Chapter 8

The Log-Euclidean Framework Applied to SPD Matrices and Polyaffine Transformations

8.1 Introduction

In this Chapter, we use what we have learned in previous chapters to describe an approach due to Arsigny, Fillard, Pennec and Ayache to define a Lie group structure and a class of metrics on symmetric, positive-definite matrices (SPD matrices) which yield a new notion of mean on SPD matrices generalizing the standard notion of geometric mean.

Recall that the set of $n \times n$ SPD matrices, $\mathbf{SPD}(n)$, is not a vector space (because if $A \in \mathbf{SPD}(n)$, then $\lambda A \notin \mathbf{SPD}(n)$ if $\lambda < 0$) but it is a convex cone.

Thus, the *arithmetic mean* of n SPD matrices, $S_1, \ldots, S_n$, can be defined as $(S_1 + \cdots + S_n)/n$, which is SPD.

However, there are many situations, especially in DTI, where this mean is not adequate.

There are essentially two problems:

(1) The arithmetic mean is not invariant under inversion, which means that if $S = (S_1 + \cdots + S_n)/n$, then in general, $S^{-1} \neq (S_1^{-1} + \cdots + S_n^{-1})/n$.

(2) The *swelling effect*: the determinant, $\det(S)$, of the mean, S, may be strictly larger than the original determinants, $\det(S_i)$.

This effect is undesirable in DTI because it amounts to introducing more diffusion, which is physically unacceptable.

To circumvent these difficulties, various metrics on SPD matrices have been proposed. One class of metrics is the *affine-invariant metrics* (see Arsigny, Pennec and Ayache [**?**]).

The swelling effect disappears and the new mean is invariant under inversion but computing this new mean has a high computational cost and, in general, there is no closed-form formula for this new kind of mean.

Arsigny, Fillard, Pennec and Ayache [**?**] have defined a new family of metrics on $\mathbf{SPD}(n)$ named *Log-Euclidean metrics* and have also defined a novel structure of Lie group on $\mathbf{SPD}(n)$ which yields a notion of mean that has the same advantages as the affine mean but is a lot cheaper to compute.

Furthermore, this new mean, called *Log-Euclidean mean*, is given by a simple closed-form formula. We will refer to this approach as the *Log-Euclidean Framework*.

The key point behind the Log-Euclidean Framework is the fact that the exponential map, $\exp\colon \mathbf{S}(n) \to \mathbf{SPD}(n)$, is a bijection, where $\mathbf{S}(n)$ is the space of $n \times n$ symmetric matrices

Consequently, the exponential map has a well-defined inverse, the *logarithm*, $\log\colon \mathbf{SPD}(n) \to \mathbf{S}(n)$.

But more is true. It turns out that $\exp\colon \mathbf{S}(n) \to \mathbf{SPD}(n)$ is a diffeomorphism.

Since exp is a bijection, the above result follows from the fact that exp is a local diffeomorphism on $\mathbf{S}(n)$, because $d\exp_S$ is non-singular for all $S \in \mathbf{S}(n)$.

In Arsigny, Fillard, Pennec and Ayache [**?**], it is proved that the non-singularity of $d\exp_I$ near 0, which is well-known, "propagates" to the whole of $\mathbf{S}(n)$.

Actually, the non-singularity of $d\exp$ on $\mathbf{S}(n)$ is a consequence of a more general result of some interest whose proof can be found in in Mmeimné and Testard [**?**], Chapter 3, Theorem 3.8.4 (see also Bourbaki [**?**], Chapter III, Section 6.9, Proposition 17, and also Theorem 6).

Let $\mathcal{S}(n)$ denote the set of all real matrices whose eigenvalues, $\lambda + i\mu$, lie in the horizontal strip determined by the condition $-\pi < \mu < \pi$. Then, we have the following theorem:

Theorem 8.1.1 *The restriction of the exponential map to $\mathcal{S}(n)$ is a diffeomorphism of $\mathcal{S}(n)$ onto its image, $\exp(\mathcal{S}(n))$. Furthermore, $\exp(\mathcal{S}(n))$ consists of all invertible matrices that have no real negative eigenvalues; it is an open subset of $\mathbf{GL}(n, \mathbb{R})$; it contains the open ball, $B(I, 1) = \{A \in \mathbf{GL}(n, \mathbb{R}) \mid \|A - I\| < 1\}$, for every norm $\| \; \|$ on $n \times n$ matrices satisfying the condition $\|AB\| \leq \|A\| \, \|B\|$.*

Part of the proof consists in showing that exp is a local diffeomorphism and for this, to prove that $d\exp_X$ is invertible for every $X \in \mathcal{S}(n)$.

This requires finding an explicit formula for the derivative of the exponential, which can be done.

With this preparation we are ready to present the natural Lie group structure on $\mathbf{SPD}(n)$ introduced by Arsigny, Fillard, Pennec and Ayache (see also Arsigny's thesis).

8.2 A Lie-Group Structure on SPD(n)

Definition 8.2.1 For any two matrices, $S_1, S_2 \in \mathbf{SPD}(n)$, define the *logarithmic product*, $S_1 \odot S_2$, by

$$S_1 \odot S_2 = \exp(\log(S_1) + \log(S_2)).$$

Obviously, the multiplication operation, $\odot$, is commutative. The following proposition is shown in Arsigny, Fillard, Pennec and Ayache [**?**] (Proposition 3.2):

Proposition 8.2.2 *The set,* $\mathbf{SPD}(n)$*, with the binary operation,* $\odot$*, is an abelian group with identity,* I*, and with inverse operation the usual inverse of matrices. Whenever* S_1 *and* S_2 *commute, then* $S_1 \odot S_2 = S_1S_2$ *(the usual multiplication of matrices).*

Actually, $(\mathbf{SPD}(n), \odot, I)$ is an abelian Lie group isomorphic to the vector space (also an abelian Lie group!) $\mathbf{S}(n)$, as shown in Arsigny, Fillard, Pennec and Ayache [**?**] (Theorem 3.3 and Proposition 3.4):

Theorem 8.2.3 *The abelian group, $(\mathbf{SPD}(n), \odot, I)$ is a Lie group isomorphic to its Lie algebra, $\mathfrak{spd}(n) = \mathbf{S}(n)$. In particular, the Lie group exponential in $\mathbf{SPD}(n)$ is identical to the usual exponential on $\mathbf{S}(n)$.*

We now investigate bi-invariant metrics on the Lie group, $\mathbf{SPD}(n)$.

8.3 Log-Euclidean Metrics on $\mathbf{SPD}(n)$

If G is a lie group, recall that we have the operations of left multiplication, L_a, and right multiplication, R_a, given by

$$L_a(b) = ab, \qquad R_a(b) = ba,$$

for all $a, b \in G$. A Riemannian metric, $\langle -, - \rangle$, on G is *left-invariant* iff dL_a is an isometry for all $a \in G$, that is,

$$\langle u, v \rangle_b = \langle (dL_a)_b(u), (dL_a)_b(v) \rangle_{ab},$$

for all $b \in G$ and all $u, v \in T_bG$.

Similarly, $\langle -, - \rangle$ is *right-invariant* iff dR_a is an isometry for all $a \in G$ and $\langle -, - \rangle$ is *bi-invariant* iff it is both left and right invariant.

In general, a Lie group does not admit a bi-invariant metric but an abelian Lie group always does because $\mathrm{Ad}_g = \mathrm{id} \in \mathbf{GL}(\mathfrak{g})$ for all $g \in G$ and so, the adjoint representation, $\mathrm{Ad}\colon G \to \mathbf{GL}(\mathfrak{g})$, is trivial (that is, $\mathrm{Ad}(G) = \{\mathrm{id}\}$) and then, the existence of bi-invariant metrics is a consequence of Proposition **??**, which we repeat here for the convenience of the reader:

Proposition 8.3.1 *There is a bijective correspondence between bi-invariant metrics on a Lie group, G, and* Ad*-invariant inner products on the Lie algebra, $\mathfrak{g}$, of G, that is, inner products, $\langle -, - \rangle$, on $\mathfrak{g}$ such that* Ad_a *is an isometry of $\mathfrak{g}$ for all $a \in G$; more explicitly, inner products such that*

$$\langle \mathrm{Ad}_a u, \mathrm{Ad}_a v \rangle = \langle u, v \rangle,$$

for all $a \in G$ and all $u, v \in \mathfrak{g}$.

Then, given any inner product, $\langle -, - \rangle$ on G, the induced bi-invariant metric on G is given by

$$\langle u, v \rangle_g = \langle (dL_{g^{-1}})_g u, (dL_{g^{-1}})_g v \rangle.$$

Now, the geodesics on a Lie group equipped with a bi-invariant metric are the left (or right) translates of the geodesics through e and the geodesics through e are given by the group exponential, as stated in Proposition **??** (3) which we repeat for the convenience of the reader:

Proposition 8.3.2 *For any Lie group, G, equipped with a bi-invariant metric, we have:*

(1) The inversion map, $\iota\colon g \mapsto g^{-1}$, is an isometry.

(2) For every $a \in G$, if I_a denotes the map given by

$$I_a(b) = ab^{-1}a, \qquad \text{for all } a, b \in G,$$

then I_a is an isometry fixing a which reverses geodesics, that is, for every geodesic, γ, through a

$$I_a(\gamma)(t) = \gamma(-t).$$

(3) The geodesics through e are the integral curves, $t \mapsto \exp(tu)$, where $u \in \mathfrak{g}$, that is, the one-parameter groups. Consequently, the Lie group exponential map, $\exp\colon \mathfrak{g} \to G$, coincides with the Riemannian exponential map (at e) from T_eG to G, where G is viewed as a Riemannian manifold.

If we apply Proposition 8.3.2 to the abelian Lie group, $\mathbf{SPD}(n)$, we find that the geodesics through S are of the form

$$\gamma(t) = S \odot e^{tV},$$

where $V \in \mathbf{S}(n)$.

But $S = e^{\log S}$, so

$$S \odot e^{tV} = e^{\log S} \odot e^{tV} = e^{\log S + tV},$$

so every geodesic through S is of the form

$$\gamma(t) = e^{\log S + tV} = \exp(\log S + tV).$$

To avoid confusion between the exponential and the logarithm as Lie group maps and as Riemannian manifold maps, we will denote the former by exp and log and their Riemannian counterparts by Exp and Log.

Here is Corollary 3.9 of Arsigny, Fillard, Pennec and Ayache [**?**]:

Theorem 8.3.3 *For any inner product, $\langle -, - \rangle$, on $\mathbf{S}(n)$, if we give the Lie group, $\mathbf{SPD}(n)$, the bi-invariant metric induced by $\langle -, - \rangle$, then the following properties hold:*

(1) For any $S \in \mathbf{SPD}(n)$, the geodesics through S are of the form

$$\gamma(t) = e^{\log S + tV}, \qquad V \in \mathbf{S}(n).$$

(2) The exponential and logarithm associated with the bi-invariant metric on $\mathbf{SPD}(n)$ are given by

$$\begin{aligned} \mathrm{Exp}_S(U) &= e^{\log S + d\log_S(U)} \\ \mathrm{Log}_S(T) &= d\exp_{\log S}(\log T - \log S), \end{aligned}$$

for all $S, T \in \mathbf{SPD}(n)$ and all $U \in \mathbf{S}(n)$.

(3) The bi-invariant metric on $\mathbf{SPD}(n)$ is given by

$$\langle U, V \rangle_S = \langle d\log_S(U), d\log_S(V) \rangle,$$

for all $U, V \in \mathbf{S}(n)$ and all $S \in \mathbf{SPD}(n)$ and the distance, $d(S, T)$, between any two matrices, $S, T \in \mathbf{SPD}(n)$, is given by

$$d(S, T) = \|\log T - \log S\|,$$

where $\| \, \|$ is the norm corresponding to the inner product on $\mathfrak{spd}(n) = \mathbf{S}(n)$.

(4) The map, $\exp\colon \mathbf{S}(n) \to \mathbf{SPD}(n)$, is an isometry.

In view of Theorem 8.3.3, part (3), bi-invariant metrics on the Lie group $\mathbf{SPD}(n)$ are called *Log-Euclidean metrics*.

Since $\exp\colon \mathbf{S}(n) \to \mathbf{SPD}(n)$ is an isometry and $\mathbf{S}(n)$ is a vector space, the Riemannian Lie group, $\mathbf{SPD}(n)$, is a complete, simply-connected and flat manifold (the sectional curvature is zero at every point) that is, a flat *Hadamard manifold* (see Sakai [**?**], Chapter V, Section 4).

Although, in general, Log-Euclidean metrics are not invariant under the action of arbitary invertible matrices, they are invariant under similarity transformations (an isometry composed with a scaling).

Recall that $\mathbf{GL}(n)$ acts on $\mathbf{SPD}(n)$, *via*,

$$A \cdot S = ASA^\top,$$

for all $A \in \mathbf{GL}(n)$ and all $S \in \mathbf{SPD}(n)$.

We say that a Log-Euclidean metric is *invariant under $A \in \mathbf{GL}(n)$* iff

$$d(A \cdot S, A \cdot T) = d(S, T),$$

for all $S, T \in \mathbf{SPD}(n)$.

The following result is proved in Arsigny, Fillard, Pennec and Ayache [**?**] (Proposition 3.11):

Proposition 8.3.4 *There exist metrics on $\mathbf{S}(n)$ that are invariant under all similarity transformations, for example, the metric $\langle S, T\rangle = \mathrm{tr}(ST)$.*

8.4 A Vector Space Structure on $\mathbf{SPD}(n)$

The vector space structure on $\mathbf{S}(n)$ can also be transfered onto $\mathbf{SPD}(n)$.

Definition 8.4.1 *For any matrix, $S \in \mathbf{SPD}(n)$, for any scalar, $\lambda \in \mathbb{R}$, define the scalar multiplication, $\lambda \circledast S$, by*

$$\lambda \circledast S = \exp(\lambda \log(S)).$$

It is easy to check that $(\mathbf{SPD}(n), \odot, \circledast)$ is a vector space with addition $\odot$ and scalar multiplication, $\circledast$.

By construction, the map, $\exp\colon \mathbf{S}(n) \to \mathbf{SPD}(n)$, is a linear isomorphism. What happens is that the vector space structure on $\mathbf{S}(n)$ is transfered onto $\mathbf{SPD}(n)$ *via* the log and exp maps.

8.5 Log-Euclidean Means

One of the major advantages of Log-Euclidean metrics is that they yield a computationally inexpensive notion of mean with many desirable properties.

If $(x_1, \ldots, x_n)$ is a list of n data points in $\mathbb{R}^m$, then it is an easy exercise to see that the mean, $\overline{x} = (x_1 + \cdots + x_n)/n$, is the unique minimum of the map

$$x \mapsto \sum_{i=1}^{n} d(x, x_i)_2^2,$$

where d_2 is the Euclidean distance on $\mathbb{R}^m$.

We can think of the quantity,

$$\sum_{i=1}^{n} d(x, x_i)_2^2,$$

as the *dispersion* of the data.

More generally, if (X, d) is a metric space, for any $\alpha > 0$ and any positive weights, $w_1, \ldots, w_n$, with $\sum_{i=1}^n w_i = 1$, we can consider the problem of minimizing the function,

$$x \mapsto \sum_{i=1}^{n} w_i d(x, x_i)^\alpha.$$

The case $\alpha = 2$ corresponds to a generalization of the notion of mean in a vector space and was investigated by Fréchet.

In this case, any minimizer of the above function is known as a *Fréchet mean*. Fréchet means are not unique but if X is a complete Riemannian manifold, certain sufficient conditions on the dispersion of the data are known that ensure the existence and uniqueness of the Fréchet mean (see Pennec [**?**]).

The case $\alpha = 1$ corresponds to a generalization of the notion of *median*.

When the weights are all equal, the points that minimize the map,

$$x \mapsto \sum_{i=1}^{n} d(x, x_i),$$

are called *Steiner points*. On a Hadamard manifold, Steiner points can be characterized (see Sakai [**?**], Chapter V, Section 4, Proposition 4.9).

In the case where $X = \mathbf{SPD}(n)$ and d is a Log-Euclidean metric, it turns out that the Fréchet mean is unique and is given by a simple closed-form formula.

We have the following theorem from Arsigny, Fillard, Pennec and Ayache [**?**] (Theorem 3.13):

Theorem 8.5.1 *Given N matrices, $S_1, \ldots, S_N \in \mathbf{SPD}(n)$, their Log-Euclidean Fréchet mean exists and is uniquely determined by the formula*

$$\mathbb{E}_{\mathrm{LE}}(S_1, \ldots, S_N) = \exp\left(\frac{1}{N}\sum_{i=1}^{N} \log(S_i)\right).$$

Furthermore, the Log-Euclidean mean is similarity-invariant, invariant by group multiplication and inversion and exponential-invariant.

Similarity-invariance means that for any similarity, A,

$$\mathbb{E}_{\mathrm{LE}}(AS_1A^\top, \ldots, AS_NA^\top) = A\mathbb{E}_{\mathrm{LE}}(S_1, \ldots, S_N)A^\top$$

and similarly for the other types of invariance.

Observe that the Log-Euclidean mean is a generalization of the notion of geometric mean.

Indeed, if $x_1, \ldots, x_n$ are n positive numbers, then their *geometric mean* is given by

$$\mathbb{E}_{\mathrm{geom}}(x_1, \ldots, x_n) = (x_1 \cdots x_n)^{\frac{1}{n}} = \exp\left(\frac{1}{n}\sum_{i=1}^{n} \log(x_i)\right).$$

The Log-Euclidean mean also has a good behavior with respect to determinants. The following theorem is proved in Arsigny, Fillard, Pennec and Ayache [**?**] (Theorem 4.2):

Theorem 8.5.2 *Given N matrices, $S_1, \ldots, S_N \in \mathbf{SPD}(n)$, we have*

$$\det(\mathbb{E}_{\mathrm{LE}}(S_1, \ldots, S_N)) = \mathbb{E}_{\mathrm{geom}}(\det(S_1), \ldots, \det(S_N)).$$

8.6 Log-Euclidean Polyaffine Transformations

The registration of medical images is an important and difficult problem.

The work described in Arsigny, Commowick, Pennec and Ayache [**?**] (and Arsigny's thesis [**?**]) makes an orginal and valuable contribution to this problem by describing a method for parametrizing a class of non-rigid deformations with a small number of degrees of freedom.

After a global affine alignment, this sort of parametrization allows a finer local registration with very smooth transformations.

This type of parametrization is particularly well adpated to the registration of histological slices, see Arsigny, Pennec and Ayache [**?**].

The goal is to fuse some affine or rigid transformations in such a way that the resulting transformation is invertible and smooth.

The direct approach which consists in blending N global affine or rigid transformations, $T_1, \ldots, T_N$ using weights, $w_1, \ldots, w_N$, does not work because the resulting transformation,

$$T = \sum_{i=1}^{N} w_i T_i,$$

is not necessarily invertible. The purpose of the weights is to define the domain of influence in space of each T_i.

The key idea is to associate to each rigid (or affine) transformation, T, of $\mathbb{R}^n$, a vector field, V, and to view T as the diffeomorphism, Φ_1^V, corresponding to the time $t = 1$, where Φ_t^V is the global flow associated with V. In other words, T is the result of integrating an ODE

$$X' = V(X, t),$$

starting with some initial condition, X_0, and $T = X(1)$.

Now, it would be highly desirable if the vector field, V, did not depend on the time parameter, and this is indeed possible for a large class of affine transformations, which is one of the nice contributions of the work of Arsigny, Commowick, Pennec and Ayache [**?**].

Recall that an affine transformation, $X \mapsto LX + v$, (where L is an $n \times n$ matrix and $X, v \in \mathbb{R}^n$) can be conveniently represented as a linear transformation from $\mathbb{R}^{n+1}$ to itself if we write

$$\begin{pmatrix} X \\ 1 \end{pmatrix} \mapsto \begin{pmatrix} L & v \\ 0 & 1 \end{pmatrix} \begin{pmatrix} X \\ 1 \end{pmatrix}.$$

Then, the ODE with constant coefficients

$$X' = LX + v,$$

can be written

$$\begin{pmatrix} X' \\ 0 \end{pmatrix} = \begin{pmatrix} L & v \\ 0 & 0 \end{pmatrix} \begin{pmatrix} X \\ 1 \end{pmatrix}$$

and, for every initial condition, $X = X_0$, its unique solution is given by

$$\begin{pmatrix} X(t) \\ 1 \end{pmatrix} = \exp\left(t \begin{pmatrix} L & v \\ 0 & 0 \end{pmatrix}\right) \begin{pmatrix} X_0 \\ 1 \end{pmatrix}.$$

Therefore, if we can find reasonable conditions on matrices, $T = \begin{pmatrix} M & t \\ 0 & 1 \end{pmatrix}$, to ensure that they have a unique real logarithm,

$$\log(T) = \begin{pmatrix} L & v \\ 0 & 0 \end{pmatrix},$$

then we will be able to associate a vector field, $V(X) = LX + v$, to T, in such a way that T is recovered by integrating the ODE, $X' = LX + v$.

Furthermore, given N transformations, $T_1, \ldots, T_N$, such that $\log(T_1), \ldots, \log(T_N)$ are uniquely defined, we can fuse $T_1, \ldots, T_N$ at the *infinitesimal level* by defining the ODE obtained by blending the vector fields, $V_1, \ldots, V_N$, associated with $T_1, \ldots, T_N$ (with $V_i(X) = L_iX + v_i$), namely

$$V(X) = \sum_{i=1}^{N} w_i(X)(L_iX + v_i).$$

Then, it is easy to see that the ODE,

$$X' = V(X),$$

has a unique solution for every $X = X_0$ defined for all t, and the fused transformation is just $T = X(1)$.

Thus, the fused vector field,

$$V(X) = \sum_{i=1}^{N} w_i(X)(L_i X + v_i),$$

yields a one-parameter group of diffeomorphisms, Φ_t. Each transformation, Φ_t, is smooth and invertible and is called a *Log-Euclidean polyaffine tranformation*, for short, *LEPT*.

Of course, we have the equation

$$\Phi_{s+t} = \Phi_s \circ \Phi_t,$$

for all $s, t \in \mathbb{R}$ so, in particular, the inverse of Φ_t is Φ_{-t}.

We can also interpret Φ_s as $(\Phi_1)^s$, which will yield a fast method for computing Φ_s.

Observe that when the weight are scalars, the one-parameter group is given by

$$\begin{pmatrix} \Phi_t(X) \\ 1 \end{pmatrix} = \exp\left(t \sum_{i=1}^{N} w_i \begin{pmatrix} L_i & v_i \\ 0 & 0 \end{pmatrix} \right) \begin{pmatrix} X \\ 1 \end{pmatrix},$$

which is the Log-Euclidean mean of the affine transformations, T_i's (w.r.t. the weights w_i).

Fortunately, there is a sufficient condition for a real matrix to have a unique real logarithm and this condition is not too restrictive in practice.

Recall that $\mathcal{S}(n)$ denotes the set of all real matrices whose eigenvalues, $\lambda + i\mu$, lie in the horizontal strip determined by the condition $-\pi < \mu < \pi$. We have the following version of Theorem 8.1.1:

Theorem 8.6.1 *The image,* $\exp(\mathcal{S}(n))$*, of* $\mathcal{S}(n)$ *by the exponential map is the set of real invertible matrices with no negative eigenvalues and* $\exp\colon \mathcal{S}(n) \to \exp(\mathcal{S}(n))$ *is a bijection.*

Theorem 8.6.1 is stated in Kenney and Laub [**?**] without proof. Instead, Kenney and Laub cite DePrima and Johnson [**?**] for a proof but this latter paper deals with complex matrices and does not contain a proof of our result either.

The injectivity part of Theorem 8.6.1 can be found in Mmeimné and Testard [**?**], Chapter 3, Theorem 3.8.4.

In fact, $\exp\colon \mathcal{S}(n) \to \exp(\mathcal{S}(n))$ is a *diffeomorphism*.

This result is proved in Bourbaki [**?**], see Chapter III, Section 6.9, Proposition 17 and Theorem 6.

Curious readers should read Gallier [**?**] for the full story.

For any matrix, $A \in \exp(\mathcal{S}(n))$, we refer to the unique matrix, $X \in \mathcal{S}(n)$, such that $e^X = A$, as the *principal logarithm* of A and we denote it as $\log A$.

Observe that if T is an affine transformation given in matrix form by

$$T = \begin{pmatrix} M & t \\ 0 & 1 \end{pmatrix},$$

since the eigenvalues of T are those of M plus the eigenvalue 1, the matrix T has no negative eigenvalues iff M has no negative eigenvalues and thus the principal logarithm of T exists iff the principal logarithm of M exists.

It is proved in Arsigny, Commowick, Pennec and Ayache that LEPT's are affine invariant, see [**?**], Section 2.3. This shows that LEPT's are produced by a truly geometric kind of blending, since the result does not depend at all on the choice of the coordinate system.

In the next section, we describe a fast method for computing due to Arsigny, Commowick, Pennec and Ayache [**?**].

8.7 Fast Polyaffine Transforms

Recall that since LEPT's are members of the one-parameter group, $(\Phi_t)_{t\in\mathbb{R}}$, we have

$$\Phi_{2t} = \Phi_{t+t} = \Phi_t^2$$

and thus,

$$\Phi_1 = (\Phi_{1/2^N})^{2^N}.$$

Observe the formal analogy of the above formula with the formula

$$\exp(M) = \exp\left(\frac{M}{2^N}\right)^{2^N},$$

for computing the exponential of a matrix, M, by the *scaling and squaring method*.

It turns out that the "scaling and squaring method" is one of the most efficient methods for computing the exponential of a matrix, see Kenney and Laub [**?**] and Higham [**?**].

The key idea is that $\exp(M)$ is easy to compute if M is close zero since, in this case, one can use a few terms of the exponential series, or better, a Padé approximant (see Higham [**?**]).

The scaling and squaring method for computing the exponential of a matrix, M, can be sketched as follows:

1. *Scaling Step*: Divide M by a factor, 2^N, so that $\frac{M}{2^N}$ is close enough to zero.
2. *Exponentiation step*: Compute $\exp\left(\frac{M}{2^N}\right)$ with high precision, for example, using a Padé approximant.
3. *Squaring Step*: Square $\exp\left(\frac{M}{2^N}\right)$ repeatedly N times to obtain $\exp\left(\frac{M}{2^N}\right)^{2^N}$, a very accurate approximation of e^M.

There is also a so-called *inverse scaling and squaring method* to compute efficiently the principal logarithm of a real matrix, see Cheng, Higham, Kenney and Laub [**?**].

Arsigny, Commowick, Pennec and Ayache made the very astute observation that the scaling and squaring method can be adpated to compute LEPT's very efficiently [**?**].

This method, called *fast polyaffine transform*, computes the values of a Log-Euclidean polyaffine transformation, $T = \Phi_1$, at the vertices of a regular n-dimensional grid (in practice, for $n = 2$ or $n = 3$). Recall that T is obtained by integrating an ODE, $X' = V(X)$, where the vector field, V, is obtained by blending the vector fields associated with some affine transformations, $T_1, \ldots, T_n$, having a principal logarithm.

Here are the three steps of the **fast polyaffine transform**:

1. *Scaling Step*: Divide the vector field, V, by a factor, 2^N, so that $\frac{V}{2^N}$ is close enough to zero.

2. *Exponentiation step*: Compute $\Phi_{1/2^N}$, using some adequate numerical integration
method.

3. *Squaring Step*: Compose $\Phi_{1/2^N}$ with itself recursively N times to obtain an accurate approximation of $T = \Phi_1$.

Of course, one has to provide practical methods to achieve step 2 and step 3.

Several methods to achieve step 2 and step 3 are proposed in Arsigny, Commowick, Pennec and Ayache [**?**].

One also has to worry about boundary effects, but this problem can be alleviated too, using bounding boxes.

To conclude our survey of the Log-Euclidean polyaffine framework for locally affine registration, we briefly discuss how the Log-Euclidean framework can be generalized to rigid and affine transformations.

8.8 A Log-Euclidean Framework for Transformations in $\exp(\mathcal{S}(n))$

Arsigny, Commowick, Pennec and Ayache observed that if T_1 and T_2 are two affine transformations in $\exp(\mathcal{S}(n))$, then we can define their distance as

$$d(T_1, T_2) = \|\log(T_1) - \log(T_2)\|,$$

where $\| \, \|$ is any norm on $n \times n$ matrices (see [**?**], Appendix A.1).

We can go a little further and make $\mathcal{S}(n)$ and $\exp(\mathcal{S}(n))$ into Riemannian manifolds in such a way that the exponential map, $\exp\colon \mathcal{S}(n) \to \exp(\mathcal{S}(n))$, is an isometry.

Since $\mathcal{S}(n)$ is an open subset of the vector space, $\mathrm{M}(n, \mathbb{R})$, of all $n \times n$ real matrices, $\mathcal{S}(n)$ is a manifold, and since $\exp(\mathcal{S}(n))$ is an open subset of the manifold, $\mathbf{GL}(n, \mathbb{R})$, it is also a manifold.

Obviously, $T_L\mathcal{S}(n) \cong \mathrm{M}(n, \mathbb{R})$ and $T_S \exp(\mathcal{S}(n)) \cong \mathrm{M}(n, \mathbb{R})$, for all $L \in \mathcal{S}(n)$ and all $S \in \exp(\mathcal{S}(n))$ and the maps, $d\exp_L \colon T_L\mathcal{S}(n) \to T_{\exp(L)} \exp(\mathcal{S}(n))$ and $d\log_S \colon T_S \exp(\mathcal{S}(n)) \to T_{\log(S)}\mathcal{S}(n)$, are linear isomorphisms.

We can make $\mathcal{S}(n)$ into a Riemannian manifold by giving it the induced metric induced by any norm, $\| \, \|$, on $\mathrm{M}(n, \mathbb{R})$, and make $\exp(\mathcal{S}(n))$ into a Riemannian manifold by defining the metric, $\langle -, - \rangle_S$, on $T_S \exp(\mathcal{S}(n))$, by

$$\langle A, B \rangle_S = \| d\log_S(A) - d\log_S(B) \| ,$$

for all $S \in \exp(\mathcal{S}(n))$ and all $A, B \in \mathrm{M}(n, \mathbb{R})$.

Then, it is easy to check that $\exp \colon \mathcal{S}(n) \to \exp(\mathcal{S}(n))$ *is indeed an isometry* and, as a consequence, the Riemannian distance between two matrices, $T_1, T_2 \in \exp(\mathcal{S}(n))$, is given by

$$d(T_1, T_2) = \| \log(T_1) - \log(T_2) \| ,$$

again called the *Log-Euclidean distance*.

Since every affine transformation, T, can be represented in matrix form as

$$T = \begin{pmatrix} M & t \\ 0 & 1 \end{pmatrix},$$

and, as we saw in section 8.6, since the principal logarithm of T exists iff the principal logarithm of M exists, we can view the set of affine transformations that have a principal logarithm as a subset of $\exp(\mathcal{S}(n+1))$.

Unfortunately, this time, even though they are both flat, $\mathcal{S}(n)$ and $\exp(\mathcal{S}(n))$ are not complete manifolds and so, the Fréchet mean of N matrices, $T_1, \ldots, T_n \in \exp(\mathcal{S}(n))$, may not exist.

However, recall that from Theorem 8.1.1 that the open ball,

$$B(I,1) = \{A \in \mathbf{GL}(n,\mathbb{R}) \mid \|A - I\|' < 1\},$$

is contained in $\exp(\mathcal{S}(n))$ for any norm, $\|\ \|'$, on matrices (not necessarily equal to the norm defining the Riemannian metric on $\mathcal{S}(n)$) such that $\|AB\|' \leq \|A\|' \, \|B\|'$ so, for any matrices $T_1, \ldots, T_n \in B(I,1)$, the Fréchet mean is well defined and is uniquely determined by

$$\mathbb{E}_{\mathrm{LE}}(T_1, \ldots, T_N) = \exp\left(\frac{1}{N}\sum_{i=1}^{N} \log(T_i)\right),$$

namely, it is their *Log-Euclidean mean*.

From a practical point of view, one only needs to ckeck that the eigenvalues, ξ, of $\frac{1}{N}\sum_{i=1}^{N}\log(T_i)$ are in the horizontal strip, $-\pi < \Im(\xi) < \pi$.

Provided that $\mathbb{E}_{\text{LE}}(T_1, \ldots, T_N)$ is defined, it is easy to show, as in the case of SPD matrices, that $\det(\mathbb{E}_{\text{LE}}(T_1, \ldots, T_N))$ is the geometric mean of the determinants of the T_i's.

The Riemannian distance on $\exp(\mathcal{S}(n))$ is not affine invariant but it is invariant under inversion, under rescaling by a positive scalar, and under rotation for certain norms on $\mathcal{S}(n)$ (see [**?**], Appendix A.2).

However, the Log-Euclidean mean of matrices in $\exp(\mathcal{S}(n))$ is invariant under conjugation by any matrix, $A \in \mathbf{GL}(n, \mathbb{R})$, since $ASA^{-1} \in \exp(\mathcal{S}(n))$ for any $S \in \exp(\mathcal{S}(n))$ and since $\log(ASA^{-1}) = A\log(S)A^{-1}$.

In particular, the Log-Euclidean mean of affine transformations in $\exp(\mathcal{S}(n+1))$ is invariant under arbitrary invertible affine transformations (again, see [**?**], Appendix A.2).

For more details on the Log-Euclidean framework for locally rigid or affine deformation, for example, about regularization, the reader should read Arsigny, Commowick, Pennec and Ayache [**?**].