

# A Riemannian Framework for Statistical Analysis of Topological Persistence Diagrams

Rushil Anirudh\*, Vinay Venkataraman\*, Karthikeyan Natesan Ramamurthy<sup>+</sup>, Pavan Turaga\*

\*Arizona State University, <sup>+</sup> IBM T. J. Watson Research Center, NY

## Abstract

Topological data analysis is becoming a popular way to study high dimensional feature spaces without any contextual clues or assumptions. This paper concerns itself with one popular topological feature, which is the number of  $d$ -dimensional holes in the dataset, also known as the Betti- $d$  number. The persistence of the Betti numbers over various scales is encoded into a persistence diagram (PD), which indicates the birth and death times of these holes as scale varies. A common way to compare PDs is by a point-to-point matching, which is given by the  $n$ -Wasserstein metric. However, a big drawback of this approach is the need to solve correspondence between points before computing the distance; for  $n$  points, the complexity grows according to  $\mathcal{O}(n^3)$ . Instead, we propose to use an entirely new framework built on Riemannian geometry, that models PDs as 2D probability density functions that are represented in the square-root framework on a Hilbert Sphere. The resulting space is much more intuitive with closed form expressions for common operations. The distance metric is 1) correspondence-free and also 2) independent of the number of points in the dataset. The complexity of computing distance between PDs now grows according to  $\mathcal{O}(K^2)$ , for a  $K \times K$  discretization of  $[0, 1]^2$ . This also enables the use of existing machinery in differential geometry towards statistical analysis of PDs such as computing the mean, geodesics, classification etc. We report competitive results with the Wasserstein metric, at a much lower computational load, indicating the favorable properties of the proposed approach.

## 1. Introduction

Topological data analysis (TDA) has emerged as a useful tool to analyze underlying properties of data before any contextual modeling assumptions kick in. Generally speaking, TDA seeks to characterize the *shape* of high dimensional data (viewed as a point-cloud in some metric space), by quantifying various topological constructs such as connected components, high-dimensional holes, level-sets and

monotonic regions of functions defined on the data [9]. In particular, the number of  $d$ -dimensional holes in a data, also known as the Betti- $d$  number, corresponds to the rank of the  $d$ -dimensional homology group. A commonly used and powerful topological feature, the *persistence diagram* (PD) summarizes the *persistence* of the Betti numbers with respect to the *scale* parameter used to analyze the data. A typical machine learning pipeline using TDA features would first estimate the PDs from the given point cloud, and define a metric on them to compare different point clouds. The Wasserstein distance measure has become ubiquitous as a metric between such PDs, as it is stable and has a well-defined metric space associated with it [21, 8]. However, computation of the Wasserstein distance involves finding a one-to-one map between the points in one persistence diagram to those in the other, which is a computationally expensive operation.

In this paper, we propose a novel representation for persistence diagrams as points on a hypersphere, by approximating them as 2D probability density functions (pdf). We perform a square-root transform of the constructed pdf, wherein the Hilbert sphere becomes the appropriate space for defining metrics [19]. This insight is used to construct closed form metrics – geodesics on the Hilbert sphere – which can be computed very efficiently, bypassing the correspondence problem entirely. The overall pipeline of operations for computing the proposed representation is given in Figure 1.

The biggest advantage of the proposed representation is that it completely works around the computationally expensive step of obtaining one-to-one correspondences between points in persistence diagrams, thereby making the distance computation between PDs extremely efficient. We show that approximating PDs as pdfs results in comparable performances to the popular and best performing  $L_1$ -Wasserstein metric, while at the same time being orders of magnitude faster. We also provide a theoretically well-grounded understanding of the geometry of the pdf representation. Additionally, the availability of closed form expressions for many important tools such as the geodesics,

This work was supported by NSF CAREER grant 1452163.

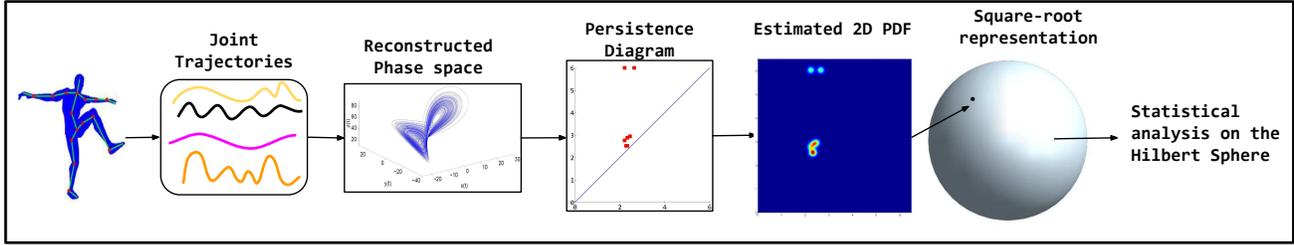


Figure 1. The overall sequence of operations leading to the proposed representation, illustrated for the application of activity analysis. Joint position data from motion capture systems are collected as 1D time series. The phase space is reconstructed from these time series via Takens’ embedding theorem. We represent the topological properties of this phase space using the persistence diagram (PD). Next, we use kernel density estimation to represent the PD itself as a 2D probability density function (pdf). Finally, we use the square-root framework to interpret these pdfs as points on a Hilbert sphere.

exponential maps etc. enables us to adapt powerful statistical tools – such as clustering, PCA, sample mean etc. – opening new possibilities for applications involving large datasets. To the best of our knowledge we are the first to propose this representation for persistence diagrams.

### Contributions

1. We present the first representation of persistence diagrams that are approximated as points on a Hilbert sphere resulting in closed-form expressions to compare two diagrams. This also completely avoids the correspondence problem, which is typically a computational bottleneck.
2. We demonstrate the ability of the proposed representation for statistical analyses, such as computing the mean persistence diagram, principal geodesic analysis (PGA), and classification using SVMs.
3. We show promising results for supervised tasks such as action recognition, and assessment of movement quality in stroke rehabilitation.
4. The space of the representation – the Hilbert sphere – is a geometrically well-understood and intuitive space, which may further promote its use in topological machine learning applications.

The rest of the paper is organized as follows. Section 2 discusses related work in more detail. Section 3 provides the necessary background on persistent homology, the space of persistence diagrams, and the square-root framework on the Hilbert space. Section 4 provides details about the proposed framework for using the new representation of persistence diagrams in a functional Hilbert space for statistical learning tasks. Section 5 describes the experiments and results. Section 6 concludes the paper.

## 2. Related Work

Persistence diagrams elegantly summarize the topological properties of point clouds, and the first algorithm for computing topological persistence was proposed in [9]. Ever since, there has been an explosion in understanding the properties of PDs, comparing PDs, and computing statistics on them. Since a PD is a multiset of off-diagonal points along with infinite number of points on the diagonal, the cost of optimal bijective assignment, also known as the Wasserstein distance, between the individual points in the PDs is a valid distance measure. The time complexity of computing the distance between two PDs with  $n$  points each is  $\mathcal{O}(n^3)$  [4]. It has been shown in [8] that the persistence diagrams of Lipschitz functions are stable with respect to  $p$ -Wasserstein distance. However, the bottleneck and  $p$ -Wasserstein distance do not allow for easy computation of statistics. Hence,  $L_p$ -Wasserstein metrics have been used to develop approaches for computing statistical summaries such as the Fréchet mean [21]. Computing Fréchet mean of PDs involves iteratively assigning points from the individual diagrams to the candidate mean diagram, recomputing the mean, and repeating till convergence. Since the mean PD obtained is not guaranteed to be unique, the authors in [16] proposed the probabilistic Fréchet mean which itself is a probability measure on the space of PDs. An investigation of the properties of the space of PDs that allow for definition of various probability measures is reported in [15], and a derivation of confidence sets that allow the separation of topological signal from noise is presented in [10].

Since operations with PDs can be computationally expensive, Bubenik proposed a new topological summary - known as the persistence landscape - derived from the PD [5]. It can be thought of as a collection of envelope functions on the points in PD based on the order of their importance. Persistence landscapes allow for fast computation of statistics since they lie in a vector space. However their practical utility has been limited since they provide decreasing importance to secondary and tertiary features in

PDs that are usually useful in terms of discriminating between data from different classes. Recently, an approach that defines a stable multi-scale kernel between persistence diagrams has been proposed [17]. The kernel is obtained by creating a surface with a Gaussian centered at each point in the PD along with a negative Gaussian centered at its reflection across the diagonal. In addition, the authors in [2] propose to compute a vector representation - the persistence image - by computing a kernel density estimate of the PD and integrating it locally. The kernel density estimate is weighted such that the points will have increasing weights as they move away from the diagonal. A similar weighting approach is used in [14] to create a persistence-weighted Gaussian kernel.

The square-root representation for 1D probability density functions (pdfs) was proposed in [19] and was used for shape classification. It has since been used in several applications, including activity recognition [22] - where the square-root velocity representation is used to model the space of time warping functions. This representation extends quite easily to arbitrary high-dimensional pdfs as well.

### 3. Mathematical Preliminaries

In this section we will introduce the concepts of a) persistence homology, b) persistent diagrams, c) our proposed representation of PDs as a 2D pdf, and the square-root framework resulting in the Hilbert sphere geometry.

#### 3.1. Persistent Homology

The homology of point cloud  $X$  can be computed by first constructing a *shape* out of the point cloud and estimating its homology. A simplex is imposed on every set of points in the point cloud that satisfy a neighborhood condition based on a scale parameter  $\epsilon$ . The collection of such simplices is known as the *simplicial complex*  $\mathcal{S}$ .  $\mathcal{S}$  can be thought of as a representation of the underlying shape of the data, and its homology can be inferred. However, homology groups obtained from  $\mathcal{S}$  depend on the scale or time parameter  $\epsilon$  based on which the complexes are constructed [9]. The homological features of the simplicial complex constructed from the point cloud that are stable across scales, i.e., that are *persistent*, are the ones that provide information about the underlying shape. Topological features that do not persist are considered to be noise. This information is represented using *persistence diagrams* as a 2D plot of birth versus death times of each homology cycle corresponding to the homology group  $\mathbb{H}_k, k = \{0, 1, \dots\}$ . The birth time is the scale at which the homological feature is born and the death time is the scale at which it ceases to exist. The homology cycle of dimension  $d$  is also referred to as a  $d$ -dimensional hole. Therefore, the PD can be considered as an object that represents the number of holes in terms of the range of scales at which they appear. Typically, PDs of the point cloud

data are obtained using *filtration* of simplicial complexes. A well-known filtration is the Vietoris-Rips (VR) filtration, where a simplex is induced between a group of points whenever the distance between each pair of points is less than the given scale  $\epsilon$  [26]. An example of point clouds and their corresponding persistence diagrams for homology groups 0 and 1 are provided in Figure 2.

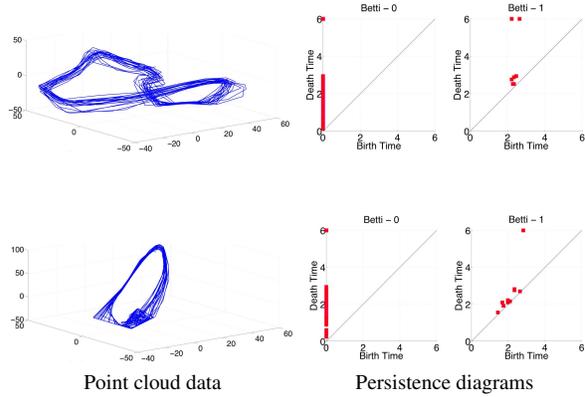


Figure 2. The above example illustrates the topological features extracted from the point cloud data with two and one one-dimensional holes. These properties are reflected well in their corresponding persistence diagrams on the right.

#### 3.2. The Space of Persistence Diagrams

Every PD is a multiset of 2D points, where each point denotes the birth and death time of a homology group. Furthermore, the diagonal is assumed to contain an infinite number of points with the same birth and death times. For any two PDs  $X$  and  $Y$ , the distance between the diagrams is usually quantified using the bottleneck distance or the  $L_q$ -Wasserstein metric [13]. In this paper, we consider only the  $L_2$ - and  $L_1$ -Wasserstein ( $d_{L_2}$  and  $d_{L_1}$ ) metrics given respectively as,

$$d_{L_2}(X, Y) = \left( \inf_{\eta: X \rightarrow Y} \sum_{x \in X} \|x - \eta(x)\|_2^2 \right)^{\frac{1}{2}}, \quad (1)$$

and

$$d_{L_1}(X, Y) = \inf_{\eta: X \rightarrow Y} \sum_{x \in X} \|x - \eta(x)\|_1. \quad (2)$$

Since each diagram contains an infinite number of points in the diagonal, this distance is computed by pairing each point in one diagram uniquely to another non-diagonal or diagonal point in the other diagram, and then computing the distance. This correspondence can be obtained via the Hungarian algorithm or its more efficient variants [13].

The space of PDs with respect to the  $L_2$ -Wasserstein metric is given by

$$\mathcal{D}_{L_2} = \{X | d_{L_2}(X, \emptyset) < \infty\}. \quad (3)$$

Turner *et. al.* [21] show this is a non-negatively curved Alexandrov space. Furthermore, the diagram on the geodesic between the PDs  $X$  and  $Y$  in this space is given as

$$\gamma(s) = (1 - s)x + s\phi(x), \quad (4)$$

where  $x$  is a point in the diagram  $X$ ,  $\phi(x)$  is a corresponding point in the diagram  $Y$  and  $s \in [0, 1]$  parametrizes the geodesic. Clearly, the points in the diagram on the geodesic can be simply obtained by linearly interpolating between the corresponding points on the candidate points  $X$  and  $Y$ . Furthermore, the Riemannian mean between two PDs is easily computed as  $\gamma(0.5)$ .

### 3.3. Square-root Framework and the Hilbert Sphere

We treat the points in a persistence diagram as samples from an underlying probability distribution function. This representation is inspired from recent work dealing with defining Mercer kernels on PDs [17], where the feature embedding obtained from the Mercer kernel closely resembles a kernel-density estimator from the given points, with an additional reflection term about the diagonal to account for boundary effects while solving a heat-equation. The fact that the feature embedding resembles a kernel density estimate is not further exploited in [17].

In our work, we more directly exploit this pdf interpretation of PDs, further endowing it with a square-root form [19] – and then utilizing the Hilbert spherical geometry that results.

Without loss of generality, we will assume that the supports for each 2D pdf is  $[0, 1]^2$ . The space of pdfs that we will consider is

$$\mathcal{P} = \{p : [0, 1] \times [0, 1] \rightarrow \mathbb{R} \ \forall x, y | p(x, y) \geq 0, \text{ and } \int_0^1 \int_0^1 p(x, y) dx dy = 1\} \quad (5)$$

It is well-known that  $\mathcal{P}$  is not a vector space [19]. Instead, it is a Riemannian manifold with the Fisher-Rao metric as the natural Riemannian metric. Geodesics under the Fisher-Rao metric are quite difficult to compute. Instead, we adopt a square-root form proposed by Srivastava *et. al.* [19] which simplifies further analysis enormously. In other words we consider the space,

$$\Psi = \{\psi : [0, 1] \times [0, 1] \rightarrow \mathbb{R} | \psi \geq 0, \int_0^1 \int_0^1 \psi^2(x, y) dx dy = 1\}. \quad (6)$$

For any two tangent vectors  $v_1, v_2 \in T_\psi(\Psi)$ , the Fisher-Rao metric is given by,

$$\langle v_1, v_2 \rangle = \int_0^1 \int_0^1 v_1(x, y) v_2(x, y) dx dy. \quad (7)$$

The above two pieces of information imply that the square-root form  $\psi = \sqrt{p}$  results in the space becoming a unit Hilbert-sphere, endowed with the usual inner-product metric. Geodesics on the unit-Hilbert sphere under the above Riemannian metric are known in closed form. In fact, the differential geometry of the Hilbert sphere results in closed form expressions for computing geodesics, exponential and inverse exponential maps [19]. Further, the square-root form with the above metric has additional favorable properties such as invariance to domain re-parameterization.

Given two points  $\psi_1, \psi_2$  the geodesic distance between them is given by

$$d_H(\psi_1, \psi_2) = \cos^{-1}(\langle \psi_1, \psi_2 \rangle), \quad (8)$$

where  $\langle \psi_1, \psi_2 \rangle_\psi = \int_0^1 \psi_1(t) \psi_2(t) dt$ . The exponential map is defined as,

$$\exp_{\psi_i}(v) = \cos(\|v\|_{\psi_i}) \psi_i + \sin(\|v\|_{\psi_i}) \frac{v}{\|v\|_{\psi_i}}, \quad (9)$$

where  $v \in T_{\psi_i}(\Psi)$  is a tangent vector at  $\psi_i$  and  $\|v\|_{\psi_i} = \sqrt{\langle v, v \rangle_{\psi_i}}$ . The logarithmic map from  $\psi_i$  to  $\psi_j$  is given by:

$$\exp_{\psi_i}^{-1}(\psi_j) = \frac{u}{\sqrt{\langle u, u \rangle}} \cos^{-1} \langle \psi_i, \psi_j \rangle, \quad (10)$$

with  $u = \psi_i - \langle \psi_i, \psi_j \rangle \psi_j$ . The geodesic on the sphere is given in closed form as

$$\pi(s) = \frac{(1 - s)\psi_1 + s\psi_2}{s^2 + (1 - s)^2 + 2s(1 - s)\langle \psi_1, \psi_2 \rangle}, \quad (11)$$

or equivalently as

$$\pi(s) = \cos(s\|v\|) \psi + \sin(s\|v\|) \frac{v}{\|v\|}. \quad (12)$$

A comparison of sampling of the geodesics in the Alexandrov space induced by the  $L_2$ -Wasserstein metric and the proposed Hilbert sphere are provided in Figure 3.

## 4. Algorithmic Details

Our proposed framework consists of reconstructing the phase space of the time series data, computing the PDs of the reconstructed phase space, transforming each PD to a 2D pdf by kernel density estimation, using the 2D pdfs in the square-root framework to estimate distances or obtain statistical summaries. The distances computed and the features obtained will be used in inference tasks. Additional details for some of the above steps are given below.

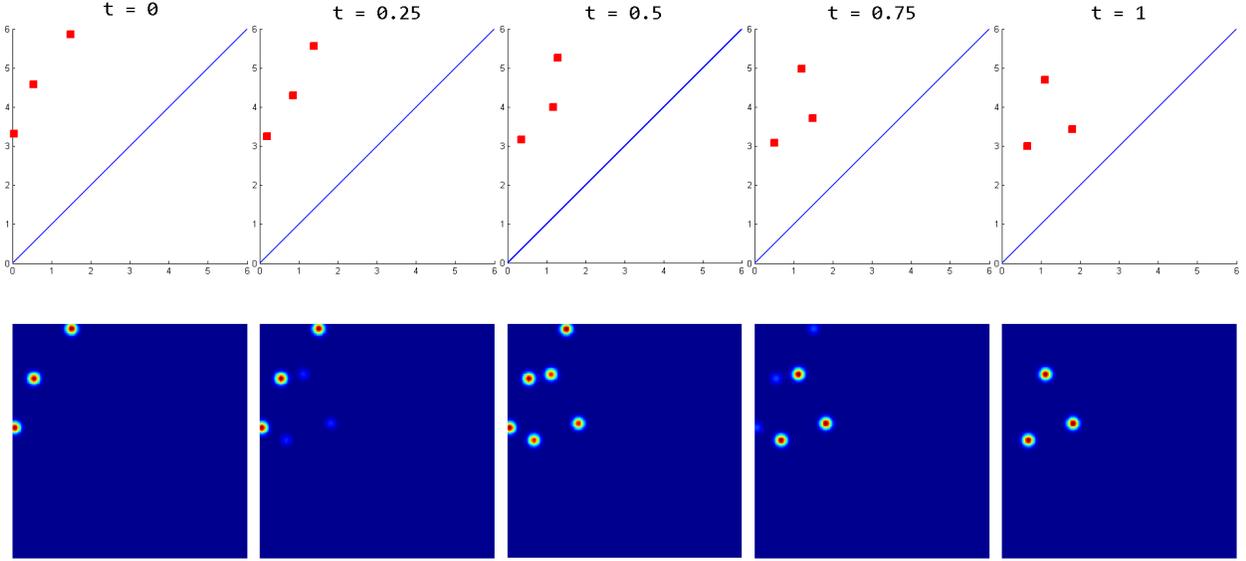


Figure 3. Comparing geodesics in the Alexandrov space (top) and the proposed Hilbert sphere (bottom), for a simple persistence diagram containing just 3 points. The Alexandrov space requires computing correspondence between points, and the geodesic involves linearly interpolating between corresponding points. The persistence pdf avoids correspondence estimation, hence the geodesics correspond to new local modes appearing and gradually increasing in intensity as the original modes decrease in intensity.

**Phase-space Reconstruction from Activity Data** In dynamical system understanding, we aim to estimate the underlying states, but we measure functions – usually projections – of the original state-space. e.g. while human movement is influenced by many joints, ligaments, muscles of the body, we might measure the location of only a few joints. One way to deal with this issue is to reconstructing the ‘phase space’ by the method of delays used commonly in non-linear dynamical analysis [20]. Reconstructing the phase-space in this manner only preserves important topological properties of the original dynamical system, and does not recover the true state-space. For a discrete dynamical system with a multidimensional phase-space, time-delay vectors (or embedding vectors) are obtained by the concatenation of delayed versions of the data points,

$$\mathbf{x}_t = [x(t), x(t + \tau), \dots, x(t + (m - 1)\tau)]^T. \quad (13)$$

where  $m$  is the embedding dimension,  $\tau$  is the embedding delay. For a sufficiently large  $m$ , the important topological properties of the unknown multidimensional system are reproduced in the reconstructed phase-space [1]. The collection of  $n$  time-delay vectors is the point cloud data under further consideration, and this is represented as  $\mathbf{X} = [\mathbf{x}_t]_{t=0}^n$ .

**Estimating the PDs** After estimating the point cloud in the reconstructed phase space, we will construct a simplicial complex  $\mathcal{S}$  using the point cloud data  $\mathbf{X}$  to compute the persistence diagrams for the Betti numbers using the

VR filtration. However, this approach considers only spatial adjacency between the points and ignores the temporal adjacency. We improve upon the existing VR filtration approach by explicitly creating temporal links between  $\mathbf{x}_{t-1}$ ,  $\mathbf{x}_t$ , and  $\mathbf{x}_{t+1}$  in the one-skeleton of  $\mathcal{S}$ , thereby creating a metric space which encodes adjacency in both space and time [23]. The persistence diagrams for homology groups of dimensions 0 and 1 are then estimated.

**Square-root Framework for Distance Estimation between PDs** Since each PD is a multiset of points in the 2D space, we start by constructing a 2D pdf from each of them using kernel density estimation using a Gaussian kernel of zero mean and variance  $\sigma^2$ . For each PD, we compute the square-root representation  $\psi$  using the approach provided in Section 3.3, and the distance between two PDs can be computed using (8).

**Dimensionality Reduction with Principal Geodesic Analysis (PGA)** We are able to extract principal components using PGA [11], adapted to our Hilbert sphere – which essentially performs classical PCA on the tangent space of the mean-point on the sphere. Given a dataset of persistence diagrams in their square-root form  $\{\psi_1, \psi_2, \dots, \psi_N\}$ , we first compute the Riemannian center of mass [12]. We use a simple approximation using an extrinsic algorithm that computes first the Euclidean mean and maps it on to the closest point on the sphere. Next, we represent each  $\psi_i$  by its tangent vector from the mean. We

then compute the principal components on the tangent space and represent a persistence diagram as a low-dimensional vector.

## 5. Experiments

We perform experiments on two datasets for human action analysis. First we perform action recognition on the MoCap dataset [3], next we show the use of the framework in quality assessment of movements in stroke rehabilitation [7]. We will describe the datasets next, followed by the evaluation settings and parameters used. In all our experiments, we performed a discretization of the 2D pdf into a  $K \times K$  grid. The choice of  $K$  indirectly affects classification performance as expected, i.e. – a smaller value of  $K$  results in a reduced ability to identify between nearby points in the PD due to lower resolution. On the other hand, a larger value of  $K$  improves resolution, but also increases computational requirements. Typical values of  $K$  used in experiments range from 50 to 100 at most, whereas typical values of  $n$  – the number of points in a PD – ranges from 20 to 50.

### 5.1. Motion Capture Data

We evaluate the performance of the proposed framework using 3-dimensional motion capture sequences of body joints [3]. The dataset is a collection of five actions: *dance*, *jump*, *run*, *sit* and *walk* with 31, 14, 30, 35 and 48 instances respectively. We generate 100 random splits having 5 test-

Method	Accuracy (%)	Time (sec)
Chaos [3]	52.44	-
VR-Complex [26]	93.68	-
DT2 [24]	93.92	-
T-VR Complex (L1) [23]	96.48	$(1.2 \pm 1.23) \times 10^3$
<b>Proposed</b> - 1NN	89.87	<b><math>(0.059 \pm 0.044)</math></b>
<b>Proposed</b> - PGA +SVM	91.68	-

Table 1. Comparison of classification rates for different methods using nearest neighbor classifier on the motion capture dataset. It is observed that on an average, the proposed metric is  $10^5$  times faster than the traditional Wasserstein metric, while achieving a comparable recognition accuracy.

ing examples from each action class and use an SVM classifier on the vector features computed with PGA, we get a performance of 91.68%. The mean recognition rates for the different methods are given in Table 1. We also compare with a 1-nearest neighbor classifier computed on the Hilbert sphere, which gives a performance of 89.87%. This is clearly competitive, when compared to approaches that use the  $L_1$ -Wasserstein metric as shown in table 1. Clearly,

the proposed metric for persistence diagrams is able to capture most of the important information, with the added benefit being free from expensive correspondence estimation. To compute the times taken using the Wasserstein metric and the proposed metric, we average across 3 samples from 5 action classes, across 57 (19 joint trajectories along 3 axes) trajectories – a total of 855 persistence diagrams. We computed a pairwise distance matrix for all these PDs, and computed the average time taken in Matlab 2014 on a standard Intel i7 PC.

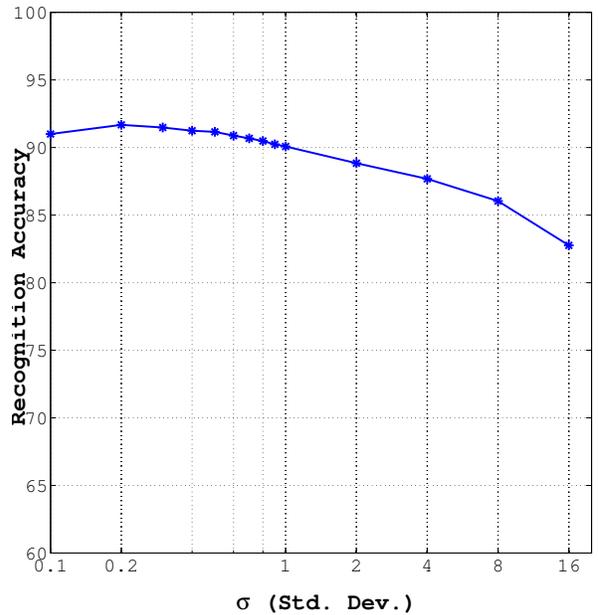


Figure 4. Recognition accuracy vs  $\sigma$ , the standard deviation of the 2D Gaussians used in the kernel density estimates.

In Figure 4, we compare the recognition accuracy with the choice of  $\sigma$ , the standard deviation used during kernel density estimation. The accuracy is generally higher for a small  $\sigma$  and drops as the  $\sigma$  increases. We note that a similar trend is also reported in [17].

### 5.2. Stroke Rehabilitation Dataset

Our aim in this experiment is to quantitatively assess the quality of movement performed by the impaired subjects during repetitive task therapy. The experimental data was collected using a heavy marker-based system (14 markers on the right hand, arm and torso) in a hospital setting from 15 impaired subjects performing reach and grasp movements to a defined target. The stroke survivors were also evaluated by the Wolf Motor Function Test (WMFT) [25] on the day of recording, which evaluates a subject’s functional ability on a scale of 1 – 5 (with 5 being least impaired and 1 being most impaired) based on predefined functional tasks. In our experiments, we only use the data corre-

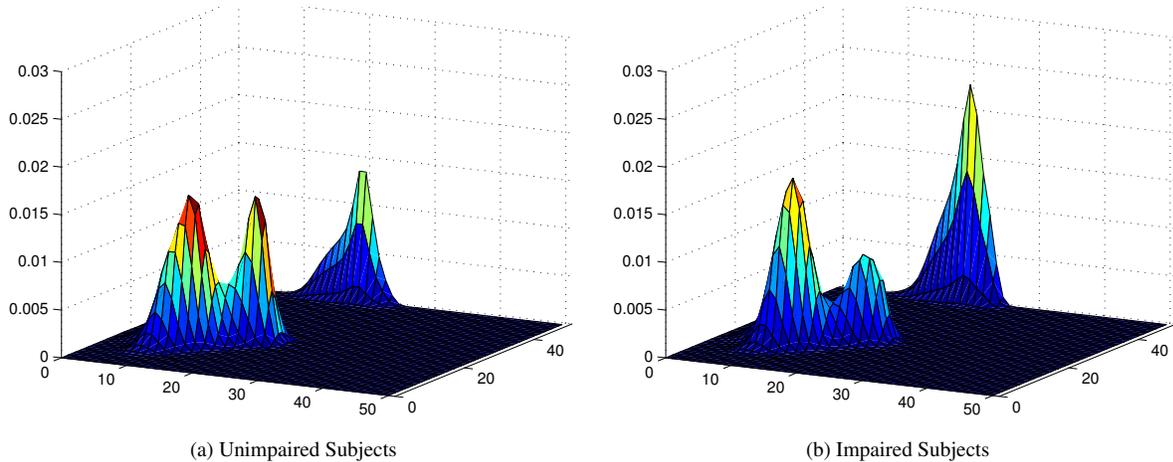


Figure 5. The mean persistence pdfs visualized as heat maps for unimpaired subjects (left) and stroke survivors (right), for the reach-and-grasp action, are shown. These means are computed as the extrinsic means (we choose extrinsic mean here for simplicity) on the Hilbert sphere using the proposed representation. The locations of the peaks are exactly the same, since the subjects perform the same general movement, however the intensity differs significantly indicating that the *quality* of movement is captured using topological tools well.

sponding to the single marker on the wrist from the heavy marker-based hospital system, which allows us to evaluate the framework in a home-based rehabilitation system setting. We utilize the WMFT scores as an approximate high-level quantitative measure for movement quality (ground-truth labels) of impaired subjects performing the reach and grasp movements.

We compute explicit features as described earlier using PGA. This allows us to represent each movement by a low-dimensional feature vector. We use the WMFT scores and split the dataset into train and test to perform regression. Since the dataset is small, we use a leave-one-out approach, where we train on all but one sample, which is used for testing, and repeat this over all the samples in the dataset. The correlation performance with the WMFT score is shown in Table 2, and it can be seen that we are comparable to the state of the art, and much better than traditional features. The predicted scores are shown in Figure 6 compared to the original scores.

The dynamical features proposed in [24] and [3], depend on describing the phase space for each movement. We compute the topological features, which are expected to be much weaker than other characteristics such as shape. However, we are still able to capture subtle information regarding movement quality. This is illustrated in Figure 5, where we see the 3D peaks associated with the average of persistence diagrams across subjects impaired with stroke and those who are unimpaired. It is clearly seen that since they are performing the same kind of movements, the peaks occur at the same location. However, the intensity is significantly different across these diagrams, perhaps indicating information regarding movement quality.

Method	Correlation with WMFT
KIM (14 markers) [6]	0.85
Lyapunov exponent (1 marker) [18]	0.60
<b>Proposed method (1 marker)</b>	<b>0.80</b>

Table 2. Comparison of correlation coefficient for different methods using leave-one-subject-out cross-validation scheme and SVM regressor on the stroke rehabilitation dataset. Even with just a single marker, We obtain comparable results to a clinical 14-marker system.

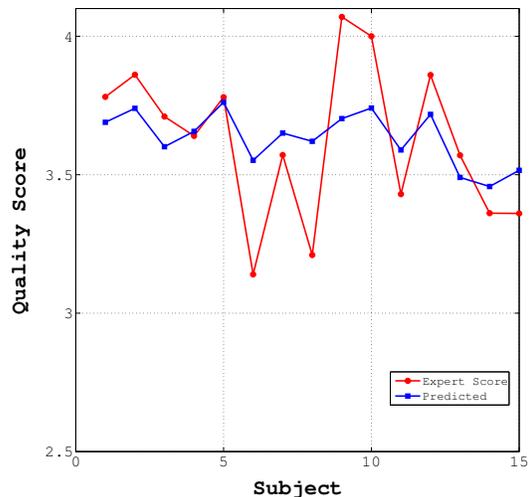


Figure 6. Predicting movement quality scores for reach-and-grasp in stroke survivors using topological features. We obtain a **0.8** correlation score with a p-value of  $5.46 \times 10^{-4}$

## 6. Discussion and Conclusion

Based on the theory and experiments presented so far, it is instructive to compare the space of PDs with respect to the different distance metrics. We will consider only the  $L_2$ -Wasserstein metric ( $d_{L_2}$ ) and the proposed Hilbert sphere metric  $d_H$ . The interpretation of PDs with respect to these two metrics is very different.  $d_{L_2}(X, Y)$  is the Earth Mover's Distance between the PDs  $X$  and  $Y$  considered as scaled discrete probability mass functions (pmfs). However,  $d_H(\psi_X, \psi_Y)$  is the geodesic distance between the square root of kernel density estimates of  $X$  and  $Y$ . Furthermore, the "length" of the persistence diagrams induced by these metrics are very different. For the Wasserstein metric, this is given by  $d_{L_2}(X, \emptyset)$  which can be arbitrarily high, whereas with the proposed metric it is constrained as the persistence diagrams live on a unit sphere. The most important distinction arises when the pairwise distances between all the points in the two PDs are sufficiently high. When the 2-dimensional pdfs obtained from kernel density estimates of the two PDs do not overlap anywhere,  $d_H(X, Y) \rightarrow 1$ . This implies that if the variance of the kernel is sufficiently small, two PDs with non-overlapping points will always have  $d_H$  close to 1. This problem can be alleviated by using kernels with multiple scales for smoothing the PDs to obtain and combining the distances obtained at each scale.

The comparison between the geodesics in the persistence space is also illuminating (Figure 3). Sampling of the geodesic in the Alexandrov space shows that the points in one diagram move towards the corresponding points in the other, as we move along the geodesic. Whereas, a similar sampling in the Hilbert sphere shows that the PDs in the middle of the geodesic contain the modes from pdfs corresponding to both the candidate PDs. This results in markedly different interpretations of the means computed from  $d_{L_2}$  and  $d_H$  even for the case of two persistence diagrams.

Several directions of future work emerge from this endeavor. On the theoretical side a more formal understanding of the variety of metrics and their relationship to the proposed one can be considered. On the applied side, the favorable computational load reduction should open new applications of TDA for massive datasets.

## References

- [1] H. D. Abarbanel. *Analysis of observed chaotic data*. New York: Springer-Verlag, 1996. 5
- [2] H. Adams, S. Chepushtanova, T. Emerson, E. Hanson, M. Kirby, F. Motta, R. Neville, C. Peterson, P. Shipman, and L. Ziegelmeier. Persistent images: A stable vector representation of persistent homology. *ArXiv preprint arXiv:1507.06217*, 2015. 3
- [3] S. Ali, A. Basharat, and M. Shah. Chaotic invariants for human action recognition. In *IEEE International Conference on Computer Vision*, pages 1–8, Oct. 2007. 6, 7
- [4] D. P. Bertsekas. A new algorithm for the assignment problem. *Mathematical Programming*, 21(1):152–171, 1981. 2
- [5] P. Bubenik. Statistical topological data analysis using persistence landscapes. *The Journal of Machine Learning Research*, 16(1):77–102, 2015. 2
- [6] Y. Chen, M. Duff, N. Lehrer, H. Sundaram, J. He, S. L. Wolf, and T. Rikakis. A computational framework for quantitative evaluation of movement during rehabilitation. In *AIP Conference Proceedings-American Institute of Physics*, volume 1371, pages 317–326, 2011. 7
- [7] Y. Chen, M. Duff, N. Lehrer, H. Sundaram, J. He, S. L. Wolf, T. Rikakis, T. D. Pham, X. Zhou, H. Tanaka, et al. A computational framework for quantitative evaluation of movement during rehabilitation. In *AIP Conference Proceedings-American Institute of Physics*, volume 1371, page 317, 2011. 6
- [8] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and Y. Mileyko. Lipschitz functions have  $l_p$ -stable persistence. *Foundations of Computational Mathematics*, 10(2):127–139, 2010. 1, 2
- [9] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28(4):511–533, 2002. 1, 2, 3
- [10] B. T. Fasy, F. Lecci, A. Rinaldo, L. Wasserman, S. Balakrishnan, A. Singh, et al. Confidence sets for persistence diagrams. *The Annals of Statistics*, 42(6):2301–2339, 2014. 2
- [11] P. T. Fletcher, C. Lu, S. M. Pizer, and S. C. Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging*, 23(8):995–1005, August 2004. 5
- [12] K. Grove and H. Karcher. How to conjugate  $C^1$ -close group actions. *Math.Z.*, 132:11–20, 1973. 5
- [13] M. Kerber, D. Morozov, and A. Nigmatov. Geometry helps to compare persistence diagrams. In *Proceedings of the Eighteenth Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 103–112, 2016. 3
- [14] G. Kusano, K. Fukumizu, and Y. Hiraoka. Persistence weighted gaussian kernel for topological data analysis. *arXiv preprint arXiv:1601.01741*, 2016. 3
- [15] Y. Mileyko, S. Mukherjee, and J. Harer. Probability measures on the space of persistence diagrams. *Inverse Problems*, 27(12):124007, 2011. 2
- [16] E. Munch, P. Bendich, K. Turner, S. Mukherjee, J. Mattingly, and J. Harer. Probabilistic Fréchet means and statistics on vineyards. *ArXiv preprint arXiv:1307.6530*, 2013. 2
- [17] J. Reininghaus, S. Huber, U. Bauer, and R. Kwitt. A stable multi-scale kernel for topological machine learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4741–4748, 2015. 3, 4, 6
- [18] M. Rosenstein, J. Collins, and C. De Luca. A practical method for calculating largest lyapunov exponents from small data sets. *Physica D: Nonlinear Phenomena*, 65(1):117–134, 1993. 7

- [19] A. Srivastava, I. Jermyn, and S. Joshi. Riemannian analysis of probability density functions with applications in vision. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007. 1, 3, 4
- [20] F. Takens. Detecting strange attractors in turbulence. *Dynamical Systems and Turbulence*, 898:366–381, 1981. 5
- [21] K. Turner, Y. Mileyko, S. Mukherjee, and J. Harer. Fréchet means for distributions of persistence diagrams. *Discrete & Computational Geometry*, 52(1):44–70, 2014. 1, 2, 4
- [22] A. Veeraraghavan, A. Srivastava, A. K. Roy-Chowdhury, and R. Chellappa. Rate-invariant recognition of humans and their activities. *Image Processing, IEEE Transactions on*, 18(6):1326–1339, 2009. 3
- [23] V. Venkataraman, K. N. Ramamurthy, and P. Turaga. Persistent homology of attractors for action recognition. *arXiv preprint arXiv:1603.05310*, 2016. 5, 6
- [24] V. Venkataraman and P. Turaga. Shape distributions of non-linear dynamical systems for video-based inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016 (accepted) ; arXiv:1601.07471. 6, 7
- [25] S. L. Wolf, P. A. Catlin, M. Ellis, A. L. Archer, B. Morgan, and A. Piacentino. Assessing wolf motor function test as outcome measure for research in patients after stroke. *Stroke*, 32(7):1635–1639, 2001. 6
- [26] A. Zomorodian. Fast construction of the Vietoris-Rips complex. *Computers & Graphics*, 34(3):263–271, 2010. 3, 6