the identity element (together with a bracket operation that we shall talk about later). To describe tangent spaces at other points, we use something called the *Maurer-Cartan form.*

Definition 1.1. The *Maurer-Cartan form* of a Lie group G is the $\mathfrak{g}$-valued one-form ω_G given by

$$(\omega_G)_g : T_g G \to \mathfrak{g} = T_e G, \ X_g \mapsto \mathrm{L}_{g^{-1}*} X_g$$

at each $g \in G$.

While its definition appears to be mere algebraic formalism, do not be fooled: the Maurer-Cartan form is one of the most deeply intuitive objects in modern differential geometry. In fact, it is so intuitive that many students will frequently use it without consciously recognizing that they are doing so. In vector calculus courses, for example, vector fields on $\mathbb{R}^n$ are often defined to be maps from $\mathbb{R}^n$ to itself; implicitly, this definition is using the Maurer-Cartan form on $\mathbb{R}^n$ to identify all of its tangent spaces with the Lie algebra $\mathbb{R}^n$. This should be profoundly exciting to anyone interested in geometric intuition, because if we already have an instinctive understanding of how to use some geometric objects—so instinctive that we often do not realize we are using them—then it stands to reason that thinking in terms of these objects might lead to similarly effortless understanding of corollary phenomena.

But what does the Maurer-Cartan form actually tell us? Let us try to elucidate its meaning in the group of Euclidean isometries. Imagine we are a pedestrian moving around on the Euclidean plane, with our configuration at time t given by $\gamma(t)$ for some smooth curve $\gamma : \mathbb{R} \to \mathrm{I}(2)$, as in Figure 10.

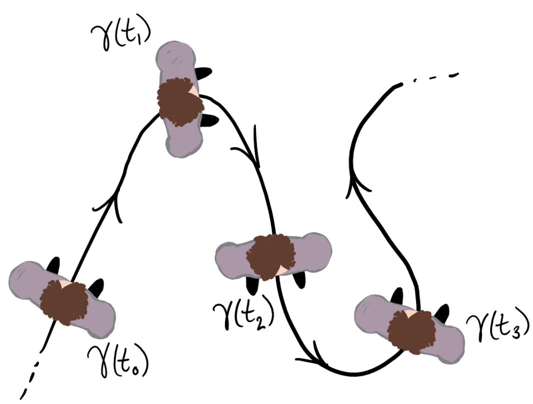

Figure 10. The pictured pedestrian moves according to the smooth curve γ in I(2)

For each time t, we have a tangent vector $\dot{\gamma}(t) \in T_{\gamma(t)}\,\mathrm{I}(2)$ describing our current velocity at $\gamma(t)$ within the space of configurations. These tangent vectors tell us how we are moving at a given instant, but they mostly lie in different tangent spaces. How do we describe these velocities in a consistent way? The elegant answer is to do what we usually do: simply describe our motion in terms of our own perspective.

Consider how we would describe our trajectory between times t_0 and t_3 if we were walking along γ. At t_0, we are walking *forward* along the curve. When we get to t_1, while we are still walking *forward*, we have also started to turn to our right (*clockwise*). By the time we reach t_2, we have stopped turning to the right and have started to—ever so slightly—turn to our left (*counterclockwise*), though we are still walking *forward* along the curve. Finally, when we arrive at t_3, we are still walking *forward*, and we have started turning left (*counterclockwise*) considerably more than before. This kind of description is exactly what the Maurer-Cartan form tells us. For example, $\omega_{\mathrm{I}(2)}(\dot{\gamma}(t_0))$ will simply be a positive scalar multiple of e_1, since $e_1 \in \mathbb{R}^2 < \mathfrak{i}(2)$ is the velocity through the identity corresponding to moving *forward* with unit speed. Meanwhile, we will have

$$\omega_{\mathrm{I}(2)}(\dot{\gamma}(t_1)) = (a_1 e_1, -b_1 [\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}])$$

for some $a_1, b_1 > 0$, since we are walking *forward* while turning *clockwise*. Likewise, for t_2 and t_3, we will have

$$\omega_{\mathrm{I}(2)}(\dot{\gamma}(t_2)) = (a_2 e_1, b_2 [\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}]) \text{ and } \omega_{\mathrm{I}(2)}(\dot{\gamma}(t_3)) = (a_3 e_1, b_3 [\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}]),$$

where $a_2, a_3, b_2, b_3 > 0$ and $\frac{b_3}{a_3} > \frac{b_2}{a_2}$, since we are walking *forward* at both times but are turning *counterclockwise* considerably more at time t_3 than at time t_2.

To summarize, the Maurer-Cartan form is a canonical coframing of the Lie group that identifies velocities in each tangent space of the Lie group with corresponding instantaneous motions in the Lie algebra. Since the Lie algebra, the space of all instantaneous motions through the identity, does not change, this gives us a constant frame of reference for describing velocities at different configurations, essentially by describing those velocities from the perspectives of those configurations.

1.5. The adjoint representation and the bracket

[Conjugation]

[In terms of Maurer-Cartan form]

Chapter 2

Euclidean Geometry

Much of this chapter is sloppy and should be rewritten. Giving the reader a strong understanding of how the Erlangen Program works for Euclidean geometry in this perspective is crucial!

To start, let us briefly summarize what we have learned so far about the Lie group I(2) of Euclidean isometries of the plane.

In addition to being the symmetry group of Euclidean geometry, we can think of I(2) as the space of configurations for ourselves as pedestrians on the Euclidean plane. Equivalently, we can think of it as the orthonormal frame bundle over $\mathbb{R}^2$, with each $\phi \in \mathrm{I}(2)$ identified with the orthonormal frame

$$\phi_* : T_0\mathbb{R}^2 \approx \mathbb{R}^2 \to T_{\phi(0)}\mathbb{R}^2$$

given by the pushforward of ϕ at 0 and with bundle map given by the natural quotient map

$$q_{\mathrm{O}(2)} : \mathrm{I}(2) \to \mathbb{R}^2 \cong \mathrm{I}(2)/\,\mathrm{O}(2),\ \ \phi \mapsto \phi(0).$$

This configuration perspective gives us an easy way of placing ourselves inside of Euclidean geometry.

The Lie-theoretic structure of I(2) also provides us with a natural way of describing how to move around within the geometry: the Maurer-Cartan form $\omega_{\mathrm{I}(2)}$ gives a canonical coframing for I(2) that identifies velocities with instantaneous motions. From the perspective of a pedestrian walking around on the Euclidean plane, this allows us to speak, for example, about "walking forward" without having to specify our current configuration.

Of course, our description of I(2) was in terms of Euclidean geometry. Now that we have learned how to move inside of a Lie group, we would like to try going the other way: we will reformulate Euclidean geometry in terms of $\mathrm{I}(2)$, or rather, in terms of the pair $(\mathrm{I}(2), \mathrm{O}(2))$.

By the end of this chapter, we should be able to talk about how the geometric structure of the Euclidean plane comes from the pair $(\mathrm{I}(2), \mathrm{O}(2))$. This will open us up to exploring, in the next chapter, how other geometries might be determined by Lie groups in a similar vein.

2.1. "Isn't geometry about circles and lines and stuff?"

As we mentioned at the end of the first lecture, when we "walk in a straight line", this really just corresponds to moving with constant translational velocity—which is to say, translational velocity that is constant with respect to the Maurer-Cartan form. Thus, a line will just be the full path of such motion projected onto the plane.

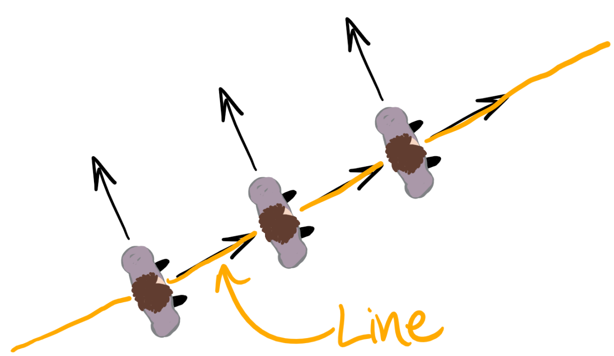

Figure 1. A line is the projection to the plane of a curve with $\omega_{\mathrm{I}(2)}$-constant translational velocity and zero angular velocity

Definition 2.1. A *line* on the Euclidean plane $\mathrm{I}(2)/\,\mathrm{O}(2) \cong \mathbb{R}^2$ is a subset of the form $q_{\mathrm{O}(2)}(g \exp(\mathbb{R}v))$ for some $g \in \mathrm{I}(2)$ and some nonzero $v \in \mathbb{R}^2 < \mathfrak{i}(2)$.

Choosing to use *right*-translations to define lines might seem odd to the uninitiated. Indeed, if you are not already familiar with Cartan geometries, then it probably seems easier to define lines as orbits of one-parameter subgroups of translations acting from the *left* on the Euclidean plane. Unfortunately, in this case, the reason for using right-translations is somewhat obscured by the fact that the subgroup of translations is normal[1] in $\mathrm{I}(2)$. By the end of the next lecture, the reason for this choice will be obvious,

[1] Recall that $N \leq G$ is *normal* in G if and only if $gN = Ng$ for all $g \in G$. In other words, when N is normal, every motion R_n with $n \in N$ has a corresponding transformation $\mathrm{L}_{n'}$ for some $n' \in N$ that behaves the same way at a given $g \in G$.

but for now, we will just say that we always want to be able to *move* along lines (and, later, geodesics and other distinguished curves).

Note that, for every $x \in \mathbb{R}^2$, we have $\tau_x \in q_{\mathrm{O}(2)}^{-1}(x)$. Thus, if $x + \mathbb{R}v$ is a line in the usual sense on $\mathbb{R}^2$, then we can write it in terms of the definition above as $q_{\mathrm{O}(2)}(\tau_{x+\mathbb{R}v}) = q_{\mathrm{O}(2)}(\tau_x \exp(\mathbb{R}v))$.

While lines themselves are defined in terms of motion along them, we will define parallelism in terms of transformations. We start with two lines, and in order to see whether they are parallel, we shift one of those lines onto the other via a translation.

Definition 2.2. Two lines ℓ and ℓ' in the Euclidean plane are *parallel* if and only if there is some $u \in \mathbb{R}^2$ such that $\tau_u(\ell) = \ell'$.

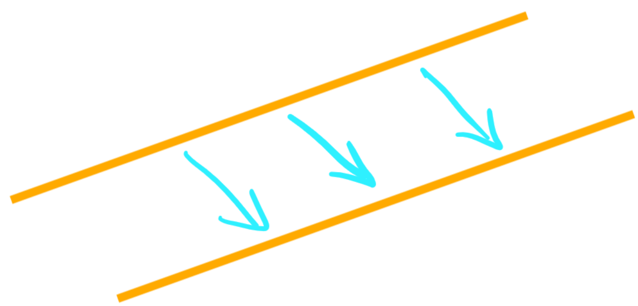

Figure 2. Two lines in the Euclidean plane are parallel when one is a translation of the other

Later, we will show that this definition is equivalent to the classical definition of parallel lines in terms of intersection.

Given two vectors in the same tangent space, we can measure the angle between them in terms of the rotation needed to move from one to the other.

Definition 2.3. Let $g \in \mathrm{I}(2)$ be such that $q_{\mathrm{O}(2)*}((\omega_{\mathrm{I}(2)})_g^{-1}(e_1))$ is a positive scalar multiple of $v \in T_x\mathbb{R}^2$. For $w \in T_x\mathbb{R}^2$, the *(oriented) angle* from v to w is the unique $\theta \in \mathbb{R}/2\pi\mathbb{Z}$ such that $q_{\mathrm{O}(2)*}((\omega_{\mathrm{I}(2)})_{g\mathrm{rot}(\theta)}^{-1}(e_1))$ is a positive scalar multiple of w.

Small correction: In order to uniquely determine the direction of the angle here, we pick both a starting vector and, implicitly, an *orientation*. Thus, these should really be called *oriented angles*, since they are not preserved by orientation

Note that this defines angles modulo 2π. If we want to talk about angles larger than 2π, or if we want to distinguish angles that are the same modulo 2π, then we would need to work with the universal cover $\mathbb{R}$ of $\mathrm{SO}(2) \simeq \mathbb{R}/2\pi\mathbb{Z}$.

An interesting feature of the geometry of the Euclidean plane is that, fixing an orientation by restricting ourselves to the identity component

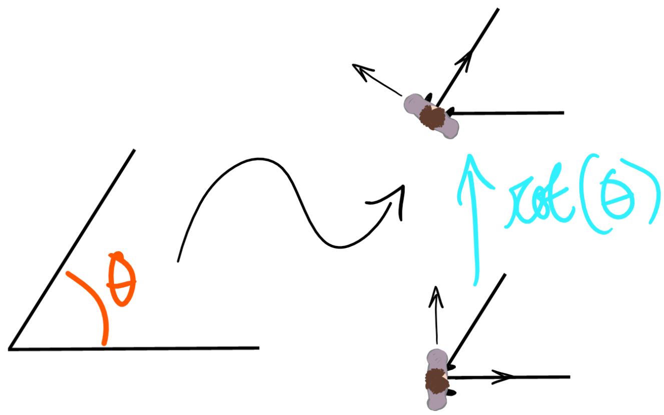

Figure 3. We can define angles between vectors in terms of the rotations needed to move between them

$I^\circ(2) \simeq \mathbb{R}^2 \rtimes SO(2)$ of $I(2)$, we have a natural way to add angles at different points. Because $\mathbb{R}^2$ is normal in $I^\circ(2)$, we get a natural homomorphic quotient map

$$\pi_{\mathbb{R}^2} : I^\circ(2) \to \mathbb{R}^2 \backslash I^\circ(2) \simeq SO(2)$$

given by $\tau_u \circ A \mapsto A$. Under the map $\pi_{\mathbb{R}^2}$, we can identify angles at different points on the plane and add them together.

An alternative way to describe angles comes from the adjoint representation $\mathrm{Ad} : I(2) \to GL(\mathfrak{i}(2))$. Conveniently, we already have some idea of what conjugation looks like, so since Ad_g is, by definition, just the pushforward at the identity of conjugation by g, we can get a fairly good picture of what the adjoint representation looks like as well.

For example, conjugating g by $\mathrm{rot}(\theta)$ gives the transformation that behaves like the motion g at the orthonormal frame corresponding to $\mathrm{rot}(\theta)$. Thus, $\mathrm{rot}(\theta) \circ \tau_{tv} \circ \mathrm{rot}(\theta)^{-1}$ is just the transformation coinciding with the motion τ_{tv} at $\mathrm{rot}(\theta)$, namely $\tau_{t\mathrm{rot}(\theta)\cdot v}$, so $\mathrm{Ad}_{\mathrm{rot}(\theta)}$ just rotates velocities by θ. The angle between two vectors $v, w \in T_x\mathbb{R}^2$ can then equivalently be described as the element $\theta \in \mathbb{R}/2\pi\mathbb{Z} \simeq SO(2)$ such that, for v a positive scalar multiple of $q_{O(2)*}((\omega_{I(2)})_g^{-1}(e_1))$, w is a positive scalar multiple of $q_{O(2)*}((\omega_{I(2)})_g^{-1}(\mathrm{Ad}_{\mathrm{rot}(\theta)}(e_1)))$.

Note that, in the first definition of angle above, we had v be a positive scalar multiple of $q_{O(2)}((\omega_{I(2)})_g^{-1}(e_1))$ and w be a positive scalar multiple of $q_{O(2)*}((\omega_{I(2)})_{g\mathrm{rot}(\theta)}^{-1}(e_1))$. Examining this second expression more closely, we have

$$\begin{aligned} q_{\mathrm{O}(2)*}((\omega_{\mathrm{I}(2)})^{-1}_{g\mathrm{rot}(\theta)}(e_1)) &= q_{\mathrm{O}(2)*}(\mathrm{L}_{g\mathrm{rot}(\theta)*}(e_1)) = q_{\mathrm{O}(2)*}(\mathrm{L}_{g*}\,\mathrm{L}_{\mathrm{rot}(\theta)*}(e_1)) \\ &= q_{\mathrm{O}(2)*}(\mathrm{R}_{\mathrm{rot}(\theta)^{-1}*}(\mathrm{L}_{g*}\,\mathrm{L}_{\mathrm{rot}(\theta)*}(e_1))) \\ &= q_{\mathrm{O}(2)*}(\mathrm{L}_{g*}(\mathrm{L}_{\mathrm{rot}(\theta)*}\,\mathrm{R}_{\mathrm{rot}(\theta)^{-1}*})(e_1)) \\ &= q_{\mathrm{O}(2)*}(\mathrm{L}_{g*}\,\mathrm{Ad}_{\mathrm{rot}(\theta)}(e_1)) \\ &= q_{\mathrm{O}(2)*}((\omega_{\mathrm{I}(2)})^{-1}_{g}(\mathrm{Ad}_{\mathrm{rot}(\theta)}(e_1))), \end{aligned}$$

where the equality in the second line follows from $q_{\mathrm{O}(2)} \circ \mathrm{R}_{\mathrm{rot}(\theta)} = q_{\mathrm{O}(2)}$ and the equality in the third line is a consequence of left-translation and right-translation commuting with each other. This verifies that the two definitions of angle are equivalent.

Finally, we get to circles. We won't use them much in the Euclidean geometry planned for this lecture, but it's still worth giving them a definition in terms of I(2), to prove that we can.

Definition 2.4. For $x \in \mathbb{R}^2$ and $v \in \mathbb{R}^2$, the *circle* centered at x with radius (the length of) v is the set

$$C_v(x) := \left\{ q_{\mathrm{O}(2)}(g\tau_v) : g \in q^{-1}_{\mathrm{O}(2)}(x) \right\}.$$

In other words, $C_v(x)$ is the set of all points that some orthonormal frame over x thinks are v away from x. Equivalently, if you stand over x and specify a radius v from your frame, and then you spin around "in a circle" until you get back to your original configuration, then you will have traced out a circle.

2.2. Two elementary results from Euclidean geometry

To demonstrate how "actual" Euclidean geometry can be done in terms of isometries, we shall prove two elementary results.

Proposition 2.5. *Suppose two lines ℓ and ℓ' intersect at a point x, determining four angles around x as in Figure 5. The angles opposite each other are congruent, so that $\theta_1 = \theta_3$ and $\theta_2 = \theta_4$.*

Using I(2), this is fairly straightforward: imagine you are occupying a frame g over x such that you are pointed along a vector tangent to ℓ used to form the angle θ_1. By definition, if we rotate ourselves by θ_1, which is to say we right-translate by $\mathrm{rot}(\theta_1)$, then we will be pointing along ℓ'. Now, imagine we are at $g \circ \mathrm{rot}(\pi) = g \circ (-\mathbb{1})$. We are still pointed along ℓ, but now in the opposite direction, along a vector we can use to form the angle θ_3. But because SO(2) is abelian, if we move by $\mathrm{rot}(\theta_1)$ from $g \circ (-\mathbb{1})$, then we'll be at the orthonormal frame $g \circ (-\mathbb{1}) \circ \mathrm{rot}(\theta_1) = (g \circ \mathrm{rot}(\theta_1)) \circ (-\mathbb{1})$,

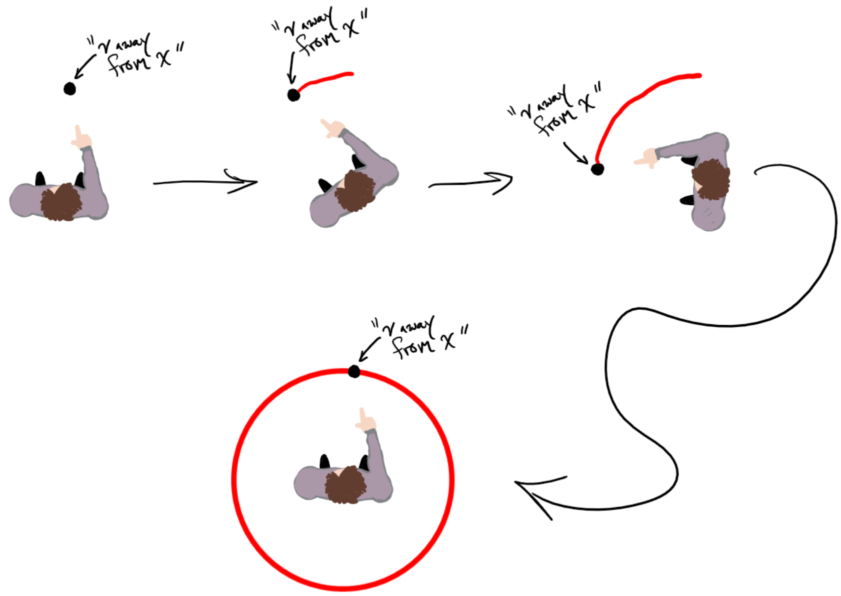

Figure 4. Tracing out a circle as in the above definition

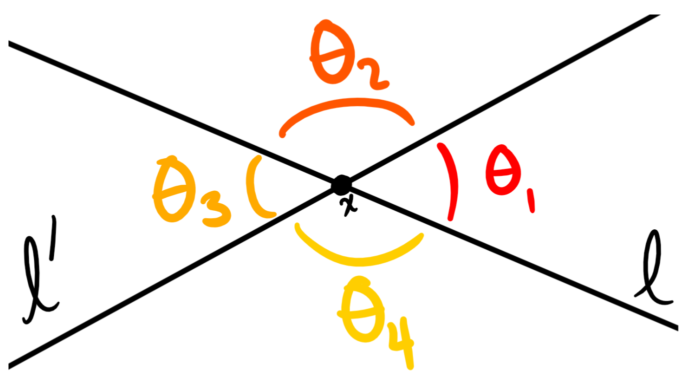

Figure 5. Two lines ℓ and ℓ' intersecting at x

which points us along ℓ' again in the opposite direction as $g \circ \mathrm{rot}(\theta_1)$. In other words, rotating by θ_1 did the same thing as rotating by θ_3, so they are equal.

Equivalently, we could just say that θ_1 and θ_3 are congruent under the isometry $\tau_x \circ (-\mathbb{1}) \circ \tau_x^{-1}$, essentially by the same reason: $-\mathbb{1}$ sends each line through 0 to itself.

Let us try another.

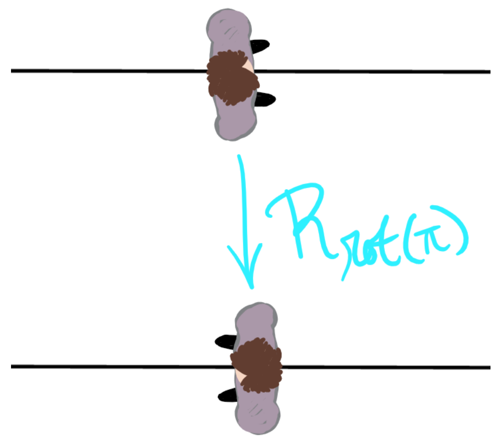

Figure 6. Turning on the spot by $\mathrm{rot}(\pi) = -\mathbb{1}$ keeps you on the same line but points you in the opposite direction

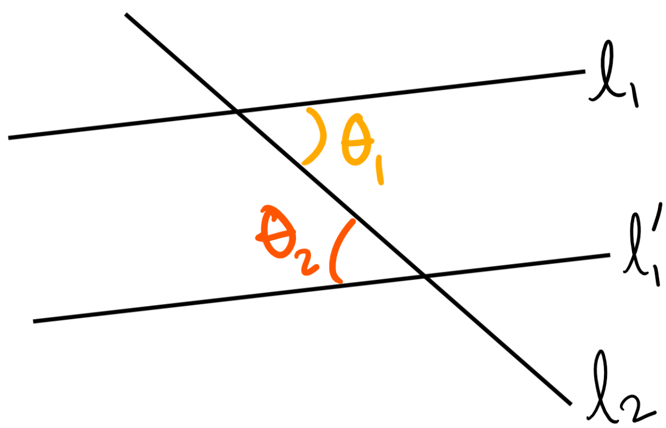

Figure 7. Parallel lines ℓ_1 and ℓ_1' intersected by a transversal ℓ_2

Proposition 2.6. *Suppose ℓ_1 and ℓ_1' are distinct parallel lines, and ℓ_2 is a line intersecting both ℓ_1 and ℓ_1', forming the angles θ_1 and θ_2 as in Figure 7. Then, θ_1 and θ_2 are congruent.*

Again, this is not too difficult in terms of I(2): because ℓ_1 and ℓ_1' are parallel, there is some $u \in \mathbb{R}^2$ such that $\tau_u(\ell_1) = \ell_1'$, and because the subgroup of translations is normal, there is some nonzero $v \in \mathbb{R}^2$ such that $\tau_v(\ell_1') = \ell_1'$. For x the point where ℓ_1 and ℓ_2 intersect and y the point where ℓ_1' and ℓ_2 intersect, there is some $t \in \mathbb{R}$ such that

$$\tau_u(x) = u + x = tv + y = \tau_{tv}(y),$$

hence $y - x = u - tv$, so $\tau_{y-x} = \tau_{-tv} \circ \tau_u$. Since τ_u sends ℓ_1 to ℓ_1' and τ_{-tv} preserves ℓ_1', this means τ_{y-x} sends ℓ_1 to ℓ_1'. Moreover, since $x, y \in \ell_2$, τ_{y-x} preserves ℓ_2, so τ_{y-x} sends θ_1 to the angle opposite θ_2 in the intersection of ℓ_1' and ℓ_2, hence they are congruent by the first proposition.

2.3. How does Euclidean geometry show up, algebraically?

Now that we have seen how to reformulate Euclidean geometry in terms of $\mathrm{I}(2)$, we are led to a natural question: why $\mathrm{I}(2)$? What about this particular Lie group gives us Euclidean geometry?

Throughout, we have relied heavily on the subgroup $\mathbb{R}^2 < \mathrm{I}(2)$ of translations. In particular, as we saw in Proposition 2.6, we explicitly used the fact that $\mathbb{R}^2$ is normal in $\mathrm{I}(2)$ to use transformations instead of motions. Implicitly, we also used the fact that $\mathbb{R}^2 < \mathrm{I}(2)$ acts simply transitively on $\mathrm{I}(2)/\mathrm{O}(2) \cong \mathbb{R}^2$, when we used the points x and y to determine the transformation τ_{y-x}. Translations were also used to describe parallelism, giving a definition that, as we now show, happens to coincide with the standard formulation in terms of intersections.

Proposition 2.7. *Two distinct lines ℓ and ℓ' in the Euclidean plane are parallel if and only if $\ell \cap \ell' = \emptyset$.*

Proof. To start, choose $x, y, v, w \in \mathbb{R}^2$ such that $\ell = q_{\mathrm{O}(2)}(\tau_{x+\mathbb{R}v})$ and $\ell' = q_{\mathrm{O}(2)}(\tau_{y+\mathbb{R}w})$. If $\ell \cap \ell' = \emptyset$, then $x \neq y$ and there are no $t, s \in \mathbb{R}$ such that

$$q_{\mathrm{O}(2)}(\tau_{x+tv}) = x + tv = y + sw = q_{\mathrm{O}(2)}(\tau_{y+sw}).$$

In particular, $y - x$ is never in the span of v and w, so because $\mathbb{R}^2$ is 2-dimensional, this means that v and w are scalar multiples of each other and

$$\tau_{y-x}(\ell) = \tau_{y-x}(q_{\mathrm{O}(2)}(\tau_{x+\mathbb{R}v})) = q_{\mathrm{O}(2)}(\tau_{y+\mathbb{R}v}) = q_{\mathrm{O}(2)}(\tau_{y+\mathbb{R}w}) = \ell'.$$

Conversely, if there exists $u \in \mathbb{R}^2$ such that $\tau_u(\ell) = \ell'$, then

$$\tau_u(q_{\mathrm{O}(2)}(\tau_{x+\mathbb{R}v})) = q_{\mathrm{O}(2)}(\tau_{u+x+\mathbb{R}v}) = q_{\mathrm{O}(2)}(\tau_{y+\mathbb{R}w}),$$

so for every $t \in \mathbb{R}$, there is an $s \in \mathbb{R}$ such that $u + x + tv = y + sw$. In particular, $u+x-y+tv$ is in the span of w for every $t \in \mathbb{R}$, so $u+x-y$ and $(u+x-y+v)-(u+x-y) = v$ are in the span of w. Thus, if $\ell \cap \ell' \neq \emptyset$, then there would be $t, s \in \mathbb{R}$ such that $x + tw = u + x + sw$, which would mean that u is in the span of w, which would mean ℓ and ℓ' were not distinct. ☐

Note that the proof of this equivalence explicitly depended on our ability to use $\mathbb{R}^2$ as a vector space. Moreover, it once again implicitly used the identification between $\mathbb{R}^2$ as the homogeneous space $\mathrm{I}(2)/\mathrm{O}(2)$ and $\mathbb{R}^2$ as

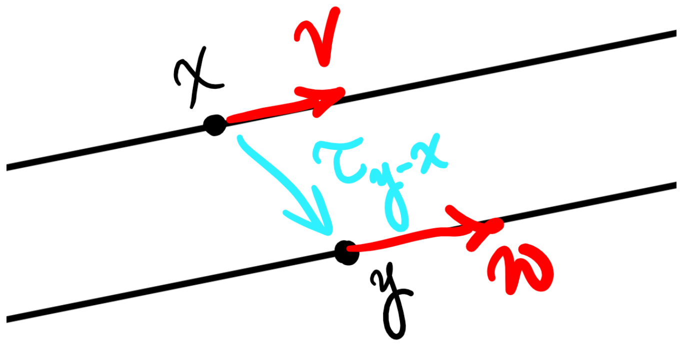

Figure 8. If $\ell \cap \ell' = \emptyset$, then for $x \in \ell$ and $y \in \ell'$, we have $\tau_{y-x}(\ell) = \ell'$

the subgroup of translations. For example, we took the points x and y of the homogeneous space and used them to create the translation τ_{y-x} from their difference, and we showed that v was in the span of w using $(u + x - y + v) - (u + x - y) = v$.

Indeed, the key feature that allows for parallelism to look the way that it does in Euclidean geometry is this simply transitive normal subgroup of translations. To see this, suppose we have another Lie group G acting transitively on the plane $\mathbb{R}^2$, and that G contains a closed normal subgroup isomorphic to $\mathbb{R}^2$ that acts simply transitively on $\mathbb{R}^2$. Again, we can decompose elements $g \in G$ as $g = \tau_{g(0)}(\tau_{g(0)}^{-1} g)$, with $\tau_{g(0)}^{-1} g$ acting linearly on $\mathbb{R}^2$ by conjugation (since conjugation must give an automorphism of $\mathbb{R}^2$ and the group of automorphisms of $\mathbb{R}^2$ is precisely $\mathrm{GL}_2 \mathbb{R}$).

We can then define lines the same way we did above, in terms of translations. Our definition of parallelism in terms of translations still makes sense as well, and by repeating the proof of the above proposition, we see that it is consistent with the usual definition in terms of intersection.

Circles and angles are a bit trickier to find in the structure of $\mathrm{I}(2)$. Of course, we can (correctly) guess that it ultimately comes from the subgroup $\mathrm{O}(2) < \mathrm{I}(2)$, but *how* is still a bit mysterious, since we used translations to describe both concepts. The key is to notice that we didn't actually need the translations for these definitions. For example, when defining circles, we used $v \in \mathbb{R}^2$ to describe the radius of $C_v(x)$ because that was more familiar, but really, every isometry $a \in \tau_v \,\mathrm{O}(2) = q_{\mathrm{O}(2)}^{-1}(v)$ determines the same circle:

$$C_v(x) = \left\{ q_{\mathrm{O}(2)}(g\tau_v) : g \in q_{\mathrm{O}(2)}^{-1}(x) \right\} = \left\{ q_{\mathrm{O}(2)}(ga) : g \in q_{\mathrm{O}(2)}^{-1}(x) \right\}.$$

Thus, the angles and circles of Euclidean geometry come from the stabilizer subgroup $\mathrm{O}(2)$ of $0 \in \mathbb{R}^2 \cong \mathrm{I}(2)/\,\mathrm{O}(2)$. In particular, if we were to replace

I(2) with a Lie group G containing a closed subgroup isomorphic to O(2), then we could use the definitions from above to talk about "circles" and "angles" in this other "geometry".

In the next lecture, we will clarify what we mean here by "geometry", and explore some famous examples.

Chapter 3

Geometry from Symmetry

Having seen that Euclidean geometry on the plane can be reformulated in terms of the pair $(\mathrm{I}(2), \mathrm{O}(2))$, we naturally want to see what else we can accomplish from this perspective. Delightfully, there turns out to be *quite a lot* we can do with this idea of describing geometry in terms of symmetry. This way of thinking, where symmetries are used to determine and categorize the structures of various geometries, is often referred to as the *Erlangen program*, which Felix Klein introduced toward the end of the nineteenth century.

In this chapter, we will show how to coax geometric structure from a pair (G, H), which we will refer to as a *model geometry*, or just *model* for short. Conveniently, much of the general structure can be seen from the example of Euclidean geometry, so our approach will be to show how each aspect of the general case extends from a particular aspect of the Euclidean case. From there, we will give a few examples of geometries that are similar enough to Euclidean geometry to be easily comparable, but different enough to highlight how the structure changes with different choices of symmetry.

By the end of this chapter, we should see that the intuition we developed for Euclidean geometry extends, with some small adjustments, fairly easily to these other geometries. In addition, we will see how, in some sense, this geometric structure comes from the Maurer-Cartan form, which will help to foreshadow some ideas in the next chapter.

3.1. Model geometries

To begin, we define the notion of geometry on which the rest of the book will be based.

Definition 3.1. A *model geometry*[1] (or simply *model*) is a pair (G, H), where G is a Lie group and $H \leq G$ is a closed subgroup such that G/H is connected. In a model (G, H), the Lie group G is called the *model group* and H is called the *isotropy* or *stabilizer subgroup*.

While the jump in abstraction might seem intimidating at first, there is really not much more going on here than there is in Euclidean geometry. For a given model (G, H), we are describing a geometric structure on the manifold G/H. We can, as we did before, think of the model group G as the bundle of configurations for ourselves as **pedestrians** wandering the geometry on G/H, with bundle map given by the natural quotient map

$$q_H : G \to G/H, g \mapsto gH.$$

The isotropy H, then, describes the space of configurations that can occur over a point of G/H; for each $g \in G$, we can reach every other configuration lying over $q_H(g)$ by right-translating by an element of H. In particular, because the isotropy subgroup H acts freely and transitively on each fiber of G over G/H by right-translation, G is a *principal H-bundle*[2] over G/H. We have organized this analogy into Table 1, and attempted to illustrate the idea in Figure 1.

Table 1. The following table clarifies how model geometries extend notions we described for Euclidean geometry on the plane

	Euclidean geometry $(\mathrm{I}(2), \mathrm{O}(2))$	Arbitrary model geometry (G, H)
Base manifold for geometry	$\mathbb{R}^2 \cong \mathrm{I}(2)/\mathrm{O}(2)$ (Euclidean plane)	G/H (homogeneous space)
Bundle of configurations	$\mathrm{I}(2)$ (orthonormal frame bundle)	G (principal H-bundle)
Quotient map for bundle	$q_{\mathrm{O}(2)} : \mathrm{I}(2) \to \mathbb{R}^2$, $\phi \mapsto \phi(0)$	$q_H : G \to G/H$, $g \mapsto gH$
Isotropy	$\mathrm{O}(2)$	H
Symmetry group	$\mathrm{I}(2)$ (as transformations)	G (as transformations)

[1] We sometimes also call these *Klein geometries*, as Sharpe does in [**5**], though in the context of this book, we will usually use this term to refer specifically to the geometric structure of the model geometry expressed in terms of a Cartan connection. This will be discussed more thoroughly in the next chapter when we introduce the Cartan machinery.

[2] Recall that a principal H-bundle $\mathscr{G}$ over M is a fiber bundle over M together with a right-action of H on $\mathscr{G}$ that is free and transitive on each fiber.

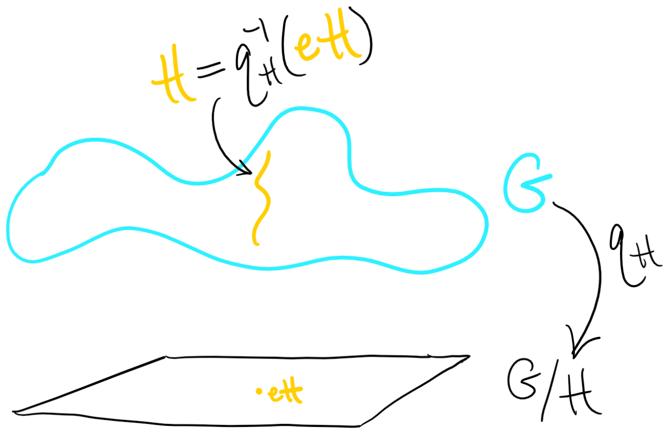

Figure 1. In a model (G, H), the model group G can be thought of as the bundle of pedestrian configurations over the space G/H, and the isotropy H runs through all the configurations lying over a given point in G/H

On top of giving us a way to place ourselves inside of the geometry, the model group G naturally acts on both itself and G/H from the left, and we take this action to *define* what symmetry means for the model geometry. In Euclidean geometry, for example, the model group is precisely the Lie group of transformations preserving the Euclidean structure: the isometry group. For a more general model geometry (G, H), the elements of G play the same role that isometries do in Euclidean geometry, acting as transformations that preserve the underlying geometric structure.

What *is* the geometric structure of (G, H)? This is the elegant idea behind Klein's Erlangen program: **the geometric structure is whatever is preserved by the symmetries of the geometry**!

"Definition" 3.2. We will say that something is *geometric* for the model (G, H) if and only if it is preserved by some action of G induced by the natural left-action of G on itself.

This definition might appear to be a bit nebulous—because it is—so let us give a few clarifying examples. In Euclidean geometry, distance is preserved by the group of Euclidean isometries, so distance is geometric. Likewise, the notions of point, line, and circle are geometric for Euclidean geometry because Euclidean isometries send points to points, lines to lines, and circles to circles. Note, however, that *specific* points, lines, and circles are not geometric, since they will not be preserved by the group of Euclidean isometries. Therefore, even though it is often a useful reference point while doing Euclidean geometry, the origin is not itself a geometric invariant of Euclidean geometry, since we cannot distinguish it from any other point;

essentially by definition, the orbit of I(2) through the origin is the whole Euclidean plane.

3.2. The definitive invariant of a model geometry

It probably seems somewhat surreal to define geometry as simply being "whatever is preserved by the symmetries". After all, we are often taught to think of geometric structures as being encoded by specific, concrete invariants, like Riemannian metrics or contact distributions. We would like the same thing here: a definitive geometric invariant that uniquely determines the geometric structure of a given model.

Of course, one thing that is *always* preserved by the model group G acting on itself from the left is its Maurer-Cartan form ω_G, since it is, by definition, left-invariant! In fact, this is precisely the type of explicit diffeo-geometric object we are looking for: the symmetries of the Maurer-Cartan form on G, where we think of G as a principal H-bundle over G/H, are precisely the left-translations by elements of G, so the Maurer-Cartan form can reconstruct the natural left-action of G on itself, which in turn defines the geometry.

In other words, the Maurer-Cartan form ω_G on the principal H-bundle G over G/H completely captures the geometric structure of (G, H). To convey this equivalence rigorously, we include the following proposition, which is more or less the same as Corollary 3.4.11 in [**5**].

Proposition 3.3. *Suppose (G, H) is a model. If $f : G \to G$ is a map such that $f^*\omega_G = \omega_G$ and $f(gh) = f(g)h$ for all $g \in G$ and $h \in H$, then there is some $a \in G$ such that $f = \mathrm{L}_a$.*

Proof. Denote by $\mu : G \times G \to G$ the group operation $(g, g') \mapsto gg'$ and by $(\cdot)^{-1} : G \to G$ the inverse operation $g \mapsto g^{-1}$.

Let $\sigma : G \to G$ be given by

$$g \mapsto f(g)g^{-1} = (\mu \circ (f, (\cdot)^{-1}))(g, g).$$

We want to show that σ is constant, because then

$$\sigma(e)g = \sigma(g)g = (f(g)g^{-1})g = f(g)$$

for all $g \in G$, so that $f = \mathrm{L}_{\sigma(e)}$. Since $f(gh) = f(g)h$,

$$\sigma(gh) = f(gh)h^{-1}g^{-1} = f(g)g^{-1} = \sigma(g)$$

for all $g \in G$ and $h \in G$, so σ is invariant under right-translation by H. In particular, since G/H is connected, σ is constant if and only if it is constant on a connected component of G, hence it suffices to prove that $\sigma^*\omega_G = 0$.

For $X \in T_gG$ and $Y \in T_{g'}G$,

$$\mu^*\omega_G(X,Y) = \mathrm{Ad}_{(g')^{-1}}(\omega_G(X)) + \omega_G(Y)$$

and

$$(\cdot)^{-1*}\omega_G(X) = -\mathrm{Ad}_g(\omega_G(X)).$$

Thus,

$$\begin{aligned}\sigma^*\omega_G(X) &= (f,(\cdot)^{-1})^*(\mu^*\omega_G)(X,X) = \mu^*\omega_G(f_*X,(\cdot)^{-1}_*X)\\ &= \mathrm{Ad}_{((\cdot)^{-1}(g))^{-1}}(\omega_G(f_*X)) + \omega_G((\cdot)^{-1}_*X)\\ &= \mathrm{Ad}_g(f^*\omega_G(X)) + (\cdot)^{-1*}\omega_G(X)\\ &= \mathrm{Ad}_g(\omega_G(X)) - \mathrm{Ad}_g(\omega_G(X)) = 0. \quad \square\end{aligned}$$

Again, this is very good news for our pursuit of geometric intuition: as we mentioned before when we defined the Maurer-Cartan form, it is one of the most intuitive objects in modern differential geometry. In particular, this means we can understand what the geometry looks like in terms of moving around inside of it.

This part needs to be clearer. We need the reader to understand the intuition of the Maurer-Cartan form before we introduce Cartan connections.

3.3. Spherical geometry

Reread this section and make sure it is satisfactory. We need spherical geometry for when we do projective geometry later.

In the previous chapter, we explained how circles and (unoriented) angles in Euclidean geometry ultimately came from the behavior of the subgroup $\mathrm{O}(2) < \mathrm{I}(2)$, which was the stabilizer of the point $0 \in \mathbb{R}^2 \cong \mathrm{I}(2)/\mathrm{O}(2)$. Using our new terminology, this was clearly referring to $\mathrm{O}(2)$ as the isotropy subgroup of the model $(\mathrm{I}(2), \mathrm{O}(2))$. Now, we will investigate what another geometry with isotropy $\mathrm{O}(2)$ looks like.

Consider an orthonormal basis $\{e_1, e_2, e_3\}$ of $\mathbb{R}^3$ with the usual Euclidean structure. The Lie group $\mathrm{O}(3)$ acts on $\mathbb{R}^3$ by linear isometries (by definition), and we get a copy of $\mathrm{O}(2)$ in $\mathrm{O}(3)$ as the subgroup stabilizing the vector e_1. This points us toward a new model geometry: $(\mathrm{O}(3), \mathrm{O}(2))$, also called (2-dimensional) *spherical geometry*.

Since $\mathrm{O}(2)$ is the stabilizer of e_1, we can identify $\mathrm{O}(3)/\mathrm{O}(2)$ with the $\mathrm{O}(3)$-orbit through e_1, which—as one might expect from our name for this geometry—happens to be precisely the unit 2-sphere S^2 in $\mathbb{R}^3$. Just like we did for $\mathrm{I}(2)$ in Euclidean geometry, we can think of $\mathrm{O}(3)$ as the orthonormal frame bundle over S^2, with bundle map $q_{\mathrm{O}(2)} : \mathrm{O}(3) \to S^2 \cong \mathrm{O}(3)/\mathrm{O}(2)$ given by $g \mapsto g \cdot e_1$. We have illustrated this in Figure 2.

Fix the colors to be consistent with the general themes for the colors. Also, maybe show how the linear frame looks first?

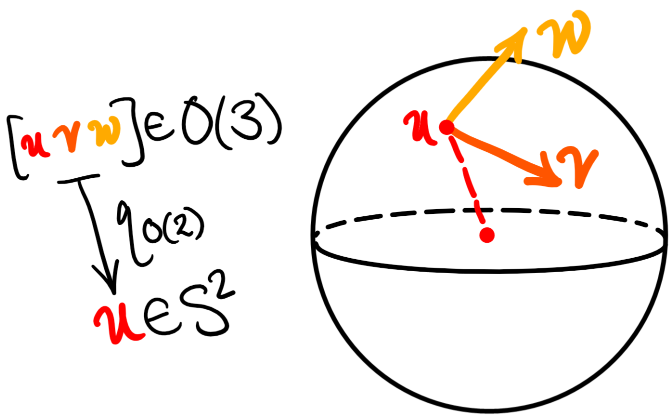

Figure 2. The Lie group O(3) thought of as the orthonormal frame bundle of $S^2 \cong \mathrm{O}(3)/\,\mathrm{O}(2)$

At the identity element $\mathbb{1} = [e_1\, e_2\, e_3] \in \mathrm{O}(3)$, we have the tangent space $T_{\mathbb{1}}\,\mathrm{O}(3)$, which we identify with the Lie algebra

$$\mathfrak{o}(3) = \left\{ \begin{bmatrix} 0 & -x & -y \\ x & 0 & -z \\ y & z & 0 \end{bmatrix} : x, y, z \in \mathbb{R} \right\}.$$

Considering $\mathfrak{i}(2)$ and $\mathfrak{o}(3)$ as O(2)-representations, using the restriction of their adjoint representations to their copy of O(2), we have an isomorphism of O(2)-representations

$$\rho_+ : \mathfrak{i}(2) \to \mathfrak{o}(3)$$

given by

$$\left(\begin{bmatrix} x \\ y \end{bmatrix}, \begin{bmatrix} 0 & -z \\ z & 0 \end{bmatrix} \right) \mapsto \begin{bmatrix} 0 & -x & -y \\ x & 0 & -z \\ y & z & 0 \end{bmatrix}.$$

In particular, O(2) behaves the same way on the subspace $\rho_+(\mathbb{R}^2)$ as it does on the subalgebra of translations $\mathbb{R}^2 < \mathfrak{i}(2)$. The subspace $\rho_+(\mathbb{R}^2)$ is not itself a subalgebra, though the one-parameter subgroups it generates can be thought of as "translations" in spherical geometry.

Writing $\{\bar{e}_1, \bar{e}_2\}$ for the usual orthonormal basis for $\mathbb{R}^2$, with bars to distinguish them from e_1 and e_2 in $\mathbb{R}^3$, we can consider the one-parameter subgroup $\exp(t\rho_+(\bar{e}_1))$ corresponding to

$$\rho_+(\bar{e}_1) = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

As we can see in Figure 3, $\exp(t\rho_+(\bar{e}_1))$ behaves as a transformation when acting on the left, rotating the sphere in a way that preserves the "equator"

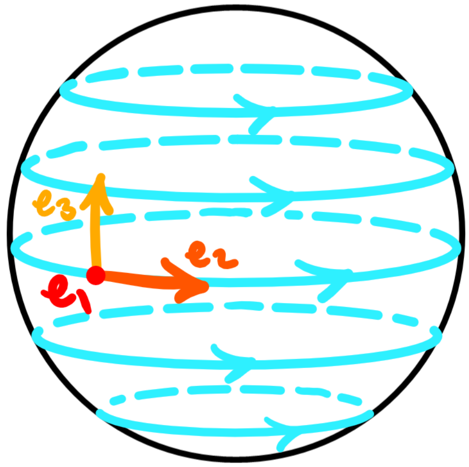

Figure 3. Acting on the left by the one-parameter subgroup of transformations $\exp\left(t\begin{bmatrix}0 & -1 & 0\\ 1 & 0 & 0\\ 0 & 0 & 0\end{bmatrix}\right)$ rotates the sphere in a way that preserves its intersection with the plane $\langle e_1, e_2\rangle$ generated by e_1 and e_2

given by the intersection of the sphere with the plane $\langle e_1, e_2\rangle$ generated by e_1 and e_2 in $\mathbb{R}^3$.

In general, the "translations" $\exp(t\rho_+(x\bar{e}_1+y\bar{e}_2))$ will preserve the great circle given by the intersection of the sphere with $\langle e_1, xe_2+ye_3\rangle$.

Definition 3.4. A *great circle* in spherical geometry is a subset given by the intersection of the sphere with a 2-dimensional subspace (through the origin) of $\mathbb{R}^3$.

Because $\mathrm{O}(3)$ acts linearly on $\mathbb{R}^3$, it sends 2-dimensional subspaces to 2-dimensional subspaces, so since $\mathrm{O}(3)$ also preserves the sphere, it sends great circles to great circles. In other words, the notion of great circle is preserved under the action of the model group $\mathrm{O}(3)$, hence great circles are geometric objects in spherical geometry. We can think of great circles as the spherical analogue of lines.

Note that only one orbit of $\exp(t\rho_+(\bar{e}_1))$ on S^2 was a great circle; all the others were intersections of the sphere with translations of $\langle e_1, e_2\rangle$ by some multiple of e_3, which are not subspaces of $\mathbb{R}^3$ as a vector space. Such intersections are also preserved by the action of $\mathrm{O}(3)$, but they are not the spherical analogue of lines.

When acting on the right, $\exp(t\rho_+(\bar{e}_1))$ behaves as a motion, which we can think of as the spherical analogue of "walking forward" in Euclidean geometry: for $g\in\mathrm{O}(3)$, $g\exp(t\rho_+(\bar{e}_1))$ is given by starting at g and moving for time t along the great circle $S^2\cap g\cdot\langle e_1, e_2\rangle$ with velocity $q_{\mathrm{O}(2)*}((\omega_{\mathrm{O}(3)})^{-1}(\rho_+(\bar{e}_1)))$, the tangent vector corresponding to $g\cdot e_2$.

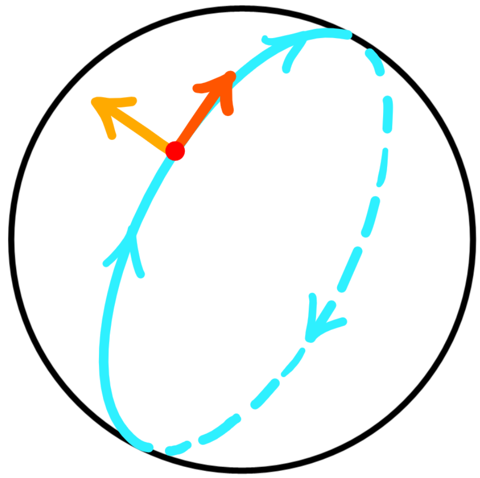

Figure 4. Acting on $g \in \mathrm{O}(3)$ from the right by the one-parameter subgroup of motions $\exp(t\rho_+(\bar{e}_1))$ moves along the great circle given by the intersection of the sphere with the plane $g \cdot \langle e_1, e_2 \rangle$

In particular, motion by a one-parameter subgroup of "translations" in spherical geometry will always trace out great circles on the sphere, so we could have defined great circles to be subsets of the form $q_{\mathrm{O}(2)}(g \exp(\mathbb{R}\rho_+(v)))$ for some $g \in \mathrm{O}(3)$ and some nonzero $v \in \mathbb{R}^2$, as we did with lines in Euclidean geometry.

Exercise 3.5. Using what "translations" look like (as motions) in spherical geometry and the discussion from Chapter 1 on visualizing the Lie bracket, verify that $\rho_+(\mathbb{R}^2)$ is not a subalgebra of $\mathfrak{o}(3)$ *without* doing any computations.

For circles, we can use the same definition as before in Euclidean geometry: pick a center $x \in S^2$, a radius given by some $a \in \mathrm{O}(3)$, and define $C_a(x) := \{q_{\mathrm{O}(2)}(ga) : g \in q_{\mathrm{O}(2)}^{-1}(x)\}$. Again, such sets are given by "all the points in S^2 that some orthonormal frame over x thinks is a away from itself". They happen to be precisely those (nonempty) intersections of the sphere with affine planes from before; in particular, great circles also happen to be special examples of circles in spherical geometry.

3.4. Hyperbolic geometry

Expand this section and make it better. We need hyperbolic geometry for representation theory later.

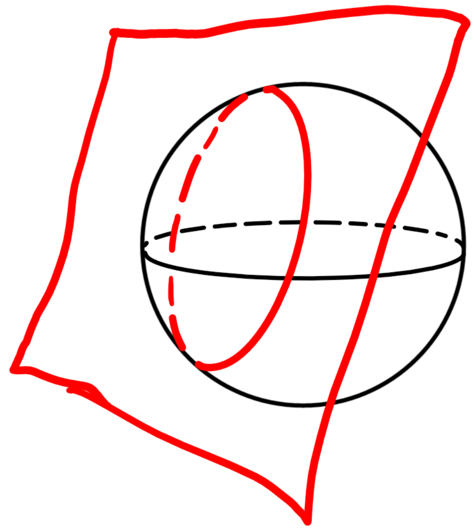

Figure 5. The intersection of an affine plane in $\mathbb{R}^3$ with S^2 gives a circle

Now, consider $\mathbb{R}^3$ with the Minkowski quadratic form Q given by

$$Q(ae_1 + be_2 + ce_3) := a^2 - b^2 - c^2,$$

for which the linear isometries are given by $\mathrm{O}(1, 2)$. When we take $\mathrm{O}(1, 2)$ and quotient by the center, generated by $-\mathbb{1}$, we get $\mathrm{PO}(1, 2)$, which naturally acts on $\mathbb{RP}^2$, the space of 1-dimensional linear subspaces of $\mathbb{R}^3$.

In $\mathbb{R}^3$, the set $Q^{-1}(1)$ is a two-sheeted hyperboloid on which $\mathrm{O}(1, 2)$ acts transitively. The two sheets are images of each other under the linear transformation $-\mathbb{1}$, so the image of $Q^{-1}(1)$ in $\mathbb{RP}^2$ is connected and $\mathrm{PO}(1, 2)$ acts transitively on it. The stabilizer of the line $\langle e_1 \rangle$ is a copy of $\mathrm{O}(2)$, which leads us to consider the model geometry $(\mathrm{PO}(1, 2), \mathrm{O}(2))$, also called (2-dimensional) hyperbolic geometry.

Again, we identify $\mathrm{PO}(1, 2)$ with the orthonormal frame bundle of $\mathbb{H}^2 \cong \mathrm{PO}(1, 2)/\,\mathrm{O}(2)$, with bundle map $q_{\mathrm{O}(2)} : \mathrm{PO}(1, 2) \to \mathbb{H}^2$ given by $g \mapsto g \cdot \langle e_1 \rangle$. Indeed, we can topologically identify $\mathbb{H}^2$ with a more familiar space: choosing a sheet of $Q^{-1}(1)$, each point of the sheet projects to a unique point of the plane $\langle e_2, e_3 \rangle$, so we can topologically identify $\mathbb{H}^2$ with $\mathbb{R}^2$.

As with spherical geometry, there is a convenient isomorphism of $\mathrm{O}(2)$-representations

$$\rho_- : \mathfrak{i}(2) \to \mathfrak{po}(1, 2)$$

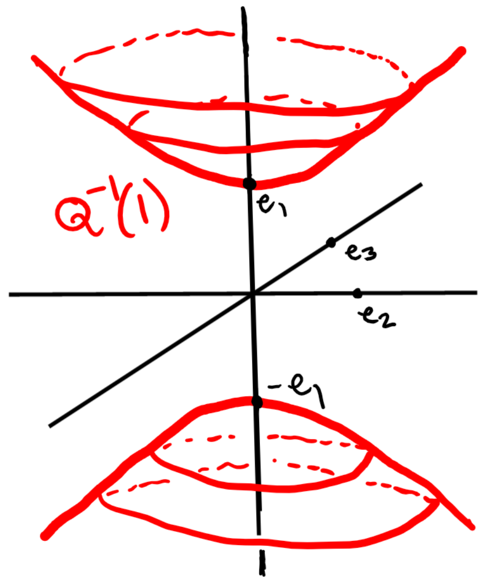

Figure 6. A drawing of the two-sheeted hyperboloid $Q^{-1}(1)$ in $\mathbb{R}^3$

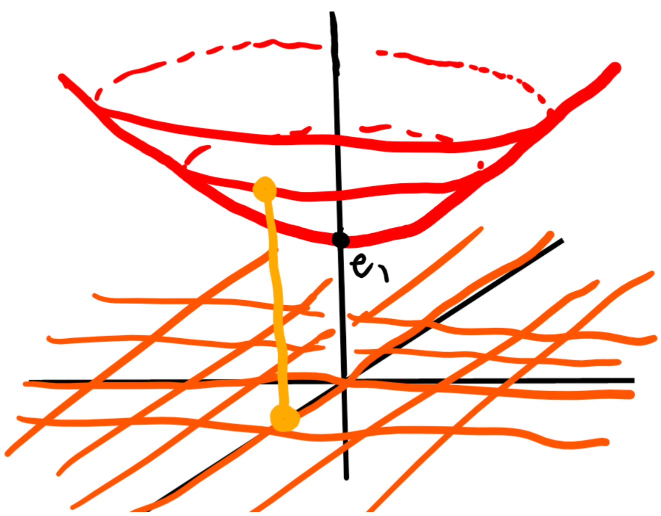

Figure 7. Each point of a sheet of $Q^{-1}(1)$ projects to a unique point of the plane $\langle e_2, e_3 \rangle$

given by

$$\left(\begin{bmatrix} x \\ y \end{bmatrix}, \begin{bmatrix} 0 & -z \\ z & 0 \end{bmatrix} \right) \mapsto \begin{pmatrix} 0 & x & y \\ x & 0 & -z \\ y & z & 0 \end{pmatrix},$$

so that $\mathrm{O}(2)$ behaves the same way on the subspace $\rho_-(\mathbb{R}^2)$ as it does on the subalgebra of translations in Euclidean geometry.

Similar to the above, we get hyperbolic analogues of lines—called *geodesics*—by taking images of (nonempty) intersections with $Q^{-1}(1)$ of 2-dimensional linear subspaces in $\mathbb{R}^3$. Naturally, these geodesics happen to be equivalent to subsets of the form $q_{\mathrm{O}(2)}(g\exp(\mathbb{R}\rho_-(v)))$ for some $g \in \mathrm{PO}(1,2)$ and some nonzero $v \in \mathbb{R}^2$. Circles follow a similar pattern to before as well.

3.5. Affine geometry

Expand this section and make it better. Attention to the change in isotropy helps solidify the general intuition for these geometries as structures on principal bundles.

At the end of the last chapter, we also described how the behavior of lines and parallelism in Euclidean geometry came from the closed normal subgroup of translations $\mathbb{R}^2$ acting simply transitively on the Euclidean plane. In other words, in a model geometry (G,H) such that the model group G has a closed normal subgroup isomorphic to $\mathbb{R}^2$ that acts simply transitively on G/H, we should get the same notions of lines and parallelism.

To give another example of this, consider the Lie group $\mathrm{Aff}(2)$ of transformations of the plane generated by translations and (not necessarily isometric) linear transformations. This is the group of *affine transformations* of the plane, and by essentially the same argument we made for $\mathrm{I}(2)$, we get an isomorphism $\mathrm{Aff}(2) \simeq \mathbb{R}^2 \rtimes \mathrm{GL}_2\,\mathbb{R}$. The model $(\mathrm{Aff}(2), \mathrm{GL}_2\,\mathbb{R})$ gives (2-dimensional) affine geometry.

Instead of the *orthonormal* frame bundle of $\mathbb{R}^2$, we identify $\mathrm{Aff}(2)$ with the full frame bundle of $\mathbb{R}^2$.

Definition 3.6. A *frame* over a point $x \in \mathbb{R}^2$ is a linear isomorphism from $\mathbb{R}^2 \approx T_0\mathbb{R}^2$ to $T_x\mathbb{R}^2$.

As before, elements $g \in \mathrm{Aff}(2)$ are identified with their pushforwards at the origin $g_* : T_0\mathbb{R}^2 \to T_{g(0)}\mathbb{R}^2$, and we get the natural bundle map $q_{\mathrm{GL}_2\,\mathbb{R}} : \mathrm{Aff}(2) \to \mathbb{R}^2 \cong \mathrm{Aff}(2)/\mathrm{GL}_2\,\mathbb{R}$ given by $g \mapsto g(0) \cong g\,\mathrm{GL}_2\,\mathbb{R}$.

It is worth spending some time thinking about what it is like to be a pedestrian on the affine plane. We are accustomed to only being able to rotate on the spot, but in affine geometry, we have a much wider range of options. For example, imagine "rotating" by the unipotent transformations $\left[\begin{smallmatrix}1 & t\\ 0 & 1\end{smallmatrix}\right]$: our notion of "forward" remains the same, but our notion of "left" skews forward by t.

Figure 8. Right-translating by $\left[\begin{smallmatrix}1&1\\0&1\end{smallmatrix}\right]$ looks quite different from just rotating on the spot

We can also rescale ourselves by right-translating by a linear transformation of the form $\lambda\mathbb{1}$ for some $\lambda > 0$. For $\lambda \in (0,1)$, such transformations shrink us, and for $\lambda \in (1,+\infty)$, they expand us.

Figure 9. Right-translating by $\frac{1}{2}\mathbb{1}$ rescales us by $1/2$

We can, again, define lines as subsets of $\mathrm{Aff}(2)/\operatorname{GL}_2\mathbb{R} \cong \mathbb{R}^2$ of the form $q_{\operatorname{GL}_2\mathbb{R}}(g\exp(\mathbb{R}v))$ for some $g \in \mathrm{Aff}(2)$ and some nonzero translational velocity $v \in \mathbb{R}^2 < \mathfrak{aff}(2)$. This definition coincides with the usual notion of line, and parallelism works the same as it does in Euclidean geometry by the argument from last lecture.

Of course, we already knew that lines and parallelism would be the same as before; that was the point. We start to see changes when we try to use the definition of circles from last time. Indeed, for $a \in \mathrm{Aff}(2)$ and $x \in \mathbb{R}^2$, consider the set

$$C_a(x) := \{q_{\operatorname{GL}_2\mathbb{R}}(ga) : g \in q^{-1}_{\operatorname{GL}_2\mathbb{R}}(x)\}.$$

For $a \in \operatorname{GL}_2\mathbb{R}$, $C_a(x)$ is just the point x. For $a \in \tau_v\operatorname{GL}_2\mathbb{R}$ for some nonzero $v \in \mathbb{R}^2$, however,

$$q_{\operatorname{GL}_2\mathbb{R}}(ga) = q_{\operatorname{GL}_2\mathbb{R}}(\tau_x \circ A \circ \tau_v) = q_{\operatorname{GL}_2\mathbb{R}}(\tau_{x+A(v)})$$

for some $A \in \operatorname{GL}_2\mathbb{R}$ such that $g = \tau_x \circ A$, hence $C_a(x)$ is the set $x + \operatorname{GL}_2\mathbb{R}\cdot v = \mathbb{R}^2 \setminus \{x\}$, the set of all points on the plane other than x. In other words, a nontrivial "affine circle" centered at a point will just be the complement of that point.

3.6. Comparing model geometries

[Extension functors?]

3.7. What makes an interesting model geometry?

Knowing that the geometric properties of models ultimately come from the behavior of their symmetries, it should not be especially surprising to learn that Lie-theoretically interesting pairs (G, H) tend to give geometrically interesting model geometries. Of course, we are then led to an obvious question: "What are some ways to identify Lie-theoretically interesting models?"

While the matter of what makes a particular Lie group or subgroup "interesting" is somewhat subjective, there are two main classes of models that are generally agreed to give compelling examples. The first, and also by far the most well-understood, are the so-called *reductive geometries*.

Definition 3.7. A model (G, H) is said to be *reductive*[3] if and only if there exists an Ad_H-invariant subspace $\mathfrak{m}$ that is complementary to $\mathfrak{h}$ in $\mathfrak{g}$, so that $\mathfrak{g} = \mathfrak{m} \oplus \mathfrak{h}$ as an H-representation.

All of the model geometries presented so far have been reductive; the subspace of "translations" that we constructed for each of them is precisely the Ad_H-invariant complement in the above definition.

The other major class of examples—currently the one that most researchers are (justifiably) excited about—are the *parabolic geometries*, which consist of a pair (G, P) where G is a semisimple Lie group and P is a parabolic subgroup. Unfortunately, these will require a considerable amount of set-up to understand intuitively, so we will wait until later in the book to give the relevant definitions.

[3]Note that this is different from requiring that G or H be reductive Lie groups.

Chapter 4

Geometry from Motion

We live within a geometric world, interacting with geometric phenomena almost constantly as part of being alive. Our minds must, then, have some way of interpreting this geometric information from the world around us. Can we describe this mechanism through which we perceive our ambient geometry, and if so, then can we encode it into diffeo-geometric terms?

This is the key question we will seek to answer in this chapter: *how do we formalize the way in which we experience geometry*?

4.1. How to walk in a straight line

To answer this question, let us start with a seemingly innocuous observation: human beings often move along geodesic paths. When we cross the street, for example, we typically walk straight across. As we move along a geodesic path, our minds must somehow interact with our ambient geometry—at the very least to tell our bodies what to do—so if we examine how we think about moving when we walk along geodesic paths, then we might be able to answer our question.

In an effort to negate any obfuscatory sensory stimuli, imagine that we blindfold ourselves, plug our ears, and then have someone take us to an unfamiliar location with plenty of open space to walk around in. For the purposes of this thought experiment, assume that we know nothing about our surroundings except that we can move freely without risk of injuring ourselves. Now, in this situation, we ask ourselves: how do we move in a straight line?

After some thought, most people seem to answer fairly similarly: to walk in a straight line, we simply "pick a direction, and then keep going in that direction (until we hit some kind of obstacle)". Crucially, the language used to describe this motion almost always implicitly assumes that the direction in which we are moving is, in some way, constant. This implies that, formally, we are somehow comparing tangent vectors from different tangent spaces in order to determine whether two velocities are pointed in the "same" direction.

Consider how we pick a given direction in the situation of our thought experiment. Since we do not know anything about our surroundings, we cannot use them to specify a direction; the only thing we can use for reference is ourselves. With this in mind, our formalization of this mechanism for how we interact with our ambient geometry should start to sound very familiar, because this forces us to name directions in terms of words like "forward" and "left" that only reference ourselves and our configuration within the geometry.

Of course, "forward" and "left" depend upon our configuration, so if we change our configuration, then "forward" and "left" change as well. This means that the mechanism formalizing how we interact with the geometry around us relies on the space $\mathscr{G}$ of configurations for ourselves over the space M that we are walking around in. On the other hand, without information from our surroundings, there is nothing to distinguish between different configurations lying over the same point of M; they should all, in some sense, be "the same". Formally, this means that we have a group H, which we might as well assume to be a Lie group, that acts transitively on all of the different configurations lying over a given point of M. Because turning in some direction should always result in a different configuration, we also want this action to be free, and since each point of M should allow for the same range of configurations, this free action should extend to all of $\mathscr{G}$, acting transitively on the space of configurations lying over each point. In other words, we can think of our configuration space $\mathscr{G}$ as a principal H-bundle over the space M that we are walking around in, with a quotient map that we will call $q_H : \mathscr{G} \to M$.

Even without knowing where we are in the geometry, we can describe our motion within the configuration space $\mathscr{G}$ using words like "forward" or "left" or "turn left (counterclockwise)", since these terms are meaningful regardless of our current configuration within the geometry. To formalize this, we consider a constant vector space $\mathfrak{g}$ of instantaneous motions and use a $\mathfrak{g}$-valued one-form ω to identify each space of velocities inside the configuration space $\mathscr{G}$ with this constant vector space, allowing us to describe vectors in arbitrary tangent spaces using terms like "forward" that we can prescribe for the constant vector space.

Within the context of our current example, we are thinking of M as a 2-dimensional manifold (the floor we are walking on), with the Lie group H corresponding to either O(2) or SO(2) to represent all the different ways we can turn on the spot as we stand atop the surface. Our vector space $\mathfrak{g}$ of instantaneous motions is then given by the sum of a 2-dimensional space $\mathbb{R}^2$ of velocities—generated by vectors that we can designate as "forward" and "leftward"—and a 1-dimensional space $\mathfrak{h} \approx \mathfrak{o}(2)$ of angular velocities; in other words, we can identify the space $\mathfrak{g}$ of instantaneous motions with $\mathfrak{i}(2)$, the Lie algebra of the isometry group of the Euclidean plane. Using this constant vector space for reference, the mechanism ω allows us to compare velocities within the configuration space with corresponding instantaneous motions in $\mathfrak{i}(2)$ the way we usually do when we, for example, walk in a straight line. If we want to walk "forward" along a geodesic path, then we simply move with, at each point in time, a velocity that ω identifies as "forward" in $\mathfrak{i}(2)$; writing e_1 for the "unit forward" vector in $\mathbb{R}^2 < \mathfrak{i}(2)$, this can be achieved by flowing along the vector field $\omega^{-1}(e_1)$ as in Figure 1.

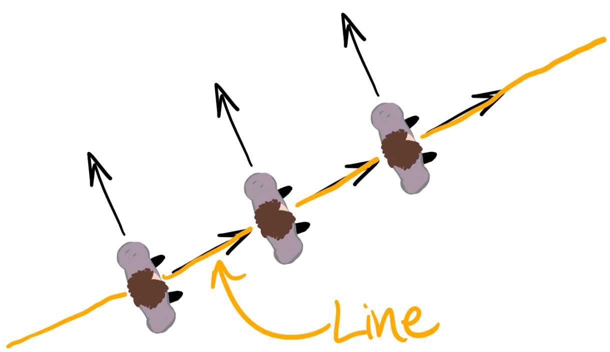

Figure 1. Walking forward along a geodesic with unit speed by flowing along the vector field $\omega^{-1}(e_1)$

4.2. The definition of Cartan geometries

At this point, we should hopefully all recognize that the mechanism ω bears striking resemblance to a Maurer-Cartan form. Indeed, it basically fulfills the same role that the Maurer-Cartan form did for a model geometry (G, H), both as the defining invariant for the geometric structure and as a way of describing how we move around within the configuration space. The key difference, naturally, is that the configuration space $\mathscr{G}$, while still a principal H-bundle over the base space M, is not necessarily a Lie group G anymore.

On the other hand, it did *look* like a Lie group; while the configuration space $\mathscr{G}$ might not have literally been I(2) in our example, the geometry clearly resembled that of the model geometry $(\mathrm{I}(2), \mathrm{O}(2))$. In fact, one might even say that the geometry determined by the pair $(\mathscr{G}, \omega)$ was *modelled* on $(\mathrm{I}(2), \mathrm{O}(2))$.

In general, we can model such *Cartan geometries* $(\mathscr{G}, \omega)$ on any model geometry (G, H) that we want. The one-form ω will, then, be an example of what we will call a *Cartan connection*, but first, we need to specify a few criteria that it must satisfy. These three criteria, while they may appear intimidating at first glance, actually just encode "common sense" things that, implicitly, we are probably already assuming about ω.

First, and most obviously, we want ω to *identify* each tangent space of $\mathscr{G}$ with the constant space $\mathfrak{g}$ of instantaneous motions. The whole point of ω is to allow us to consistently describe arbitrary velocities in the configuration space in terms of corresponding elements of $\mathfrak{g}$; it would be silly to somehow forget what "forward" means at certain points, for example. Formally, this just means that, for each configuration $\mathscr{g} \in \mathscr{G}$, the map

$$\omega_{\mathscr{g}} : T_{\mathscr{g}}\mathscr{G} \to \mathfrak{g}$$

should be a linear isomorphism.

Second, when we turn on the spot, the ground beneath us should stay where it is, even though our description of it changes. To see what we mean by this, imagine we are standing on top of an arrow painted onto the ground, facing toward its tip as in Figure 2.

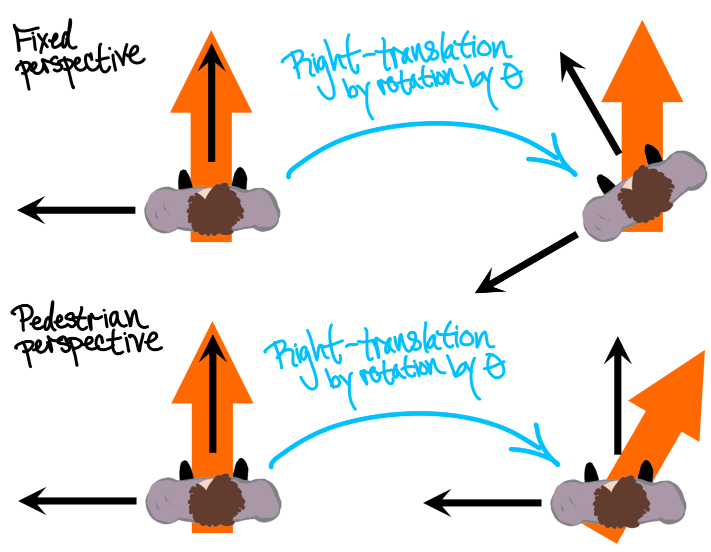

Figure 2. If we start pointed in the direction indicated by an orange arrow painted onto the ground, and then we rotate by some angle θ, then our description using ω of the direction indicated by the arrow is rotated in the opposite direction by the same angle θ

If we turn ourselves on the spot by an angle of θ, then the painted arrow does not move, but we do. In particular, the arrow no longer points "forward" for us; from our new perspective, the arrow has rotated in the opposite direction by the same angle θ. Formally, if we started at $\mathscr{g} \in \mathscr{G}$, then the arrow specifies a vector $v \in T_{q_H(\mathscr{g})}M$, and we can lift this to a vector $\hat{v} \in (q_H)_{*\mathscr{g}}^{-1}(v) \subseteq T_{\mathscr{g}}\mathscr{G}$. If we then right-translate by $h \in H$, pushing $\hat{v} \in T_{\mathscr{g}}\mathscr{G}$ to $\mathrm{R}_{h*}\,\hat{v} \in T_{\mathscr{g}h}\mathscr{G}$, then our description $\omega_{\mathscr{g}h}(\mathrm{R}_{h*}\,\hat{v})$ of $\mathrm{R}_{h*}\,\hat{v}$ using ω is given by applying the inverse of h to the original description $\omega_{\mathscr{g}}(\hat{v})$. More succinctly, for each $h \in H$, $\mathrm{R}_h^*\,\omega = \mathrm{Ad}_{h^{-1}}\,\omega$.

Third, *moving* with a given angular velocity should always correspond to *rotating* with that angular velocity. In the context of our example modelled on $(\mathrm{I}(2), \mathrm{O}(2))$, consider the curve $\gamma : \mathbb{R} \to \mathscr{G}$ given by $t \mapsto \mathscr{g}\exp(t\left[\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right])$, where $\left[\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right] \in \mathfrak{o}(2) \approx \mathfrak{h}$ is the "unit counterclockwise" angular velocity and $\mathscr{g} \in \mathscr{G}$. As before, we would describe the motion along γ within the principal $\mathrm{O}(2)$-bundle $\mathscr{G}$ as *rotating* on the spot with constant "unit counterclockwise" angular velocity for time t, and since ω is there to describe motion within the configuration space $\mathscr{G}$ in precisely this way, we want its description of the velocities of γ to agree with this. Therefore, we want $\omega(\dot{\gamma}(t)) = \left[\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right]$ for all t, since γ is *moving* with constant "unit counterclockwise" angular velocity. In general, we will have a subalgebra $\mathfrak{h} \leq \mathfrak{g}$ corresponding to the isotropy H, and we want ω to use the elements of $\mathfrak{h}$ to describe their corresponding right-actions on the principal H-bundle $\mathscr{G}$. Formally, this means that, for $Y \in \mathfrak{h}$, flowing for time t along the vector field $\omega^{-1}(Y)$ should be equivalent to right-translation by $\exp(tY)$: $\exp(t\omega^{-1}(Y)) = \mathrm{R}_{\exp(tY)}$.

With these three criteria, we can finally define Cartan geometries.

Definition 4.1. Let (G, H) be a model and M be a smooth manifold. A *Cartan geometry of type* (G, H) *over* M is a pair $(\mathscr{G}, \omega)$, where $\mathscr{G}$ is a principal H-bundle over M with quotient map $q_H : \mathscr{G} \to M$ and ω is a $\mathfrak{g}$-valued one-form—called the *Cartan connection*—on $\mathscr{G}$ satisfying the following three criteria:

- For every $\mathscr{g} \in \mathscr{G}$, $\omega_{\mathscr{g}} : T_{\mathscr{g}}\mathscr{G} \to \mathfrak{g}$ is a linear isomorphism.
- For every $h \in H$, $\mathrm{R}_h^*\,\omega = \mathrm{Ad}_{h^{-1}}\,\omega$.
- For every $Y \in \mathfrak{h}$ and $t \in \mathbb{R}$, $\exp(t\omega^{-1}(Y)) = \mathrm{R}_{\exp(tY)}$.

To summarize, a Cartan geometry $(\mathscr{G}, \omega)$ of type (G, H) over M consists of a principal H-bundle $\mathscr{G}$, which we think of as the space of configurations for ourselves as pedestrians on or within the space M, together with a Cartan connection ω, which is a $\mathfrak{g}$-valued one-form on $\mathscr{G}$ that describes velocities within the configuration space in terms of the constant vector space $\mathfrak{g}$ of instantaneous motions. The Cartan connection ω is required to satisfy three criteria that, while initially impressive-looking, are actually just formalizing certain "common sense" stipulations that we likely would have assumed about ω anyway. Finally, the geometric structure determined by the pair $(\mathscr{G}, \omega)$ *looks* like the geometry of (G, H), with $\mathscr{G}$ resembling G, $M \cong \mathscr{G}/H$ resembling G/H, and ω resembling the Maurer-Cartan form of G.

4.3. Klein geometries and curvature

[Use ω_G as the Cartan connection to get an easy example of a Cartan geometry]

[Comparing Cartan geometries using geometric maps]

[Bumps and hills?]

Chapter 5

How to Pretend to Do Lie Theory

In the previous chapter, we defined Cartan geometries and saw that they, in some sense, resemble their archetypal model geometries. Specifically, for a Cartan geometry $(\mathscr{G}, \omega)$ of type (G, H), we said that the principal H-bundle $\mathscr{G}$ *looked* like the Lie group G. This comparison has, so far, been fairly superficial. After all, $\mathscr{G}$ is generally not going to be a Lie group, so we cannot talk about doing Lie theory on $\mathscr{G}$ the way that we can with G.

But what if we could? What if we *pretended* that $\mathscr{G}$ was actually the Lie group G, and that ω was actually the Maurer-Cartan form ω_G?

In this chapter, as fanciful and implausible as it may seem, we will describe how to do this rather judiciously. The analogy between Lie groups and Cartan geometries turns out to be surprisingly robust, and as long as we take the holonomy and topology into account, we can—kind of—treat Cartan geometries as if they were their model groups.

5.1. Development

Remember to point toward the appendix! We want to completely eliminate any barriers to collaboration between Cartan geometers and the people studying locally homogeneous geometric structures, so we need them to see that we're talking about the same things in different ways.

5.2. Holonomy

5.3. Subgroups and cosets

5.4. Automorphisms of Cartan geometries

Chapter 6

Riemannian Geometry

At some point, we need to give the reader a chapter to put the ideas so far into more concrete practice. Doing this with Riemannian geometry here seems like a good choice, though we probably want to do some stuff with affine geometry too?
We also don't want to overwhelm the reader with Riemannian geometry stuff; they're probably reading these notes because they want to understand parabolic geometries.

Shortlist of potential topics to include:

- Torsion (getting the unique Levi-Civita connection motivates normality later)
- Model mutations and extension functors (being able to prove the classification of space forms in three lines is a fairly good indication of the power of the machinery)
- Isometries?
- Recovering the Riemannian metric and distances?

Chapter 7

How to Pretend to Do Representation Theory

[Tractor bundles and tractor connections]

[Levi-Civita and affine connections as examples]

Chapter 8

Interpreting the Killing Form

Again, rewrite this so that it looks better and fits well within the book.

The Killing form is an *exceptionally* powerful idea in Lie theory. It is the key to understanding much of the structure theory of Lie algebras, and as such, it will be a vital part of our exploration into parabolic geometries.

Unfortunately, the Killing form is quite tricky to understand intuitively. This lecture and its sequel are the result of nearly a decade of trying to understand the intuition behind the Killing form; I am still not entirely satisfied—perhaps after another eight years I'll have even better answers—but I hope that, by sharing this, I can help you avoid struggling with it as much as I did.

To be clear, I do not consider "it's just the trace form" to be intuition. Given a pair of elements in a Lie algebra, we can obviously compute the Killing form applied to that pair, but that ultimately just gives us a number. Visually, what is that number telling us?

The lecture should proceed as follows:

- Review the definition of the Killing form
- Rediscover a convincing reason for the conic section terminology in the classification of elements of $\mathfrak{sl}_2\mathbb{R}$
- Compare the Killing form on $\mathfrak{sl}_2\mathbb{R}$ to the notion of eccentricity for conic sections
- Learn how to interpret the Killing form for general Lie algebras

As we said above, a fundamental understanding of the Killing form will be crucial for the lectures to come. In the next lecture, we will present the Killing form in a more geometric context, after which we will finally be ready to define parabolic subgroups.

8.1. Introduction

To start, let's define the Killing form.

Definition 8.1. The *Killing form* $\mathfrak{B}$ on a Lie algebra $\mathfrak{g}$ is the symmetric bilinear form given by $\mathfrak{B}(X, Y) := \operatorname{tr}(\operatorname{ad}_X \circ \operatorname{ad}_Y)$.

As an example, let us look at $\mathfrak{sl}_2\mathbb{R}$. For $\left[\begin{smallmatrix} a & b \\ c & -a \end{smallmatrix}\right], \left[\begin{smallmatrix} z & y \\ x & -z \end{smallmatrix}\right] \in \mathfrak{sl}_2\mathbb{R}$,

$$\begin{aligned}\left[\begin{bmatrix} a & b \\ c & -a \end{bmatrix}, \left[\begin{bmatrix} z & y \\ x & -z \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}\right]\right] &= \left[\begin{bmatrix} a & b \\ c & -a \end{bmatrix}, \begin{bmatrix} y & 0 \\ -2z & -y \end{bmatrix}\right] \\ &= \begin{bmatrix} -2bz & -2by \\ 2cy + 4az & 2bz \end{bmatrix},\end{aligned}$$

$$\begin{aligned}\left[\begin{bmatrix} a & b \\ c & -a \end{bmatrix}, \left[\begin{bmatrix} z & y \\ x & -z \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}\right]\right] &= \left[\begin{bmatrix} a & b \\ c & -a \end{bmatrix}, \begin{bmatrix} 0 & -2y \\ 2x & 0 \end{bmatrix}\right] \\ &= \begin{bmatrix} 2(bx + cy) & -4ay \\ -4ax & -2(bx + cy) \end{bmatrix},\end{aligned}$$

and

$$\begin{aligned}\left[\begin{bmatrix} a & b \\ c & -a \end{bmatrix}, \left[\begin{bmatrix} z & y \\ x & -z \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}\right]\right] &= \left[\begin{bmatrix} a & b \\ c & -a \end{bmatrix}, \begin{bmatrix} -x & 2z \\ 0 & x \end{bmatrix}\right] \\ &= \begin{bmatrix} -2cz & 4az + 2bx \\ -2cx & 2cz \end{bmatrix},\end{aligned}$$

so

$$\begin{aligned}\mathfrak{B}\left(\left[\begin{smallmatrix} a & b \\ c & -a \end{smallmatrix}\right], \left[\begin{smallmatrix} z & y \\ x & -z \end{smallmatrix}\right]\right) &= \operatorname{tr}\left(\operatorname{ad}_{\left[\begin{smallmatrix} a & b \\ c & -a \end{smallmatrix}\right]} \circ \operatorname{ad}_{\left[\begin{smallmatrix} z & y \\ x & -z \end{smallmatrix}\right]}\right) \\ &= 2cy + 4az + 2(bx + cy) + 4az + 2bx \\ &= 8az + 4(bx + cy).\end{aligned}$$

In particular, note that the elements $\left[\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right]$, $\left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right]$, and $\left[\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right]$ are orthogonal with respect to $\mathfrak{B}$, with

$$\mathfrak{B}(\left[\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right], \left[\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right]) = \mathfrak{B}(\left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right], \left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right]) = 8$$

and

$$\mathfrak{B}(\left[\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right], \left[\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right]) = -8,$$

so that $\mathfrak{B}$ is nondegenerate on $\mathfrak{sl}_2\mathbb{R}$ with signature $(2, 1)$.

Arguably one of the main reasons that the Killing form is so remarkably useful is that it is intrinsic to the Lie algebra itself, so that it does not depend on any particular description of its elements. Formally, this just means that the Killing form is automorphism-invariant.

Proposition 8.2. *The Killing form is invariant under automorphisms of the Lie algebra. In other words, if ϕ is an automorphism of $\mathfrak{g}$ and $X, Y \in \mathfrak{g}$, then $\mathfrak{B}(\phi(X), \phi(Y)) = \mathfrak{B}(X, Y)$.*

Proof. Because ϕ is an automorphism, $[\phi(X), Y] = \phi([X, \phi^{-1}(Y)])$, hence $\mathrm{ad}_{\phi(X)} = \phi \circ \mathrm{ad}_X \circ \phi^{-1}$. Thus,

$$\begin{aligned}\mathfrak{B}(\phi(X), \phi(Y)) &= \mathrm{tr}(\mathrm{ad}_{\phi(X)} \circ \mathrm{ad}_{\phi(Y)}) \\ &= \mathrm{tr}((\phi \circ \mathrm{ad}_X \circ \phi^{-1}) \circ (\phi \circ \mathrm{ad}_Y \circ \phi^{-1})) \\ &= \mathrm{tr}(\phi \circ (\mathrm{ad}_X \circ \mathrm{ad}_Y) \circ \phi^{-1}) \\ &= \mathrm{tr}(\mathrm{ad}_X \circ \mathrm{ad}_Y) = \mathfrak{B}(X, Y). \quad \square\end{aligned}$$

Note that, for every $g \in G$, Ad_g is an automorphism of $\mathfrak{g}$, so for $X, Y, Z \in \mathfrak{g}$, $\mathfrak{B}(\mathrm{Ad}_{\exp(tX)}(Y), \mathrm{Ad}_{\exp(tX)}(Z)) = \mathfrak{B}(Y, Z)$. Differentiating this, we get another useful property of the Killing form.

Corollary 8.3. *For $X, Y, Z \in \mathfrak{g}$, $\mathfrak{B}(\mathrm{ad}_X(Y), Z) + \mathfrak{B}(Y, \mathrm{ad}_X(Z)) = 0$.*

Even with just this information so far, the Killing form lets us prove very interesting things.

Proposition 8.4. *The Lie algebra $\mathfrak{sl}_2\mathbb{R}$ is isomorphic to $\mathfrak{o}(1, 2)$.*

Proof. Consider the symmetric bilinear form $-\mathfrak{B}$ on the 3-dimensional vector space $\mathfrak{sl}_2\mathbb{R}$. By Corollary 8.3, for $X, Y, Z \in \mathfrak{sl}_2\mathbb{R}$,

$$-\mathfrak{B}(\mathrm{ad}_X(Y), Z) - \mathfrak{B}(Y, \mathrm{ad}_X(Z)) = 0,$$

so under the adjoint representation, $\mathfrak{sl}_2\mathbb{R}$ maps into the Lie algebra $\mathfrak{o}(-\mathfrak{B}) \approx \mathfrak{o}(1, 2)$. Since the adjoint representation of $\mathfrak{sl}_2\mathbb{R}$ is injective and $\dim(\mathfrak{sl}_2\mathbb{R}) = 3 = \dim(\mathfrak{o}(1, 2))$, this means that the adjoint representation gives an isomorphism of $\mathfrak{sl}_2\mathbb{R}$ with $\mathfrak{o}(1, 2)$. $\square$

An explicit realization of this isomorphism $\rho : \mathfrak{sl}_2\mathbb{R} \to \mathfrak{o}(1, 2)$ is given by

$$\begin{bmatrix} a & b \\ c & -a \end{bmatrix} \mapsto \begin{bmatrix} 0 & b+c & 2a \\ b+c & 0 & c-b \\ 2a & b-c & 0 \end{bmatrix},$$

with inverse $\rho^{-1} : \mathfrak{o}(1, 2) \to \mathfrak{sl}_2\mathbb{R}$ given by

$$\begin{bmatrix} 0 & r & s \\ r & 0 & -t \\ s & t & 0 \end{bmatrix} \mapsto \frac{1}{2} \begin{bmatrix} s & r+t \\ r-t & -s \end{bmatrix}.$$

We'll use this to relate elements of $\mathfrak{sl}_2\mathbb{R}$ to hyperbolic geometry.

8.2. Classification of elements of $\mathfrak{sl}_2\mathbb{R}$

Elements of $\mathfrak{sl}_2\mathbb{R}$ have a well-known classification using the terminology of conic sections: every $X \in \mathfrak{sl}_2\mathbb{R}$ is either hyperbolic, parabolic, or elliptic.

Definition 8.5. Suppose $X \in \mathfrak{sl}_2\mathbb{R}$, viewed as a linear endomorphism of $\mathbb{R}^2$.

- If X is diagonalizable over $\mathbb{R}$, then we say that X is *hyperbolic*.
- If X is nilpotent, then we say that X is *parabolic*.
- If X has purely imaginary eigenvalues, then we say that X is *elliptic*.

It's not too difficult to see that every nonzero $X \in \mathfrak{sl}_2\mathbb{R}$ falls into exactly one of these three categories: because X has trace 0 by definition, the complex eigenvalues of X must be λ and $-\lambda$ for some $\lambda \in \mathbb{C}$. Since X is a real matrix, the eigenvalues must be complex conjugates of each other if they are not real, so if they are not real, then λ and $-\lambda = \bar{\lambda}$ must be purely imaginary. If $\lambda = 0 = -\lambda$ and $X \neq 0$, then X must be nilpotent. Finally, if λ and $-\lambda$ are real and nonzero, then X is diagonalizable over $\mathbb{R}$ by definition.

Of course, none of this explains why we're using this conic section terminology. Where does this terminology come from?

Recall that our model for the hyperbolic plane was $(\mathrm{PO}(1,2), \mathrm{O}(2))$, so that elements of $\mathfrak{o}(1,2)$ determine one-parameter subgroups of hyperbolic isometries. We also had a projection map $\mathrm{pr} : \mathbb{H}^2 \to \mathbb{R}^2$, given by identifying $\mathbb{H}^2 \cong \mathrm{PO}(1,2)/\,\mathrm{O}(2)$ with the sheet through e_1 of the hyperboloid $Q^{-1}(1)$ for $Q(ae_1 + be_2 + ce_3) = a^2 - b^2 - c^2$ and then projecting to the plane $\langle e_2, e_3\rangle$, which allowed us to topologically identify $\mathbb{H}^2$ with $\mathbb{R}^2$.

Utilizing the isomorphism $\rho : \mathfrak{sl}_2\mathbb{R} \to \mathfrak{o}(1,2)$, let us look at some one-parameter subgroups of hyperbolic isometries corresponding to elements of $\mathfrak{sl}_2\mathbb{R}$.

The parabolic element $\left[\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right]$ of $\mathfrak{sl}_2\mathbb{R}$ maps to

$$\rho(\left[\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right]) = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix},$$

so

$$\exp(t\rho(\left[\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right])) = \begin{pmatrix} 1 + \frac{t^2}{2} & t & -\frac{t^2}{2} \\ t & 1 & -t \\ \frac{t^2}{2} & t & 1 - \frac{t^2}{2} \end{pmatrix}.$$

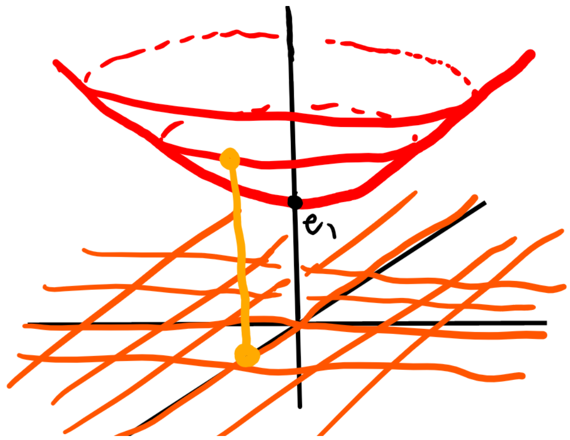

Figure 1. The map $\mathrm{pr} : \mathbb{H}^2 \to \mathbb{R}^2$ identifies the point $ae_1 + be_2 + ce_3 \in Q^{-1}(1)$ with the point $\left[\begin{smallmatrix} b \\ c \end{smallmatrix}\right] \in \mathbb{R}^2$

Applying this to e_1, thought of as the identity coset of $\mathrm{PO}(1,2)/\,\mathrm{O}(2)$, and looking at the image under the projection pr, we get

$$\mathrm{pr}(\exp(t\rho(\left[\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right])) \cdot e_1) = \begin{bmatrix} t \\ t^2/2 \end{bmatrix}.$$

In particular, the orbit of this one-parameter subgroup through e_1 projects to a **parabola**! Indeed, all of its orbits on $\mathbb{H}^2$ project to parabolas!

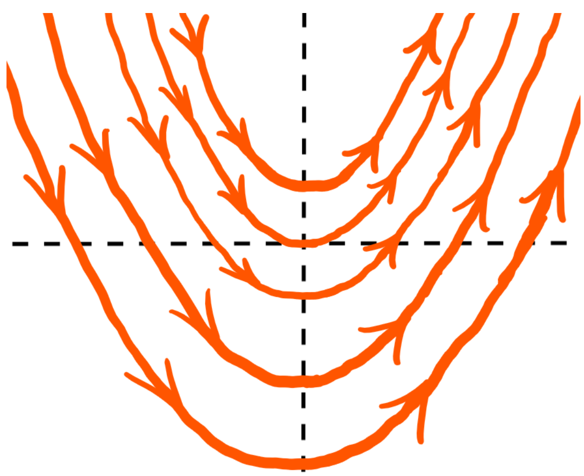

Figure 2. The orbits of the one-parameter subgroup $\exp(t\rho(\left[\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right]))$ all project to parabolas!

More generally, let's consider the element $\frac{1}{2}\left[\begin{smallmatrix} 0 & r+1 \\ r-1 & 0 \end{smallmatrix}\right]$. For $r^2 > 1$, this element has eigenvalues $\lambda = \frac{1}{2}\sqrt{r^2 - 1}$ and $-\lambda$, so it is diagonalizable over

$\mathbb{R}$, hence it is hyperbolic. In this case, we get corresponding one-parameter subgroups

$$\exp\left(\frac{t}{2}\rho(\left[\begin{smallmatrix}0 & r+1\\ r-1 & 0\end{smallmatrix}\right])\right) = \begin{pmatrix} \frac{r^2\cosh(\sqrt{r^2-1}t)-1}{r^2-1} & r\frac{\sinh(\sqrt{r^2-1}t)}{\sqrt{r^2-1}} & -r\frac{\cosh(\sqrt{r^2-1}t)-1}{r^2-1} \\ r\frac{\sinh(\sqrt{r^2-1}t)}{\sqrt{r^2-1}} & \cosh(\sqrt{r^2-1}t) & -\frac{\sinh(\sqrt{r^2-1}t)}{\sqrt{r^2-1}} \\ r\frac{\cosh(\sqrt{r^2-1}t)-1}{r^2-1} & \frac{\sinh(\sqrt{r^2-1}t)}{\sqrt{r^2-1}} & \frac{r^2-\cosh(\sqrt{r^2-1}t)}{r^2-1} \end{pmatrix},$$

and the orbit of this through e_1 projects to (the connected component through 0 of) the **hyperbola** given by

$$\left(y + \frac{r}{r^2-1}\right)^2 - \frac{x^2}{r^2-1} = \left(\frac{r}{r^2-1}\right)^2.$$

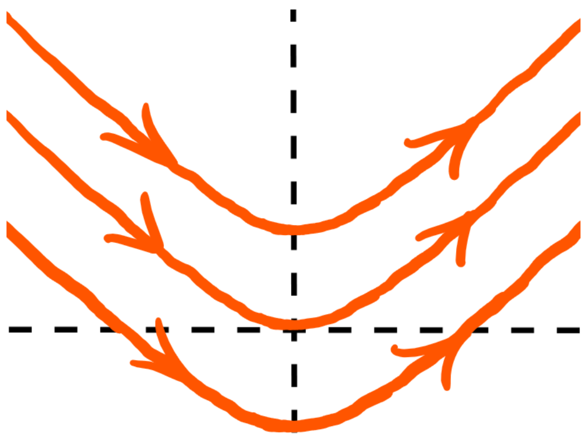

Figure 3. When $r^2 > 1$, the orbits of the one-parameter subgroup $\exp(\frac{t}{2}\rho(\left[\begin{smallmatrix}0 & r+1\\ r-1 & 0\end{smallmatrix}\right]))$ all project to hyperbolas with eccentricity $|r|$

In fact, every orbit of this one-parameter subgroup projects to (a connected component of) a hyperbola with eccentricity[1] $|r|$.

Finally, as you may have guessed by now, the element $\frac{1}{2}\left[\begin{smallmatrix}0 & r+1\\ r-1 & 0\end{smallmatrix}\right]$ is elliptic for $r^2 < 1$. Such elements give us one-parameter subgroups

$$\exp\left(\frac{t}{2}\rho(\left[\begin{smallmatrix}0 & r+1\\ r-1 & 0\end{smallmatrix}\right])\right) = \begin{pmatrix} \frac{1-r^2\cos(\sqrt{1-r^2}t)}{1-r^2} & r\frac{\sin(\sqrt{1-r^2}t)}{\sqrt{1-r^2}} & -r\frac{1-\cos(\sqrt{1-r^2}t)}{1-r^2} \\ r\frac{\sin(\sqrt{1-r^2}t)}{\sqrt{1-r^2}} & \cos(\sqrt{1-r^2}t) & -\frac{\sin(\sqrt{1-r^2}t)}{\sqrt{1-r^2}} \\ r\frac{1-\cos(\sqrt{1-r^2}t)}{1-r^2} & \frac{\sin(\sqrt{1-r^2}t)}{\sqrt{1-r^2}} & \frac{\cos(\sqrt{1-r^2}t)-r^2}{1-r^2} \end{pmatrix},$$

whose orbits project to **ellipses** of eccentricity $|r|$ (except for the orbit that

It's probably worth noting as well, as I remembered the night before the lecture, that elliptic and hyperbolic one-parameter subgroups for $\mathrm{SL}_2\mathbb{R}$ also trace out ellipses and hy-

[1]Recall that the *eccentricity* of a conic is the ratio of the distance of a given point from a point called the "focus" and the distance of that same point from a line called the "directrix". The actual definition isn't really that important though; the significance of the eccentricity is that it completely determines a conic section up to similarity transformations, with ellipses of eccentricity in $[0, 1)$, parabolas of eccentricity 1, and hyperbolas of (finite) eccentricity greater than 1.

consists of the fixed point of the one-parameter subgroup). In particular, for $r \neq 0$, the orbit of $\exp\left(\frac{t}{2}\rho(\left[\begin{smallmatrix} 0 & r+1 \\ r-1 & 0 \end{smallmatrix}\right])\right)$ through e_1 projects to the ellipse determined by the equation

$$\left(y - \frac{r}{1-r^2}\right)^2 + \frac{x^2}{1-r^2} = \left(\frac{r}{1-r^2}\right)^2.$$

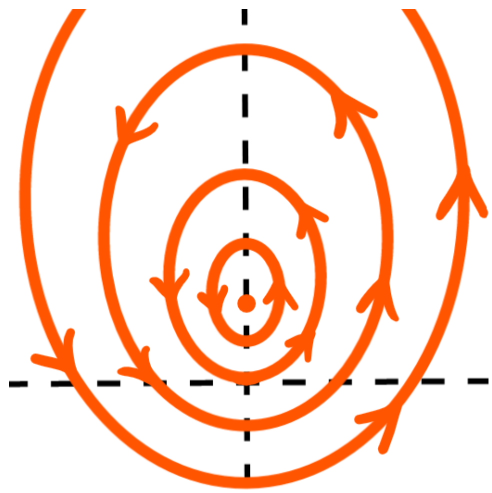

Figure 4. When $r^2 < 1$, the orbits of the one-parameter subgroup $\exp(\frac{t}{2}\rho(\left[\begin{smallmatrix} 0 & r+1 \\ r-1 & 0 \end{smallmatrix}\right]))$ all project to ellipses with eccentricity $|r|$, except for the one orbit corresponding to the fixed point

In general, elements of $\mathfrak{sl}_2\mathbb{R}$ are elliptic, parabolic, or hyperbolic according to whether the orbits of their one-parameter subgroups project to conic sections of eccentricity in $[0, 1)$, equal to 1, or greater than 1, respectively.

8.3. The Killing form on $\mathfrak{sl}_2\mathbb{R}$

Eccentricity gives us a parameter that uniquely determines a conic section up to similarity transformations. We would like something similar for elements of $\mathfrak{sl}_2\mathbb{R}$: a real number that completely characterizes an element of $\mathfrak{sl}_2\mathbb{R}$ up to automorphism. As it turns out, the Killing form gives us such a parameter.

Theorem 8.6. *For nonzero $X, Y \in \mathfrak{sl}_2\mathbb{R}$, there is an automorphism ϕ such that $\phi(X) = Y$ if and only if $\mathfrak{B}(X, X) = \mathfrak{B}(Y, Y)$.*

Proof. This is actually much less daunting than it seems. To start, automorphisms of $\mathfrak{sl}_2\mathbb{R}$ are exactly conjugations by elements of $\mathrm{GL}_2\,\mathbb{R}$, so this is just a fancy way of saying that X and Y are conjugate over $\mathbb{R}$ whenever $\mathfrak{B}(X, X) = \mathfrak{B}(Y, Y)$. To show this, we just find representatives of each conjugacy class and evaluate $\mathfrak{B}$ on them; since $\mathfrak{B}$ is invariant under automorphisms, the choice of representative does not matter.

The Jordan decomposition tells us that every nonzero $X \in \mathfrak{sl}_2\mathbb{R}$ is conjugate over $\mathbb{C}$ to precisely one matrix of the form $\left[\begin{smallmatrix} \lambda & 0 \\ 0 & -\lambda \end{smallmatrix}\right]$, $\left[\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right]$, or $\left[\begin{smallmatrix} i\lambda & 0 \\ 0 & -i\lambda \end{smallmatrix}\right]$ for some $\lambda > 0$. Since real matrices that are conjugate over $\mathbb{C}$ are conjugate over $\mathbb{R}$ and $\left[\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right]$ is conjugate to $\left[\begin{smallmatrix} i & 0 \\ 0 & -i \end{smallmatrix}\right]$ over $\mathbb{C}$, this means that every nonzero $X \in \mathfrak{sl}_2\mathbb{R}$ is conjugate over $\mathbb{R}$ to exactly one element of the form $\left[\begin{smallmatrix} \lambda & 0 \\ 0 & -\lambda \end{smallmatrix}\right]$, $\left[\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right]$, or $\left[\begin{smallmatrix} 0 & -\lambda \\ \lambda & 0 \end{smallmatrix}\right]$ for some $\lambda > 0$.

Thus, because we have $\mathfrak{B}(\left[\begin{smallmatrix} \lambda & 0 \\ 0 & -\lambda \end{smallmatrix}\right], \left[\begin{smallmatrix} \lambda & 0 \\ 0 & -\lambda \end{smallmatrix}\right]) = 8\lambda^2$, $\mathfrak{B}(\left[\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right], \left[\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right]) = 0$, and $\mathfrak{B}(\left[\begin{smallmatrix} 0 & -\lambda \\ \lambda & 0 \end{smallmatrix}\right], \left[\begin{smallmatrix} 0 & -\lambda \\ \lambda & 0 \end{smallmatrix}\right]) = -8\lambda^2$, which never coincide for $\lambda > 0$, each conjugacy class is uniquely determined by the value of the Killing form. □

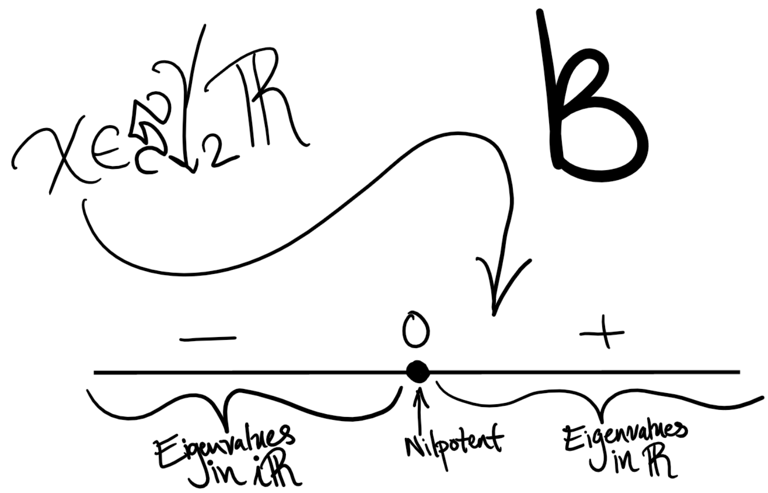

Figure 5. We can imagine $\mathfrak{B}(X, X)$ for $X \in \mathfrak{sl}_2\mathbb{R}$ to be a parameter describing X on a continuum where negative values correspond to diagonalizability over $\mathbb{C}$ with imaginary eigenvalues, 0 corresponds to nilpotence, and positive values correspond to diagonalizability over $\mathbb{R}$

From the above proof, we see that an element $X \in \mathfrak{sl}_2\mathbb{R}$ is elliptic if and only if $\mathfrak{B}(X, X) < 0$, parabolic if and only if $\mathfrak{B}(X, X) = 0$, and hyperbolic if and only if $\mathfrak{B}(X, X) > 0$. In particular, we can imagine $\mathfrak{B}(X, X)$ to be a parameter describing X on a continuum where negative values correspond to diagonalizability over $\mathbb{C}$ with imaginary eigenvalues, 0 corresponds to nilpotence, and positive values correspond to diagonalizability over $\mathbb{R}$.

8.4. What does the Killing form tell us?

For general Lie algebras $\mathfrak{g}$, the Killing form obviously isn't going to completely determine elements up to automorphism the way it does for $\mathfrak{sl}_2\mathbb{R}$. Nevertheless, $\mathfrak{B}$ can still tell us a lot about how elements of $\mathfrak{g}$ behave, if we look at it the right way.

To start, note that $\mathfrak{B}(X, X) = \operatorname{tr}(\operatorname{ad}_X^2)$ is the sum of the squares of the eigenvalues of ad_X. This tells us, in particular, that if we want to understand $\mathfrak{B}$, then we need to look at elements, as well as notions like diagonalizability and nilpotence, from the perspective of the adjoint representation. Going back to the special case of $\mathfrak{sl}_2\mathbb{R}$, for example, we can reclassify elements in terms of the adjoint representation.

Definition 8.7. Suppose $X \in \mathfrak{sl}_2\mathbb{R}$.

- X is *hyperbolic* if and only if ad_X is diagonalizable over $\mathbb{R}$.
- X is *parabolic* if and only if it is ad-nilpotent.
- X is *elliptic* if and only if ad_X is diagonalizable over $\mathbb{C}$ with eigenvalues in $i\mathbb{R}$.

Even in a general real Lie algebra $\mathfrak{g}$, if ad_X is diagonalizable over $\mathbb{R}$, then we will have $\mathfrak{B}(X, X) > 0$. Of course, we won't necessarily get the converse as we do for $\mathfrak{sl}_2\mathbb{R}$, but if $\mathfrak{B}(X, X) > 0$, then we *can* say that the sum of the squares of the real parts of the eigenvalues of ad_X is bigger than the sum of the squares of the imaginary parts. In other words, $\mathfrak{B}(X, X) > 0$ if and only if the real parts of the eigenvalues contribute the most to the behavior of ad_X, in which case we can think of it as "mostly scaling".

Similarly, when $\mathfrak{B}(X, X) < 0$, the sum of the squares of the imaginary parts of the eigenvalues of ad_X is more than the sum of the squares of the real parts. In particular, if ad_X is diagonalizable over $\mathbb{C}$ with all eigenvalues in $i\mathbb{R}$, then $\mathfrak{B}(X, X) < 0$, and while we can't get a true converse in general, we can think of ad_X in this case as being "mostly rotation". Notably, we will have $\mathfrak{B}(X, X) < 0$ whenever X comes from a subalgebra corresponding to a compact subgroup and X isn't central.

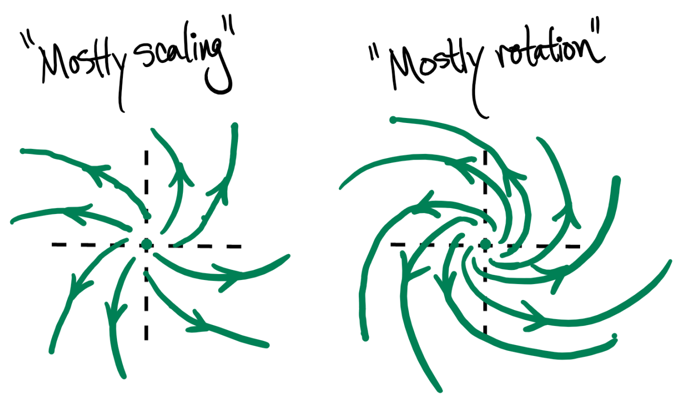

Figure 6. Attempted illustrations of the terms "mostly scaling" and "mostly rotation"

Finally, ad-nilpotent elements X will satisfy $\mathfrak{B}(X,X)=0$. Again, unlike in the case of $\mathfrak{sl}_2\mathbb{R}$, $\mathfrak{B}(X,X)=0$ doesn't necessarily guarantee that ad_X is nilpotent, but it does mean that the sum of the squares of the eigenvalues of ad_X is 0. We can kind of think of this as meaning that "the compact and scaling parts of ad_X cancel out".

Exercise. Suppose K is a compact Lie group. What can we say about the Killing form on the Lie algebra $\mathfrak{k}$? (Note that K could have nontrivial center.)

Exercise. Suppose N is a nilpotent Lie group. What can we say about the Killing form on the Lie algebra $\mathfrak{n}$?

Exercise. Using that elements of the ideal of translations in $\mathfrak{i}(2)$ are ad-nilpotent, describe the Killing form on $\mathfrak{i}(2)$ without performing any computations.

Of course, we would also like to be able to say things about $\mathfrak{B}(X,Y)$ for $X \neq Y$. Using polarization,

$$\mathfrak{B}(X,Y)=\frac{1}{2}(\mathfrak{B}(X+Y,X+Y)-\mathfrak{B}(X,X)-\mathfrak{B}(Y,Y)).$$

This is particularly useful when X and Y are ad-nilpotent, in which case $\mathfrak{B}(X,Y)=\frac{1}{2}\mathfrak{B}(X+Y,X+Y)$. Thus, for ad-nilpotent X and Y, $\mathfrak{B}(X,Y)>0$ when ad_{X+Y} is "mostly scaling" and $\mathfrak{B}(X,Y)<0$ when ad_{X+Y} is "mostly rotation". For example, $\mathfrak{B}(\left[\begin{smallmatrix}0&1\\0&0\end{smallmatrix}\right],\left[\begin{smallmatrix}0&0\\1&0\end{smallmatrix}\right])>0$ because $\left[\begin{smallmatrix}0&1\\0&0\end{smallmatrix}\right]+\left[\begin{smallmatrix}0&0\\1&0\end{smallmatrix}\right]=\left[\begin{smallmatrix}0&1\\1&0\end{smallmatrix}\right]$ is "mostly scaling", and $\mathfrak{B}(\left[\begin{smallmatrix}0&-1\\0&0\end{smallmatrix}\right],\left[\begin{smallmatrix}0&0\\1&0\end{smallmatrix}\right])<0$ because $\left[\begin{smallmatrix}0&-1\\1&0\end{smallmatrix}\right]$ is "mostly rotation".

Next time, we will focus on $\mathfrak{g}$ where $\mathfrak{B}$ is nondegenerate, in which case we say that $\mathfrak{g}$ is *semisimple*. For semisimple Lie algebras, the behavior of $\mathfrak{B}$ described above suggests a particular form for $\mathfrak{g}$: there should be an ad-diagonalizable part together with ad-nilpotent elements occurring in pairs on which the Killing form is nonzero.

Chapter 9

Noncompact Riemannian Symmetric Spaces

Again, rewrite this so that it looks better and fits well within the book. Specifically, discuss the maximal compact subgroup stuff a bit more? Also, move the stuff about parabolic subgroups to the next chapter? Alternatively, consider keeping it here and reframing the next chapter to just cover the filtration stuff?

Whenever G is a Lie group with finitely many connected components, a fundamental result (see, for example, Theorem 14.1.3 of [**4**]) in the general structure theory of Lie groups tells us that G has a maximal compact subgroup, and that all maximal compact subgroups are conjugate to each other. Moreover, maximal compact subgroups contain all of the nontrivial aspects of the topology of G: for $K \leq G$ a maximal compact subgroup, G is diffeomorphic—though generally not isomorphic—to the product of K with a vector space.

We saw this, for example, in Euclidean geometry: $\mathrm{I}(2)$, viewed as the orthonormal frame bundle over $\mathbb{R}^2$, was clearly diffeomorphic to $\mathbb{R}^2 \times \mathrm{O}(2)$; indeed, it was isomorphic to $\mathbb{R}^2 \rtimes \mathrm{O}(2)$. Here, $\mathrm{O}(2)$ is a maximal compact subgroup of $\mathrm{I}(2)$, and all other maximal compact subgroups of $\mathrm{I}(2)$ are conjugate to $\mathrm{O}(2)$.

As we said earlier, if we want to find geometrically interesting models, then it makes sense to look for Lie-theoretically interesting models. Choosing our isotropy to be a maximal compact subgroup K is amongst the most Lie-theoretically interesting choices we can make in this case, and if, moreover, we choose our model group G to be a semisimple Lie group, then the underlying geometry is, as we would expect, remarkably deep. Such models (G, K) correspond to Riemannian symmetric spaces of noncompact type, and the Killing form gives several key tools for studying them, including:

- A $\mathfrak{B}$-orthogonal decomposition of the Lie algebra of the model group, called the *Cartan decomposition*
- A canonical Riemannian metric and notion of distance on G/K
- A convenient description of the stabilizers of "points at infinity"

While this doesn't directly help us understand the Killing form unless we already have experience with symmetric spaces, it does let us visualize several important interactions between the Killing form and the underlying representation theory. In particular, it will give us insight into what parabolic subgroups look like, and in the next lecture, we will explore our algebraic definition of parabolic subgroups and connect it to the more immediately geometric idea of "points at infinity".

9.1. Riemannian symmetric spaces of noncompact type

For the rest of the lecture, let us fix a model (G, K), where G is a semisimple Lie group with finitely many connected components such that the identity component G° has finite center and K is a maximal compact subgroup. In this case, the notion of *Killing perpendiculars* gives us a very convenient description of the topological decomposition of G as a product of K with a vector space.

Definition 9.1. For a subspace $V \subseteq \mathfrak{g}$, its *Killing perpendicular* (or *Killing perp*) is the subspace

$$V^\perp := \{X \in \mathfrak{g} : \mathfrak{B}(X, v) = 0 \text{ for each } v \in V\}.$$

In our case, the Lie algebra $\mathfrak{g}$ decomposes, as a vector space, as a $\mathfrak{B}$-orthogonal direct sum $\mathfrak{k}^\perp \oplus \mathfrak{k}$, where the subspace $\mathfrak{k}^\perp$ is the Killing perp of the Lie subalgebra $\mathfrak{k}$ corresponding to K. The exponential map restricts to an *embedding* on $\mathfrak{k}^\perp$, so that $\exp(\mathfrak{k}^\perp)$ is diffeomorphic to $\mathfrak{k}^\perp$, and moreover, the map $\mu : \exp(\mathfrak{k}^\perp) \times K \to G$ given by applying the group operation $(\exp(X), k) \mapsto \exp(X)k$ is a diffeomorphism. In particular, the usual quotient map $q_K : G \to G/K$ restricts to a diffeomorphism from $\exp(\mathfrak{k}^\perp)$ to G/K, and we get a projection map

$$\mathrm{pr} : G \to \exp(\mathfrak{k}^\perp),$$

which induces a section of q_K.

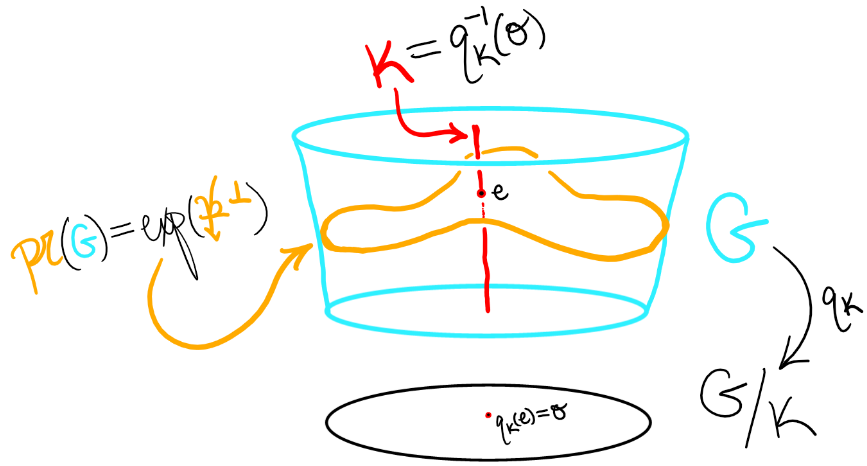

Figure 1. The projection $\mathrm{pr} : G \to \exp(\mathfrak{k}^\perp)$ gives a section to the natural quotient map $q_K : G \to G/K$

The decompositions $\mathfrak{g} = \mathfrak{k}^\perp \oplus \mathfrak{k}$ for the Lie algebra and $G = \exp(\mathfrak{k}^\perp)K$ for the Lie group are both called the *Cartan decomposition* corresponding to K. At the level of Lie algebras, we can think of this as a decomposition into symmetric and skew-symmetric elements.

We've actually seen this before with hyperbolic geometry. In that case, our model (G, K) had $G = \mathrm{PO}(1, n)$ and $K \simeq \mathrm{O}(n)$, and we had a nice projection map that we used to give a diffeomorphism from $\mathrm{PO}(1, n)/\mathrm{O}(n) = \mathbb{H}^n$ to $\mathbb{R}^n$ after recognizing a subspace

$$\mathfrak{k}^\perp = \left\{ \begin{pmatrix} 0 & v^\top \\ v & 0 \end{pmatrix} : v \in \mathbb{R}^n \right\}$$

vaguely analogous to translations in the Euclidean case. It turns out that the symmetric space projection map $\mathrm{pr} : \mathrm{PO}(1, n) \to \exp(\mathfrak{k}^\perp)$ for hyperbolic geometry is given by

$$\begin{pmatrix} a & \alpha \\ x & R \end{pmatrix} \mapsto \begin{pmatrix} a & x^\top \\ x & \mathbb{1} + \frac{1}{1+a} x x^\top \end{pmatrix},$$

and the image of this projection is uniquely determined by $x \in \mathbb{R}^n$. After identifying $\mathrm{pr}(\left(\begin{smallmatrix} a & \alpha \\ x & R \end{smallmatrix}\right))$ with x, the induced map from $\mathrm{PO}(1, n)/\mathrm{O}(n)$ to $\mathbb{R}^n$ happens to coincide with that nice projection map that we used to identify $\mathbb{H}^n \cong \mathbb{R}^n \cong \exp(\mathfrak{k}^\perp)$.

For Cartan decompositions, our heuristic for the Killing form works exactly as expected: on the subalgebra $\mathfrak{k}$ corresponding to the maximal compact subgroup, $\mathscr{B}$ is negative-definite, and on the subspace $\mathfrak{k}^\perp$, whose elements generate scaling transformations in the adjoint representation, $\mathfrak{B}$ is positive-definite. This gives us an easy way of describing the pushforward projection $\mathrm{pr}_* : \mathfrak{g} \to \mathfrak{k}^\perp$ at the identity: for $X \in \mathfrak{g}$, the projection $\mathrm{pr}_*(X) \in \mathfrak{k}^\perp$ is the element $X' \in X + \mathfrak{k}$ for which $\mathfrak{B}(X', X')$ is maximal. Moreover,

$$\mathfrak{B}(X,Y) = \mathfrak{B}(\mathrm{pr}_*(X), \mathrm{pr}_*(Y)) + \mathfrak{B}(X - \mathrm{pr}_*(X), Y - \mathrm{pr}_*(Y)),$$

so we can genuinely decompose $\mathfrak{B}(X,X)$ as the sum of the "scaling part" and the "compact part", and for $X \in \mathfrak{k}^\perp$, $\mathfrak{B}(X,Y) = \mathfrak{B}(X, \mathrm{pr}_*(Y))$.

Since pr_* induces an isomorphism between $\mathfrak{g}/\mathfrak{k}$ and $\mathfrak{k}^\perp$, we can identify the tangent bundle $T(G/K) \cong G \times_K \mathfrak{g}/\mathfrak{k}$ with the homogeneous vector bundle $G \times_K \mathfrak{k}^\perp$. This isomorphism also gives us a canonical choice of Riemannian metric on G/K: $\mathfrak{B}$ is positive-definite on $\mathfrak{k}^\perp$, so for $X, Y \in T_{gK}(G/K)$, we can define a Riemannian metric $\mathrm{pr}^*\mathfrak{B}$ by

$$\mathrm{pr}^*\mathfrak{B}(X,Y) := \mathfrak{B}(\mathrm{pr}_*(\mathrm{L}_{g^{-1}*}X), \mathrm{pr}_*(\mathrm{L}_{g^{-1}*}Y)).$$

By construction, this is invariant under the canonical left-action of G, so it is a geometric object for the model.

Of course, for Riemannian manifolds, we get an associated notion of geodesic. As we did before with Euclidean geometry and hyperbolic geometry, though, we'll define geodesics in terms of motion rather than the Riemannian metric. Specifically, we can think of $\mathfrak{k}^\perp$ as being analogous to the subspace of translations in Euclidean geometry, and we define geodesics as (projections of) left-translates of one-parameter subgroups generated by elements of $\mathfrak{k}^\perp$.

Definition 9.2. A *geodesic* for the model geometry (G, K) is a curve $\gamma : \mathbb{R} \to G/K$ of the form $t \mapsto q_K(g \exp(tX))$ for some $g \in G$ and $X \in \mathfrak{k}^\perp$.

As before, this corresponds to starting at some configuration $g \in G$, picking a velocity $X \in \mathfrak{k}^\perp$, and at each point in time moving with the velocity that the Maurer-Cartan form identifies with X, so that we move with "constant velocity"; by construction, every left-translate of a geodesic is again a geodesic, so geodesics are geometric for (G, K). In this case, geodesics in our sense coincide with geodesics in the Riemannian sense.

This is, of course, not a very thorough introduction to the topic of Riemannian symmetric spaces of noncompact type. For such an introduction, we highly recommend [**3**].

9.2. Asymptotic behavior of geodesics

Another concept that makes sense for Riemannian manifolds is the *distance* between two points. Indeed, if we know the projection map $\mathrm{pr} : G \to \exp(\mathfrak{k}^\perp)$, then distance is fairly straightforward to find in this case:

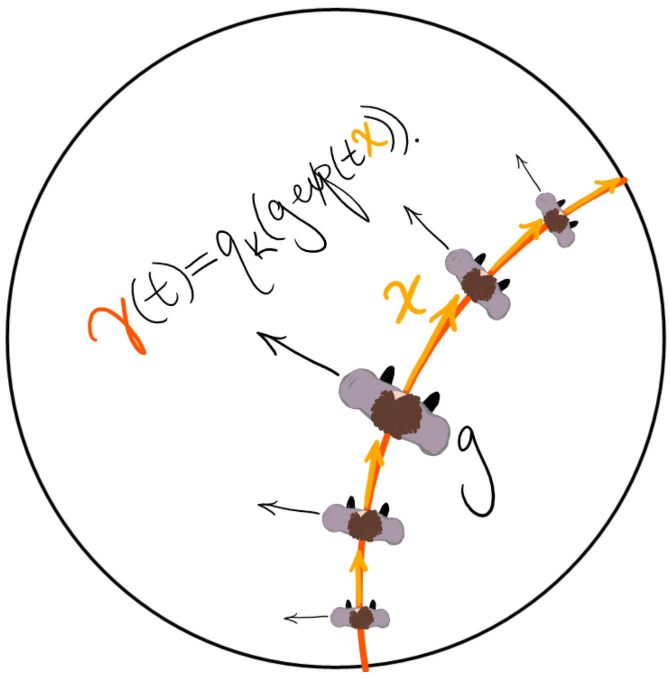

Figure 2. Geodesic motion corresponds to starting at some configuration in G, then moving with ω_G-constant velocity in $\mathfrak{k}^\perp$

for elements $g_0, g_1 \in G$, there is a unique $X \in \mathfrak{k}^\perp$ such that $\exp(X) = \mathrm{pr}(g_0^{-1}g_1)$, and the distance $\mathrm{dist}(q_K(g_0), q_K(g_1))$ from $q_K(g_0)$ to $q_K(g_1)$ is just $\sqrt{\mathfrak{B}(X,X)}$.

For us, the main use for this is to describe the asymptotic behavior of geodesics, since this will lead us to parabolic subgroups.

Definition 9.3. Suppose γ_1 and γ_2 are unit-speed geodesics in G/K. We say that γ_1 and γ_2 are *asymptotic* if and only if the distance $\mathrm{dist}(\gamma_1(t), \gamma_2(t))$ is bounded for $t \geq 0$.

This defines an equivalence relation on geodesics, and an equivalence class of asymptotic geodesics is called a *point at infinity*.

Definition 9.4. An equivalence class of asymptotic geodesics is called a *point at infinity*. For a geodesic γ, we denote its corresponding point at infinity by $\gamma(+\infty)$.

Topologically, we can identify the space of all points at infinity with the unit sphere in $\mathfrak{k}^\perp$, since each $Z \in \mathfrak{k}^\perp$ with $\mathfrak{B}(Z,Z) = 1$ uniquely determines a unit-speed geodesic of the form $t \mapsto q_K(\exp(tZ))$, and every point at infinity corresponds to exactly one such geodesic.

Figure 3. Several asymptotic geodesics and their corresponding point at infinity

9.3. Prelude to parabolic subgroups

In the trichotomy for elements of $\mathfrak{sl}_2\mathbb{R}$ in terms of conic sections from last time, a considerably more well-known characterization of parabolic transformations is as transformations that fix a single point at infinity for the hyperbolic plane. With this in mind, it *almost* wouldn't be ridiculous to call the stabilizer subgroup of a point at infinity, or more generally a finite-index subgroup of such a stabilizer, a *parabolic subgroup*.

While this does give a mostly valid[1] definition for parabolic subgroups, it would be kind of annoying to use in practice. Imagine we found a closed subgroup of G and we wanted to check whether it was parabolic; without more information, we'd basically have to start checking geodesics to see whether their asymptotic behavior was preserved by our subgroup. We'd like a more direct definition, preferably one that comes from the structure of the Lie algebra.

In an attempt to ascertain such a definition, let's start with a point $\gamma(+\infty)$ at infinity and try to find its stabilizer. As we mentioned above, we may assume that our geodesic γ is of the form $t \mapsto q_K(\exp(tZ))$ for some $Z \in \mathfrak{k}^\perp$.

An element $g \in G$ fixes $\gamma(+\infty)$ if and only if $g\gamma$ is asymptotic to γ. Thus, we want to find $g \in G$ such that

$$g\gamma(t) = q_K(g\exp(tZ)) = q_K\left(\exp(tZ)(\exp(tZ)^{-1}g\exp(tZ))\right)$$

is a bounded distance away from $\gamma(t) = q_K(\exp(tZ))$ for all $t \geq 0$, which amounts to showing that $\mathrm{pr}(\exp(tZ)^{-1}g\exp(tZ))$ is bounded. In particular,

[1] It should be noted that some larger subgroups, such as G itself, would be considered parabolic by most representation theorists, even though they don't fix a point at infinity. Our algebraic definition below accounts for this.

for $g = \exp(X)$ for some $X \in \mathfrak{g}$, we have

$$\begin{aligned}\exp(tZ)^{-1} g \exp(tZ) &= \exp(tZ)^{-1} \exp(X) \exp(tZ) \\ &= \exp(\mathrm{Ad}_{\exp(tZ)^{-1}}(X)),\end{aligned}$$

so we want to find $X \in \mathfrak{g}$ such that $\mathrm{pr}_*(\mathrm{Ad}_{\exp(tZ)^{-1}}(X))$ is bounded for $t \geq 0$.

Because ad_Z is diagonalizable over $\mathbb{R}$, we can decompose $\mathfrak{g}$, as a vector space, as $\mathfrak{g}_- \oplus \mathfrak{g}_0 \oplus \mathfrak{p}_+$, where $\mathfrak{g}_-$ is the sum of all the negative eigenspaces for ad_Z, $\mathfrak{g}_0$ is the centralizer $\mathfrak{z}_{\mathfrak{g}}(Z)$ of Z, and $\mathfrak{p}_+$ is the sum of all the positive eigenspaces for ad_Z. Equivalently, we could define $\mathfrak{g}_-$ as

$$\mathfrak{g}_- := \left\{X \in \mathfrak{g} : \mathrm{Ad}_{\exp(tZ)}(X) \to 0 \text{ as } t \to +\infty\right\}$$

and $\mathfrak{p}_+$ as

$$\mathfrak{p}_+ := \left\{X \in \mathfrak{g} : \mathrm{Ad}_{\exp(tZ)^{-1}}(X) \to 0 \text{ as } t \to +\infty\right\}.$$

Because $\mathrm{Ad}_{\exp(tZ)}$ is an automorphism, both $\mathfrak{g}_-$ and $\mathfrak{p}_+$ are subalgebras, which homogeneous dynamicists would call the *contracting* and *expanding horospherical subalgebras* of Z, respectively. (See, for example, the excellent book [**7**].)

Crucially, note that $\mathrm{Ad}_g(X)$ has the same eigenvalues in the adjoint representation as X, so elements of the expanding and contracting horospherical subalgebras only have 0 as an eigenvalue under the adjoint representation, which means that elements of these subalgebras are always ad-nilpotent. In particular, $\mathfrak{g}_-$ and $\mathfrak{p}_+$ are always nilpotent subalgebras of $\mathfrak{g}$.

Writing $X = X_- + X_0 + X_+$, with $X_- \in \mathfrak{g}_-$, $X_0 \in \mathfrak{g}_0$, and $X_+ \in \mathfrak{p}_+$, we see that

$$\mathrm{pr}_*(\mathrm{Ad}_{\exp(tZ)^{-1}}(X)) = \mathrm{pr}_*(\mathrm{Ad}_{\exp(tZ)^{-1}}(X_-) + X_0 + \mathrm{Ad}_{\exp(tZ)^{-1}}(X_+))$$

is bounded for all $t \geq 0$ if and only if $X_- = 0$. Thus, the Lie subalgebra of the stabilizer subgroup for $\gamma(+\infty)$ is precisely $\mathfrak{p} := \mathfrak{g}_0 + \mathfrak{p}_+$.

Before moving on, it's well-worth trying to visualize this decomposition, since it will be very important from here onward.

Let us once again imagine ourselves as observers in the model group G, moving geodesically using right-translation by $\exp(tZ)$. At each configuration in G, we can use the Maurer-Cartan form ω_G to decompose the tangent spaces according to the decomposition $\mathfrak{g}_- + \mathfrak{g}_0 + \mathfrak{p}_+$ for $\mathfrak{g}$. Since $\mathfrak{g}_-$, $\mathfrak{g}_0$, and $\mathfrak{p}_+$ are subalgebras, the corresponding distributions are integrable, with integral submanifolds corresponding to the left-cosets of the connected subgroups generated by each subalgebra.

Let us denote by G_- and P_+ the connected subgroups generated by $\mathfrak{g}_-$ and $\mathfrak{p}_+$, respectively. Then, the integral submanifold for $\omega_G^{-1}(\mathfrak{g}_-)$ through $g \in G$ is precisely gG_-, and similarly, gP_+ is the integral submanifold for

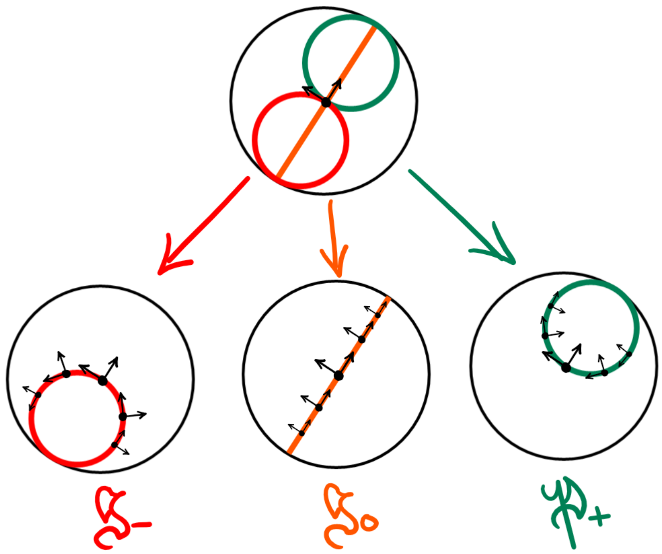

Figure 4. Using the Maurer-Cartan form ω_G, we can decompose tangent spaces of G as sums of the integrable distributions $\omega_G^{-1}(\mathfrak{g}_-)$, $\omega_G^{-1}(\mathfrak{g}_0)$, and $\omega_G^{-1}(\mathfrak{p}_+)$

$\omega_G^{-1}(\mathfrak{p}_+)$ through g. As one might imagine from the term "horospherical subalgebra", these left-cosets for G_- and P_+ project to horospheres under the quotient map q_K.

Consider a starting configuration $g \in G$ and an element $p \in P_+$, so that g and gp lie on the same integral submanifold for $\omega_G^{-1}(\mathfrak{p}_+)$. Then, moving by $\exp(tZ)$ at both these points, g goes to $g\exp(tZ)$ and gp goes to $gp\exp(tZ) = g\exp(tZ)(\exp(tZ)^{-1}p\exp(tZ))$. Essentially by definition of P_+, $\exp(tZ)^{-1}p\exp(tZ)$ will converge to the identity element as $t \to +\infty$, so $g\exp(tZ)$ and $gp\exp(tZ)$ get closer and closer together for larger and larger t. In other words, motion by $\exp(tZ)$ contracts the leaves gP_+.

Similarly, motion by $\exp(tZ)$ expands the leaves gG_- of the distribution $\omega_G^{-1}(\mathfrak{g}_-)$. We call the foliation from the distribution $\omega_G^{-1}(\mathfrak{p}_+)$ the *stable foliation* for $\mathrm{R}_{\exp(tZ)}$ and the foliation from $\omega_G^{-1}(\mathfrak{g}_-)$ the *unstable foliation* for $\mathrm{R}_{\exp(tZ)}$.

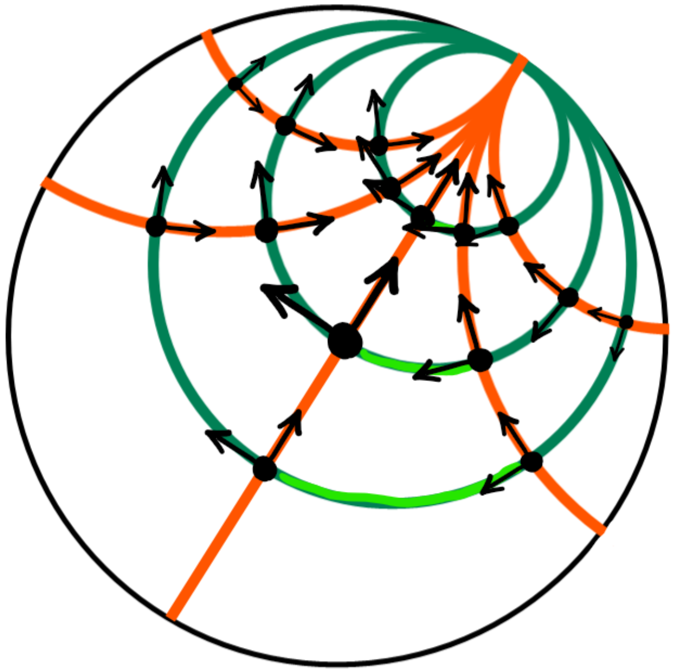

Figure 5. Moving by $\exp(tZ)$ contracts the leaves of the stable foliation generated by $\omega_G^{-1}(\mathfrak{p}_+)$

We should remark that homogeneous dynamicists are typically interested in the behavior of elements of G as *transformations*. Since we want to consider elements of G in terms of *motions* here—acting on the right so that we preserve left-invariance—the roles of the expanding and contracting horospherical subgroups are reversed: the left-cosets of the "expanding" horospherical subgroup P_+ are contracted by moving by $\exp(tZ)$, and the left-cosets of the "contracting" horospherical subgroup G_- are expanded.

Since Z is centralized by $\mathfrak{g}_0$, motion by $\exp(tZ)$ doesn't affect the distribution $\omega_G^{-1}(\mathfrak{g}_0)$: for every $X \in \mathfrak{g}_0$, $\mathrm{R}_{\exp(tZ)*}\,\omega_G^{-1}(X) = \omega_G^{-1}(X)$. We call the foliation generated by $\omega_G^{-1}(\mathfrak{g}_0)$ the *neutral foliation*. Since Z obviously centralizes itself, the leaf of this foliation through $g \in G$ will contain the full geodesic trajectory $g\exp(\mathbb{R}Z)$ of g. This allows us to imagine these leaves as "tubes" of asymptotic geodesic trajectories.

For X_1 and X_2 eigenvectors of ad_Z with respective eigenvalues λ_1 and λ_2, we have

$$0 = \mathfrak{B}(\mathrm{ad}_Z(X_1), X_2) + \mathfrak{B}(X_1, \mathrm{ad}_Z(X_2)) = (\lambda_1 + \lambda_2)\mathfrak{B}(X_1, X_2),$$

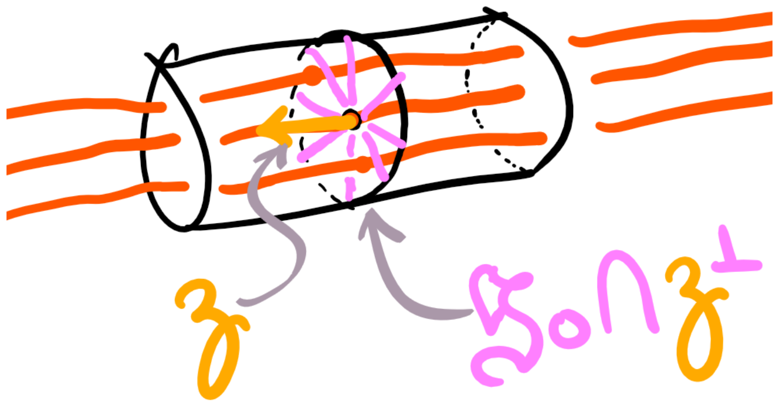

Figure 6. Each leaf of the neutral foliation generated by $\omega_G^{-1}(\mathfrak{g}_0)$ consists of a tube of asymptotic geodesic trajectories of the form $g\exp(\mathbb{R}Z)$

so $\mathfrak{b}(X_1, X_2) = 0$ unless $\lambda_1 + \lambda_2 = 0$. In particular, since $\mathfrak{p}_+$ is the sum of the positive eigenspaces, this tells us that $\mathfrak{p}_+$ is $\mathfrak{b}$-orthogonal to both itself and $\mathfrak{g}_0$, and similarly, $\mathfrak{g}_-$ is $\mathfrak{b}$-orthogonal to both itself and $\mathfrak{g}_0$. Moreover, because $\mathfrak{b}$ is nondegenerate, this also tells us that the eigenvalues of ad_Z must occur in pairs $\pm\lambda$, with the eigenspace for λ dual to the eigenspace for $-\lambda$ with respect to $\mathfrak{b}$, and $\mathfrak{b}$ must remain nondegenerate when restricted to the 0-eigenspace $\mathfrak{g}_0$.

To summarize the picture, we have a tube of asymptotic geodesic trajectories around each configuration $g \in G$, corresponding to the leaf of the neutral foliation through g, together with two left-cosets gP_+ and gG_-, corresponding to leaves of the stable and unstable foliation respectively, that are $\mathfrak{b}$-orthogonal to the tube. Under the natural quotient map q_K, these left-cosets project to horospheres, with $q_K(gP_+)$ a horosphere "centered" at the point at infinity given by following $t \mapsto q_K(g\exp(tZ))$ as $t \to +\infty$ and $q_K(gG_-)$ a horosphere "centered" at the point at infinity given by following $t \mapsto q_K(g\exp(tZ))$ as $t \to -\infty$. These horospheres are tangent at $q_K(g)$, and both are transverse to the image of the leaf of the neutral foliation.

In the semisimple case, we can get a lot of useful intuition for $\mathfrak{b}$ from the duality between $\mathfrak{p}_+$ and $\mathfrak{g}_-$. For each eigenspace $\mathfrak{g}_\lambda$ of ad_Z with positive eigenvalue λ, there is another eigenspace $\mathfrak{g}_{-\lambda}$ with negative eigenvalue $-\lambda$, and they pair together under the Killing form. In the picture above, $\omega_G^{-1}(\mathfrak{g}_\lambda)$ is tangent to the stable foliation and $\omega_G^{-1}(\mathfrak{g}_{-\lambda})$ is tangent to the unstable foliation, and they project to the same subspace of the tangent space under the natural quotient map q_K, so the pairing can sort of be seen from the canonical Riemannian metric being positive-definite.

For us, though, the crucial takeaway from this duality is what it tells us about the horospherical subalgebra $\mathfrak{p}_+$. We've already seen that $\mathfrak{p}_+$ is $\mathfrak{B}$-orthogonal to both itself and $\mathfrak{g}_0$. Because each element $Y \in \mathfrak{p}_+$ has an element $X \in \mathfrak{g}_-$ for which $\mathfrak{B}(X, Y) \neq 0$, this then tells us that $\mathfrak{p}_+^\perp$ is precisely the Lie subalgebra $\mathfrak{p}$ of the stabilizer of the point at infinity. By nondegeneracy of $\mathfrak{B}$, we therefore have $\mathfrak{p}^\perp = \mathfrak{p}_+$.

We should think of the existence of these nilpotent, horospherical subalgebras $\mathfrak{p}_+$ that are $\mathfrak{B}$-orthogonal to all of $\mathfrak{p}$ as the defining characteristic of parabolic subalgebras. Indeed, their existence is precisely the property that we will use to define parabolicity.[2]

Definition 9.5. A subalgebra $\mathfrak{p} \leq \mathfrak{g}$ is *parabolic* if and only if $\mathfrak{p}^\perp$ is a nilpotent subalgebra. A *parabolic subgroup* $P \leq G$, then, is a closed subgroup whose Lie subalgebra $\mathfrak{p}$ is parabolic.

Next time, we will take this definition and attempt to build some useful tools for working with parabolic subgroups. Additionally, we will construct a fixed point at infinity for each (proper) parabolic subgroup.

[2]Since this is a somewhat less well-known definition of parabolicity, I should note that I essentially learned this while perusing [**1**], which attributes it to Grothendieck and Burstall in papers that I was unable to find.

Chapter 10

What is a Parabolic Subgroup?

Again, rewrite this so that it looks better and fits well within the book. Also, consider whether introducing parabolic subgroups in the previous chapter and then just doing the structure theory here is better?

Let us first recall the definition of parabolic subgroups from last time.

Definition 10.1. For a subspace $V \subseteq \mathfrak{g}$, its *Killing perp* is the subspace

$$V^{\perp} := \{X \in \mathfrak{g} : \mathfrak{B}(X, v) = 0 \text{ for each } v \in V\}.$$

Definition 10.2. A subalgebra $\mathfrak{p} \leq \mathfrak{g}$ is *parabolic* if and only if $\mathfrak{p}^{\perp}$ is a nilpotent subalgebra. A *parabolic subgroup* $P \leq G$, then, is a closed subgroup whose Lie subalgebra $\mathfrak{p}$ is parabolic.

In the last lecture, we spent considerable effort to introduce and motivate these parabolic subgroups in a directly geometric way, as (finite-index subgroups of) stabilizers of points at infinity for a model (G, K). Toward the end, we showed that the Lie algebras of such stabilizers satisfy the above algebraic condition. This time, we will verify that these notions of parabolicity are essentially the same. Along the way, we will introduce some incredibly useful tools from representation theory, including:

- A filtration of a semisimple Lie algebra $\mathfrak{g}$ canonically associated to a parabolic subalgebra $\mathfrak{p}$
- An automorphism θ, called a *Cartan involution*, that swaps horospherical subalgebras

- A grading of a semisimple Lie algebra $\mathfrak{g}$ underlying the canonical filtration

Next time, we will see how these tools interact with the geometry of a model (G, P), where G is semisimple and P is parabolic. In particular, we will be able to get a vague picture of the shape of a general parabolic model geometry.

10.1. A few examples

For a bit of amusement, it is perhaps worth noting that $\mathrm{O}(2) < \mathrm{I}(2)$ technically satisfies our definition of parabolicity, since $\mathfrak{o}(2)^\perp$ is the abelian (hence nilpotent) subalgebra of translations. However, while we *can* define parabolic subgroups for arbitrary Lie groups, most would consider the idea of parabolic subgroups to be specific to semisimple Lie groups, for which the Killing form is nondegenerate. Henceforth, we will focus on semisimple model groups.

In $\mathrm{SL}_2\,\mathbb{R}$, recall that we had a closed subgroup B, which we called a *Borel subgroup*[1], defined by

$$B := \left\{ \begin{bmatrix} a & b \\ 0 & a^{-1} \end{bmatrix} : a \in \mathbb{R}^\times, b \in \mathbb{R} \right\},$$

with Lie subalgebra

$$\mathfrak{b} := \left\{ \begin{bmatrix} a & b \\ 0 & -a \end{bmatrix} : a, b \in \mathbb{R} \right\}.$$

Further, recall that the Killing form on $\mathfrak{sl}_2\mathbb{R}$ is given by

$$\mathfrak{Z}\left(\left[\begin{smallmatrix} a & b \\ c & -a \end{smallmatrix}\right], \left[\begin{smallmatrix} z & y \\ x & -z \end{smallmatrix}\right]\right) = 8az + 4(bx + cy).$$

Thus,

$$\mathfrak{Z}\left(\left[\begin{smallmatrix} a & b \\ 0 & -a \end{smallmatrix}\right], \left[\begin{smallmatrix} z & y \\ x & -z \end{smallmatrix}\right]\right) = 8az + 4bx,$$

which vanishes for all $\left[\begin{smallmatrix} a & b \\ 0 & -a \end{smallmatrix}\right] \in \mathfrak{b}$ if and only if $x = z = 0$; in other words, $\mathfrak{b}^\perp = \langle \left[\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right] \rangle =: \mathfrak{b}_+$. Since a 1-dimensional subalgebra is necessarily abelian, this shows that $\mathfrak{b}^\perp$ is a nilpotent subalgebra, hence $\mathfrak{b}$ is parabolic.

Similarly, in $\mathrm{SL}_3\,\mathbb{R}$, we can define a Borel subgroup B as

$$B := \left\{ \begin{bmatrix} r & p & q \\ 0 & s & u \\ 0 & 0 & (rs)^{-1} \end{bmatrix} : \ r, s \in \mathbb{R}^\times, p, q, u \in \mathbb{R} \right\},$$

[1] For a real semisimple Lie group, the term "Borel subgroup" refers to either an arbitrary minimal parabolic subgroup or a specific type of minimal parabolic subgroup that complexifies in a particularly nice way. I used to believe that the former was the better interpretation, but I'm reading more stuff by representation theorists working over $\mathbb{Q}$ and $\mathbb{C}$, and now I'm not sure. Here, the usage of the term is correct regardless of which definition we choose.

the subgroup of upper triangular matrices, with corresponding Lie subalgebra

$$\mathfrak{b} := \left\{ \begin{bmatrix} r & p & q \\ 0 & s & u \\ 0 & 0 & -(r+s) \end{bmatrix} : \ r,s,p,q,u \in \mathbb{R} \right\}.$$

The Killing form on $\mathfrak{sl}_3\mathbb{R}$ is given by $\mathfrak{Z}(R,S) = 6\,\mathrm{tr}(RS)$, where the elements $R,S \in \mathfrak{sl}_3\mathbb{R}$ are considered as linear endomorphisms of $\mathbb{R}^3$ under the "usual" representation of $\mathrm{SL}_3\,\mathbb{R}$. Thus,

$$\begin{aligned} \mathfrak{Z}\left(\left[\begin{smallmatrix} r & p & q \\ 0 & s & u \\ 0 & 0 & -(r+s) \end{smallmatrix}\right], \left[\begin{smallmatrix} m & a & b \\ x & n & c \\ z & y & -(m+n) \end{smallmatrix}\right]\right) &= 6\Big(rm + px + qz + sn + uy \\ &\qquad + (r+s)(m+n)\Big) \\ &= 6\Big(r(2m+n) + s(m+2n) \\ &\qquad + px + qz + uy\Big), \end{aligned}$$

which vanishes for every $r,s,p,q,u \in \mathbb{R}$ if and only if $x = y = z = 0$ and $2m + n = m + 2n = 0$, which means that $m = n = 0$ as well. Thus,

$$\mathfrak{b}^\perp = \left\{ \begin{bmatrix} 0 & a & b \\ 0 & 0 & c \\ 0 & 0 & 0 \end{bmatrix} : a,b,c \in \mathbb{R} \right\},$$

the nilpotent subalgebra of strictly upper triangular matrices, and $\mathfrak{b}$ is parabolic.

For now, the important thing to note is that we get a $\mathfrak{b}$-invariant filtration[2]

$$\mathfrak{sl}_3\mathbb{R} = \mathfrak{g}^{-2} \supset \mathfrak{g}^{-1} \supset \cdots \supset \mathfrak{g}^2 \supset \{0\}$$

of $\mathfrak{sl}_3\mathbb{R}$ given by

$$\begin{aligned} \mathfrak{g}^{-1} &:= \left\{ \begin{bmatrix} m & a & b \\ x & n & c \\ 0 & y & -(m+n) \end{bmatrix} : \ m,n,a,b,c,x,y \in \mathbb{R} \in \mathbb{R} \right\}, \\ \mathfrak{g}^0 &:= \mathfrak{b} = \left\{ \begin{bmatrix} m & a & b \\ 0 & n & c \\ 0 & 0 & -(m+n) \end{bmatrix} : \ m,n,a,b,c \in \mathbb{R} \in \mathbb{R} \right\}, \\ \mathfrak{g}^1 &:= \mathfrak{b}^\perp = \left\{ \begin{bmatrix} 0 & a & b \\ 0 & 0 & c \\ 0 & 0 & 0 \end{bmatrix} : \ a,b,c \in \mathbb{R} \in \mathbb{R} \right\}, \\ \mathfrak{g}^2 &:= [\mathfrak{b}^\perp, \mathfrak{b}^\perp] = \left\langle \left[\begin{smallmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{smallmatrix}\right] \right\rangle. \end{aligned}$$

[2]The weird choice of direction for these filtrations bothered me at first too; there's a very good reason for it, though, so just go with it.

For each i and j, one can show that $[\mathfrak{g}^i, \mathfrak{g}^j] \subseteq \mathfrak{g}^{i+j}$. As it turns out, every parabolic subalgebra has such a filtration canonically associated to it.

10.2. The canonical filtration

We can generalize our observations from $\mathfrak{b} < \mathfrak{sl}_3\mathbb{R}$ to arbitrary parabolic subalgebras by using the following theorem.

Theorem 10.3. *For $\mathfrak{p}$ a parabolic subalgebra of a semisimple Lie algebra $\mathfrak{g}$, we get a canonical filtration*

$$\mathfrak{g} = \mathfrak{g}^{-k} \supset \mathfrak{g}^{-k+1} \supset \cdots \supset \mathfrak{g}^{k} \supset \{0\}$$

of $\mathfrak{g}$, defined by $\mathfrak{g}^0 = \mathfrak{p}$, $\mathfrak{g}^1 = \mathfrak{p}^\perp$, $\mathfrak{g}^i = [\mathfrak{p}^\perp, \mathfrak{g}^{i-1}]$ for each $i > 1$, and $\mathfrak{g}^{-j} = (\mathfrak{g}^{j+1})^\perp$ for all j, such that $[\mathfrak{g}^i, \mathfrak{g}^j] \subseteq \mathfrak{g}^{i+j}$ for all i and j.

Proof. The basic idea is to first show that the subspaces $\mathfrak{g}^i$ satisfy $[\mathfrak{g}^i, \mathfrak{g}^j] \subseteq \mathfrak{g}^{i+j}$ without showing that they give a filtration, and then use this to verify that we get a filtration.

For $j > 0$, $[\mathfrak{p}^\perp, \mathfrak{g}^j] = \mathfrak{g}^{j+1}$ by definition, and since the Killing form satisfies $\mathfrak{B}([X,Y],Z) = \mathfrak{B}(X,[Y,Z])$, we get $[\mathfrak{p}^\perp, \mathfrak{g}^j] \subseteq \mathfrak{g}^{j+1} = (\mathfrak{g}^{-j})^\perp$ for $j \leq 0$ as well. Thus, by the Jacobi identity, $[\mathfrak{g}^i, \mathfrak{g}^j] \subseteq \mathfrak{g}^{i+j}$ whenever $i > 0$. Using this and the invariance of $\mathfrak{B}$ again, it follows that

$$[\mathfrak{g}^{-i}, \mathfrak{g}^{-j}] = [(\mathfrak{g}^{i+1})^\perp, (\mathfrak{g}^{j+1})^\perp] \subseteq \mathfrak{g}^{-i-j} = (\mathfrak{g}^{i+j+1})^\perp$$

for $i, j \geq 0$ as well, so $[\mathfrak{g}^i, \mathfrak{g}^j] \subseteq \mathfrak{g}^{i+j}$ for all i and j.

Then, to show that $\mathfrak{g}^i \supseteq \mathfrak{g}^{i+1}$ for all i, note that

$$[\mathfrak{g}^{1-i}, [\mathfrak{g}^{i+1}, \mathfrak{g}^j]] \subseteq [\mathfrak{g}^{1-i}, \mathfrak{g}^{i+j+1}] \subseteq \mathfrak{g}^{j+2},$$

so $\mathrm{ad}_X \circ \mathrm{ad}_Y$ is nilpotent for each $X \in \mathfrak{g}^{1-i}$ and $Y \in \mathfrak{g}^{i+1}$, hence we get $\mathfrak{g}^{i+1} \subseteq (\mathfrak{g}^{1-i})^\perp = \mathfrak{g}^i$. □

Since $[\mathfrak{p}^\perp, \mathfrak{p}] \subseteq \mathfrak{p}^\perp$, the subalgebra $\mathfrak{p}^\perp$ is a nilpotent ideal of $\mathfrak{p}$. Moreover, since $[\mathfrak{p}^\perp, \mathfrak{g}^j] \subseteq \mathfrak{g}^{j+1}$, every element of $\mathfrak{p}^\perp$ is ad-nilpotent for $\mathfrak{g}$. This nilpotent ideal $\mathfrak{p}^\perp$ is precisely the horospherical subalgebra $\mathfrak{p}_+$ of $\mathfrak{p}$ that we discussed last time. To define the other horospherical subalgebra $\mathfrak{g}_-$ and the neutral subalgebra $\mathfrak{g}_0$, though, we'll need to define a grading, which will require one more tool from the symmetric space perspective.

10.3. Cartan involutions

Given a maximal compact subgroup $K \leq G$, recall from last time that we can decompose $\mathfrak{g}$ (as a vector space) as $\mathfrak{k}^\perp + \mathfrak{k}$, where the Killing form is positive-definite on $\mathfrak{k}^\perp$ and negative-definite on $\mathfrak{k}$. Using this, we can specify a remarkable linear endomorphism θ of $\mathfrak{g}$, called a *Cartan involution*, by defining $\theta|_{\mathfrak{k}^\perp} = -\mathrm{id}_{\mathfrak{k}^\perp}$ and $\theta|_{\mathfrak{k}} = \mathrm{id}_{\mathfrak{k}}$; visually, this just corresponds to reversing geodesic trajectories through e.

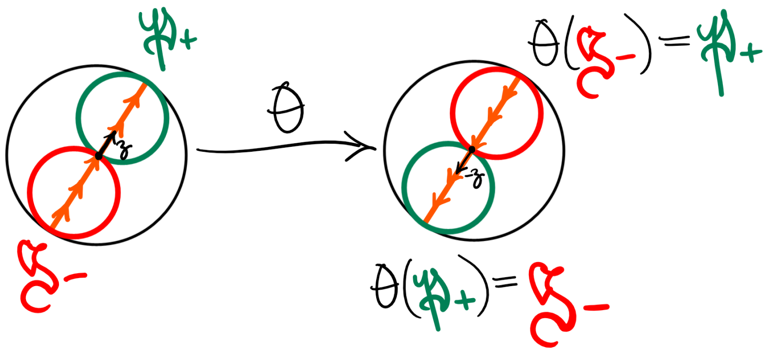

Figure 1. The Cartan involution θ reverses geodesic trajectories through e and swaps the horospherical subalgebras associated to each $Z \in \mathfrak{k}^\perp$

The map θ has several useful properties for representation theory. Perhaps chief among these useful properties is that θ happens to be an automorphism of $\mathfrak{g}$, hence an isometry for $\mathfrak{Z}$. From its definition, we can also see that $\theta^2 = \mathrm{id}_{\mathfrak{g}}$. Using the decomposition $\mathfrak{k}^\perp + \mathfrak{k}$, we can even see that the symmetric bilinear form $\mathfrak{Z}_\theta$ given by $\mathfrak{Z}_\theta := \mathfrak{Z}(\theta(X), Y)$ is negative-definite on all of $\mathfrak{g}$, which allows us to define things like orthogonal projections.

For now, our main interest in the Cartan involution θ associated to K comes from its behavior on horospherical subalgebras. To see what this behavior is, imagine that we have a point $\gamma(+\infty)$ at infinity, where we can once again take γ to be of the form $t \mapsto q_K(\exp(tZ))$ for some $Z \in \mathfrak{k}^\perp$. For Y an eigenvector of ad_Z with eigenvalue λ,

$$[Z, \theta(Y)] = [-\theta(Z), \theta(Y)] = -\theta([Z, Y]) = -\theta(\lambda Y) = -\lambda\theta(Y),$$

so $\theta(Y)$ is an eigenvector of ad_Z with eigenvalue $-\lambda$. In particular, the Cartan involution θ swaps the horospherical subalgebras $\mathfrak{g}_-$ and $\mathfrak{p}_+$ that we defined last time.

10.4. Gradings and a fixed point at infinity

With a given Cartan involution θ, we can use the canonical filtration $\mathfrak{g}^i$ to construct an underlying grading

$$\mathfrak{g} = \mathfrak{g}_{-k} + \mathfrak{g}_{-k+1} + \cdots + \mathfrak{g}_k$$

given by $\mathfrak{g}_i := \mathfrak{g}^i \cap \theta(\mathfrak{g}^{-i})$.

Theorem 10.4. *The grading $\sum_i \mathfrak{g}_i$ satisfies the following properties.*

(1) *For each i, $\theta(\mathfrak{g}_i) = \mathfrak{g}_{-i}$.*

(2) *For each* i, $\mathfrak{g}^i = \sum_{j \geq i} \mathfrak{g}_j$.

(3) *For each* i *and* j, $[\mathfrak{g}_i, \mathfrak{g}_j] \subseteq \mathfrak{g}_{i+j}$.

Proof. For (1), note that $\theta(\mathfrak{g}_i) = \theta(\mathfrak{g}^i) \cap \theta^2(\mathfrak{g}^{-i}) = \theta(\mathfrak{g}^i) \cap \mathfrak{g}^{-i} = \mathfrak{g}_{-i}$.

For (2), let $\mathfrak{g}^k$ be the smallest nonzero filtration component. Then, $\mathfrak{g}^{k+1} = \{0\}$, so $\mathfrak{g}^{-k} = (\mathfrak{g}^{k+1})^{\perp} = \mathfrak{g}$, so $\mathfrak{g}_k = \mathfrak{g}^k \cap \theta(\mathfrak{g}^{-k}) = \mathfrak{g}^k$. Proceeding by induction, suppose $\mathfrak{g}^{i+1} = \sum_{j \geq i+1} \mathfrak{g}_j$; we want to prove that

$$\mathfrak{g}^i = \sum_{j \geq i} \mathfrak{g}_j = \mathfrak{g}_i + \mathfrak{g}^{i+1}.$$

To do this, let $\pi : \mathfrak{g}^i \to \mathfrak{g}^{i+1}$ be the B_θ-orthogonal projection map. Then, we can decompose each $X \in \mathfrak{g}^i$ as $X = (X - \pi(X)) + \pi(X)$, where $\pi(X) \in \mathfrak{g}^{i+1}$ and $X - \pi(X) \in \theta(\mathfrak{g}^{i+1})^{\perp} = \theta(\mathfrak{g}^{-i})$, so $\mathfrak{g}^i = \mathfrak{g}_i + \mathfrak{g}^{i+1}$.

Finally, for (3), we know that $[\mathfrak{g}_i, \mathfrak{g}_j] \subseteq [\mathfrak{g}^i, \mathfrak{g}^j] \subseteq \mathfrak{g}^{i+j}$, and since θ is an automorphism,

$$[\mathfrak{g}_i, \mathfrak{g}_j] \subseteq [\theta(\mathfrak{g}^{-i}), \theta(\mathfrak{g}^{-j})] = \theta([\mathfrak{g}^{-i}, \mathfrak{g}^{-j}]) \subseteq \theta(\mathfrak{g}^{-i-j}).$$

Thus, $[\mathfrak{g}_i, \mathfrak{g}_j] \subseteq \mathfrak{g}^{i+j} \cap \theta(\mathfrak{g}^{-i-j}) = \mathfrak{g}_{i+j}$. □

These properties of the grading tell us several useful things. First, $[\mathfrak{g}_0, \mathfrak{g}_0] \subseteq \mathfrak{g}_{0+0} = \mathfrak{g}_0$, so $\mathfrak{g}_0$ is a subalgebra of $\mathfrak{g}$. Indeed, this $\mathfrak{g}_0$ coincides with the neutral subalgebra introduced last time. Similarly, the subspaces $\mathfrak{p}_+ := \sum_{i>0} \mathfrak{g}_i = \mathfrak{g}^1 = \mathfrak{p}^{\perp}$ and $\mathfrak{g}_- := \sum_{i<0} \mathfrak{g}_i$ are subalgebras, and coincide with the horospherical subalgebras from last time.

Delightfully, this grading can also help us find a point at infinity for the model (G, K) fixed by our parabolic subgroup $P < G$, retrieving the more directly geometric definition we mentioned in the last lecture. To do this, we define the *grading derivation* $\delta_{\mathrm{gr}} : \mathfrak{g} \to \mathfrak{g}$ to be the linear endomorphism given by $\delta_{\mathrm{gr}}(X) = iX$ for each i and each $X \in \mathfrak{g}_i$.

For $X \in \mathfrak{g}_i$ and $Y \in \mathfrak{g}_j$,

$$\begin{aligned}\delta_{\mathrm{gr}}([X,Y]) &= (i+j)[X,Y] = i[X,Y] + j[X,Y] \\ &= [iX, Y] + [X, jY] = [\delta_{\mathrm{gr}}(X), Y] + [X, \delta_{\mathrm{gr}}(Y)].\end{aligned}$$

This means that δ_{gr} is a *derivation* on $\mathfrak{g}$, meaning a linear map $\delta : \mathfrak{g} \to \mathfrak{g}$ such that $\delta([X,Y]) = [\delta(X), Y] + [X, \delta(Y)]$. The space $\mathrm{Der}(\mathfrak{g})$ of all derivations of $\mathfrak{g}$ is a Lie algebra under the commutator bracket, and contains the image of ad as a subalgebra. For semisimple Lie algebras, it turns out that *all* derivations are of the form ad_X for some $X \in \mathfrak{g}$.

Lemma 10.5. *For semisimple* $\mathfrak{g}$, $\mathrm{Der}(\mathfrak{g}) = \mathrm{ad}_{\mathfrak{g}}$.

Proof. For $\delta \in \mathrm{Der}(\mathfrak{g})$ and $X \in \mathfrak{g}$, $[\delta, \mathrm{ad}_X] = \mathrm{ad}_{\delta(X)}$, so $\mathrm{ad}_{\mathfrak{g}} \trianglelefteq \mathrm{Der}(\mathfrak{g})$. Thus, the Killing form on $\mathrm{Der}(\mathfrak{g})$ restricts to the Killing form on $\mathrm{ad}_{\mathfrak{g}}$. Since the Killing form on $\mathrm{ad}_{\mathfrak{g}} \approx \mathfrak{g}$ is nondegenerate because $\mathfrak{g}$ is semisimple, we get $\mathrm{Der}(\mathfrak{g}) = \mathrm{ad}_{\mathfrak{g}} \oplus \mathrm{ad}_{\mathfrak{g}}^{\perp}$ as a vector space. But for $\delta \in \mathrm{ad}_{\mathfrak{g}}^{\perp}$, this means

$$[\delta, \mathrm{ad}_X] = \mathrm{ad}_{\delta(X)} \in \mathrm{ad}_{\mathfrak{g}} \cap \mathrm{ad}_{\mathfrak{g}}^{\perp} = \{0\}$$

for all $X \in \mathfrak{g}$, so $\delta = 0$ because $\mathfrak{z}(\mathfrak{g}) = \{0\}$. □

In particular, this tells us that the derivation δ_{gr} is of the form $\mathrm{ad}_{E_{\mathrm{gr}}}$ for some $E_{\mathrm{gr}} \in \mathfrak{g}$. We call this element E_{gr} the *grading element* for the grading on $\mathfrak{g}$. By definition, $\mathrm{ad}_{E_{\mathrm{gr}}} = \delta_{\mathrm{gr}}$ is diagonalizable over $\mathbb{Z} < \mathbb{R}$ on $\mathfrak{g}$ and satisfies $\theta \circ \mathrm{ad}_{E_{\mathrm{gr}}} = -\mathrm{ad}_{E_{\mathrm{gr}}} \circ \theta$, so $E_{\mathrm{gr}} \in \mathfrak{k}^{\perp}$.

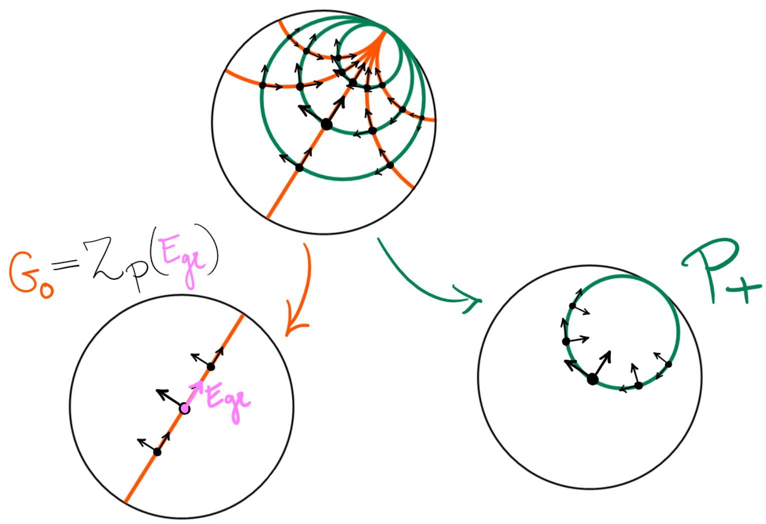

Figure 2. A depiction of a parabolic subgroup, decomposed into G_0, the part that centralizes E_{gr}, and P_+, the part generated by the horospherical subalgebra $\mathfrak{p}_+$

The geodesic $t \mapsto q_K(\exp(tE_{\mathrm{gr}}))$ generated by E_{gr} determines a point at infinity fixed by P. This follows directly from our next theorem, which is essentially just Theorem 3.1.3 of [**2**] and whose proof we consider optional for our current endeavor.

Theorem 10.6. *If P is a parabolic subgroup, then it is of the form $\mathrm{Z}_P(E_{gr})P_+$, where $\mathrm{Z}_P(E_{gr}) = \{p \in P : \mathrm{Ad}_p(E_{gr}) = E_{gr}\}$ and P_+ is the connected subgroup generated by $\mathfrak{p}_+ = \mathfrak{p}^{\perp}$.*

Proof. Suppose $p \in P$. The adjoint action Ad_p on $\mathfrak{g}$ preserves the canonical filtration, so it induces an automorphism $\phi_{\mathrm{gr}}(p)$ of the graded Lie algebra associated to the filtration, so that $\mathrm{Ad}_p Y - \phi_{\mathrm{gr}}(p) \cdot Y \in \mathfrak{g}^{i+1}$ for each $Y \in \mathfrak{g}_i$. In particular, our grading element $E_{\mathrm{gr}} \in \mathfrak{g}_0$ satisfies $\mathrm{Ad}_p E_{\mathrm{gr}} - \phi_{\mathrm{gr}}(p) \cdot E_{\mathrm{gr}} \in \mathfrak{g}^1$, so $\mathrm{Ad}_{p^{-1}}(\phi_{\mathrm{gr}}(p) \cdot E_{\mathrm{gr}}) \in E_{\mathrm{gr}} + \mathfrak{g}^1$.

Let Z_1 be the $\mathfrak{g}_1$-component of $\mathrm{Ad}_{p^{-1}}(\phi_{\mathrm{gr}}(p) \cdot E_{\mathrm{gr}})$, so that

$$\mathrm{Ad}_{p^{-1}}(\phi_{\mathrm{gr}}(p) \cdot E_{\mathrm{gr}}) \in E_{\mathrm{gr}} + Z_1 + \mathfrak{g}^2.$$

Then, $\mathrm{ad}_{Z_1}(E_{\mathrm{gr}} + Z_1) = -Z_1$ and $\mathrm{ad}^2_{Z_1}(E_{\mathrm{gr}} + Z_1) = 0$, so

$$\mathrm{Ad}_{\exp(Z_1)} \circ \mathrm{Ad}_{p^{-1}}(\phi_{\mathrm{gr}}(p) \cdot E_{\mathrm{gr}}) \in E_{\mathrm{gr}} + \mathfrak{g}^2.$$

Recursively, we define Z_i to be $\frac{1}{i}$ times the $\mathfrak{g}_i$-component of

$$\mathrm{Ad}_{\exp(Z_{i-1})} \circ \cdots \circ \mathrm{Ad}_{\exp(Z_1)} \circ \mathrm{Ad}_{p^{-1}}(\phi_{\mathrm{gr}}(p) \cdot E_{\mathrm{gr}}) \in E_{\mathrm{gr}} + \mathfrak{g}^i,$$

so that

$$\mathrm{Ad}_{\exp(Z_i)} \circ \mathrm{Ad}_{\exp(Z_{i-1})} \circ \cdots \circ \mathrm{Ad}_{\exp(Z_1)} \circ \mathrm{Ad}_{p^{-1}}(\phi_{\mathrm{gr}}(p) \cdot E_{\mathrm{gr}}) \in E_{\mathrm{gr}} + \mathfrak{g}^{i+1}.$$

Eventually, there is some k such that $\mathfrak{g}^{k+1} = \{0\}$, so that

$$\mathrm{Ad}_{\exp(Z_k)} \circ \cdots \circ \mathrm{Ad}_{\exp(Z_1)} \circ \mathrm{Ad}_{p^{-1}}(\phi_{\mathrm{gr}}(p) \cdot E_{\mathrm{gr}}) = E_{\mathrm{gr}},$$

hence $\phi_{\mathrm{gr}}(p) \cdot E_{\mathrm{gr}} = \mathrm{Ad}_{p\exp(-Z_1)\cdots\exp(-Z_k)}(E_{\mathrm{gr}})$. But, recall that $\phi_{\mathrm{gr}}(p)$ is an automorphism of the graded Lie algebra, so that

$$[\phi_{\mathrm{gr}}(p) \cdot E_{\mathrm{gr}}, \phi_{\mathrm{gr}}(p) \cdot Y] = \phi_{\mathrm{gr}}(p) \cdot [E_{\mathrm{gr}}, Y] = i\, \phi_{\mathrm{gr}}(p) \cdot Y$$

for each $Y \in \mathfrak{g}_i$. In particular, $\phi_{\mathrm{gr}}(p) \cdot E_{\mathrm{gr}}$ must agree with the grading element because E_{gr} is the unique element with $\mathrm{ad}_{E_{\mathrm{gr}}} = \delta_{\mathrm{gr}}$. Thus,

$$\phi_{\mathrm{gr}}(p) \cdot E_{\mathrm{gr}} = \mathrm{Ad}_{p\exp(-Z_1)\cdots\exp(-Z_k)}(E_{\mathrm{gr}}) = E_{\mathrm{gr}},$$

hence $p\exp(-Z_1)\cdots\exp(-Z_k) \in \mathrm{Z}_P(E_{\mathrm{gr}})$. $\square$

10.5. Parabolic model geometries

As one might guess, we can now define a model to be *parabolic* when its model group is semisimple and its isotropy is parabolic.

Definition 10.7. We say that a model geometry (G, P) is *parabolic* when G is a semisimple Lie group and P is a parabolic subgroup.

These parabolic model geometries are the core of the study of parabolic geometries. With our current pacing through the course, we probably won't get to talk much about the general "curved" case, but this was always meant to be more of an invitation to the topic anyway. Next time, we will investigate what these parabolic models look like, in general.

Chapter 11

The Anatomy of a Parabolic Geometry

Again, rewrite this so that it looks better and fits well within the book. Specifically, rewrite the unipotent tilts part at the end to talk about nice line bundle perspective, and clean up the cell decomposition stuff. Also, be very clear about the curvature stuff!

Last time, we finished by defining a parabolic model geometry to be a model (G, P) with G semisimple and P parabolic. Now, we will begin exploring what these parabolic models look like. This can, of course, seem overwhelming at first, since—even topologically—these geometries are quite a bit more involved than just frames on a plane.

However, it turns out that these model geometries aren't *that* much more complicated than frames on a plane when we separate them into manageable pieces. In this lecture, we'll be learning how to imagine ourselves as observers in a parabolic model geometry with the following tools:

- A large open subset of G/P over which (G, P) looks like a frame bundle over a vector space
- A way of dissecting the base manifold G/P, cutting it into manageable pieces
- A method for visualizing the higher-order parts of G

By the end of the lecture, we should have a decent grasp of what to expect visually when we encounter a parabolic model geometry. This will prepare us for the next two lectures, which will cover the specific examples of projective geometry and conformal Riemannian geometry.

11.1. Open cells

Previously, we saw that the semisimple Lie algebra $\mathfrak{g}$ decomposes as $\mathfrak{g} = \mathfrak{g}_- + \mathfrak{g}_0 + \mathfrak{p}_+ = \mathfrak{g}_- + \mathfrak{p}$. For sufficiently small open neighborhoods V and W of the identity in G_- and P, respectively, this tells us that $\exp(V)\exp(W)$ is an open neighborhood of the identity in G. In particular, for each $up \in G_-P$, the open neighborhood $u\exp(V)\exp(W)p$ of up is also contained in G_-P, so G_-P is an open subset of G. Since q_P is a submersion—hence an open map—it follows that $q_P(G_-) = q_P(G_-P)$ is an open subset of G/P, and since $G_- \cap P = \{e\}$, $q_P|_{G_-}$ is an embedding of G_- into G/P.

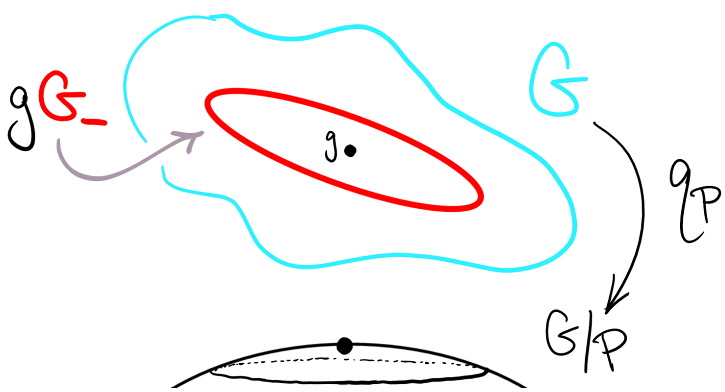

Figure 1. The natural quotient map $q_P : G \to G/P$ restricts to an embedding on each gG_-

We saw this *open cell* when we first encountered the parabolic model geometry $(\mathrm{SL}_2\,\mathbb{R}, B)$; in that case, the open cell $q_P(G_-)$ corresponded to a copy of the affine line.

The horospherical subgroup G_- is simply connected and nilpotent. In particular, this tells us that the exponential map $\exp : \mathfrak{g}_- \to G_-$ is a diffeomorphism, so that G_- is topologically equivalent to a vector space. The subgroup $G_0 := \mathrm{Z}_P(E_{\mathrm{gr}})$ acts on the subalgebra $\mathfrak{g}_-$ by the adjoint representation, so under this topological identification between $\mathfrak{g}_-$ and G_- given by the exponential map, the conjugation action of G_0 on G_- is linear. This puts us in a situation with which we should be fairly comfortable: G_-G_0 has G_0 as a closed subgroup acting linearly on the normal subgroup $G_- \trianglelefteq G_-G_0$, just like how $\mathrm{I}(2)$ has $\mathrm{O}(2)$ as a closed subgroup acting linearly on the normal subgroup $\mathbb{R}^2$ of translations. In short, we can think of G_-G_0 as a space of particular frames over G_-.

Note that G_-G_0 is another parabolic subgroup of G. Indeed, its Lie subalgebra $\mathfrak{g}_- + \mathfrak{g}_0$ satisfies $(\mathfrak{g}_- + \mathfrak{g}_0)^\perp = \mathfrak{g}_-$, and for θ a Cartan involution used to obtain the grading, $\theta(\mathfrak{p}_+) = \mathfrak{g}_-$ and $\theta(\mathfrak{g}_0) = \mathfrak{g}_0$, so $\mathfrak{g}_- + \mathfrak{g}_0 = \theta(\mathfrak{p})$.

We call it the *opposite parabolic* to P; note that there might be a different choice of opposite parabolic for a different choice of Cartan involution θ determining the grading.

Often, the geometry of the model (G_-G_0, G_0) is a kind of affine analogue of the geometry of (G, P). In the model $(\mathrm{PGL}_{m+1}\mathbb{R}, P)$ for projective geometry, for example, $G_- \simeq \mathbb{R}^m$ and $G_0 \simeq \mathrm{GL}_m\mathbb{R}$, so that (G_-G_0, G_0) is equivalent to $(\mathrm{Aff}(m), \mathrm{GL}_m\mathbb{R})$, the model for affine geometry. We'll see this in a bit more detail in the next lecture.

Conveniently, the open subset $G_-P = q_P^{-1}(G_-)$ of G is topologically a product $G_- \times P$, since $G_- \cap P = \{e\}$. As we saw last time, P itself is of the form G_0P_+, and since $G_0 \cap P_+ = \{e\}$, it is also topologically a product $G_0 \times P_+$. Altogether, this tells us that $G_-P = q_P^{-1}(G_-)$ looks like $G_-G_0 \times P_+$, so over $q_P(G_-)$, the geometry looks like a kind of frame bundle G_-G_0 over G_-, together with some "higher-order frames" from P_+ on top. We'll give some insight into what these "higher-order frames" look like later in this lecture.

For each configuration $g \in G$ over G/P, we get a copy of G_- as the left-coset gG_-. Since these are just left-translations of G_- by g, meaning they are images of G_- under the transformation given by g, the geometry looks the same on gG_- as it does on G_-. In other words, wherever we are at in G, we can give ourselves a convenient open subset on which the geometry looks like a "higher-order frame bundle" over a copy of G_-.

Of course, all of this makes G_- a prime candidate for an analogue of the translation subgroup in $\mathrm{I}(2)$, so we can get a notion of geodesic inside our current copy of G_- by using one-parameter subgroups generated by elements of $\mathfrak{g}_-$. In the case of projective geometry, these will just be affine geodesics inside the current affine patch. These types of distinguished curves generally aren't as consistent in the base manifold as the other types of geodesics we've dealt with so far; we'll see this most prominently when we talk about conformal geometry. However, this type of motion is always available and meaningful from our observer perspective in the model group G.

11.2. Filling in the rest of G/P

As we saw above, the horospherical subgroup G_- essentially lets us reduce the local picture of (G, P) to that of frames on a vector space. However, we'd still like to have an idea of what G/P looks like globally.

Thankfully, the open cell $q_P(G_-)$ often takes up a large portion of G/P. We saw this in the case of $(\mathrm{SL}_2\mathbb{R}, B)$, for example, where the affine line took up all of $\mathrm{SL}_2\mathbb{R}/B$ except for a "point at infinity". Unlike in the case of symmetric spaces, we do not need to describe this point at infinity in terms of asymptotic boundedness; it is literally the limit of an affine line embedded into $\mathrm{SL}_2\,\mathbb{R}/B$.

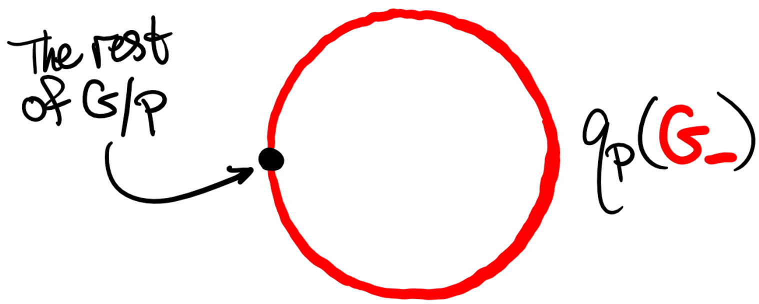

Figure 2. The open subset $q_P(G_-)$ often fills up a large portion of G/P

We'd like to have a way of breaking G/P into smaller, topologically simple pieces, similar to the case of $(\mathrm{SL}_2\,\mathbb{R}, B)$. It turns out that we can do this, through a generalization of something called the *Bruhat decomposition*.

Inside of our parabolic subgroup P, let us choose a minimal parabolic subgroup $B \le P$. Since B is parabolic, we get a corresponding filtration subordinate to the filtration from P, and by using the same Cartan involution θ, we get a grading of $\mathfrak{g}$ subordinate to the grading from P. Let us denote by Z the grading element for this new grading.

As before, we can decompose $\mathfrak{g}$ into the centralizer $\mathfrak{b}_0 := \mathfrak{z}_{\mathfrak{g}}(Z)$ and two horospherical subalgebras

$$\mathfrak{b}_- := \{X \in \mathfrak{g} : \mathrm{Ad}_{\exp(tZ)}(X) \to 0 \text{ as } t \to +\infty\}$$

and

$$\mathfrak{b}_+ := \mathfrak{b}^\perp = \{X \in \mathfrak{g} : \mathrm{Ad}_{\exp(tZ)}(X) \to 0 \text{ as } t \to -\infty\},$$

so that $\mathfrak{g} = \mathfrak{b}_- + \mathfrak{b}_0 + \mathfrak{b}_+$ and $\mathfrak{b} = \mathfrak{b}_0 + \mathfrak{b}_+$. Let B_- be the connected subgroup generated by $\mathfrak{b}_-$.

Since $\mathfrak{b} \le \mathfrak{p}$, we must have $\mathfrak{p}^\perp = \mathfrak{p}_+ \le \mathfrak{b}_+ = \mathfrak{b}^\perp$, and similarly, $\theta(\mathfrak{p}_+) = \mathfrak{g}_- \le \mathfrak{b}_- = \theta(\mathfrak{b}_+)$. In other words, because B is smaller than P, the horospherical part of B must be larger than the horospherical part of P. Moreover, since $\mathfrak{b}_+ \le \mathfrak{b} \le \mathfrak{p}$, we must also have that $\theta(\mathfrak{b}_+) = \mathfrak{b}_- \le \mathfrak{g}_- + \mathfrak{g}_0 = \theta(\mathfrak{p})$. Thus, $G_- \le B_- \le G_-G_0$, and in particular, $q_P(G_-) = q_P(B_-)$.

Of course, every element of G/P lies in some orbit of B_-, but it turns out that there are often only finitely many B_--orbits (when G/P is compact).

Theorem 11.1. *Given a parabolic model (G, P), G/P decomposes as a disjoint union of cells*

$$G/P = \bigsqcup_{\sigma \in \mathfrak{W}^P} B_- q_P(\sigma),$$

where $\mathfrak{W}^P = \mathrm{N}_G(\mathfrak{b}_0)/\mathrm{N}_{G_0}(\mathfrak{b}_0)$. *Moreover, if* G/P *is compact, then* $\mathfrak{W}^P$ *is finite.*

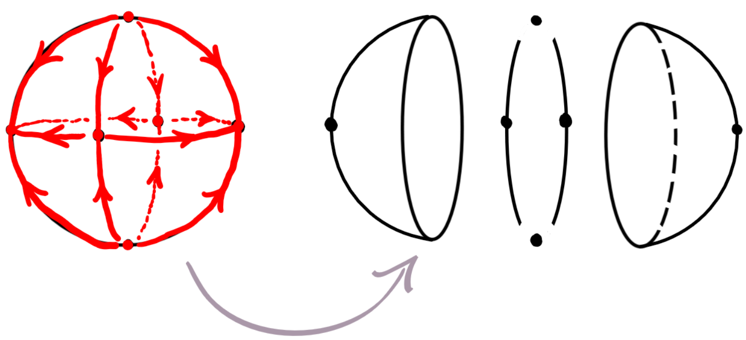

Figure 3. The cell decomposition for a parabolic geometry, corresponding to a decomposition into stable manifolds for the action of the grading element of a minimal parabolic subgroup

In the classical, algebraic case over $\mathbb{C}$, $\mathfrak{W}^B = \mathrm{N}_G(\mathfrak{b}_0)/B_0$ is a finite group called the *Weyl group*.

The proof, which will hopefully be part of an upcoming joint work between Rachel and me, requires quite a few technical results from representation theory. However, the idea of the proof is fairly straightforward: consider the left-action of $\exp(tZ)$ on G/P. The fixed points of this flow will correspond to the points of $\mathfrak{W}^P$, and the stable manifolds for these fixed points will be their B_--orbits.

This decomposition is, geometrically, a bit fragile. In the general "curved" case, it often doesn't work. **If, however, the holonomy happens to be unipotent, and one happens to have a way of describing "curved" cosets...** Well, more on that later.

Ultimately, what this usually looks like is a big cell coming from the open subset $q_P(G_-) = q_P(B_-)$ together with some collections of "points at infinity" that compactify it.

11.3. How do we see P_+?

Above, we showed that the open cell lets us reduce to the open set $G_-P = (G_-G_0)P_+$ to get the local picture of (G, P). We already have a fairly good picture of G_-G_0, as a particular frame bundle over G_-, so all that really remains is to figure out the P_+ part. There are three perspectives that I find useful for this purpose; all three are useful in different situations, and together they give a fairly satisfying picture of what's going on.

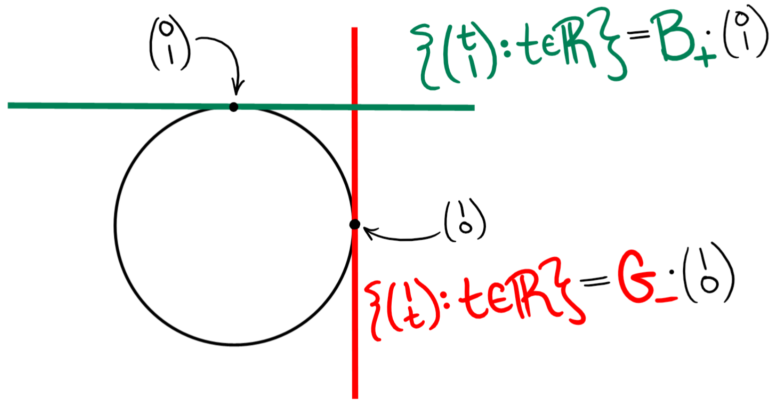

Figure 4. In $(\mathrm{SL}_2\,\mathbb{R}, B)$, G_- acts by translations on the affine line through $\binom{1}{0}$, while B_+ acts by translations on the affine line through $\binom{0}{1}$

First, we can think of P_+ as a kind of "dual" translation subgroup to G_-. Just like in the case of G_-, the left-action of the horospherical subgroup P_+ determines an open cell[1] on G/P. We saw this in the case of $(\mathrm{SL}_2\,\mathbb{R}, B)$, where the subgroup G_- acted by translations along one affine line, and B_+ acted by translations along another affine line through the point at infinity of the first. On the open cell determined by P_+, it acts as G_- does on its own open cell through $q_P(e)$. This is, of course, quite useful for seeing P_+ as a group of transformations of G/P, but the global nature of it kind of defeats the purpose of restricting to the local picture in the first place.

The second way of seeing P_+ comes from using the Killing form $\mathfrak{k}$. Recall that $\mathfrak{p} = \mathfrak{p}_+^\perp$. Because $\mathfrak{k}$ is nondegenerate, this gives us a duality between $\mathfrak{g}/\mathfrak{p} = \mathfrak{g}/\mathfrak{p}_+^\perp$ and $\mathfrak{p}_+$. In particular, the dual space $(\mathfrak{g}/\mathfrak{p})^\vee$ is isomorphic to $\mathfrak{p}_+$ as a P-representation, and hence the cotangent bundle $T^\vee(G/P)$ satisfies

$$T^\vee(G/P) \cong G \times_P (\mathfrak{g}/\mathfrak{p})^\vee \cong G \times_P \mathfrak{p}_+.$$

Since $\mathfrak{g}_- \approx (\mathfrak{g}_- + \mathfrak{g}_0 + \mathfrak{p}_+)/(\mathfrak{g}_0 + \mathfrak{p}_+) = \mathfrak{g}/\mathfrak{p}$ as G_0-representations, this recovers the duality between $\mathfrak{g}_-$ and $\mathfrak{p}_+$ that we've mentioned before: each element $\alpha \in \mathfrak{g}_-^\vee$ corresponds to a unique element $\alpha^\mathfrak{k} \in \mathfrak{p}_+$ such that $\alpha(X) = \mathfrak{k}(\alpha^\mathfrak{k}, X)$ for every $X \in \mathfrak{g}_-$. This again lets us think of $\mathfrak{p}_+$ as a subalgebra of "dual" translations to $\mathfrak{g}_-$. Algebraically, this is a convenient perspective, though it is a bit difficult to give a visual depiction of it.

The third way to see P_+ generalizes the "unipotent tilt" perspective that we described when we discussed $(\mathrm{SL}_2\,\mathbb{R}, B)$. As we mentioned above, we have a copy gG_- of G_- through each $g \in G$, which allows us to give a local

[1]Specifically, it is of the form $P_+q_P(\sigma)$ for a particular element of $\mathfrak{W}^P$ of "maximal length", though we don't really need to know what that means right now.

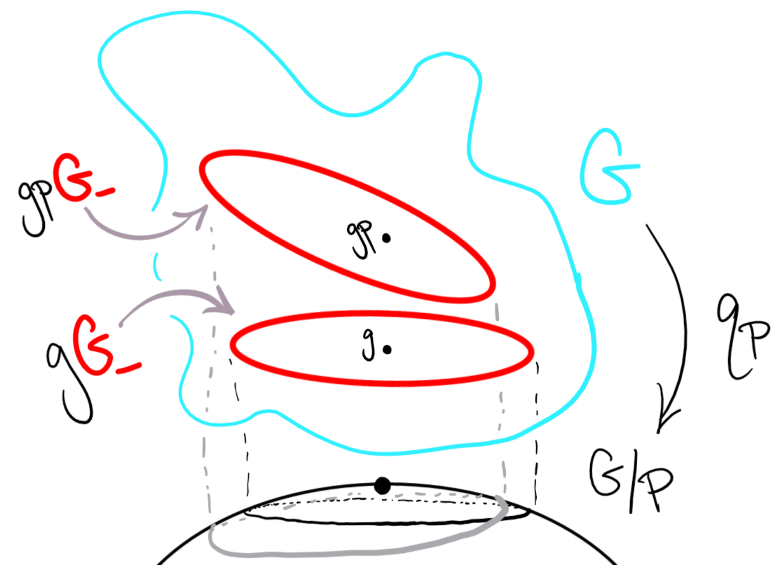

Figure 5. The right-action of P_+ is by "unipotent tilts" of the copies of G_-

picture of (G, P) as a kind of "higher-order frame bundle" $g(G_-G_0)P_+$ over gG_-. Visually, when we right-translate by some $p \in P_+$ and then consider the corresponding copy gpG_- of G_-, the result is a kind of "tilting" of gG_-. Since this is difficult to describe abstractly, we'll return to this in the next two lectures when we talk about explicit examples.

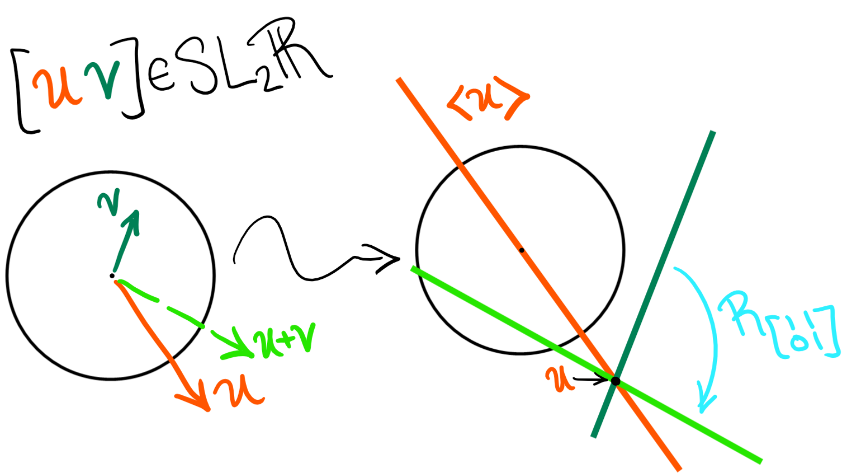

Figure 6. Right-translating $[u\ v] \in \mathrm{SL}_2\,\mathbb{R}$ by a unipotent tilt takes the affine line determined by v and tilts it along the line determined by u

11.4. Curvature?

Chapter 12

Projective Geometry

Again, rewrite this so that it looks better and fits well within the book.

Have you ever wondered what it is like to move around inside a painting? It's a fun and evocative exercise in the imagination, and it was something I remember thinking about often as a child. However, as a child, I was ill-equipped to understand the geometry of the situation, or even what geometry means in this case.

In today's lecture, we'll be exploring this two-dimensional *projective geometry* and its higher-dimensional analogues. In outline, our plan is the following:

- Explain what geometry means for a painting
- Verify that the geometry is parabolic
- Describe how to move around within projective geometry
- Discuss what geodesics in the geometry look like

By the end of the lecture, we should have a decent idea of what it's like to move around inside the geometry of a painting. In particular, we'll have a better picture of what parabolic model geometries look like; we will further supplement this picture next time, when we talk about conformal geometry.

12.1. The geometry of a painting

Let's imagine a landscape painting: there is a pond with a small pine tree next to it, and behind these is a majestic mountain range, beyond which is a setting sun.

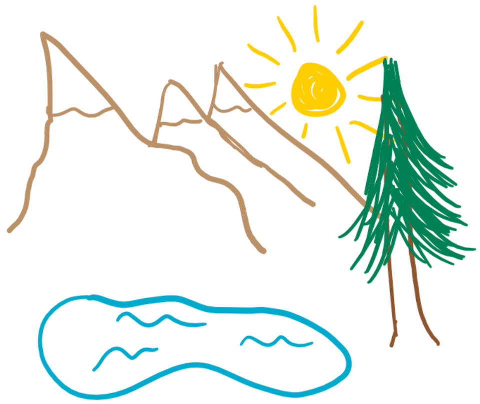

Figure 1. A two-dimensional image depicting a three-dimensional scene, with a pond and pine tree next to each other and a mountain range in the background, behind which is a setting sun

Our use of the words "behind" and "beyond" here suggests that, while the painting itself is two-dimensional, we think of the scene depicted as occurring in three dimensions. How do we get the two-dimensional image from the three-dimensional scene?

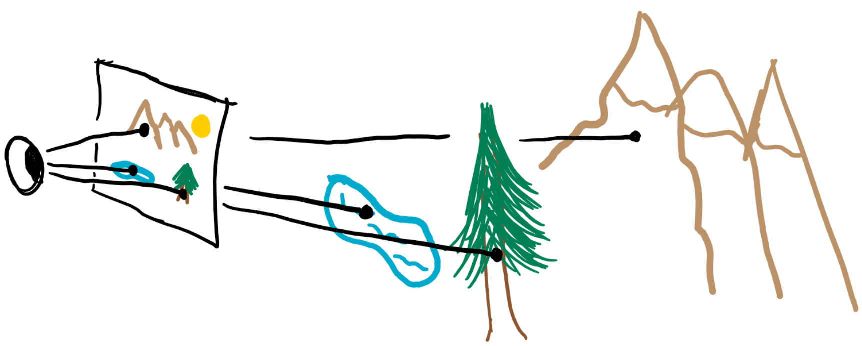

Figure 2. Each point of the canvas corresponds to a sight-line between the eye of the painter and a point in the scene

Let's suppose that the landscape occurs inside of $\mathbb{R}^3$, and that the painter observes the scene through their eye, which we place at the origin in $\mathbb{R}^3$. For each point $x \in \mathbb{R}^3$ in the scene, there is a unique line $\langle x \rangle = \mathbb{R}x$ through the origin that also contains this point; we call this line the *sight-line* through x. When the painter commits this scene to their canvas, they are identifying each point of their canvas with a corresponding sight-line, effectively *projecting* the three-dimensional scene down to a two-dimensional image.

Since the geometry comes from these sight-lines, its symmetries will be those that preserve them. A transformation that preserves lines in $\mathbb{R}^3$ *and* preserves the origin (where the eye is) is going to be an element of $\mathrm{GL}_3\,\mathbb{R}$. However, since the sight-lines are what we're really interested in and the center $\mathrm{Z}(\mathrm{GL}_3\,\mathbb{R}) = \mathbb{R}^\times\mathbb{1}$ sends each line through 0 to itself, we want to ignore these central elements. Thus, the model group of this geometry is $\mathrm{GL}_3\,\mathbb{R}/\mathbb{R}^\times\mathbb{1} = \mathrm{PGL}_3\,\mathbb{R}$.

The model group $\mathrm{PGL}_3\,\mathbb{R}$ acts transitively on the projective plane $\mathbb{RP}^2$, also known as the space of sight-lines in $\mathbb{R}^3$. Thus, defining P to be the stabilizer of the sight-line through $\left[\begin{smallmatrix}1 & 0 & 0\end{smallmatrix}\right]^\top$, we get a bijection between $\mathrm{PGL}_3\,\mathbb{R}/P$ and $\mathbb{RP}^2$. In short, our model for 2-dimensional projective geometry is $(\mathrm{PGL}_3\,\mathbb{R}, P)$.

More generally, we can consider m-dimensional projective geometry, which corresponds to the geometry of sight-lines inside of $\mathbb{R}^{m+1}$. In that case, our model is $(\mathrm{PGL}_{m+1}\,\mathbb{R}, P)$, where

$$P := \left\{\begin{pmatrix} a & \alpha \\ 0 & A\end{pmatrix} \in \mathrm{PGL}_{m+1}\,\mathbb{R} : a \in \mathbb{R}^\times, \alpha^\top \in \mathbb{R}^m, A \in \mathrm{GL}_m\,\mathbb{R}\right\}$$

is, again, the stabilizer of the sight-line through $\left[\begin{smallmatrix}1 & 0 & \cdots & 0\end{smallmatrix}\right]^\top$.

12.2. Parabolicity of projective geometry

The Killing form on $\mathfrak{pgl}_{m+1}\mathbb{R} := \mathfrak{gl}_{m+1}\mathbb{R}/\mathbb{R}\mathbb{1}$, where by definition elements of $\mathfrak{pgl}_{m+1}\mathbb{R}$ are equivalent if and only if they differ by a scalar multiple of the identity matrix, is given by

$$\mathfrak{B}\Big(\left(\begin{smallmatrix} -\mathrm{tr}(R) & \alpha \\ v & R\end{smallmatrix}\right), \left(\begin{smallmatrix} -\mathrm{tr}(S) & \beta \\ w & S\end{smallmatrix}\right)\Big) = 2(m+1)\Big(\mathrm{tr}(R)\mathrm{tr}(S) + \mathrm{tr}(RS) + \alpha(w) + \beta(v)\Big),$$

so

$$\mathfrak{B}\Big(\left(\begin{smallmatrix} -\mathrm{tr}(R) & \alpha \\ 0 & R\end{smallmatrix}\right), \left(\begin{smallmatrix} -\mathrm{tr}(S) & \beta \\ w & S\end{smallmatrix}\right)\Big) = 2(m+1)\Big(\mathrm{tr}(R)\mathrm{tr}(S) + \mathrm{tr}(RS) + \alpha(w)\Big),$$

which vanishes for all $\left(\begin{smallmatrix} -\mathrm{tr}(R) & \alpha \\ 0 & R\end{smallmatrix}\right) \in \mathfrak{p}$ precisely when $S = 0$ and $w = 0$. Thus, $\mathfrak{p}^\perp$ is the abelian subalgebra $\{\left(\begin{smallmatrix} 0 & \alpha \\ 0 & 0\end{smallmatrix}\right) : \alpha^\top \in \mathbb{R}^m\}$, hence $\mathfrak{p}$ is parabolic.

Choosing our Cartan involution θ to be given by $X \mapsto -X^\top$, so that

$$\theta\Big(\left(\begin{smallmatrix} r & \alpha \\ v & R\end{smallmatrix}\right)\Big) = \left(\begin{smallmatrix} -r & -v^\top \\ -\alpha^\top & -R^\top\end{smallmatrix}\right),$$

we get a grading of $\mathfrak{pgl}_{m+1}\mathbb{R}$ given by

$$\mathfrak{g}_{-1} = \mathfrak{g}_- := \left\{\left(\begin{smallmatrix} 0 & 0 \\ v & 0\end{smallmatrix}\right) \in \mathfrak{pgl}_{m+1}\mathbb{R} : v \in \mathbb{R}^m\right\},$$

$$\begin{aligned}\mathfrak{g}_0 &:= \left\{\left(\begin{smallmatrix} r & 0 \\ 0 & R\end{smallmatrix}\right) \in \mathfrak{pgl}_{m+1}\mathbb{R} : r \in \mathbb{R},\, R \in \mathfrak{gl}_m\mathbb{R}\right\} \\ &= \left\{\left(\begin{smallmatrix} -\mathrm{tr}(S) & 0 \\ 0 & S\end{smallmatrix}\right) \in \mathfrak{pgl}_{m+1}\mathbb{R} : S \in \mathfrak{gl}_m\mathbb{R}\right\},\end{aligned}$$

and

$$\mathfrak{g}_1 = \mathfrak{p}_+ := \left\{ \left(\begin{smallmatrix} 0 & \alpha \\ 0 & 0 \end{smallmatrix}\right) \in \mathfrak{pgl}_{m+1}\mathbb{R} : \alpha^\top \in \mathbb{R}^m \right\}.$$

It is worth drawing attention to the fact that $\mathfrak{g}_-$ and $\mathfrak{p}_+$ are abelian in this case, so that we only have three grading components. The horospherical subgroups

$$G_- := \{ \left(\begin{smallmatrix} 1 & 0 \\ v & \mathbb{1} \end{smallmatrix}\right) \in \mathrm{PGL}_{m+1}\mathbb{R} : v \in \mathbb{R}^m \}$$

and

$$P_+ := \left\{ \left(\begin{smallmatrix} 1 & \alpha \\ 0 & \mathbb{1} \end{smallmatrix}\right) \in \mathrm{PGL}_{m+1}\mathbb{R} : \alpha^\top \in \mathbb{R}^m \right\}$$

are, in particular, also abelian.

The grading element for this grading is given by

$$E_{\mathrm{gr}} := \tfrac{1}{m+1}\left(\begin{smallmatrix} m & 0 \\ 0 & -\mathbb{1} \end{smallmatrix}\right) = \tfrac{1}{m+1}\left(\begin{smallmatrix} m & 0 \\ 0 & -\mathbb{1} \end{smallmatrix}\right) + \tfrac{1}{m+1}\left(\begin{smallmatrix} 1 & 0 \\ 0 & \mathbb{1} \end{smallmatrix}\right) = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 0 \end{smallmatrix}\right),$$

where again, elements of $\mathfrak{pgl}_{m+1}\mathbb{R}$ are equivalent whenever they differ by a scalar multiple of the identity matrix. From this, we can see that

$$\begin{aligned} G_0 := \mathrm{Z}_P(E_{\mathrm{gr}}) &= \left\{ \left(\begin{smallmatrix} a & 0 \\ 0 & A \end{smallmatrix}\right) \in \mathrm{PGL}_{m+1}\mathbb{R} : a \in \mathbb{R}^\times,\ A \in \mathrm{GL}_m\mathbb{R} \right\} \\ &= \left\{ \begin{pmatrix} \frac{1}{\det(S)} & 0 \\ 0 & S \end{pmatrix} \in \mathrm{PGL}_{m+1}\mathbb{R} : S \in \mathrm{GL}_m\mathbb{R} \right\} \end{aligned}$$

is the neutral subgroup.

Momentarily, we will also be interested in a specific normal subgroup $G_0^{\mathrm{ss}} \trianglelefteq G_0$, the **s**emi**s**imple part of G_0, given by

$$G_0^{\mathrm{ss}} := \left\{ \left(\begin{smallmatrix} 1 & 0 \\ 0 & A \end{smallmatrix}\right) \in \mathrm{PGL}_{m+1}\mathbb{R} : A \in \mathrm{SL}_m^{\pm}\mathbb{R} \right\},$$

where $\mathrm{SL}_m^{\pm}\mathbb{R}$ is the Lie group of linear transformations of $\mathbb{R}^m$ with determinant either $+1$ or -1. This subgroup G_0^{ss} is normal in G_0, and moreover, G_0 decomposes as $G_0 = \exp(\mathbb{R}E_{\mathrm{gr}})G_0^{\mathrm{ss}}$, where $\exp(\mathbb{R}E_{\mathrm{gr}})$ is the image of the one-parameter subgroup generated by E_{gr}. In particular, $G_0/G_0^{\mathrm{ss}} \simeq P/(G_0^{\mathrm{ss}}P_+) \simeq \exp(\mathbb{R}E_{\mathrm{gr}})$.

12.3. The pedestrian perspective

As before, we'd like to think of our model group $\mathrm{PGL}_{m+1}\mathbb{R}$ as the space of configurations for ourselves as pedestrians[1] inside of the geometry of our model. To do this, let's return to the idea of walking around inside a painting.

Paintings don't typically depict what is going on both in front and behind the painter, so it makes sense to assume the scene in $\mathbb{R}^{m+1}$ depicted in an m-dimensional painting takes place in some half-space $\{x \in \mathbb{R}^{m+1} : \alpha(x) > 0\}$

[1] Note that I'm returning to the term "pedestrian" rather than "observer" here, since the "observer perspective" could easily be confused with the perspective of the eye viewing the sight-lines.

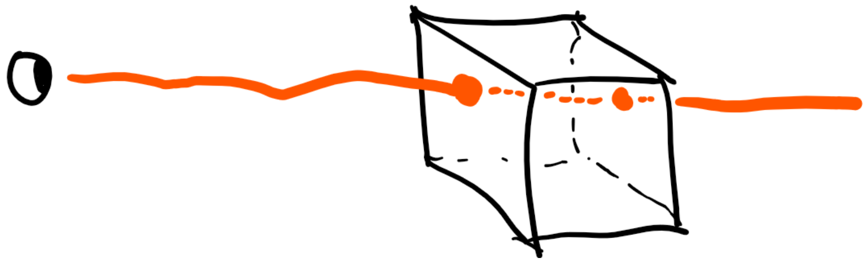

Figure 3. We can reconstruct the $(m+1)$-dimensional scene in an m-dimensional painting by undoing the identification of each sight-line with a point; the result is an embedding of the $(m+1)$-dimensional scene into the principal $\exp(\mathbb{R}E_{\mathrm{gr}})$-bundle $\mathrm{PGL}_{m+1}\mathbb{R}/G_0^{\mathrm{ss}}P_+$ over $\mathrm{PGL}_{m+1}\mathbb{R}/P \cong \mathbb{RP}^m$

for some $\alpha \in (\mathbb{R}^{m+1})^\vee$. Every sight-line for the scene intersects this half-space in a ray, so if we wanted to recreate the $(m+1)$-dimensional scene, then we could just imagine it as occurring in the $\mathbb{R}_+$-bundle over $\mathrm{PGL}_{m+1}\mathbb{R}/P \cong \mathbb{RP}^m$ given by undoing the identification of the ray with a point. This $\mathbb{R}_+$-bundle corresponds to the quotient by $\{\pm 1\}$ of the canonical $\mathbb{R}^\times$-bundle $\mathbb{R}^{m+1} \setminus \{0\}$ over $\mathbb{RP}^m = (\mathbb{R}^{m+1} \setminus \{0\})/\mathbb{R}^\times$; since half-spaces embed into this $\mathbb{R}_+$-bundle by inclusion into $\mathbb{R}^{m+1} \setminus \{0\}$, we lose nothing by assuming the scene happens in this bundle.

Again, $\mathrm{PGL}_{m+1}\mathbb{R}$ acts transitively on the space of sight-lines. The subgroup of $\mathrm{GL}_{m+1}\mathbb{R}$ fixing the sight-line through $\begin{bmatrix} 1 & 0 & \cdots & 0 \end{bmatrix}^\top$ *pointwise* is $\left\{\begin{bmatrix} 1 & \alpha \\ 0 & A \end{bmatrix} : \alpha^\top \in \mathbb{R}^m,\ A \in \mathrm{GL}_m\mathbb{R}\right\}$, and under the quotient by $\mathbb{R}^\times \mathbb{1}$, the image of this subgroup in $\mathrm{PGL}_{m+1}\mathbb{R}$ is precisely $G_0^{\mathrm{ss}}P_+$. Thus, the $\mathbb{R}_+$-bundle over $\mathbb{RP}^m$ given by undoing the identification of sight-lines with points on the canvas is precisely the principal $\exp(\mathbb{R}E_{\mathrm{gr}})$-bundle $\mathrm{PGL}_{m+1}\mathbb{R}/G_0^{\mathrm{ss}}P_+$ over $\mathrm{PGL}_{m+1}\mathbb{R}/P \cong \mathbb{RP}^m$.

In other words, we imagine the $(m+1)$-dimensional scene depicted in an m-dimensional painting as occurring within the space of the principal $\exp(\mathbb{R}E_{\mathrm{gr}})$-bundle $\mathrm{PGL}_{m+1}\mathbb{R}/G_0^{\mathrm{ss}}P_+$, with right-translation by $\exp(tE_{\mathrm{gr}})$ corresponds to moving closer to the eye if $t < 0$ and farther away if $t > 0$. Note that, from the perspective of the eye, we look smaller when we move farther away and bigger when we move closer.

By knowing the position of the canvas, we can also describe a family of affine hyperplanes in $\mathrm{PGL}_{m+1}\mathbb{R}/G_0^{\mathrm{ss}}P_+$ that we imagine to be parallel to the canvas (and, in particular, transverse to the sight-line through each point of the scene). Through each point of the scene is one of these affine hyperplanes, and thinking of ourselves as pedestrians within the scene, we can configure ourselves along the affine hyperplane through our point. The choice of configuration gives an element of the principal G_0^{ss}-bundle $\mathrm{PGL}_{m+1}\mathbb{R}/P_+$

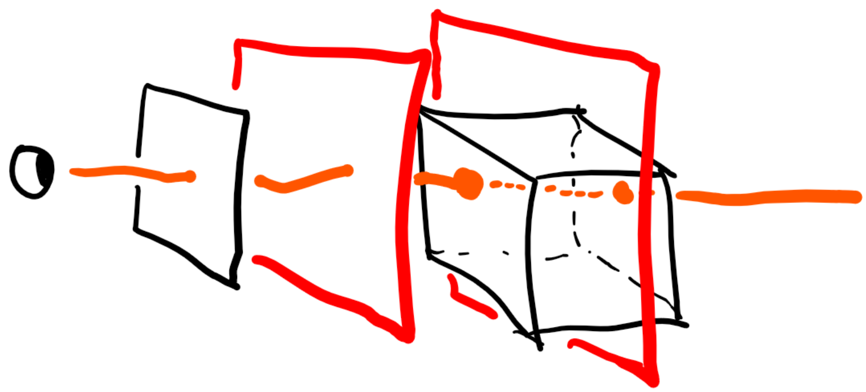

Figure 4. Knowing the positioning of the canvas gives us a family of affine hyperplanes in $\mathrm{PGL}_{m+1}\mathbb{R}/G_0^{\mathrm{ss}}P_+$, and our configuration along one of these hyperplanes corresponds to an element of the principal G_0-bundle $\mathrm{PGL}_{m+1}\mathbb{R}/P_+$ over $\mathrm{PGL}_{m+1}\mathbb{R}/P$

over $\mathrm{PGL}_{m+1}\mathbb{R}/G_0^{\mathrm{ss}}P_+$, the space where the scene takes place; this space $\mathrm{PGL}_{m+1}\mathbb{R}/P_+$ is then also a principal G_0-bundle over $\mathrm{PGL}_{m+1}\mathbb{R}/P$.

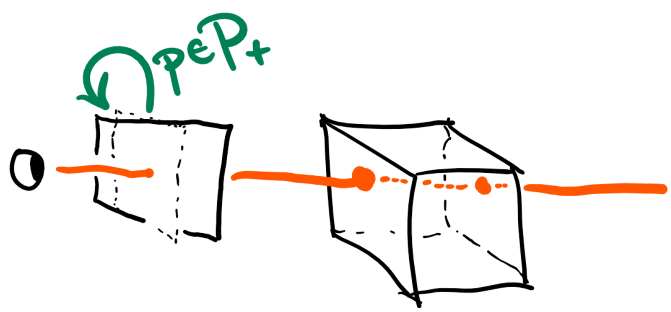

Figure 5. Right-translating by an element of P_+ in $\mathrm{PGL}_{m+1}\mathbb{R}$ amounts to tilting the choice of affine hyperplane through a point in the scene

Finally, different ways of positioning the canvas result in different families of affine hyperplanes. Given an initial choice of hyperplane through a given point, though, all of the other choices of hyperplane can be obtained by "tilting" the initial one. The "unipotent tilts" $p \in P_+$ run through all of the different choices of hyperplane transverse to the sight-line, so the space of choices of hyperplane is the principal P_+-bundle $\mathrm{PGL}_{m+1}\mathbb{R}$ over $\mathrm{PGL}_{m+1}\mathbb{R}/P_+$.

Thus, we have arrived at the principal P-bundle $\mathrm{PGL}_{m+1}\mathbb{R}$ over $\mathrm{PGL}_{m+1}\mathbb{R}/P$. Having built it up from these smaller bundles, we have a fairly good picture of what a configuration within $\mathrm{PGL}_{m+1}\mathbb{R}$ looks like, and from this, it's not hard to see how motion works in this case.

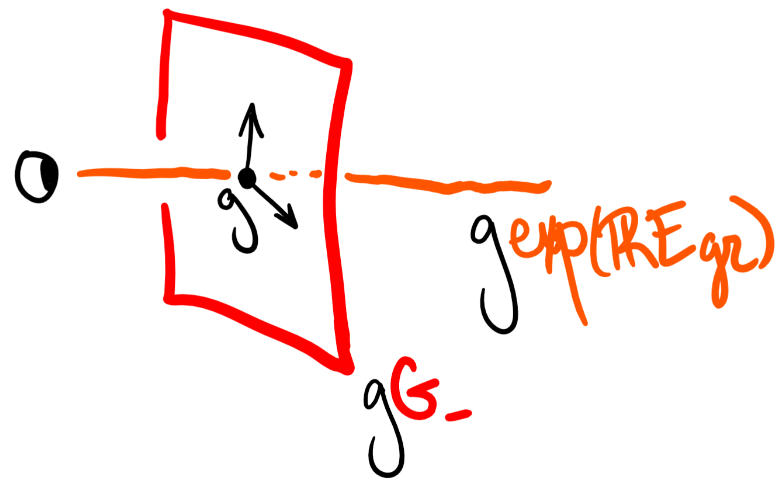

Figure 6. Right-translating by elements of G_{-} amounts to translation within the affine hyperplane through the point in the scene, while right-translating by an element of $\exp(\mathbb{R}E_{\mathrm{gr}})$ amounts to moving along the sight-line

An element $g \in \mathrm{PGL}_{m+1}\mathbb{R}$ determines a choice of affine hyperplane within the $(m+1)$-dimensional scene of an m-dimensional painting, and as we might guess from last time, translation along this affine hyperplane amounts to right-translation by elements of G_{-}. Then, right-translating by elements of G_0^{ss} corresponds to changing our frame within this affine hyperplane, while right-translating by an element of $\exp(\mathbb{R}E_{\mathrm{gr}})$ amounts to moving along the sight-line through our point in the scene; from the perspective of the eye, right-translation by an element of $\exp(\mathbb{R}E_{\mathrm{gr}})$ also corresponds to rescaling the affine hyperplane. Finally, right-translating by an element of P_{+} tilts our choice of affine hyperplane.

12.4. Geodesics

The choice of affine hyperplane described above corresponds to a choice of affine patch in projective space. Inside of a given affine patch are affine geodesics, corresponding to the images $t \mapsto q_P(g\exp(tv))$ for $v \in \mathfrak{g}_{-}$. However, the images of these aren't going to be full copies of geodesics on $\mathbb{RP}^m$ in the sense we'd usually mean, since they are restricted to an affine patch.

A full (unparametrized) geodesic in $\mathbb{RP}^m$ corresponds to a choice of (two-dimensional) plane through the origin inside of $\mathbb{R}^{m+1}$. More specifically, thinking of $\mathbb{RP}^m$ as the space of one-dimensional subspaces of $\mathbb{R}^{m+1}$, a geodesic is the set of all one-dimensional subspaces lying in a given two-dimensional subspace. This is analogous, and in fact related, to the situation in spherical geometry, where we defined great circles to be intersections of the unit sphere with two-dimensional subspaces.

Note, as we did with spherical geometry, that such a definition is geometric for the model: elements of $\mathbb{R}^\times \mathbb{1}$ preserve every subspace of $\mathbb{R}^{m+1}$, and elements of $\mathrm{GL}_{m+1}\,\mathbb{R}$, being invertible linear transformations, send two-dimensional subspaces to two-dimensional subspaces, so $\mathrm{PGL}_{m+1}\,\mathbb{R}$ sends two-dimensional subspaces to two-dimensional subspaces.

Conveniently, the affine geodesics inside a given affine patch are the intersections of full geodesics with that affine patch. Indeed, identifying $\mathrm{PGL}_{m+1}\,\mathbb{R}/P$ with $\mathbb{RP}^m$,

$$q_P\left(g \exp\left(t\left(\begin{smallmatrix} 0 & 0 \\ v & 0 \end{smallmatrix}\right)\right)\right) = g \cdot q_P\left(\left(\begin{smallmatrix} 1 & 0 \\ tv & \mathbb{1} \end{smallmatrix}\right)\right) \in \left\langle g \cdot \left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right), g \cdot \left(\begin{smallmatrix} 0 \\ v \end{smallmatrix}\right)\right\rangle .$$

This situation is noteworthy; in other parabolic model geometries, the geodesics corresponding to one-parameter subgroups generated by elements of $\mathfrak{g}_-$ might not correspond to particularly meaningful curves in the base manifold at all. The corresponding motion in the model group will always be geometrically meaningful though, so it is still worthwhile if we view things from our observer perspective.

Chapter 13

Conformal Geometry and Tanaka Prolongation

Again, rewrite this so that it looks better and fits well within the book. Also, consider whether moving the general Tanaka prolongation procedure to the next chapter is better, since it is more relevant there?

Last time, we briefly explored projective geometry, which we thought of as the geometry of sight-lines. As we shall soon see, there are many similarities between the models for projective and conformal geometry, many of which can be found in all parabolic model geometries.

However, there are some key differences as well. Perhaps foremost among these differences is that the model for conformal geometry is significantly less obvious than in projective geometry. Recall that a *conformal structure* on a manifold M is an equivalence class $[\mathrm{g}]$, corresponding to all Riemannian metrics of the form $f\mathrm{g}$ for some (smooth) function $f : M \to \mathbb{R}_+$. Deciding what should count as the stabilizer of this structure is a bit trickier than just preserving sight-lines, especially if we don't necessarily know what the base manifold should be either.

Fortunately, this difficulty also gives us a convenient opportunity to showcase a useful algebraic construction called *Tanaka prolongation*, which happens to solve this issue. As such, our plan for the lecture is as follows:

- Motivate what Tanaka prolongation does geometrically
- Explain the construction in the case of conformal geometry

- Describe conformal motion from an observer perspective
- (Appendix) Sketch how Tanaka prolongation works in general

By the end of this lecture, we should have a good idea of what the model for conformal geometry looks like. While there are additional nuances to more general parabolic models, many of the main ideas are similar, so the reader will hopefully be prepared to encounter other parabolic geometries on their own. In the next lecture, we will finally see what a Cartan connection is, and why they're so easy to work with.

13.1. From similarity to conformality

Let us imagine that we don't already know what our model (G, P) for conformal geometry should be. Where's a good place to start exploring what this model could be?

We're looking for a Lie group G corresponding to symmetries that preserve a Riemannian metric up to scale. As such, a good initial candidate would be the Lie group of similarity transformations of $\mathbb{R}^m$. A *similarity transformation* of $\mathbb{R}^m$ is an affine transformation that preserves the underlying Euclidean metric up to scale. The Lie group of such transformations is isomorphic to $\mathbb{R}^m \rtimes \mathbb{R}_+ \operatorname{O}(m)$, where $\mathbb{R}_+ \operatorname{O}(m)$ is the group of linear transformations that are positive scalar multiples of orthogonal transformations; for $(u, A) \in \mathbb{R}^m \rtimes \mathbb{R}_+ \operatorname{O}(m)$, as with Euclidean geometry, the action on $\mathbb{R}^m$ is just $(u, A) \cdot v = u + A(v)$. The geometry of the model $(\mathbb{R}^m \rtimes \mathbb{R}_+ \operatorname{O}(m), \mathbb{R}_+ \operatorname{O}(m))$ is called *similarity geometry.*

This geometry looks a lot like Euclidean geometry from an observer perspective. Thinking of $\mathbb{R}^m \rtimes \mathbb{R}_+ \operatorname{O}(m) \simeq \operatorname{I}(m) \rtimes \mathbb{R}_+$ as a bundle of perspectives for ourselves as observers within the geometry, it's essentially the same as the Euclidean case except that we can now also right-translate by elements of the subgroup $\mathbb{R}_+ \mathbb{1}$ to rescale ourselves.

Let us suggestively denote by G_- the subgroup of translations $\mathbb{R}^m$, and by G_0 the subgroup $\mathbb{R}_+ \operatorname{O}(m)$, so that the model for similarity geometry is $(G_- G_0, G_0)$.

Inside the Lie subalgebra $\mathfrak{g}_0$, we can preemptively denote by E_{gr} the element with one-parameter subgroup $\exp(tE_{\text{gr}}) = e^{-t}\mathbb{1}$; note that $\operatorname{ad}_{E_{\text{gr}}}$ restricts to multiplication by -1 on $\mathfrak{g}_-$ and vanishes on $\mathfrak{g}_0$. We can also define a homomorphism $\lambda : G_0 \to \mathbb{R}_+$ given by $rA \mapsto |r|$ for $A \in \operatorname{O}(m)$ and $r \in \mathbb{R}^\times$. The kernel of this homomorphism is just $\operatorname{O}(m)$, and we can decompose G_0 as $\exp(\mathbb{R}E_{\text{gr}}) \ker(\lambda) = \mathbb{R}_+ \operatorname{O}(m)$.

Fixing an inner product g_0 on $T_0\mathbb{R}^m$, corresponding to the "usual" one for Euclidean geometry, we can determine a new inner product $\varphi \cdot \mathrm{g}_0$ on

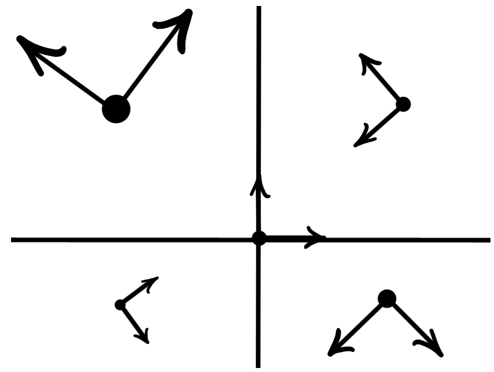

Figure 1. As with Euclidean geometry, we can think of $\mathbb{R}^m \rtimes \mathbb{R}_+ \operatorname{O}(m)$ as a bundle of perspectives for ourselves as observers inside similarity geometry

$T_{\varphi(0)}\mathbb{R}^m = T_{q_{G_0}(\varphi)}(G_-G_0/G_0)$ for each $\varphi \in G_-G_0$ by

$$\varphi \cdot \mathrm{g}_0(v, w) := \mathrm{g}_0(\varphi_*^{-1}(v), \varphi_*^{-1}(w)).$$

In particular, the "usual" Riemannian metric for Euclidean geometry is given by $\mathrm{g}_x := x \cdot \mathrm{g}_0$ for each $x \in \mathbb{R}^m = G_-$.

Since G_- is a normal subgroup of G_-G_0 and $G_- \cap G_0 = \{e\}$, we have a natural quotient homomorphism $\pi_{G_-} : G_-G_0 \to G_0$. When used together with the homomorphism $\lambda : G_0 \to \mathbb{R}_+$, this gives us a convenient way of describing the inner product $\varphi \cdot \mathrm{g}_0$ for arbitrary $\varphi \in G_-G_0$: the subgroup $G_- \ker(\lambda) = \mathrm{I}(m)$ acts by isometries, so

$$\varphi \cdot \mathrm{g}_0 = \lambda(\pi_{G_-}(\varphi))^{-2}\mathrm{g}_{\varphi(0)}.$$

Because a Riemannian metric $\tilde{\mathrm{g}}$ conformal to g is, by definition, of the form $\tilde{\mathrm{g}} = f\mathrm{g}$ for some (smooth) function $f : \mathbb{R}^m \to \mathbb{R}_+$, we can identify a choice of metric conformal to the Euclidean one with a choice of section $\sigma_f : G_-G_0/G_0 \to G_-G_0/\ker(\lambda)$. Explicitly, for each $x \in G_-$, we define $\sigma_f(x) := x(\lambda^{-1}(\sqrt{f(x)}))^{-1}$, so that

$$\sigma_f(x) \cdot \mathrm{g}_0 = \lambda(\lambda^{-1}(\sqrt{f(x)})^{-1})^{-2}\mathrm{g}_x = f(x)\mathrm{g}_x = \tilde{\mathrm{g}}_x.$$

In other words, similar to what we saw with projective geometry last time, we can think of the conformal structure on the base manifold G_-G_0/G_0 as something that "lives" in the principal $\exp(\mathbb{R}E_{\mathrm{gr}})$-bundle $G_-G_0/\ker(\lambda)$ over G_-G_0/G_0.

Intuitively, this means that the fibers of $G_-G_0/\ker(\lambda)$ over the base manifold G_-G_0/G_0 are like the sight-lines from projective geometry. Indeed, the key invariant of conformal geometry is the choice of metric up to scale, and by the above, these fibers essentially correspond to the space of all

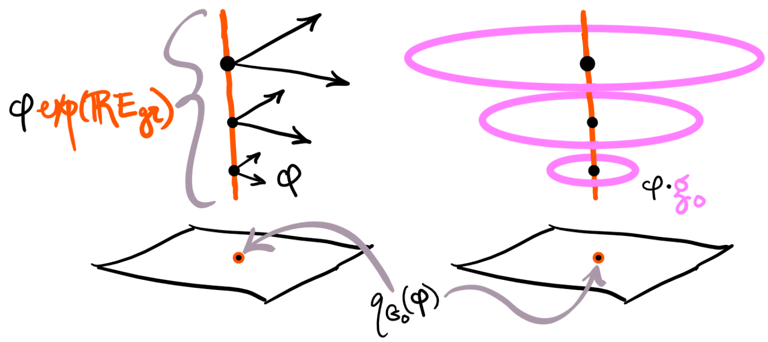

Figure 2. Different elements of $G_-G_0/\ker(\lambda)$ over a given point in G_-G_0/G_0 correspond to different choices of inner product over that point, depicted here as the unit disks determined by these inner products

such choices over a given point, so it makes sense that these are what the geometry wants to keep preserved.

Now, we want to give ourselves a notion of "conformal frame", so that we can meaningfully place ourselves inside of this geometry. As with projective geometry, it is convenient to build up these frames in steps. To start, we consider the principal $\exp(\mathbb{R}E_{\mathrm{gr}})$-bundle $G_-G_0/\ker(\lambda)$ over G_-G_0/G_0, which we think of as the space where the conformal structure actually lives and whose elements correspond to choices of scale for the Euclidean metric at the underlying point on the base manifold. From here, we can naturally include the orthonormal frames for the metric at each choice of scale; this amounts to moving up to the principal $\ker(\lambda)$-bundle G_-G_0 over the space $G_-G_0/\ker(\lambda)$, which (unsurprisingly) makes G_-G_0 into a principal G_0-bundle over the base manifold G_-G_0/G_0. Here, $\ker(\lambda)$ accounts for stabilizer motion that preserves the metric and scale, and $\exp(\mathbb{R}E_{\mathrm{gr}})$ lets us rescale directly; neither of these changes the fact that motion from G_- preserves the scale. Thus, the final step is to include "higher-order frames" corresponding to changes of perspective where motion from G_- can alter the scale.

This leads to an obvious question: what can such "higher-order frames" be?

13.2. Tanaka prolongation in the conformal case

Recall that, given a parabolic model (G, P), we can often think of the geometry of (G_-G_0, G_0) as a kind of affine analogue of the geometry of (G, P). *Tanaka prolongation* gives a way to reverse this analogy, so that given the "affine version", we can (usually) build the corresponding parabolic structure.

Back in the conformal case, let's work at the level of Lie algebras. We already have a Lie algebra $\mathfrak{g}_- + \mathfrak{g}_0$, where $\mathfrak{g}_-$ is the subalgebra of translations $\mathbb{R}^m$ and $\mathfrak{g}_0$ is the subalgebra $\mathbb{R}\mathbb{1} + \mathfrak{o}(m)$. Writing $\mathfrak{g}_{-1} := \mathfrak{g}_-$ and $\mathfrak{g}_{-\ell} := \{0\}$ for all $\ell > 1$, this gives a (somewhat boring) graded Lie algebra structure on $\mathfrak{g}_- + \mathfrak{g}_0$: for all $i, j \leq 0$, we have $[\mathfrak{g}_i, \mathfrak{g}_j] \subseteq \mathfrak{g}_{i+j}$. Algebraically, the goal of Tanaka prolongation is to extend this to a new graded Lie algebra $\mathfrak{g} := \mathfrak{g}_{-1} + \mathfrak{g}_0 + \mathfrak{g}_1 + \cdots$. From the above, we know that the subalgebra $\sum_{\ell>0} \mathfrak{g}_\ell$, which we will preemptively call $\mathfrak{p}_+$, should correspond to changes of perspective that allow motion from $\mathfrak{g}_-$ to change the scale.

To build this new Lie algebra, let us start by considering what $\mathfrak{g}_1$ must do to $\mathfrak{g}_-$: to maintain the graded structure, we need to have $[\mathfrak{g}_1, \mathfrak{g}_{-1}] \subseteq \mathfrak{g}_0$. Moreover, because we want the end result to be a Lie algebra, it should satisfy the Jacobi identity, so for $v, w \in \mathfrak{g}_{-1}$ and $\alpha \in \mathfrak{g}_1$, we should have

$$0 = [\alpha, [v, w]] = [[\alpha, v], w] + [v, [\alpha, w]].$$

Deciding[1] that $\alpha \in \mathfrak{g}_1$ should be uniquely determined by the action of ad_α on $\mathfrak{g}_-$, we can therefore identify all of the possible choices for elements of $\mathfrak{g}_1$ with the space of linear maps $\alpha \in \mathfrak{g}_-^\vee \otimes \mathfrak{g}_0$ such that $[\alpha(v), w] + [v, \alpha(w)] = \alpha(v)w - \alpha(w)v = 0$ for all $v, w \in \mathfrak{g}_-$.

For context, let us give some algebraic definitions. Given an arbitrary Lie algebra $\mathfrak{h}$ and $\mathfrak{h}$-representation V, define

$$\partial : \mathfrak{h}^\vee \otimes V \to \Lambda^2\mathfrak{h}^\vee \otimes V$$

by $\partial\alpha(X \wedge Y) := X \cdot \alpha(Y) - Y \cdot \alpha(X) - \alpha([X, Y])$. With this, we can further define the space $\mathrm{Der}(\mathfrak{h}; V) := \{\alpha \in \mathfrak{h}^\vee \otimes V : \partial(\alpha) = 0\}$ of *derivations* from $\mathfrak{h}$ to V. In our case, we are identifying $\mathfrak{g}_1$ with the subspace of $\mathrm{Der}(\mathfrak{g}_-; \mathfrak{g}_- + \mathfrak{g}_0)$ with images contained in $\mathfrak{g}_0$.

The space of all linear maps $\alpha : \mathfrak{g}_- \to \mathfrak{g}_0$ satisfying

$$\alpha(v)w - \alpha(w)v = 0$$

for all $v, w \in \mathfrak{g}_-$ is determined by the $\dim(\mathfrak{g}_-^\vee)\dim(\Lambda^2\mathfrak{g}_-)$ independent linear equations

$$\mathrm{g}_0(e_k, \alpha(e_i)e_j - \alpha(e_j)e_i) = 0.$$

In other words, we can identify the component $\mathfrak{g}_1$ with a subspace of $\mathfrak{g}_-^\vee \otimes \mathfrak{g}_0$ of dimension

$$\begin{aligned}
\dim(\mathfrak{g}_1) &= \dim(\mathfrak{g}_-^\vee \otimes \mathfrak{g}_0) - \dim(\mathfrak{g}_-^\vee)\dim(\Lambda^2\mathfrak{g}_-) \\
&= \dim(\mathfrak{g}_-^\vee)(\dim(\mathfrak{g}_0) - \dim(\Lambda^2\mathfrak{g}_-)) \\
&= \dim(\mathfrak{g}_-^\vee).
\end{aligned}$$

[1] Adding in central elements isn't interesting in this case; if $[\alpha, \mathfrak{g}_{-1}] = \{0\}$, then it isn't doing anything.

Indeed, it turns out there is a convenient identification between $\mathfrak{g}_-^\vee$ and $\mathfrak{g}_1$: for $\alpha \in \mathfrak{g}_-^\vee$, the corresponding map $\mathrm{ad}_\alpha|_{\mathfrak{g}_-}$ is given by

$$v \mapsto -\alpha(v)\mathbb{1} - \alpha \otimes v + \mathrm{g}_0(v, \cdot) \otimes \alpha^\sharp,$$

where $\alpha^\sharp \in \mathfrak{g}_-$ is the unique element such that $\mathrm{g}_0(\alpha^\sharp, \cdot) = \alpha$. In other words, $\mathrm{ad}_\alpha(v)w = -\alpha(v)w - \alpha(w)v + \mathrm{g}_0(v, w)\alpha^\sharp$.

Now that we have $\mathfrak{g}_1$, we can try the same thing with $\mathfrak{g}_2$. We want elements $\beta \in \mathfrak{g}_2$ to satisfy $0 = \beta([v, w]) = [\beta(v), w] + [v, \beta(w)]$, with $\beta(v), \beta(w) \in \mathfrak{g}_1 \approx \mathfrak{g}_-^\vee$ to preserve the graded structure. From here, the computations get a bit heinous, but the key thing to note is that, when $\dim(\mathfrak{g}_-) > 2$, we must have $\mathfrak{g}_2 = \{0\}$, so the construction stops with $\mathfrak{g}_{-1} + \mathfrak{g}_0 + \mathfrak{g}_1$. Thinking of $\alpha \in \mathfrak{g}_1$ as a linear map from $\mathfrak{g}_-$ to $\mathfrak{g}_0$, we define

$$[\alpha, R] = \alpha \circ R - \mathrm{ad}_R \circ \alpha$$

for each $R \in \mathfrak{g}_0$ and define $\mathfrak{g}_1$ to be abelian. This gives $\mathfrak{g} = \mathfrak{g}_{-1} + \mathfrak{g}_0 + \mathfrak{g}_1$ a graded Lie algebra structure, and a faithful representation of $\mathfrak{g}$ is given by

$$(v, r\mathbb{1} + R, \alpha) \mapsto \begin{bmatrix} -r & (\alpha^\sharp)^\top & 0 \\ v & R & -\alpha^\sharp \\ 0 & -v^\top & r \end{bmatrix}.$$

Letting G be the Lie group with Lie algebra $\mathfrak{g}$ such that $G_- G_0 \leq G$ and G/G_0 is connected, we can define $P := G_0 P_+$ for P_+ the connected subgroup generated by $\mathfrak{p}_+ = \mathfrak{g}_1$, and our model for conformal geometry becomes (G, P). It is not too difficult to check that G is isomorphic to $\mathrm{PO}(1, m+1)$, with corresponding quadratic form Q on $\mathbb{R}^{m+2}$ given by

$$Q\left(\begin{bmatrix} x_0 \\ x \\ x_{m+1} \end{bmatrix}\right) = -2x_0 x_{m+1} + \sum_{i=1}^{m} x_i^2.$$

What about when $\dim(\mathfrak{g}_-) = m = 2$? Well, in that case, $\mathfrak{g}_2$, and more generally each $\mathfrak{g}_\ell$ with $\ell > 1$, is not trivial, so the $\mathfrak{g}$ we construct is infinite-dimensional. This is one of the issues with Tanaka prolongation: sometimes, the information you put into it isn't sufficient to return a finite-dimensional Lie algebra. We could probably have expected some sort of problem here, though; in dimension two, all holomorphic maps are conformal wherever their derivatives don't vanish, so we were never going to construct a finite-dimensional model symmetry algebra for two-dimensional conformal geometry.

13.3. Moving around in the conformal sphere

Thinking of G as the Lie group of $\mathrm{PO}(\mathbb{R}^{m+2}, Q)$, with Q as in the previous section, we can see that G acts transitively on the projectivized null-cone

$\{\langle x\rangle \in \mathbb{RP}^{m+1} : Q(x) = 0\}$, and that

$$\mathrm{Stab}_G\begin{pmatrix}1\\0\\0\end{pmatrix} = \left\{\begin{pmatrix}a & p & -p(p^\top)/2\\ 0 & A & -p^\top\\ 0 & 0 & a^{-1}\end{pmatrix} : a \in \mathbb{R}^\times,\ p^\top \in \mathbb{R}^m,\ A \in \mathrm{O}(m)\right\}$$

corresponds to the closed subgroup P, so we can identify G/P with the projectivized null-cone.

Topologically, this projectivized null-cone is a sphere. To see this, let $x \in Q^{-1}(0) \setminus \{0\}$. We're trying to understand the projectivized null-cone, so we only care about x up to scale. In particular, we can rescale x so that $x_0 + x_{m+1} = \sqrt{2}$. Under this rescaling, we can define $y := \frac{x_0 - x_{m+1}}{\sqrt{2}}$, so that

$$1 + y = \frac{(x_0 + x_{m+1}) + (x_0 - x_{m+1})}{\sqrt{2}} = \sqrt{2}x_0$$

and

$$1 - y = \frac{(x_0 + x_{m+1}) - (x_0 - x_{m+1})}{\sqrt{2}} = \sqrt{2}x_{m+1}.$$

Thus, $1 - y^2 = (1+y)(1-y) = 2x_0x_{m+1}$, so

$$Q(x) = -2x_0x_{m+1} + \sum_{i=1}^{m} x_i^2 = -1 + y^2 + \sum_{i=1}^{m} x_i^2 = 0,$$

hence we can identify the projectivized null-cone with the space with $y^2 + \sum_{i=1}^{m} x_i^2 = 1$, namely the m-sphere. Because of this, we call G/P the *conformal sphere.*

Alternatively, we can think of G/P as the one-point compactification of $G_-G_0/G_0 \cong \mathbb{R}^m$. The subgroup G_- takes the form

$$\left\{\begin{pmatrix}1 & 0 & 0\\ v & \mathbb{1} & 0\\ -v^\top v/2 & -v^\top & 1\end{pmatrix} : v \in \mathbb{R}^m\right\},$$

so

$$G_-\begin{pmatrix}1\\0\\0\end{pmatrix} = \left\{\begin{pmatrix}1\\ v\\ -v^\top v/2\end{pmatrix} : v \in \mathbb{R}^m\right\}$$

gives us our open cell. The complement of this open cell inside the projectivized null-cone is the single point $\left(\begin{smallmatrix}0\\0\\1\end{smallmatrix}\right)$, which we can think of as the "point at infinity" for the copy of Euclidean space given by the open cell.

As for movement within G, things are much the same as with projective geometry. At each configuration $g \in G$, we determine a copy $q_P(gG_-)$ of Euclidean space corresponding to motion from G_-. Using $\exp(\mathbb{R}E_{\mathrm{gr}})$, we can control the scale of that copy of Euclidean space, and with $\ker(\lambda)$, we can move amongst the different orthonormal frames for that point in Euclidean space. Finally, the subgroup P_+ gives us "unipotent tilts", which let us tilt between different copies of Euclidean space through our underlying point $q_P(g) \in G/P$.

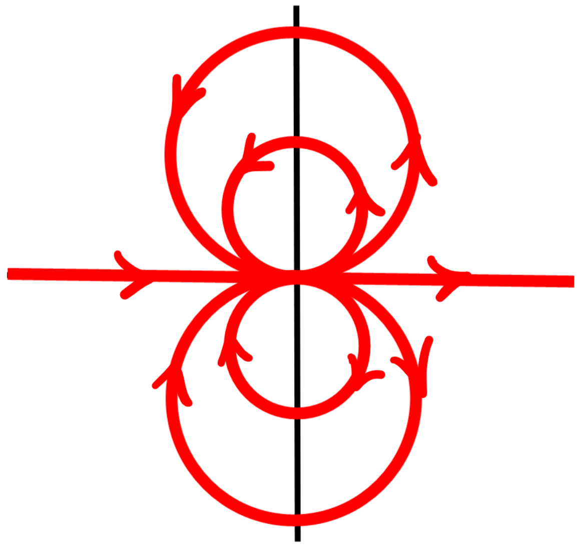

Figure 3. In conformal geometry, the trajectories of motion from one-parameter subgroups in G_- are not uniquely determined by an initial velocity in the base manifold

Perhaps the main difference here is that motion from G_- no longer determines consistent curves on the base manifold up to reparametrization. However, again, the motion is consistent and meaningful inside of G.

Appendix: Tanaka prolongation in general

Both conformal and projective geometry are $|1|$-graded, meaning that the grading of $\mathfrak{g}$ determined by a Cartan involution θ and the parabolic $\mathfrak{p}$ is of the form $\mathfrak{g}_{-1} + \mathfrak{g}_0 + \mathfrak{g}_1$. This makes many algebraic aspects of these geometries fairly simplistic compared to the general case. In particular, Tanaka prolongation is a bit more involved when there are multiple negative grading components.

Let's imagine we have a graded nilpotent Lie algebra

$$\mathfrak{g}_- = \mathfrak{g}_{-k} + \cdots + \mathfrak{g}_{-1},$$

with $[\mathfrak{g}_{-i}, \mathfrak{g}_{-j}] \subseteq \mathfrak{g}_{-i-j}$, together with another Lie algebra $\mathfrak{g}_0$ that both acts on $\mathfrak{g}_-$ by derivations—so $R \cdot [v, w] = [R \cdot v, w] + [v, R \cdot w]$ for all $R \in \mathfrak{g}_0$ and $v, w \in \mathfrak{g}_-$—and preserves the grading—so $R \cdot \mathfrak{g}_{-i} \subseteq \mathfrak{g}_{-i}$ for every $R \in \mathfrak{g}_0$ and each $i \geq 1$. We consider the semidirect sum $\mathfrak{g}_- \ni \mathfrak{g}_0$, which we will write as just $\mathfrak{g}_- + \mathfrak{g}_0$; what is the corresponding Tanaka prolongation?

In essence, we follow the same idea as before: build up positive grading components piece by piece. To start, we want $\mathfrak{g}_1$, consisting of elements that act as derivations that send each $\mathfrak{g}_{-i}$ to $\mathfrak{g}_{-i+1}$. In other words, we are looking for

$$\mathfrak{g}_1 := \{\alpha \in \mathrm{Der}(\mathfrak{g}_-; \mathfrak{g}_- + \mathfrak{g}_0) : \alpha(\mathfrak{g}_{-i}) \subseteq \mathfrak{g}_{-i+1} \text{ for each } i > 0\}.$$

Naturally, $\mathfrak{g}_- + \mathfrak{g}_0 + \mathfrak{g}_1$ is a representation of $\mathfrak{g}_- + \mathfrak{g}_0$: given $\alpha \in \mathfrak{g}_1$, we have $v \cdot \alpha = -\alpha(v)$ for $v \in \mathfrak{g}_-$ and $R \cdot \alpha = \mathrm{ad}_R \circ \alpha - \alpha \circ \mathrm{ad}_R$ for $R \in \mathfrak{g}_0$.

Next, we want to build a grading component $\mathfrak{g}_2$ of degree 2, so that its elements act as derivations sending each $\mathfrak{g}_{-i}$ to $\mathfrak{g}_{-i+2}$. Symbolically, this means

$$\mathfrak{g}_2 := \{\beta \in \mathrm{Der}(\mathfrak{g}_-; \mathfrak{g}_- + \mathfrak{g}_0 + \mathfrak{g}_1) : \beta(\mathfrak{g}_{-i}) \subseteq \mathfrak{g}_{-i+2} \text{ for each } i > 0\}.$$

Again, $\mathfrak{g}_- + \mathfrak{g}_0 + \mathfrak{g}_1 + \mathfrak{g}_2$ is a representation of $\mathfrak{g}_- + \mathfrak{g}_0$: for $\beta \in \mathfrak{g}_2$, $v \cdot \beta = -\beta(v)$ and $R \cdot \beta = \mathrm{ad}_R \circ \beta - \beta \circ \mathrm{ad}_R$ as before.

We continue this process, recursively defining

$$\mathfrak{g}_\ell := \left\{\zeta \in \mathrm{Der}\left(\mathfrak{g}_-; \sum_{j<\ell} \mathfrak{g}_j\right) : \zeta(\mathfrak{g}_{-i}) \subseteq \mathfrak{g}_{-i+\ell} \text{ for each } i > 0\right\}$$

for each $\ell > 0$. Letting $\mathfrak{g}$ be the representation of $\mathfrak{g}_- + \mathfrak{g}_0$ given by the sum

$$\mathfrak{g} := \mathfrak{g}_- + \mathfrak{g}_0 + \sum_{\ell>0} \mathfrak{g}_\ell$$

of all of these grading components, we imbue it with a Lie algebra structure as follows. First, the bracket agrees with the bracket of $\mathfrak{g}_- + \mathfrak{g}_0$ when restricted there, and for each $\alpha \in \sum_{\ell>0} \mathfrak{g}_\ell$ and $X \in \mathfrak{g}_- + \mathfrak{g}_0$, $[X, \alpha] := X \cdot \alpha$, where $\cdot$ denotes the representation action of $\mathfrak{g}_- + \mathfrak{g}_0$ on $\mathfrak{g}$. From here, we want to continue defining the bracket in a way that satisfies the Jacobi identity, so that

$$[\alpha, \beta](v) = [[\alpha, \beta], v] = [[\alpha, v], \beta] + [\alpha, [\beta, v]] = [\alpha(v), \beta] + [\alpha, \beta(v)]$$

for $\alpha, \beta \in \sum_{\ell>0} \mathfrak{g}_\ell$ and $v \in \mathfrak{g}_-$. Conveniently, this gives us a way to construct the bracket recursively as well: for $\alpha, \alpha' \in \mathfrak{g}_1$, we define $[\alpha, \alpha'] \in \mathfrak{g}_2 \subseteq \mathrm{Der}(\mathfrak{g}_-; \mathfrak{g}_- + \mathfrak{g}_0 + \mathfrak{g}_1)$ to be the unique element of the form

$$[\alpha, \alpha'](v) = [\alpha(v), \alpha'] + [\alpha, \alpha'(v)]$$

for each $v \in \mathfrak{g}_-$. Since $\alpha(v), \alpha'(v) \in \mathfrak{g}_- + \mathfrak{g}_0$, we already know that $[\alpha(v), \alpha'] = \alpha(v) \cdot \alpha'$ and $[\alpha, \alpha'(v)] = -\alpha'(v) \cdot \alpha$, so this bracket is well-defined. Then, for arbitrary $\beta \in \mathfrak{g}_i$ and $\zeta \in \mathfrak{g}_j$, we can recursively define $[\beta, \zeta] \in \mathfrak{g}_{i+j}$ to be the unique element such that

$$[\beta, \zeta](v) = [\beta(v), \zeta] + [\beta, \zeta(v)]$$

for each $v \in \mathfrak{g}_-$. Since $\beta(v) \in \sum_{\ell<i} \mathfrak{g}_\ell$ and $\zeta(v) \in \sum_{\ell<j} \mathfrak{g}_\ell$, if we know how to form brackets with elements in grading components of lesser degree, then these brackets are well-defined as well.

Thus, we get a Lie algebra $\mathfrak{g}$. If $\mathfrak{g}_\ell = \{0\}$ for every ℓ greater than some k, then $\mathfrak{g}$ is finite-dimensional and $\mathfrak{p}_+ := \sum_{\ell>0} \mathfrak{g}_\ell$ is a nilpotent subalgebra. Because $[\mathfrak{g}_i, [\mathfrak{g}_j, \mathfrak{g}_\ell]] \subseteq \mathfrak{g}_{i+j+\ell}$, we must have $\mathfrak{z}(\mathfrak{g}_i, \mathfrak{g}_j) = \{0\}$ unless $i + j = 0$, so if $\mathfrak{g}$ is semisimple, then $\mathfrak{g}_-$ and $\mathfrak{p}_+$ must be $\mathfrak{z}$-dual. In particular, if $\mathfrak{g}$ is semisimple, then $\mathfrak{p} := \mathfrak{g}_0 + \mathfrak{p}_+$ must satisfy $\mathfrak{p}^\perp = \mathfrak{p}_+$, hence $\mathfrak{p}$ must be parabolic.

Chapter 14

Strictly Pseudoconvex CR Geometry

Appendix A

Locally Homogeneous Geometric Structures

[Explain how to convert between flat Cartan geometries and locally homogeneous geometric structures]

Bibliography

[1] Calderbank, D. M. J., Noppakaew, P.: *Parabolic subalgebras, parabolic buildings and parabolic projection*. (2017) arXiv:1607.00370

[2] Čap, A., Slovák, J.: *Parabolic Geometries I: Background and General Theory*. Mathematical Surveys and Monographs, vol. 154. Amer. Math. Soc., Providence, RI (2009)

[3] Eberlein, P. B.: *Geometry of Nonpositively Curved Manifolds*. Chicago Lectures in Mathematics, University of Chicago Press (1996)

[4] Hilgert, J., Neeb, K.-H.: *Structure and Geometry of Lie Groups*. Springer Monographs in Mathematics, Springer-Verlag New York (2012)

[5] Sharpe, R. W.: *Differential Geometry: Cartan's Generalization of Klein's Erlangen Program*. Graduate Texts in Mathematics, vol. 166, Springer-Verlag, New York (1997)

[6] Spivak, M.: *A Comprehensive Introduction to Differential Geometry*. Vol. 1, 3rd Edition, Publish or Perish, Houston (1999)

[7] Starkov, A. N.: *Dynamical Systems on Homogeneous Spaces*. Translations of Mathematical Monographs, vol. 190. Amer. Math. Soc., Providence, RI (2000)