

Riemannian Time Warping: Multiple Sequence Alignment in Curved Spaces

Julian Richter^{1,2}, Christopher A. Erdős¹, Christian Scheurer¹, Jochen J. Steil² and Niels Dehio¹




Fig. 1. Schematic overview of RTW: our approach enables the temporal alignment of multiple signals, represented through data on a Riemannian manifold \mathcal{M} . Given N signals \mathbf{x}_n with temporal shift (color encoding), the exponential and logarithmic maps are utilized to apply time warping within the local tangent spaces \mathcal{T} . RTW utilizes windowed sinc interpolation (bold segments) and parameterized warping functions $\Omega = \{\omega_n\}_n^N$, to temporally shift the data within these tangent spaces, i.e., \mathbf{u}_n becomes $\hat{\mathbf{u}}_n$. Finally, the mean signal $\hat{\mathbf{y}}$ is computed from all warped signals $\hat{\mathbf{x}}_n$. Gradient-based optimization is utilized to find warping functions that minimize the distance between all warped signals and their mean. See subsection IV-D for further remarks on this figure.

Abstract—Temporal alignment of multiple signals through time warping is crucial in many fields, such as classification within speech recognition or robot motion learning. Almost all related works are limited to data in Euclidean space. Although an attempt was made in 2011 to adapt this concept to unit quaternions, a general extension to Riemannian manifolds remains absent. Given its importance for numerous applications in robotics and beyond, we introduce Riemannian Time Warping (RTW). This novel approach efficiently aligns multiple signals by considering the geometric structure of the Riemannian manifold in which the data is embedded. Extensive experiments on synthetic and real-world data, including tests with an LBR iiwa robot, demonstrate that RTW consistently outperforms state-of-the-art baselines in both averaging and classification tasks.

Index Terms—AI-Based Methods, Learning from Demonstration, Multiple Sequence Alignment, Riemannian Manifold

I. INTRODUCTION

Time warping describes the temporal alignment of several signals – a fundamental operation for processing sequential data commonly used in speech recognition [1], [2] and human action recognition [3], [4]. In robotics, temporal alignment is crucial for learning movement primitives: it helps significantly to extract a consistent motion pattern from multiple human demonstrations that are naturally provided with temporal variations. Given multiple signals with varying lengths describing

the same underlying process, each signal is warped with the goal to improve their temporal alignment [5], [6].

Dynamic Time Warping (DTW) [7] finds optimal alignment through Dynamic Programming [8] but suffers from exponential complexity that prevents application to larger datasets. Alternative algorithms sacrifice global optimality to accelerate the search process by several orders of magnitude [9]–[11].

Almost all related works focus on Euclidean data. However, many robotics problems involve temporal phenomena on Riemannian manifolds (cf. Figure 1). For example, demonstrated trajectories naturally lie on the unit sphere \mathbb{S}^3 for orientations or on symmetric positive definite matrices \mathbb{S}_{++}^D describing a D -dimensional stiffness or inertia. Applying Euclidean time warping can lead to geometrically invalid interpolations (e.g., averaging quaternions in Euclidean space yields non-unit vectors) and suboptimal alignment that ignores geometric constraints [12]. While Quaternion Dynamic Time Warping (QDTW) [13] extends DTW to unit quaternions, it retains exponential complexity. To the best of our knowledge, no generic time warping method handles structured non-Euclidean data, partly because sum and scalar operations used in related works are not defined on Riemannian manifolds [14].

As our main contribution, we present *Riemannian Time Warping (RTW)* for efficient temporal alignment of multiple signals with varying lengths while explicitly considering the Riemannian manifold \mathcal{M} on which the data resides. We evaluate RTW on datasets in \mathbb{R}^D , \mathbb{S}^D and \mathbb{S}_{++}^D . RTW consistently outperforms state-of-the-art methods [10], [11] w.r.t. alignment quality and mean representation at linear complexity, and provides significant benefits for learning movement primitives. We validate our approach on real-robot experiments using a

Manuscript received: April 23, 2025; Revised: July 11, 2025; Accepted: August 5, 2025

This paper was recommended for publication by Editor Jens Kober upon evaluation of the Associate Editor and Reviewers' comments.

¹Technology and Innovation Center (TIC), KUKA Deutschland GmbH, Germany, e-mail: Julian.Richter@kuka.com

²Institute of Robotics and Process Control (IRP), Technische Universität Braunschweig, Germany

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

KUKA LBR iiwa manipulator in $\mathbb{R}^3 \times \mathbb{S}^3$. A video is available at https://youtu.be/v_k2LwEDH9I.

Notation: In the remainder, we will use square bracket notation for discrete values, e.g., $\mathbf{x}[t] \in \mathbb{R}^D$, $\mathbf{x}[t] \in \mathcal{M}$, and round bracket notation for continuous functions, e.g., $\phi(z)$.

II. RELATED WORK

Most related works focus on aligning $N = 2$ signals of length T . Despite its age, DTW [7] remains widely used due to its simplicity and effectiveness. Extending DTW to $N > 2$ signals introduces significant computational challenges. Multiple Multi-Dimensional Dynamic Time Warping (MMDDTW) [15] computes the distance matrix over all N dimensions. While globally optimal through Dynamic Programming, it becomes unfeasible for $N \gg 2$ due to complexity $\mathcal{O}(T^N)$. The Non-Linear Alignment and Averaging Framework (NLAFF) [16] iteratively aligns pairs of signals while adjusting previous solutions, with complexity $\mathcal{O}(NT^2)$.

Generalized Time Warping (GTW) [9], Trainable Time Warping (TTW) [10], and Neural Time Warping (NTW) [11] are newer methods that align several signals in linear time. GTW uses pre-defined basis functions, which is cumbersome for complex warplings. TTW and NTW transform discrete signals into continuous time via sinc interpolation, apply an index shift, and sample back at Z uniformly distributed time points. TTW estimates discrete sine transform coefficients, while NTW uses a neural network. Both outperform GTW [10], [11].

To the best of our knowledge, QDTW [13] is the only warping scheme that explicitly deals with a Riemannian manifold, focusing exclusively on unit quaternions. However, in recent years, Riemannian manifolds have gained popularity in various domains – especially in robotics [17]–[20], as they provide a convenient way to generalize methods originally developed for the Euclidean space. This letter presents the first approach for time warping on generic Riemannian manifolds, that applies to datasets with $N \gg 2$. Table I compares related works.

III. TIME WARPING IN THE CONTINUOUS TIME DOMAIN

Here we first introduce a simplified version of our approach dedicated to Euclidean data $\mathbf{x} \in \mathbb{R}^D$. The full *Riemannian Time Warping (RTW)* algorithm considering curved spaces with $\mathbf{x} \in \mathcal{M}$ is then proposed in the next section.

A. Problem Statement

Consider N discrete signals $\mathbf{X} = \{\mathbf{x}_n\}_{n=1}^N$ of lengths T_n , i.e. $\mathbf{x}_n = \{\mathbf{x}_n[t]\}_{t=1}^{T_n}$, where $\mathbf{x}_n[t] \in \mathbb{R}^D$. Signals

with temporal modifications $\hat{\mathbf{X}} = \{\hat{\mathbf{x}}_n\}_{n=1}^N$ of new length $Z \geq T_{\max} = \max\{T_n\}_{n=1}^N$, with $\hat{\mathbf{x}}_n = \{\hat{\mathbf{x}}_n[z]\}_{z=1}^Z$ and $\hat{\mathbf{x}}_n[z] \in \mathbb{R}^D$ are obtained by applying a set of parameterized warping functions $\Omega = \{\omega_n : \{1, \dots, Z\} \mapsto [0, 1]\}_{n=1}^N$. The *Multiple Sequence Alignment* problem consists of

$$\Omega_{opt} = \underset{\Omega}{\operatorname{argmin}} \mathcal{D}_{\Omega}(\hat{\mathbf{X}}) \quad \text{s.t.} \quad (2a), (2b), (2c), \quad (1)$$

where the warping functions ω_n must satisfy three constraints

$$\text{Boundary Constraint:} \quad \omega_n[1] = 0, \omega_n[Z] = 1 \quad (2a)$$

$$\text{Continuity Constraint:} \quad \omega_n[z + 1] - \omega_n[z] \leq \frac{1}{T_{\max}} \quad (2b)$$

$$\text{Monotonicity Constraint:} \quad \omega_n[z] \leq \omega_n[z + 1] \quad (2c)$$

and $\mathcal{D}_{\Omega}(\hat{\mathbf{X}})$ is a distance measure between all warped signals $\hat{\mathbf{x}}_n$ and their mean

$$\hat{\mathbf{y}} = \frac{1}{N} \sum_{n=1}^N \hat{\mathbf{x}}_n. \quad (3)$$

The challenge is to find optimal and feasible warping functions efficiently. This is non-trivial, as \mathbf{X} is a set of discrete signals.

B. Proposed Time Warping Algorithm in \mathbb{R}^D

We transform \mathbf{X} into the continuous time domain to compute the aligned signals $\hat{\mathbf{X}}$ through sinc interpolation³

$$\hat{\mathbf{x}}_n[z] = \sum_{m=\lfloor \omega_n[z]T_n \rfloor - \nu}^{\lfloor \omega_n[z]T_n \rfloor + \nu} \mathbf{x}_n[m] \operatorname{sinc}(m - \omega_n[z]T_n), \quad (4)$$

where $\lfloor \cdot \rfloor$ denotes the floor-function and $\nu \in \mathbb{N}_0$ is the window size. In contrast to NTW [11], we adapt a windowed sinc interpolation⁴ as in TTW [10], resulting in a computational complexity of $\mathcal{O}(NZ)$. Since the amplitudes in the sinc function decay over time, $\nu = 10$ captures more than 99% of its power [10]. Similar to NTW [11], we utilize N orthogonal basis vectors $\mathbf{e}_k \in \mathbb{R}^N$ that are derived from the Gram-Schmidt process or QR-decomposition with $\mathbf{e}_1 = \frac{1}{\sqrt{N}} [1, \dots, 1]^T$, and model the warping functions ω_n through

$$\omega_n[z] = \hat{z}\sqrt{N}\mathbf{e}_1 + \hat{z}(1 - \hat{z}) \sum_{k=1}^{N-1} [\phi_{\theta}(\hat{z})]_k \mathbf{e}_{k+1}, \quad (5)$$

where the θ -parameterized function $\phi_{\theta} : [0, 1] \mapsto \mathbb{R}^{N-1}$ is modeled as a neural network, and $\hat{z} = \frac{z-1}{Z-1}$. The boundary constraint (2a) is guaranteed by design. During warping, indices of a signal may be duplicated to improve the temporal alignment, requiring a longer signal, i.e., $Z \geq T_{\max}$. Hence, the continuity constraint (2b) is fulfilled with $Z = NT_{\max}$ [11]. To enforce the monotonicity constraint (2c), a λ -weighted penalty term is added to the loss as in [11]

$$\Omega_{opt} = \underset{\Omega}{\operatorname{argmin}} \left(\mathcal{D}_{\Omega}(\hat{\mathbf{X}}) + \lambda p(\Omega) \right), \quad (6)$$

$$p(\Omega) = \sum_{n=1}^N \sum_{z=1}^{Z-1} \max(\omega_n[z] - \omega_n[z + 1], 0). \quad (7)$$

³Sinc interpolation provides optimal reconstruction for signals sampled at or above Nyquist rate, i.e., twice the maximum frequency of the data [21].

⁴We apply clamping to always receive valid indices for the full interpolation window. Given a data point $\mathbf{x}_n[m]$, we set $m = \max(1, \min(T_n, m))$.

TABLE I
COMPARISON BETWEEN STATE-OF-THE-ART TIME WARPING METHODS

Method	N	\mathbb{S}^3	\mathcal{M}	Complexity
DTW [7]	= 2	×	×	$\mathcal{O}(T^2)$
QTDW [13]	= 2	✓	×	$\mathcal{O}(T^2)$
MMDDTW [15]	> 2	×	×	$\mathcal{O}(T^N)$
TTW [10]	$\gg 2$	×	×	$\mathcal{O}(NZ)$
NTW [11]	$\gg 2$	×	×	$\mathcal{O}(NZT)$
RTW [proposed]	$\gg 2$	✓	✓	$\mathcal{O}(NZ)$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 2. Visualization of the Gaussian-weighted window loss between two signals $\hat{\mathbf{x}}$ (green) and $\hat{\mathbf{y}}$ (red) as defined in (8). The example shows one segment at index $z = 50$ with window size $\kappa = 4$ and step size $\rho = 5$, where red, yellow and green mark the areas covered by one, two and three times the variance σ of the Gaussian. By considering a segment around index z , the first peak of both signals becomes distinguishable from the second.

The unconstrained optimization problem (6) approximates (1) and is solved using state-of-the-art optimization techniques.

Finally, we define the distance measure $\mathcal{D}_\Omega(\hat{\mathbf{X}})$ as a Gaussian-weighted window loss of point segments

$$\mathcal{D}_\Omega(\hat{\mathbf{X}}) = \frac{1}{NZ} \sum_{n=1}^N \sum_{z=1}^Z \sum_{k=-\kappa}^{\kappa} \hat{g}(v, z) \text{dist}(\hat{\mathbf{x}}_n[v], \hat{\mathbf{y}}[v]), \quad (8)$$

$$\hat{g}(v, z) = \frac{g(v)}{\sum_{i=-\kappa}^{\kappa} g(z + i\rho)}, \quad g(j) = \exp\left(-\frac{j^2}{2\sigma^2}\right),$$

with $v = z + k\rho$, where the window size $\kappa \in \mathbb{N}_0$ describes how many additional indices are considered at index z and the step size $\rho \in \mathbb{N}$ defines the distance between them, resulting in the indices v that form a segment⁵. Each index within a segment is weighted with the Gaussian kernel $\hat{g}(v, z)$. Given the empirical rule, we set $\sigma = \frac{\kappa\rho}{3} + \epsilon$ with a small $\epsilon > 0$ for numerical reasons to ensure a weighting, such that all indices lie within 99,7% of the Gaussian's power, emphasizing central indices as visualized in Figure 2. Note that by setting $\kappa = 0$, we receive a point-to-point metric similar to TTW or NTW. For data in \mathbb{R}^D , we utilize the Euclidean distance metric

$$\text{dist}_{\mathbb{R}^D}(\hat{\mathbf{x}}[z], \hat{\mathbf{y}}[z]) = \|\hat{\mathbf{x}}[z] - \hat{\mathbf{y}}[z]\|. \quad (9)$$

IV. RTW EXTENSION FOR RIEMANNIAN MANIFOLDS

Extending RTW to data on a Riemannian manifold \mathcal{M} is not straightforward as there is no sum and scalar multiplication defined, i.e., it is not a vector space. Accordingly, the classical time warping approaches yield geometrically invalid results, because the warped signals and their mean would not reside on the manifold. Specifically, sinc interpolation (4), mean computation (3) and distance evaluation (9) would provide outcomes that do not respect the manifold structure. To address this, our key idea is to utilize Euclidean tangent spaces \mathcal{T} that locally preserve distances. Then, we adjust the sinc interpolation (see IV-A), mean computation (see IV-B) and distance evaluation (see IV-C), to incorporate the curved manifold structure by operating within the tangent spaces. These novel modifications enable temporal alignment in the continuous time domain to be applied on Riemannian manifolds without affecting the computational complexity $\mathcal{O}(NZ)$: RTW scales linearly with N and Z .

⁵As for the sinc interpolation, we apply clamping to always receive valid indices for a full segment, i.e., we set $v = \max(1, \min(Z, v))$.

More specifically, a manifold \mathcal{M} is a D -dimensional smooth space [14], for which there exists a tangent space \mathcal{T}_p for each point $p \in \mathcal{M}$ on the manifold, with p being the origin of \mathcal{T}_p . A manifold with a positive definite inner product defined on the tangent space is called Riemannian. The exponential map $\text{Exp}_p : \mathcal{T}_p \mapsto \mathcal{M}$ maps a point from the tangent space onto the manifold. The inverse operation is the logarithmic map $\text{Log}_p : \mathcal{M} \mapsto \mathcal{T}_p$. These mappings are distance-preserving, i.e., the geodesic distance between the point p and any other point is maintained. This feature allows indirect computation to be performed on the manifold by applying Euclidean operations in the tangent space. The formulation of the exponential and logarithmic map is manifold dependent.

A. Windowed Sinc Interpolation on Riemannian Manifolds

The windowed sinc interpolation in (4) for a time index z constitutes a weighted sum, which is not a closed operation under \mathcal{M} . Therefore, we project the original $2\nu + 1$ data points from the manifold onto the tangent space $\mathcal{T}_{p_n[z]}$ through $\mathbf{u}_{n,z}[m] = \text{Log}_{p_n[z]}(\mathbf{x}_n[m])$, where we warp and interpolate

$$\hat{\mathbf{u}}_n[z] = \sum_{m=\lfloor \omega_n[z]T_n \rfloor - \nu}^{\lfloor \omega_n[z]T_n \rfloor + \nu} \mathbf{u}_{n,z}[m] \text{sinc}(m - \omega_n[z]T_n). \quad (10)$$

Then, the result for the time index z is projected from the specific tangent space $\mathcal{T}_{p_n[z]}$ back onto the manifold via

$$\hat{\mathbf{x}}_n[z] = \text{Exp}_{p_n[z]}(\hat{\mathbf{u}}_n[z]). \quad (11)$$

Note that the choice of the tangent point $p_n[z] \in \mathcal{M}$ is important to minimize undesired distortions [22]. Therefore, we initialize the origin of the tangent space at the center of the interpolation window as $p_n[z] = \mathbf{x}_n[\lfloor \omega_n[z]T_n \rfloor]$, and refine it through few Gauss-Newton iterations $p_n[z] \leftarrow \hat{\mathbf{x}}_n[z]$ to approximate the optimal tangent point that minimizes distortions, i.e., we compute a weighted Fréchet mean within the interpolation window. Due to the continuity constraint of the warping functions ω_n , the Fréchet mean is constrained to lie on the sinc interpolated line connecting the three adjacent discrete data points $\mathbf{x}_n[\lfloor \omega_n[z]T_n \rfloor - 1]$, $\mathbf{x}_n[\lfloor \omega_n[z]T_n \rfloor]$ and $\mathbf{x}_n[\lfloor \omega_n[z]T_n \rfloor + 1]$. Further note that the optimal tangent space is reevaluated for each time index z and signal n , instead of projecting all data points onto one single tangent space.

B. Signal Mean Computation on Riemannian Manifolds

Similar to the sinc interpolation, the mean computation (3) is adjusted by utilizing Gauss-Newton iterations. Starting from an initial estimate $\hat{\mathbf{y}}[z] \in \mathcal{M}$, e.g., $\hat{\mathbf{y}}[z] = \hat{\mathbf{x}}_1[z]$, the z -th data point $\{\hat{\mathbf{x}}_n[z] \in \mathcal{M}\}_{n=1}^N$ of each signal is projected onto the tangent space $\mathcal{T}_{\hat{\mathbf{y}}[z]}$ through $\hat{\mathbf{u}}_n[z] = \text{Log}_{\hat{\mathbf{y}}[z]}(\hat{\mathbf{x}}_n[z])$. We obtain the center of these projections as

$$\hat{\mathbf{v}}[z] = \frac{1}{N} \sum_{n=1}^N \hat{\mathbf{u}}_n[z]. \quad (12)$$

This center is projected back from the tangent space $\mathcal{T}_{\hat{\mathbf{y}}[z]}$ onto the manifold \mathcal{M} to update the mean at time index z through

$$\hat{\mathbf{y}}[z] \leftarrow \text{Exp}_{\hat{\mathbf{y}}[z]}(\hat{\mathbf{v}}[z]). \quad (13)$$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

This process, (12) and (13), is repeated until convergence, which is typically reached after few iterations, e.g., see [12].

C. Distance Computation on Riemannian Manifolds

The distance metric (9) is adapted as well. We utilize the geodesic distance on Riemannian manifolds

$$\text{dist}_{\mathcal{M}}(\hat{\mathbf{x}}[z], \hat{\mathbf{y}}[z]) = \left\| \mathbf{Log}_{\hat{\mathbf{y}}[z]}(\hat{\mathbf{x}}[z]) \right\|_{\hat{\mathbf{y}}[z]}, \quad (14)$$

where the norm is computed with respect to the Riemannian metric at $\hat{\mathbf{y}}[z]$, which depends on the manifold's geometry.

D. Further Remarks on the Overview provided in Fig. 1

Figure 1 provides a schematic overview, summarizing the proposed RTW approach. Three similar signals with significant temporal shifts are considered in \mathbb{S}^2 , a unit sphere in \mathbb{R}^3 . The color gradient from blue to red indicates the progression of time along the trajectories. We illustrate the mapping of one signal onto the tangent space for two different origins $\mathbf{p}_n[z]$, where the tangent space is a two-dimensional plane. As expected, distortion is noticeable only for data points far from the origin. For visualization purposes, we choose a large interpolation window ν , indicated by the bold segments. Note that these parts are close to the origin of the tangent space but not spatially centered around it, as the window is centered in the time domain instead. Furthermore, RTW only projects data points from a single signal onto each tangent space during interpolation, while related works often project all data points onto them, e.g., see [12], [19], [23]. Applying warping functions ω_n does not alter the shape of a signal in the tangent spaces; rather, signals are temporally adjusted, as shown by the new color coding. RTW results in three signals on the manifold that are well aligned. After a few Gauss-Newton iterations, we obtain the signal mean, respecting the unit sphere structure.

V. SIMULATIONS AND EXPERIMENTS

This section compares the proposed RTW approach with baselines in various settings considering datasets in \mathbb{R}^D , \mathbb{S}^D and \mathbb{S}_{++}^D . We also apply RTW as a pre-processing step for learning movement primitives for a real-robot teaching scenario in $\mathbb{R}^3 \times \mathbb{S}^3$ using an LBR iiwa manipulator.

RTW, TTW and NTW have been implemented in PyTorch, utilizing the Adam optimizer with a learning rate of 0.01 and Autograd for gradient computations. The best model over 256 epochs was selected during the comparisons. For RTW, we use a neural network to model $\phi_{\theta}(\hat{z})$ with four fully connected linear layers⁶, ReLU activations and skip connections of size 1-512-512-1025-(N-1), as suggested in [11]. However, unlike NTW, we do not initialize the weights with zero, as preliminary trials have shown that this significantly degrades the performance. Instead, we apply symmetry breaking [24], initializing weights with the Xavier initialization method [25].

⁶We also evaluated simpler model architectures, such as estimating coefficients of a polynomial, which achieved comparable results at faster runtime for our experiments. However, for generalization purposes, and to maintain a fair comparison with NTW, we stick to the neural network architecture.

We use a window size of $\nu = 10$ for the sinc interpolation in RTW/TTW (10). For the Gaussian-weighted window loss in RTW (8), we choose window size $\kappa = 5$ and step size $\rho = 5$, which had negligible impacts on the runtime. The penalty term (6) for RTW/NTW is weighted with $\lambda = 100$.

The number of learnable parameters θ in NTW and RTW is $263680 + 1025(N - 1) + (N - 1)$, i.e., the number of weights and biases from all layers of the neural network, whereas it is only NK for TTW due to the design of the warping function

$$\omega_n^{ttw}[z] = \hat{z} + \sum_{k=1}^K \alpha_{n,k} \sin(\pi k \hat{z}) \quad , \quad \hat{z} = \frac{z - 1}{Z - 1}, \quad (15)$$

where $K \in \mathbb{N}_+$ defines how many components of the discrete sine transform are used to model the warping functions.

We also recommend watching the accompanying video.

A. Extensive Benchmark on the UCR Time Series Archive

We evaluate the performance of the simplified RTW (as presented in III-B) w.r.t. TTW/NTW in \mathbb{R}^1 . Therefore, we perform *DTW averaging* and *classification tasks* on the UCR Time Series Archive [26], replicating the experiments in [10], [11] to enable consistent benchmarking and fair comparisons. This archive consists of 128 datasets covering various applications; each is pre-split into train $\mathbf{x}_{n,c}^{\text{train}}$ and test data $\mathbf{x}_{n,c}^{\text{test}}$, containing $N \in \{16, \dots, 16800\}$ signals classified into $c \in \{2, \dots, 60\}$ classes, with $T \in \{15, \dots, 2844\}$ data points, where most signals are of the same length, i.e., $T_1 = \dots = T_N$.

1) *DTW Averaging Task*: For all datasets, we perform temporal alignment on the training data $\mathbf{x}_{n,c}^{\text{train}}$ for each class c separately, to estimate its mean signal $\hat{\mathbf{y}}_c^{\text{train}}$. We then compute the accumulated DTW distance $\sum_{n=1}^N \mathcal{D}_{dtw}(\hat{\mathbf{x}}_{n,c}^{\text{train}}, \hat{\mathbf{y}}_c^{\text{train}})$ [7] for all three methods and perform a paired t-test ($\alpha = 0.05$), to find statistically significant differences. Table II-1. lists the results, where each entry describes for how many datasets a method significantly outperformed the other: RTW and NTW are comparable, whereas RTW clearly outperforms TTW.

2) *Classification Task*: We conduct classification tasks on the remaining test data $\mathbf{x}_{n,c}^{\text{test}}$. Given the estimated means $\hat{\mathbf{y}}_c^{\text{train}}$ for each class c of a dataset, we predict the class label \hat{c}_n for all test signals $\mathbf{x}_{n,c}^{\text{test}}$ utilizing nearest centroid classification as $\hat{c}_n = \text{argmin}_c \mathcal{D}_{dtw}(\mathbf{x}_{n,c}^{\text{test}}, \hat{\mathbf{y}}_c^{\text{train}})$. The results of the paired t-tests ($\alpha = 0.05$) between TTW, NTW and RTW are shown in Table II-2. RTW learns complex non-linear warpings, outperforming both TTW and NTW significantly. Compared to a neural network, modeling the warping functions through the discrete sine transform as in TTW (15) limits their flexibility.

TABLE II
STATISTICAL COMPARISON ON THE UCR TIME SERIES ARCHIVE

	1. Averaging	2. Classification
TTW significantly better than RTW	7.34%	5.47%
RTW significantly better than TTW	39.84%	54.69%
No significant difference	52.82%	39.84%
NTW significantly better than RTW	7.81%	6.25%
RTW significantly better than NTW	8.59%	35.16%
No significant difference	83.60%	58.59%

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 3. Classification performance of TTW (red), NTW (yellow), RTW (green) on all 128 datasets from the UCR Time Series Archive. Each point in the chart represents the TTW/NTW accuracy (x-axis) and RTW accuracy (y-axis). Due to overlapping points, a histogram shows the accuracy-distribution at the top.

Without applying symmetry breaking in NTW, the neural network only learns a linear model, which limits its warping functions to a quadratic structure. Figure 3 visualizes the classification performance for all 128 datasets.

B. Inverted Time Warping in \mathbb{S}^1

Signal alignment is typically performed in an unsupervised setting, making it difficult to evaluate the result w.r.t. the optimal solution, as it is usually not available for datasets with $N \gg 2$ signals [11]. Here, we mitigate this issue through an inverted time warping approach¹¹, which we apply to data in \mathbb{S}^1 . Exponential and logarithmic maps are defined as

$$\text{Exp}_p^{\mathbb{S}^D}(\mathbf{u}) = p \cos(\|\mathbf{u}\|) + \frac{\mathbf{u}}{\|\mathbf{u}\|} \sin(\|\mathbf{u}\|), \quad (16)$$

$$\text{Log}_p^{\mathbb{S}^D}(\mathbf{x}) = \arccos(p^T \mathbf{x}) \frac{\mathbf{x} - p^T \mathbf{x} p}{\|\mathbf{x} - p^T \mathbf{x} p\|}. \quad (17)$$

For the unit sphere \mathbb{S}^D , the tangent space is Euclidean at any base point, hence the geodesic distance simplifies to

$$\text{dist}_{\mathbb{S}^D}(\mathbf{x}, p) = \|\text{Log}_p^{\mathbb{S}^D}(\mathbf{x})\| = \arccos(p^T \mathbf{x}). \quad (18)$$

More precisely, starting from a single pre-defined signal $\mathbf{x}_{\text{original}}$, we create N synthetic trajectories $\mathbf{X} = \{\mathbf{x}_n\}_{n=1}^N$ of lengths $T_n = 100$, i.e. $\mathbf{x}_n = \{\mathbf{x}_n[t]\}_{t=1}^{T_n}$, where $\mathbf{x}_n[t] \in \mathbb{S}^1$. This is achieved by generating N random warping functions ω_n , obtained by uniformly sampling from both (5) and (15) with random parameters, as well as from random spline interpolations of various degrees. Sampled warping functions that do not satisfy constraints (2a)-(2c) are rejected. Then, starting from the original signal $\mathbf{x}_{\text{original}}$, the synthetic signals \mathbf{X} are obtained through (10)-(11), resulting in a diverse dataset. The goal of this evaluation is to reverse this process, by temporally re-aligning the generated data \mathbf{X} , i.e., we aim for aligned signals with $\mathbf{x}_{\text{original}} \stackrel{!}{=} \hat{\mathbf{x}}_n, \forall n \in \{1, \dots, N\}$.

⁷Modified version utilizing (14) for data on Riemannian manifolds.

⁸Modified version utilizing (10)-(14) for data on Riemannian manifolds.

⁹CPU used for the experiments is an AMD® Ryzen 5 3500U.

¹⁰GPU used for the experiments is an NVIDIA GeForce RTX 4080 Super.

¹¹Warping functions that satisfy constraints (2a)-(2c) can also be used for creating an artificial temporal shift of a given signal. Then, optimal warping functions exist to reverse the temporal shift.

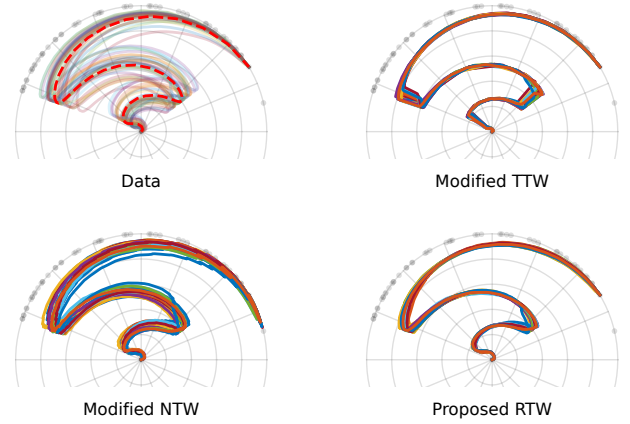


Fig. 4. Time warping on the \mathbb{S}^1 Riemannian manifold. All data points lie on the unit circle (grey dots), however, for visualization purposes we vary the radius to also indicate the time information between 0 and 1. Starting from an original signal $\mathbf{x}_{\text{original}}$ (dashed line), $N = 30$ random signals (transparent lines) are generated through an inverted time warping approach (see Sec. V-B).

We evaluate RTW in two settings, with $N = 4$ and $N = 30$ signals. For transparent results, we measure the computation time on a high-grade GPU and a standard CPU. Furthermore, we evaluate three well-established metrics

$$\text{Restoration Accuracy: } \sum_{n=1}^N \mathcal{D}_{dtw}(\hat{\mathbf{x}}_n, \mathbf{x}_{\text{original}}) \quad (19a)$$

$$\text{Barycenter Loss: } \sum_{n=1}^N \mathcal{D}_{dtw}(\mathbf{x}_n, \hat{\mathbf{y}}) \quad (19b)$$

$$\text{Alignment Quality: } \sum_{n=1}^N \mathcal{D}_{dtw}(\hat{\mathbf{x}}_n, \hat{\mathbf{y}}) \quad (19c)$$

where (18) is used within each \mathcal{D}_{dtw} distance. The *Restoration Accuracy* describes how effectively random warping functions are reversed to restore the original signal $\mathbf{x}_{\text{original}}$. The *Barycenter Loss* measures how accurately the mean $\hat{\mathbf{y}}$ of warped signals $\hat{\mathbf{x}}_n$ represents the data \mathbf{x}_n . The *Alignment Quality* indicates closeness of aligned signals $\hat{\mathbf{x}}_n$ to their mean $\hat{\mathbf{y}}$. The average for each metric is computed from 100 experiments.

1) *Evaluation for $N = 4$* : The literature does not offer time warping schemes for $N > 2$ signals with data on Riemannian manifolds. Hence, we modify MMDDTW [15] to use the

TABLE III
INVERTED TIME WARPING ON \mathbb{S}^1 FOR $N = 4, T = 100$

Metric	MMDDTW ⁷	TTW ⁸	NTW ⁸	RTW
Restoration Accuracy	0.0601	0.1712	0.0200	0.0197
Barycenter Loss	0.0552	0.1801	0.0229	0.0206
Alignment Quality	0.0403	0.0878	0.0243	0.0215
CPU ⁹ Runtime in [s]	3834.3273	3.3734	10.7157	8.4883
GPU ¹⁰ Runtime in [s]	1505.9386	1.5650	1.6639	1.7574

TABLE IV
INVERTED TIME WARPING ON \mathbb{S}^1 FOR $N = 30, T = 100$

Metric	p-DTW ⁷	TTW ⁸	NTW ⁸	RTW
Restoration Accuracy	1.0806	0.6458	0.0254	0.0293
Barycenter Loss	2.1259	0.7049	0.2954	0.0689
Alignment Quality	1.2591	0.4332	0.3043	0.0749
CPU ⁹ Runtime in [s]	19.5264	18.8754	226.9687	103.1399
GPU ¹⁰ Runtime in [s]	10.6953	9.4830	10.9298	10.2900

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Riemannian distance (18), as well as TTW and NTW to adapt (10)-(14). As Table III shows, RTW achieves the best performance across all metrics. While TTW is the fastest, the computational overhead of RTW is well-justified by its results.

2) *Evaluation for $N = 30$* : MMDDTW is not capable of handling such a large number of signals in a reasonable time. Instead, we utilize a pairwise approach by iteratively aligning two signals, denoted as p-DTW. The result is shown in Figure 4 and Table IV. RTW performs best across all metrics and the difference to NTW becomes more pronounced.

C. Multi-dimensional Riemannian manifolds

Next, we extend our experiments to higher-dimensional data by applying RTW to the Riemannian manifolds \mathbb{S}^3 and \mathbb{S}_{++}^D .

1) *Evaluation in \mathbb{S}^3* : To further validate our results from section V-B, we compare RTW with QDTW on signals with unit quaternions, i.e., \mathbb{S}^3 . Similar to [12], we artificially generate $N = 4$ valid signals of length $T = 200$ from the RCFS [27] dataset. This involves projecting handwritten motion data from \mathbb{R}^2 onto the unit sphere of \mathbb{S}^3 .

In Table V we observe that the mean signal obtained from RTW is an improved representation of the original data compared to QDTW. For the alignment quality, both methods yield similar results, however, RTW is significantly faster.

2) *Evaluation in \mathbb{S}_{++}^D* : We temporally align signals from the Riemannian manifold of symmetric positive definite (SPD) $D \times D$ matrices. The exponential and logarithmic maps are

$$\text{Exp}_{\mathbf{p}}^{\mathbb{S}_{++}^D}(\mathbf{u}) = \sqrt{\mathbf{p}} \exp\left(\sqrt{\mathbf{p}}^{-1} \mathbf{u} \sqrt{\mathbf{p}}^{-1}\right) \sqrt{\mathbf{p}}, \quad (20)$$

$$\text{Log}_{\mathbf{p}}^{\mathbb{S}_{++}^D}(\mathbf{x}) = \sqrt{\mathbf{p}} \log\left(\sqrt{\mathbf{p}}^{-1} \mathbf{x} \sqrt{\mathbf{p}}^{-1}\right) \sqrt{\mathbf{p}}, \quad (21)$$

where $\sqrt{\mathbf{p}}$ and $\sqrt{\mathbf{p}}^{-1}$ denote the matrix square root and its inverse, respectively. Note that the Riemannian metric on the SPD manifold is Euclidean only at the identity matrix. Hence, for distance computation under the affine-invariant-Riemannian metric, we use the canonical form where the Frobenius norm $\|\cdot\|_F$ yields the geodesic distance

$$\text{dist}_{\mathbb{S}_{++}^D}(\mathbf{x}, \mathbf{p}) = \left\| \log\left(\sqrt{\mathbf{p}}^{-1} \mathbf{x} \sqrt{\mathbf{p}}^{-1}\right) \right\|_F. \quad (22)$$

We compare RTW with two modified versions of MMD-DTW. The first version, denoted as MMDDTW_M, computes the distance between two SPD matrices according to the geodesic distance defined in (22). The second version, denoted as MMDDTW_C, utilizes the Cholesky decomposition $\text{chol}(\cdot)$ to linearize the manifold, and computes a distance through

$$\text{dist}_{\text{chol}}(\hat{\mathbf{x}}[z], \hat{\mathbf{y}}[z]) = \|\text{chol}(\hat{\mathbf{x}}[z]) - \text{chol}(\hat{\mathbf{y}}[z])\|_F. \quad (23)$$

¹²Modified version with N -dimensional cost matrix as in MMDDTW [15].

TABLE V
TIME WARPING ON \mathbb{S}^3 FOR $N = 4, T = 200$

Metric	QDTW ¹²	RTW
Barycenter Loss	0.0359	0.0339
Alignment Quality	0.0338	0.0338




Fig. 5. A planar robot (solid lines) follows endeffector trajectories (dashed lines), illustrated at $t = 50$ and $t = 100$ for each signal respectively (top-left and top-right). The configuration-dependent manipulability is computed for each time step and shown before and after applying RTW (bottom).

The data for this experiment was generated from the RCFS [27] dataset, i.e., a planar robot is controlled to follow the handwritten character G. We create a dataset for \mathbb{S}_{++}^2 by computing the manipulability as in [28], as shown in Figure 5, as well as \mathbb{S}_{++}^8 and \mathbb{S}_{++}^{32} by computing the mass matrix for a robot with 8 and 32 links, respectively, resulting in a dataset of $N = 3$ signals with a length of $T = 200$ each.

We evaluate the Barycenter Loss (19b) and the Alignment Quality (19c) including the geodesic distance given in (22). Table VI presents the results, supporting our previous findings: RTW outperforms MMDDTW in terms of Barycenter Loss and alignment quality with remarkable speed.

D. Real-Robot Teaching with Time Warping in $\mathbb{R}^3 \times \mathbb{S}^3$

Finally, we apply time warping in real-robot teaching scenarios, utilizing a KUKA LBR iiwa manipulator. We demonstrated $N = 4$ signals through hand-guiding control [29], and recorded the end-effector motion at 30 Hz, resulting in varying signal lengths $T_n \in \{2048, \dots, 3667\}$. Subsequently, the mean signal $\hat{\mathbf{y}}$ obtained from temporal alignment is exe-

TABLE VI
TIME WARPING WITH SPD MATRICES FOR $N = 3, T = 200$

Metric	D	MMDDTW _M	MMDDTW _C	RTW
Barycenter Loss	2	42.1131	43.8988	33.5702
Alignment Quality	2	42.3271	44.6875	23.0621
Barycenter Loss	8	20.8496	21.3290	18.8587
Alignment Quality	8	20.9473	21.4635	16.6931
Barycenter Loss	32	22.2595	23.3155	20.7979
Alignment Quality	32	22.6282	23.5338	18.7008

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

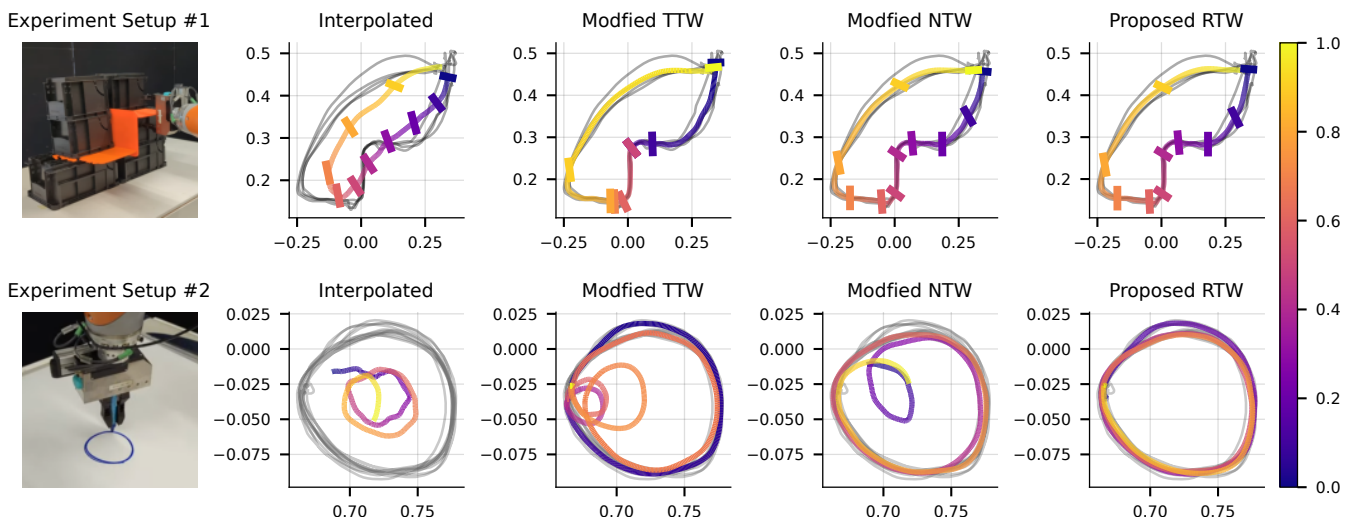


Fig. 6. Comparison of the results for both robot teaching tasks with naive interpolation, modified TTW/NTW and the proposed RTW approach. The hand-guided demonstration data \mathbf{x}_n is drawn in grey and the time evolution of the mean signal $\hat{\mathbf{y}}$ is visualized using the color map. Top: the robot has to follow a constrained path, which requires precise positioning and orientation. The results are shown from a side perspective and the bars on the trajectory indicate the end-effector's orientation (drawn ten times), which has to be perpendicular to the motion direction. Bottom: the robot has to draw three circles, where the difficulty lies in identifying and aligning the repetitive characteristics correctly. The results are shown from a top-down perspective.

cuted employing a state-of-the-art QP-controller [30], taking hardware limits and manipulator dynamics into account.

For the trajectory teaching, we denote an end-effector pose as $\mathbf{x} = \begin{bmatrix} \mathbf{x}_{\text{pos}} \\ \mathbf{x}_{\text{quat}} \end{bmatrix} \in \{\mathbb{R}^3 \times \mathbb{S}^3\}$, consisting of position $\mathbf{x}_{\text{pos}} \in \mathbb{R}^3$ and orientation $\mathbf{x}_{\text{quat}} \in \mathbb{S}^3$. Similar to [23], we concatenate the logarithmic and exponential map into single operations

$$\text{Log}_p^{\mathbb{R}^3 \times \mathbb{S}^3}(\mathbf{x}) = \begin{bmatrix} \mathbf{x}_{\text{pos}} \\ \text{Log}_{\mathbb{S}^3}(\mathbf{x}_{\text{quat}}) \end{bmatrix} = \begin{bmatrix} \mathbf{u}_{\text{pos}} \\ \mathbf{u}_{\text{quat}} \end{bmatrix} = \mathbf{u}, \quad (24)$$

$$\text{Exp}_p^{\mathbb{R}^3 \times \mathbb{S}^3}(\mathbf{u}) = \begin{bmatrix} \mathbf{u}_{\text{pos}} \\ \text{Exp}_{\mathbb{S}^3}(\mathbf{u}_{\text{quat}}) \end{bmatrix} = \begin{bmatrix} \mathbf{x}_{\text{pos}} \\ \mathbf{x}_{\text{quat}} \end{bmatrix} = \mathbf{x}. \quad (25)$$

In this experiment, we compare RTW with naive interpolation of the recorded data to the same length, as well as with modified TTW⁸ and NTW⁸ utilizing (24) and (25). Evaluation is done on two tasks, (i) following a constrained path, and (ii) repetitively drawing three circles, both shown in Figure 6.

1) *Following a constrained path:* We instruct the robot to follow a constrained path with its end-effector, requiring precise control of position and orientation. Due to unintended time shifts within the recorded demonstration data, naive interpolation does not suffice and the resulting mean signal deviates from the desired path, resulting in a collision with the environment. TTW⁸, NTW⁸ and RTW succeed in following the constrained path without collision, however, the warping functions generated from TTW⁸ (15) highly compress the motion in certain sections while stretching others, resulting in a mean signal that becomes difficult to execute by the robot, as high velocities are required at the compressed sections.

2) *Repetitively drawing three circles:* We demonstrate end-effector trajectories consisting of three circles on a whiteboard, which allows us to investigate the time warping performance for repetitive segments within provided signals. Again, naive interpolation is not sufficient. Even though both, TTW⁸ and NTW⁸, perform better, they do not align all parts of the trajectory properly. By only considering the distance between

single points, TTW⁸ and NTW⁸ fail to match repetitive patterns correctly. This results in severe deviations from the desired motion at the start and end of the trajectory for both TTW⁸ and NTW⁸. Furthermore, TTW⁸ again suffers from fast velocity changes. Only RTW successfully aligns all parts of the demonstration data. Repetitive characteristics of the signal are correctly identified due to the Gaussian-weighted window loss (8) that considers entire segments. When executing the mean signal $\hat{\mathbf{y}}$ obtained from the proposed RTW approach, the robot's end-effector successfully draws the three consecutive circles on the whiteboard, as shown in Figure 6.

To further validate the benefits of RTW, we evaluate its impact on the learning performance across three imitation learning approaches: Dynamic Movement Primitives (DMP) [31], Gaussian Mixture Models (GMM) [6], Gaussian Processes (GP). Using a leave-one-out cross-validation approach, we train each method on three demonstrations and evaluate the reproduction quality on the fourth unseen demonstration. The

TABLE VII
AVERAGE REPRODUCTION ERROR (CONSTRAINED PATH TASK)

Method	DTW _{pos} distance		DTW _{ori} distance	
	Raw Data	RTW Data	Raw Data	RTW Data
DMP	0.0392	0.0227	0.0825	0.0371
GMM	0.0302	0.0172	0.0639	0.0334
GP	0.0316	0.0175	0.0695	0.0338

TABLE VIII
AVERAGE REPRODUCTION ERROR (REPETITIVE CIRCLES TASK)

Method	DTW _{pos} distance		DTW _{ori} distance	
	Raw Data	RTW Data	Raw Data	RTW Data
DMP	0.0379	0.0222	0.0293	0.0244
GMM	0.0323	0.0220	0.0245	0.0239
GP	0.0329	0.0240	0.0245	0.0235

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

procedure is repeated four times per task to ensure each demonstration is reproduced once. Table VII and Table VIII present the average reproduction errors for both position and orientation. The results demonstrate that RTW significantly improves the learning performance across all methods. This validates that the benefits from RTW extend beyond the time warping process itself to the entire learning pipeline, making it a valuable preprocessing tool for imitation learning in robotics.

VI. CONCLUSION AND FUTURE WORK

This letter introduces *Riemannian Time Warping (RTW)*, a novel approach that efficiently aligns multiple signals while considering the geometric structure of the Riemannian manifold the data resides in, a feature that is essential for many robotic applications. We conducted extensive experiments on synthetic and real world data, proving that RTW consistently outperforms state-of-the-art baselines on both classification and averaging tasks. Furthermore, we have shown that preprocessing motion data with RTW for training a robot primitive significantly improves the learning result. The computational complexity of RTW scales linearly with the number of signals and data points, which constitutes a significant advancement over Dynamic Programming techniques.

While RTW demonstrates strong performance across various tasks, there are limitations to be addressed in future work. RTW has only been evaluated with closed-form exponential and logarithmic maps, which do not exist for all Riemannian manifolds, wherefore numerical alternatives may be utilized that result in instabilities. Additionally, while RTW shows clear benefits for classical motion primitive learning, it remains unclear whether RTW provides similar advantages for recent imitation learning schemes like [32] or [33], which encode multi-modal motion strategies based on demonstration data with large spatial variances.

ACKNOWLEDGMENTS

This work was partly supported by KUKA Deutschland GmbH and the state of Bavaria through the OPERA project DIK-2107-0004/DIK0374/01.

REFERENCES

[1] H. Shimodaira, K.-i. Noma, M. Nakai, and S. Sagayama, "Dynamic time-alignment kernel in support vector machine," in *Advances in Neural Information Processing Systems* (T. Dietterich, S. Becker, and Z. Ghahramani, eds.), vol. 14, MIT Press, 2001.

[2] L. Muda, M. Begam, and I. Elamvazuthi, "Voice recognition algorithms using mel frequency cepstral coefficient (MFCC) and dynamic time warping (DTW) techniques," *arXiv:1003.4083*, 2010.

[3] Y. Sheikh, M. Sheikh, and M. Shah, "Exploring the space of a human action," in *Proc. IEEE Int. Conf. Computer Vision*, pp. 144–149, 2005.

[4] S. Sempena, N. U. Maulidevi, and P. R. Aryan, "Human action recognition using dynamic time warping," in *Proc. Int. Conf. Electrical Engineering and Informatics*, pp. 1–5, 2011.

[5] M. Muhlig, M. Gienger, S. Hellbach, J. J. Steil, and C. Goerick, "Task-level imitation learning using variance-based movement optimization," in *IEEE Int. Conf. Robotics and Automation*, pp. 1177–1184, 2009.

[6] S. Calinon, F. Guenter, and A. Billard, "On learning, representing, and generalizing a task in a humanoid robot," *IEEE Trans. Systems, Man, and Cybernetics, Part B*, vol. 37, no. 2, pp. 286–298, 2007.

[7] H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," *IEEE Trans. Acoustics, Speech, and Signal Processing*, vol. 26, no. 1, pp. 43–49, 1978.

[8] R. Bellman, "The theory of dynamic programming," *Bulletin of the American Mathematical Society*, vol. 60, no. 6, pp. 503–515, 1954.

[9] F. Zhou and F. De la Torre, "Generalized time warping for multi-modal alignment of human motion," in *IEEE Conf. Computer Vision and Pattern Recognition*, pp. 1282–1289, 2012.

[10] S. Khorram, M. G. McInnis, and E. M. Provost, "Trainable time warping: Aligning time-series in the continuous-time domain," in *IEEE Int. Conf. Acoustics, Speech and Signal Processing*, pp. 3502–3506, 2019.

[11] K. Kawano, T. Kutsuna, and S. Koide, "Neural time warping for multiple sequence alignment," in *IEEE Int. Conf. Acoustics, Speech and Signal Processing*, pp. 3837–3841, 2020.

[12] S. Calinon, "Gaussians on Riemannian manifolds: Applications for robot learning and adaptive control," *IEEE Robotics & Automation Magazine*, vol. 27, no. 2, pp. 33–45, 2020.

[13] B. Jablonski, "Quaternion dynamic time warping," *IEEE Trans. Signal Processing*, vol. 60, no. 3, pp. 1174–1183, 2011.

[14] J. M. Lee, *Riemannian Manifolds: An Introduction to Curvature*, vol. 176, Springer, 2006.

[15] P. Sanguansat, "Multiple multidimensional sequence alignment using generalized dynamic time warping," *WSEAS Trans. Mathematics*, vol. 11, no. 8, pp. 668–678, 2012.

[16] L. Gupta, D. L. Molfese, R. Tammana, and P. G. Simos, "Nonlinear alignment and averaging for estimating the evoked potential," *IEEE Trans. Biomedical Engineering*, vol. 43, no. 4, pp. 348–356, 1996.

[17] N. Jaquier, L. Rozo, D. G. Caldwell, and S. Calinon, "Geometry-aware manipulability learning, tracking and transfer," *Int. Journal of Robotics Research*, vol. 40, no. 2-3, pp. 624–650, 2021.

[18] N. Jaquier and T. Asfour, "Riemannian geometry as a unifying theory for robot motion learning and control," in *Int. Symp. Robotics Research*, pp. 395–403, 2022.

[19] M. Saveriano, F. J. Abu-Dakka, and V. Kyrki, "Learning stable robotic skills on Riemannian manifolds," *Robotics and Autonomous Systems*, vol. 169, p. 104510, 2023.

[20] T. Löw, P. Abbet, and S. Calinon, "GAFRO: Geometric algebra for robotics," *IEEE Robotics & Automation Magazine*, 2024.

[21] C. E. Shannon, "Communication in the presence of noise," *Proc. IRE*, vol. 37, no. 1, pp. 10–21, 1949.

[22] N. Jaquier, L. Rozo, and T. Asfour, "Unraveling the single tangent space fallacy: An analysis and clarification for applying Riemannian geometry in robot learning," in *IEEE Int. Conf. Robotics and Automation*, pp. 242–249, 2024.

[23] M. J. Zeestraten, I. Havoutis, J. Silvério, S. Calinon, and D. G. Caldwell, "An approach for imitation learning on Riemannian manifolds," *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1240–1247, 2017.

[24] H. Tanaka and D. Kunin, "Noether's learning dynamics: Role of symmetry breaking in neural networks," *Advances in Neural Information Processing Systems*, vol. 34, pp. 25646–25660, 2021.

[25] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. Int. Conf. Artificial Intelligence and Statistics*, pp. 249–256, 2010.

[26] H. A. Dau, E. Keogh, K. Kamgar, C.-C. M. Yeh, Y. Zhu, S. Gharghabi, C. A. Ratanamahatana, Y. Chen, B. Hu, N. Begum, A. Bagnall, A. Mueen, G. Batista, and Hexagon-ML, "The UCR time series classification archive," October 2018. https://www.cs.ucr.edu/~eamonn/time_series_data_2018.

[27] S. Calinon, "Learning and optimization in robotics, lecture notes," tech. rep., Idiap Research Institute, 2023.

[28] N. Jaquier and S. Calinon, "Gaussian mixture regression on symmetric positive definite matrices manifolds: Application to wrist motion estimation with semg," in *2017 IEEE/RSS Int. Conf. on Intelligent Robots and Systems*, pp. 59–64, 2017.

[29] J. D. M. Osorio, F. Allmendinger, M. D. Fiore, U. E. Zimmermann, and T. Ortmaier, "Physical human-robot interaction under joint and cartesian constraints," in *IEEE Int. Conf. Advanced Robotics*, pp. 185–191, 2019.

[30] K. Bouyarmane, K. Chappellet, J. Vaillant, and A. Kheddar, "Quadratic programming for multirobot and task-space force control," *IEEE Trans. Robotics*, vol. 35, no. 1, pp. 64–77, 2018.

[31] S. Schaal, "Dynamic movement primitives - a framework for motor control in humans and humanoid robots," in *The International Symposium on Adaptive Motion of Animals and Machines*, (Kyoto, Japan, March 4-8, 2003), Mar. 2003. clmc.

[32] Y. Lee, "MMP++: Motion manifold primitives with parametric curve models," *IEEE Trans. Robotics*, 2024.

[33] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," *arXiv:2304.13705*, 2023.