PARABOLIC GEOMETRIES
FOR PEOPLE THAT LIKE PICTURES

LECTURE 1: HOW DO WE MOVE INSIDE A LIE GROUP?

JACOB W. ERICKSON

As we will see later, Cartan geometries are a sophisticated way of making a principal bundle over a manifold resemble a particular Lie group. Therefore, if we want to talk about geometric intuition for these things, then we really need to start with Lie groups. In this lecture, we will answer some basic questions regarding visualization of Lie groups, including:

- What do left-translation and right-translation look like?
- What does conjugation look like?
- What is the Maurer-Cartan form?
- Why is the Maurer-Cartan form so easy to use for geometry?

By the end of this lecture, 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 lecture, we will practice using this to do geometry.

1. Picturing the group of Euclidean isometries

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

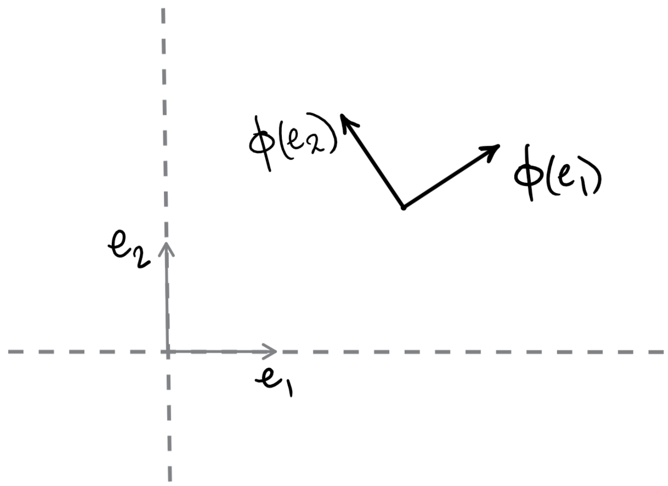

Figure 1. We can depict an orthonormal frame ϕ on $\mathbb{R}^2$ by the pair of tangent vectors $(\phi(e_1), \phi(e_2))$

Consider the plane $\mathbb{R}^2$ with the usual Euclidean structure. An *orthonormal frame* over $u \in \mathbb{R}^2$ is just a linear isometry ϕ from $\mathbb{R}^2 \approx T_0\mathbb{R}^2$ to the tangent space $T_u\mathbb{R}^2$. 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 as the pair of tangent vectors $(\phi(e_1), \phi(e_2))$ on the plane.

While using pairs of vectors is useful for drawing 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. Thus, 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.

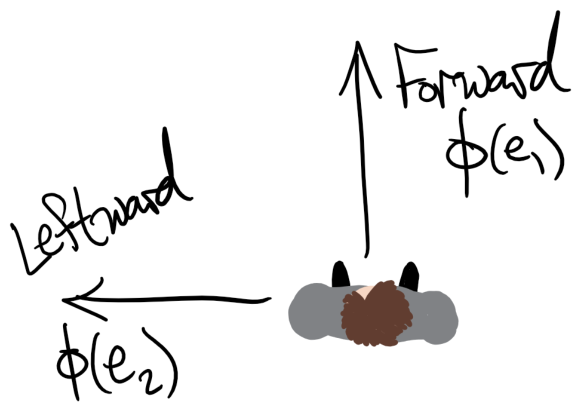

FIGURE 2. Each orthonormal frame corresponds to a unique configuration for ourselves as pedestrians on the Euclidean plane

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

[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.

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

$\phi \in \mathrm{I}(2)$ 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

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

Elements of the orthogonal group $\mathrm{O}(2)$ 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 $\phi \in \mathrm{I}(2)$ 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.

2. Transformation and Motion

Throughout, we will use L_a to denote left-translation $g \mapsto ag$ and R_a to denote right-translation $g \mapsto ga$. 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.

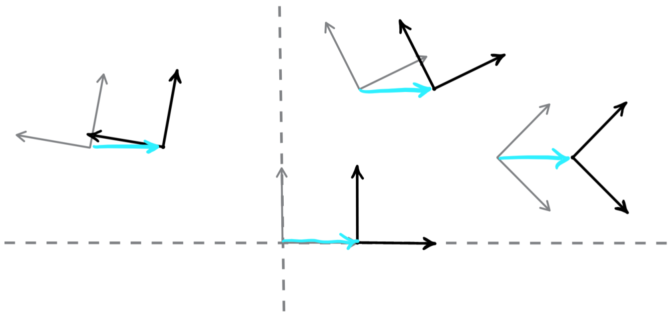

FIGURE 3. Left-translating by τ_{e_1} shifts all orthonormal frames uniformly by e_1, as if we were applying it as a transformation to the plane and the orthonormal frames were thought of as being inside of the plane

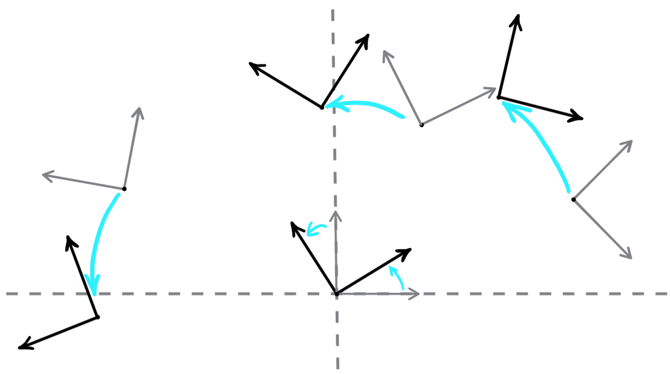

FIGURE 4. Left-translating by the rotation given by $\mathrm{rot}(\theta) := \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} \in \mathrm{O}(2)$ uniformly rotates all orthonormal frames by θ around 0

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 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 does 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

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

Using translations to identify each tangent space of $\mathbb{R}^2$ with the tangent space at the identity, this means that right-translation by τ_{e_1} shifts each orthonormal frame by the vector to which that orthonormal frame maps e_1.

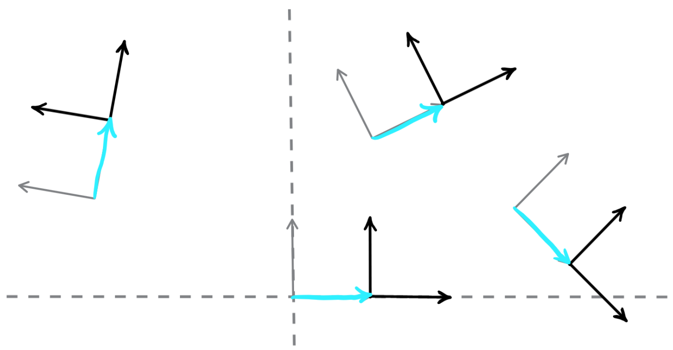

FIGURE 5. Right-translating by τ_{e_1} shifts each orthonormal frame by the vector to which that orthonormal frame maps e_1

Under the 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**.

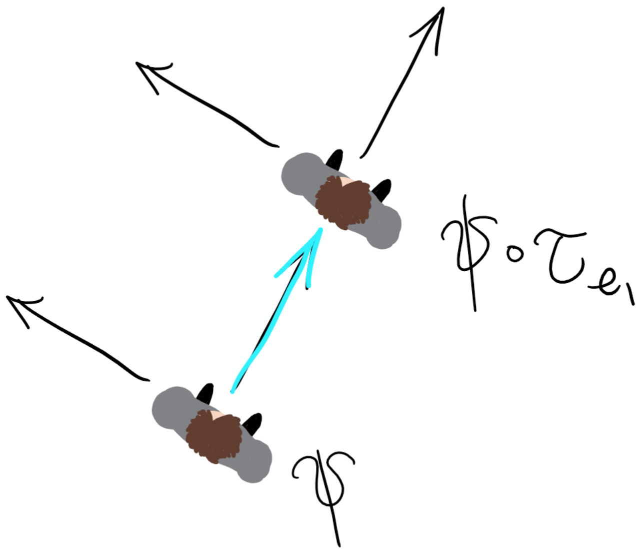

FIGURE 6. Right-translating by τ_{e_1} corresponds to walking forward one unit

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, right-translation by the rotation $\mathrm{rot}(\theta) := \left[\begin{smallmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{smallmatrix}\right]$ just turns us on the spot by θ, rather than necessarily rotating us around 0.

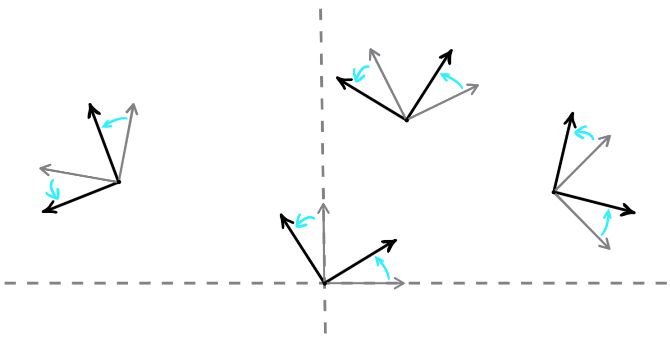

FIGURE 7. Right-translating by $\mathrm{rot}(\theta)$ turns each orthonormal frame on the spot by θ

Let us try to summarize this intuition in a few words. Under left-translation, elements act as *transformations*, effecting everything uniformly according to the action on the homogeneous space $\mathrm{I}(2)/\mathrm{O}(2)$. Under right-translation, 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.

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 u.

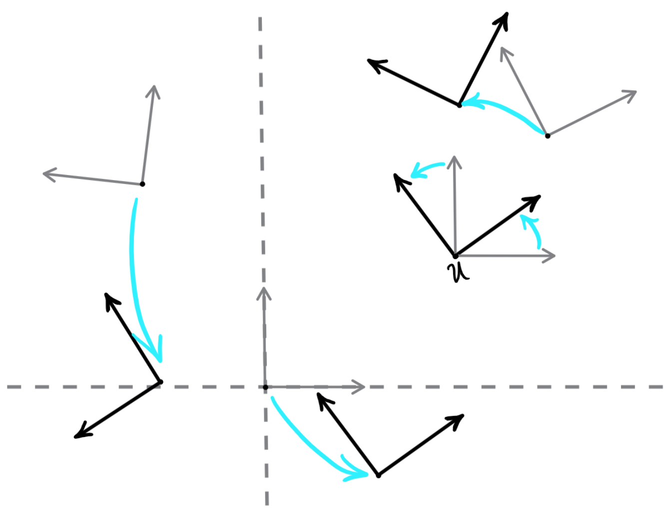

Figure 8. $\tau_u \mathrm{rot}(\theta)\tau_u^{-1}$ rotates the plane by θ around u

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 $\mathrm{rot}(\theta) \cdot u$.

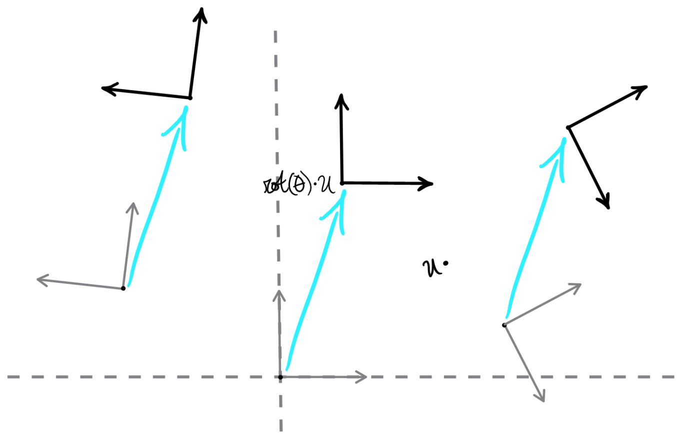

FIGURE 9. $\mathrm{rot}(\theta)\tau_u\mathrm{rot}(\theta)^{-1}$ shifts the plane by the vector to which the orthonormal frame corresponding to $\mathrm{rot}(\theta)$ sends u

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.

4. The Maurer-Cartan form: your new best friend

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. 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 $\mathrm{I}(2)$ is the tangent space of $\mathrm{I}(2)$ at the identity element (together with a bracket operation that we'll talk about later). To describe tangent spaces at other points, we use something called the *Maurer-Cartan form*. While its definition appears to be just algebraic formalism, do not be fooled: the Maurer-Cartan form is one of the most deeply intuitive objects in modern differential geometry.

Definition 4.1. The *Maurer-Cartan form* of a Lie group G is the $\mathfrak{g}$-valued one-form ω_G given by

$$(\omega_G)_g : T_gG \to \mathfrak{g} = T_eG, \;\; X_g \mapsto \mathrm{L}_{g^{-1}*} X_g$$

at each $g \in G$.

Let us try to elucidate what the Maurer-Cartan form is trying to tell us. For $g \in \mathrm{I}(2)$, consider the tangent vector $\mathrm{L}_{g*}(e_1) \in T_g\,\mathrm{I}(2)$; this is tangent to the curve $t \mapsto g \exp(te_1)$ at $t = 0$, so since right-translating by $\exp(te_1) = \tau_{te_1}$ simply corresponds to walking forward with unit speed for time t, the tangent vector $\mathrm{L}_{g*}(e_1)$ gives "unit forward velocity" at g. By definition, the Maurer-Cartan form maps $\mathrm{L}_{g*}(e_1)$ to

$$\omega_{\mathrm{I}(2)}(\mathrm{L}_{g*}(e_1)) = \mathrm{L}_{g^{-1}*}(\mathrm{L}_{g*}(e_1)) = e_1,$$

which is the "unit forward velocity" at the identity element. The Maurer-Cartan form, in other words, takes the "unit forward velocity" at each $g \in \mathrm{I}(2)$ and identifies it with the "unit forward velocity" in the Lie algebra, so that we have a constant notion of "unit forward velocity" that does not depend on which tangent space we are currently at. Of course, we obtain similar results for "unit leftward velocity", "unit positive angular velocity", and more generally, any element of the Lie algebra. Thus, **the Maurer-Cartan form $\omega_{\mathrm{I}(2)}$ is a canonical coframing of a Lie group in terms of motion**, identifying each tangent space of $\mathrm{I}(2)$ with the Lie algebra $\mathfrak{i}(2)$ in a way that is consistent with how elements of the Lie algebra determine one-parameter subgroups of motions.

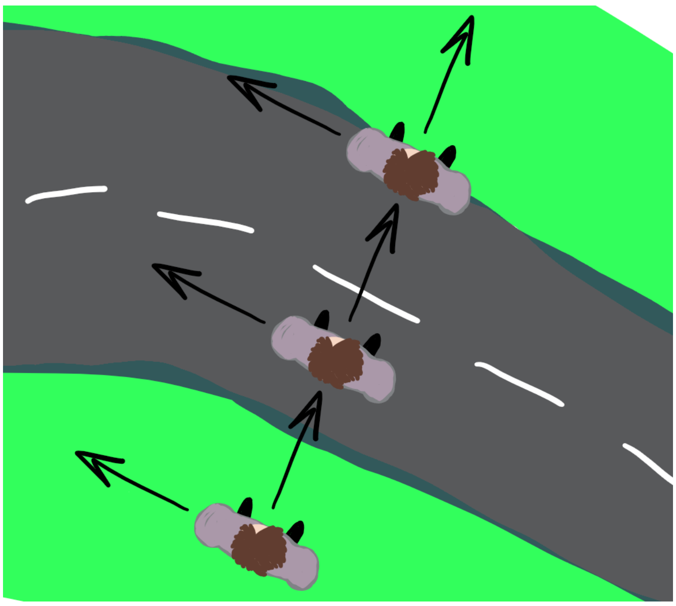

FIGURE 10. While walking across a street, a pedestrian is "walking forward (without turning)"

You, the reader, are likely far more familiar with this idea than you realize. For example, think about the last time you crossed a street; the path you took was probably a geodesic segment, right? We could try to convince ourselves that we did this with a Riemannian metric, first by using it to construct a covariant derivative, an operator that takes in tangent vectors and outputs a differential operator on the infinite-dimensional space of vector fields, and then using this covariant derivative to construct a path in which all of the acceleration is tangent to the path, but this seems unlikely with a bit of introspection. Really, when you think about it, you are just telling yourself “walk forward (without turning) until you get to the other side”, and this is precisely what a Maurer-Cartan form allows us to do: at each point in time along the path, you are specifying that your (translational) velocity is “forward” and your angular velocity is 0.

Of course, the problem with this is that our ambient geometry is not inherently homogeneous, so we do not necessarily get a Maurer-Cartan form. On the other hand, whatever it is that we are using to think about the geometry of the world around us, it does *look like* a Maurer-Cartan form, and as it turns out, this is precisely the kind of thing we will be studying in this course.