

# One-shot action recognition in challenging therapy scenarios

Alberto Sabater<sup>1,2</sup> Laura Santos<sup>1,4</sup> José Santos-Victor<sup>1</sup> Alexandre Bernardino<sup>1</sup>  
Luis Montesano<sup>2,3</sup> Ana C. Murillo<sup>2</sup>

<sup>1</sup>ISR-Lisboa, Instituto Superior Técnico, Universidade de Lisboa, Portugal,

<sup>2</sup>DIIS-I3A, Universidad de Zaragoza, <sup>3</sup>Bitbrain Technologies, Spain

<sup>4</sup>Politecnico di Milano, Italy

## Abstract

*One-shot action recognition aims to recognize new action categories from a single reference example, typically referred to as the anchor example. This work presents a novel approach for one-shot action recognition in the wild that computes motion representations robust to variable kinematic conditions. One-shot action recognition is then performed by evaluating anchor and target motion representations. We also develop a set of complementary steps that boost the action recognition performance in the most challenging scenarios. Our approach is evaluated on the public NTU-120 one-shot action recognition benchmark, outperforming previous action recognition models. Besides, we evaluate our framework on a real use-case of therapy with autistic people. These recordings are particularly challenging due to high-level artifacts from the patient motion. Our results provide not only quantitative but also online qualitative measures, essential for the patient evaluation and monitoring during the actual therapy.*

## 1. Introduction

Human action recognition is a challenging problem with high relevance for many application fields such as human-computer interaction, virtual reality or medical therapy analysis. Certain works use raw appearance information for the action recognition tasks [4, 25]. Most recent works use skeleton-based representations [23, 13]. These representations encode the human pose independently of appearance, surroundings and being more robust to occlusions.

Real world scenarios often require the capability to recognize new action categories that cannot be learned, due to their creation on the fly or data limitations. Action classifiers that handle this problem [21, 11] of learning from limited data are based on encoding motion sequences into meaningful descriptors. Then, to recognize a new target sequence, they evaluate the similarity of the target descrip-

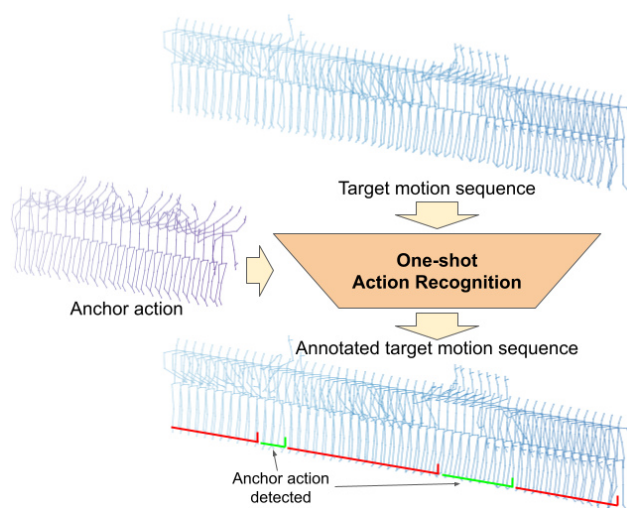


Figure 1: Overview of the proposed action recognition method. An anchor action is processed and compared against a target motion sequence. Output results show when the anchor action has been performed within the target motion sequence.

tor with the descriptors from one (one-shot) or few (few-shot) anchor labeled actions. Learning discriminative action encodings and their application in real scenarios are still open challenges. This is partially due to variable motion recording set-ups and unconstrained action execution rules (unsegmented action sequences, multiple consecutive executions, heterogeneous action executions, etc.).

This work tackles the problem of one-shot action recognition in unconstrained motion sequences. Our novel skeleton-based solution<sup>1</sup> is summarized in Fig. 1. In the proposed workflow, a stream of skeleton poses is encoded in an online manner, generating a descriptor at each time-step. Comparing the descriptors from a given reference anchor

<sup>1</sup>Code, learned models and supplementary video can be found in: <https://sites.google.com/a/unizar.es/filovi/>

action with descriptors from a target video, we can detect when the anchor action has been performed within the target sequence.

The main components of our work are two fold: 1) a motion encoder, based on geometric information, which is robust to heterogeneous movement kinematics; 2) the actual one-shot action recognition step, based on evaluating the similarity between an anchor action and target motion encodings. We propose a set of improvements to this final action recognition step designed to achieve a more accurate action recognition in the wild. These improvements include an extended anchor action representation and a dynamic threshold that discriminates challenging action sequences. Besides, the proposed action recognition approach can be easily extended to a few-shot problem, if multiple anchor actions are available.

The presented approach is validated on a generic and public one-shot action recognition benchmark, the NTU RGB+D 120 dataset [10], where it outperforms the existing baseline results. Besides, we exhaustively analyze the performance of our system on data from a real application, automatic analysis of therapies with autistic people (based on gesture and action imitation games). Evaluation shows our proposed improvements to work in the wild, managing to overcome challenging motion artifacts that are particular for this real environment. Final outcome provides online and real-time quantitative and qualitative results, essential to evaluate the patient attention and coordination.

## 2. Related Work

This section summarizes the most relevant contributions for skeleton-based action recognition, making special emphasis on N-shot action recognition methods.

### 2.1. Skeleton representations for action recognition

Although many skeleton-based action recognition approaches use the raw Cartesian pose coordinates with no special pre-processing [17, 10], other works research on skeleton data transformations to achieve view-invariant coordinates or compute new geometric features. The approach in [23] trains a Variational Autoencoder to estimate the most suitable observation view-points, and transforms the skeletons to those view-points. Other approach applies view-invariant transformations to the skeleton coordinates [20]. Regarding the computation of new skeleton features, earlier work [2] computes multiple geometric features including joint distances and orientation, distances between joints, lines and planes, velocity and acceleration. More recent approaches describe a set of geometric features based on distances between joints and lines [24] or propose to use pair-wise euclidean joint distances and two-scale motion speeds [22].

Our approach uses a skeleton representation that combines a view-invariant transformation of the skeleton Cartesian coordinates with the computation of a set of additional geometric features.

### 2.2. Skeleton-based neural networks for action recognition

Recurrent neural networks have been widely applied to learn temporal dependencies for action recognition approaches. [20] uses a Variational Autoencoder with bidirectional GRU layers for unsupervised training, and work in [12] presents a recurrent attention mechanism that improves iteratively. Other approaches rely on 1D Convolutional Neural Networks (CNN), such as [22], that uses a ResNet backbone, and [23] in which a CNN is combined with LSTM networks. Another common architecture within action recognition approaches is Graph Convolutional Networks, such as the models proposed in [13] and [3], that reduce the computational complexity with flexible receptive fields.

In our method, a Temporal Convolutional Network (TCN) ([1, 15]) is chosen to encode temporal action segments into fixed-length representations. TCNs have already shown a good action recognition performance, easing the interpretability of their results [8]. TCNs use one-dimensional dilated convolutions to learn long-term dependencies from variable length input sequences. Convolutions allow parallelizing computations allowing fast inference and performing equally or even better than RNNs in sequence modeling tasks by exhibiting longer memory [1].

### 2.3. N-shot action recognition

N-shot action recognition is still an active area of research. We find earlier methods like [5] and [9] that use HOF and HOG features for one-shot action recognition in RGB+D videos. More recently, [10] and [11] use a 2D Spatio-Temporal LSTM Network, the latter with a bidirectional pass. Other applications join visual and semantic information to perform zero-shot action recognition, such as [6], that uses hierarchical LSTMs and word2vec [14], and [7], that uses a Relation Network, a Spatial Temporal Graph Convolution Network (ST-GCN) and sent2vec. [16].

Different from these methods, we propose the simple yet effective TCN described above as a feature extractor for one-shot action recognition. Additionally, we show how a robust variation of the approach can boost the final recognition performance in challenging real world scenarios.

## 3. Proposed Framework

Figure 2 summarizes our motion description approach. It first normalizes input streams of skeleton data (computed using standard techniques [26]), it generates sets of pose features that are encoded into motion descriptors by a




Figure 2: Motion descriptor generation. Input skeleton coordinates  $X_n$  are normalized  $\bar{X}_n$  and pose features  $M$  are calculated. Pose features are processed by a TCN to generate motion descriptors  $z$ .

TCN. One-shot action recognition is performed evaluating the similarity between motion descriptors from anchor and target sequences.

### 3.1. Pose normalization

A human movement is defined by a set of  $N$  poses  $X = \{X_1, \dots, X_N\}$ . Each pose  $X_n$  is defined as a set of  $J$  3D body keypoint coordinates,  $X_n = \{x_1^n, \dots, x_J^n\}, x_j^n \in \mathbb{R}^3$ , composing what we name a skeleton.

Human actions are frequently recorded in dynamic scenarios that involve different view-points and people moving and interacting freely. To achieve a better action recognition generalization, we normalize skeleton data by applying a per-frame coordinate transformation from the original coordinate system  $W$  to a new one  $H$ , obtaining new **view** and **location-invariant** coordinate sets. As represented in the Fig. 3,  $H$  is set to have its origin at the middle point of the vector composed by the two hip keypoints. This hip vector is aligned to the new  $X$  axis and oriented to always have left and right hip  $X$  coordinates negative and positive values respectively. Similarly, the vector composed by the spine keypoint and the origin becomes the  $Y$  axis, leaving the  $Z$  axis to describe the depth information by being orthogonal to  $X$ - $Y$ . The corresponding transformation  $T^{HW}$  is applied to each 3D point in a pose  $X_n$  to obtain the new set of 3D keypoint coordinates  $\hat{X}_n$  as:

$$\hat{X}_n = T_n^{HW} X_n \quad (1)$$

Regardless of the camera configuration, action sequences can be performed by different people with heterogeneous heights. To get **scale-invariant** coordinates  $\bar{X}_n$ , each skeleton  $\hat{X}_n$  is scaled to a predefined size. In particular, since the joints defining the torso usually present little




Figure 3: Skeleton representation.  $W$  and  $H$  refer to the original and transformed skeleton coordinate systems respectively.  $H$  axis are aligned with the vectors (dashed lines) that cross the human hip and spine.  $\varphi$  and  $\theta$  angles refer to the elevation and azimuth calculation in a bone from the leg.

noise, we scale each skeleton to have a fixed torso length  $\bar{L}$ :

$$\bar{X}_n = \hat{X}_n * \frac{\bar{L}}{L_n} \quad (2)$$

where  $L_n$  is the length of the corresponding torso and  $\bar{L}$  is set as the average ratio between the torso length and the height (calculated from the NTU RGB+D 120 Dataset).

Figure 4 shows original and normalized skeleton coordinates after applying the two proposed normalization steps.

### 3.2. Pose Features

The presented approach includes as final pose features the normalized coordinates  $\bar{X}_n$ , described above, and the following additional geometric features:

- **Pair-wise keypoint distances**  $P_n$ , calculated as the Euclidean distance between each possible pair of skeleton joints, encode the pose relative to the  $J$  skeleton joints. This set of features has the size of  $\binom{J}{2}$ .
- **Bone angles**  $B_n$  from the original coordinates, calculated as the elevation  $\varphi$  and azimuth  $\theta$  (see Fig. 3) of each vector composed by two connected joints. These angles encode the orientation of each bone relative to the world. This set of features has the size of  $b \times 2$ , being  $b$  the number of bones within the skeleton.

### 3.3. Motion descriptor generation from a TCN

In order to generate motion representations based on the temporal context and not only static pose features, we use a

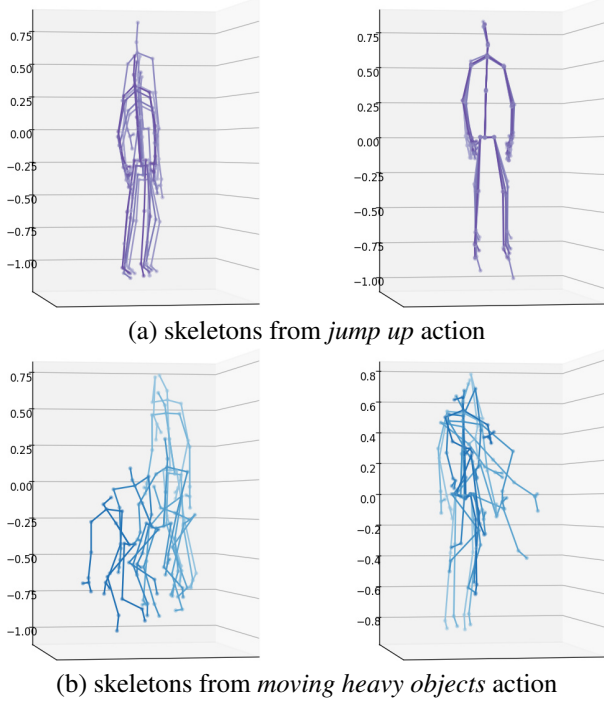


Figure 4: Pose normalization. Left plots show original skeleton coordinates of different actions from the NTU-120 RGB-D dataset. Right plots represent the same skeletons after applying the proposed pose normalization.

Temporal Convolutional Network (TCN) [1, 15]. The TCN processes, as illustrated in Fig. 2, streams of pose features  $M = \{\bar{X}, P, B\}$ , and obtains motion embeddings  $z$  (or descriptors). This motion generation works in an online fashion, creating embeddings  $z_n = TCN(M_{n-w:n})$  that encode, at the time  $n$ , all the motion from the last  $w$  frames. This receptive field (memory)  $w$  is implicitly defined by the TCN hyperparameters (details in Section 4.1.2).

### 3.4. One-shot action recognition

We formulate the one-shot action recognition as a simple similarity evaluation between the anchor embedding  $z_a$  calculated to describe an anchor action and a stream of target embeddings  $z_T$  extracted from a full video sequence.

For the **anchor action description**, we use the embedding associated to the last frame of the anchor action, i.e.  $z_a$ , assuming it encodes all the relevant previous motion information. Then, the **evaluation distance** at time  $n$  is given by the distance between the anchor embedding and the target embedding  $z_T(n)$  computed at time  $n$ :

$$d_1(n) = D(z_a, z_T(n)) \quad (3)$$

For the **embedding distance computation**  $D$  we have explored several options. The cosine distance (*cos*) and the

Jensen-Shannon divergence (*JS*) are the two best performing ones:

$$D_{cos}(z_1, z_2) = 1 - \frac{z_1 \cdot z_2}{\|z_1\| \|z_2\|}, \quad (4)$$

$$D_{JS}(z_1, z_2) = \sqrt{\frac{KL(z_1 \parallel \bar{z}_{1,2}) + KL(z_2 \parallel \bar{z}_{1,2})}{2}} \quad (5)$$

where  $z_1$  and  $z_2$  are two motion embeddings,  $KL$  is the *Kullback-Leibler divergence* and  $\bar{z}_{1,2}$  is the pointwise mean of  $z_1$  and  $z_2$ . Both functions are bounded between 0 and 1, being this last value the lowest similarity between two movement descriptors.

The **final action recognition** is performed by thresholding the calculated distance. If  $d_1(n)$  is below the acceptance threshold  $\alpha$ , we consider that the anchor action has been detected by frame  $n$ . This threshold value is set by evaluating the precision/recall curve over an evaluation dataset, as detailed in Section 4.3.1.

### 3.5. Improving action recognition in the wild

Real action recognition applications (e.g. real medical therapies) involve artifacts that hinder the motion description and recognition. These issues are intensified in the one-shot recognition set-up, where the available labeled data is limited. In the following, we describe different improvements for the action recognition described previously, to get better performance in the wild.

**Extended anchor action representation.** Anchor actions are not only scarce but frequently also hard to consistently segment in a video sequence. Therefore, using just the last embedding generated for them can lead to noisy action descriptions. In order to get a better anchor action representation, we use a set of descriptors  $z_A = \{z_1, \dots, z_m\}$  composed by the ones generated at their last  $m$  frames. The distance between a target embedding  $z_T(n)$  and the anchor action is then set as the minimum distance to each element of the set of anchor embeddings ( $z_A$ ):

$$d_m(n) = \min_{\forall z_a \in z_A} D(z_a, z_T(n)) \quad (6)$$

**Few-shot recognition.** In case more than one anchor sequence is available, the set of anchor embeddings  $z_A$  can be easily augmented with the embeddings generated for the different anchor sequences.

**Dynamic threshold.** Most challenging scenarios can include actions consisting of maintaining a static pose or performing subtle movements. The descriptors generated from this type of actions are very similar to the descriptors from target idle poses (e.g. when no specific action is being performed). If information of idle positions is available, we propose to use it to set a dynamic threshold that can better





(a) NTU RGB+D 120 dataset frames.



(b) Therapy dataset frames.

Figure 5: Sample frames from the training (a) dataset (*kick something*, *handshaking* and *shoot at the basket* actions) and evaluation (b) datasets (*big*, *give* and *high* gestures).

discriminate idle poses from actual actions. New threshold value is set to be the minimum between the original threshold  $\alpha$  and the 10<sup>th</sup> percentile  $P_{10\%}$  of all the distances computed between a given idle target sub-sequence (identified within the target sequence) and the anchor action representation:

$$\alpha = \min(\alpha, P_{10\%} \{d(n)\}), \quad (7)$$

$$n=a \dots b$$

where  $a$  and  $b$  are the initial and final time steps of the idle interval, and  $d$  is computed as in Eq. (4) or Eq. (6) depending on the anchor representation used.

## 4. Experiments

This section details experimental validation of the proposed action recognition approach, including implementation and training details.

### 4.1. Experimental setup

#### 4.1.1 Datasets

The proposed method is designed for one-shot online action recognition in real scenarios. In particular, our motivation is to automate the analysis of medical therapies. Since the available data in this setting is scarce, real therapy data is only used for evaluation, while a large public action dataset (NTU RGB+D 120) is used to learn the action encoding model as well as to validate the motion representation framework in a public benchmark.

The **NTU RGB+D 120** dataset [10] (see Fig. 5a) is so far the largest dataset available for skeleton-based action classification. It contains 114,480 action samples belonging to 120 different categories with different complexities. Action sequences are recorded at 30 fps with different users, setups and points of view.

The **therapy dataset**<sup>2</sup> (see Fig. 5b) has been acquired in real medical settings and consists of 57 different imitation games where the patient has to imitate the therapist [19, 18] action and gestures. Each imitation game is composed of at least one anchor action and one patient imitation, but frequently more anchors (up to 3) or imitations appear in the games. Due to the nature of this data, patient motion is characterized by strong artifacts such as highly variable length imitations, uncommon resting poses and poor quality imitations. Due to limited computational resources, sessions are recorded with a variable low frame rate of 12 fps on average.

#### 4.1.2 Implementation and training details

Due to the nature of the therapy dataset, we have specified the memory length  $w$  of our TCN to be 32 frames long by using a convolutional kernel size of 4, stacks of 2 residual blocks, and dilations of 1, 2 and 4 for the layers within each convolutional block. We also use skip connections and 256 filters in each convolutional layer, which finally generates motion descriptors of size 256.

This TCN is pretrained for a classification task (categorical cross-entropy loss) in the NTU RGB+D 120 dataset by adding an output classification layer to the TCN backbone and applying the following random data augmentation to the original coordinates:

- **Movement speed variation.** Joint coordinates are randomly scaled by interpolation along the time dimension to simulate varying movement speeds.
- **Skip frames.** Since the training dataset is recorded at more than twice the frame rate of the evaluation dataset, for each training sample we just use one out of either two or three frames, discarding the rest.
- **Horizontal flip.** Coordinates that correspond to the left or right part of the body are randomly flipped. Vertical and depth dimensions remain as originally.
- **Random cropping.** Actions longer than the receptive field are randomly cropped to fit the TCN memory length. Shorter ones are pre-padded with zeros.

Finally, after calculating and stacking all the pose features  $M$  the input training data has a size of  $32 \times 423$ .

<sup>2</sup><http://vislab.isr.tecnico.ulisboa.pt/datasets/#autism>

Method	Accuracy
ST-LSTM [11] + Average Pooling	42.9%
ST-LSTM [11] + Fully Connected	42.1%
ST-LSTM [11] + Attention	41.0%
APSR[10]	45.3%
<b>Ours (one-shot, <math>m=1</math>)</b>	<b>46.5 %</b>

Table 1: NTU RGB+D 120 Dataset One-shot evaluation

After training, the output classification layer is discarded to extract the 256-dimensional motion descriptors straight from the TCN backbone.

## 4.2. Validation of the system on a generic one-shot action recognition benchmark

Table 1 shows the classification accuracy of our framework for one-shot action recognition on the NTU-120 one-shot 3D action recognition problem, where we outperform the results of previous action recognition models [10]. For this benchmark, we use the same implementation described in Section 4.1.2 and the same training procedure, but with the train/test splits specified in the original paper [10]. Additionally, we also set the frame skipping to 2 both for training and evaluation and we suppress the random horizontal flipping during training. Descriptors are evaluated with the cosine distance, but other distance functions such as the Jensen-Shannon divergence report identical or comparable performance. Since this problem is about classifying action segments and not identifying anchor actions in the wild, classification is performed by assigning, for each evaluation sample, the class with the lowest similarity distance among the set of anchor action segments. Note that, unlike the other action recognition models that work in an offline fashion, our model works in an online and real-time fashion and only uses half of the available data (alternate frames).

## 4.3. Validation and discussion on real therapies data

This experiment evaluates the performance of our action recognition approach on the imitation games from the therapy dataset. Here, the actions from the therapist are taken as anchor to later detect their imitation at each time-step of the patient motion stream.

### 4.3.1 Quantitative results

**Metrics.** Action recognition is evaluated with common metrics, i.e. precision, recall and F1. However, since each target motion descriptor  $z_T(n)$  does not refer only to a single time-step  $n$  but to the motion from the  $w$  previous frames, we have defined the following terms as follows:

- **True Positive (TP).** An action is correctly detected when the ground truth action is being executed or has




Figure 6: Precision/recall curves for the two evaluation distances used: cosine distance (cos) and Jensen-Shannon divergence (js). Solid lines refer to regular one-shot evaluations ( $m = 1$ ) and dashed lines refer to few-shot recognition with extended anchor representations ( $m = 3$ ). Red crosses refer to the optimal F1 values.

been recently executed within the TCN receptive field  $w$ . A groundtruth action referenced by many detections only counts as one TP.

- **False Positive (FP).** An action is incorrectly detected when no groundtruth action has been recently executed. Consecutive detections inside the TCN receptive field  $w$  only count as one FP.
- **False Negative (FN).** A groundtruth action is missed if it has not been referenced by any action detection. Each miss-detected action counts as a single FN.

**Precision-Recall trade off.** The threshold  $\alpha$  selected to be used over the distance scores to perform action recognition is the one that optimizes the trade-off between precision and recall defined by the F1 metric. In our experiments, we calculate different thresholds to optimize the results achieved when using one descriptor per anchor ( $m = 1$ ) and with extended anchor representations ( $m = 3$ ). As seen in the precision-recall curves from Fig. 6, the best trade-off is achieved at the cost of lowering the precision of the framework, especially when working with single anchor descriptors.

**Analysis of variations of our approach.** The effect of the different variations proposed to make our action recognition more robust are summarized in Table 2. First three rows show the performance of the one-shot version, i.e., using as anchor the last action performed by the therapist in each imitation game. The use of extended anchor representations ( $m = 3$ ) allows us to have a more restrictive threshold  $\alpha$ , which reduces false positives, increasing the detection precision. The increase in precision is also achieved by

Recognition Modality	m	Dynamic Threshold	Cosine distance				Jensen-Shannon divergence			
			$\alpha$	Precision	Recall	F1	$\alpha$	Precision	Recall	F1
one-shot	1		0.48	0.655	0.922	0.697	0.50	0.689	0.902	0.714
one-shot	3		0.40	0.773	0.837	0.724	0.48	0.765	0.853	0.728
one-shot	3	✓	$\leq 0.40$	0.803	0.837	<b>0.751</b>	$\leq 0.48$	0.779	0.853	<b>0.739</b>
few-shot*	1		0.48	0.784	0.827	0.731	0.50	0.785	0.846	0.744
few-shot*	3		0.40	0.760	0.876	0.753	0.48	0.749	0.892	0.755
few-shot*	3	✓	$\leq 0.40$	0.790	0.876	<b>0.781</b>	$\leq 0.48$	0.763	0.892	<b>0.765</b>

\* few-shot stands for the use of 1, 2 or 3 anchor sequences depending on the available reference data for each test, as explained in Section 4.1.1

Table 2: Comparison of the studied variations of our action recognition approach on the evaluation therapy dataset.

making the threshold dynamic. New threshold values can be set differently in each imitation game by evaluating the anchor representation and the target motion up the end of the anchor execution (period in which the patient is not performing any specific action). The dynamic threshold avoids false positives related to "static" anchor actions, at no recall cost. Finally, as expected, the performance improves when more anchor actions (up to 3 in the experiments) are available to run a few-shot recognition. Regarding the embedding similarity computation, both the cosine distance and the JS divergence report similar performance when using extended anchor representations. However, the JS divergence performs better in simpler set-ups and the cosine distance excels when using a dynamic threshold.

The results can be further scrutinized in Table 3, that compares the recognition performance by action category. Extended anchor representations ( $m = 3$ ) and few-shot recognition significantly overcome the limitations of the simple one-shot recognition for challenging classes (last 8 rows of the table refer to actions with softer and subtler arm movements or actions that consist in staying on specific positions). This is further improved when using a dynamic threshold (DT).

**Influence of different pose feature sets.** Table 4 reports the influence of different pose features proposed in Section 3.2. Using only normalized coordinates achieves lower performance than using just the original body coordinates. Normalized coordinates gain in invariance, but they lose certain discriminative power. The proposed geometric feature set achieves good performance on its own. However, the best performance comes from the combination of normalized coordinates and geometric features, which suggests that the geometric features work better along with the invariance provided by normalized coordinates.

#### 4.3.2 Qualitative results

Figure 7 shows the timeline of an imitation game sampled from the therapy dataset. This exercise involves the action of raising and moving the arms, which is performed once

Class name	One-shot (m=1)	Few-shot (m=3)	Few-shot+DT (m=3)
<b>big</b>	1.00	1.00 (+0.00)	1.00 (+0.00)
<b>high</b>	1.00	1.00 (+0.00)	1.00 (+0.00)
<b>happy</b>	0.96	0.96 (+0.00)	0.96 (+0.00)
<b>waving</b>	0.89	0.95 (+0.06)	0.95 (+0.06)
<b>pointing</b>	0.89	0.89 (+0.00)	0.89 (+0.00)
<b>giving</b>	0.76	0.76 (+0.00)	0.76 (+0.00)
<b>small</b>	0.68	0.89 (+0.21)	0.96 (+0.28)
<b>coming</b>	0.65	0.69 (+0.04)	0.69 (+0.04)
<b>waiting</b>	0.63	0.73 (+0.10)	0.73 (+0.10)
<b>hungry</b>	0.47	0.35 (-0.12)	0.42 (-0.05)
<b>where</b>	0.47	0.61 (+0.14)	0.61 (+0.14)
<b>down</b>	0.45	0.54 (+0.08)	0.54 (+0.08)
<b>me</b>	0.44	0.67 (+0.22)	0.67 (+0.22)
<b>angry</b>	0.31	0.30 (-0.01)	0.50 (+0.19)

Table 3: Per-class F1 evaluation comparison of different variations of our action recognition approach (using the cosine distance).

Coordinates	Normalized Coordinates	Geometric Features	F1
✓			0.711
	✓		0.689
		✓	0.757
✓		✓	0.695
	✓	✓	<b>0.781</b>

Table 4: F1 evaluation of different pose features running our top-performing action recognition (few-shot, cosine distance, extended anchor representations and dynamic threshold).

(anchor) by the therapist and repeated twice (target) by the patient. Detection results from the timeline represent both the location and quality of both action repetitions. The latter is estimated according to the calculated similarity between anchor and target motion. The first action imitation is performed quickly (worse quality) and the second one slower and more detailed (better quality). For a better understanding, the whole execution of the experiment is shown in the

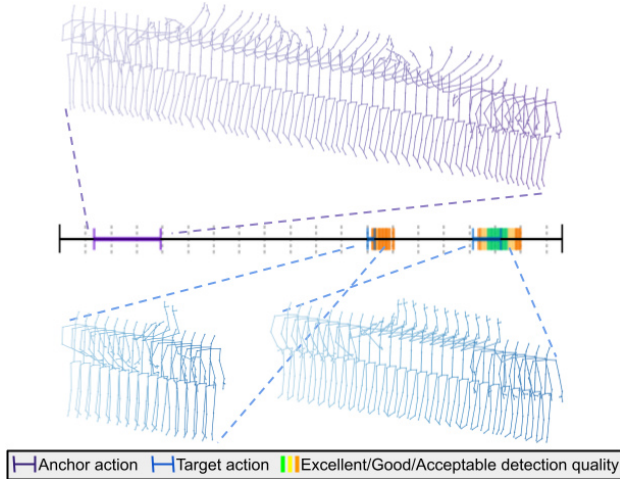


Figure 7: Action recognition on one imitation game from the therapy dataset. Central horizontal line represents the execution timeline, with one anchor action (top/purple) performed by the therapist, followed by two imitations (bottom/blue) performed by the patient. Dashed line segments indicate the groundtruth action time segmentations. Vertical small solid segments, ranging from green to orange, refer to time steps where the action has been detected, and the color gives an estimation of the repetition quality.

supplementary video<sup>3</sup>.

#### 4.3.3 Time performance

Temporal Convolutional Networks are frequently lightweight models, making them suitable for low-resource environments like in the therapies described in Section 4.1.1. With the settings described in Section 4.1.2, the action recognition (except the skeleton extraction from RGB-D data) takes 0.08 ms per time-step using the GPU and 0.1 ms just using the CPU<sup>4</sup>. Time performance has been calculated with the therapy dataset skeleton sequences. Note that due to the TCN parallel computing, the longer the motion sequences are, the faster per-frame processing we can get.

## 5. Acknowledgments

This research has been funded by FEDER/Ministerio de Ciencia, Innovación y Universidades/Agencia Estatal de Investigación RTC-2017-6421-7 and PGC2018-098817-A-I00, DGA T45 17R/FSE, the Office of Naval Research Global project ONRG-NICOP-N62909-19-1-2027, Universidad de Zaragoza, Fundación Bancaria Ibercaja and Fundación CAI IT 17/19, and the Fundação para a Ciência

<sup>3</sup><https://sites.google.com/a/unizar.es/filovi/>

<sup>4</sup>Performing speed has been calculated with a NVIDIA GeForce RTX 2080 Ti (GPU) and a Intel Core i7-9700K (CPU)

e Tecnologia (FCT) project UIDB/50009/2020 and PhD scholarship SFRH/BD/145040/2019.

## 6. Conclusions

This work presents a novel skeleton-based framework for one-shot action recognition in the wild. Our method generates motion descriptors robust different motion artifacts and variable kinematics. We achieve accurate action recognition by combining the proposed set of pose features and an efficient architecture based on a TCN that encodes these pose features. Besides, we demonstrate the effectiveness of several simple steps included in our method to boost the action recognition performance in challenging real world scenarios, i.e. with limited reference data and noisy repetitions of the actions to be recognized. Our base one-shot recognition approach is evaluated on the public NTU RGB+D 120 dataset, outperforming previous action recognition methods, by using only half of the available data. We also demonstrate the suitability of our framework to analyze videos from real therapies with autistic people. These recordings are characterized by having extreme motion artifacts. Evaluation results provide both quantitative and qualitative measures about the actions recognized, which is essential for the patient evaluation and monitoring. Moreover, our approach can run online, being able to provide immediate feedback to the therapist and making the sessions more dynamic.

## References

- [1] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *Universal Language Model Fine-tuning for Text Classification*, 2018.
- [2] Cheng Chen, Yueting Zhuang, Feiping Nie, Yi Yang, Fei Wu, and Jun Xiao. Learning a 3d human pose distance metric from geometric pose descriptor. *IEEE Transactions on Visualization and Computer Graphics*, 17(11):1676–1689, 2010.
- [3] Ke Cheng, Yifan Zhang, Xiangyu He, Weihang Chen, Jian Cheng, and Hanqing Lu. Skeleton-based action recognition with shift graph convolutional network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 183–192, 2020.
- [4] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015.
- [5] Sean Ryan Fanello, Ilaria Gori, Giorgio Metta, and Francesca Odone. One-shot learning for real-time action recognition. In *Iberian Conference on Pattern Recognition and Image Analysis*, pages 31–40. Springer, 2013.
- [6] Meera Hahn, Andrew Silva, and James M Rehg. Action2vec: A crossmodal embedding approach to action learning. 2019.



- [7] Bhavan Jasani and Afshaan Mazagonwalla. Skeleton based zero shot action recognition in joint pose-language semantic space. *arXiv preprint arXiv:1911.11344*, 2019.
- [8] Tae Soo Kim and Austin Reiter. Interpretable 3d human action analysis with temporal convolutional networks. In *2017 IEEE conference on computer vision and pattern recognition workshops (CVPRW)*, pages 1623–1631. IEEE, 2017.
- [9] Jakub Konečný and Michal Hagara. One-shot-learning gesture recognition using hog-hof features. *The Journal of Machine Learning Research*, 15(1):2513–2532, 2014.
- [10] Jun Liu, Amir Shahroudy, Mauricio Lisboa Perez, Gang Wang, Ling-Yu Duan, and Alex Kot Chichung. Ntu rgb+d 120: A large-scale benchmark for 3d human activity understanding. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [11] Jun Liu, Amir Shahroudy, Dong Xu, Alex C Kot, and Gang Wang. Skeleton-based action recognition using spatio-temporal lstm network with trust gates. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):3007–3021, 2017.
- [12] Jun Liu, Gang Wang, Ping Hu, Ling-Yu Duan, and Alex C Kot. Global context-aware attention lstm networks for 3d action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1647–1656, 2017.
- [13] Ziyu Liu, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, and Wanli Ouyang. Disentangling and unifying graph convolutions for skeleton-based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 143–152, 2020.
- [14] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- [15] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016.
- [16] Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. Unsupervised learning of sentence embeddings using compositional n-gram features. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 528–540, 2018.
- [17] Mauricio Perez, Jun Liu, and Alex C Kot. Interaction relational network for mutual action recognition. *IEEE Transactions on Multimedia*, 2021.
- [18] Laura Santos, Silvia Annunziata, Alice Geminiani, Elena Brazzoli, Arianna Caglio, José Santos-Victor, Alessandra Pedrocchi, and Ivana Olivieri. Interactive social games with a social robot (IOGIOCO): Communicative gestures training for preschooler children with autism spectrum disorder. In *I Congresso Annuale Rete IRCCCS Neuroscienze e Neuroriabilitazione*, 2020.
- [19] Laura Santos, Alice Geminiani, Ivana Olivieri, José Santos-Victor, and Alessandra Pedrocchi. Copyrobot: Interactive mirroring robotics game for asd children. In Jorge Henriques, Nuno Neves, and Paulo de Carvalho, editors, *XV Mediterranean Conference on Medical and Biological Engineering and Computing – MEDICON 2019*, pages 2014–2027, Cham, 2020. Springer International Publishing.
- [20] Kun Su, Xiulong Liu, and Eli Shlizerman. Predict & cluster: Unsupervised skeleton based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9631–9640, 2020.
- [21] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638, 2016.
- [22] Fan Yang, Yang Wu, Sakriani Sakti, and Satoshi Nakamura. Make skeleton-based action recognition model smaller, faster and better. In *Proceedings of the ACM Multimedia Asia*, pages 1–6, 2019.
- [23] Pengfei Zhang, Cuiling Lan, Junliang Xing, Wenjun Zeng, Jianru Xue, and Nanning Zheng. View adaptive neural networks for high performance skeleton-based human action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1963–1978, 2019.
- [24] Songyang Zhang, Xiaoming Liu, and Jun Xiao. On geometric features for skeleton-based action recognition using multilayer lstm networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 148–157. IEEE, 2017.
- [25] Yimeng Zhang, Xiaoming Liu, Ming-Ching Chang, Weina Ge, and Tsuhan Chen. Spatio-temporal phrases for activity recognition. In *European Conference on Computer Vision*, pages 707–721. Springer, 2012.
- [26] Zhengyou Zhang. Microsoft kinect sensor and its effect. *IEEE multimedia*, 19(2):4–10, 2012.