Signatures in Shape Analysis: An Efficient Approach to Motion Identification

Elena Celledoni[ID], Pål Erik Lystad[ID], and Nikolas Tapia[(✉)][ID]

Department of Mathematical Sciences, NTNU Trondheim, Trondheim, Norway
`nikolas.tapia@ntnu.no`

Abstract. Signatures provide a succinct description of certain features of paths in a reparametrization invariant way. We propose a method for classifying shapes based on signatures, and compare it to current approaches based on the SRV transform and dynamic programming.

Keywords: Shape analysis · Signature · Motion identification

1 Introduction

Shape analysis is a broad and growing subject addressing the analysis of different types of data ranging from surfaces, landmarks, animation data etc. In this paper shapes are unparametrized curves. Mathematically a shape is an equivalence class of curves under reparameterization, that is, two curves $c_0, c_1 : [0,1] \to M$ are equivalent and determine the same shape if there exists a strictly increasing smooth bijection $\varphi : [0,1] \to [0,1]$ such that $c_1 = c_0 \circ \varphi$. For a given curve c we denote by $[c]$ the corresponding shape.

The similarity between two shapes $[c_0], [c_1]$ is then defined by creating a distance function $d_{\mathcal{S}}$ on the space of shapes $\mathcal{S}$,

$$d_{\mathcal{S}}([c_0], [c_1]) := \inf_{\varphi} d_{\mathcal{P}}(c_0, c_1 \circ \varphi) \tag{1}$$

where $d_{\mathcal{P}}$ is a suitable reparameterization invariant Riemannian distance on the manifold of parametrized curves.

Finding the optimal reparameterization φ is however computationally demanding, and in many applications simply unnecessary. This is specifically the case of applications where the optimal parametrization is not explicitly used for further calculations, e.g. problems of identification and classification. Ways of circumventing this step are therefore of great interest.

In recent years, after extensive work by Terry Lyons and collaborators, the theory of rough paths has gained considerable importance as a toolbox for mathematical analysis and for mathematical modeling in applications. In this context, the signature map provides a faithful representation of paths, capturing their essential global properties. A fundamental property of the signature is its invariance under reparameterization, surmising its importance for shapes.

In this paper, we define a measure of similarity between shapes in $\mathcal{S}$ by means of the signature. We define a distance directly on $\mathcal{S}$. We test the viability of this approach and use it to classify motion capture animations from the CMU motion capture database [7]. Indeed, this leads to an efficient technique that delivers results comparable to what is obtainable with methodologies based on the SRV transform, but at a much lower computational cost.

2 Shape Analysis on Lie Groups

In the following, G will denote a finite-dimensional Lie group under multiplication with identity element denoted by e. We let $\mathfrak{g}$ denote the corresponding right Lie algebra $\mathfrak{g} := \mathcal{L}_R(G)$. For a fixed $g \in G$, left and right translation by g will be denoted $L_g(h) = g \cdot h$ and $R_g(h) = h \cdot g$ respectively.

2.1 Shape Space

We consider the space $C^\infty([0,1], G)$ of parameterized smooth curves on G, i.e. smooth maps $c : [0,1] \to G$. To model the curves as unparameterized, or independent of parameterization, we define the *shape space* $\mathcal{S}$ as the quotient space

$$\mathcal{S} = C^\infty([0,1], G)/\mathrm{Diff}^+, \tag{2}$$

where Diff^+ is the group of orientation preserving diffeomorphisms of the parameter space $[0,1]$. The elements of $\mathcal{S}$ are equivalence classes of curves. The elements of the same class are curves which can be mapped to one another by changing their parameterization, that is, two curves $c_0, c_1 \in C^\infty(I, G)$ are equal in shape space if there exists $\varphi \in \mathrm{Diff}^+$ such that $c_1 = c_0 \circ \varphi$.

In the setting of our application, the search for optimal time parametrizations can be viewed as syncing up the animations, removing disturbances due to small pauses, different periodicity, or asynchronous starting and stopping, by shifting the movement of one character to match the other as closely as possible.

2.2 Geodesic Distances on Shape Space

Our goal is to introduce a meaningful and computable distance $d_\mathcal{S}$ on $\mathcal{S}$ to estimate the similarity between two shapes. This area of research started with the efforts of Younes [16]. We will restrict the space of curves to the space of immersions, i.e. curves with non-vanishing first derivative, which we denote by

$$\mathcal{P} = \mathrm{Imm}([0,1], G). \tag{3}$$

Let $d_\mathcal{P}$ be a pseudo-metric on $\mathcal{P}$. We define $d_\mathcal{S}$, for two elements $[c_0], [c_1] \in \mathcal{S}$, by

$$d_\mathcal{S}([c_0], [c_1]) := \inf_{\varphi \in \mathrm{Diff}^+} d_\mathcal{P}(c_0, c_1 \circ \varphi). \tag{4}$$

As shown in [3, Lemma 3.4], d_S will be a pseudo-metric on S if d_P is a *reparameterization invariant* or, in other words, if for any two $c_0, c_1 \in \mathcal{P}$ and any $\varphi \in \text{Diff}^+$ we have that

$$d_P(c_0 \circ \varphi, c_1 \circ \varphi) = d_P(c_0, c_1). \tag{5}$$

An obvious choice of metric on $\mathcal{P}$ is the familiar L_2-metric. However, as shown by Michor and Mumford [13], this metric leads to vanishing geodesic distance which renders it useless. They further show in [14] that one solution to this problem is to consider metrics based on arc-length derivatives, creating a class of Sobolev-type metrics.

There are multiple possible metrics in this class. One option is based on what is usually referred to as the *Square Root Velocity Transform* (SRVT). This transform and accompanying metric was first introduced, in the context of shape analysis, by Srivastava et al. [15], who used the transformation when working with curves in Euclidian spaces. The transformation has later been adopted to more general shapes. Of particular interest is the formulation for shapes that are represented as Lie-group valued curves [3].

We define the SRVT $\mathcal{R} : \mathcal{P} \rightarrow C^\infty([0, 1], \mathfrak{g} \setminus \{0\})$ by

$$\mathcal{R}(c)(t) := \frac{R_{c(t)*}^{-1}(\dot{c}(t))}{\sqrt{\|\dot{c}(t)\|}}. \tag{6}$$

This transformation has the following useful properties [3, Lemma 3.6]:

1. For every $c \in \mathcal{P}$ and $\varphi \in \text{Diff}^+$, the following equivariant property holds:

$$\mathcal{R}(c \circ \varphi) = \mathcal{R}(c) \circ \varphi \cdot \sqrt{\dot{\varphi}}. \tag{7}$$

2. It is translation invariant: for all $c \in \mathcal{P}$ and $g \in G$

$$\mathcal{R}(R_g(c)) = \mathcal{R}(c).$$

A similar result is true for shapes with values in Euclidean spaces [15].

Further, one can obtain a Riemannian metric $d_{\mathcal{P}_*}$ that coincides with the geodesic distance on a submanifold $\mathcal{P}_* \subset \mathcal{P}$ by using the SRVT to pull back the L_2-metric on $C^\infty(I, \mathfrak{g} \setminus \{0\})$ [3]. Further restricting the immersion space to $\mathcal{P}_* = \{c \in \mathcal{P} : c(0) = e\}$, where e is the identity element in G, the distance $d_{\mathcal{P}_*}$ turns out to be reparameterization invariant.

This invariance implies, in particular, that it will also yield a geodesic distance on $\mathcal{S}_* := \mathcal{P}_*/\text{Diff}^+$ [2]. The restriction to $\mathcal{P}_*$ isn't very troublesome as any curve can be transferred to this space by right translation by the inverse of its initial value, that is $R_{c(0)^{-1}}$ [3].

Using the equivariant property for the SRVT from Eq. (7) and defining $q_i = \mathcal{R}(c_i)$ for $i = 0, 1$, the problem of calculating the metric for the shape space $\mathcal{S}_*$ in Eq. (4) can be written as

$$d_{\mathcal{S}_*}(c_0, c_1) = \inf_{\varphi \in \text{Diff}^+(I)} \sqrt{\int_I \|q_0(t) - q_1(\varphi(t)) \cdot \sqrt{\dot{\varphi}}\|^2 dt}. \tag{8}$$

Finding this infimum will generally be very difficult. The usual approach is therefore to discretize the curves and solve instead a finite dimensional optimization problem. The most common methods used to solve this problem in shape analysis [15] are based on either the gradient descent method or a dynamic programming algorithm (DP). In our experiments we use the DP approach described in [1].

3   Signatures

Signatures, introduced by Chen [4] for smooth paths and later generalized by Lyons [11] under the name of geometric rough paths, are an important tool for the study of the solutions of controlled differential equations, but have also proved useful for solving classification problems of time series, Machine Learning and Topological Data Analysis [6].

In the usual framework, signatures are defined for paths taking values in a Banach space. From a geometric point of view, and in light of our purposes, this setting has to be adapted. Luckily, Chen also considered signatures for curves taking values on a smooth manifold [4]. This definition is quite general and relies on the selection of a frame bundle. For Lie groups there is a canonical choice: the Maurer–Cartan form. This is the unique right-invariant one form ω such that $\omega_e = \mathrm{id}_{\mathfrak{g}}$, i.e. $\omega(v) = (R_g^{-1})_* v$ for $v \in T_g(G)$ [8, p. 311].

Below we denote, for a finite-dimensional vector space V of dimension $d = \dim V$, the tensor algebra over V,

$$T(V) := \bigoplus_{n \geq 0} V^{\otimes n}.$$

We observe that $T(V)$ is always infinite-dimensional. Its dual space is denoted by $T((V)) := T(V)^*$, and it may be identified with the ring of formal power series in d noncommuting variables $\{e_1, \dots, e_d\}$.

Definition 1. *Let G be a d-dimensional Lie group and $\alpha \in C^\infty([0,1], G)$ be a smooth curve and ω the Maurer-Cartan form on G. The signature $S(\alpha)$ of α is the family of linear maps on $T(\mathbb{R}^d)$ recursively defined by $\langle S(\alpha)_{s,t}, 1 \rangle := 1$ and*

$$\langle S(\alpha)_{s,t}, e_{i_1 \cdots i_p} \rangle := \int_s^t \langle S(\alpha)_{s,u}, e_{i_1 \cdots i_{p-1}} \rangle \, \omega^{i_p}_{\alpha(u)}(\dot\alpha(u)) \, \mathrm{d}u.$$

In this definition, the notation $\omega^j_g(v)$ denotes the j-th component of the vector $\omega_g(v) \in \mathfrak{g}$ in a basis of the Lie algebra $\mathfrak{g}$ of G.

The signature provides a compact description of certain features of a path [5]. One of its main advantages in our context is its reparameterization invariance: for any orientation-preserving diffeomorphism φ on $[s,t]$ we have that

$$S(\alpha \circ \varphi)_{s,t} = S(\alpha)_{s,t}.$$

Other fundamental properties include:

1. For each $0 \leq s < t \leq 1$, the signature $S(x)_{s,t}$ belongs to the set of group-like elements of $T((\mathbb{R}^d))$, and for any $0 \leq s \leq 1$, $S(x)_{s,s} = 1$, the neutral element in the group.
2. **Chen's rule:** For any three $0 \leq s < u < t \leq 1$ we have

$$S(x)_{s,u} \otimes S(x)_{u,t} = S(x)_{s,t}.$$

Using these properties, signatures may be efficiently computed for some restricted classes of paths. For example, if x is a straight line in $\mathbb{R}^d$ with base point $a \in \mathbb{R}^d$ direction $b \in \mathbb{R}^d$, i.e. $x_t = a + tb$ for $t \in [0, 1]$, then

$$
\begin{aligned}
S(x)_{s,t} &= \exp_\otimes ((t - s)b) \\
&= 1 + (t - s)b + \frac{(t - s)^2}{2} b \otimes b + \frac{(t - s)^3}{6} b \otimes b \otimes b + \cdots .
\end{aligned}
\tag{9}
$$

A similar statement is true for geodesic curves on a finite-dimensional compact Lie group.

We may think of signatures as an infinite vector indexed by *words* over the alphabet $\{1, \ldots, d\}$. In particular, for a piecewise linear path the above formula means that if we want to know the component in (9) corresponding to the word $w = i_1 \cdots i_k$ then

$$\langle S(x)_{s,t}, e_w \rangle = \frac{(t - s)^k}{k!} \prod_{j=1}^{k} b_{i_j}$$

For a general piecewise linear path x, we may use the above formula and Chen's rule to deduce that

$$S(x)_{s,t} = \exp_\otimes (\Delta t_1 b_1) \otimes \exp_\otimes (\Delta t_2 b_2) \otimes \cdots \otimes \exp_\otimes (\Delta t_m b_m)$$

where $\Delta t_k = t_k - t_{k-1}$ are the length of the time intervals where the path is sampled and $b_1, \ldots, b_k$ are the slopes of the path in each of these intervals. The entries of this expression may be computed by using a Baker–Campbell–Hausdorff-type formula, for example.

Finally, we remark that the signature possesses another interesting property, namely it is an homomorphism from path space with concatenation to the tensor algebra $T((\mathbb{R}^d))$. This means that if we are given two paths $x \colon [0, 1] \to G$ and $y \colon [0, 1] \to G$, and we concatenate them to form a new path $x \cdot y$, then

$$S(x \cdot y)_{0,1} = S(x)_{0,1} \otimes S(y)_{0,1}.$$

Moreover, if we reverse the path x, i.e. we define $\overleftarrow{x}(t) := x(1 - t)$ then

$$S(\overleftarrow{x})_{0,1} = S(x)_{0,1}^{-1}$$

where the inverse is taken in the group-like elements of the tensor algebra.

It can be shown that actually, as a function of time the signature satisfies the differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t}S(x)_{s,t} = S(x)_{s,t} \otimes \dot{x}_t, \quad S(x)_{s,s} = 1$$

in the tensor algebra. From this point of view, the signature map corresponds to the flow map of the vector field given by the base path. Thus, the signature belongs to an infinite-dimensional Lie group whose Lie algebra is the free Lie algebra over $\mathbb{R}^d$ which we denote by $\mathfrak{L}(\mathbb{R}^d)$. It does not, however, constitute a one-parameter subgroup. Therefore, for each fixed time interval $[s,t]$ we can map the signature to the free Lie algebra via a logarithm map, and we define

$$\Lambda(x)_{s,t} = \log(S(x)_{s,t}) \in \mathfrak{L}(\mathbb{R}^d).$$

This element, called the log-signature in the literature, provides a minimal description of the path, which is equivalent to the full signature.

There are many ways in which signatures can be used to compare shapes, but the essential feature is that since the map S is reparameterization invariant, one obtains a way of directly comparing shapes instead of parameterized curves. For our experiments we chose a particular distance on $T((\mathbb{R}^d))$ (see next section for the precise formula), but this is by no means the only possible choice.

In making this choice one has to truncate the signature to obtain a finite-dimensional object. Due to the factorial decay of iterated integrals little information is lost in the process; still, some level has to be chosen and usually this done by running experiments. Once the truncation level is chosen, several choices of metric are available: the truncated tensor algebra becomes finite-dimensional so it has a nice linear structure and we are free to choose norms on it subject to some compatibility restrictions. There is also the notion of homogeneous norm on group-like elements, which takes into account the geometry of this group. Finally, the logarithm in this group maps signatures into a linear space (the free Lie algebra) in a bijective way, so no information is lost, but there is a substantial dimensional reduction.

According to our observations, is the last option which represents the most robust choice in terms of noise sensitivity, while also providing an accurate way of comparing signatures.

4   Experiments

Motion capture animations are usually recorded as the angle of every joint in a skeleton for every frame in an animation. A natural setting for the rotating joints is the Lie group of 3D rotations, $SO(3)$. Every frame consists of d independently rotating joints so the frame can be modeled as an element in $SO(3)^d$, where $SO(3)^d$ is the Cartesian product of d copies of $SO(3)$. Interpolating between the frames will then allow us to model the animation as a parameterized curve.

We use an interpolation scheme in which one uses the log map to linearly interpolate on the Lie algebra, and then pull back to the Lie group with the

exponential map. Let $A, B \in SO(3)$, we define the interpolation $\kappa : [0,1] \rightarrow SO(3)$ between A and B as

$$\kappa(s) := \exp\left(s\log\left(B \cdot A^T\right)\right) \cdot A.$$

Notice that $\kappa(0) = A$ and $\kappa(1) = B$. Applying this interpolation component-wise to the frames in $SO(3)^d$ will enable us to construct a piece-wise interpolation between the frames of the animation. The Maurer–Cartan form along the interpolation is piece-wise constant, making it easy to compute SRV representations, $d_{\mathcal{P}_*}$-metrics, and signatures.

To test the effectiveness of the proposed frameworks we check whether they are able to identify different types of character motion. We have selected animations from the CMU motion capture database with descriptions "walk", "run/jog" and "forward jump". These are similar in length, and should produce results that conform with human intuition.

The test will calculate a distance matrix using the proposed similarity measures. From the distance matrix we produce a multidimensional scaling plot (MDS), depicting how similar, or dissimilar, the animations are. MDS tries to place the data points in 2-dimensional scatter plot while preserving the distances given by the distance matrix. See Kruskal [9] for more information on this method.

In Fig. 2a we calculate the distance matrix using the metric $d_{\mathcal{P}_*}$ on interpolation curves in $\mathcal{P}_*$, and in Fig. 2b we use the metric $d_{\mathcal{S}_*}$, Eq. (8), on the shapes generated by the curves in $\mathcal{S}_*$, where the optimal reparameterization is calculated with a DP algorithm. There are little to no patterns when projecting to the space $(\mathcal{P}_*, d_{\mathcal{P}_*})$, as seen in Fig. 2a. In Fig. 2b however, we observe that modelling
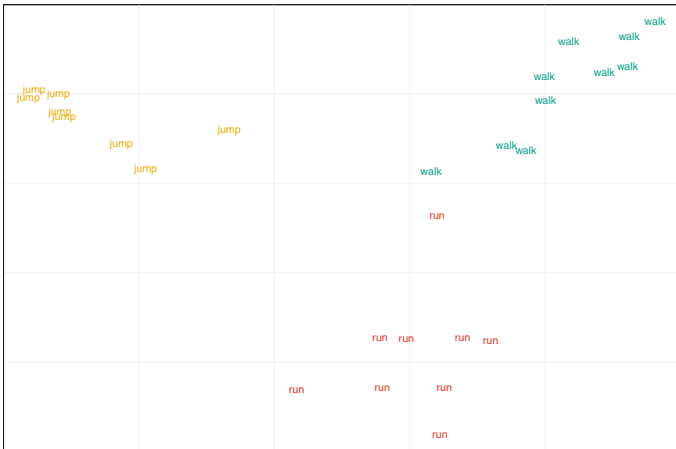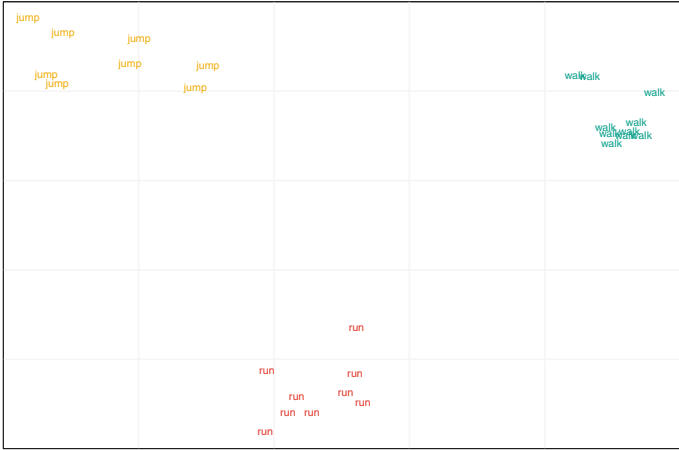

Fig. 1. Multi dimensional scaling plot of distance matrix calculated from by projecting animations to the space $\mathcal{S}_*$ equipped with the distance function d_{sig}. In this plot we have taken animation with descriptions "run/jog", "forward jump" and "walk" from the CMU Motion Capture Database [7].

the curves as being parameterization invariant yields three easily distinguishable clusters of animations. Compared to Fig. 2a we see a big benefit from this model assumption.

(a) Animations projected to $\mathcal{P}_*$ with distance matrix calculated with the metric $d_{\mathcal{P}_*}$.

(b) Animations projected to $\mathcal{S}_*$ with distance matrix calculated with metric $d_{\mathcal{S}_*}$ using a DP algorithm.

Fig. 2. Multi dimensional scaling plots of distance matrix based on geodesic distances calculated in $\mathcal{P}_*$ and $\mathcal{S}_*$, figure (a) and (b) respectively. In this plot we have taken animation with descriptions "run/jog", "forward jump" and "walk" from the CMU Motion Capture Database [7].

In Fig. 1 the animations are projected to the shape space $\mathcal{S}$ equipped with the distance function $d_{\text{sig}}(c_0, c_1) = \left\| \frac{\log S(c_0)}{\|\log S(c_0)\|} - \frac{\log S(c_1)}{\|\log S(c_1)\|} \right\|$. While this figure does reveal the same structure as seen in Fig. 2b, the clusters exhibit both a higher internal and a lower external variability. An important take away from this experiment is that this distance function in fact does preserve some of the structure of the shape space.

5   Concluding Remarks

Our preliminary experiments, show that classifying animations using a distance function on $\mathcal{S}_*$ based on signatures produces very encouraging results. The proposed method is computationally very efficient, even though somewhat less accurate than known methods in shape analysis.

The Riemannian metric (4) requires calculating the optimal reparameterizations between every pair of animations. The proposed signature method instead only requires calculating the signature once for every animation, and then compares animations by computing inexpensive norms. The optimisation procedure is no longer necessary.[1]

In our experiments, the signature method outperformed the optimal reparameterization metric by a factor of $\sim$2000 when classifying animations. A more precise comparison with the SRVT approach and other methods, see e.g. [10] goes beyond the scope of this work and will be considered in future work. Still our preliminary experiments give an idea of the possible performance benefits gained with the signature approach.

Increasing the accuracy of the signature method might also be possible by defining a more precise similarity measure. Nonetheless, our results can be seen as proof of concept for using signatures as an efficient way of classifying shapes.

---

[1] A more thorough analysis of the run time complexities associated with these algorithms has been left out due to space constraints.

References

1. Bauer, M., Eslitzbichler, M., Grasmair, M.: Landmark-guided elastic shape analysis of human character motions. Inverse Probl. Imaging **11**, 601–621 (2015). https://doi.org/10.3934/ipi.2017028
2. Bruveris, M., Michor, P., Mumford, D.: Geodesic completeness for Sobolev metrics on the space of immersed plane curves. Forum Math. Sigma [electronic only] **2** (2013). https://doi.org/10.1017/fms.2014.19
3. Celledoni, E., Eslitzbichler, M., Schmeding, A.: Shape analysis on lie groups with applications in computer animation. J. Geom. Mech. **8**(3), 273–304 (2016)
4. Chen, K.T.: Iterated integrals and exponential homomorphisms. Proc. Lond. Math. Soc. **s3-4**(1), 502–512 (1954). https://doi.org/10.1112/plms/s3-4.1.502
5. Chen, K.T.: Integration of paths-a faithful representation of paths by noncommutative formal power series. Trans. Am. Math. Soc. **89**(2), 395–407 (1958). http://www.jstor.org/stable/1993193
6. Chevyrev, I., Nanda, V., Oberhauser, H.: Persistence paths and signature features in topological data analysis. IEEE Trans. Pattern Anal. Mach. Intell. (2018, to appear). https://doi.org/10.1109/TPAMI.2018.2885516
7. CMU Graphics Lab: CMU Graphics Lab Motion Capture Database. http://mocap.cs.cmu.edu/. Accessed 10 Dec 2018
8. Hilgert, J., Neeb, K.H.: Structure and Geometry of Lie Groups. Springer, New York (2012). https://doi.org/10.1007/978-0-387-84794-8
9. Kruskal, J., Wish, M.: Multidimensional Scaling. Sage Publications, Newbury Park (1978)
10. Lahiri, S., Robinson, D., Klassen, E.: Precise matching of PL curves in $\mathbb{R}^N$ in the square root velocity framework. Geom. Imaging Comput. **2** (2015). https://doi.org/10.4310/GIC.2015.v2.n3.a1
11. Lyons, T.J.: Differential equations diriven by rough signals. Revista Matemática Iberoamericana **14**(2), 215–310 (1998). https://doi.org/10.4171/RMI/240
12. Lystad, P.E.: Signatures in shape analysis: an efficient approach to motion identification. Master's thesis, Department of Mathematical Sciences, NTNU, Trondheim, Norway (2019, manuscript submitted for publication)
13. Michor, P., Mumford, D.: Vanishing geodesic distance on spaces of submanifolds and diffeomorphisms. Doc. Math. **10**, 217–245 (2005)
14. Michor, P., Mumford, D.: An overview of the Riemannian metrics on spaces of curves using the Hamiltonian approach. Appl. Comput. Harmon. Anal. **23**, 74–113 (2006). https://doi.org/10.1016/j.acha.2006.07.004
15. Srivastava, A., Klassen, E., Joshi, S.H., Jermyn, I.H.: Shape analysis of elastic curves in Euclidean spaces. IEEE Trans. Pattern Anal. Mach. Intell. **33**(7), 1415–1428 (2011)
16. Younes, L.: Computable elastic distances between shapes. SIAM J. Appl. Math. **58**, 565–586 (1998)