

Parallel Transport on Kendall Shape Spaces

Nicolas Guigui¹[0000-0002-7901-0732], Elodie Maignant^{1,2}, and Alain Trouvé²
Xavier Pennec¹

¹ Université Côte d’Azur, Inria Epione project team, France

² Centre Borelli, ENS Paris Saclay

Abstract. Kendall shape spaces are a widely used framework for the statistical analysis of shape data arising from many domains, often requiring the parallel transport as a tool to normalise time series data or transport gradient in optimisation procedures. We present an implementation of the pole ladder, an algorithm to compute parallel transport based on geodesic parallelograms and compare it to methods by integration of the parallel transport ordinary differential equation.

Keywords: Parallel Transport · Shape Spaces.

1 Introduction

Kendall shape spaces are a ubiquitous framework for the statistical analysis of data arising from medical imaging, computer vision, biology, chemistry and many more domains. The underlying idea is that a shape is what is left after removing the effects of rotation, translation and re-scaling, and to define a metric accounting for those invariances. This involves defining a Riemannian submersion and the associated quotient structure, resulting in non trivial differential geometries with singularities and curvature, requiring specific statistical tools to deal with such data.

In this context parallel transport is a fundamental tool to define statistical models and optimisation procedures, such as the geodesic or spline regression [6,12] and for the normalisation of time series of shapes [8,1]. However parallel transport is defined by an ordinary differential equation (ODE) and there is usually no closed-form solution. Approximation methods have therefore been derived, either by direct integration [6], or by integration of the geodesic equation to approximate Jacobi fields (the fanning scheme [9]). Another class of approximations referred to as ladder methods, relies on iterative constructions of geodesic parallelograms that only require approximate geodesics [3].

In this work, we present an implementation of the pole ladder that leverages the quotient structure of Kendall shape spaces and strongly relies on the open-source Python package `geomstats`. We compare it to the method of Kim et al. [6] by approximate integration.

We first recall the quotient structure of Kendall shape spaces and its use by Kim et al. to compute parallel transport in sec. 2, then in sec. 3 we recall the pole ladder scheme and the main result from [3] on its convergence properties. Numerical simulations to compare the two methods are reported in sec. 4.

2 The quotient structure of Kendall shape spaces

We first describe the shape space, as quotient of the space of configurations of k points in \mathbb{R}^m (called landmarks) by the groups of translations, re-scaling and rotations of \mathbb{R}^m . For a thorough treatment this topic, we refer the reader to [4,2].

2.1 The pre-shape space

We define the space of k landmarks of \mathbb{R}^m as the space of $m \times k$ matrices $M(m, k)$. For $x \in M(m, k)$, let x_i denote the columns of x , i.e. points of \mathbb{R}^m and let \bar{x} be their barycentre. We remove the effects of translation by considering the matrix with columns $x_i - \bar{x}$ instead of x . We further remove the effects of scaling by dividing x by its Frobenius norm (written $\|\cdot\|$). This defines the pre-shape space $\mathcal{S}_m^k = \{x \in M(m, k) \mid \sum_{i=1}^k x_i = 0, \|x\| = 1\}$, which is identified with the hypersphere of dimension $m(k-1) - 1$. The pre-shape space is therefore a differential manifold whose tangent space at any $x \in \mathcal{S}_m^k$ is given by $T_x \mathcal{S}_m^k = \{w \in M(m, k) \mid \sum_{i=1}^k w_i = 0, \text{Tr}(w^T x) = 0\}$.

The ambient Frobenius metric $\langle \cdot, \cdot \rangle$ thus defines a Riemannian metric on the pre-shape space, with constant sectional curvature and known geodesics: let $x, y \in \mathcal{S}_m^k$ with $x \neq y$, and $w \in T_x \mathcal{S}_m^k$

$$x_w = \exp_x(w) = \cos(\|w\|)x + \sin(\|w\|)\frac{w}{\|w\|}, \quad (1)$$

$$\log_x(y) = \arccos(\langle y, x \rangle) \frac{y - \langle y, x \rangle x}{\|y - \langle y, x \rangle x\|}. \quad (2)$$

Moreover, this metric is invariant to the action of the rotation group $SO(m)$. This allows to define the shape space as the quotient $\Sigma_m^k = \mathcal{S}_m^k / SO(m)$.

2.2 The shape space

To remove the effect of rotations, we define the equivalence relation \sim on \mathcal{S}_m^k by $x \sim y \iff \exists R \in SO(m)$ such that $y = Rx$. For $x \in \mathcal{S}_m^k$, let $[x]$ denote its equivalence class for \sim . This equivalence relation results from the group action of $SO(m)$ on \mathbb{R}^m . This action is smooth, proper but not free everywhere when $m \geq 3$. This makes the orbit space $\Sigma_m^k = \{[x] \mid x \in \mathcal{S}_m^k\}$ a differential manifold with singularities where the action is not free.

One can describe these singularities explicitly : they correspond to the matrices of \mathcal{S}_m^k of rank $m - 2$ or less [7]. For $k \geq 3$, the spaces Σ_1^k and Σ_2^k are always smooth. Moreover, as soon as $m \geq k$, the manifold acquires boundaries. As an example, while the space Σ_2^3 of 2D triangles is identified with the sphere $S^2(1/2)$, the space Σ_3^3 of 3D triangles is isometric to a 2-ball [7].

Away from the singularities, the canonical projection map $\pi : x \mapsto [x]$ is a Riemannian submersion, and plays a major role in defining the metric on the shape space. Let $d_x \pi$ be its differential map at $x \in \mathcal{S}_m^k$, whose kernel defines the

vertical tangent space, which corresponds to the tangent space of the submanifold $\pi^{-1}([x])$, called fiber above $[x]$:

$$\text{Ver}_x = \{Ax \mid A \in \text{Skew}(m)\} = \text{Skew}(m) \cdot x$$

where $\text{Skew}(m)$ is the space of skew-symmetric matrices of size m .

2.3 The quotient metric

The Frobenius metric on the pre-shape space allows to define the horizontal spaces as the orthogonal complements to the vertical spaces:

$$\begin{aligned} \text{Hor}_x &= \{w \in T_x \mathcal{S}_m^k \mid \text{Tr}(Axw^T) = 0 \forall A \in \text{Skew}(m)\} \\ &= \{w \in T_x \mathcal{S}_m^k \mid xw^T \in \text{Sym}(m)\} \end{aligned}$$

where $\text{Sym}(m)$ is the space of symmetric matrices of size m . Lemma 1 from [12] allows to compute the vertical component of any tangent vector:

Lemma 1. *For any $x \in \mathcal{S}_m^k$ and $w \in T_x \mathcal{S}_m^k$, the vertical component of w can be computed as $\text{Ver}_x(w) = Ax$ where A solves the Sylvester equation:*

$$Axx^T + xx^T A = wx^T - xw^T \quad (3)$$

If $\text{rank}(x) \geq m - 1$, A is the unique skew-symmetric solution of (3).

In practice, the Sylvester equation can be solved by an eigenvalue decomposition of xx^T . This defines ver_x , the orthogonal projection on Ver_x . As $T_x \mathcal{S}_m^k = \text{Ver}_x \oplus \text{Hor}_x$, any tangent vector w at $x \in \mathcal{S}_m^k$ may be decomposed into a horizontal and a vertical component, by solving (3) to compute $\text{ver}_x(w)$, and then $\text{hor}_x(w) = w - \text{ver}_x(w)$.

Furthermore, as $\text{Ver}_x = \ker(d_x \pi)$, $d_x \pi$ is a linear isomorphism from Hor_x to $T_{[x]} \Sigma_m^k$. The metric on Σ_m^k is defined such that this isomorphism is an isometry. Note that the metric does not depend on the choice of the y in the fiber $\pi^{-1}([x])$ since all y in $\pi^{-1}([x])$ may be obtained by a rotation of x , and the Frobenius metric is invariant to the action of rotations. This makes π a Riemannian submersion. Additionally, π is surjective so for every vector field on Σ_m^k there is a unique horizontal lift, i.e. a vector field on \mathcal{S}_m^k whose vertical component is null everywhere. The tangent vectors of Σ_m^k can therefore be identified with horizontal vectors of \mathcal{S}_m^k . One of the main characteristics of Riemannian submersions was proved by O'Neill [13]:

Theorem 1 (O'Neill). *Let $\pi : M \rightarrow B$ be a Riemannian submersion. If γ is a geodesic in M such that $\dot{\gamma}(0)$ is a horizontal vector, then $\dot{\gamma}$ is horizontal everywhere and $\pi \circ \gamma$ is a geodesic of B of the same length as γ .*

Remark 1. We emphasise that an equivalent proposition cannot be derived for the parallel transport of a tangent vector. Indeed the parallel transport of a horizontal vector field along a horizontal geodesic may not be horizontal. This will be detailed in the next subsection and constitutes a good example of metric for which computing geodesics is easier than computing parallel transport, although the former is a variational problem and the latter is a linear ODE.

Furthermore, the Riemannian distances d on \mathcal{S}_m^k and d_Σ on Σ_m^k are related by

$$d_\Sigma(\pi(x), \pi(y)) = \inf_{R \in SO(m)} d(x, Ry). \quad (4)$$

The optimal rotation R between any x, y is unique in a subset U of $\mathcal{S}_m^k \times \mathcal{S}_m^k$, which allows to define the *align* map $\omega : U \rightarrow \mathcal{S}_m^k$ that maps (x, y) to Ry . In this case, $d_\Sigma(\pi(x), \pi(y)) = d(x, \omega(x, y))$ and $x\omega(x, y)^T \in \text{Sym}(m)$. It is useful to notice that $w(x, y)$ can be directly computed by a pseudo-singular value decomposition of xy^T [5]. Finally, x and $\omega(x, y)$ are joined by a horizontal geodesic.

2.4 Implementation in `geomstats`

The `geomstats` library [10], available at <https://geomstats.ai>, implements classes of manifolds equipped with Riemannian metrics. It contains an abstract class for quotient metrics, that allows to compute the Riemannian distance, exponential and logarithm maps in the quotient space from the ones in the top space.

In the case of the Kendall shape spaces, the quotient space cannot be seen as a submanifold of some \mathbb{R}^N . Moreover, the projection π and its total derivative $d\pi$ can't be computed explicitly. However, the align map amounts to identifying the shape space with a local horizontal section of the pre-shape space, and thanks to the characteristics of Riemannian submersions mentioned in the previous subsections, all the computations can be done in the pre-shape space.

Recall that \exp , \log , and d denote the operations of the pre-shape space \mathcal{S}_m^k and are given in (1). We obtain from theorem 1 for any $x, y \in \mathcal{S}_m^k$ and $v \in T_x\mathcal{S}_m^k$

$$\begin{aligned} \exp_{\Sigma, [x]}(d_x\pi v) &= \pi(\exp_x(\text{hor}_x(v))), \\ \log_{\Sigma, [x]}([y]) &= d_x\pi \log_x(\omega(x, y)), \\ d_\Sigma([x], [y]) &= d(x, \omega(x, y)). \end{aligned}$$

2.5 Parallel transport in the shape space

As noticed in Remark 1, one cannot use the projection of the parallel transport in the pre-shape space \mathcal{S}_m^k to compute the parallel transport in the shape space Σ_m^k . Indeed [6] proved the following

Proposition 1 (Kim et al. [6]). *Let γ be a horizontal C^1 -curve in \mathcal{S}_m^k and v be a horizontal tangent vector at $\gamma(0)$. Assume that $\text{rank}(\gamma(s)) \geq m - 1$ except for finitely many s . Then the vector field $s \mapsto v(s)$ along γ is horizontal and the projection of $v(s)$ to $T_{[\gamma(s)]}\Sigma_m^k$ is the parallel transport of $d_x\pi v$ along $[\gamma(s)]$ if and only if $s \mapsto v(s)$ is the solution of*

$$\dot{v}(s) = -\text{Tr}(\dot{\gamma}(s)v(s)^T)\gamma(s) + A(s)\gamma(s), \quad v(0) = v \quad (5)$$

where for every s , $A(s) \in \text{Skew}(m)$ is the unique solution to

$$A(s)\gamma(s)\gamma(s)^T + \gamma(s)\gamma(s)^T A(s) = \dot{\gamma}(s)v(s)^T - v(s)\dot{\gamma}(s)^T. \quad (6)$$




Fig. 1. Schematic representation of the pole ladder

Eq. (5) means that the covariant derivative of $s \mapsto v(s)$ along γ must be a vertical vector at all times, defined by the matrix $A(s) \in \text{Skew}(m)$. These equations can be used to compute parallel transport in the shape space. To compute the parallel transport of $d_x \pi w$ along $[\gamma]$, [6] propose the following method: one first chooses a discretization time-step $\delta = \frac{1}{n}$, then repeat for every $s = \frac{i}{n}, i = 0 \dots n$

1. Compute $\gamma(s)$ and $\dot{\gamma}(s)$,
2. Solve the Sylvester equation (6) to compute $A(s)$ and the r.h.s. of (5),
3. Take a discrete Euler step to obtain $\tilde{v}(s + \delta)$
4. Project $\tilde{v}(s + \delta)$ to $T_{\gamma(s)} \mathcal{S}_m^k$ to obtain $\hat{v}(s + \delta)$,
5. Project to the horizontal subspace: $v(s + \delta) \leftarrow \text{hor}(\hat{v}(s + \delta))$
6. $s \leftarrow s + \delta$

We notice that this method can be accelerated by a higher-order integration scheme, such as Runge-Kutta (RK) by directly integrating the system $\dot{v} = f(v, s)$ where f is a smooth map given by (5) and (6). In this case, steps 4. and 5. are not necessary. The precision and complexity of this method is then bound to that of the integration scheme used. As ladder methods rely only on geodesics, which can be computed in closed-form and their convergence properties are well understood [3], we compare this method by integration to the pole ladder. We focus on the case where γ is a horizontal geodesic.

3 The Pole ladder algorithm

3.1 Description

The pole ladder is a modification of the Schild's ladder [11] proposed by [8]. The pole ladder is more precise and cheaper to compute as shown by [3]. It is also exact in symmetric spaces [14]. We thus focus on this method. We describe it here in a Riemannian manifold $(M, \langle \cdot, \cdot \rangle)$.

Consider a geodesic curve $\gamma : t \mapsto \gamma(t) \in M$, with initial conditions $x = \gamma(0) \in M$ and $w = \dot{\gamma}(0) \in T_x M$. In order to compute the parallel transport of $v \in T_x M$ along γ , between times x and $y = \gamma(1)$, the pole ladder consists in first dividing the main geodesic γ in n segments of equal length and computing the geodesic from x with initial velocity $\frac{v}{n^\alpha}$, obtaining $x_v = \text{Exp}_x(\frac{v}{n^\alpha})$. Then for each segment to repeat the following construction (see figure 1):

1. Compute the midpoint of the segment $m = \text{Exp}_x(\frac{w}{2n})$ and the initial speed of the geodesic from m to x_v : $a = \text{Log}_m(x_v)$.
2. Extend this diagonal geodesic by the same length to obtain $z = \text{Exp}_m(-a)$.
3. Repeat steps 2 and 3 with $x_v \leftarrow z$ and $m \leftarrow \text{Exp}_m(\frac{w}{n})$.

After n steps, compute $\tilde{v} = n^\alpha(-1)^n \text{Log}_y(z)$. According to [3], $\alpha \geq 1$ can be chosen, and $\alpha = 2$ is optimal. This vector is an approximation of the parallel transport of v along γ , $\Pi_x^{x_w} v$. This is illustrated on the 2-sphere and in the case $k = m = 3$ on Figure 2.

3.2 Properties

The pole ladder is studied in depth in [3]. We give here the two main properties. Beside its quadratic convergence speed, the main advantage is that this method is available as soon as geodesics are known (even approximately). It is thus applicable very easily in the case of quotient metrics.

Theorem 2. – *The pole ladder converges to the exact parallel transport when the number of steps n goes to infinity, and the error decreases in $O(\frac{1}{n^2})$, with rate related to the covariant derivative of the curvature tensor.*

– *If M is a symmetric space, then the pole ladder is exact in just one step.*

For instance, Σ_2^3 is symmetric, making pole ladder exact in this case.

3.3 Complexity

The main drawback of ladder schemes is that logarithms are required. Indeed the Riemannian logarithm is only locally defined, and often solved by an optimisation problem when geodesics are not known in closed form.

In the case of Kendall shape spaces, it only requires to compute an alignment step, through a singular value decomposition, with usual complexity $O(m^3)$, then the log of the hypersphere, with linear complexity. Moreover, the result of $\log \circ \omega$ is horizontal, so the vertical component needs not be computed for the exponential of step 2, and only the exp of the hypersphere, also with linear complexity, needs to be computed. The vertical projection needs to be computed for the first step. Solving the Sylvester equation through an eigenvalue decomposition also has complexity m^3 . For n rungs of the pole ladder, the overall complexity is thus $O((n+1)(m^3 + 2mk) + mk + m^3) = O(nm^3)$.




Fig. 2. Visualisation of the pole ladder on S^2 (left) and Σ_3^3 (middle and right)

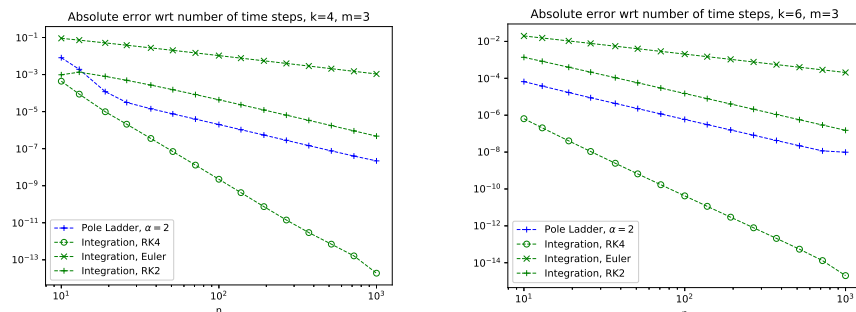


Fig. 3. Error of the parallel transport of v along the geodesic with initial velocity w where v and w are orthonormal.

On the other hand, the method by integration doesn't require logarithms but requires solving a Sylvester equation and a vertical decomposition at every step. The overall complexity is thus $O(2nm^3 + mk)$. Both algorithms are thus comparable in terms of computational cost for a single step.

4 Numerical Simulations and Results

We draw a point x at random in the pre-shape space, along with two orthogonal horizontal unit tangent vectors v, w , and compute the parallel transport of $d_x \pi v$ along the geodesic with initial velocity $d_x \pi w$. We use a number of steps n between 10 and 1000 and the result with $n = 1100$ as the reference value to compute the error made by lower numbers of steps. The results are displayed on Figure 3 for the cases $k = 4, 6$ and $m = 3$ in log-log plots. As expected, the method proposed by [6] converges linearly, while RK schemes of order two and four show significant acceleration. The pole ladder converges with quadratic speed and thus compares with the RK method of order two, although the complexity of the RK method is multiplied by its order.

5 Conclusion and future work

We presented the Kendall shape space and metric, highlighting the properties stemming from its quotient structure. This allows to compute parallel transport with the pole ladder using closed-form solution for the geodesics. This off-the-shelf algorithm can now be used in learning algorithms such as geodesic regression or local non-linear embedding. This will be developed in future works.

6 Acknowledgments

This work was partially funded by the ERC grant Nr. 786854 G-Statistics from the European Research Council under the European Union's Horizon 2020 research and innovation program. It was also supported by the French government

through the 3IA Côte d’Azur Investments ANR-19-P3IA-0002 managed by the National Research Agency.

References

1. Cury, C., Lorenzi, M., Cash, D., Nicholas, J., Routier, A., Rohrer, J., Ourselin, S., Durrleman, S., Modat, M.: Spatio-Temporal Shape Analysis of Cross-Sectional Data for Detection of Early Changes in Neurodegenerative Disease. In: SeSAMI 2016 - First International Workshop Spectral and Shape Analysis in Medical Imaging. LNCS 10126, pp. 63 – 75. Springer (Sep 2016).
2. Dryden, I.L., Mardia, K.V.: Statistical Shape Analysis: With Applications in R. John Wiley & Sons (Sep 2016)
3. Guigui, N., Pennec, X.: Numerical Accuracy of Ladder Schemes for Parallel Transport on Manifolds (Jul 2020).
4. Kendall, D.G.: Shape Manifolds, Procrustean Metrics, and Complex Projective Spaces. *Bulletin of the London Mathematical Society* **16**(2), 81–121 (1984).
5. Kendall, W. S., Le, H.: Statistical Shape Theory. *New Perspectives in Stochastic Geometry*, 348-373 (2009).
6. Kim, K.R., Dryden, I.L., Le, H., Severn, K.E.: Smoothing splines on Riemannian manifolds, with applications to 3D shape space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* (Dec 2020).
7. Le, H., Kendall, D.G.: The Riemannian Structure of Euclidean Shape Spaces: A Novel Environment for Statistics. *The Annals of Statistics* **21**(3), 1225–1271 (1993),
8. Lorenzi, M., Pennec, X.: Efficient Parallel Transport of Deformations in Time Series of Images: From Schild to Pole Ladder. *J Math-Imaging Vis* **50**(1), 5–17 (Sep 2014).
<https://www.overleaf.com/project/601fe6b8ac2c5d67d48ca497>
9. Louis, M., Charlier, B., Jusselin, P., Pal, S., Durrleman, S.: A Fanning Scheme for the Parallel Transport Along Geodesics on Riemannian Manifolds. *SIAM Journal on Numerical Analysis* **56**(4), 2563–2584 (2018).
10. Miolane, N., Guigui, N., Brigant, A.L., Mathe, J., Hou, B., Thanwerdas, Y., Heyder, S., Peltre, O., Koep, N., Zaatiti, H., Hajri, H., Cabanes, Y., Gerald, T., Chauchat, P., Shewmake, C., Brooks, D., Kainz, B., Donnat, C., Holmes, S., Pennec, X.: Geomstats: A Python Package for Riemannian Geometry in Machine Learning. *Journal of Machine Learning Research* **21**(223), 1–9 (Dec 2020)
11. Misner, C.W., Thorne, K.S., Wheeler, J.A.: *Gravitation*. Princeton University Press (1973),
12. Nava-Yazdani, E., Hege, H.C., Sullivan, T.J., von Tycowicz, C.: Geodesic Analysis in Kendall’s Shape Space with Epidemiological Applications. *J Math Imaging Vis* **62**(4), 549–559 (May 2020).
13. O’Neill, B.: *Semi-Riemannian Geometry With Applications to Relativity*. Academic Press (Jul 1983).
14. Pennec, X.: Parallel Transport with Pole Ladder: a Third Order Scheme in Affine Connection Spaces which is Exact in Affine Symmetric Spaces (May 2018),