

STATISTICAL ANALYSIS OF TRAJECTORIES ON RIEMANNIAN MANIFOLDS: BIRD MIGRATION, HURRICANE TRACKING AND VIDEO SURVEILLANCE

BY JINGYONG SU^{*}, SEBASTIAN KURTEK[†], ERIC KLASSEN[‡]
AND ANUJ SRIVASTAVA[‡]

Texas Tech University^{}, Ohio State University[†] and Florida State University[‡]*

We consider the statistical analysis of trajectories on Riemannian manifolds that are observed under arbitrary temporal evolutions. Past methods rely on cross-sectional analysis, with the given temporal registration, and consequently may lose the mean structure and artificially inflate observed variances. We introduce a quantity that provides both a cost function for temporal registration and a proper distance for comparison of trajectories. This distance is used to define statistical summaries, such as sample means and covariances, of synchronized trajectories and “Gaussian-type” models to capture their variability at discrete times. It is invariant to identical time-warpings (or temporal reparameterizations) of trajectories. This is based on a novel mathematical representation of trajectories, termed transported square-root vector field (TSRVF), and the \mathbb{L}^2 norm on the space of TSRVFs. We illustrate this framework using three representative manifolds— \mathbb{S}^2 , $SE(2)$ and shape space of planar contours—involving both simulated and real data. In particular, we demonstrate: (1) improvements in mean structures and significant reductions in cross-sectional variances using real data sets, (2) statistical modeling for capturing variability in aligned trajectories, and (3) evaluating random trajectories under these models. Experimental results concern bird migration, hurricane tracking and video surveillance.

1. Introduction. The need to summarize and model trajectories arises in many statistical procedures. An important issue in this context is that trajectories are often observed at random times. If this temporal variability is not accounted for in the analysis, then the resulting statistical summaries will not be precise. The mean trajectory may not be representative of individual trajectories and the cross-sectional variance will be artificially inflated. This, in turn, will greatly reduce the effectiveness of any subsequent modeling or analysis based on the estimated mean and covariance. As a simple example consider the trajectory on \mathbb{S}^2 shown in the top panel of Figure 1(a). We simulate a set of random, discrete observation times and generate observations of this trajectory at these random times. These simulated trajectories are identical in terms of the points traversed but their evolutions, or parameterizations, are quite different. If we compute the cross-sectional mean

Received January 2013; revised November 2013.

Key words and phrases. Riemannian manifold, time warping, variance reduction, temporal trajectory, rate invariant, parallel transport.

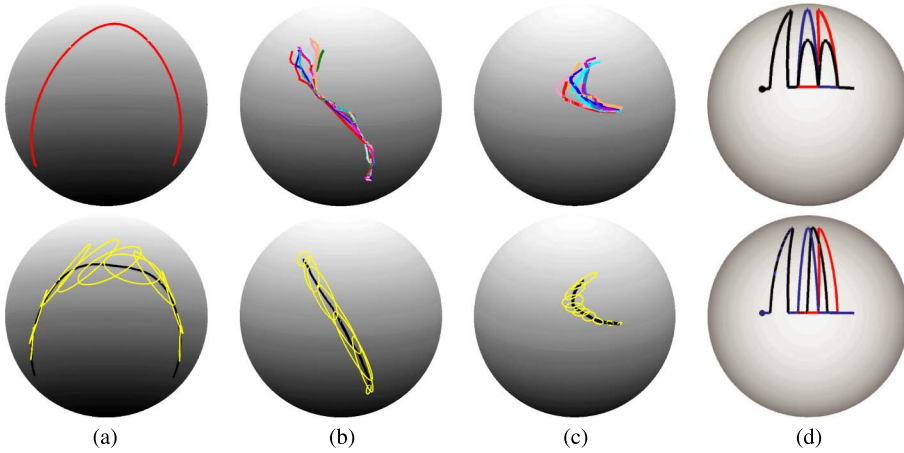


FIG. 1. Summary of trajectories on S^2 : (a) a simulated example; (b) bird migration paths; (c) hurricane tracks; (d) cross-sectional mean of two trajectories without (top) and with (bottom) registration.

and variance, the results are shown in the bottom panel. We draw the sample mean trajectory in black and the sample variance at discrete times using tangential ellipses. Not only is the mean fairly different from the original curve, the variance is purely due to randomness in observation times and is somewhat artificial. If we have observed the trajectory at fixed, synchronized times, this problem would not exist.

To motivate further, consider the phenomenon of bird migration which is the regular seasonal journey undertaken by many species of birds. There are variabilities in migration trajectories, even within the same species, including the variability in their rates of travels. In other words, either birds can travel along different paths or, even if they travel the same path, different birds (or subgroups) may fly at different speed patterns along the path. This results in variability in observation times of migration paths for different birds and artificially inflates the cross-sectional variance in the data. Another issue is that such trajectories are naturally studied as paths on a unit sphere which is a nonlinear manifold. We will study the migration data for Swainson’s Hawk, with some example paths shown in the top panel of Figure 1(b). Swainson’s Hawk inhabits North America mainly in the spring and summer, and winters in South America. It shows perhaps the longest migration of any North American raptor, with durations in excess of two months. Owen and Moore (2008) discovered that Swainson’s Hawk in migratory disposition exhibits reduced immune system functions. Therefore, it becomes important to investigate and summarize such travels. The bottom panel in Figure 1(b) shows the cross-sectional sample mean and variance of the trajectories.

Another motivating application comes from hurricane tracking, where one is interested in studying the shapes of hurricane tracks in certain geographical regions. The statistical summaries and models of hurricane tracks can prove very useful

for monitoring and issuing warnings. Hurricanes potentially evolve at variable dynamical rates and any statistical analysis to these tracks should be invariant of the evolution rates. As in the previous application, the hurricane tracks also are naturally treated as trajectories on a unit sphere. The top panel of Figure 1(c) shows a set of hurricane tracks originating from the Atlantic region. The sample mean and variance of these trajectories are adversely affected by phase variability, as shown in the bottom row of Figure 1(c).

As the last motivating example, consider two synthetic trajectories, drawn in red and blue in the top of Figure 1(d). These two trajectories have the same shape, that is, two bumps each, and a curve representing their mean is also expected to have two bumps. A simple cross-sectional mean, shown by the black trajectory in the same picture, has three bumps. If we solve for the optimal temporal alignment, then such inconsistencies are avoided and the black trajectory in the bottom panel shows the mean obtained using the method proposed in this paper, which accounts for the time-warping variability.

Although there has been progress in the removal of temporal variability, often termed *phase variability*, in Euclidean spaces including Trouvé and Younes (2000), Kneip and Ramsay (2008), Srivastava et al. (2011a) and Tucker, Wu and Srivastava (2013), there has not been any treatment of trajectories on Riemannian manifolds. There are many other applications involving analysis of trajectories on Riemannian manifolds. For example, human activity recognition has attracted tremendous interest in recent years because of its potential in applications such as surveillance, security and human body animation. There are several survey articles, for example, Aggarwal and Cai (1999) and Gavrilu (1999), that provide a detailed review of research in this area. Here each observed activity is represented by a sequence of silhouettes in video frames, each silhouette being an element of the shape space of planar contours. The shape sequences have also been called *shape curves* or curves on shape spaces [Kenobi, Dryden and Le (2010), Le (2003)]. Since activities can be performed at different execution rates, their corresponding shape curves will exhibit distinct evolution rates. Veeraraghavan et al. (2009) accounted for the time-warping variability but their method has some fundamental problems, as explained later. [Briefly, the method is based on equation (2) which is not a proper distance. In fact, it is not even symmetric.] Another motivating application is in pattern analysis of vehicle trajectories at a traffic intersection using surveillance videos, where the instantaneous motion of a vehicle is denoted by the position and orientation on the road. The movements of vehicles typically fall into predictable categories—left turn, right turn, *U* turn, straight line—but the instantaneous speeds can vary depending on the traffic. In order to classify these movements, one has to temporally align the trajectories, thus removing the effects of travel speeds, and then compare them.

Now we describe the problem in mathematical terms. Let $\alpha : [0, 1] \rightarrow M$, where M is a Riemannian manifold, be a differentiable map; it denotes a trajectory on M . We will study such trajectories as elements of an appropriate subset of $M^{[0,1]}$.

Rather than observing a trajectory α directly, say, in the form of time observations $\alpha(t_1), \alpha(t_2), \dots$, we instead observe the time-warped trajectory $\alpha(\gamma(t_1)), \alpha(\gamma(t_2)), \dots$, where $\gamma: [0, 1] \rightarrow [0, 1]$ is an unknown time-warping function (a function with certain constraints described later) that governs the rate of evolution. The mean and variance of $\{\alpha_1(t), \alpha_2(t), \dots, \alpha_n(t)\}$ for any t , where n is the number of observed trajectories, are termed the cross-sectional mean and variance at that t . If we use the observed samples $\{\alpha_i(\gamma_i(t)), i = 1, 2, \dots, n\}$ for analysis, the cross-sectional variance is inflated due to random γ_i . Our hypothesis is that this problem can be mitigated by temporally registering the trajectories. Thus, we are interested in the following four tasks:

1. *Temporal registration*: This is a process of establishing a one-to-one correspondence between points along multiple trajectories. That is, given any n trajectories, say, $\alpha_1, \alpha_2, \dots, \alpha_n$, we are interested in finding functions $\gamma_1, \gamma_2, \dots, \gamma_n$ such that the points $\alpha_i(\gamma_i(t))$ are matched optimally for all t .
2. *Metric-based comparisons*: We want to develop a metric that is invariant to different evolution rates of trajectories. Specifically, we want to define a distance $d(\cdot, \cdot)$ such that for arbitrary evolution functions γ_1, γ_2 and arbitrary trajectories α_1 and α_2 , we have $d(\alpha_1, \alpha_2) = d(\alpha_1 \circ \gamma_1, \alpha_2 \circ \gamma_2)$.
3. *Statistical summary*: The main use of this metric will be in defining and computing a (Karcher) mean trajectory $\mu(t)$ and a cross-sectional variance function $\hat{\rho}(t)$, associated with any given set of trajectories. The reason for performing registration is to reduce the cross-sectional variance that is artificially introduced in the data due to random observation times. The reduction in variance is quantified using $\hat{\rho}$.
4. *Statistical modeling and evaluation*: We will use the estimated mean and covariance of registered trajectories to define a ‘‘Gaussian-type’’ model on random trajectories. This model will then be used to evaluate p -values associated with new trajectories. Here the p -value implies the proportion of trajectories with smaller density than the current trajectory under the given model.

For performing comparison and summarization of trajectories, we need a metric and, at first, we consider a more conventional solution. Since M is a Riemannian manifold, we have a natural distance d_m between points on M . Using d_m , one can compare any two trajectories: $\alpha_1, \alpha_2: [0, 1] \rightarrow M$, as

$$(1) \quad d_x(\alpha_1, \alpha_2) = \int_0^1 d_m(\alpha_1(t), \alpha_2(t)) dt.$$

Although this quantity represents a natural extension of d_m from M to $M^{[0,1]}$, it suffers from the problem that $d_x(\alpha_1, \alpha_2) \neq d_x(\alpha_1 \circ \gamma_1, \alpha_2 \circ \gamma_2)$ in general. It is not preserved even when the same γ is applied to both trajectories, that is, $d_x(\alpha_1, \alpha_2) \neq d_x(\alpha_1 \circ \gamma, \alpha_2 \circ \gamma)$ in general. If we have equality in the last case, for all γ , then we can develop a fully invariant distance and use it to properly register

trajectories, as described later. So, the failure to have this equality is a key issue that forces us to look for other solutions in situations where trajectories are observed at random temporal evolutions. When a trajectory α is observed as $\alpha \circ \gamma$, for an arbitrary temporal re-parameterization γ , we call this perturbation *compositional noise*. In these terms, d_x is not useful in comparing trajectories observed under compositional noise.

Our goal is to take time-warping into account, derive a warping-invariant metric, and generate statistical summaries (sample mean, covariance, etc.) for trajectories on a set M . The fact that M is a Riemannian manifold presents a formidable challenge in developing a comprehensive framework. But this is not the only challenge. To clarify, how has this registration and analysis problem been handled for trajectories in Euclidean spaces? In case $M = \mathbb{R}$, that is, if one is interested in registration and modeling of real-valued functions under random time-warpings, the problem has been studied by many authors, including [Srivastava et al. \(2011a\)](#), [Liu and Müller \(2004\)](#), [Kneip and Ramsay \(2008\)](#) and [Tucker, Wu and Srivastava \(2013\)](#). In case $M = \mathbb{R}^2$, where the problem involves registration and shape analysis of planar curves, the solution is discussed in [Michor and Mumford \(2007\)](#), [Younes et al. \(2008\)](#), [Shah \(2008\)](#) and [Sundaramoorthi et al. \(2011\)](#). [Srivastava et al. \(2011b\)](#) proposed a solution that applies to curves in arbitrary \mathbb{R}^n . One can also draw solutions from problems in image registration where 2D and 3D images are registered to each other using a spatial warping instead of a temporal warping [see, e.g., LDDMM technique, [Beg et al. \(2005\)](#)]. A majority of the existing methods in Euclidean spaces formulate an objective function of the type

$$\min_{\gamma} \left(\int_0^1 |\alpha_1(t) - \alpha_2(\gamma(t))|^2 dt + \lambda \mathcal{R}(\gamma) \right),$$

where $|\cdot|$ is the Euclidean norm, \mathcal{R} is a regularization term on the warping function γ , and $\lambda > 0$ is a constant. In the case of a Riemannian manifold, one can modify the first term to obtain

$$(2) \quad \min_{\gamma} \left(\int_0^1 d_m(\alpha_1(t), \alpha_2(\gamma(t)))^2 dt + \lambda \mathcal{R}(\gamma) \right),$$

where $d_m(\cdot, \cdot)$ is the geodesic distance on the manifold. The main problem with this procedure is that (a) it is not symmetric, that is, the registration of α_1 to α_2 is not the same as that of α_2 to α_1 , as pointed out by [Christensen and Johnson \(2001\)](#), among others, and (b) the minimum value is not a proper distance, so it cannot be used to compare trajectories. This sums up the fundamental dilemma in trajectory analysis—equation (1) provides a metric between trajectories but does not perform registration, while equation (2) performs registration but is not a metric.

Another potential approach is to map trajectories onto a vector space, for example, the tangent space at a point, using the inverse exponential map, and then compare the mapped trajectories using the Euclidean solutions in the vector space.

While this idea is feasible, the results may not be consistent since the inverse exponential map is a local and highly nonlinear operator. For example, on a sphere, under stereographic map two points near the south pole will map to two distant points in the tangent space at the north pole, and their distance will be highly distorted. In contrast, the solution proposed here transports vector fields associated with trajectories, rather than trajectories themselves, into a standard tangent space and this provides a more stable alternative.

We would like an objective function for alignment that (a) is a proper distance, that is, it is symmetric, positive definite and satisfies the triangle inequality, (b) is invariant to simultaneous warping of two trajectories by the same warping function, and (c) leads to minimal cross-sectional variance for sample trajectories. For real-valued functions, a Riemannian framework has already been presented in Kurtek, Wu and Srivastava (2011) and Srivastava et al. (2011a), but to our knowledge this framework has not been generalized to manifolds.

In this paper we develop a framework for automated registration of multiple trajectories and obtain improvements in statistical summaries of time-warped trajectories on Riemannian manifolds. This framework is based on a novel mathematical representation called the transported square-root vector field (TSRVF) and the \mathbb{L}^2 norm between TSRVFs. The setup satisfies the invariance property mentioned earlier, that is, an identical time-warping of TSRVFs representing two trajectories preserves the \mathbb{L}^2 norm of their difference and, therefore, this difference is used to define a warping-invariant distance between trajectories. The resulting distance is found useful in registration, comparison and summarization of trajectories on manifolds. To illustrate these ideas, we take three manifolds, \mathbb{S}^2 , $\text{SE}(2)$ and the shape space of planar closed curves, and provide simulated and real examples. Our paper can also be viewed as an extension, albeit not a trivial one, of the work of Kurtek, Wu and Srivastava (2011) and Srivastava et al. (2011a) from $M = \mathbb{R}$ to Riemannian manifolds.

The paper is organized as follows. In Section 2 we introduce a general mathematical framework for analyzing trajectories on Riemannian manifolds and demonstrate the use of this framework in registration, comparison, summarization, modeling and evaluation. We also provide algorithms for performing these tasks. In Section 3 we specialize this framework to \mathbb{S}^2 and consider two applications. In Section 4 we apply it to pattern analysis of vehicle trajectories on $\text{SE}(2)$. In Section 5 we provide details for time-warping invariant analysis of trajectories on the shape space of planar closed curves, with applications to activity recognition.

2. Mathematical framework. Let α denote a smooth trajectory on a Riemannian manifold M endowed with a Riemannian metric $\langle \cdot, \cdot \rangle$. Let \mathcal{M} denote the set of all such trajectories: $\mathcal{M} = \{\alpha : [0, 1] \rightarrow M \mid \alpha \text{ is smooth}\}$. Also, define Γ to be the set of all orientation preserving diffeomorphisms of $[0, 1]$: $\Gamma = \{\gamma : [0, 1] \rightarrow [0, 1] \mid \gamma(0) = 0, \gamma(1) = 1, \gamma \text{ is a diffeomorphism}\}$. Note that Γ forms a group under the composition operation. If α is a trajectory on M , then $\alpha \circ \gamma$ is a trajectory

that follows the same sequence of points as α but at the evolution rate governed by γ . More technically, the group Γ acts on \mathcal{M} according to $(\alpha, \gamma) = \alpha \circ \gamma$.

Given two smooth trajectories $\alpha_1, \alpha_2 \in \mathcal{M}$, we want to register points along the trajectories and compute a time-warping invariant distance between them. As mentioned earlier, the quantity given in equation (2) would be a natural choice for this purpose, but it fails for several reasons, including the fact that it is not symmetric. Fundamentally, this and other quantities used in previous literature are not appropriate for solving the registration problem because they are not measuring registration in the first place. To highlight this issue, take the registration of points between the pair (α_1, α_2) and the pair $(\alpha_1 \circ \gamma, \alpha_2 \circ \gamma)$, for any $\gamma \in \Gamma$. It can be seen that the pairs (α_1, α_2) and $(\alpha_1 \circ \gamma, \alpha_2 \circ \gamma)$ have exactly the same registration of points. In fact, any identical time-warping of two trajectories does not change the registration of points between them. But the quantities given in equations (1) and (2) provide different values for these pairs, despite the same registration. Hence, they are not good measures of registration. We emphasize that the invariance under identical time-warping is a key property needed in the desired framework.

We introduce a new representation of trajectories that will be used to compare and register them. We assume that for any two points $p, q \in M$, we have an expression for parallel transporting any vector $v \in T_p(M)$ along the shortest geodesic from p to q , denoted by $(v)_{p \rightarrow q}$. As long as p and q do not fall in the cut loci of each other, the geodesic between them is unique and the parallel transport is well defined. The measure of the set of cut locus on the manifolds of our interest is typically zero. So, the practical implications of this limitation are negligible. Let c be a point in M that we designate as a reference point. We assume that none of the observed trajectories pass through the cut locus of c to avoid the problem mentioned above.

DEFINITION 1. For any smooth trajectory $\alpha \in \mathcal{M}$, the transported square-root vector field (TSRVF) is a parallel transport of a scaled velocity vector field of α to a reference point $c \in M$ according to

$$h_\alpha(t) = \frac{\dot{\alpha}(t)_{\alpha(t) \rightarrow c}}{\sqrt{|\dot{\alpha}(t)|}} \in T_c(M),$$

where $|\cdot|$ denotes the norm related to the Riemannian metric on M .

Since α is smooth, so is the vector field h_α . Let $\mathcal{H} \subset T_c(M)^{[0,1]}$ be the set of smooth curves in $T_c(M)$ obtained as TSRVFs of trajectories in M , $\mathcal{H} = \{h_\alpha | \alpha \in \mathcal{M}\}$. If $M = \mathbb{R}^n$ with the Euclidean metric, then h is exactly the square-root velocity function defined in [Srivastava et al. \(2011b\)](#).

The choice of the reference point c used in Definition 1 is important and can affect the results. The choice typically depends on the application, the data and the manifold under study. In case all the trajectories pass through a point or pass close

to a point, then that point is a natural candidate for c . This would be true, for example, in the case of hurricane tracks, if we are focused on all hurricanes starting from the same region. Another remark is that instead of parallel transporting of scaled velocity vectors along geodesics, one can transport them along trajectories themselves, as was done by Jupp and Kent (1987), but that requires c to be a common point of all trajectories. While the choice of c can, in principle, affect distances, our experiments suggest that the results of registration, distance-based clustering and classification are quite stable with respect to this choice. An example is presented later in Figure 2.

We represent a trajectory $\alpha \in \mathcal{M}$ with the pair $(\alpha(0), h_\alpha) \in M \times \mathcal{H}$. Given this representation, we can reconstruct the path, an element of \mathcal{M} , as follows. For any time t , let V_t be a time-varying tangent vector-field on M obtained by parallel transporting $h_\alpha(t)$ over M [except for the cut locus of $\alpha(t)$], that is, for any $p \in M$, $V_p(t) = (h_\alpha(t))_{c \rightarrow p}$. Then, define an integral curve β such that $\dot{\beta}(t) = |V_{\beta(t)}(t)| V_{\beta(t)}(t)$ with the starting point $\beta(0) = \alpha(0) \in M$. This resulting curve β will be exactly the same as the original curve α .

The starting points of different curves can be compared using the Riemannian distance d_m on M . However, these points do not play an important role in the alignment of trajectories since they are already assumed to be matched to each other. Therefore, the main focus of analysis, in terms of alignment and comparison, is on TSRVFs. Since a TSRVF is a path in $T_c(M)$, one can use the \mathbb{L}^2 norm to compare such paths.

DEFINITION 2. Let α_1 and α_2 be two smooth trajectories on M and let h_{α_1} and h_{α_2} be the corresponding TSRVFs. The distance between them is

$$d_h(h_{\alpha_1}, h_{\alpha_2}) = \left(\int_0^1 |h_{\alpha_1}(t) - h_{\alpha_2}(t)|^2 dt \right)^{1/2}.$$

The distance d_h , being the standard \mathbb{L}^2 norm, satisfies symmetry, positive definiteness and triangle inequality. Also, due to the invertibility of the mapping from \mathcal{M} to $M \times \mathcal{H}$, one can use d_h (along with d_m) to define a distance on \mathcal{M} . The main motivation of this setup—TSRVF representation and \mathbb{L}^2 norm—comes from the following fact. If a trajectory α is warped by γ , to result in $\alpha \circ \gamma$, the TSRVF of $\alpha \circ \gamma$ is given by

$$\begin{aligned} h_{\alpha \circ \gamma}(t) &= \frac{(\dot{\alpha}(\gamma(t))\dot{\gamma}(t))_{\alpha(\gamma(t)) \rightarrow c}}{\sqrt{|\dot{\alpha}(\gamma(t))\dot{\gamma}(t)|}} = \frac{(\dot{\alpha}(\gamma(t)))_{\alpha(\gamma(t)) \rightarrow c} \sqrt{\dot{\gamma}(t)}}{\sqrt{|\dot{\alpha}(\gamma(t))|}} \\ &= h_\alpha(\gamma(t)) \sqrt{\dot{\gamma}(t)}, \end{aligned}$$

which is also denoted as $(h_\alpha, \gamma)(t)$. We will often write (h_α, γ) to denote $h_{\alpha \circ \gamma}$. As stated earlier, we need a distance for registration that is invariant to identical time-warpings of trajectories. Next, we show that d_h satisfies this property.

THEOREM 1. *For any $\alpha_1, \alpha_2 \in \mathcal{M}$ and $\gamma \in \Gamma$, the distance d_h satisfies $d_h(h_{\alpha_1 \circ \gamma}, h_{\alpha_2 \circ \gamma}) = d_h(h_{\alpha_1}, h_{\alpha_2})$. In geometric terms, this implies that the action of Γ on \mathcal{H} under the \mathbb{L}^2 metric is by isometries.*

The proof is given below:

$$\begin{aligned} d_h(h_{\alpha_1 \circ \gamma}, h_{\alpha_2 \circ \gamma}) &= \left(\int_0^1 |h_{\alpha_1}(\gamma(t))\sqrt{\dot{\gamma}(t)} - h_{\alpha_2}(\gamma(t))\sqrt{\dot{\gamma}(t)}|^2 dt \right)^{1/2} \\ &= \left(\int_0^1 |h_{\alpha_1}(s) - h_{\alpha_2}(s)|^2 ds \right)^{1/2} = d_h(h_{\alpha_1}, h_{\alpha_2}), \end{aligned}$$

where $s = \gamma(t)$.

Next we define a quantity that can be used as a distance between trajectories while being invariant to their temporal variability. To set up this definition, we first introduce an equivalence relation between trajectories. For any two trajectories α_1 and α_2 , we define them to be equivalent, $\alpha_1 \sim \alpha_2$, when:

1. $\alpha_1(0) = \alpha_2(0)$, and
2. there exists a sequence $\{\gamma_k\} \in \Gamma$ such that $\lim_{k \rightarrow \infty} h_{(\alpha_1 \circ \gamma_k)} = h_{\alpha_2}$ under the \mathbb{L}^2 metric.

In other words, any two trajectories are equivalent if they have the same starting point and the TSRVF of one can be time-warped into the TSRVF of the other using a sequence of warpings. It can be easily checked that \sim forms an equivalence relation on \mathcal{H} (and, correspondingly, \mathcal{M}).

Since we want our distance to be invariant to time-warpings of trajectories, we wish to compare trajectories by comparing their equivalence classes. Thus, our next step is to inherit the distance d_h to the set of such equivalence classes. Toward this goal, we introduce the set $\tilde{\Gamma}$ as the set of all nondecreasing, absolutely continuous functions $\gamma : [0, 1] \rightarrow [0, 1]$ such that $\gamma(0) = 0$ and $\gamma(1) = 1$. This set $\tilde{\Gamma}$ is a semigroup with the composition operation (it is not a group because the elements do not have inverses). The group Γ is a subset of $\tilde{\Gamma}$. The elements of $\tilde{\Gamma}$ warp the time axis of trajectories in M in the same way as elements of Γ , except they allow certain singularities. For a TSRVF $h_\alpha \in \mathcal{H}$, its equivalence class, or *orbit* under $\tilde{\Gamma}$, is given by $[h_\alpha] = \{(h_\alpha, \gamma) | h_\alpha \in \mathcal{H}, \gamma \in \tilde{\Gamma}\}$.

It can be shown that the orbits under $\tilde{\Gamma}$ are exactly the same as the closures of the orbits of Γ , defined as $[h_\alpha]_0 = \{(h_\alpha, \gamma) | \gamma \in \Gamma\}$, as long as α has nonvanishing derivatives almost everywhere. (The last condition is not restrictive since we can always re-parameterize α by the arc-length.) The closure is with respect to the \mathbb{L}^2 metric on \mathcal{H} . Please refer to [Robinson \(2012\)](#) for a detailed description of a similar construction for trajectories in \mathbb{R} .

Now we define the quantity that will serve both as the cost function for registration and distance for comparison. This quantity is essentially d_h measured between equivalence classes.

DEFINITION 3. The distance d_s on \mathcal{H}/\sim (or \mathcal{M}/\sim) is the shortest d_h distance between equivalence classes in \mathcal{H} , given as

$$\begin{aligned}
 & d_s([h_{\alpha_1}], [h_{\alpha_2}]) \\
 (3) \quad &= \inf_{\gamma_1, \gamma_2 \in \tilde{\Gamma}} d_h((h_{\alpha_1}, \gamma_1), (h_{\alpha_2}, \gamma_2)) \\
 &= \inf_{\gamma_1, \gamma_2 \in \tilde{\Gamma}} \left(\int_0^1 |h_{\alpha_1}(\gamma_1(t))\sqrt{\dot{\gamma}_1(t)} - h_{\alpha_2}(\gamma_2(t))\sqrt{\dot{\gamma}_2(t)}|^2 dt \right)^{1/2}.
 \end{aligned}$$

THEOREM 2. The distance d_s is a proper distance on \mathcal{H}/\sim .

PROOF. The symmetry of d_s comes directly from the symmetry of d_h . For positive definiteness, we need to show that $d_s([h_{\alpha_1}], [h_{\alpha_2}]) = 0 \Rightarrow [h_{\alpha_1}] = [h_{\alpha_2}]$. Suppose that $d_s([h_{\alpha_1}], [h_{\alpha_2}]) = 0$, by definition, it then follows immediately that for all $\varepsilon > 0$, there exists a $\gamma \in \Gamma$ such that $d_h(h_{\alpha_1}, (h_{\alpha_2}, \gamma)) < \varepsilon$. From this, it follows that h_{α_1} is in the orbit h_{α_2} . Since we are assuming that orbits are closed, it follows that $h_{\alpha_1} \in [h_{\alpha_2}]$, so $[h_{\alpha_1}] = [h_{\alpha_2}]$.

To establish the triangle inequality, we need to prove

$$d_s([h_{\alpha_1}], [h_{\alpha_3}]) \leq d_s([h_{\alpha_1}], [h_{\alpha_2}]) + d_s([h_{\alpha_2}], [h_{\alpha_3}])$$

for any $h_{\alpha_1}, h_{\alpha_2}, h_{\alpha_3} \in \mathcal{H}$. For a contradiction, suppose

$$d_s([h_{\alpha_1}], [h_{\alpha_3}]) > d_s([h_{\alpha_1}], [h_{\alpha_2}]) + \tilde{d}([h_{\alpha_2}], [h_{\alpha_3}]).$$

Let

$$\varepsilon = \frac{1}{3}(d_s([h_{\alpha_1}], [h_{\alpha_3}]) - d_s([h_{\alpha_1}], [h_{\alpha_2}]) - d_s([h_{\alpha_2}], [h_{\alpha_3}])).$$

By our supposition, $\varepsilon > 0$. From the definition of ε , it follows that

$$d_s([h_{\alpha_1}], [h_{\alpha_3}]) = d_s([h_{\alpha_1}], [h_{\alpha_2}]) + d_s([h_{\alpha_2}], [h_{\alpha_3}]) + 3\varepsilon.$$

By the definition of d_s , we can choose $\gamma_1, \gamma_2 \in \Gamma$, such that

$$d_h((h_{\alpha_1}, \gamma_1), h_{\alpha_2}) \leq d_s([h_{\alpha_1}], [h_{\alpha_2}]) + \varepsilon$$

and

$$d_h(h_{\alpha_2}, (h_{\alpha_3}, \gamma_2)) \leq d_s([h_{\alpha_2}], [h_{\alpha_3}]) + \varepsilon.$$

Now, by the triangle inequality for d_h , we know that

$$\begin{aligned}
 d_h((h_{\alpha_1}, \gamma_1), (h_{\alpha_3}, \gamma_2)) &\leq d_h((h_{\alpha_1}, \gamma_1), h_{\alpha_2}) + d_h(h_{\alpha_2}, (h_{\alpha_3}, \gamma_2)) \\
 &\leq d_s([h_{\alpha_1}], [h_{\alpha_2}]) + d_s([h_{\alpha_2}], [h_{\alpha_3}]) + 2\varepsilon.
 \end{aligned}$$

It follows that

$$d_s([h_{\alpha_1}], [h_{\alpha_3}]) \leq d_s([h_{\alpha_1}], [h_{\alpha_2}]) + d_s([h_{\alpha_2}], [h_{\alpha_3}]) + 2\varepsilon.$$

But this contradicts the fact that

$$d_s([h_{\alpha_1}], [h_{\alpha_3}]) = d_s([h_{\alpha_1}], [h_{\alpha_1}]) + d_s([h_{\alpha_2}], [h_{\alpha_3}]) + 3\varepsilon.$$

Hence, our supposition that

$$d_s([h_{\alpha_1}], [h_{\alpha_3}]) > d_s([h_{\alpha_1}], [h_{\alpha_2}]) + d_s([h_{\alpha_2}], [h_{\alpha_3}])$$

must be false. The triangle inequality follows. \square

Now, since Γ is dense in $\tilde{\Gamma}$, for any $\delta > 0$, there exists a γ^* such that

$$(4) \quad |d_h(h_{\alpha_1}, h_{\alpha_2 \circ \gamma^*}) - d_s([h_{\alpha_1}], [h_{\alpha_2}])| < \delta.$$

This γ^* may not be unique but any such γ^* is sufficient for our purpose. Furthermore, since $\gamma^* \in \Gamma$, it has an inverse that can be used in further analysis. The minimization over Γ in equation (4) is performed in practice using the dynamic programming (DP) algorithm [Bertsekas (2007)]. Here one samples the interval $[0, 1]$ using T discrete points and then restricts to only piecewise linear γ 's that pass through that $T \times T$ grid. The search for the optimal trajectory on this grid is accomplished in $O(T^2)$ steps.

2.1. Metric-based comparison of trajectories. Our goal of warping-invariant comparison of trajectories is achieved using d_s . For any $\gamma_1, \gamma_2 \in \Gamma$ and $\alpha_1, \alpha_2 \in \mathcal{M}$, we have

$$[h_{\alpha_1 \circ \gamma_1}] = [h_{\alpha_1}], \quad [h_{\alpha_2 \circ \gamma_2}] = [h_{\alpha_2}]$$

and, therefore, we get $d_s([h_{\alpha_1 \circ \gamma_1}], [h_{\alpha_2 \circ \gamma_2}]) = d_s([h_{\alpha_1}], [h_{\alpha_2}])$. Examples of this metric are presented later.

2.2. Pairwise temporal registration of trajectories. The next goal is to perform registration of points along trajectories. Let our approximation to the optimal warping be as defined in equation (4). This allows for the registration between α_1 and α_2 , in that the point $\alpha_1(t)$ on the first trajectory is optimally matched to the point $\alpha_2(\gamma^*(t))$ on the second trajectory.

If we compare equation (3) with equation (2), we see the advantages of the proposed framework. Both equations present a registration problem between α_1 and α_2 , but only the minimum value resulting from equation (3) is a proper distance. Also, in equation (2) we have two separate terms for matching and regularization, with an arbitrary weight λ , but in equation (3) the two terms have been merged into a single natural form. Recall that the change in TSRVF h due to the time-warping of α by γ is given by $(h, \gamma) = (h \circ \gamma)\sqrt{\dot{\gamma}}$, and the distance d_s is based on these warped TSRVFs. The term $\sqrt{\dot{\gamma}}$ provides an intrinsic regularization on γ in the matching process. It provides an elastic penalty against excessive warping since $\dot{\gamma}$ becomes large at those places. Lastly, the optimal registration in equation (3) remains the same if we change the order of the input functions. That is, the registration process is inverse consistent.

2.3. *Summarization and registration of multiple trajectories.* An additional advantage of this framework is that one can compute an average of several trajectories and use it as a *template* for future classification. Furthermore, this template can be used for registering multiple trajectories. We use the notion of the Karcher mean to define and compute average trajectories. Given a set of sample trajectories $\alpha_1, \dots, \alpha_n$ on M , we represent them using the corresponding pairs $(\alpha_1(0), h_{\alpha_1}), (\alpha_2(0), h_{\alpha_2}), \dots, (\alpha_n(0), h_{\alpha_n})$. We compute the Karcher means of each component in their respective spaces: (1) the Karcher mean of $\alpha_i(0)$ is computed with respect to d_m in M , and (2) the Karcher mean of h_{α_i} with respect to d_s in \mathcal{H}/\sim . The latter Karcher mean is defined as

$$h_\mu = \operatorname{argmin}_{[h_\alpha] \in \mathcal{H}/\sim} \sum_{i=1}^n d_s([h_\alpha], [h_{\alpha_i}])^2.$$

Note that $[h_\mu]$ is an equivalence class of trajectories and one can select any element of this mean class to help in the alignment of multiple trajectories. The standard algorithm to compute the Karcher mean proposed by [Le and Kume \(2000\)](#) is adapted to this problem as follows:

ALGORITHM 1 (Karcher mean of multiple trajectories). *Compute the Karcher mean of $\{\alpha_i(0)\}$ and set it to be $\mu(0)$.*

1. *Initialization step: select μ to be one of the original trajectories and compute its TSRVF h_μ .*
2. *Align each $h_{\alpha_i}, i = 1, \dots, n$, to h_μ according to equation (4). That is, solve for γ_i^* using the DP algorithm and set $\tilde{\alpha}_i = \alpha_i \circ \gamma_i^*$.*
3. *Compute TSRVFs of the warped trajectories, $h_{\tilde{\alpha}_i}, i = 1, 2, \dots, n$, and update h_μ as a curve in $T_c(M)$ according to $h_\mu(t) = \frac{1}{n} \sum_{i=1}^n h_{\tilde{\alpha}_i}(t)$.*
4. *Define μ to be the integral curve associated with a time-varying vector field on M generated using h_μ , that is, $\frac{d\mu(t)}{dt} = |h_\mu(t)|(h_\mu)(t)_{c \rightarrow \mu(t)}$, and the initial condition $\mu(0)$.*
5. *Compute $E = \sum_{i=1}^n d_s([h_\mu], [h_{\alpha_i}])^2 = \sum_{i=1}^n d_h(h_\mu, h_{\tilde{\alpha}_i})^2$ and check for convergence. If not converged, return to step 2.*

It can be shown that the cost function decreases iteratively and, as zero is a natural lower bound, $\sum_{i=1}^n d_s([h_\mu], [h_{\alpha_i}])^2$ will always converge. This algorithm provides two sets of outputs: an average trajectory denoted by the final μ and the set of aligned trajectories $\tilde{\alpha}_i$. Therefore, this solves the problem of aligning multiple trajectories too.

For computing and analyzing the second and higher moments of a sample trajectory, the tangent space $T_{\mu(t)}(M)$, for $t \in [0, 1]$, is used. This is convenient because it is a vector space and one can apply more traditional methods here. First, for each aligned trajectory $\tilde{\alpha}_i(t)$ at time t , the vector $v_i(t) \in T_{\mu(t)}(M)$ is computed

such that a geodesic that goes from $\mu(t)$ to $\tilde{\alpha}_i(t)$ in unit time has the initial velocity $v_i(t)$. This is also called the *shooting vector* from $\mu(t)$ to $\tilde{\alpha}_i(t)$. Let $\hat{K}(t)$ be the sample covariance matrix of all shooting vectors from $\mu(t)$'s to $\tilde{\alpha}_i(t)$'s. The sample Karcher covariance at time t is given by $\hat{K}(t) = \frac{1}{n-1} \sum_{i=1}^n v_i(t)v_i(t)^T$, with the trace $\hat{\rho}(t) = \text{trace}(\hat{K}(t))$. This $\hat{\rho}(t)$ represents a quantification of the cross-sectional variance, as a function of t , and can be used to study the level of alignment of trajectories. Also, for capturing the essential variability in the data, one can perform Principal Component Analysis (PCA) of the shooting vectors. The basic idea is to compute the Singular Value Decomposition (SVD) $\hat{K}(t) = U(t)\Sigma(t)U^T(t)$, where $U(t)$ is an orthogonal matrix and $\Sigma(t)$ is the diagonal matrix of singular values. Assuming that the entries along the diagonal in $\Sigma(t)$ are organized in a nonincreasing order, the functions $U_1(t), U_2(t), \dots$ represent the dominant directions of variability in the data.

2.4. Modeling and evaluation of trajectories. An important use of means and covariances of trajectories is in devising probability models for capturing the observed statistical variability, and for using these models in evaluating p -values of future observations. By p -values we mean the proportion of random trajectories that will have lower probability density under a given model when compared to the test trajectory. Several models are possible in this situation, but since our main focus is on temporal registration of trajectories, we will choose a simple model to demonstrate our ideas. After the registration, we treat a trajectory α as a discrete-time process, composed of m points as $\{\alpha(t_1), \alpha(t_2), \dots, \alpha(t_m)\}$, for a fixed partition $\{0 = t_1, t_2, \dots, t_m = 1\}$ of $[0, 1]$. Given the mean and the covariance at each t_j , we model the points $\alpha(t_j) \in M, j = 1, 2, \dots, m$ independently, and obtain the joint density by taking the product. The difficulty in this step comes from the fact that M is a nonlinear manifold but one can use the tangent space $T_{\mu(t_j)}(M)$, instead, to impose a probability model since this is a vector space. We impose a multivariate normal density on the tangent vector $v(t_j) = \exp_{\mu(t_j)}^{-1}(\alpha(t_j))$, with mean zero and variance given by $\hat{K}(t_j)$ (as defined above). It is analogous to the model of additive white Gaussian noise when $M = \mathbb{R}$. Then, for any trajectory α , one can compute a joint probability of the full trajectory as $P(\alpha) = \prod_{j=1}^m f(\alpha(t_j)) \equiv \prod_{j=1}^m N(v(t_j); 0, \hat{K}(t_j))$. This model is potentially useful for many situations: (1) It can be used to simulate new trajectories via random sampling. Given $\{(\mu(t_j), \hat{K}(t_j)) | t \in [0, 1]\}$, we can simulate the tangent vectors and compute the corresponding trajectory points $\alpha(t_j)$, for the desired t_j . (2) Given a trajectory, we evaluate its p -value under the imposed model. This measures how likely is the occurrence of the trajectory by chance assuming the null hypothesis H_0 , where H_0 represents that imposed model.

Since we are interested in studying the effects of temporal registration, we demonstrate these ideas with the following experiment. We compute p -values of

trajectories using the parametric bootstrap under two situations: without registration and with registration. In each situation, we first take a set of trajectories as the training set and estimate the mean and covariance at each discrete time, and then impose a ‘‘Gaussian-type’’ model on the tangent spaces of M at the mean values at those times. This becomes the imposed model H_0 . The evaluation of p -values requires Monte Carlo sampling. We generate a large number, say, $N = 10,000$, of trajectories from the model, denoted as $X_i, i = 1, 2, \dots, N$. Then we compute the proportion that are less likely than our test trajectory and denote it as $p(\alpha) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{[P(X_i) < P(\alpha)]}$.

In the following sections we consider three examples of M and present experimental results to validate our framework.

3. Trajectories on \mathbb{S}^2 . Statistical methods for unit vectors in three-dimensional space have been studied extensively in directional statistics [Mardia and Jupp (2000)]. In the landmark-based shape analysis of objects, including Dryden and Mardia (1998), Jupp and Kent (1987) and Kume, Dryden and Le (2007), where 2D objects are represented by configurations of salient points or landmarks, the set of all such configurations after removing translation and scale is a real sphere \mathbb{S}^{2n-3} (for configurations with n landmarks). To illustrate this framework, in a simple setting, we start with $M = \mathbb{S}^2$, with the standard Euclidean Riemannian metric. For any two points $p, q \in \mathbb{S}^2$ ($p \neq -q$) and a tangent vector $v \in T_p(\mathbb{S}^2)$, the parallel transport $(v)_{p \rightarrow q}$ along the shortest geodesic (i.e., great circle) from p to q is given by $v - \frac{2\langle v, q \rangle}{|p+q|^2} (p + q)$.

Registration of trajectories: As mentioned earlier, for any two trajectories on \mathbb{S}^2 , we use their TSRVFs and DP algorithm in equation (4) to find the optimal registration between them. In Figure 2 we show one example of registering such trajectories. The parameterization of trajectories is displayed using colors. In the top row, the left column shows the trajectories α_1 and α_2 , the middle column shows α_1 and $\alpha_2 \circ \gamma^*$ and the right column shows γ^* using $c = [0, 0, 1]$. The correspondences between the two trajectories are depicted by black lines connecting points along them. Due to the optimization of γ in equation (4), the d_h value between them reduces from 1.67 to 0.36 and the correspondences become more natural after the alignment. We also consider different choices of c ($c = [0, 0, -1], [-1, 0, 0], [0, 1, 0]$). In all cases the registration results are very close, as shown in the bottom row.

In the following, we consider two specific applications, bird migration and hurricane tracks, and show how the cross-sectional variance of the mean trajectories is reduced by registration. For both applications, we use the mean of the starting points of the trajectories as the reference point c in Definition 1.

Bird migration data: This data set has 35 migration trajectories of Swainson’s Hawk, observed during the period 1995 to 1997, each having geographic coordinates measured at some random times. Several sample paths are shown at the

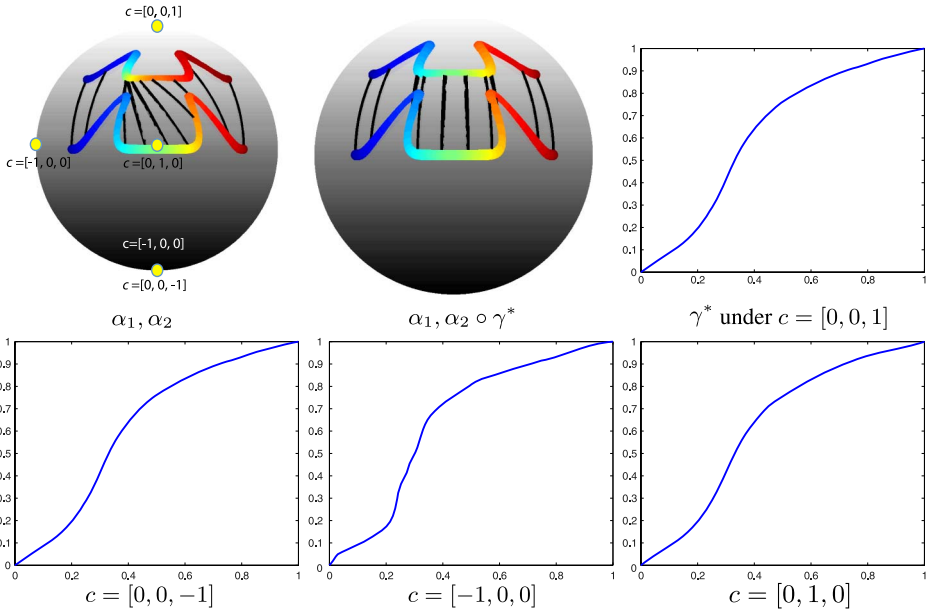


FIG. 2. Registration of trajectories on \mathbb{S}^2 .

top row in Figure 3(a). In the bottom panel of Figure 3(a), we show the optimal warping functions $\{\gamma_i^*\}$ used in aligning them and this clearly highlights a significant temporal variation present in the data. In Figure 3(b) and (c), we show the

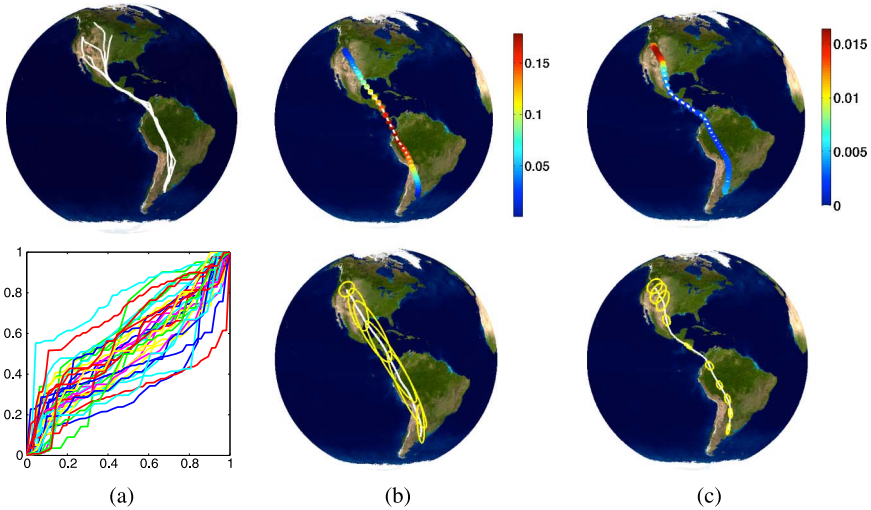


FIG. 3. Swainson's Hawk migration: (a) $\{\alpha_i\}$ (top) and $\{\gamma_i^*\}$ (bottom); (b) μ and $\hat{\rho}$ without registration; (c) μ and $\hat{\rho}$ with registration.

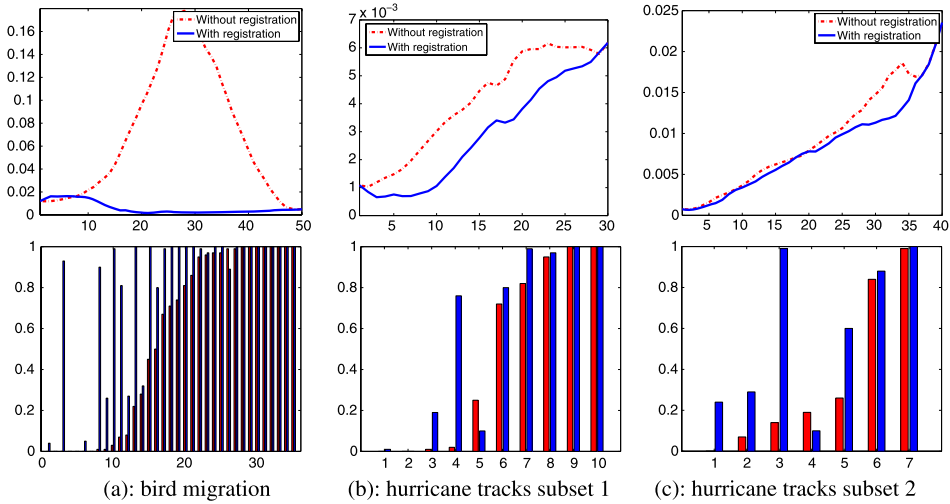


FIG. 4. $\hat{\rho}$ (1st row) and p -values (2nd row) for each trajectory without (red) and with (blue) registration.

Karcher mean μ and the cross-sectional variance $\hat{\rho}$ without and with registration, respectively. In the top row, μ is displayed using colors, where red areas correspond to higher $\hat{\rho}$ value. In the bottom row, the principal modes of variation are displayed by ellipses on tangent spaces. We use the first and second principal tangential directions as the major and minor axes of ellipses, and the corresponding singular values as their sizes. We observe that (1) the mean after registration better preserves the shapes of trajectories, and (2) the variance ellipses before registration have their major axes along the trajectory while the ellipses after registration exhibit a smaller, actual variability in the data. Most of the variability after registration is limited to the top end where the original trajectories indeed have differences. The top row of Figure 4(a) shows a decrease in the function $\hat{\rho}$ due to the registration.

Next we construct a ‘‘Gaussian-type’’ model for these trajectories using estimated summaries for two cases (with and without temporal registration), as described previously, and compute p -values of individual trajectories using Monte Carlo simulation. The results are shown in the bottom of Figure 4(a), where we note a general increase in the p -values for the original trajectories after the alignment. This is attributed to a reduced variance in the model due to temporal alignment and the resulting movement of individual samples closer to the mean values.

Hurricane tracks: We choose two subsets of Atlantic Tracks File 1851-2011, available on the National Hurricane Center website.¹ The first subset has 10 tracks

¹<http://www.nhc.noaa.gov/pastall.shtml>.

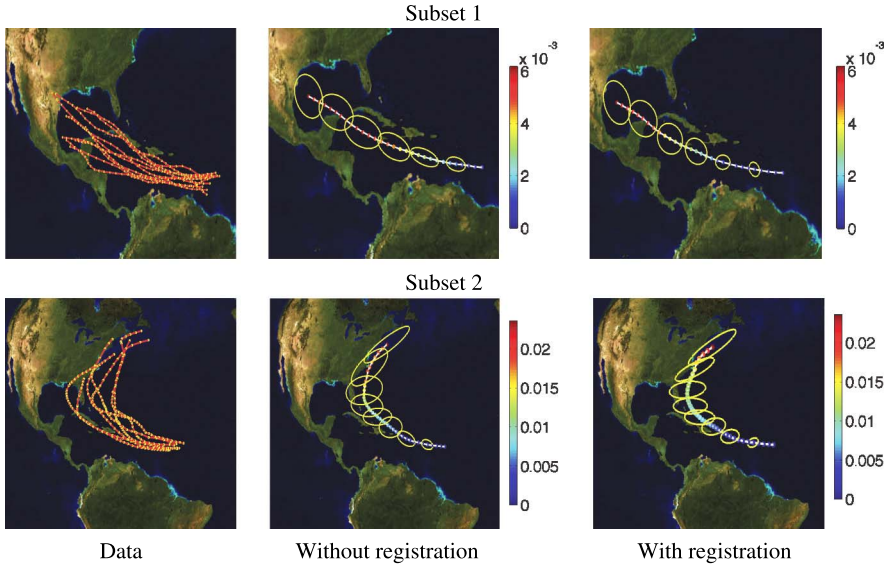


FIG. 5. Summary of hurricane tracks without and with temporal registration.

and another has 7 tracks, with observations at six-hour separation. In Figure 5 we show the data, their Karcher mean and variance without and with registration for each subset. The decrease in the value of $\hat{\rho}$ is shown in the top of Figure 4(b) and (c). Although the decrease here is not as large as the previous example, we observe about 20% reduction in $\hat{\rho}$ on average due to registration. In the bottom plots of Figure 4(b) and (c), it is also seen that there is a general increase of the p -values after registration, although they decreased in a few cases. This is because those trajectories are closer to the mean without registration.

4. Vehicle trajectories on SE(2). Here we study the problem of classifying vehicle trajectories into broad motion patterns using data obtained from traffic videos. While the general motion of a vehicle at a traffic intersection is predictable—left turn, right turn, U turn or straight line—the travel speeds of vehicles may be different in distinct instances due to traffic variations. Since we are interested in tracking position and orientation of a vehicle, we consider individual tracks as parameterized trajectories on $SE(2)$, which is a semidirect product of $SO(2)$ and \mathbb{R}^2 , that is, $SE(2) = SO(2) \ltimes \mathbb{R}^2$. For the rotation component $O \in SO(2)$ and tangent vectors $X_1, X_2 \in T_O(SO(2))$, the standard Riemannian metric is given by $\langle X_1, X_2 \rangle = \text{trace}(X_1^T X_2)$, while we use the Euclidean metric for \mathbb{R}^2 . We choose the rotation component of c as the identity matrix and the translation component as $[0, 0]$. We found that the results of registration, clustering and classification are quite stable with respect to different choices of c . For a tangent vector $W \in T_O(SO(2))$, the parallel transport of W from O to $I_{2 \times 2}$ is $O^T W$. The formulae for the \mathbb{R}^2 component are standard.

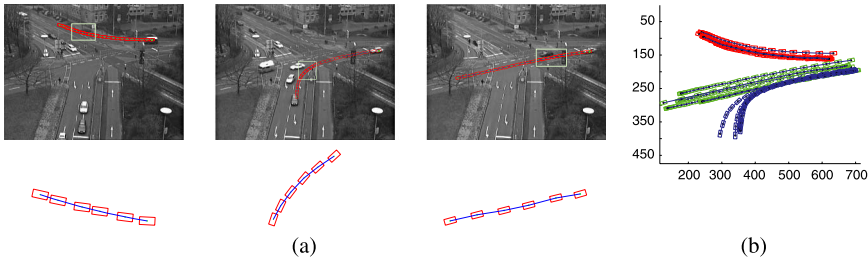


FIG. 6. (a): Real trajectories in $SE(2)$ obtained from a traffic video. (b): Trajectories used for clustering.

Registration of trajectories: The data for this experiment comes from traffic videos available at the Image Sequence Server website.² In Figure 6(a) we show an example trajectory for each of the three classes: right turn (first panel), left turn (second panel), and straight line (third panel). In this small experiment, the total data includes 14 trajectories with 5 trajectories corresponding to right turn indexed from 1 to 5, 5 trajectories of straight line indexed from 6 to 10, and 4 trajectories of left turn indexed from 11 to 14.

Next, in Figure 7 we show two examples of temporally aligning trajectories described above. In Example 1 we first choose a trajectory as α_1 , apply to it a simulated γ and consider this time-warped trajectory as α_2 . The right plot of γ^{-1} (dashed) and γ^* shows that we are able to recover the simulated time-warping using the proposed framework. In Example 2 we show alignment results for trajectories coming from different classes. In this case the distance d_h between the

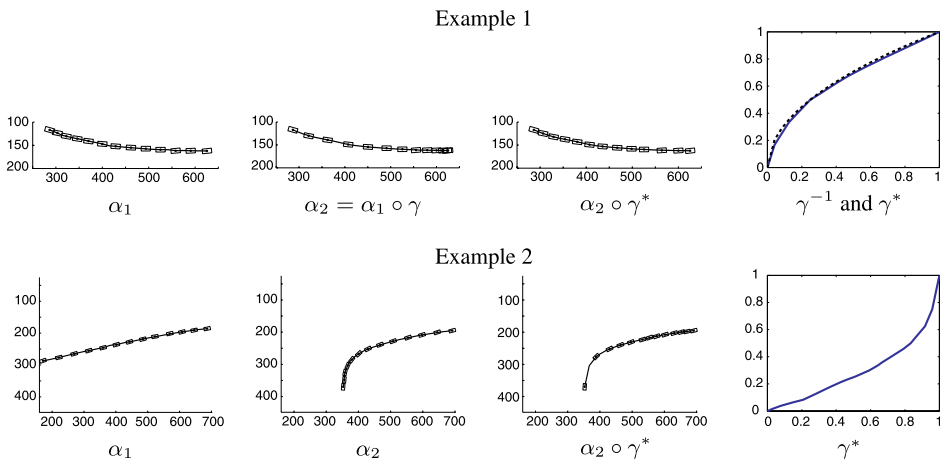


FIG. 7. Registration of trajectories on $SE(2)$.

²http://i21www.ira.uka.de/image_sequences/.

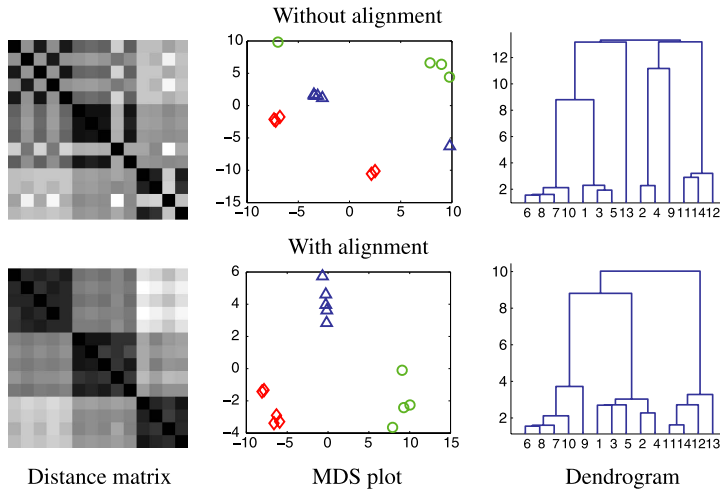


FIG. 8. Clustering without and with alignment.

trajectories is large, since they are from different classes, but it decreases from 14.2 to 10.8 after registration. Furthermore, the registration result is quite intuitive since it matches as much of the common features (straight line part) as possible.

Clustering and classification: Here we study the effects of temporal alignment on clustering and classification. In the first example, we introduce simple speed variations in the vehicle motions; these variations represent either fast-slow or slow-fast movements of a vehicle and apply them randomly to the 14 given trajectories, shown in Figure 6(b). In Figure 8 we display the resulting pairwise distance matrices, multidimensional scaling (MDS) plots and dendrograms computed with and without temporal alignment. The temporal alignment helps in revealing the underlying patterns of the data. Also, it greatly improves the clustering performance.

In the second experiment, we introduce more drastic, random speed variations, corresponding to multiple stop-and-go patterns of a vehicle. We again apply them to the given trajectories and compute the distance matrices with and without temporal alignment. In Table 1 we report the classification performances based on

TABLE 1
Classification rates without and with alignment

Classification rate	1-NN	3-NN	5-NN
Without alignment	64.3%	64.3%	50%
With alignment	100%	100%	93%

1-, 3- and 5-nearest neighbor (NN) classifiers. The method described in this paper produces superior classification of driving patterns. In particular, we can achieve a 100% classification rate using the 1-NN classifier.

5. Shape space of planar contours. Motivated by the problem of analyzing human activities using video data, we are interested in alignment, comparison and averaging of trajectories on the shape space of planar, closed curves. There are several mathematical representations available for this analysis, and we use the representation of Srivastava et al. (2011b). The benefits of using this representation over other methods are discussed there. We provide a very brief description and refer the reader to the original paper for details. Let $\beta : \mathbb{S}^1 \mapsto \mathbb{R}^2$ denote a planar closed curve. Its corresponding q -function is defined as

$$q(s) = \frac{\dot{\beta}(s)}{\sqrt{|\dot{\beta}(s)|}}, \quad s \in \mathbb{S}^1.$$

A major advantage of using q -functions to represent shapes of curves is that the translation variability is automatically removed (q only depends on $\dot{\beta}$). To remove the scaling variability, we re-scale all curves to be of unit length. This restriction translates to the following condition for q -functions: $\int_{\mathbb{S}^1} |\dot{\beta}(s)| ds = \int_{\mathbb{S}^1} |q(s)|^2 ds = 1$. Therefore, the q -functions associated with unit length curves are elements of a unit hypersphere in the Hilbert space $\mathbb{L}^2(\mathbb{S}^1, \mathbb{R}^2)$. In order to study shapes of closed curves, we impose an additional condition, which ensures that the curve starts and ends at the same point. This condition is given by $\int_{\mathbb{S}^1} q(s)|q(s)| ds = 0$. Using these two conditions and the q -function representation, we can define the pre-shape space of unit length, closed curves as

$$\mathcal{C} = \left\{ q \in \mathbb{L}^2(\mathbb{S}^1, \mathbb{R}^2) \mid \int_{\mathbb{S}^1} |q(s)|^2 ds = 1, \int_{\mathbb{S}^1} q(s)|q(s)| ds = 0 \right\}.$$

The shape space of these curves is obtained by removing the re-parameterization group Ψ , the set of diffeomorphisms from \mathbb{S}^1 to itself, and rotation, that is, $\mathcal{S} = \mathcal{C}/(\Psi \times \text{SO}(2))$. A unit circle is used as the standard shape and c in Definition 1 is given by its q -representation. For algorithms on computing parallel transports of tangent vectors along geodesic trajectories in the shape space \mathcal{S} , we refer the reader to Srivastava et al. (2011b).

To illustrate our framework, we apply it to real sequences in the UMD common activities data set. We use a subset of 8 classes from this data set with 10 instances in each class. Each instance consists of 80 consecutive planar closed curves. As a first step, we down-sample each of these trajectories to 17 contours.

Registration: An example of registering two trajectories of planar closed curves from the same class is shown in Figure 9. The distance d_h between the two trajec-

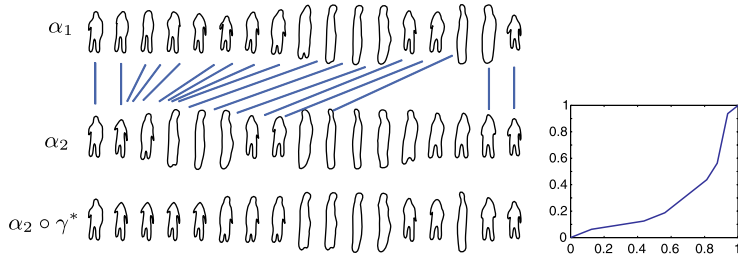


FIG. 9. Registration of two trajectories on the shape space of planar contours.

jectories decreases from 4.27 to 3.26. The optimal γ^* for this registration is shown in the right panel.

Statistical summaries: We give an example of averaging and registration of multiple trajectories using Algorithm 1 in Figure 10. The aligned sample trajectories within the same class are much closer to each other than before temporal alignment. The energy when computing the Karcher mean converges quickly, as shown at the left bottom corner in Figure 10. The right bottom plot shows that the cross-sectional variance $\hat{\rho}$ is significantly reduced after temporal registration.

Classification: For this activity data set we computed the full pairwise distance matrix for trajectories, using d_h (without registration) and d_s (with registration).

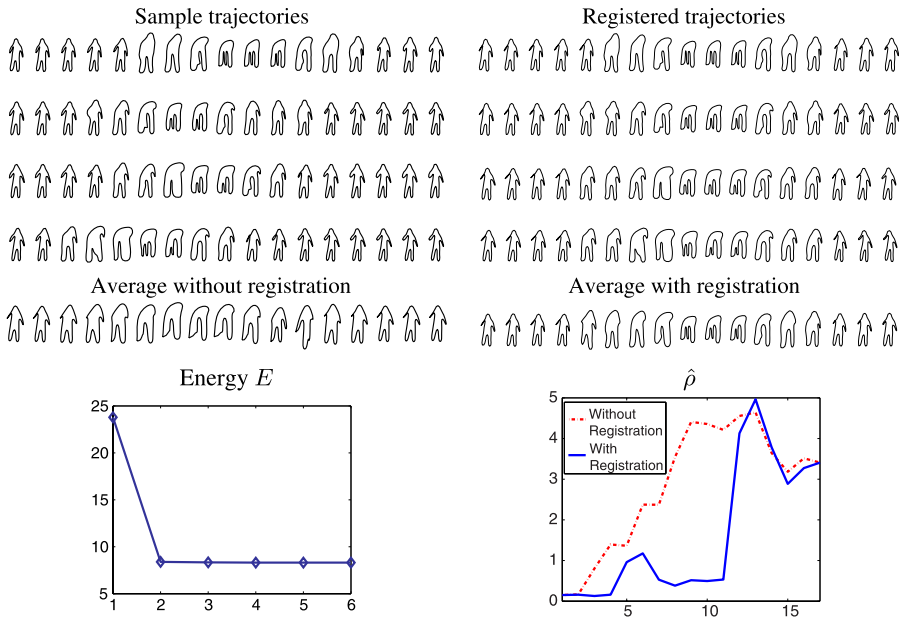


FIG. 10. Registration and summary of multiple trajectories.

The leave-one-out nearest neighbor classification rate (1-NN as described earlier) for d_s is 95% as compared to only 87.5% when using d_h .

6. Conclusion. Statistical analysis of trajectories on nonlinear manifolds is important in many areas, including medical imaging and computer vision. In this paper we have provided a framework for registering, comparing, summarizing and modeling trajectories on \mathbb{S}^2 , $SE(2)$ and shape space of planar contours under invariance to time-warping. Specifically, we have defined a proper metric, which allows us to register trajectories and compute their sample means and covariances. For future work, we would like to extend the framework to other applications with other underlying manifolds. In addition, we encourage further efforts on the statistical modeling of such trajectories.

REFERENCES

- AGGARWAL, J. K. and CAI, Q. (1999). Human motion analysis: A review. *Comput. Vis. Image Underst.* **73** 428–440.
- BEG, M. F., MILLER, M. I., TROUVE, A. and YOUNES, L. (2005). Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vis.* **61** 139–157.
- BERTSEKAS, D. P. (2007). *Dynamic Programming and Optimal Control*, 3rd ed. Athena Scientific, Belmont, MA.
- CHRISTENSEN, G. E. and JOHNSON, H. J. (2001). Consistent image registration. *IEEE Trans. Med. Imag.* **20** 568–582.
- DRYDEN, I. L. and MARDIA, K. V. (1998). *Statistical Shape Analysis*. Wiley, Chichester. [MR1646114](#)
- GAVRILA, D. M. (1999). The visual analysis of human movement: A survey. *Comput. Vis. Image Underst.* **73** 82–98.
- JUPP, P. E. and KENT, J. T. (1987). Fitting smooth paths to spherical data. *J. R. Stat. Soc. Ser. C Appl. Stat.* **36** 34–46. [MR0887825](#)
- KENOBI, K., DRYDEN, I. L. and LE, H. (2010). Shape curves and geodesic modelling. *Biometrika* **97** 567–584. [MR2672484](#)
- KNEIP, A. and RAMSAY, J. O. (2008). Combining registration and fitting for functional models. *J. Amer. Statist. Assoc.* **103** 1155–1165. [MR2528838](#)
- KUME, A., DRYDEN, I. L. and LE, H. (2007). Shape-space smoothing splines for planar landmark data. *Biometrika* **94** 513–528. [MR2410005](#)
- KURTEK, S., WU, W. and SRIVASTAVA, A. (2011). Signal estimation under random time warpings and its applications in nonlinear signal alignments. *Adv. Neural Inf. Process. Syst.* **24** 676–683.
- LE, H. (2003). Unrolling shape curves. *J. Lond. Math. Soc. (2)* **68** 511–526. [MR1994697](#)
- LE, H. and KUME, A. (2000). The Fréchet mean shape and the shape of the means. *Adv. in Appl. Probab.* **32** 101–113. [MR1765168](#)
- LIU, X. and MÜLLER, H.-G. (2004). Functional convex averaging and synchronization for time-warped random curves. *J. Amer. Statist. Assoc.* **99** 687–699. [MR2090903](#)
- MARDIA, K. V. and JUPP, P. E. (2000). *Directional Statistics*, 3rd ed. Wiley, Chichester. [MR1828667](#)
- MICHOR, P. W. and MUMFORD, D. (2007). An overview of the Riemannian metrics on spaces of curves using the Hamiltonian approach. *Appl. Comput. Harmon. Anal.* **23** 74–113. [MR2333829](#)
- OWEN, J. C. and MOORE, F. R. (2008). Swainson’s thrushes in migratory disposition exhibit reduced immune function. *Journal of Ethology* **26** 383–388.

- ROBINSON, D. (2012). Functional analysis and partial matching in the square root velocity framework. Ph.D. thesis, Florida State Univ., Tallahassee, FL.
- SHAH, J. (2008). H^0 -type Riemannian metrics on the space of planar curves. *Quart. Appl. Math.* **66** 123–137. [MR2396654](#)
- SRIVASTAVA, A., WU, W., KURTEK, S., KLASSEN, E. and MARRON, J. S. (2011a). Registration of functional data using Fisher–Rao metric. Preprint. Available at [arXiv:1103.3817v2](#).
- SRIVASTAVA, A., KLASSEN, E., JOSHI, S. H. and JERMYN, I. H. (2011b). Shape analysis of elastic curves in Euclidean spaces. *IEEE Trans. Pattern Anal. Mach. Intell.* **33** 1415–1428.
- SUNDARAMOORTHY, G., MENNUCCI, A., SOATTO, S. and YEZZI, A. (2011). A new geometric metric in the space of curves, and applications to tracking deforming objects by prediction and filtering. *SIAM J. Imaging Sci.* **4** 109–145. [MR2792407](#)
- TROUVÉ, A. and YOUNES, L. (2000). On a class of diffeomorphic matching problems in one dimension. *SIAM J. Control Optim.* **39** 1112–1135. [MR1814269](#)
- TUCKER, J. D., WU, W. and SRIVASTAVA, A. (2013). Generative models for functional data using phase and amplitude separation. *Comput. Statist. Data Anal.* **61** 50–66. [MR3063000](#)
- VEERARAGHAVAN, A., SRIVASTAVA, A., ROY-CHOWDHURY, A. K. and CHELLAPPA, R. (2009). Rate-invariant recognition of humans and their activities. *IEEE Trans. Image Process.* **18** 1326–1339. [MR2742162](#)
- YOUNES, L., MICHOR, P. W., SHAH, J. and MUMFORD, D. (2008). A metric on shape space with explicit geodesics. *Atti Accad. Naz. Lincei Cl. Sci. Fis. Mat. Natur. Rend. Lincei (9) Mat. Appl.* **19** 25–57. [MR2383560](#)

J. SU
DEPARTMENT OF MATHEMATICS AND STATISTICS
TEXAS TECH UNIVERSITY
LUBBOCK, TEXAS 79409
USA
E-MAIL: jingyong.su@ttu.edu

S. KURTEK
DEPARTMENT OF STATISTICS
OHIO STATE UNIVERSITY
COLUMBUS, OHIO 43210
USA

E. KLASSEN
A. SRIVASTAVA
DEPARTMENT OF MATHEMATICS
FLORIDA STATE UNIVERSITY
TALLAHASSEE, FLORIDA 32306
USA