Comprehensive Riemannian Framework for Open Curves

The underlying theme in this thesis is the use of the full set of physical features – shape, scale, orientation, and position – in open curve analysis. This information can significantly aid the clustering, classifying, labeling and data analysis of open curves. In this chapter, we describe a flexible Riemannian framework for open curves which defines joint feature spaces or manifolds to study combinations of these features in a consistent way.

1 Introduction

White matter fibers reconstructed from DT-MRI images can be described as 3-dimensional open, continuous curves. Sulci can also be reconstructed to give a similar geometric description. A physical description of these structures would involve shape, scale, orientation, and position, the physical features associated with the curves. Of these features, shape is most commonly used in medical image analysis.

Several recent papers have proposed the use of a formal Riemannian framework for shape analysis of continuous curves [1, 2]. This type of framework has many advantages: (i) It provides techniques for comparing, matching, and deforming shapes of curves under the chosen metric. The correspondences for these tasks are established automatically. (ii) It also provides tools for defining and computing statistical summaries of sample shapes for different shape classes [3].

The Riemannian framework for open curves described in this chapter uses the same core ideas that were developed for elastic shape analysis of continuous closed curves. With this framework, we can compare and quantify differences between open curves in a coherent way. These comparisons are based on different feature combinations and each such combination constitutes a manifold. The five feature combinations we consider are:

1. Shape, orientation, scale and position: $\mathcal{S}_1$
2. Shape, orientation, scale: $\mathcal{S}_2$
3. Shape and scale: $\mathcal{S}_3$
4. Shape and orientation: $\mathcal{S}_4$
5. Shape: $\mathcal{S}_5$

These manifolds are united by a set of considerations and mathematical techniques; we describe this common methodology in Section 2. We then give mathematical descriptions for each of these manifolds (Sections 3.1–3.5).

This chapter serves as a mathematical background for the discussions in later chapters, particularly those in Part II of this thesis. The material presented is based directly on Mani et al. [4]. The methodology was developed by Anuj Srivastava and his collaborators; their publications, in particular those on Riemannian analysis of elastic curves, are other reference sources [2, 5, 6, 7, 8]. Concepts from differential geometry and group theory are discussed informally. For a more formal and complete treatment, the reader is referred to do Carmo [9, 10] for differential geometry and Rotman [11] for group theory.

2 General Methodology

The mathematical framework we describe draws from ideas in differential geometry, algebra and functional analysis and is an extension of the mathematical techniques developed for elastic shape analysis of continuous closed curves. Joshi et al. [2, 5, 12] introduced the square-root velocity function (SRVF) for the analysis of closed curves and Balov et al. [6] describe the analysis of open curves using the square-root function (SRF). The same general principles presented in these papers, i.e. the use of an elastic metric for shape analysis and a path-straightening approach to construct geodesics, apply here. The elastic metric, which is a Riemannian metric originally proposed by Younes [13], allows a curve to stretch and bend as it deforms along a geodesic. The path-straightening method [14] uses an arbitrary path to initialize the geodesic between two curves; the geodesic is iteratively computed using variational methods.

The mathematical tools we use enable us to define a representation space for open curves, ensure that the curves are invariant to certain shape-preserving transformations and compute geodesic distances between two curves. Each space defined uses a unique combination of features and, by extension, invariances but they share common procedural steps. These are outlined below:

1. Each fiber or sulcal curve is represented as an open continuous parameterized curve, β, as seen in Figure 1.

2. In order to compare curves we use elastic curve matching and for this the parameterized curve β is represented by a function. With the choice of function representation we can make the curve invariant to translation. We can also, at this stage, make the curve invariant to scale by manually scaling it.

3. The representation space described above (2) is a preshape space, $\mathcal{C}$, because many curves that are rotated and reparameterized versions of each other can actually be different elements. To unify these different representations of the same curve, one defines an equivalence class of functions. The set of all these equivalence classes is called the shape space $\mathcal{S}$.
(*Note that the terms preshape space and shape space are used in a general sense for the representation spaces in manifolds $\mathcal{S}_1 - \mathcal{S}_5$. The methodology we use is an extension of work done in shape analysis where these terms were first applied and for consistency we use the same nomenclature for all the feature spaces.*)

4. A Riemannian structure is imposed on $\mathcal{S}$ and curves are compared by computing geodesics between their representations in $\mathcal{S}$. The geodesic length is a quantification of the difference between two curves in the joint feature space under consideration. The geodesics are computed using numerical algorithms.

We detail these steps describing the representations, the invariances and the computation of distances in the sections below. Figure 3 illustrates these steps and Figure 4 the relationship between the manifolds.

2.1 Representation Space

The curve is represented in an $\mathbb{L}^2$ space of square integrable functions first in a preshape space and then in a shape space which is a quotient space of the preshape space. There are thus two levels of representation. In this section, we will discuss the representation space under four headings: the curve representation, the function representation, the preshape space and the shape space.

2.1.1 Curve Representation

Let $\beta : [0,1] \to \mathbb{R}^3$ be an open continuous parameterized curve such that its speed $\|\dot{\beta}(t)\|$ is non-zero everywhere. The parameterization determines the rate at which the curve is traversed. The norm $\|\beta\| = \sqrt{\int_0^1 \|\beta(t)\|^2 dt}$ is a square-integrable function so we can refer to the space of all such curves as $\mathbb{L}^2([0,1], \mathbb{R}^3)$ or more simply $\mathbb{L}^2$.

An advantage to using continuous curves instead of landmarks which are reference points along the curve is that we avoid the problems associated with the selection of these points which many of the existing approaches to shape analysis have to contend with. As a trade-off, tools from functional analysis are needed.

2.1.2 Function Representation

The curve β is represented either by the *square-root velocity function* (SRVF) or the *square-root function* (SRF). These two function representations are important for the following reasons: First, the use of the $\mathbb{L}^2$ metric on the space of SRVFs and SRFs generates an elastic metric and a framework for elastic shape comparisons; Second, these functions preserve the $\mathbb{L}^2$ metric under reparameterization (see Section 2.3); Lastly, a versatile framework can be built around these representations–the shape space can be easily modified and different combinations of features included.

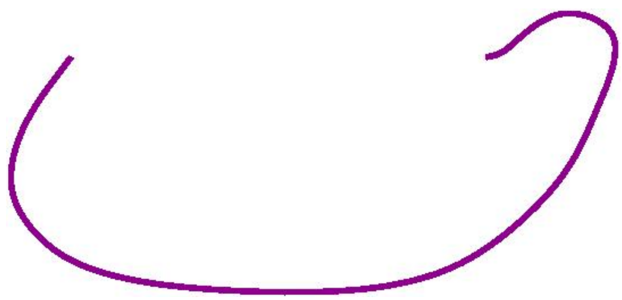

Figure 1: $\beta : [0,1] \to \mathbb{R}^3$ is an open continuous parameterized curve. A sulcal curve or white matter fiber can be described as an open curve β.

SRVF The square-root velocity function of the curve β is

$$q(t) = \frac{\dot{\beta}(t)}{\sqrt{\|\dot{\beta}(t)\|}} \ , \quad q : [0,1] \to \mathbb{R}^3 \ . \tag{1}$$

From Eqn. 1 we see that q is only dependant on the velocity term $\dot{\beta}$. It is therefore called the ***square root velocity function*** and the norm $||q||$ is the square root of the instantaneous speed along the curve β. It is possible to recover the original curve β, within a translation, using $\beta(t) = \int_0^t ||q(s)||q(s)ds$. The function q is invariant to translation in $\mathbb{R}^3$. By manually scaling the curves to the same length we can also remove the scale information at this stage.

SRF The square-root function is defined thus:

$$h(t) = \sqrt{\|\dot{\beta}(t)\|}\beta(t) \ , \quad h : [0,1] \to \mathbb{R}^3 \ . \tag{2}$$

Since h is dependant on β, the full set of attributes of the open curve, including global translation, are incorporated in the expression. The $\sqrt{\|\dot{\beta}(t)\|}$ term is a scalar that ensures invariance to reparameterization.

It is not easy to recover the curve β from the SRF because this involves solving a higher-order ordinary differential equation. Consequently, we cannot draw geodesic paths between two curves or compute sample statistics, though we can still compute distances between the curves. Because of this limitation, we only apply the SRF to cases where we want to include position as a feature; for all other cases, we use the SRVF.

2.1.3 Preshape Space

The space of all square-root or square-root velocity representations of curves $\mathcal{C}$ is an infinite dimensional vector space of all functions in $\mathbb{L}^2([0,1],\mathbb{R}^3)$. It has the form:

$$\mathcal{C} = \{f : [0,1] \rightarrow \mathbb{L}^2([0,1],\mathbb{R}^3)\},$$

where f is the function representation of the curve.

An inner product can be defined on $T_f(\mathcal{C})$, the tangent space of $\mathcal{C}$ at the point f:

$$\langle w_1, w_2\rangle = \int_0^1 \langle w_1(t), w_2(t)\rangle \, dt, \qquad w_1, w_2 \in T_f(\mathcal{C}).$$

This metric, defined infinitesimally using elements of the tangent space at a point, is the ***Riemannian metric***. The use of the $\mathbb{L}^2$ metric on the space of square-root velocity representations generates an elastic metric; the Riemannian metric is therefore an elastic metric in this space.

The differentiable manifold $\mathcal{C}$ with this Riemannian metric is a ***Riemannian manifold***. With this Riemannian structure we can derive the following: (i) geodesic paths between curves; (ii) the exponential map; (iii) the inverse exponential map.

The elements of $\mathcal{C}$ do not represent the shape of a curve uniquely. A reparametrization of β, using an element $\gamma \in \Gamma$, where Γ is the group of ***diffeomorphisms*** (a smooth bijective map with smooth inverses) from $[0,1]$ to itself, results in a different square-root velocity function while preserving its shape. Similarly, any rigid rotation of β changes q but not its shape. Since $\mathcal{C}$ has many elements of the same shape, we call it the ***preshape*** space of open curves. (As noted before, we use the term preshape space in a general sense; it applies to the preshape space of $\mathcal{S}_1 - \mathcal{S}_4$, as well as to $\mathcal{S}_5$.)

In this framework, the three different preshape spaces we use are:

1. $\mathcal{C}_1 = \{h : \mathbb{R}^3 \rightarrow \mathbb{R}^3\}$. This is the preshape space for the $\mathcal{S}_1$ manifold. The h-function (SRF) is used here.

2. $\mathcal{C}_2 = \{q : \mathbb{R}^3 \rightarrow \mathbb{R}^3\}$. This is the preshape space for the $\mathcal{S}_2$ and $\mathcal{S}_3$ manifolds. The q-function (SRVF) is used since the space is invariant to translation.

3. $\mathcal{C}_3 = \{q : [0,1] \rightarrow \mathbb{R}^3 | \int_0^1 \|q(t)\|^2 dt = 1\}$. This is the space of all unit-length, open, elastic curves and is the preshape space for $\mathcal{S}_4$ and $\mathcal{S}_5$. The q-function is used and the curves are scaled to remove variability due to scale. The space is thus invariant to translation and uniform scaling.

2.1.4 Shape Space

To unify all the different representations of the same curve, we define an ***equivalence class*** or ***orbit*** of functions. In cases where invariance to reparameterization and orientation is sought, the orbit has the form

$$[f] = \{(\gamma, Of) | \gamma \in \Gamma, O \in SO(3)\}.$$

Here f is the function representation, i.e., the h-function or the q-function; Γ is the reparameterization group; $SO(3)$ is the 3-dimensional rotation group and (γ, Of) is the reparameterized and rotated curve. Equivalence relations are a by-product of group operations. Group actions with which we achieve invariance to reparameterization or orientation are more fully described in Section 2.2.

The elements of the orbit are considered equivalent, but the orbits themselves are distinct and do not intersect; collectively they define a disjoint set which is the ***shape*** space: $\mathcal{S} = \mathcal{C}/(\Gamma \times SO(3))$. This is now a ***quotient*** space of the preshape space $\mathcal{C}$ and has unique elements.
(*As with preshape space, the term shape space is generic to all five feature spaces,* $\mathcal{S}_1 - \mathcal{S}_5$*. These manifolds, discussed in Sections 3.1–3.5, are differentiated by the preshape spaces and equivalence classes that define them; only* $\mathcal{S}_5$ *is an actual shape space*).

There is a second type of equivalence class we are interested in. For manifolds $\mathcal{S}_1$, $\mathcal{S}_2$ and $\mathcal{S}_4$, invariance to reparameterization but not to orientation is desired. The orbit in this case has the form

$$[f] = \{(\gamma, f) | \gamma \in \Gamma\}$$

and the corresponding quotient space is $\mathcal{S} = \mathcal{C}/\Gamma$.

Since the shape space $\mathcal{S}$ is a quotient space of the preshape space $\mathcal{C}$, it inherits the Riemannian metric from $\mathcal{C}$. With this structure, we can compute a geodesic distance in $\mathcal{S}$ between two orbits $[f_1]$ and $[f_2]$. More details on geodesic distance computations are provided in Section 2.3.

2.2 Invariances, Equivalence Relationships

Since the shape of an object does not change when it is translated, scaled or rotated, an important requirement in shape analysis is that metrics be invariant to these transformations. Parameterized curves require an additional invariance. A reparameterization only changes the speed with which a curve is traversed, not its shape. Reparameterization is thus another shape preserving transformation. Metrics should be invariant to it as well.

There are two ways in which invariances to transformations are achieved. The first comes directly from the function representation. With the choice of function we can remove translation or scale information. In the $\mathcal{S}_2$ and $\mathcal{S}_3$ manifolds, for instance, we remove translation and in the $\mathcal{S}_4$ and $\mathcal{S}_5$ manifolds we remove both translation and scale. The second way to achieve invariance to transformations is by establishing equivalence classes. We elaborate on this in Sections 2.2.3 and 2.2.4.

We apply the four transformations, *translation, nonrigid uniform scaling, rigid rotation* and *reparameterization* to the shape manifold, $\mathcal{S}_5$. These transformations are discussed below. The invariance requirements for the four other manifolds, $\mathcal{S}_1 - \mathcal{S}_4$, are different but they are dealt with in a similar manner.

2.2.1 Translation

The *shape+orientation+scale+position* manifold, $\mathcal{S}_1$, includes position information and so we want the function representation we use to reflect this. For this manifold, the value of the function is designed to change with the translation or the change in position of the curve.
The four other manifolds, $\mathcal{S}_2 - \mathcal{S}_5$, do not include position as a feature and are invariant to translation. The function representation, q, used in those cases, is expressed exclusively in terms of the time derivative of the curve, β, and does not have a position component.

2.2.2 Scale

In order to remove the influence of the scales of curves in the quantitative analysis, we can rescale them to be of the same length. The rescaling, which is done with the preshape representation, leaves the function representation, q, and the Riemannian metric unchanged. The preshape space, however, is reduced. It is computationally convenient to scale the curve, β, to unit length; $\int_0^1 ||\dot{\beta}(t)||dt = \int_0^1 ||f(t)||^2 dt = 1$ then holds. Since they have unit norm, the set of all functions associated with curves of length one are elements of a hypersphere in $\mathbb{L}^2$. The differential geometry of a sphere is well-known and because of this, analysis and computations in subsequent steps is greatly simplified.

Two of the manifolds we consider, *shape*, $\mathcal{S}_5$, and *shape+orientation*, $\mathcal{S}_4$, are invariant to scale and receive this treatment. The shape attribute is by definition invariant to scale. In the case of the *shape+orientation* space, orientation, not scale, is added to the shape representation. Thus, this manifold retains the scale invariance of the shape space.

2.2.3 Reparameterization

The curve β is parameterized. This parameterization introduces an additional source of variability since arbitrary parameterizations are included in the representation. For any two curves, β_1 and β_2, different parameterizations, in general, result in different distances between them. We account for this variability by applying ideas from group theory. The shape of the curve does not change due to reparameterization so we would like to treat this mapping as we do rigid rotations and other shape-preserving transformations. We define a reparameterization group (i.e. the set of all reparameterized curves) and the action of this group on the preshape space results in reparameterization orbits. The elements of these orbits are equivalent. This enables us to compare orbits of curves, a comparison that is now independent of parameterization.

Γ, as noted before, is the set of all orientation-preserving diffeomorphisms (i.e., the direction at different points along the curve do not change due to the diffeomorphism). For the SRVF, i.e. the q-function, the map, $\mathcal{C} \times \Gamma \rightarrow \mathcal{C}$, is a group action defined by:

$$(q, \gamma) \rightarrow \sqrt{\dot{\gamma}}(qo\gamma).$$

The new function (q, γ) is the reparameterized curve and its distance from the original function q, is, in general, non-zero. We define two elements q_1 and q_2 as equivalent (i.e. $q_1 \sim q_2$) if for some $\gamma \in \Gamma$, $q_2 = (q_1, \gamma)$. A set of equivalent elements constitute an equivalence class:

$$[q] = \{(q, \gamma) | \gamma \in \Gamma\}.$$

We see that the equivalence class, which is obtained as a result of the group action Γ on $\mathcal{C}$, enables us to achieve invariance to reparameterization. The equivalence classes partition $\mathcal{C}$ into disjoint sets. The quotient space that results is $\mathcal{S} = \mathcal{C}/\Gamma$.

For the SRF, i.e. the h-function, the group action $\mathcal{C} \times \Gamma \to \mathcal{C}$ is defined by:

$$(h, \gamma) \to \sqrt{\dot{\gamma}}(h o \gamma).$$

Since the parameterized curve, β, is the starting point for all the manifolds, $\mathcal{S}_1 - \mathcal{S}_5$, we need to ensure invariance to reparameterization for each one of them.

2.2.4 Orientation

Let $O \in SO(3)$ be a rotation matrix and let $SO(3)$, the set of all possible rotations in $\mathbb{R}^3$, act isometrically (i.e., it is a distance-preserving map) on $\mathcal{C}$ as follows: $SO(3) \times \mathcal{C} \to \mathcal{C}$. $(O, q) = \{Oq(t) | t \in [0, 1]\}$ is the rotated curve that results from this group action. In cases where we want the analysis to be invariant to this group action we define the orbit of q under $SO(3)$ as

$$[q]_o = \{(Oq) | O \in SO(3)\} \subset \mathcal{C}.$$

The elements of $[q]_o$ are equivalent; they are rotated versions of each other.

The action of $SO(3)$ is combined with the action of Γ since, as noted in Section 2.2.4, the quotient spaces $\mathcal{S}_1 - \mathcal{S}_5$, need to also be invariant to reparameterization. The joint action of Γ and $SO(3)$ gives the larger orbit

$$[q] = \{(\gamma, Oq) | \gamma \in \Gamma, O \in SO(3)\} \subset \mathcal{C}. \tag{3}$$

The corresponding quotient space is $\mathcal{S} = \mathcal{C}/(\Gamma \times SO(3))$. We will consider how the two group actions, Γ and $SO(3)$, interact next.

2.2.5 The Product Group $(\Gamma \times SO(3))$

There are two important properties associated with the product group $(\Gamma \times SO(3))$:

1. The actions of $SO(3)$ and Γ on $\mathcal{C}$ commute. It is due to this that we can form an equivalence class (Eqn. 3) and define the action of the product group. We also make use of this property when we find distances and need to optimize iteratively (see Section 2.3.1).

2. The joint action of $(\Gamma \times SO(3))$ on $\mathcal{C}$ is by isometries with respect to the distance metric. This implies that the inner product on $T_{[q]}$, the tangent space to $\mathcal{S}$, is independent of the choice of $\tilde{q} \in [q]$. The geodesic distance between two points in $\mathcal{S}$ is given by:

$$d_{\mathcal{S}}([q]_0, [q]_1) = \min_{\tilde{q}_1 \in [q_1]} d_{\mathcal{C}}(q_0, \tilde{q}_1). \tag{4}$$

2.3 Geodesics and Distances

We can compare curves by quantifying their differences using a distance function. The standard $\mathbb{L}^2$ metric, often used in quantitative analysis of fibers, is given by

$$\|\beta_1 - \beta_2\|_2 = \sqrt{\int_0^1 \|\beta_1(t) - \beta_2(t)\|^2 dt}.$$

This metric is not, in general, invariant to reparameterizations, i.e.

$$\|\beta_1 \circ \gamma - \beta_2 \circ \gamma\| \neq \|\beta_1 - \beta_2\|.$$

Since we require invariance to reparameterization, we solve this issue by introducing the SRF and the SRVF representations. These functions preserve the $\mathbb{L}^2$ distance under reparameterization so that for any two curves β_1, β_2, with the corresponding functions f_1 and f_2, and any $\gamma \in \Gamma$, we have that $\|(f_1, \gamma) - (f_2, \gamma)\|_2 = \|f_1 - f_2\|_2$. Because of this equality, we are able to define the distance between the two curves as

$$d(\beta_1, \beta_2) = \gamma^* = \operatorname*{argmin}_{\gamma \in \Gamma} \left(\|f_1 - (f_2, \gamma)\|_2 \right) \ . \tag{5}$$

The minimization is performed by taking the path-straightening approach and using the standard dynamic programming algorithm for the computations. This is described in Section 2.3.1.

For the $\mathcal{S}_3$ and $\mathcal{S}_5$ manifolds, invariance to orientation is sought. Since we need to take the variability due to reparameterization also into account, the distance function now becomes

$$d(\beta_1, \beta_2) = \gamma^* = \operatorname*{argmin}_{\gamma \in \Gamma, O \in SO(3)} \|f_1 - O(f_2, \gamma)\|_2 \ . \tag{6}$$

We solve this with a joint optimization described in Section 2.3.1.

In our open curve analysis, we are interested in the distance function in $\mathcal{S}$, the shape space. This distance function is inherited from $\mathcal{C}$. The distance between two curves β_1 and β_2 is the distance between their orbits $[\beta_1]$ and $[\beta_2]$ and is the pairwise shortest distance between elements in these two orbits. Once the optimal reparametrization and/or orientation of f_2 are obtained, we can compute the geodesic path between the orbits $[f_1]$ and $[f_2]$. Figure 2 shows the geodesic path between two curves in the $\mathcal{S}_2 - \mathcal{S}_5$ manifolds.

2.3.1 Optimization

The optimal distance is obtained by minimizing the cost function. For optimization over reparameterization, Γ, we use dynamic programming or gradient descent. When we need to optimize over both orientation and reparameterization, we perform a *joint optimization*; the optimal Γ is obtained by dynamic programming and the optimal orientation over $SO(3)$ is computed using Procrustes alignment. We briefly describe these procedures in this section.

Dynamic Programming Dynamic programming (DP) is a numerical optimization algorithm where we obtain an optimal path by solving the problem sequentially. At each stage we select an optimal trajectory from all the possible trajectories by optimizing (minimizing) the cost function. For this analysis, the cost function is the distance function (Eqns. 5 and 6) that matches the point $f_2(\gamma(t))$ with the point $f_1(t)$. Since the cost function is defined by the $\mathbb{L}^2$ norm, it is additive over the path $(t, \gamma(t))$ and can be cast as a DP problem. The algorithm forms a graph from $(0, 0)$ to $(1, 1)$ in $\mathbb{R}^2$ and searches over all the paths on that grid, such that the slope of the graph is strictly between 0^o and 90^o. This constraint is placed so that $0 < \dot{\gamma}^* < \infty$. The cumulative cost over the entire grid gives us an approximation to γ^*. The DP algorithm gives an exact solution.

A gradient-based optimization is an alternative to DP. The gradient-descent gives a local solution. Its estimate for γ^* is less accurate than DP but, as a trade-off, it is computationally less expensive.

Joint Optimization Consider the optimization problem:

$$(\gamma^*, O^*) = \underset{\gamma \in \Gamma, O \in SO(3)}{\operatorname{argmin}} \|f_1 - \sqrt{\dot{\gamma}} O f_2(\gamma)\|^2 \ . \tag{7}$$

We encounter this problem with the $\mathcal{S}_3$ and $\mathcal{S}_5$ manifolds where the curves are invariant to both orientation and reparameterization and it requires a joint optimization solution. The individual solutions for optimal orientation and registration are given below:

1. **Optimal Rotation**: For a fixed $\gamma \in \Gamma$, the optimization problem over $SO(3)$ in Eqn. 7 is solved by Procrustes alignment. We find $O^* = UV^T$ where O^* is the optimal rotation for aligning two curves and USV^T is the singular valued decomposition of $A = \int_0^1 f_1(t)(\sqrt{\dot{\gamma}} f_2(\gamma(t)))^T dt$. (In cases where the determinant of A is negative, one needs to modify V by multiplying the last column by -1. We have $O^* = U\tilde{V}^T$ in this case. This is a known result from rigid alignment of objects when the points across objects are already registered.)

2. **Optimal Registration**: For a fixed O, the optimization problem in Eqn. 7 over Γ can be solved using the DP algorithm described above. The cost function is defined by the $\mathbb{L}^2$ norm and, thus, is additive over the path $(t, \gamma(t))$. The algorithm forms a finite-dimensional grid in $[0, 1]^2$ and searches over all the paths on that grid, satisfying the required constraints, to obtain an approximation to γ^*.

Since we have algorithms for optimizing over the two components individually, we can go back and forth between the two steps till convergence is reached.

$\mathcal{S}_2$: shape+orientation+scale

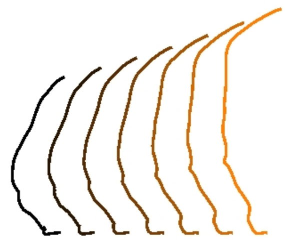

$\mathcal{S}_3$: shape+scale

$\mathcal{S}_4$: shape+orientation

$\mathcal{S}_5$: shape

Figure 2: **Evolution of one curve into another along the geodesic path.** The two curves are DTI fibers that vary in shape, scale and orientation as seen in the $\mathcal{S}_2$ manifold. The reparameterization allows for bending and stretching of curves. This elastic matching results for a smooth and natural transition between curves. In $\mathcal{S}_3$, the two curves are oriented similarly so we see the evolution of the shape and scale. In $\mathcal{S}_4$, they have the same scale so we see the transformation of shape and orientation. The two curves have the same orientation and scale in $\mathcal{S}_5$, and so we see one shape transform into the other. Figure credit: Anuj Srivastava/FSU.

3 Description of the Riemannian Manifolds, $\mathcal{S}_1 - \mathcal{S}_5$

In this section, we present five Riemannian manifolds which capture five different combinations of the physical features of interest. We also provide metrics for comparing curves based on these features. In particular, for each manifold we provide: (i) a geodesic distance between curves that depends only on selected features (and is independent of the parameterization of curves), and (ii) a geodesic path between the two curves. We begin with $\mathcal{S}_1$, a space where all the features are utilized. Each subsequent manifold is restricted to fewer features and $\mathcal{S}_5$, the last of these, uses only shape. Table 1 is a helpful summary of the manifolds and their preshape spaces and shape spaces.

3.1 Shape, orientation, scale and position, $\mathcal{S}_1$

We start by considering a situation where we are interested in comparing curves using all the four physical features – shape, scale, position and orientation. This might be useful when one of the four feature vectors (say position) predominates and the other three features are used to fine-tune the classification of curves. We shall see an example of this in Appendix **??** where we attempt to cluster and label sulci.

Preshape space: In this feature space, we use the square-root function (SRF) to represent the curve β:

$$h(t) = \sqrt{\|\dot{\beta}(t)\|}\beta(t) \ , \quad h : [0,1] \to \mathbb{R}^3 \ .$$

The SRFs are elements of the full $\mathbb{L}^2$ space so the preshape space is:

$$\mathcal{C}_1 = \{h \in \mathbb{L}^2([0,1], \mathbb{R}^3)\}.$$

Shape space: Since orientation, scale and translation are included in the representation, the only computation we need consider is invariance to reparametrization. The shape space is thus, $\mathcal{S}_1 = \mathcal{C}_1/\Gamma$.

When comparing two curves, we look at the difference between their SRFs in the shape space. For a curve β, the curve $\tilde{\beta}(t) \equiv \beta(\gamma(t))$ is simply the old curve with a new parameterization. For this reparameterized curve $\tilde{\beta}$, the SRF is given by $\tilde{h}(t) = \sqrt{\dot{\gamma}(t)}h(\gamma(t))$. We use (h, γ) to denote this reparameterized SRF. Now, it can be shown that for any two curves, β_1, β_2, with the corresponding SRFs, h_1 and h_2, and any $\gamma \in \Gamma$, we have that $\|(h_1, \gamma) - (h_2, \gamma)\|_2 = \|h_1 - h_2\|_2$. Because of this equality, we can define a distance between the two curves as:

$$d_1(\beta_1, \beta_2) = \min_{\gamma \in \Gamma} \left(\|h_1 - (h_2, \gamma)\|_2\right) \ . \tag{8}$$

This minimization is performed using the standard dynamic programming (DP) algorithm, and it results in a quantification of differences in curves that is associated with the aforementioned four features. The geodesic path between the curves is the straight line:

$$\psi(\tau) = (1 - \tau)h_1 + \tau(h_2, \gamma^*) \ , \tag{9}$$

where γ^* is the optimal reparameterization obtained earlier in minimization using DP.

3.2 Shape, orientation and scale, $\mathcal{S}_2$

Here we consider the case where we compare curves using all the feature information except position. The need for such computations which use shape, scale and orientation distances may arise when the focus of our study is a localized section of brain anatomy.

Preshape space: For this feature space we use the q-function to represent β:

$$q(t) = \frac{\dot{\beta}(t)}{\sqrt{\|\dot{\beta}(t)\|}} \ , \quad q : [0,1] \to \mathbb{R}^3 \ .$$

This function is different from the square-root function used in Section 3.1 in that this definition is based only on the velocity function $\dot{\beta}$. Thus we call it the *square-root velocity function* (SRVF) [2]. Since this function is invariant to a global translation of β, analysis based on it will not depend on the global coordinates of the curves. The SRVFs are elements of the full space but since we use q instead of h we define a new preshape space:

$$\mathcal{C}_2 = \{q \in \mathbb{L}^2([0,1], \mathbb{R}^3)\}.$$

Shape space: Since we want to include orientation and scale in the computations, the only invariance we need to consider is reparameterization. The shape space is thus, $\mathcal{S}_2 = \mathcal{C}_2/\Gamma$.

The SRVF of the reparameterized curve is given by $(q, \gamma) \equiv \sqrt{\dot{\gamma}(t)}q(\gamma(t))$, where q is the SRVF of the original curve. As before, it can be shown that for any two curves β_1, β_2, with the corresponding SRVFs q_1 and q_2, and for any $\gamma \in \Gamma$, we have that $\|(q_1, \gamma) - (q_2, \gamma)\|_2 = \|q_1 - q_2\|_2$. Once again, we can define a distance between the two curves as:

$$d_2(\beta_1, \beta_2) = \min_{\gamma \in \Gamma} \|q_1 - (q_2, \gamma)\|_2 \ . \tag{10}$$

This minimization is performed using the DP algorithm, and it results in a quantification in differences in curves based on the three features–shape, orientation, and scale. The geodesic path between the two curves is a straight line given by:

$$\psi(\tau) = (1 - \tau)q_1 + \tau(q_2, \gamma^*) \ , \tag{11}$$

where γ^* is the optimal reparameterization obtained by DP.

3.3 Shape and scale, $\mathcal{S}_3$

This combination of features allows us to compare curves based on their shapes and scales. For WM analysis, we note that the lengths of WM tracts are determined by the brain regions they connect. Two prominent fiber bundles, the inferior longitudinal fasciculus (ilf) and inferior fronto-occipital fasciculus (ifo) are both long. They connect the occipital lobe to the temporal lobe in the first case and to the frontal lobe in the second case and can be separated from each other by their shape. Short association fibers connecting adjacent gyri intermingle with these fiber bundles but can be separated out because of their length. Sulcal curves also come in different lengths; the central sulcus, which is a primary sulcus, is long, the tertiary sulci are much shorter. Being able to discriminate sulci and fibers on the basis of their shape and length is important.

Preshape space: For the *shape+scale* manifold, $\mathcal{S}_3$, we use the same SRVF representation as the $\mathcal{S}_2$ (*shape+orientation+scale*) manifold. Since the SRVFs are elements of the full space the preshape space is once again $\mathcal{C}_2 = \{q \in \mathbb{L}^2([0,1], \mathbb{R}^3)\}$.

Shape space: Curves in the $\mathcal{S}_3$ feature space are invariant to orientation. To remove the rigid motions from the $\mathcal{S}_2$ representation, we rotate the curve β by a rotation matrix $O \in SO(3)$, where $SO(3)$ the set of all possible rotations in $\mathbb{R}^3$. The shape space in this case is $\mathcal{S}_3 = \mathcal{C}_2/(\Gamma \times SO(3))$.

The SRVF of the rotated curve is given by Oq where q is the SRVF of the original curve. We also need to ensure invariance to reparameterization. Consequently, the SRVF of a rotated and reparameterized curve is given by $\sqrt{\dot{\gamma}(t)}Oq(\gamma(t))$. The distance function is a joint optimization over $SO(3)$ and the group of orientation preserving diffeomorphisms, Γ. It is given by:

$$d_3(\beta_1, \beta_2) = \min_{\gamma \in \Gamma, O \in SO(3)} \|q_1 - O(q_2, \gamma)\|_2 \ . \tag{12}$$

Let γ^* and O^* be the reparameterization and the rotation that minimize the right side in this equation. Then, the geodesic path between any two curves, which once again is a straight line, is:

$$\psi(\tau) = (1-\tau)q_1 + \tau(O^* q_2, \gamma^*) \ . \tag{13}$$

3.4 Shape and orientation, $\mathcal{S}_4$

The short U-fibers–white matter that connect adjacent gyri–are similar in shape but have different orientation. There are other structures that are identical in most respects but because of the bilateral symmetry in the brain, are oppositely oriented. Orientation is thus an important feature in brain image analysis.

We describe *shape+orientation* distances in this section. In Chapter 10, we use these distances along with shape distances to detect changes in the curvature of the corpus callosum along the median plane.

Preshape space: For the *shape+orientation* manifold, $\mathcal{S}_4$, we begin with the SRVF of the $\mathcal{S}_2$ (*shape+orientation+scale*) manifold and remove the scale component. To do this, we manually rescale the curves to be of the same length as described in Section 2.2.2. The mathematical representation of the SRVF remains the same as in $\mathcal{S}_2$ but the space of SRVFs, reduces from the full $\mathbb{L}^2$ space to a hypersphere as a result. This preshape space, the space of all unit-length (the curves are usually scaled to unit length) elastic curves, is defined as $\mathcal{C}_3 = \{q : [0,1] \to \mathbb{R}^3 \mid \int_0^1 ||q(t)||^2 dt = 1\}$.

The advantage of using a sphere is that the differential geometry is well-known and analytical expressions can be found for exponential maps, inverse exponential maps and geodesic paths. For example, if q_1 and q_2 are two elements of a unit hypersphere, the geodesic distance between them is given by the length of the shortest arc connecting them on the sphere. This length is $d_{\mathcal{C}_3} = \cos^{-1}(\int_0^1 \langle q_1(t), q_2(t) \rangle \, dt)$.

Shape space: As with the $\mathcal{S}_1$, $\mathcal{S}_2$, and $\mathcal{S}_3$ manifolds, invariance to reparameterization is achieved by the action of the group of diffeomorphisms, Γ, on $\mathcal{C}_3$. We shall denote this by $(q, \gamma) = \sqrt{\dot{\gamma(t)}}q(\gamma(t))$ and define the equivalence class of q as $[q] = \{(q, \gamma) \mid q \in S_\infty, \gamma \in \Gamma\}$. The shape space is $\mathcal{S}_4 = \mathcal{C}_3/(\Gamma)$. Note that we do not remove the $SO(3)$ group action on $\mathcal{C}_3$.

The distance between two curves does not depend on the reparameterization of the curves, i.e., for any q_1, q_2 and γ,

$$\cos^{-1}(\int_0^1 \langle q_1(t), q_2(t)\rangle\, dt) = \cos^{-1}(\int_0^1 \langle (q_1,\gamma)(t), (q_2,\gamma)(t)\rangle\, dt).$$

This leads to the definition of a distance between two curves which depends only on their shapes and orientations. The geodesic distance is calculated by minimizing the following:

$$d_4(\beta_1,\beta_2) = \min_{\gamma\in\Gamma}\left(\cos^{-1}(\int_0^1 \langle (q_1,\gamma)(t), (q_2,\gamma)(t)\rangle\, dt)\right) . \tag{14}$$

$\mathcal{S}_4$ is a sphere so the geodesic or shortest path between the two curves is a great circle. It can be specified analytically by:

$$\psi(\tau) = \frac{1}{\sin(\theta)}[\sin(\theta-\tau\theta)q_1 + \sin(\tau\theta)(q_2,\gamma^*)] , \tag{15}$$

where $\theta = d_4(\beta_1,\beta_2)$.

3.5 Shape manifold, $\mathcal{S}_5$

White matter fiber bundles have a well-defined shape and structure which is determined by the regions they connect and the constraints of the surrounding anatomy. Shape analysis of white matter fibers is an active area of study and is a starting point for our investigations in geometric modeling.

Preshape space: If we are interested in analyzing only the shape of 3-dimensional open curves, we must achieve invariance to translation, scaling, rotation and reparametrization. Since the q-function is defined in its entirety by the derivative of β, translational invariance is automatically removed. In order to achieve scale invariance, we manually scale all curves to be of unit length. This allows us to define the preshape space, $\mathcal{C}_3 = \{q : [0,1] \to \mathbb{R}^3 | \int_0^1 ||q(t)||^2 dt = 1\}$. This is the space of all unit length, elastic curves. We define a Riemannian metric and the tangent space on this manifold.

Shape space: We define the shape space to be:

$$\mathcal{S}_5 = \mathcal{C}_3/(\Gamma \times SO(3)).$$

The elements of $\mathcal{S}_5$ are the orbits of the type:

$$[q] = \{\sqrt{\dot\gamma} O q(\gamma) | q \in S_\infty, \gamma \in \Gamma, O \in SO(3)\} .$$

The shape space $\mathcal{S}_5$ inherits a Riemannian structure from the preshape space $\mathcal{C}_3$. The geodesic distance between any two orbits $[q_1]$ and $[q_2]$ is given by:

$$d_5(\beta_1,\beta_2) = d_{\mathcal{S}_5}([q_1],[q_2]) = \min_{\gamma\in\Gamma, O\in SO(3)} d_{\mathcal{C}_3}(q_1, \sqrt{\dot\gamma} O q_2(\gamma)) , \tag{16}$$

where $d_{\mathcal{C}_3}$ is the distance in the preshape space. A closer look at that distance function reveals the following:

$$\operatorname*{argmin}_{\gamma\in\Gamma, O\in SO(3)} \cos^{-1}\left\langle q_1, \sqrt{\dot\gamma} O q_2(\gamma)\right\rangle = \operatorname*{argmin}_{\gamma\in\Gamma, O\in SO(3)} \|q_1 - \sqrt{\dot\gamma} O q_2(\gamma)\|^2 , \tag{17}$$

Table 1: Summary of the Comprehensive Riemannian Framework Manifolds

Manifold	function representation	preshape space	quotient space
shape + orientation + scale + position	$h(t) = \sqrt{\|\dot{\beta}(t)\|}\beta(t)$	$\mathcal{C}_1$ $\mathbb{L}^2$ space	$\mathcal{S}_1 = \mathcal{C}_1/(\Gamma)$
shape + orientation + scale	$q(t) = \frac{\dot{\beta}(t)}{\sqrt{\|\dot{\beta}(t)\|}}$	$\mathcal{C}_2$ $\mathbb{L}^2$ space	$\mathcal{S}_2 = \mathcal{C}_2/(\Gamma)$
shape + scale	$q(t) = \frac{\dot{\beta}(t)}{\sqrt{\|\dot{\beta}(t)\|}}$	$\mathcal{C}_2$ $\mathbb{L}^2$ space	$\mathcal{S}_3 = \mathcal{C}_2/(\Gamma \times SO(3))$
shape + orientation	$q(t) = \frac{\dot{\beta}(t)}{\sqrt{\|\dot{\beta}(t)\|}}$	$\mathcal{C}_3$ hypersphere	$\mathcal{S}_4 = \mathcal{C}_3/(\Gamma)$
shape	$q(t) = \frac{\dot{\beta}(t)}{\sqrt{\|\dot{\beta}(t)\|}}$	$\mathcal{C}_3$ hypersphere	$\mathcal{S}_5 = \mathcal{C}_3/(\Gamma \times SO(3))$

The manifolds, $\mathcal{S}_1$–$\mathcal{S}_5$, are quotient spaces that result when the reparameterization group, Γ, or the product group, $(\Gamma \times SO(3))$, are removed from the respective preshape spaces. The optimizations are done using DP and Procrustes alignment. $\mathcal{C}_1$ and $\mathcal{C}_2$ are different $\mathbb{L}^2$ spaces since the curves have different function representations. $\mathcal{C}_3$ is a unit hypersphere because the curves have been scaled to unit length to remove the effects of scale.

where the $||.||$ is simply the $\mathbb{L}^2$ norm on the representation space. This equality implies that minimizing the arc length on a unit sphere is the same as minimizing the chord length. If one is minimized then so is the other. We can therefore use the $\mathbb{L}^2$ norm since it is computationally more efficient.

The actual geodesic between $[q_1]$ and $[q_2]$ in $\mathcal{S}_5$ is given by $[\psi_t]$, where ψ_t is the geodesic in $\mathcal{C}_3$ between q_1 and $\sqrt{\dot{\gamma}^*}O^* q_2(\gamma^*)$. Here (O^*, γ^*) are the optimal transformations of q_2 that minimize the right side in Eqn. 16. For $\theta = d_5(\beta_1, \beta_2)$, the geodesic path is a great circle given by:

$$\psi(\tau) = \frac{1}{\sin(\theta)}[\sin(\theta - \tau\theta)q_1 + \sin(\tau\theta)(O^* q_2, \gamma^*)] \ .$$

References

[1] L. Younes, P. W. Michor, J. Shah, and D. Mumford, "A metric on shape space with explicit geodesics," *Rend. Lincei Mat. Appl.*, vol. 9, pp. 25–57, 2008.

[2] S. H. Joshi, E. Klassen, A. Srivastava, and I. Jermyn, "A novel representation for Riemannian analysis of elastic curves in R^{n}," in *CVPR*, 2007.

[3] A. Srivastava, S. Joshi, W. Mio, and X. Liu, "Statistical shape analysis: Clustering, learning and testing," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 4, pp. 590–602, 2005.

[4] M. Mani, S. Kurtek, C. Barillot, and A. Srivastava, "A Comprehensive Riemannian Framework for the Analysis of White Matter Fiber Tracts," in *International Symposium on Biomedical Imaging (ISBI)*, (Rotterdam), pp. 1101–1104, IEEE, April 2010.

[5] S. Joshi, E. Klassen, A. Srivastava, and I. Jermyn, "Removing shape-preserving transformations in square-root elastic (SRE) framework for shape analysis of curves," in *Energy Minimization Methods in Computer Vision and Pattern Recognition*, vol. 4679 of *Lecture Notes in Computer Science*, pp. 387–398, 2007.

[6] N. Balov, A. Srivastava, C. Li, and Z. Ding, "Shape analysis of open curves in R^3 with applications to study of fiber tracts in DT-MRI data," in *EMMCVPR*, pp. 399–413, 2007.

[7] W. Liu, A. Srivastava, and E. Klassen, "Joint shape and texture analysis of objects boundaries in images using a Riemannian approach," in *Asilomar Conferences on Signals, Systems, and Computers*, 2008.

[8] E. Klassen, A. Srivastava, W. Mio, and S. Joshi, "Analysis of planar shapes using geodesic paths on shape spaces," *IEEE Patt. Analysis and Machine Intell.*, vol. 26, no. 3, pp. 372–383, 2004.

[9] M. P. do Carmo, *Differential Geometry of Curves and Surfaces*. Englewood Cliffs, NJ: Prentice-Hall, 1976.

[10] M. P. do Carmo, *Riemannian Geometry*. Boston, MA: Birkhäuser, 1992.

[11] J. Rotman, *An introduction to the theory of groups*. New York: Springer-Verlag, 1994.

[12] S. H. Joshi, *Inferences in Shape Space with Applications to Image Analysis and Computer Vision*. PhD thesis, Florida State University, Tallahassee, Florida, USA, 2007.

[13] L. Younes, "Computable elastic distances between shapes," *SIAM Journal on Applied Mathematics*, vol. 58, no. 2, pp. 565–586, 1998.

[14] E. Klassen and A. Srivastava, "Geodesics between 3D closed curves using path-straightening," in *ECCV 2006* (A. Leonardis, H. Bischof, and A. Pinz, eds.), vol. 3951 of *Lecture Notes in Computer Science*, pp. 95–106, Springer Berlin / Heidelberg, 2006.

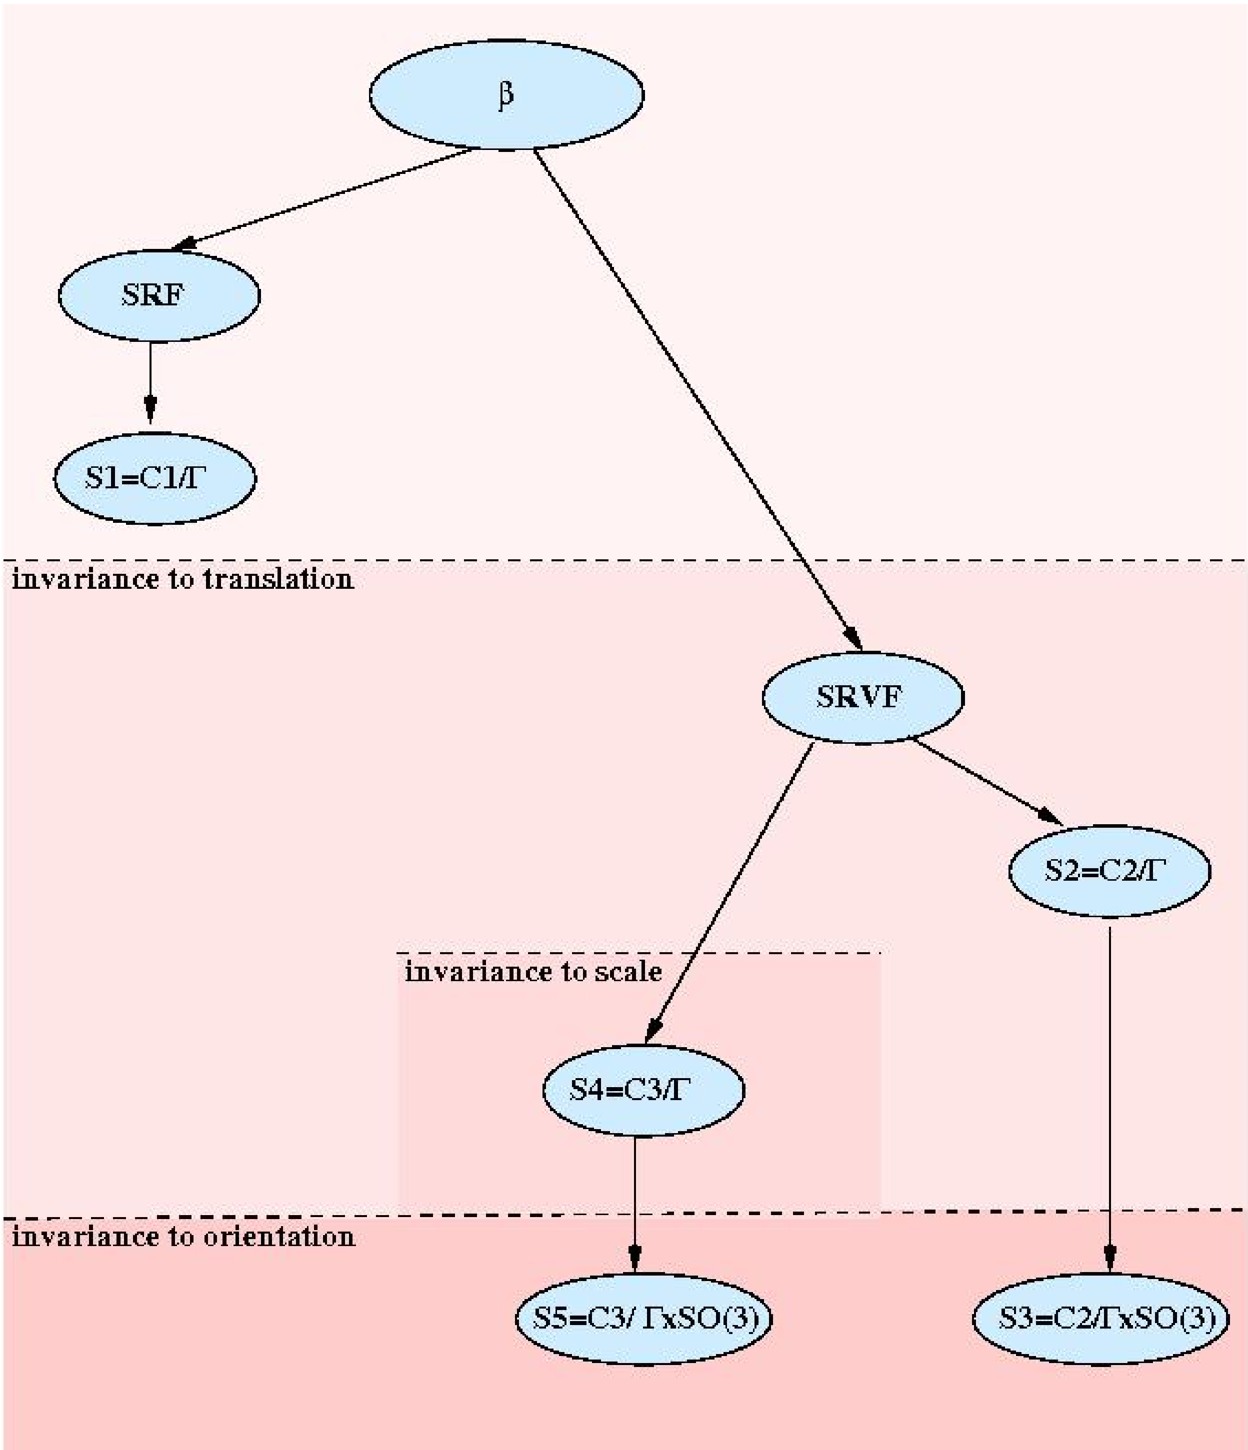

Figure 3: **Relationship between manifolds $\mathcal{S}_1 - \mathcal{S}_5$.** There is a hierarchy of steps that takes us from the curve β to its representation in the shape space. $\mathcal{S}_4$ and $\mathcal{S}_5$, for instance, follow the same mathematical procedures but the $\mathcal{S}_5$ space is in addition also invariant to orientation.

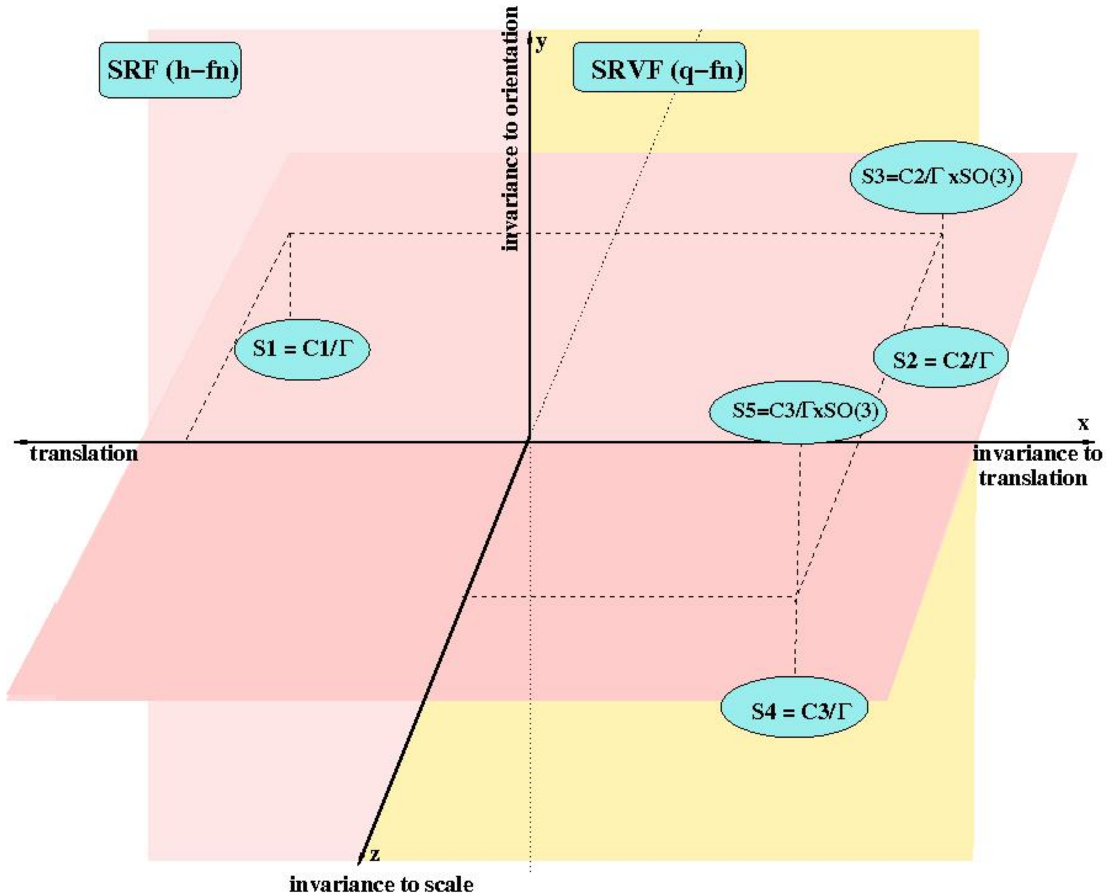

Figure 4: **Manifolds $\mathcal{S}_1 - \mathcal{S}_5$ and the invariances associated with them.** The comprehensive Riemannian framework allows for flexible feature combinations. Each of these constitute a Riemannian manifold and is associated with a subset of the four shape-preserving invariances: translation, orientation, scale and reparameterization. To visualize this, translation, orientation and scale are superimposed on an xyz-coordinate system. In the figure, which is multiperspective, since $\mathcal{S}_1$ includes position (translation), orientation and scale as features along with a SRF representation, it is located on the $-ve$ side of the $x-$, $y-$ and $z-$axes. $\mathcal{S}_2$, which is invariant to translation and uses the SRVF representation, is on the $+ve$ side of the $x-$axis. $\mathcal{S}_3$, invariant to translation and orientation, is on the $+ve$ $x-y$ quadrant; $\mathcal{S}_4$, invariant to translation and scale, is on the $+ve$ $x-z$ quadrant. Finally $\mathcal{S}_5$, the *shape* manifold, is invariant to translation, orientation and scale and has $+ve$ x, y, z values.