PARABOLIC GEOMETRIES FOR PEOPLE THAT LIKE PICTURES

LECTURE 4: WHAT IS A (HOMOGENEOUS) TRACTOR BUNDLE?

JACOB W. ERICKSON

In the last two lectures, we talked about how symmetries determine geometries, so that Lie-theoretic properties of G and $H \leq G$ determine geometric properties of the model (G, H). Of course, many important properties of Lie groups come from their representations; in this lecture, we will start exploring representation theory through the lens of the model $(\mathrm{SL}_2 \mathbb{R}, B)$, which we described in the warm-up. Along the way, we will attempt to do the following:

- Introduce tractor bundles
- Review the representation theory of $\mathrm{SL}_2 \mathbb{R}$
- Learn a geometric construction for the finite-dimensional irreducible representations of $\mathrm{SL}_2 \mathbb{C}$
- Foreshadow some stuff about Lie algebra homology and BGG operators

The model geometry $(\mathrm{SL}_2 \mathbb{R}, B)$ is our first example of a *parabolic geometry*. In order to talk about other parabolic geometries, we will need to understand what semisimple Lie groups look like, so in the next lecture, we will take our experience here with $\mathrm{SL}_2 \mathbb{R}$ to explain what the Killing form is trying to tell us, which will give us a better picture of general semisimple Lie groups.

1. Homogeneous vector bundles

Recall from last time that, in a model geometry (G, H), geometric objects were precisely those things that were preserved under an action induced by the natural left-action of G. We focused primarily on the action of G on G/H and on G itself, but there are other objects that seem like they should probably be called "geometric", like the tangent bundle of G/H.

Indeed, the tangent bundle *is* geometric in the sense we've described, albeit in a kind of boring, tautological way: if we take our action to be given by the pushforward of the left-action of G on G/H, then the tangent bundle is obviously preserved under this action.

This perspective is actually much more useful than it originally sounds, because it allows us to describe every tangent space of G/H

as some left-translate of the tangent space $T_{eH}(G/H) \approx \mathfrak{g}/\mathfrak{h}$ at the identity coset eH. Not every element of G moves the tangent space at the identity coset, though: if we left-translate by an element of H, then we stay in the tangent space at the identity coset. In particular, this determines an action of H on $T_{eH}(G/H)$, which coincides with the induced action on $\mathfrak{g}/\mathfrak{h}$ from the adjoint representation of G restricted to H.

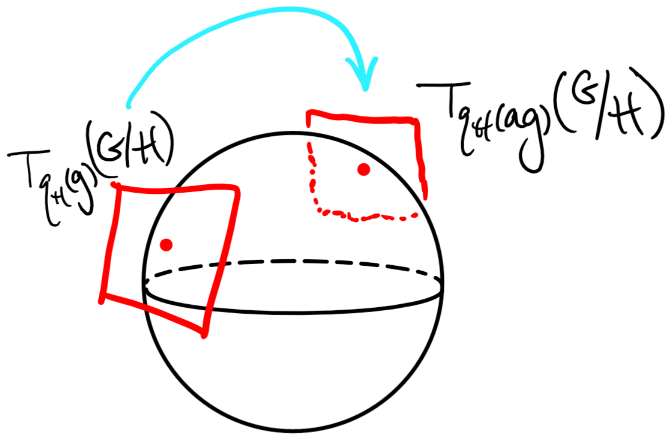

FIGURE 1. Pushing tangent spaces around using left-translation

More generally, whenever we have an H-representation V, we can consider the corresponding *homogeneous vector bundle.*

Definition 1.1. The *homogeneous vector bundle* corresponding to an H-representation V is the associated vector bundle $G \times_H V$ over G/H, whose elements are ordered pairs $(g, v) \in G \times V$ modulo the H-action $h \star (g, v) = (gh^{-1}, h \cdot v)$, and whose bundle map $\pi_V : G \times_H V \to G/H$ is given by $(g, v) \mapsto q_H(g)$.

As a homogeneous vector bundle over G/H, we can identify the tangent bundle $T(G/H)$ as $G \times_H \mathfrak{g}/\mathfrak{h}$. This essentially just restates the description above: we take the tangent space at the identity coset and push it to the rest of G/H using the left-action of G, and acting on the tangent space at the identity coset by an element of H corresponds to the representation action of H on $\mathfrak{g}/\mathfrak{h}$, since for each $h \in H$ and $X \in \mathfrak{g}$, $(h, X + \mathfrak{h}) = (e, \mathrm{Ad}_h(X) + \mathfrak{h})$ as elements of $G \times_H \mathfrak{g}/\mathfrak{h}$.

It is worth thinking about what this looks like for the tangent bundle $\mathrm{I}(2) \times_{\mathrm{O}(2)} \mathfrak{i}(2)/\mathfrak{o}(2)$ of the Euclidean plane. Here, $\mathfrak{i}(2)/\mathfrak{o}(2) \approx \mathbb{R}^2$, with the induced action of $\mathrm{O}(2)$ corresponding to the usual action of $\mathrm{O}(2)$ on $\mathbb{R}^2$. Each isometry $g \in \mathrm{I}(2)$ determines a unique linear isomorphism from $T_0\mathbb{R}^2 \approx \mathfrak{i}(2)/\mathfrak{o}(2) \approx \mathbb{R}^2$ to the tangent space $T_{g(0)}\mathbb{R}^2$; this is precisely the orthonormal frame perspective from before, with the vector $(g, v) \in \mathrm{I}(2) \times_{\mathrm{O}(2)} \mathbb{R}^2$ corresponding to the tangent vector $g_*(v) \in T_{g(0)}\mathbb{R}^2$. In particular, somewhat tautologically, we can think of different orthonormal frames over the same point of the plane as different framings of the tangent space over that point.

Because tangent vectors over G/H "live" on G/H, they do not really interact with right-translation by elements of the isotropy group H. However, our *descriptions* of them, from our perspective as observers, do change under right-translation by H. To see what we mean by this, consider the tangent vector $(g, v) \in \mathrm{I}(2)\times_{\mathrm{O}(2)}\mathbb{R}^2$ corresponding to $g_*(v)$. By picking the representative (g, v) for this tangent vector, we are describing it from the perspective of g. If we want to describe the same tangent vector from a rotated perspective $g\mathrm{rot}(\theta)$, then our *description* of the tangent vector must be rotated in the opposite direction because the tangent vector itself didn't change. This is where the equivalence relation on the pairs comes from: $(g\mathrm{rot}(\theta), \mathrm{rot}(\theta)^{-1} \cdot v)$ is the same tangent vector as (g, v), just described from the perspective of $g\mathrm{rot}(\theta)$ rather than g.

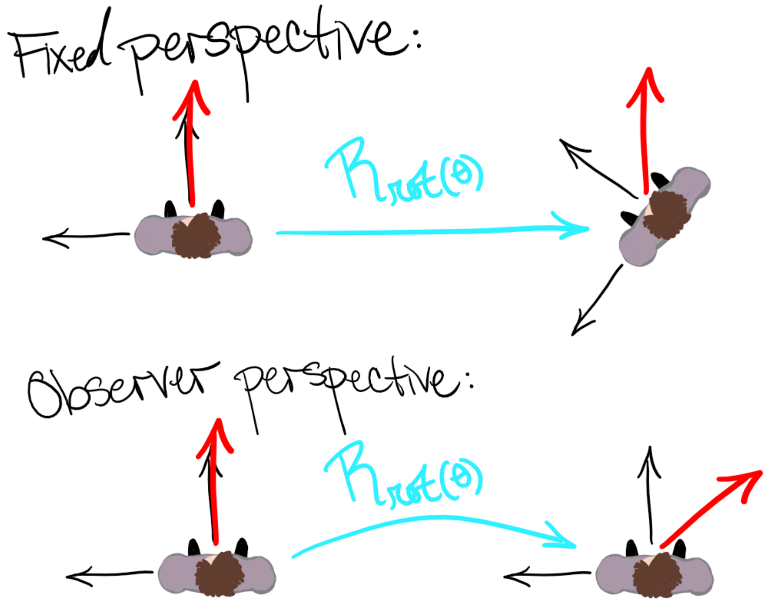

FIGURE 2. Tangent vectors over G/H stay in the same place when we change perspective, so if we rotate on the spot, then they appear to rotate in the opposite direction; in other words, $(g, v) = (g\mathrm{rot}(\theta), \mathrm{rot}(\theta)^{-1} \cdot v)$

This intuition, that $(g, v) \in G \times_H V$ is the vector that an observer at the element $g \in G$ would describe as $v \in V$, gives us a particularly easy way to deal with sections of homogeneous vector bundles. By definition, a section of the homogeneous vector bundle $G \times_H V$ is a map $\sigma : G/H \to G \times_H V$ such that $\pi_V \circ \sigma$ is the identity map on G/H. For a section σ, though, let us define $\tilde{\sigma} : G \to V$ to be the map that takes each element $g \in G$ to its description of $\sigma(q_H(g))$ in V, so that if $\sigma(q_H(g)) = (g, v)$, then $\tilde{\sigma}(g) = v$.

First, note that $\tilde{\sigma}$ uniquely determines σ: given $\tilde{\sigma}$, we can simply define σ by $q_{_H}(g) \mapsto (g, \tilde{\sigma}(g))$. Second, $\tilde{\sigma}$ is an H-equivariant map, so that $\tilde{\sigma}(gh) = h^{-1} \cdot \tilde{\sigma}(g)$ for each $h \in H$, because

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

and whenever we have an H-equivariant map $f : G \to V$, it uniquely determines a section by $q_{_H}(g) \mapsto (g, f(g))$, the same way $\tilde{\sigma}$ determined σ. Thus, we can identify sections of $G \times_H V$ with their descriptions as H-equivariant maps from G to V. This is convenient, since we have an abundance of tools to deal with maps from manifolds into vector spaces.

2. Parallelism on tractor bundles

As we saw above, given an H-representation V, we have a corresponding homogeneous vector bundle $G \times_H V$ for (G, H). When this H-representation comes from the restriction $\rho|_H : H \to \mathrm{GL}(V)$ of a G-representation $\rho : G \to \mathrm{GL}(V)$, we call such a homogeneous vector bundle a *(homogeneous) tractor bundle*[1] for (G, H).

As an example, consider again the tangent bundle $\mathrm{I}(2) \times_{\mathrm{O}(2)} \mathbb{R}^2$ of the Euclidean plane. Because the subgroup of translations is normal, we have a natural representation $\rho : \mathrm{I}(2) \to \mathrm{O}(2) < \mathrm{GL}_2 \mathbb{R}$ given by the quotient map $\tau_u \circ A \mapsto A$, and $\rho|_{\mathrm{O}(2)}$ coincides with the usual representation of $\mathrm{O}(2)$ on $\mathbb{R}^2$, so *in Euclidean geometry*, the tangent bundle is a tractor bundle.

Careful!
I am definitely *not* saying that the tangent bundle is a tractor bundle for other model geometries. Indeed, it usually isn't!

Recall that two tangent vectors on the Euclidean plane were said to be *parallel* whenever one of them was the image of the other under the pushforward of a translation. Let us reexamine what this looks like when we consider the tangent bundle as a tractor bundle for $(\mathrm{I}(2), \mathrm{O}(2))$.

Suppose we have two vectors (ϕ, v) and (ψ, w) in $\mathrm{I}(2) \times_{\mathrm{O}(2)} \mathbb{R}^2$, corresponding to $\phi_*(v)$ and $\psi_*(w)$ in $T\mathbb{R}^2$, respectively. For the tangent vector $\phi_*(v)$ to be parallel to $\psi_*(w)$, there must be some $u \in \mathbb{R}^2$ such that $(\tau_u)_*(\phi_*(v)) = \psi_*(w)$; in terms of elements of the homogeneous vector bundle, this just means that $\tau_u \cdot (\phi, v) = (\tau_u \circ \phi, v) = (\psi, w)$.

Note that this doesn't necessarily mean that v has to be equal to w, since $(\phi \circ A, \rho(A)^{-1}v) = (\phi, v)$ for every $A \in \mathrm{O}(2)$. To check whether $(\tau_u \circ \phi, v)$ and (ψ, w) are equal, we need to find $A \in \mathrm{O}(2)$ such that $\tau_u \circ \phi \circ A = \psi$, and then checking whether $(\tau_u \circ \phi, v) = (\tau_u \circ \phi \circ A, \rho(A)^{-1}v)$ and (ψ, w) are equal amounts to checking whether $\rho(A)^{-1}v = w$.

[1]Later, we will also use the term "tractor bundle" to refer to the vector bundles analogous to these homogeneous ones in the correspondence between Cartan geometries and their models.

This is still rather inelegant though, so let us try to make a few simplifications. First, note that we can rewrite $\tau_u \circ \phi \circ A$ as

$$(\phi \circ \phi^{-1}) \circ \tau_u \circ \phi \circ A = \phi \circ (\phi^{-1} \circ \tau_u \circ \phi) \circ A = \phi \circ \tau_{\rho(\phi)^{-1}(u)} \circ A,$$

so $\tau_u \circ \phi \circ A = \psi$ if and only if $\tau_{\rho(\phi)^{-1}(u)} \circ A = \phi^{-1} \circ \psi$. Second, recall that translations lie in the kernel of ρ, so $\rho(A) = \rho(\tau_{\rho(\phi)^{-1}(u)} \circ A)$. In particular, $\rho(A) = \rho(\phi^{-1} \circ \psi)$, so

$$(\tau_u \circ \phi \circ A, \rho(A)^{-1}v) = (\phi \circ (\phi^{-1} \circ \psi), \rho(\phi^{-1} \circ \psi)^{-1}v).$$

With these simplifications, we now see that (ϕ, v) and (ψ, w) are parallel if and only if $\rho(\phi^{-1} \circ \psi)^{-1}v = w$. More generally, the same idea allows us to define parallelism on every (homogeneous) tractor bundle.

Definition 2.1. Suppose (g_1, v_1) and (g_2, v_2) are vectors in the (homogeneous) tractor bundle $G \times_H V$. Then, (g_1, v_1) and (g_2, v_2) are *parallel* if and only if $\rho(g_1^{-1}g_2)^{-1}v_1 = v_2$, or equivalently, if and only if $\rho(g_1)v_1 = \rho(g_2)v_2$. Moreover, we say that a section σ of $G \times_H V$ is *parallel* if and only if the corresponding H-equivariant map satisfies $\tilde{\sigma}(g) = \rho(g)^{-1}\tilde{\sigma}(e)$ for all $g \in G$.

Because parallel sections σ are uniquely determined by the value of $\tilde{\sigma}(e) \in V$, we get a copy of V inside of the space of sections of $G \times_H V$ by sending $v \in V$ to the section corresponding to the H-equivariant map $\tilde{v} : g \mapsto \rho(g)^{-1}v$. Moreover, for $a, g \in G$,

$$(a \cdot \tilde{v})(g) = \tilde{v}(a^{-1}g) = \rho(a^{-1}g)^{-1}v = \rho(g)^{-1}(\rho(a)v) = (\widetilde{\rho(a)v})(g),$$

so the action of G on sections of $G \times_H V$ induced by the natural left-action of G on itself coincides with the representation action of G on V when we restrict to parallel sections. In particular, the notion of parallel section on a tractor bundle is geometric.

3. Review of the representation theory of SL_2

The classification of the irreducible representations of SL_2 by highest weights is arguably one of the most well-known ideas in mathematics, so I'm going to speed through most of this. There are tons of good explanations of these ideas elsewhere for the uninitiated; I personally like [3] and [4].

To start, every representation of $\mathrm{SL}_2\mathbb{R}$ induces a representation of $\mathfrak{sl}_2\mathbb{R}$, so we can restrict our attention to $\mathfrak{sl}_2\mathbb{R}$. We have a useful basis $\{Y_-, Z, Y_+\}$ for $\mathfrak{sl}_2\mathbb{R}$ given by

$$Y_- = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}, \ Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \text{ and } Y_+ = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix},$$

with brackets $[Z, Y_+] = 2Y_+$, $[Z, Y_-] = -2Y_-$, and $[Y_+, Y_-] = Z$.

Because ad_Z is diagonalizable over $\mathbb{R}$ on $\mathfrak{sl}_2\mathbb{R}$, it turns out that the image $\rho(Z)$ of Z under every finite-dimensional representation $\rho : \mathfrak{sl}_2\mathbb{R} \to \mathfrak{gl}(V)$ is diagonalizable over $\mathbb{R}$ as well. Similarly, because

ad_{Y_+} and ad_{Y_-} are nilpotent on $\mathfrak{sl}_2\mathbb{R}$, the images $\rho(Y_+)$ and $\rho(Y_-)$ are also nilpotent.

Now, suppose we have a finite-dimensional representation V of $\mathfrak{sl}_2\mathbb{R}$. Because $\rho(Z)$ is diagonalizable over $\mathbb{R}$, we can decompose V as a sum of eigenspaces for $\rho(Z)$. Let us denote by V_α the eigenspace with eigenvalue α.

For each eigenvector $v \in V_\alpha$, we have

$$\begin{aligned}\rho(Z)(\rho(Y_+)(v)) &= \rho(Z)\rho(Y_+)(v) - \rho(Y_+)\rho(Z)(v) + \rho(Y_+)\rho(Z)(v)\\ &= [\rho(Z), \rho(Y_+)](v) + \rho(Y_+)(\rho(Z)(v))\\ &= \rho([Z, Y_+])(v) + \rho(Y_+)(\alpha v)\\ &= 2\rho(Y_+)(v) + \alpha\rho(Y_+)(v) = (\alpha+2)\rho(Y_+)(v),\end{aligned}$$

so $\rho(Y_+)(V_\alpha) \subseteq V_{\alpha+2}$. Similarly, $\rho(Y_-)(V_\alpha) \subseteq V_{\alpha-2}$.

Because $\rho(Y_+)$ is nilpotent on the representation V, or alternatively because V is finite-dimensional, there is some $\rho(Z)$-eigenvector v in some V_λ such that $\rho(Y_+)(v) = 0$. We call such a vector v a *highest weight vector* of V. Using that

$$\begin{aligned}\rho(Y_+)\rho(Y_-)^k(v) &= \rho(Y_+)\rho(Y_-)\rho(Y_-)^{k-1}(v) - \rho(Y_-)\rho(Y_+)\rho(Y_-)^{k-1}(v)\\ &\quad + \rho(Y_-)\rho(Y_+)\rho(Y_-)^{k-1}(v)\\ &= \rho([Y_+, Y_-])\rho(Y_-)^{k-1}(v) + \rho(Y_-)\left(\rho(Y_+)\rho(Y_-)^{k-1}(v)\right)\\ &= \rho(Z)\rho(Y_-)^{k-1}(v) + \rho(Y_-)\left(\rho(Y_+)\rho(Y_-)^{k-1}(v)\right)\end{aligned}$$

for each $k \geq 1$, we can inductively show that

$$\rho(Y_+)\rho(Y_-)^k(v) = k(\lambda - k + 1)\rho(Y_-)^{k-1}(v).$$

In particular, the subspace of V spanned by the vectors

$$v, \rho(Y_-)(v), \rho(Y_-)^2(v), \ldots$$

is a $\mathfrak{sl}_2\mathbb{R}$-subrepresentation of V.

We know that $\rho(Y_-)$ is nilpotent on V as well, so there must be some smallest n such that $\rho(Y_-)^n(v) = 0$. By the formula for $\rho(Y_+)\rho(Y_-)^k(v)$ above, we have

$$0 = \rho(Y_+)(\rho(Y_-)^n(v)) = n(\lambda - n + 1)\rho(Y_-)^{n-1},$$

so $n = \lambda + 1$. In particular, λ is a non-negative integer.

The long and short of it is that each non-negative integer λ determines a unique irreducible representation $V(\lambda)$ with a highest weight vector of $\rho(Z)$-eigenvalue λ, and every finite-dimensional irreducible representation of $\mathrm{SL}_2\,\mathbb{R}$ is of the form $V(\lambda)$ for some λ. Again, this is fairly basic representation theory, so I'm just going to assume that we're all at least a bit familiar with this and move on to the fun stuff.

4. Parallelism for $(\mathrm{SL}_2\,\mathbb{R}, B)$

Let us start by looking at the tractor bundle corresponding to the "usual" representation $V(1) = \mathbb{R}^2$ of $\mathrm{SL}_2\,\mathbb{R}$. As before, we take a copy of $\mathbb{R}^2$ at the identity coset $q_B(\mathbb{1}) \cong \left(\begin{smallmatrix}1\\0\end{smallmatrix}\right) \in \mathbb{RP}^1 \cong \mathrm{SL}_2\,\mathbb{R}/B$ and push it around to all the other points of $\mathbb{RP}^1$.

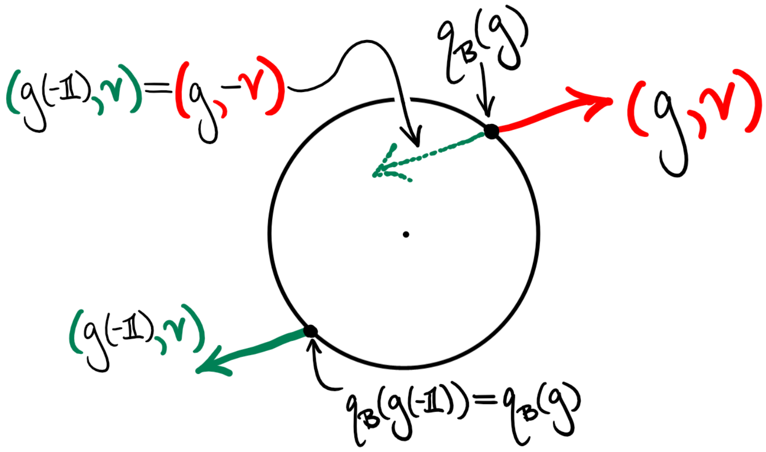

Figure 3. Depicting $\mathbb{RP}^1$ as the circle with antipodal points identified, the vectors (g, v) over one point and $(g(-\mathbb{1}), v)$ over the antipodal point are negatives of each other in the vector bundle $\mathrm{SL}_2\,\mathbb{R} \times_B \mathbb{R}^2$, since $-\mathbb{1} \in B$ and $(g(-\mathbb{1}), v) = (g, -\mathbb{1} \cdot v) = (g, -v)$

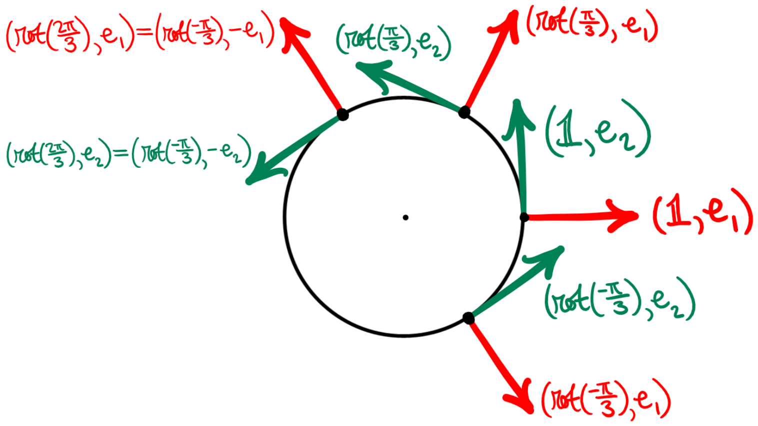

Figure 4. Drawing of some vectors in $\mathrm{SL}_2\,\mathbb{R} \times_B \mathbb{R}^2$ of the form $(\mathrm{rot}(\theta), e_1)$ and $(\mathrm{rot}(\theta), e_2)$ over the projective line $\mathbb{RP}^1$; note that the vectors point in the directions indicated by the "frames" $\mathrm{rot}(\theta) \in \mathrm{SL}_2\,\mathbb{R}$

Of course, the point of having a tractor bundle is that we can use parallelism to get a geometric version of the representation $V(1)$: for each $v \in \mathbb{R}^2$, we can construct the parallel section corresponding to the B-equivariant map $\tilde{v} : \mathrm{SL}_2\,\mathbb{R} \to \mathbb{R}^2$ given by $A \mapsto A^{-1}(v)$, and as before, we have the relation $\widetilde{A(v)} = A \cdot \tilde{v}$.

As examples, for $\{e_1, e_2\}$ our usual basis for $\mathbb{R}^2$, we have

$$\widetilde{e_1} : \begin{bmatrix} a & b \\ c & d \end{bmatrix} \mapsto \begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} e_1 = \begin{bmatrix} d & -b \\ -c & a \end{bmatrix} e_1 = \begin{bmatrix} d \\ -c \end{bmatrix}$$

and

$$\widetilde{e_2} : \begin{bmatrix} a & b \\ c & d \end{bmatrix} \mapsto \begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} e_2 = \begin{bmatrix} -b \\ a \end{bmatrix}.$$

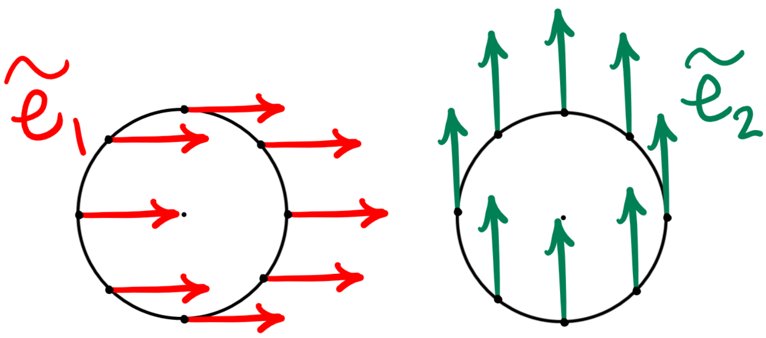

FIGURE 5. Depiction of the parallel sections corresponding to $\widetilde{e_1}$ and $\widetilde{e_2}$ over the projective line $\mathbb{RP}^1$; note how they always point in the same direction that e_1 and e_2 usually point

Every element of $\mathrm{SL}_2\,\mathbb{R}$ is in a coset of the form $\mathrm{rot}(\theta)B$ for some $\theta \in \mathbb{R}/2\pi\mathbb{Z}$, so B-equivariant functions are uniquely determined by their values on $\mathrm{SO}(2)$. In particular, this means that we can think of each B-equivariant function f on $\mathrm{SL}_2\,\mathbb{R}$ as a function $f : \theta \mapsto f(\theta)$. In the case of $V(1) \approx \mathbb{R}^2$, the parallel section corresponding to the vector $v = \left[\begin{smallmatrix} v_1 \\ v_2 \end{smallmatrix}\right] \in \mathbb{R}^2$ is given by

$$\begin{aligned} \tilde{v}(\theta) := \tilde{v}(\mathrm{rot}(\theta)) &= \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}^{-1} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} \\ &= \begin{bmatrix} \cos(\theta)v_1 + \sin(\theta)v_2 \\ -\sin(\theta)v_1 + \cos(\theta)v_2 \end{bmatrix}. \end{aligned}$$

Notice that this is of the form $\left[\begin{smallmatrix} -\partial_\theta y \\ y \end{smallmatrix}\right]$, so that we only really need to specify y, the coordinate of the *lowest weight space* of $\mathbb{R}^2$, to determine the section.

Let's do a bit of representation theory to see whether we can make sense of this. Consider a finite-dimensional representation V of $\mathrm{SL}_2\,\mathbb{R}$. The *Borel subalgebra* $\mathfrak{b}$ corresponding to the closed subgroup B is given by $\mathfrak{b} = \langle Z, Y_+ \rangle$. Inside of $\mathfrak{b}$ is an ideal $\mathfrak{b}_+ := \langle Y_+ \rangle$, whose elements send each weight space V_k to the weight space V_{k+2}. Because $\mathfrak{b}_+$ is preserved by the adjoint action of B, the representation ρ of $\mathrm{SL}_2\mathbb{R}$ makes both V and $\rho(\mathfrak{b}_+)V$ into B-representations, so that we may consider the quotient B-representation

$$\mathrm{H}_0(\mathfrak{b}_+; V) := V/\rho(\mathfrak{b}_+)V,$$

also known as the *zeroth homology group* of $\mathfrak{b}_+$ with coefficients in V, which corresponds to the sum of the *lowest weight spaces* of the representation V. We can sort of think of $\mathrm{H}_0(\mathfrak{b}_+; V)$ as the "tangential part" of V.

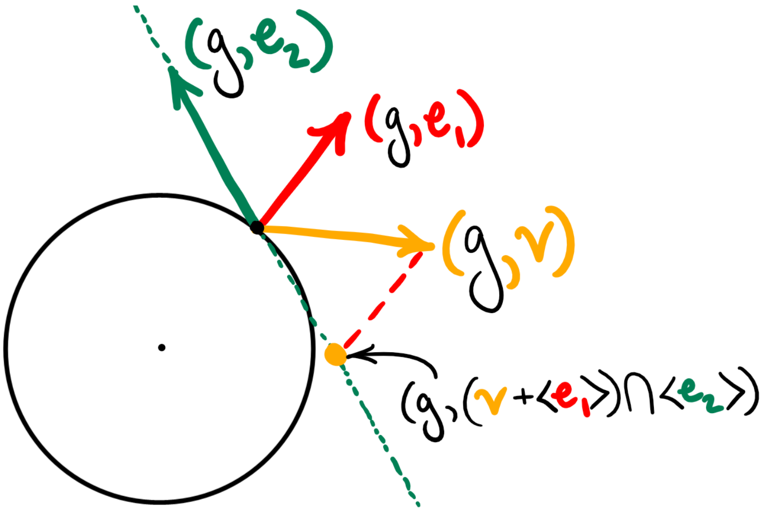

FIGURE 6. For a vector $(g, v) \in \mathrm{SL}_2\mathbb{R} \times_B V$, we can sort of think of the projection of v down to $V/\rho(\mathfrak{b}_+)V$ as the "tangential part" of v

As a B-representation, $\mathrm{H}_0(\mathfrak{b}_+; V(n))$ is just a 1-dimensional vector space on which each $\left[\begin{smallmatrix} a & b \\ 0 & a^{-1} \end{smallmatrix}\right] \in B$ acts by multiplication by a^{-n}. Since $\mathrm{H}_0(\mathfrak{b}_+; V(n))$ is a B-representation, we can form the corresponding homogeneous vector bundle $\mathrm{SL}_2\mathbb{R} \times_B \mathrm{H}_0(\mathfrak{b}_+; V(n))$, which is just a line bundle over $\mathbb{RP}^1$.

It turns out[2] that, if we complexify everything, so that we're instead considering the complex line bundle $\mathrm{SL}_2\mathbb{C} \times_{B_\mathbb{C}} (\mathrm{H}_0(\mathfrak{b}_+; V(n)) \otimes \mathbb{C})$ over $\mathbb{CP}^1$, then the space of holomorphic sections of this line bundle, equipped with the action of $\mathrm{SL}_2\mathbb{C}$ induced by the natural left-action of $\mathrm{SL}_2\mathbb{C}$ on itself, is isomorphic to $V(n) \otimes \mathbb{C}$ as an $\mathrm{SL}_2\mathbb{C}$-representation! Intuitively, this amounts to showing that each holomorphic section of $\mathrm{SL}_2\mathbb{C} \times_{B_\mathbb{C}} (\mathrm{H}_0(\mathfrak{b}_+; V(n)) \otimes \mathbb{C})$ *lifts* to an overlying parallel section of $\mathrm{SL}_2\mathbb{C} \times_{B_\mathbb{C}} (V(n) \otimes \mathbb{C})$.

Of course, this construction for the complex case relies heavily on the rigidity of holomorphic functions. If we had tried the same thing

[2]This is basically just a baby version of the Bott-Borel-Weil theorem from Kostant's perspective. See, for example, Theorems 3.3.5 and 3.3.8 of [1].

with smooth or real-analytic sections of $\mathrm{SL}_2\,\mathbb{R}\times_B \mathrm{H}_0(\mathfrak{b}_+; V(n))$, then we would not get a finite-dimensional representation. On the other hand, there's clearly going to be a *copy* of $V(n)$ in the space of such sections, since we can just take the parallel sections and project down from $V(n)$ to $\mathrm{H}_0(\mathfrak{b}_+; V(n))$. Couldn't we just... pick out the representation somehow?

In fact, we can do this, though the ideas involved will probably seem quite mysterious until later. In addition to $\mathrm{H}_0(\mathfrak{b}_+; V(n))$, let us consider another 1-dimensional B-representation $\mathrm{H}_1(\mathfrak{b}_+; V(n))$, which for now we will simply define to be a copy of $\mathbb{R}$ on which each $\left[\begin{smallmatrix} a & b \\ 0 & a^{-1}\end{smallmatrix}\right] \in B$ acts by multiplication by a^{n+2}. Again, we can consider the homogeneous vector bundles $\mathrm{SL}_2\,\mathbb{R}\times_B \mathrm{H}_0(\mathfrak{b}_+; V(n))$ and $\mathrm{SL}_2\,\mathbb{R}\times_B \mathrm{H}_1(\mathfrak{b}_+; V(n))$.

We define a differential operator $\mathcal{D}^{V(n)}$ for each $V(n)$ as follows: for

$$f : \mathrm{SL}_2\,\mathbb{R} \to \mathrm{H}_0(\mathfrak{b}_+; V(n)) \approx \mathbb{R},$$

the new B-equivariant function

$$\mathcal{D}^{V(n)}(f) : \mathrm{SL}_2\,\mathbb{R} \to \mathrm{H}_1(\mathfrak{b}_+; V(n)) \approx \mathbb{R}$$

is given by[3]

$$\mathcal{D}^{V(2k)}(f) = \frac{1}{(2k)!}\left(\prod_{i=1}^{k}(\partial_\theta^2 + (2i)^2)\right)(\partial_\theta f)$$

for $n = 2k$ and

$$\mathcal{D}^{V(2k+1)}(f) = \frac{-1}{(2k+1)!}\left(\prod_{i=0}^{k}(\partial_\theta^2 + (2i+1)^2)\right)(f)$$

for $n = 2k+1$. This $\mathcal{D}^{V(n)}$ is called the *first Bernstein-Gelfand-Gelfand (BGG) operator* for $V(n)$, and in this homogeneous case, $\mathcal{D}^{V(n)}(f) = 0$ precisely when f lifts to a parallel section of $\mathrm{SL}_2\,\mathbb{R}\times_B V(n)$.

Going back to our earlier computation of the parallel sections of $\mathrm{SL}_2\,\mathbb{R}\times_B V(1)$, we have $\mathcal{D}^{V(1)}(f) = -(\partial_\theta^2 + 1)(f) = -\partial_\theta^2(f) - f$, which vanishes precisely when $f(\theta) = r_1\cos(\theta) + r_2\sin(\theta)$. In this case, the section corresponding to f lifts to a parallel section given by

$$\hat{f} : \theta \mapsto \begin{bmatrix} -(\partial_\theta f)(\theta) \\ f(\theta)\end{bmatrix},$$

as we noted earlier.

Unfortunately, we are not yet even remotely ready to explain where these BGG operators came from; we will probably get to this toward the end of the course. For now, just keep these mysterious differential operators in the back of your mind.

---

[3]While we obtained these formulae by recursion via the usual method (using the equivariant version of the Kostant Laplacian), it is worth noting that similar formulae can be found in [2].

References

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

[2] Eastwood, M. G., Gover, A. R.: *The BGG complex on projective space*. SIGMA, vol. 7 (2011)

[3] Fulton, W., Harris, J.: *Representation Theory: A First Course*. Graduate Texts in Mathematics, Springer-Verlag New York (1991)

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