

A Generalized Lyapunov Feature for Dynamical Systems on Riemannian Manifolds

Rushil Anirudh
ranirudh@asu.edu

Vinay Venkataraman
vvenka18@asu.edu

Pavan Turaga
pturaga@asu.edu

School of Electrical, Computer and
Energy Engineering
School of Arts, Media, and Engineering
Arizona State University
Tempe, Arizona, USA.

Abstract

Dynamic phenomena such as human activities, dynamic scenes, and moving crowds are commonly observed through visual sensors, resulting in feature trajectories sampled in time. Such phenomena can be accurately modeled by taking the temporal variations and changes into account. For problems where the trajectories are sufficiently different, elastic metrics can provide distances that are invariant to speed, but for more complex problems such as fine grained activity classification, one needs to exploit higher order dynamical properties. For features in the Euclidean space, applications such as crowd monitoring, dynamic scene recognition and human movement quality analysis have found a lot of success this way. In this paper we propose the largest Riemannian Lyapunov exponent (L-RLE), which is the first generalization of the largest Lyapunov exponent to Riemannian manifolds. The largest Lyapunov exponent is a classic measure to quantify the amount of chaos within signals in the Euclidean space, and allows us to exploit higher order dynamics for various applications. We show the effectiveness of the L-RLE on two manifolds - the Grassmann and the $SO(3)$ lie group. By modeling human actions as dynamic processes evolving on Riemannian manifolds, we show that L-RLE can measure the amount of chaos within each action accurately. We show that our measure is a good generalization of largest Euclidean Lyapunov exponent (L-ELE), and is less susceptible to arbitrary distortions.

1 Introduction

Dynamic phenomena such as human activities are commonly observed through visual sensors, typically resulting in feature trajectories sampled in time. Accurate metrics on such trajectories are those that take its temporal nature into account. For example, in human action recognition it is well known that accounting for temporal re-parametrization improves the distance metric between two actions, resulting in significantly improved recognition performance [24, 28]. For problems where the elasticity of the metric does not suffice, one has to go a step further and study the properties of the dynamical system that generates the trajectory. A diverse set of applications have benefited from dynamics based metrics such as –

human action recognition [4, 57], bio-mechanics [23], dynamics of crowds [9], and dynamic scene recognition [22]. It has also been shown that such properties can help in fine grained classification between similar kinds of human movement [60]. Exploiting the dynamics is relatively easy when the concerned feature space is Euclidean, but the last few years have seen an increased interest in modeling features that lie on non Euclidean spaces such as Riemannian manifolds. Some examples of such features are – shapes on Kendall’s shape space [45], Histogram features [40], skeletal features [29], and Covariance features [27]. More recently, there has also been an interest in modeling trajectories on such non-linear manifolds using elastic metrics [6, 24]. However, the study of dynamical invariants has remained unexplored.

In this paper, we address the problem of uncovering the properties of dynamical processes evolving on Riemannian manifolds, for applications in human action analysis. While the problem of modeling Riemannian trajectories is recent, to the best of the authors’ knowledge there has not been a study to exploit Riemannian dynamics in computer vision. In this regard, we propose the largest-Riemannian Lyapunov exponent (L-RLE), which is a generalization of the largest Euclidean Lyapunov exponent [64], a widely used feature that measures chaos in a time series. Traditional chaotic invariant measures determine the average rate of divergence (or convergence) between nearby states of a system over time. We show that the proposed measure can be used to quantify the amount of chaos within a Riemannian dynamical process. Further, we show that for human action analysis from silhouettes on the Grassmann manifold and curves in $SO(3)$, the representational manifold itself is a good candidate for the *phase space*. Experiments indicate that the L-RLE correlates well with the largest Lyapunov exponent extracted on Euclidean features for human action analysis, while also retaining the robustness advantages and invariances of the manifold-valued features.

More broadly, the chaotic properties of time series data have been found useful in modeling temporal data in several applications. Measures such as the largest Euclidean Lyapunov exponent (L-ELE) allow us to quantify the intrinsic dynamical nature of these phenomenon. Existing methods to compute the L-ELE assume the time series is Euclidean, i.e., the underlying metric is the commonly used ℓ_2 - norm. Whereas many state-of-the-art representations or features involve features that are non-Euclidean. We generalize the notion of a Lyapunov exponent to two different manifolds - the special orthogonal group denoted as $SO(3)$, which is a lie group containing all 3×3 rotation matrices. These occur commonly in data collected from smartphones, and fitness trackers, which measure orientation information. Next, we represent shape silhouettes of humans performing different activities, and model the shapes as affine invariant subspaces, which naturally lie on the Grassmann manifold [46]. This shape representation provides invariance to affine transformations in addition to scale and rotation changes - which is useful in modeling small camera viewpoint changes.

2 Related Work

Recent years have seen advancements in understanding different properties of Riemannian trajectories, motivated by the increasing availability temporal data from videos. For example, [24] proposed a rate invariant representation known as the Transport Square Root Velocity Function (TSRVF) for trajectories, such that the final metric remains unchanged to identical time warping. Next, [6] showed that the TSRVF could be used to exploit statistical properties of the trajectories to obtain a lower dimensional embedding. Another study [56] models

human actions as Riemannian trajectories, by transporting all the points in the trajectory to the starting point, this representation is used to learn a subspace that preserves geodesics. The development of such tools to manipulate, and represent such non-linear trajectories provides a foundation to explore even higher order properties such as dynamics.

While the idea of studying the dynamics of Riemannian trajectories remains to be addressed, the idea of using differential geometry to understand the phase space obtained from Euclidean time series exists [10]. A closely related theoretical piece of work proposed the generalization of the finite time Lyapunov exponent (FTLE) and Lagrangian coherent structures (LCS) to Riemannian manifolds [16]. We differentiate our work by proposing an algorithm to compute the largest Lyapunov exponent, a different measure of chaos than FTLE. We also validate our work on real and synthetic data, with varying degrees of chaos.

Traditional dynamical modeling approaches for Euclidean space data include parametric methods such as Hidden Markov Models (HMMs) and Linear Dynamical Systems (LDSs), which have been used for computer vision applications like action recognition [3, 5] and gait analysis [8, 14]. Recent work by Ali *et al.* proposed the use of nonparametric modeling approach using ideas from chaos theory to model the dynamics in human actions [4]. The authors use Rosenstein’s algorithm [19] to estimate largest Lyapunov exponent from trajectories of action data as part of their feature representation. We propose an extension of Rosenstein’s algorithm that computes it for Riemannian manifolds.

3 Preliminaries

We describe the geometric properties of the Grassmann and the $SO(3)$ manifolds next. For a primer on differential geometry and Riemannian manifolds, we refer the interested reader to excellent resources on the topic [1, 9].

3.1 The Grassmann manifold

We represent the shape silhouette of a human extracted from an image as a point on the Grassmann manifold. The silhouette is represented as $L = [(x_1, y_1), (x_2, y_2) \dots, (x_m, y_m)]$, where (x_i, y_i) represent the zero centered landmarks on the silhouette. Affine transforms of a base shape L_{base} can be expressed as $L_{affine}(A) = L_{base}A^T$, and this multiplication by a full rank matrix on the right preserves the column-space of the matrix, L_{base} . Thus the 2D subspace of \mathbb{R}^m spanned by the columns of the shape, i.e. $span(L_{base})$ is invariant to affine transforms of the shape. Subspaces such as these can be identified as points on a Grassmann manifold [26]. A given d -dimensional subspace of \mathbb{R}^m , \mathcal{Y} can be associated with a idempotent rank- d projection matrix $P = YY^T$, where Y is a $m \times d$ orthonormal matrix such as $span(Y) = \mathcal{Y}$. The space of $m \times m$ projectors of rank d , denoted by $\mathbb{P}_{m,d}$ can be embedded into the set of all $m \times m$ matrices - $\mathbb{R}^{m \times m}$ - which is a vector space. Using the embedding $\Pi : \mathbb{R}^{m \times m} \rightarrow \mathbb{P}_{m,d}$ we can define a distance function on the manifold using the metric inherited from $\mathbb{R}^{m \times m}$.

$$d^2(P_1, P_2) = tr(P_1 - P_2)^T(P_1 - P_2) \quad (1)$$

The projection $\Pi : \mathbb{R}^{m \times m} \rightarrow \mathbb{P}_{m,d}$ is given by:

$$\Pi(M) = UU^T \quad (2)$$

where $M = USV^T$ is the d -rank SVD of M .

3.2 The Special Orthogonal Group - $SO(3)$

The special orthogonal group, denoted by $SO(3)$ is a Lie group, containing the set of all 3×3 rotation matrices about the origin of three-dimensional Euclidean space \mathbb{R}^3 under the operation of composition. The 3×3 identity matrix I_3 is an element of $SO(3)$ and is the identity element of the group. The exponential map, which is defined as $\exp_{SO(3)} : \mathfrak{so}(3) \rightarrow SO(3)$ and the inverse exponential map, defined as $\log_{SO(3)} : SO(3) \rightarrow \mathfrak{so}(3)$ are used to traverse between the manifold and the tangent space respectively. The exponential and inverse exponential maps for $SO(3)$ are simply the matrix exponential and matrix logarithms respectively, from the identity element I_3 . For efficient implementations of a general exponential and inverse exponential maps between any two arbitrary points, we refer the reader to [18]. The tangent space at I_3 of a $SO(3)$ is called the Lie algebra of $SO(3)$, denoted by $\mathfrak{so}(3)$. It is a 3-dimensional space formed by matrices of the form:

$$B = \begin{bmatrix} 0 & -u_3 & u_2 \\ u_3 & 0 & -u_1 \\ -u_2 & u_1 & 0 \end{bmatrix}, \quad (3)$$

An equivalent representation of B is $\text{vec}(B) = [u_1, u_2, u_3]$ which lies on \mathbb{R}^3 .

4 Dynamical Systems on Geometric Spaces

Dynamical systems are mathematical models which simulate a physical phenomenon of states evolving over time. Chaos theory studies the behavior of nonlinear dynamical systems, that are highly sensitive to initial conditions. Any perturbation to the initial conditions of such systems yield widely diverging dynamics. This behavior is known as deterministic chaos. Convincing evidence for existence of deterministic chaos has been provided from a variety of research experiments [21, 25]. Due to the inherent variability in human movement, tools from chaos theory have found wide applications in the bio-mechanics community for analysis of human actions [23].

In the most general sense, a dynamical system is the tuple $\langle \mathcal{M}, f, \mathcal{T} \rangle$, where \mathcal{M} is a manifold, \mathcal{T} is non-negative time and f is a diffeomorphism that governs the evolution of trajectories, defined as $f : \mathcal{M} \times \mathcal{T} \rightarrow \mathcal{M}$. In the Euclidean space, one can learn the parametric representation of the function f . When it is difficult to estimate the function directly, one can estimate properties of the function. One such property is the largest Euclidean Lyapunov exponent (L-ELE), which has seen a lot of success for dynamical analysis of Euclidean signals. There is currently no existing method to estimate f directly for trajectories Riemannian manifolds, therefore we propose an algorithm to first generalize the L-ELE to manifolds.

The largest Lyapunov exponent, denoted as λ , is a measure of average rate of divergence (or convergence) of initially closely-spaced trajectories over time [1, 22]. A positive value indicates orbital divergence and hence chaos in the system. A negative value indicates orbital convergence and hence a dissipative system. A practical method for estimating the largest Lyapunov exponent from a time series proposed by Rosenstein [19] quantifies chaos by monitoring the rate of divergence of closely spaced trajectories over time. The algorithm is fast, easy to implement and robust to changes in embedding dimension, size of dataset,

embedding delay and noise level. We refer to the Euclidean space largest Lyapunov exponent as the largest Euclidean Lyapunov exponent (L-ELE) to differentiate it with our measure, the largest Riemannian Lyapunov exponent (L-RLE). More formally, the Lyapunov exponent is defined as follows:

$$d_j(i) = d_j(0)e^{\lambda_1(i\Delta t)}, \quad (4)$$

where $d_j(0)$ is the initial separation in the phase space and $d_j(i)$ is the separation after i time steps of Δt .

4.1 Largest Riemannian Lyapunov Exponent (L-RLE)

The L-ELE is computed as follows [19]: the embedding parameters lag and dimension are estimated using the Fast Fourier Transform (FFT), which are used to construct the phase space. Next, in the phase space the nearest neighbors are calculated constrained on temporal separation. This is used to estimate how far two points have diverged in the phase space as the attractor evolves over time. In generalizing this to Riemannian manifolds, we first describe how the manifold itself can be treated as the phase space next.

The manifold as a phase space: The phase space is defined as an approximation to the high dimensional state space of the dynamic system that governs the observed time series. Obtaining the phase space directly is challenging because we often do not have access to all the information required to reconstruct it, instead many algorithms resort to reconstructing the phase space. However, reconstruction of the phase space requires estimating the period using the FFT, which do not generalize well to manifolds. On the other hand, action features such as shape silhouettes and stick figures are already high dimensional, and contain a lot of information. For example, the states in a action sequences may be closely related to the poses of the human, which are naturally points on an appropriate high-dimensional shape manifold. Therefore, we propose that the underlying manifold can be treated as the phase space of the system, where each time sample behaves as a “state”. We show in our experiments that the manifold behaves similar to the phase space, and therefore is a good approximation.

Computing the L-RLE: In the phase space, the next step involves measuring how far two nearby points have diverged over time. With the geodesic distance, we first perform a k-nearest neighbor and then compute the quantities $d_j(0), d_j(i)$ from (4) for a given point and its nearest neighbor. To compute λ from (4), it is useful to rearrange as follows.

$$\ln(d_j(i)) \approx \ln(d_j(0)) + \lambda(i \Delta t) \quad (5)$$

Equation (4) represents a set of approximately parallel lines for different points in the phase space. The largest Lyapunov exponent is calculated as the slope of the “average” line. The procedure to estimate the L-RLE is outlined in algorithm 1.

5 Experimental validation

To evaluate the proposed dynamical measure, we apply it to human actions to study their dynamic properties. We use the UMD actions dataset [28] which contains 10 actions such as *walk, run, squat, throw a ball, talk on the cell phone, push an object, and batting*. These are performed 10 times, giving a total of 100 actions in the dataset. The relatively static background allows us to extract the shape silhouettes easily, which we represent as a subspace

Algorithm 1 Largest Lyapunov exponent on Manifolds

Input: $\alpha(t) \in \mathcal{M}, t = 1 \dots T$

Assume the manifold is the phase space.

for $j = 1 \rightarrow T$ **do**

Find K Nearest neighbors constrained in time [19].

for $i = 1 \rightarrow K$ **do**

$d_j(i) = \min_k d_{\mathcal{M}}(\alpha(j), \alpha(i)),$

$\ln(d_j(i)) = \ln(d_j(0)) + \lambda_1(i \Delta t)$

end for

Fit a line, L_j , for each set of d_j s, compute its slope m_j .

end for

Average slope gives a robust estimate of L-RLE $\lambda = \frac{1}{T} \sum_j m_j$

which is a point on the Grassmann manifold. This results in actions being represented as trajectories on the Grassmann manifold. In our first experiment, we motivate the manifold as the phase space, before computing the RLE. We show quantitative and qualitative results indicating the advantage of the proposed measure. Apart from directly working with the shape trajectories, we consider alternate representations to understand and validate our measure.

Alternate Representations: We first motivate the idea of using the manifold itself as the phase space before computing the Largest Lyapunov exponent. Since the alternative representations are extracted from the same actions in the same dataset, and there is a severe lack of “ground truth”, we make the assumption that the dynamics remains unchanged across the features. That is, the dynamics of human actions remains unchanged when represented using features that are Euclidean or non-Euclidean. To the best of our knowledge, we are not aware of any work that can claim otherwise.


Multivariate Embedding (MVE) – We compare with an established algorithm to embed multivariate time series data in the Euclidean space, for time-delay reconstruction of phase space known as the multivariate embedding [20]. This simple yet powerful extension of univariate embedding as proposed by Cao et al. [20] has proven to be useful in computer vision applications such as action synthesis and dynamic texture synthesis [6]. Recent theoretical and empirical findings have demonstrated that multivariate embedding of time series data by simple concatenation of individual univariate embedding vectors achieves good state space reconstruction as evaluated by the shape and dynamics distortion measures [6]. The embedding method only works with Euclidean time series data, and hence we consider the 2D landmarks on the silhouette per frame as our feature for each action. This results in each action being represented as an $(N \times 2) \times T$, where N is the number of landmarks on each silhouette, and T is the total number of frames. Using this data, we perform uniform multivariate embedding. Given multivariate time series data $\{x_{i,t}\}_{t=1}^T, i = 1, \dots, p$, where p is the dimension of time series data, the reconstructed phase space vector is of the form

$$\begin{aligned} \mathbf{z}_t = & [x_{1,t}, x_{1,t+\tau_1}, \dots, x_{1,t+(m_1-1)\tau_1}, \\ & x_{2,t}, x_{2,t+\tau_2}, \dots, x_{2,t+(m_2-1)\tau_2}, \\ & \dots, \\ & x_{p,t}, x_{p,t+\tau_p}, \dots, x_{p,t+(m_p-1)\tau_p}] \end{aligned} \quad (6)$$

where m_i and τ_i are respectively the embedding dimension and time delay for each of the p -dimension in the multivariate time series data.

Vector Field Parallel Transport (VFPT) – We also use an intermediate representation, where we represent each trajectory as a collection of tangents. For an action j , $\mathcal{F}_j = \{\dot{\alpha}_{t \rightarrow t+1}(t) | \forall i = 1, \dots, T\}$, where $\dot{\alpha}_{t \rightarrow t+1}(t)$ represents the tangent that goes from $\alpha(t)$ to $\alpha(t+1)$ in unit time. We perform a parallel transport on all the tangent vectors and bring them to a common point at the Riemannian center of mass (RCM) [13]. We treat the transported tangent bundle as the phase space in this case. This feature takes the geometry into account while also giving us a Euclidean representation, which we can exploit for visualizing the phase space. We use the largest Euclidean Lyapunov exponent (L-ELE) algorithm [19] on this feature.

Viewpoint invariance: Figure 1a shows the phase space for the walking action, in 3D, after performing dimensionality reduction using Laplacian Eigenmaps [20]. It is seen that in both cases, the cyclic pattern of the action is captured even after dimensionality reduction. Since the Grassmann manifold can afford us affine invariance, we artificially shear the shapes to simulate minor camera viewpoint changes. Since the multivariate embedding uses the coordinate locations in each frame, the phase space estimated from the sheared data is significantly distorted. The phase space obtained from the Grassmann representation remains unchanged, as shown in figures 1b.



(a) Phase space estimates on shape data

(b) Phase space estimates after change in viewpoint.

Figure 1: The Grassmann manifold for shapes, as the phase space for human activities provides invariance to commonly observed problems such as viewpoint, scale and shift. The resulting phase space is significantly distorted in the case of multivariate embedding, but remains unchanged in our case. We perform dimensionality reduction to facilitate easy visualization using the Laplacian-eigen maps [20].

Phase space type	Correlation
Landmarks + Multivariate Embedding + L-ELE [19] (reference)	1.00
Vector Field Parallel Transport + L-ELE	0.52
Shape feature + L-RLE	0.76
Landmarks + Multivariate Embedding + L-ELE [19] (sheared)	0.27
Vector Field Parallel Transport + L-ELE (sheared)	0.40
Shape feature + L-RLE (sheared)	0.763

Table 1: The proposed Riemannian Lyapunov exponent (RLE) on the Grassmann manifold closely relates to the estimate obtained from the multivariate embedding on the landmarks of the silhouette. It is also much more robust to affine transforms, compared to the Euclidean measure. Here we assume the standard largest Euclidean Lyapunov exponent (L-ELE) without any shearing to be the reference standard.

We estimate the L-RLE on 100 actions in the UMD Actions dataset and report the correlation between our measure and the Lyapunov from Multivariate Embedding on Euclidean features in table 1. It is seen that our measure compares well with the vector space version – indicating that for clean data the L-RLE is a good generalization of the L-ELE. Further, when we artificially shear the data to simulate minor viewpoint changes, the multivariate embedding algorithm and L-ELE algorithm fail severely due to the distortion in the data whereas the L-RLE remains robust to such changes.

5.1 Validation on standard attractors

A challenging aspect to generalizing the largest Lyapunov exponent to Riemannian manifolds is validation. In the Euclidean space, a common way to evaluate a chaotic measure is to test it on different attractors arising from closed-form dynamical equations such as the Lorenz [17] and the Rossler [20] systems. Any chaotic measure must be as close to zero as possible for perfectly periodic time series. Unfortunately, the Rossler and Lorenz systems do not generalize easily to manifolds. We approximate these systems by generating them on a tangent plane and wrapping them onto the manifold. Once again, we assume that the wrapping action from the tangent space to the manifold does not affect the chaotic nature of the time series. This maybe a restricting assumption in general, but is valid for small deviations from the pole of the tangent space. We choose the special orthogonal group $SO(3)$, which is a Lie group, since it is 3-dimensional and allows us to naturally embed a 3-D time series generated on its Lie algebra. The properties of the Largest Euclidean Lyapunov Exponent (L-ELE) which we expect to observe here are the following: 1) The value for periodic signals must be zero, 2) The L-ELE is direct measure of the chaotic nature of the signal, i.e. the higher the chaos within a signal, the higher its measured value. This pattern is clearly observed in table 2, where a periodic signal gives us a value thats close to zero. We also compare the L-ELE values as a reference, it is observed that even though the values are not exactly the same, the trend is clear. A trajectory thats more chaotic has a higher L-RLE value, similar to the Euclidean case.

Attractor	L-Euclidean LE	L-Riemannian LE
Lorenz [17] (higher chaos)	1.50	21.02
Rossler [20](lower chaos)	0.09	7.26
Periodic (zero chaos)	0	0.008

Table 2: **Validating the L-RLE:** We embed the standard attractors into the $SO(3)$ lie group and evaluate the L-RLE using the proposed algorithm. It is seen that the nature of the L-RLE is consistent with the L-ELE, higher chaos implies a higher value, and periodicity implies a very low value.

6 Discussion and Conclusion

We presented a formulation to study the invariant properties of dynamical systems evolving on Riemannian manifolds. Such systems occur frequently in problems such as human movement analysis, action recognition, and crowd analysis in computer vision. The invariant properties of Euclidean dynamical systems have been useful in characterizing temporal events for such applications in computer vision. However, there is a lack of such methods for dynamical systems on non-Euclidean spaces. To address this, we proposed a generalization of the largest Lyapunov exponent, a classic chaotic measure, to Riemannian manifolds. Towards this end, we use the ambient manifold as the phase space and compute the largest

Riemannian Lyapunov exponent (L-RLE). We show that it correlates well with the analogous measure for Euclidean dynamics. By estimating the L-RLE on standard attractors such as the Lorenz and Rossler, we show that our L-RLE measures the chaotic properties accurately. We have validated the L-RLE under the assumption that the dynamical latent properties of temporal events remains unchanged when observed in different feature spaces. A direction of future study could be to further investigate how the dynamical properties are preserved when the same event is observed in different modalities. While the presented work is primarily empirical, a theoretical analysis of Riemannian dynamical invariants and associated estimation algorithms, such as the proposed one, may be fruitful areas of future work.

Acknowledgement

This work was supported in part by NSF CAREER grant 1452163 and NSF 1320267.

References

- [1] Henry DI Abarbanel. *Analysis of observed chaotic data*. New York: Springer-Verlag, 1996.
- [2] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008. ISBN 978-0-691-13298-3.
- [3] S. Ali and M. Shah. A lagrangian particle dynamics approach for crowd flow segmentation and stability analysis. *CVPR*, pages 1–6, 2007.
- [4] Saad Ali, Arslan Basharat, and Mubarak Shah. Chaotic invariants for human action recognition. In *ICCV*, pages 1–8, 2007.
- [5] Rushil Anirudh, Pavan Turaga, Jingyong Su, and Anuj Srivastava. Elastic functional coding of human actions: From vector-fields to latent variables. In *CVPR*, 2015.
- [6] Arslan Basharat and Mubarak Shah. Time series prediction by chaotic modeling of nonlinear dynamical systems. In *IEEE International Conference on Computer Vision*, pages 1941–1948, 2009.
- [7] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15:1373–1396, 2003.
- [8] Alessandro Bissacco, Alessandro Chiuso, Yi Ma, and Stefano Soatto. Recognition of human gaits. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 52–57, 2001.
- [9] W. M. Boothby. *An Introduction to Differentiable Manifolds and Riemannian Geometry. Revised 2nd Ed.* Academic, New York, 2003.
- [10] Thomas J Bridges and Sebastian Reich. Computing lyapunov exponents on a stiefel manifold. *Physica D: Nonlinear Phenomena*, 156(3):219–238, 2001.
- [11] Liangyue Cao, Alistair Mees, and Kevin Judd. Dynamics from multivariate time series. *Physica D: Nonlinear Phenomena*, 121(1):75–88, 1998.
- [12] R. Chaudhry, A. Ravichandran, G. Hager, and R. Vidal. Histograms of oriented optical flow and Binet-Cauchy kernels on nonlinear dynamical systems for the recognition of human actions. In *CVPR*, June 2009.
- [13] K. Grove and H. Karcher. How to conjugate C^1 -close group actions. *Math.Z.*, 132:11–20, 1973.
- [14] Amit Kale, Aravind Sundaresan, AN Rajagopalan, Naresh P Cuntoor, Amit K Roy-Chowdhury, Volker Kruger, and Rama Chellappa. Identification of humans using gait. *IEEE Transactions on Image Processing*, 13(9):1163–1173, 2004.
- [15] D. G. Kendall. Shape manifolds, procrustean metrics and complex projective spaces. *Bulletin of London Mathematical society*, 16:81–121, 1984.

- [16] Francois Lekien and Shane D Ross. The computation of finite-time lyapunov exponents on unstructured meshes and for non-euclidean manifolds. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20(1): 017505, 2010.
- [17] Edward N Lorenz. Deterministic nonperiodic flow. *Journal of the atmospheric sciences*, 20(2):130–141, 1963.
- [18] Richard M Murray, Zexiang Li, and S Shankar Sastry. *A mathematical introduction to robotic manipulation*. CRC press, 1994.
- [19] M.T. Rosenstein, J.J. Collins, and C.J. De Luca. A practical method for calculating largest lyapunov exponents from small data sets. *Physica D: Nonlinear Phenomena*, 65(1):117–134, 1993.
- [20] Otto E RöSSLer. An equation for continuous chaos. *Physics Letters A*, 57(5):397–398, 1976.
- [21] J-C Roux, Reuben H Simoyi, and Harry L Swinney. Observation of a strange attractor. *Physica D: Nonlinear Phenomena*, 8(1):257–266, 1983.
- [22] Aswin C. Sankaranarayanan, Pavan K. Turaga, Richard G. Baraniuk, and Rama Chellappa. Compressive acquisition of dynamic scenes. In *ECCV (1)*, pages 129–142, 2010.
- [23] Nicholas Stergiou. *Innovative Analyses of Human Movement*. Human Kinetics, first edition, 2003.
- [24] Jingyong Su, Sebastian Kurtek, Eric Klassen, and Anuj Srivastava. Statistical analysis of trajectories on Riemannian manifolds: Bird migration, hurricane tracking, and video surveillance. *Annals of Applied Statistics*, 8(1), 2014.
- [25] Harry L Swinney. Observations of order and chaos in nonlinear systems. *Physica D: Nonlinear Phenomena*, 7(1):3–15, 1983.
- [26] Pavan K. Turaga, Ashok Veeraraghavan, Anuj Srivastava, and Rama Chellappa. Statistical computations on Grassmann and Stiefel manifolds for image and video-based recognition. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 33(11):2273–2286, 2011.
- [27] Oncel Tuzel, Fatih Porikli, and Peter Meer. Region covariance: A fast descriptor for detection and classification, 2006.
- [28] A. Veeraraghavan, R. Chellappa, and A. K. Roy-Chowdhury. The function space of an activity. *IEEE CVPR*, pages 959–968, 2006.
- [29] Raviteja Vemulapalli, Felipe Arrate, and Rama Chellappa. Human action recognition by representing 3d skeletons as points in a lie group. In *(CVPR), 2014*, June 2014.
- [30] Vinay Venkataraman, Pavan K. Turaga, Nicole Lehrer, Michael Baran, Thanassis Rikakis, and Steven L. Wolf. Attractor-shape for dynamical analysis of human movement: Applications in stroke rehabilitation and action recognition. In *CVPR Workshops*, 2013.
- [31] I Vlachos and D Kugiumtzis. State space reconstruction from multiple time series. In *Topics on Chaotic Systems: Selected Papers from Chaos 2008 International Conference*, page 378. World Scientific, 2009.
- [32] Garnett P Williams. *Chaos theory tamed*. Joseph Henry Press, 1997.
- [33] Andrew David Wilson and Aaron F Bobick. Learning visual behavior for gesture analysis. In *IEEE International Symposium on Computer Vision*, pages 229–234, Nov. 1995.
- [34] Alan Wolf, Jack B Swift, Harry L Swinney, and John A Vastano. Determining lyapunov exponents from a time series. *Physica D: Nonlinear Phenomena*, 16(3):285–317, 1985.
- [35] Junji Yamato, Jun Ohya, and Kenichiro Ishii. Recognizing human action in time-sequential images using hidden markov model. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 379–385, June 1992.
- [36] S. Yi and H. Krim. A subspace learning of dynamics on a shape manifold: A generative modeling approach. In *Geometric Science of Information - First International Conference, GSI*, 2013.
- [37] Sheng Yi, Hamid Krim, and Larry K. Norris. Human activity as a manifold-valued random process. *IEEE Transactions on Image Processing*, 21(8):3416–3428, 2012.