PARABOLIC GEOMETRIES FOR PEOPLE THAT LIKE PICTURES

LECTURE 8: THE ANATOMY OF A PARABOLIC MODEL GEOMETRY

JACOB W. ERICKSON

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.

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.

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.

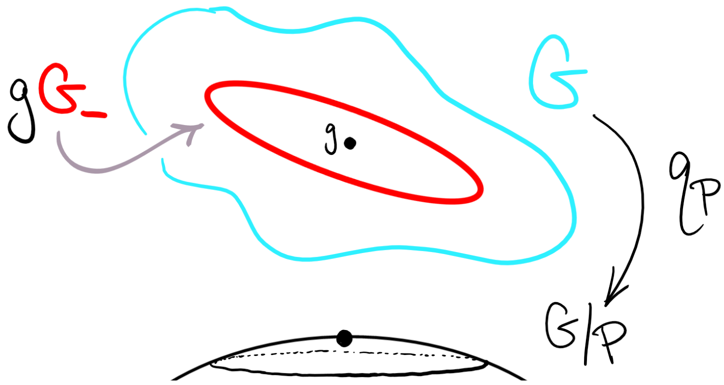

FIGURE 1. The natural quotient map $q_P : G \to G/P$ restricts to an embedding on each gG_-

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 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.

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 $\mathrm{SL}_2\,\mathbb{R}/B$.

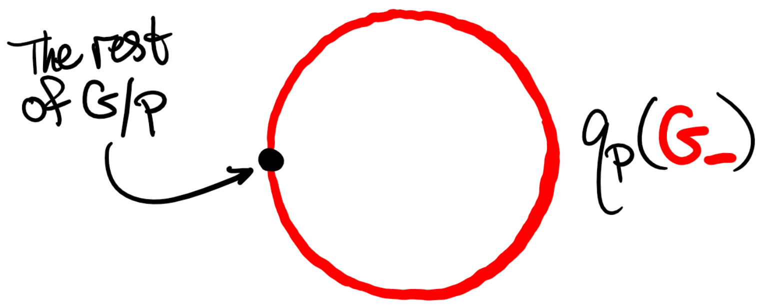

Figure 2. The open subset $q_P(G_-)$ often fills up a large portion of G/P

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 $B \leq P$. 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 $\mathfrak{b} \leq \mathfrak{p}$, we must have $\mathfrak{p}^\perp = \mathfrak{p}_+ \leq \mathfrak{b}_+ = \mathfrak{b}^\perp$, and similarly, $\theta(\mathfrak{p}_+) = \mathfrak{g}_- \leq \mathfrak{b}_- = \theta(\mathfrak{b}_+)$. 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 $\mathfrak{b}_+ \leq \mathfrak{b} \leq \mathfrak{p}$, we must also have that $\theta(\mathfrak{b}_+) = \mathfrak{b}_- \leq \mathfrak{g}_- + \mathfrak{g}_0 = \theta(\mathfrak{p})$. Thus, $G_- \leq B_- \leq G_-G_0$, and in particular, $q_{\scriptscriptstyle P}(G_-) = q_{\scriptscriptstyle P}(B_-)$.

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

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

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.

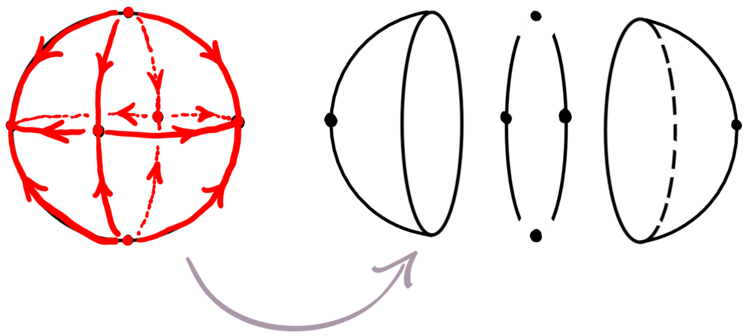

FIGURE 3. The cell decomposition for a parabolic geometry, corresponding to a decomposition into stable manifolds for the action of the grading element of a minimal parabolic subgroup

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.

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.

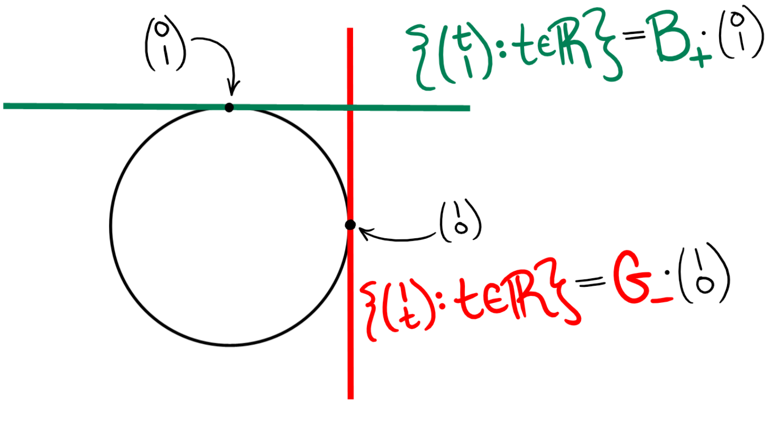

Figure 4. In $(\mathrm{SL}_2\,\mathbb{R}, B)$, G_- acts by translations on the affine line through $\binom{1}{0}$, while B_+ acts by translations on the affine line through $\binom{0}{1}$

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.

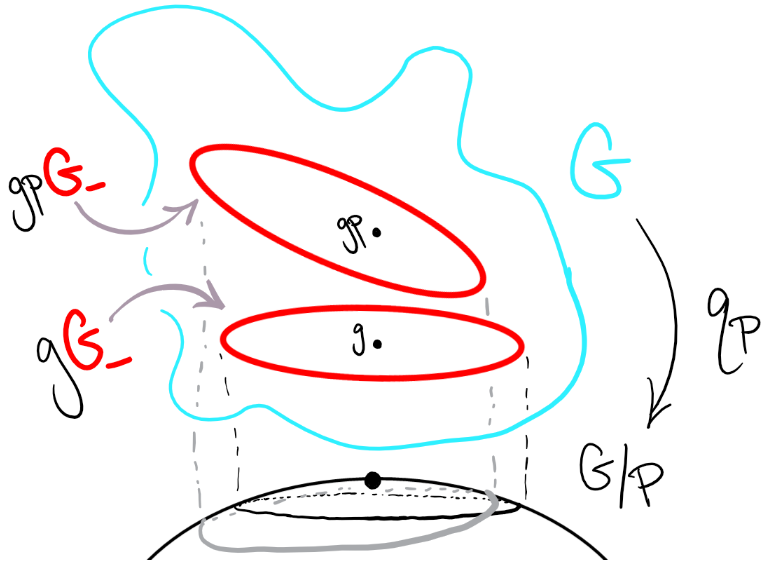

FIGURE 5. The right-action of P_+ is by "unipotent tilts" of the copies of G_-

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

[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.

above, we have a copy gG_- of G_- through each $g \in G$, which allows us to give a local 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.

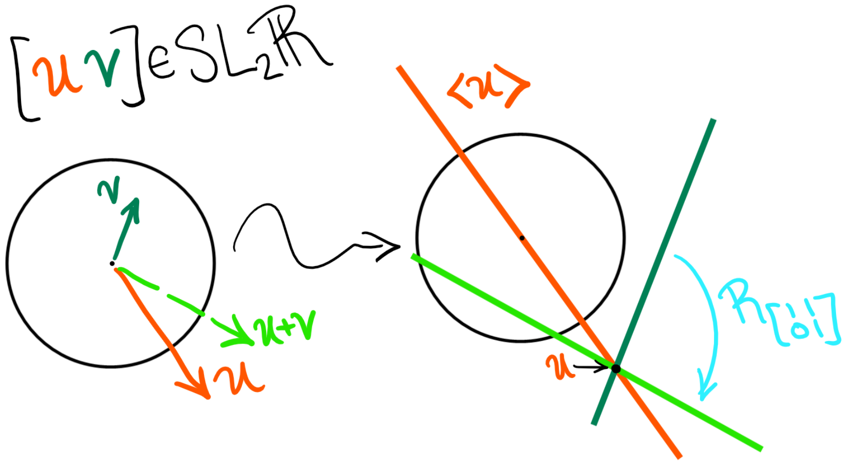

FIGURE 6. Right-translating $[u\ v] \in \mathrm{SL}_2\,\mathbb{R}$ by a unipotent tilt takes the affine line determined by v and tilts it along the line determined by u