

Temporal Transformer Networks: Joint Learning of Invariant and Discriminative Time Warping

Suhas Lohit Qiao Wang Pavan Turaga

Geometric Media Lab, Arizona State University

{slohit, qiao.wang, pturaga}@asu.edu

Abstract

Many time-series classification problems involve developing metrics that are invariant to temporal misalignment. In human activity analysis, temporal misalignment arises due to various reasons including differing initial phase, sensor sampling rates, and elastic time-warps due to subject-specific biomechanics. Past work in this area has only looked at reducing intra-class variability by elastic temporal alignment. In this paper, we propose a hybrid model-based and data-driven approach to learn warping functions that not just reduce intra-class variability, but also increase inter-class separation. We call this a temporal transformer network (TTN). TTN is an interpretable differentiable module, which can be easily integrated at the front end of a classification network. The module is capable of reducing intra-class variance by generating input-dependent warping functions which lead to rate-robust representations. At the same time, it increases inter-class variance by learning warping functions that are more discriminative. We show improvements over strong baselines in 3D action recognition on challenging datasets using the proposed framework. The improvements are especially pronounced when training sets are smaller.

1. Introduction

Guaranteed invariances of machine learning algorithms to nuisance parameters is an important design consideration in critical applications. Classically, invariances can only be guaranteed under a model-based approach. Learned representations however have not been able to guarantee invariances, except by empirical tests [8]. Learning invariant representations that build on analytical models of phenomena may hold the cue to bridge this gap, and can also help lend explainability to the model.

However, deep learning presents a challenge for learning explainable invariants, primarily due to incompatibility

between the mathematical approaches that underly invariant design, and the architectures prevalent in deep learning. There have been recent attempts at leveraging model-based and data-driven approaches to learn invariant representations across spatial transforms [16, 14, 38], illumination [33, 21], and viewpoints [20]. On the other hand, learning invariant/robust representations to temporal rate-variations has received significantly less attention. If tackled well, many applications of human activity modeling will benefit, including more robust recognition algorithms for human-robot interaction, richer synthesis of human motion for computer-generated imagery, and health applications.

Hybrid model- and data-based approach: In this paper, our chosen application is activity classification from RGBD devices, where skeleton data are available. Activities such as walking can be performed at different rates by different subjects owing to physiological and biomechanical factors. We would like to design representations that provide robust classification against such nuisance factors. To do this, we adopt a model-based approach, and constrain certain layers in deep network using the model. The model for temporal variability is adopted from past work in elastic temporal alignment which considers temporal variability as a result of a temporal diffeomorphism acting on a given time-series [31]. The space of such diffeomorphisms has a group structure, and can be converted to simpler geometric constraints by exploiting contemporary square-root forms to represent the diffeomorphic maps [29].

Compatibility with deep architectures: We design a novel module, which we refer to as a Temporal Transformer Network (TTN). The hallmark of this module is that it can be easily integrated into existing time-series classifiers such as Temporal Convolution Networks (TCNs) [17] and Long Short-Term Memory (LSTM) networks [9]. TTN is explainable in the sense that it is designed so as to interact with the classification network in a predefined, predictable and visualizable manner. TTN is a trainable network added at the beginning of the classifier and operates on the input sequence by performing selective temporal warping of the input sequences. As such, it has the ability to factor out

* Qiao Wang is now at SRI International.

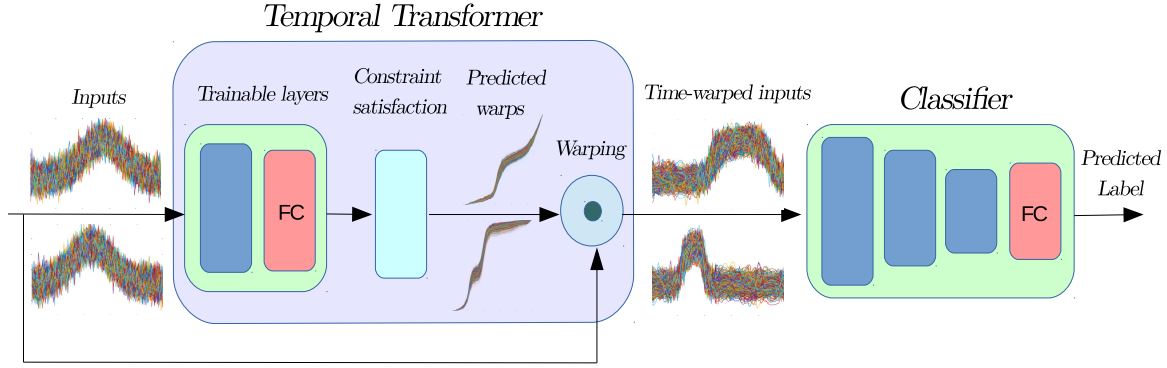


Figure 1. The Temporal Transformer Network (TTN) is a trainable module that is easily added at the beginning of a time-series classifier. Its function is to warp the input sequences so as to maximize the classification performance, as shown in the figure for two classes of waveforms which become more discriminative after passing through the TTN. The sub-modules of the TTN are explained in Section 4.

rate variations, if present in the data, as well as increase the inter-class separation by learning to align sequences in dissimilar classes away from each other.

Application impact: Recognition of human activities from motion capture (mocap) systems such as OptiTrack or depth sensors like Microsoft Kinect and Intel RealSense has been gathering a great deal of interest in the recent past. The cost of these sensors is ever-reducing and the increasing effectiveness of pose estimation algorithms (e.g. [34]) makes 3D skeletons an important sensing modality for action recognition. As the problem of action recognition presents a large amount of variability both inter-class as well as intra-class, we choose it as the focus of this paper.

Contributions

- We propose the Temporal Transformer Network (TTN), which performs joint representation learning as well as class-aware discriminative alignment for time-series classification including action trajectories.
- We design the TTN to generate highly expressive non-parametric, order-preserving diffeomorphisms.
- The TTN exploits the non-uniqueness of the optimal alignment (between equivalence classes) to generate discriminative warps for improved classification.
- The proposed TTN can be easily integrated into existing time-series classification architectures, with just a single line of code to employ the warping module.

We validate our contributions by demonstrating improved performance on small and large, and real and synthetic datasets, for action recognition from 3D pose obtained from two different modalities – Kinect and mocap. The combined architecture of the TTN and the classifier consistently yields improved classification performance compared to several baseline classifiers.

2. Related work

Deep learning of invariant representations: One of the main inspirations for this work is the paper by Jaderberg et

al. [11] on Spatial Transformer Networks (STNs) where a smaller network first predicts a geometric transform of the input grid parameterized by affine transforms or thin plate splines. The transformation is then applied to the input before feeding it to the classification network. A recent work is that of Detlefsen et al. [4] who improve the performance of spatial transformers by replacing affine transforms and thin plate splines with a richer class of parameterized diffeomorphic transforms called continuous piecewise-affine transforms, but at the expense of complex implementation and considerably longer training times. Both these works are aimed at building invariances to spatial geometric transforms of images. Capsule networks by Sabour et al. [25] expand the expressive capacity of CNNs by allowing them to learn explicit spatial relationships. An interesting recent work by Tallec and Ollivier [32] show that LSTM networks have the capability to learn to warp input sequences. Our experiments show that by integrating LSTMs with the module designed in this paper, the performance can be further increased, as the proposed framework can also lead to more discriminative representations.

In this paper, we design a module to predict warping functions in the temporal domain which when applied to the input sequences lead to higher classification performance. This requires the predicted warping functions to satisfy the order-preserving property. Moreover, in our case, the predicted warping functions are non-parameterized, can span the entire space of rate-modifying transforms and are much more expressive than earlier related works. The warping is also elastic, as opposed to rigid deformations which is the case with STNs. We note that these are important design requirements in the case of temporally varying signals, that are different from transforms in the 2D spatial domain.

Alignment of time-series data: The most commonly used method to align time-series data is Dynamic Time Warping (DTW) [2, 26]. DTW tries to minimize the \mathbb{L}^2 distance between two time series after a time warping is

applied to one of them, and is agnostic of class information. To address some of DTW’s shortcomings, new methods have been proposed recently, including the elastic functional data and shape analyses [30, 31] which defined proper metrics that are invariant to time warping, and soft-DTW which is a differentiable loss function that may be integrated into neural networks [3]. There are also several works on Canonical Time Warping (CTW) for multimodal data where the time series from different streams are projected to a common space (similar to canonical correlation analysis) before aligning them [43]. Deep learning versions of CTW have also been proposed recently [36, 35, 12].

One of the major differences between our proposed method and the aforementioned optimization-based time warping methods is that our approach performs discriminative warping based on class information and does not need signal templates. This is further discussed in Section 4. During the preparation of this document, we came across the paper by Oh et al. [24] who propose a similar architecture for clinical time series data classification but restricted to the space of linear time scaling and offsets.

3D action recognition using deep learning: As sensing systems like Microsoft Kinect, Intel RealSense and camera-based mocap systems are getting more effective at acquiring depth and at human pose estimation with centimeter-to millimeter-level accuracy, research and commercial interest in employing 3D pose data for action recognition has understandably increased. Recent experiments suggest that for small datasets, recognition accuracies are better with 3D pose information compared to video frames [7]. That simple landmark-based or skeleton-based action recognition can be effective is supported by evidence from works in psychology which show that humans are excellent at recognizing actions only from a few points on the human body [13].

Recurrent neural architectures, especially Long Short-Term Memory (LSTM) networks have been used to perform 3D action recognition e.g. [5, 27]. Song et al. propose including layers for spatial and temporal attention (STA-LSTM) [28] which greatly improves the recognition performance. For majority of the experiments in this paper, we will use the Temporal Convolution Network (TCN) with residual connections [18] as they are effective, simple to build and faster to train compared to LSTM-based networks. Additionally, Kim and Reiter have shown excellent results on using TCNs for 3D action recognition [17]. This network outperforms STA+LSTM [28] for 3D action recognition. They further show that TCNs can learn both spatial and temporal attention without the need for special attention layers. Also, the network filter activations are interpretable by design because of the residual connections. We also note that the TCN architecture presented in [17] incorporates pooling mechanisms inside the network.

We note that more recently, newer architectures have

proposed modifications to baseline architectures by using graph convolutions to better take into account the spatial structure of the joints in the human body [39]. However, it has a much higher computational load. Other representations include image-based ones [15, 20], and fusing skeleton information with velocity information [19] etc. Our contributions in this paper are orthogonal to these works, and the main focus of the paper is to design a specialized module for learning rate-robust discriminative representations. As such, for our experiments, we choose two effective widely used simple-to-implement architectures as our baselines – TCNs and LSTMs – and demonstrate improvements in recognition performance over these frameworks.

3. Diffeomorphic models for rate variation

A continuous time-series can be represented as a single-parameter curve, which we denote by $\alpha(\cdot)$, where $t \in [0, 1]$ is the parameter. In our case, t is time and we assume $\alpha(t) \in \mathbb{R}^N$. Another curve β is a *resampling* of α if $\beta = \alpha \circ \gamma$, where \circ is a function composition, and γ is the *resampling/warping function*. We focus on a specific set Γ of these warping functions (defined below), and two curves α_1, α_2 differing only by a change of rate of execution obey the equation $\alpha_1 = \alpha_2 \circ \gamma$, for some $\gamma \in \Gamma$. Figure 2 shows

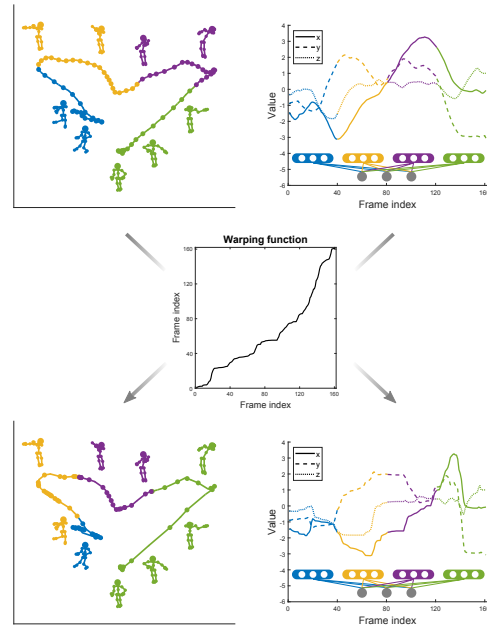


Figure 2. Top-left, Bottom-left: the trajectories and sampling points of the same action (“wearing jacket”) before and after a time warping (Center). The parameterized trajectories are visualized in \mathbb{R}^3 by using the sums of x, y, z coordinates of all joints. Notice how the time series of x, y, z , which are the inputs of a neural network (Top-right and Bottom-right), are quite different despite the action being the same. Here the action is segmented and colored to highlight the rate variation.

an example of such time warping, illustrated in \mathbb{R}^3 .

Given a 1-differentiable function γ defined on the domain $[0, T]$, for γ to be an element of Γ , γ needs to satisfy the following conditions:

$$\gamma(0) = 0, \gamma(1) = 1, \text{ and } \gamma(t_1) < \gamma(t_2), \text{ if } t_1 < t_2 \quad (1)$$

The above conditions fix the boundary conditions, and imply that any $\gamma \in \Gamma$ is a monotonically increasing function. This property is also called order-preserving which is important to the current discussion of action recognition as actions are critically dependent on sequencing/ordering of poses/frames. It is easy to show that

- $\forall \gamma_1, \gamma_2 \in \Gamma, \gamma_1 \circ \gamma_2 \in \Gamma$,
- $\gamma_{Id} \in \Gamma$,
- $\forall \gamma \in \Gamma, \exists \gamma^{-1} \in \Gamma$ s.t. $\gamma \circ \gamma^{-1} = \gamma_{Id}$, where $\gamma_{Id}(t) = t$, the identity warping function.

These properties imply that Γ admits a group structure with the group action being function composition. We denote by $\dot{\gamma}$, the first derivative of $\gamma \in \Gamma$, or

$$\gamma(t) = \int_0^t \dot{\gamma}(t) dt, \int_0^1 \dot{\gamma}(t) dt = \gamma(1) - \gamma(0) = 1 \quad (2)$$

Further, due to the monotonically increasing property of γ , we have $\dot{\gamma} > 0$. This, in conjunction with (2), implies that $\dot{\gamma}$ has the properties of a probability distribution function (positive, and integrates to 1), and the corresponding γ is thus equivalent to a cumulative distribution function.

As we work with digitized signals from sensors such as Kinect and mocap, we represent a discrete time series (either signals or features) by $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$. In this paper, we will work with time series in \mathbb{R}^N . Each $\mathbf{x}_t \in \mathbb{R}^N$ is called a frame of the sequence X . More clearly, we have $\alpha(t) = \mathbf{x}_t, t \in \{1, 2, \dots, T\}$. The warping function in the case of a discrete time signal is a discretized version of $\gamma \in \Gamma$, which we represent using γ with a slight abuse of notation. The derivative $\dot{\gamma}$ can be approximated by first order numerical differencing. Thus, (2) now becomes

$$\gamma(t) = \sum_{i=1}^t \dot{\gamma}(i) \quad \text{and} \quad \frac{1}{T} \sum_{i=1}^T \dot{\gamma}(i) = 1. \quad (3)$$

Two sequences α and β are said to be *equivalent* if there exists a $\gamma \in \Gamma$ such that $\alpha = \beta \circ \gamma$, and the set $\{\alpha \circ \gamma \mid \gamma \in \Gamma\}$ is called the equivalence class of α under rate variations and is denoted by $[\alpha]$. In classical elastic alignment, given two signals, a metric between sequences is defined as the minimal distance between their equivalence classes. This approach can be used to develop class-specific templates, and phase-amplitude separation [23] that reduces intra-class variance, but does not promote inter-class separation. Once the equivalence classes are defined, metrics are designed to compute distances between equivalence classes and develop methods to compute statistical measures such as mean and variance, which can be used to compute optimal alignments [31].

4. Temporal transformers for learning discriminative warping functions

The main idea presented in this paper is to use a specialized module, which we call a *Temporal Transformer Network* (TTN) for neural network-based classification which, given an input test sequence X , generates a warping function γ used to warp the input sequence by computing $X \circ \gamma$ and feed it to the classification network. It is important to note that the warping is carried out using linear interpolation. This makes it possible to train both the TTN and the classifier jointly end-to-end as the entire pipeline is (sub-) differentiable. Another notable aspect of this framework is that the warping functions are predicted without a “class-template”. Even though this sounds paradoxical, we will soon show that this allows our framework to jointly learn features as well as achieve discriminative warps. This capability makes our framework more powerful than template-based matching techniques like Dynamic Time Warping (DTW) and variants [2, 26].

Key Insight: Given two input sequences X_1 and X_2 such that they differ only by a warping transform, the trained TTN would ideally predict γ_1 and γ_2 , corresponding to X_1 and X_2 respectively such that $X_1 \circ \gamma_1 = X_2 \circ \gamma_2$. However, we note that γ_1 and γ_2 are not unique because $X_1 \circ \gamma_1 \circ \gamma = X_2 \circ \gamma_2 \circ \gamma, \forall \gamma \in \Gamma$. We believe this to be an important and often overlooked fact, and is formally referred to as invariance to group-action. The non-uniqueness of the warping functions is not fully exploited in the temporal alignment literature. The goal in past works in alignment has been to learn a class-specific template using an iterative optimization method, which is not unique due to invariance to group-action. Non-uniqueness here presents an opportunity that can be exploited to develop discriminative warps for classification problems.

The non-uniqueness of the warping functions can be advantageous as it expands the expressive capacity for classification. Minimizing intra-class variations — rate variations in our case — is only one part of the problem. For classification, we would also like to maximize inter-class variations. For example, if we have four sequences X_1, X_2, X_3 and X_4 such that X_1, X_2 belong to Class A, and X_3, X_4 belong to class B, the TTN has the capacity to predict $\gamma_1, \gamma_2, \gamma_3$ and γ_4 such that

- $d(X_1 \circ \gamma_1, X_2 \circ \gamma_2) < d(X_1, X_2)$
- $d(X_3 \circ \gamma_3, X_4 \circ \gamma_4) < d(X_3, X_4)$
- $d(X_i \circ \gamma_i, X_j \circ \gamma_j) > d(X_i, X_j), i \in \{1, 2\} \text{ \& } j \in \{3, 4\}$,

where $d(\cdot)$ is the Euclidean distance between sequences. However, we do not explicitly train the networks to achieve the above. Both the TTN and the classifier are trained so as to maximize classification performance by minimizing the cross-entropy loss between the predicted and true distribution over the class labels given the input sequence. The TTN can be divided into three sub-modules:

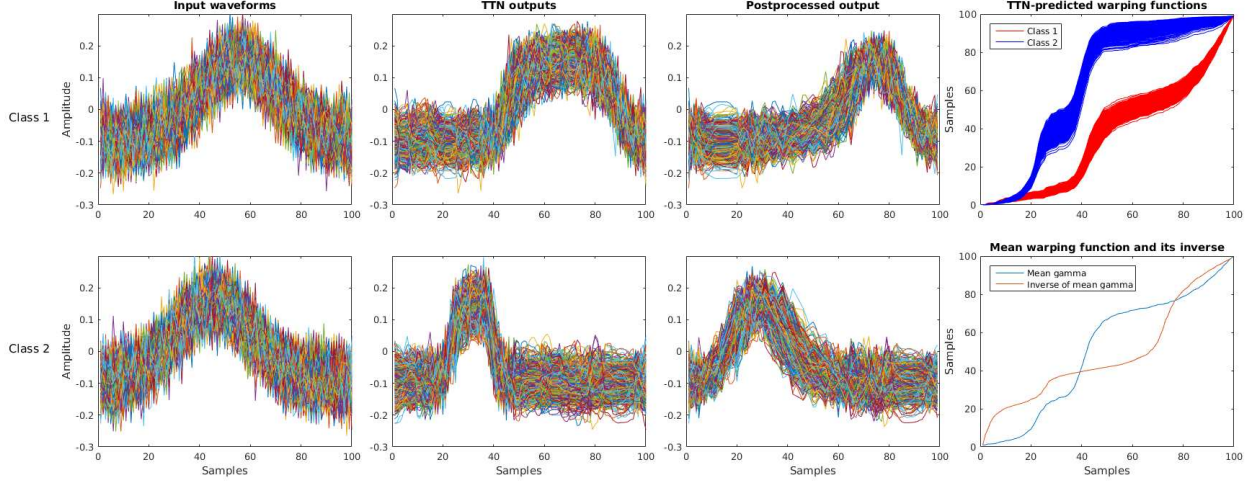


Figure 3. Results on synthetic dataset 1. Rows 1 and 2 show waveforms corresponding to classes 1 and 2 respectively. Columns 1 and 2 show the test inputs and the TTN outputs respectively. It is clear by comparing these columns that the TTN outputs are much better discriminated after warping. The TTN-predicted warping functions also show that the TTN performs class-dependent warping. Column 3 is a better visualization of column 2 after some post-processing making the mean of the generated warping functions $\gamma_\mu = \gamma_{Id}$ (see text).

Trainable layers: As shown in Figure 1, the input to the TTN trainable layers is an input sequence. The input is then passed through a few layers of convolutions and fully-connected layers. The network outputs a vector of length T , such that the first element is set to be zero. T is the length/number of frames in the input sequence. Let us denote this vector by $\mathbf{v} \in \mathbb{R}^T$.

Constraint satisfaction layers: The output \mathbf{v} is unconstrained, and hence, we need to convert it into a valid warping function that satisfies Equation (1).

To this end, we first divide \mathbf{v} by its norm to get a unit-vector, followed by squaring each of its entries. This has the effect of converting the vector into a point on the probability simplex. Thus, we use the following mappings:

$$\dot{\gamma} = \frac{\mathbf{v}}{\|\mathbf{v}\|} \odot \frac{\mathbf{v}}{\|\mathbf{v}\|}, \quad \text{and} \quad \gamma(t) = T \cdot \sum_{i=1}^t \dot{\gamma}(i), \quad (4)$$

where \odot is the Hadamard product (element-wise multiplication). This is treated as the network's estimate of the derivative of the warping function, denoted by $\dot{\gamma}$. We compute the cumulative sum and multiply it by the length of the input sequence, T , in order to form the warping function γ as shown in Equation (4).

Differentiable temporal resampling: The warping function γ is then applied to the input sequence using linear interpolation. We assume that the sampling rate of the signal is high enough in relation to the speed of the activity that simple linear interpolation of the frames is sufficient to get intermediate skeletons to look realistic (in practice, 20 frames/sec is plenty for most common action recognition applications). Warping is performed using the equation $Y(t_t) = X(t_s) = X(\gamma(t_t))$, where X and Y are the input and output sequences respectively, and t_s and

t_t are the source and target regular grid indices respectively. The frames of Y are to be defined at regular intervals $t_t = 1, 2, \dots, T$. As the t_s corresponding to t_t may not be integers, we use linear interpolation to find the values of $X(\gamma(t_t))$. This operation is sub-differentiable, as in the case of STN. Thus, we can write the expressions of the required gradients as follows (these expressions are adapted from Jaderberg et al. [11]). If X^j is the input sequence of the j^{th} joint, X_τ^j is the value at time index τ for channel j , Y^j is the warped sequence output by the TTN module and $i \in \{1, 2, \dots, T\}$ is the time index, we have

$$\frac{\partial Y_i^j}{\partial X_\tau^j} = \max(0, 1 - |t_i^s - \tau|) \quad \text{and} \quad (5)$$

$$\frac{\partial Y_i^j}{\partial t_i^s} = \sum_{\tau=1}^T X_\tau^j \cdot \begin{cases} 0, & \text{if } |t_i^s - \tau| \geq 1 \\ 1, & \text{if } \tau \geq t_i^s \\ -1, & \text{if } \tau < t_i^s \end{cases} \quad (6)$$

5. Experimental results

All networks in this paper are trained and tested using Tensorflow [1]. Due to space constraints, some training and testing details and results are provided in the supplement.

5.1. Synthetic datasets

(1) Demonstrating discriminative properties of TTN:

We consider a two-class classification problem where the two classes are one-dimensional time series signals of length 100. Let us denote each sequence in the dataset by $X \in \mathbb{R}^{100}$. All the signals are Gaussian functions with varying amplitude. Signals in class 1 are centered at $t = 0.55$ while signals in class 2 are centered at $t = 0.45$. Further, we corrupt the function with additive Gaussian

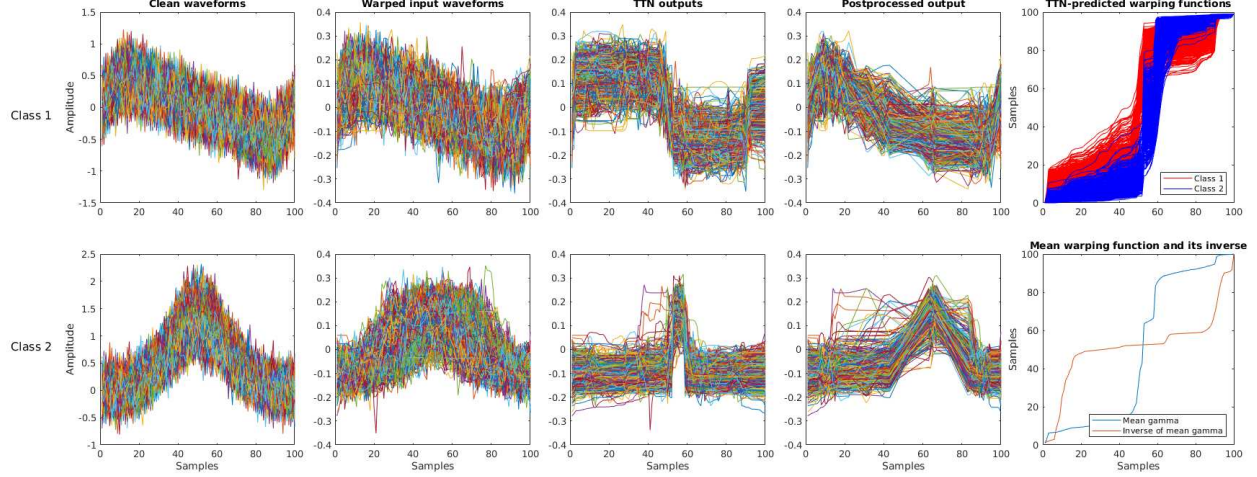


Figure 4. Results on synthetic dataset 2. Rows 1 and 2 show waveforms corresponding to classes 1 and 2 respectively. Columns 1, 2 and 3 show the clean waveforms, test inputs (after random warping) and the TTN outputs respectively. It is clear by comparing these columns that the TTN outputs are much more closely clustered especially for class 2, showing that the TTN outputs are robust to rate-variations. Column 4 is a better visualization of column 3 after some post-processing making the mean of the generated warping functions $\gamma_\mu = \gamma_{1d}$.

noise ($\mathcal{N}(0, 0.2)$). Samples of these functions are shown in Figure 3. We generate 8000 training and 2000 test sequences evenly balanced between classes 1 and 2. We use a simple classifier with a one-layer fully connected layer. The TTN is a 2-layer network with 1 convolutional layer producing 1 feature map with a filter of size 8, and 1 fully-connected layer. We train the networks for 10^3 iterations using Adam optimizer with an initial learning rate of 10^{-4} for the classifier. The weights of the TTN are updated at one-tenth the learning rate of the classifier. Figure 3 shows the test signals, corresponding outputs of the TTN, as well as the TTN-generated warping functions for every test input. It is clear from the figures that the TTNs predict class-specific warping functions in order to separate the peaks in the signals which makes them more discriminative. Note that this behavior arises automatically by minimizing the cross-entropy loss. In order to visualize the TTN outputs better, we perform post-processing by warping the TTN outputs with γ_μ^{-1} , where $\gamma_\mu = \sum_{i=1}^N \gamma_i$, where N is the size of the test set. This experiment clearly shows that TTNs are effective at increasing inter-class variations, as desired.

(2) Demonstrating rate-invariance of TTN: Here, we construct a dataset such that rate variations in the signals are the major nuisance parameter. In this scenario, intuitively, minimizing classification error should lead to the following: different signals belonging to the same class, but differing only by a γ should come closer to each other after passing through a trained TTN module. In class 1, we have signals which are N-waves with random warping applied and additive Gaussian noise added to them. Signals in class 2 are similar except that they are Gaussian functions. As before, we generated 8000 training sequences and 2000 test sequences evenly balanced between classes 1 and 2. These

are shown before (Column 1) and after random warping (Column 2) in Figure 4. The TTN, classifier, training and testing protocols are the same as in dataset (1) above. From column 3 in Figure 4, it is clear the TTN leads to reduction in intra-class rate variations.

5.2. ICL First-Person Hand Action dataset

In this section, we conduct experiments on a recently released real-world dataset of hand actions [7]. The dataset contains 3D hand pose sequences with 21 joint locations per frame of 45 daily hand action categories interacting with 26 objects, such as “pour juice”, “put tea bag” and “read paper”. These sequences are performed by 6 subjects and are acquired using an accurate mocap system. For our experiments, we use the training/test splits suggested by the authors of the dataset [7], with subjects 1,3,4 used for training and the rest for testing. The training set contains 600 sequences and the test set contains 575 sequences. As the sequences are of varying lengths, we uniformly sample the sequences such that all sequences contain 50 samples. If the sequences are shorter than 50 samples, we use zero-padding. As there are 21 joints per frame, each input sequence is of dimension 50×63 ($21 \times 3 = 63$). We normalize the sequences such that the wrist position of the first frame is at the origin. We conduct our experiments with two different types of classifiers widely used for action recognition: (1) Temporal Convolutional Network (TCN) and (2) 2-layer LSTM, showing that the proposed TTN framework can yield better results for both the classifier architectures.

The TTN module consists of 3 fully connected layers with \tanh nonlinearity and hidden states of dimensions 16 and 16. The final FC layer produces a vector of length 50 (equal to the input sequence length), with the first element set to zero (see Section 4).

The TCN architecture contains 1 temporal convolutional layer with 16, 32 or 64 feature maps, and 1 FC layer. We refer to these networks as TCN-16, TCN-32 and TCN-64 respectively. We ran the algorithm 5 times with different initializations of the TTN and report the mean and standard deviation in Table 1.

The LSTM architecture is similar to the one proposed in [7] containing two layers of LSTMs with a state dimension of 1024 and a dropout probability of 0.2. We use momentum optimizer with a momentum of 0.9 for training.

The results obtained are shown in Table 1. In addition to our experiments, we have reported results given in [7] for other important algorithms used for 3D pose-based action recognition including JOULE-pose [10], Moving Pose [40], Hierarchical Recurrent Neural Networks (HBRNN) [5], Transition Forests (TF) [6], and Lie Groups [37] and the Gram Matrix method [42], the last two of which also used DTW for sequence alignment, as well as non-Euclidean features to help improve performance. Among the baseline neural networks, TCN-32 led to the best results for this dataset, and addition of more layers did not yield better performance. We observe that addition of the TTN consistently improves performance over the baseline networks by 3.8 percentage points (TCN-16), 1.0 point (TCN-32), and 2.2 points (TCN-64). In the case of the LSTM classifier, we observe an improvement of 2.25 points using TTN + LSTM over just the LSTM. For the TCN-32 model, we performed K-means clustering (with $\#clusters = \#classes$ and averaged over 100 runs) on the features at the output of the conv. layer of the classifier for the test set, learned with and without TTN. Then, we computed *cluster purity* (CP), *homogeneity* (H) and *completeness* (C). Without the TTN, $CP = 0.519, H = 0.656, C = 0.705$. Adding the TTN module, we get improved scores of $CP = 0.530, H = 0.664, C = 0.709$.

Introducing distortions in the data: As datasets are usually collected in lab settings, there are relatively “clean” and do not contain many rate variations. Now, we introduce artificial rate variations in the data in order to better illustrate the usefulness of the TTN module. Here, we set the sequence length to 100 such that the original sequences of length 50 range from $t = 25$ to 75, and the rest of the values are set to zero. Now, we apply random “affine warps” to the training and test data. By an affine warp, we mean a warping function of the form $\gamma(t) = at + b, t = 25$ to 75, which is a linear time-scaling with an offset. We use $a \in [0.75, 1.25]$ and $b \in [0, 1, \dots, 49]$.

We observe that the induced distortion leads to a huge drop in performance of TCN-32 from 81.74% to 70.43%. With the TTN, the performance drop is much lower — from 82.75% to 78.26% and TTN+TCN-32 performs about 8 percentage points higher than TCN-32. Furthermore, from Figure 6, which shows the inputs, generated warping func-

Method	Accuracy (%)
Moving Pose [40]	56.34
JOULE-pose [10]	74.60
HBRNN [5]	77.40
TF [6]	80.69
Lie Group [37]	82.69
Gram Matrix [42]	85.39
2-layer LSTM	76.17
2-layer LSTM + TTN	78.43
TCN-16	76.28 \pm 0.29
TCN-16 + TTN	80.14 \pm 0.33
TCN-64	79.10 \pm 0.76
TCN-64 + TTN	81.32 \pm 0.36
TCN-32	81.74 \pm 0.27
TCN-32 + TTN	82.75 \pm 0.31
TCN-32 (affine warp)	70.43
TCN-32 + TTN (affine warp)	78.26

Table 1. Action recognition results on the ICL hand action dataset showing that LSTM+TTN and TCN+TTN frameworks consistently outperform LSTM and TCN baselines.




Figure 5. t-SNE plots for the test set features for the affine-warped ICL dataset with and without TTN. When the TTN is employed, we see a better separation of the clusters. This is also reflected in the accuracy scores.

tions and the TTN outputs, it can be readily observed that the TTN performs alignment of the sequences which then makes the classification problem much easier. This experiment shows that addition of the TTN enhances the interpretability of the network, and also delivers superior performance when there are larger rate variations in the data. As before, we performed K-means clustering. Without the TTN, $CP = 0.327, H = 0.487, C = 0.545$. Adding the TTN module, we get improved scores of $CP = 0.476, H = 0.611, C = 0.677$. We also ran t-SNE [22] to visualize the features in 2D. More pure and separated clusters are seen with the addition of TTN, as shown in the Figure 5.

5.3. NTU RGB-D dataset

In this section, we conduct experiments on a large-scale dataset of human actions known as the NTU RGB-D dataset [27] which contains about 56000 sequences of 3D skeleton positions acquired by a Microsoft Kinect. 25 joint locations are provided for each skeleton. The dataset contains actions belonging to 60 human activities performed by 45 subjects, with some actions containing two actors. The

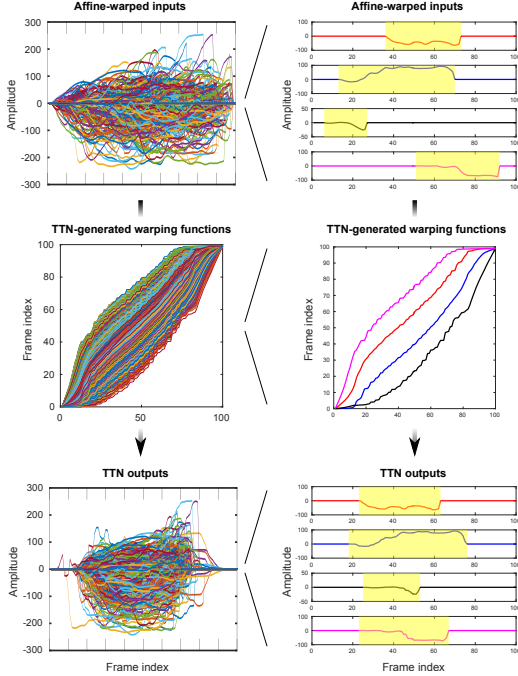


Figure 6. Visualizations of results of TCN-32 + TTN on ICL action dataset with induced rate variations. In the left column are waveforms corresponding to joint 1 of all test sequences. In the right column, 4 of those sequences are shown for clarity. We see clearly that the generated warping functions undo the affine-warp distortion in the test data, and the TTN outputs are nearly perfectly aligned leading to much better classification results.

data are acquired using a Microsoft Kinect. We sample 50 frames per sequence uniformly. We conduct two sets of experiments for this dataset – Cross Subject (CS) and Cross View (CV) – as per the protocol suggested by the authors in [27] using the same training and testing splits.

We construct a TTN module with 2 temporal convolutional layers and 3 FC layers with ReLU non-linearity. We use a filter size of 8 and 16 output feature maps in each conv layer. The FC layers produce hidden representations of sizes 16, 16 and 50 respectively.

We use the Temporal Convolution Network (TCN) described in [17]. The network consists of 10 convolutional layers with batch normalization and ReLU non-linearity. While training, the TTN parameters are updated at one-tenth the learning rate of the TCN.

The results obtained for this dataset are shown in Table 2. For cross-subject experiments, we observe that the addition of the TTN module results in a performance improvement of about 1 percentage point over the baseline TCN. We also found that using 2 parallel TTNs and concatenating the TTN outputs results in further improvement with a final performance of 77.80%. The addition of the TTN module leads to less improvement in the case of cross-view experiment. This can be explained by the fact that there is

Method	CS (%)	CV (%)
Lie Groups [37]	50.08	52.76
FTP Dynamic Skeletons [10]	60.23	65.22
HBRNN [5]	59.07	63.97
2-layer part-LSTM [27]	62.93	70.27
STA-LSTM [28]	73.40	81.20
VA-LSTM [41]	79.40	87.60
STA-GCN [39]	81.50	88.30
TCN [17]	76.54	83.98
TCN + TTN	77.55	84.25

Table 2. Action recognition results on the NTU RGB-D dataset showing that TCN+TTN frameworks outperforms the TCN.

likely less rate variation in the case of cross-view protocol compared to cross-subject. The ablation studies for TTN and TCN are provided in the supplement. As is the case of the ICL dataset, we performed K-means clustering (with $\#clusters = \#classes$ and averaged over 100 runs) on the features obtained at the penultimate layer of the TCN, and computed the same clustering metrics. Without the TTN, $CP = 0.466$, $H = 0.575$, $C = 0.597$. Adding the TTN module, we get improved scores of $CP = 0.493$, $H = 0.596$, $C = 0.621$. The corresponding t-SNE plots are provided in the supplement showing improved clusters with the addition of the TTN module.

6. Discussion and future work

In this work, we have proposed the Temporal Transformer Network (TTN) which can be readily integrated into classification pipelines. TTN has the ability to generate rate-invariant as well as discriminative warping functions for general time-series classification. We have shown improved classification results using different types of classifiers – TCNs and LSTMs – on challenging 3D action recognition datasets acquired using different modalities – Kinect and mocap. We have demonstrated the rate-invariant and discriminative properties of the TTN.

In the future, we would like to apply the ideas presented in this paper to video action recognition. However, it is not immediately clear how to perform temporal warping for videos as the currently widely used features for video frames may not be well suited for interpolation. One possible solution is to jointly train the image-level features and the action classification pipeline along with the TTN module. Temporal transformers can also be applied in general time-series classification which includes recognition from wearables, speech, EEG data, etc. Unsupervised pattern discovery with inbuilt warp-invariant metrics will be another interesting direction for further research.

Acknowledgements: This work was supported in part by NSF grant 1617999, and ARO grant number W911NF-17-1-0293

References

- [1] M. Abadi et al. Tensorflow: A system for large-scale machine learning. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation (OSDI)*, pages 265–283, 2016. [5](#)
- [2] R. Bellman and R. Kalaba. On adaptive control processes. *IRE Transactions on Automatic Control*, 4(2):1–9, November 1958. [2, 4](#)
- [3] M. Cuturi and M. Blondel. Soft-DTW: a differentiable loss function for time-series. In *International Conference on Machine Learning*, pages 894–903, 2017. [3](#)
- [4] N. S. Dettlefsen, O. Freifeld, and S. Hauberg. Deep diffeomorphic transformer networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4403–4412, 2018. [2](#)
- [5] Y. Du, W. Wang, and L. Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1110–1118, 2015. [3, 7, 8](#)
- [6] G. Garcia-Hernando and T.-K. Kim. Transition forests: Learning discriminative temporal transitions for action recognition and detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 407–415. IEEE, 2017. [7](#)
- [7] G. Garcia-Hernando, S. Yuan, S. Baek, and T.-K. Kim. First-person hand action benchmark with RGB-D videos and 3D hand pose annotations. In *Proceedings of Computer Vision and Pattern Recognition (CVPR)*, 2018. [3, 6, 7](#)
- [8] I. Goodfellow, H. Lee, Q. V. Le, A. Saxe, and A. Y. Ng. Measuring invariances in deep networks. In *Advances in neural information processing systems*, pages 646–654, 2009. [1](#)
- [9] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. [1](#)
- [10] J.-F. Hu, W.-S. Zheng, J. Lai, and J. Zhang. Jointly learning heterogeneous features for RGB-D activity recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5344–5352, 2015. [7, 8](#)
- [11] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015. [2, 5](#)
- [12] C. Jia, M. Shao, and Y. Fu. Sparse canonical temporal alignment with deep tensor decomposition for action recognition. *IEEE Transactions on Image Processing*, 26(2):738–750, 2017. [3](#)
- [13] G. Johansson. Visual perception of biological motion and a model for its analysis. *Perception & psychophysics*, 14(2):201–211, 1973. [3](#)
- [14] A. Kanazawa, A. Sharma, and D. Jacobs. Locally scale-invariant convolutional neural networks. *arXiv preprint arXiv:1412.5104*, 2014. [1](#)
- [15] Q. Ke, M. Bennamoun, S. An, F. Sohel, and F. Bousaid. A new representation of skeleton sequences for 3D action recognition. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 4570–4579. IEEE, 2017. [3](#)
- [16] R. Khasanova and P. Frossard. Graph-based isometry invariant representation learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1847–1856. JMLR. org, 2017. [1](#)
- [17] T. S. Kim and A. Reiter. Interpretable 3D human action analysis with temporal convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017. [1, 3, 8](#)
- [18] C. Lea, R. Vidal, A. Reiter, and G. D. Hager. Temporal convolutional networks: A unified approach to action segmentation. In *European Conference on Computer Vision*, pages 47–54. Springer, 2016. [3](#)
- [19] C. Li, Q. Zhong, D. Xie, and S. Pu. Skeleton-based action recognition with convolutional neural networks. In *2017 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pages 597–600. IEEE, 2017. [3](#)
- [20] M. Liu, H. Liu, and C. Chen. Enhanced skeleton visualization for view invariant human action recognition. *Pattern Recognition*, 68:346–362, 2017. [1, 3](#)
- [21] S. Lohit and P. Turaga. Learning invariant Riemannian geometric representations using deep nets. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 1329–1338, 2017. [1](#)
- [22] L. v. d. Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(Nov):2579–2605, 2008. [7](#)
- [23] J. S. Marron, J. O. Ramsay, L. M. Sangalli, and A. Srivastava. Functional data analysis of amplitude and phase variation. *Statistical Science*, 30(4):468–484, 11 2015. [4](#)
- [24] J. Oh, J. Wang, and J. Wiens. Learning to exploit invariances in clinical time-series data using sequence transformer networks. *Machine Learning for Healthcare*, 2018. [3](#)
- [25] S. Sabour, N. Frosst, and G. E. Hinton. Dynamic routing between capsules. In *Advances in Neural Information Processing Systems*, pages 3856–3866, 2017. [2](#)
- [26] H. Sakoe and S. Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1):43–49, Feb 1978. [2, 4](#)
- [27] A. Shahroury, J. Liu, T.-T. Ng, and G. Wang. NTU RGB+D: A large scale dataset for 3D human activity analysis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. [3, 7, 8](#)
- [28] S. Song, C. Lan, J. Xing, W. Zeng, and J. Liu. An end-to-end spatio-temporal attention model for human action recognition from skeleton data. In *Proc. of AAAI Conference on Artificial Intelligence*, 2017. [3, 8](#)
- [29] A. Srivastava, I. Jermyn, and S. Joshi. Riemannian analysis of probability density functions with applications in vision. In *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*, pages 1–8. IEEE, 2007. [1](#)
- [30] A. Srivastava, E. Klassen, S. H. Joshi, and I. H. Jermyn. Shape analysis of elastic curves in Euclidean spaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(7):1415–1428, 2011. [3](#)
- [31] A. Srivastava and E. P. Klassen. *Functional and shape data analysis*. Springer, 2016. [1, 3, 4](#)

- [32] C. Tallec and Y. Ollivier. Can recurrent neural networks warp time? *International Conference on Learning Representations*, 2018. 2
- [33] Y. Tang, R. Salakhutdinov, and G. Hinton. Deep lambertian networks. In *Proceedings of the 29th International Conference on International Conference on Machine Learning*, pages 1419–1426. Omnipress, 2012. 1
- [34] A. Toshev and C. Szegedy. DeepPose: Human pose estimation via deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1653–1660, 2014. 2
- [35] G. Trigeorgis, M. A. Nicolaou, B. W. Schuller, and S. Zafeiriou. Deep canonical time warping for simultaneous alignment and representation learning of sequences. *IEEE transactions on pattern analysis and machine intelligence*, 40(5):1128–1138, 2018. 3
- [36] G. Trigeorgis, M. A. Nicolaou, S. Zafeiriou, and B. W. Schuller. Deep canonical time warping. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5110–5118, 2016. 3
- [37] R. Vemulapalli, F. Arrate, and R. Chellappa. Human action recognition by representing 3D skeletons as points in a Lie group. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 588–595, 2014. 7, 8
- [38] Y. Xu, T. Xiao, J. Zhang, K. Yang, and Z. Zhang. Scale-invariant convolutional neural networks. *arXiv preprint arXiv:1411.6369*, 2014. 1
- [39] S. Yan, Y. Xiong, and D. Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proc. of AAAI Conference on Artificial Intelligence*, 2018. 3, 8
- [40] M. Zanfir, M. Leordeanu, and C. Sminchisescu. The moving pose: An efficient 3D kinematics descriptor for low-latency action recognition and detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2752–2759, 2013. 7
- [41] P. Zhang, C. Lan, J. Xing, W. Zeng, J. Xue, and N. Zheng. View adaptive recurrent neural networks for high performance human action recognition from skeleton data. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2136–2145. IEEE, 2017. 8
- [42] X. Zhang, Y. Wang, M. Gou, M. Sznaiar, and O. Camps. Efficient temporal sequence comparison and classification using Gram matrix embeddings on a Riemannian manifold. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4498–4507, 2016. 7
- [43] F. Zhou and F. De la Torre. Generalized canonical time warping. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):279–294, 2016. 3