The Symmetric Space for $SL_n(\boldsymbol{R})$

Rich Schwartz

November 27, 2013

The purpose of these notes is to discuss the symmetric space X on which $SL_n(\boldsymbol{R})$ acts. Here, as usual, $SL_n(\boldsymbol{R})$ denotes the group of $n \times n$ matrices of determinant 1.

1 The Group Action

The space X denotes the sets of $n \times n$ real matrices M such that

- M is symmetric: $M^t = M$.
- $\det(M) = 1$.
- M is positive definite: $M(v) \cdot v > 0$ for all $v \neq 0$.

The action of $SL_n(\boldsymbol{R})$ on X is given by

$$g(M) = gMg^t. \tag{1}$$

One can easily verify that $g(M) \in X$. Only the third step has anything to it:

$$gMg^t(v) \cdot v = M(g^t(v)) \cdot g^t(v) = M(w) \cdot w > 0.$$

The action we have described really is a group action, because

$$g(h(M)) = g(hMh^t)g^t = (gh)M(gh)^t = gh(M).$$

2 The Tangent Space

The identity matrix I is the natural origin for X. The tangent space $T_I(X)$ consists of derivatives of the form

$$V = \frac{dM}{ds}(0), \tag{2}$$

where M_s is a path of matrices in X and $M_0 = I$. Clearly V is a symmetric matrix. Also, to first order, we have $M_s = I + sV$, and

$$\det(I + sV) = 1 + s \ \mathrm{trace}(V) + ... \tag{3}$$

Hence $\mathrm{trace}(V) = 0$. Conversely, if $\mathrm{trace}(V) = 0$ and V is symmetric, then $I + sV$ belongs to X to first order. More precisely

$$M_s = \exp(sV) = I + sV + \frac{s^2}{2!}V^2 + \frac{s^3}{3!}V^3 + \cdots \tag{4}$$

gives a path in X, for small s, having V as derivative. In short, $T_I(X)$ is the vector space of $n \times n$ symmetric matrices of trace 0.

The differential action of $SL_n(\boldsymbol{R})$ on $T_I(X)$ works very conveniently. If $g \in SL_n(\boldsymbol{R})$ we have

$$dg(V) = \frac{d}{ds}\Big(g(M_s)\Big) = \frac{d}{ds}(gM_sg^t) = gVg^t.$$

In short

$$dg(V) = gVg^t. \tag{5}$$

In other words, the action of g on the tangent space $T_I(X)$ is really just the same as the action of g on X.

Lemma 2.1 *The stabilizer of I is $SO(n)$.*

Proof: The stabilizer of I consists of those matrices $g \in SL_n(\boldsymbol{R})$ such that $g(I) = I$. This happens if and only of $gg^t = I$. These are precisely the orthogonal matrices. That is $g \in SO(n)$. To see this, note that

$$g(v) \cdot g(w) = g^t g(v) \cdot w.$$

So, g preserves the dot product if and only if $g^t g = I$, which is the same as $gg^t = I$. ♠

3 Lemmas about Symmetric Matrices

Lemma 3.1 *Every $M \in X$ has the form $M = gDg^{-1}$ where $g \in SO(n)$ and D is diagonal.*

Proof: Introduce the quadratic form $Q(v,v) = M(v) \cdot v$. The lemma is equivalent to the claim that there is an orthonormal basis which diagonalizes Q. To see this, choose a unit vector w_1 which maximizes $Q(*,*)$. Next, consider the restriction of M to the subspace $W = (w_1)^{\perp}$. If $w_2 \cdot w_1 = 0$ then $M(w_2) \cdot w_1 = 0$ because

$$0 = \frac{d}{ds} Q(w_1 + sw_2, w_1 + sw_2) = 2M(w_1) \cdot w_2. \tag{6}$$

But this means that W is an invariant subspace for M, and $M|_W$ is the matrix defining the quadratic form $Q|_W$. By induction, $Q|_W$ is diagonalizable, with basis $\{w_2, ..., w_n\}$. Then $\{w_1, ..., w_n\}$ is the desired basis on $\boldsymbol{R}^n$. ♠

Corollary 3.2 *Every $V \in T_I(X)$ has the form $T = gEg^{-1}$ where $g \in SO(n)$ and E is diagonal.*

Proof: By the previous result, and by definition, we can find a smooth path $q_s \in SO(n)$ and a smooth path D_s of diagonal matrices so that

$$V = \frac{d}{ds}(g_s D_s g_s^t).$$

By the product rule,

$$\frac{d}{ds}(g_s D_s g_s^t)\bigg|_0 = gD'g^t + \Big(g'g^t + g(g')^t\Big).$$

Here we have set $g = g_0$ and $D' = D'(0)$, etc. The last two terms cancel, because

$$0 = \frac{dI}{dt} = \frac{d}{ds}(g_s g_s^t) = gD'g^t + g'g^t + g(g')^t.$$

Hence

$$V = gEg^t, \qquad g = g_0, \qquad E = D'.$$

This completes the proof ♠

4 The Invariant Riemannian Metric

Given $A, B \in T_I(X)$, we define

$$\langle A, B\rangle = \mathrm{trace}(AB). \tag{7}$$

Lemma 4.1 *The metric is invariant under the action of $SO(n)$.*

Proof: Given $g \in SO(n)$, we have $g^t = g^{-1}$. Hence

$$\langle g(A), g(B)\rangle = \langle gAg^{-1}, gBg^{-1}\rangle = \mathrm{trace}(gAg^{-1}gBg^{-1}) =$$

$$\mathrm{trace}(gABg^{-1}) = \mathrm{trace}(AB) = \langle A, B\rangle.$$

This completes the proof. ♠

Lemma 4.2 *The metric is positive definite.*

Proof: We can write $V = gEg^{-1}$, where E is diagonal and g is orthogonal. For this reason

$$\langle V, V\rangle = \mathrm{trace}(E^2) > 0.$$

The point is that E^2 has non-negative entries, not all of which are 0. ♠

5 The Diagonal Slice

There is one subspace of X on which it is easy to understand the Riemannian metric. Let $\Delta \subset X$ denote the subset consisting of the diagonal matrices with positive entries. We can identify Δ with the subspace

$$\boldsymbol{R}_0^n = \{(x_1, ..., x_n) | \sum x_i = 0\}. \tag{8}$$

The identification carries a matrix to the sequence of logs of its diagonal entries. This makes sense because all the diagonal entries of elements of Δ are positive.

In the log coordinates, the diagonal subgroup of $SL_n(\boldsymbol{R})$ acts by translations. Moreover, our metric agrees with the standard dot product on $\boldsymbol{R}_0^n$

at the origin, by symmetry. Hence, our log coordinates give an isometry between Δ and $\boldsymbol{R}_0^n$ equipped with its usual Euclidean metric.

It remains to show that Δ is geodesically embedded. Before we do this, we mention a cautionary case: In hyperbolic space, if we restrict the hyperbolic metric to a horosphere, we get the Euclidean metric. However, horospheres are not geodesically embedded, and so the zero curvature is a consequence of the distorted embedding. We want to rule out this kind of thing for Δ.

Lemma 5.1 *Δ is geodesically embedded in X.*

Proof: We'll show that the shortest paths connecting two points in Δ belong to Δ. Choose some $M_1 \in \Delta$ and let M_s be a path connecting M_1 to $I = M_0$. We can choose continuous paths g_s and D_s such that

- $g_0 = g_1 = I$.
- g_s and D_s vary smoothly. Here g_s is orthogonal and D_s is diagonal.
- $M_s = g_s D_s g_s^t$ for all $s \in [0, 1]$.

The path D_s also connects M_0 to M_1. We just have to show that D_s is not a longer path than M_s.

We compute

$$\frac{d}{ds} M_s = A_s + B_s,$$

where

$$A_s = g_s D_s' g_s^t, \qquad B_s = g_s' D_s g_s^t + g_s D_s (g_s')^t. \tag{9}$$

Here the primed terms are the derivatives. Note that

$$\langle A_s, A_s \rangle_{M_s} = \langle D_s', D_s' \rangle_{D_s} \tag{10}$$

because the action of g_s is an isometry of our metric. To finish the proof, we just have to check that A_s and B_s are orthogonal. To so this, we check that $dM_s^{-1}(A_s)$ and $dM_s^{-1}(B_s)$ are orthogonal. These are two tangent vectors in $T_I(X)$, and we know how to compute their inner product.

Dropping the subscript s and using $g^t = g^{-1}$, we compute

$$dM^{-1}(A) = (gD^{-1}g^{-1})(gD'g^{-1})(gD^{-1}g^{-1}) =$$

$$dG^{-1}D'D^{-1}g^{-1} = g\Omega g^{-1}, \tag{11}$$

where Ω is some diagonal matrix. Similarly

$$dM^{-1}(B) = g(\Psi_1 + \Psi_2)g^{-1},$$

$$\Psi_1 = D^{-1}g^{-1}g', \qquad \Psi_2 = (g^{-1})'gD^{-1} = -g^{-1}g'D^{-1}. \tag{12}$$

The last equality comes from Equation 6.

We compute

$$\langle A, B\rangle_M = \mathrm{trace}\Big(g(\Omega\Psi_1 + \Omega\Psi_2)g^{-1}\Big) = \mathrm{trace}(\Omega\Psi_1 + \Omega\Psi_2) =$$

$$\mathrm{trace}(\Omega D^{-1}g^{-1}g') - \mathrm{trace}(\Omega g^{-1}g'D^{-1}) =_1$$
$$\mathrm{trace}(\Omega D^{-1}g^{-1}g') - \mathrm{trace}(D^{-1}\Omega g^{-1}g') =_2$$
$$\mathrm{trace}(\Omega D^{-1}g^{-1}g') - \mathrm{trace}(\Omega D^{-1}g^{-1}g') = 0. \tag{13}$$

Equality 1 is comes from the fact that XY and YX in have the same trace for any matrices X and Y. Equality 2 comes from the fact that D^{-1} and Ω, both diagonal matrices, commute. ♠

6 Maximal Flats

Now we know that $\Delta \subset X$ is a totally geodesic slice isometric to $\boldsymbol{R}^{n-1}$. By symmetry, any subset $g(\Delta) \subset X$, for $g \in SL_n(\boldsymbol{R})$, has the same properties. We call these objects the *maximal flats*. As the name suggests, these are the totally geodesic Euclidean slices of maximal dimension. (We will not prove this last assertion, but will stick with the traditional terminology just the same.)

Lemma 6.1 *Any two points in X are contained in a maximal flat.*

Proof: By symmetry, it suffices to consider the case when one of the points is I and the other point is some M. We have $M = gDg^t$ where g is orthogonal and D is diagonal. But then $g(\Delta)$ contains both points. ♠

It is not true that any two points lie in a unique maximal flat. For generic choices of points, this is true. However, sometimes a pair of points can lie in infinitely many maximal flats. This happens, for instance, when the two points are I and D, and D has some repeated eigenvalues.

The maximal flats are naturally in bijection with the subgroups conjugate in $SL_n(\boldsymbol{R})$ to the diagonal subgroup. There is a nice way to picture this correspondence geometrically. Let $\boldsymbol{P}$ denote the real projective space of dimension $n-1$. Say that a *simplex* is a collection of n general position points in $\boldsymbol{P}$. The diagonal subgroup preserves the simplex whose vertices are $[e_1], ..., [e_n]$, the projectivizations of the vectors in the standard basis. In general, a conjugage group preserves some other simplex; the vertices of the simplex are fixed by all elements of the subgroup. Thus, there is a bijection between the maximal flats and the simplices in $\boldsymbol{P}$.

7 The Weyl Group and Weyl Chambers

The Weyl group is usually defined in terms of the Lie Algebra, but here is a rough and ready geometric description. The Weyl group associated to X is the subgroup $W \subset O(n)$ which acts isometrically on the diagonal slice Δ. When n is odd, we can take $W \subset SO(n)$. In general, W is generated by the permutation matrices.

W contains a number of reflections, for instance, the permutation matrix which swaps the first two coordinates. Each such reflection $g \in W$ fixes some hyperplane $H_g \subset \Delta$. The complement $\Delta - \bigcup H_g$ is a finite union of convex cones. Each such cone is called a *Weyl chamber*.

Here is the first nontrivial example. When $n = 3$ there are 3 reflections. In log coordinates, the corresponding lines in $\boldsymbol{R}^3_0$ are the intersections with the coordinate planes. Geometrically, these lines branch out along the 6th roots of unity, and the Weyl chambers fit together like 6 slices of pizza. In general, the points in the interior of the Weyl chambers correspond to matrices having no repeated eigenvalues.

Using the action of the diagonal group on Δ, we define, for $p \in \Delta$, the union C_p of convex cones to be the translation of the Weyl chambers to p. Thus, we think of the Weyl chambers as something akin to the way we think of a light cone in Minkowski space: The cone is something that really lives at every point, in the tangent space at that point.

We say that a line L in a maximal flat F is *regular* if, for some $p \in L$, the line L points into the interior of some chamber of C_p. This definition is independent of the choice of p. If L is not regular, we call L *singular*. For

$n = 3$, the singular lines in Δ are parallel to the 6th roots of unity, when we use log coordinates.

Lemma 7.1 *If $L \subset F$ is regular, then L lies in a unique maximal flat.*

Proof: We will argue by contradiction. Let $F_1 = F$ and let F_2 be some second maximal flat containing L. Let S_1 and S_2 be the corresponding simplices in $\boldsymbol{P}$. We can assume by symmetry that $F_1 = \Delta$ and L contains the origin. Then the matrices of $L - I$ have no repeating eigenvalues. Hence, the fixed points of these matrices determine S_1 and S_2. But then $S_1 = S_2$. Hence $F_1 = F_2$. ♠

Lemma 7.2 *If $L \subset F$ is singular, L is contained in infinitely many flats.*

Proof: Again, we normalize so that L is contained in Δ and goes through I. In this case, the matrices of $L - I$ have some repeated entries, and the same entries repeat for all the matrices. Moreover, all the elements of $L - I$ have common eigenspaces. Let $\Sigma \subset \boldsymbol{P}$ denote the union of the projectivizations of the eigenspaces. There are infinitely many simplices S which are compatible with Σ in the sense that each point of S is contained in some projectivized eigenspace. Any such simplex corresponds to a flat containing L. ♠

Remark: The last proof is a little bit opaque, so consider the example when $n = 3$ and L consists of matrices having the first two entries equal. Then Σ is a union of the origin O_0 and the line Λ at infinity in the projective plane. If we choose any two distinct points $O_1, O_2 \subset \Lambda$, then the triangle (O_0, O_1, O_2) is compatible with Σ.

So, the maximal flats form a network of Euclidean slices in X. The maximal flats intersect along the boundaries of the Weyl chambers. This picture suggests a kind of higher dimensional graph, called a *building*, and indeed when one considers $SL_n(\boldsymbol{Q}_p)$, the linear group over the p-adics, the corresponding space X_p is indeed known as a building.

8 Hyperbolic Slices

It is worth mentioning that X contains some totally geodesic copies of $\boldsymbol{H}^2$, the hyperbolic plane. When $n = 2$, the corresponding space X_2 is isometric to the hyperbolic plane. We can fit X_2 inside X_n by considering diagonal matrices. Precisely, we can make the upper 2×2 block an arbitrary element of X_2, and then we can put (1)s for the other diagonal entries. This embeds X_2 isometrically inside X_n.

Lemma 8.1 *X_2 is a totally geodesic subspace of X_n.*

Proof: Let $G \subset SL_n(\boldsymbol{R})$ be the subgroup of block diagonal matrices having (1)'s for the first two entries and then some $(n-2) \times (n-2)$ diagonal matrix. The action of G fixes X_2 pointwise. Moreover, for any point $p \in X_n - X_2$, there is some $g \in G$ such that $g(p) \neq p$.

Suppose that $x, y \in X_2$ are two points, connected by a geodesic segment γ. Choosing x and y close enough together, we can arrange that γ is the unique distance minimizing geodesic connecting x and y. But if $\gamma \not\in X_2$ then we can find some $g \in G$ such that $g(\gamma) \neq \gamma$. But then γ and $g(\gamma)$ are two distinct geodesic segments, having the same length, connecting x to y. This is a contradiction. Hence $\gamma \subset X_2$. Since geodesics locally stay in X_2, the subspace X_2 is geodesically embedded. ♠

Once we have one totally geodesic copy of $\boldsymbol{H}^2$ in X_n, we get others by using the action of $SL_n(\boldsymbol{R})$. Thus $g(X_2) \subset X_n$ is a totally geodesic embedded copy of $\boldsymbol{H}^2$ as well. We call these the *hyperbolic slices*.

It turns out that X has non-positive sectional curvature, and the Euclidean slices are the slices of maximum curvature (zero) and the hyperbolic slices are the slices of minimum curvature.