

# Temporal Attention-Augmented Graph Convolutional Network for Efficient Skeleton-Based Human Action Recognition

Negar Heidari and Alexandros Iosifidis

Department of Electrical and Computer Engineering, Aarhus University, Denmark  
Emails: {negar.heidari, alexandros.iosifidis}@eng.au.dk

**Abstract**—Graph convolutional networks (GCNs) have been very successful in modeling non-Euclidean data structures, like sequences of body skeletons forming actions modeled as spatio-temporal graphs. Most GCN-based action recognition methods use deep feed-forward networks with high computational complexity to process all skeletons in an action. This leads to a high number of floating point operations (ranging from 16G to 100G FLOPs) to process a single sample, making their adoption in restricted computation application scenarios infeasible. In this paper, we propose a temporal attention module (TAM) for increasing the efficiency in skeleton-based action recognition by selecting the most informative skeletons of an action at the early layers of the network. We incorporate the TAM in a light-weight GCN topology to further reduce the overall number of computations. Experimental results on two benchmark datasets show that the proposed method outperforms with a large margin the baseline GCN-based method while having  $\times 2.9$  less number of computations. Moreover, it performs on par with the state-of-the-art with up to  $\times 9.6$  less number of computations.

## I. INTRODUCTION

Human action recognition has been a very popular research topic in recent years. RGB videos and different types of modalities such as depth, optical flow and human body skeletons can be used for this task [1], [2], [3]. Compared to other data modalities, human body skeletons encode compact and high level information representing the human pose and joints' motion while being invariant to viewpoint variations, motion speed, human appearance, and body scale [4], and it is robust to context noise. Considering the success of existing pose estimation techniques [5], [6], [7] and the availability of depth cameras [8], obtaining the human body skeleton data is much easier than before. Thus, skeleton-based human action recognition has attracted an increasing research interest in recent years and many deep learning-based methods were proposed to model both the spatial and temporal evolution of skeletons in a sequence. Some methods use recurrent neural network (RNN) architectures, like the Long Short-Term Memory (LSTM) network [9], which are suitable to model temporal dynamics [1], [2], [10], [11], [12], [13] and utilize the skeleton data as vectors formed by the body joints coordinates. Methods employing Convolutional Neural Networks (CNN) [14], [15], [16], [17], [18], [19] reorganize the body joints' coordinates of each pose to a 2D map which is a suitable

input format for CNNs. The high model complexity in all these methods make their training and inference processes very time consuming [16], [20]. Besides, these methods are not able to completely benefit from the non-euclidean structure of the skeleton data which represents the spatial body joints connections.

Graph Convolutional Networks (GCNs) have been very successful when applied to many pattern recognition tasks [21], [22], [23], [24], [25], [26], [27], [28], by generalizing the convolution operation from grid data into graph data structures. Recently, significant results have been obtained by employing GCNs for skeleton-based human action recognition [29], [30], [31], [32], [33], [34], [35]. In these methods, an action is represented as a sequence of body poses and each body pose is represented by a skeleton. The skeleton data is treated as a graph which models the spatial relationship between different body joints and the temporal dynamics in an action are expressed by a sequence of skeletons.

Most of the recently proposed GCN-based methods use deep feed forward networks to model the spatio-temporal features of body skeletons. Considering that these methods process all the body skeletons in a sequence depicting the performed action, this approach is not efficient in terms of memory consumption and computation time. While memory efficiency can be increased by employing network compression, weight pruning, quantization and low-rank approximation approaches [36], [37], [38], [39], by processing a large number of body skeletons in each sequence, the number of floating point operations (FLOPs) is still large. Thus, to address both memory and computational efficiency in skeleton-based human action recognition, not only we need more compact and lightweight network architectures, but also the number of FLOPs should be minimized. Hence, processing fewer body skeletons for action recognition is a large step towards increasing the computational efficiency of both training and inference processes.

In this paper, we argue that all body skeletons in a temporal sequence are not equally important for recognizing actions. For each action class, there exist body poses which are the most informative for the action, and we can extract sufficient information for action recognition by focusing on the skeletons of these body poses only. Our goal is to increase computational

efficiency while performing on par, or even better, compared to models which utilize all the body skeletons in a sequence for action recognition. In this regard, we propose a GCN-based model which is capable to select a subset of body skeletons for human action recognition. To select the most informative skeletons, we propose a trainable temporal attention module which measures the importance of each skeleton in a sequence and we employ this module in the GCN-based spatio-temporal model to increase its efficiency. The main contributions of our work are the following:

- We propose a temporal attention module (TAM) to extract the most informative skeletons in an action sequence, leading to increased computational efficiency in both the training phase and inference.
- We experimentally show that employing our proposed TAM with a light-weight GCN topology leads to state-of-the-art performance levels in two widely adopted datasets for skeleton-based action recognition.
- We show that a subset of skeletons is as informative as the full skeleton sequence, since our method performs on par with the state-of-the-art methods while increasing computational efficiency by a factor of up to  $\times 9.6$ . This is a major advantage of our method compared to the state-of-the-art, as it is suitable for real time action recognition under restricted computation scenarios.

The remainder of the paper is organized as follows. Section II discusses the related works and section III introduces the baseline method ST-GCN [30]. Section IV describes the proposed method. The conducted experiments and results are presented in section V, and the concluding remarks are drawn in section VI.

## II. RELATED WORK

Data driven skeleton-based human action recognition methods which use deep learning models are mainly categorized into RNN-based, CNN-based and GCN-based methods. RNN-based methods [1], [2], [10], [11], [12], [13] mostly employ LSTM networks [9] to model the temporal dynamics of skeletons' sequence. In these methods, the sequence of skeletons is introduced to the model as a sequence of vectors. Each vector is formed by the concatenated 3D coordinates of all body joints of skeleton. ST-LSTM [2] is one representative RNN-based method which captures the body joints' relationship in both temporal and spatial domain. CNN-based methods [40], [14], [15], [16], [17], [18], [19] convert the sequence of skeletons into pseudo-images by reorganizing the joints' coordinates into a 2D map and employ a state-of-the-art CNN model like ResNet [41], [42] to extract temporal and spatial features. Although the CNN-based methods are easier to train than RNN-based methods, they employ deep network architectures with large receptive fields to perceive the semantics of the input map. Besides, since all these RNN-based and CNN-based methods convert the skeletons into a regular grid or a sequence, they cannot utilize the complex, irregular and non-Euclidean structure of the skeleton data.

Recently, many GCN-based methods for skeleton-based human action recognition have been proposed which treat the human skeleton data as a graph structure representing the body joints (graph nodes) and their natural connections with bones (graph edges). Since the GCNs are able to capture the embedded features in irregular structured data, the GCN-based methods achieve higher performance in skeleton-based human action recognition compared to RNN-based and CNN-based methods. Spatio-temporal graph convolutional network (ST-GCN) [30] is the first GCN-based method proposed for action recognition. It receives the sequence of body skeletons directly as input and employs the GCNs' [21] aggregation rule to extract the spatial features of each skeleton in a sequence, while the temporal dynamics are modeled by using a temporal graph with fixed connections. Several methods have been proposed based on ST-GCN for skeleton based human action recognition which mostly focus on exploiting adaptive spatial graphs. 2s-AGCN [29] is one of the state-of-the-art methods proposed on top of ST-GCN [30]. This method adaptively learns the topology of the graph in different layers for each sequence of skeletons in an end-to-end manner. It updates the graph structure of the skeleton with a spatial attention and also a data dependent graph. Besides, it uses a two-stream framework to utilize both joints features and bones features in parallel. It introduces joints and bones data into two different models and, finally, the SoftMax scores of the two model are added to obtain the fused score and predict the action label. In DGNN [34] the skeleton data is represented as a directed acyclic graph to benefit from the relationship between joints and bones based on their kinematic dependency and it improved 2s-AGCN [29] by utilizing the motion data and also making the graph structure adaptive. GCN-NAS [31] follows a neural architecture search approach to explore an optimal GCN-based model in terms of graph structures. That is, it explores the search space to determine the best graph structure at each layer of the network, while it still employs the ST-GCN network topology. DPRL+GCNN [32] aims to select the most representative skeletons from the input sequence using deep reinforcement learning. It adjusts the chosen skeletons progressively by evaluating the models' performance in action recognition. Actional-Structural GCN (AS-GCN) [33] proposed an inference module based on encoder-decoder structure to capture richer action-specific correlations and also structural links between the joints in skeletons.

## III. SPATIAL TEMPORAL GRAPH CONVOLUTIONAL NETWORK (ST-GCN)

In this section, we describe ST-GCN model [30] as our baseline. It is formulated over a sequence of skeleton graphs with multiple layers of spatio-temporal graph convolution operations. The model receives the joint coordinate vectors and the constructed spatio-temporal graph as input and applies GCN to capture the embedded patterns of the graph and extract high level features of the skeletons. Each convolution layer integrates the human body information along both the spatial and temporal dimensions. The extracted features are

then passed to a fully connected layer which classifies the skeleton sequence to the action class using a standard SoftMax classifier. The entire model is trained with backpropagation to minimize the classification loss. The network architecture of ST-GCN [30] is constructed by 9 layers of spatio-temporal convolution units. In each layer the Resnet module is also applied and the output is followed by a batch normalization layer. In the following, the graph construction, and the spatial and temporal convolution units are described.

### A. Graph Construction

Each skeleton is represented as a graph which forms the hierarchical representation of skeleton sequences and the spatial patterns and temporal dynamics of these graphs are captured automatically. The spatio-temporal graph is constructed using the 2D or 3D human body joint coordinates as nodes, and spatial and temporal edges. The spatial edges correspond to the natural connectivity of body joints in the skeleton, and the temporal edges connect the same body joints across consecutive skeletons. Fig. 1 (right) shows the spatio-temporal skeleton graph.

Formally, an undirected spatio-temporal graph on a sequence of skeletons is denoted as  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where the set of nodes:

$$\mathcal{V} = \{\nu_{ti} \mid t, i \in \mathbb{Z}, 1 \leq t \leq T, 1 \leq i \leq N\}$$

indicates  $N$  body joints of a skeleton in a sequence of  $T$  time steps and  $\mathcal{E}$  is the set of spatial (intra-skeleton) and temporal (inter-skeleton) connections. The graph structure is captured by the adjacency matrix  $\mathbf{A} \in \mathbb{R}^{N \times N}$  which is a symmetric binary matrix defined as  $\mathbf{A}_{ij} = 1$  if there is a connection between nodes  $\nu_{ti}$  and  $\nu_{tj}$  in time step  $t$ , otherwise  $\mathbf{A}_{ij} = 0$ . Given  $\mathbf{A}$  as the spatial graph adjacency matrix which represents the joint connections in a single skeleton, the normalized adjacency matrix  $\hat{\mathbf{A}}$  with self connections is computed as:

$$\hat{\mathbf{A}} = \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}}, \quad (1)$$

where  $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}_N$  and  $\tilde{\mathbf{D}}$  is the diagonal degree matrix of  $\tilde{\mathbf{A}}$ , i.e.  $\tilde{\mathbf{D}}_{ii} = \sum_j \tilde{\mathbf{A}}_{ij}$ .

The body motions are grouped as concentric and eccentric, the spatially localized structure of a single skeleton is used to design the spatial partitioning process which divides the neighbor nodes of each body joint into three subsets: 1) the root node itself; 2) the nodes connected to the root which are closer to the skeletons' center of gravity than the root node; 3) the remaining roots' neighbors. Fig. 1 (left) shows the neighboring sets of each graph node which are presented with different colors [30], [29].

### B. Spatial Graph Convolution

Given the input spatial features of  $i^{th}$  joint,  $f_{in}(\nu_i)$ , the spatial graph convolution is defined as [30]:

$$f_{out}(\nu_i) = \sum_{\nu_j \in \mathcal{N}_i} \frac{1}{Z_{ij}} f_{in}(\nu_j) w(l_i(\nu_j)), \quad (2)$$

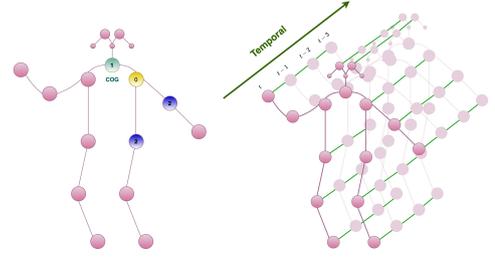


Fig. 1. Illustration of Spatio-temporal graph (right) in TA-GCN, and the spatially partitioned skeleton (left). Each partition set is illustrated in a different color.

where  $\mathcal{N}_i$  is the set of intra-frame neighbor vertices of  $\nu_i$  which are directly connected to it. As mentioned in III-A, ST-GCN uses a partitioning method to divide the neighbor vertices of each node into three subsets. According to this partitioning process,  $l_i$  in (2) maps each neighbor vertex into one of these three subsets and  $w$  is the weight matrix which is unique for each neighbor set and maps the target nodes' features into a new subspace. Note that the number of neighbor subsets and weight matrices are fixed while the number of neighbor nodes in each subset is varied for each node.  $Z_{ij}$  is a normalization factor which balances the contribution of each set of neighbor nodes in producing  $f_{out}(\nu_i)$  which contains the output features of the target node  $\nu_i$ .

In terms of implementation, the  $C$ -dimensional input feature vector for the node  $\nu_i$  is denoted as  $\mathbf{x}_i = f_{in}(\nu_i)$  and  $\mathbf{X} \in \mathbb{R}^{C_{in} \times T \times N}$  represents the input feature tensor for a sequence of skeletons, where  $C_{in}$  denotes the number of input channels,  $T$  is the number of skeletons and  $N$  is the number of body joints in each skeleton. The model receives the feature tensor  $\mathbf{X}$  as input and updates the nodes' feature vectors by applying the spatial convolution to produce  $\mathbf{X}' \in \mathbb{R}^{C_{out} \times T \times N}$  with  $C_{out}$  channels as output. By employing the layer-wise propagation rule of GCNs proposed in [21], the spatial convolution is formally defined as [30]:

$$\mathbf{X}' = ReLU \left( \sum_p (\hat{\mathbf{A}}_p \otimes \mathbf{M}_p) \mathbf{X} \mathbf{W}_p \right), \quad (3)$$

where  $\otimes$  is the element-wise product of two matrices. According to the partitioning process described in section III-A, each node has 3 subsets of neighbors. Therefore, the adjacency matrix  $\tilde{\mathbf{A}}$  is defined as the summation of 3 different adjacency matrices which are indexed by  $p$  as follows:

$$\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}_N = \sum_p \mathbf{A}_p \quad (4)$$

$\mathbf{A}_0 = \mathbf{I}_N$  represents the nodes' self connections and the normalized adjacency matrix for each subset is defined as:

$$\hat{\mathbf{A}}_p = \mathbf{D}_p^{-\frac{1}{2}} \mathbf{A}_p \mathbf{D}_p^{-\frac{1}{2}}, \quad (5)$$

where  $\mathbf{D}_{(ii)_p} = \sum_j \mathbf{A}_{(ij)_p} + \varepsilon$ , with  $\varepsilon = 0.001$  used to avoid empty rows in degree matrix.  $\mathbf{M}_p \in \mathbb{R}^{N \times N}$  is a learnable attention map which highlights the elements of

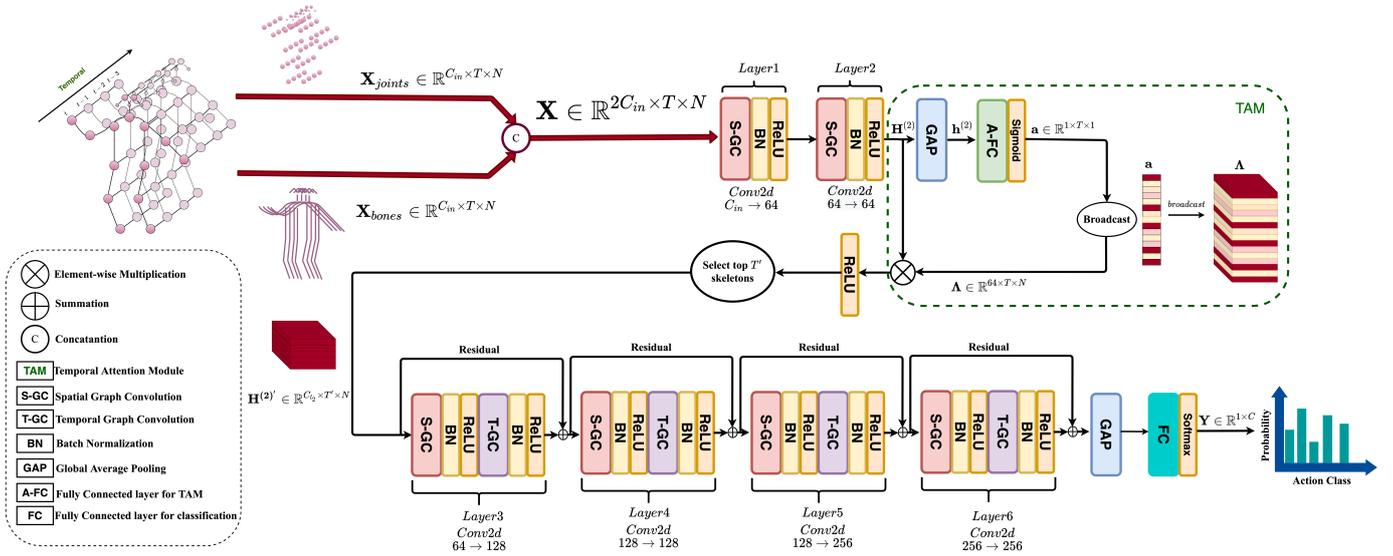


Fig. 2. Illustration of the proposed model diagram. The whole model architecture is composed of 6 spatio-temporal layers and the temporal attention module places after the first two layers to select the most informative skeletons. The highlighted skeletons' indices are sorted in descending manner and then extracted from the  $\mathbf{H}^{(2)}$ . The 3<sup>rd</sup> to 6<sup>th</sup> layers employ both spatial and temporal convolution.

each adjacency matrix and it is initialized as an all-one matrix.  $\mathbf{W}_p \in \mathbb{R}^{C_{out} \times C_{in}}$  denotes the weight matrix which transforms the node features of each partition. In practice, a standard 2D convolution is used which performs  $C_{out} \times 1 \times 1$  convolutions, then the resulting tensor is multiplied with the masked normalized adjacency matrix on the last dimension  $N$ .

### C. Temporal Graph Convolution

To capture the temporal dynamics in a skeleton sequence and consider the motions taking place in an action, we need to propagate information related to the nodes' features in other skeletons of the sequence. To this end, a temporal convolution is applied on the output tensor of the spatial convolution step. In temporal dimension, each node is only connected to its corresponding node in its previous and next skeletons, so the number of neighbors for each node is fixed to 2. Therefore, a 2D convolution with a predefined temporal kernel size  $K_t$  is applied to the output of (3),  $\mathbf{X}'$ , to aggregate the features of each body joint in different time steps.

## IV. TEMPORAL ATTENTION-AUGMENTED GRAPH CONVOLUTIONAL NETWORK (TA-GCN)

In this section we describe an end-to-end temporal attention-augmented GCN model for efficient skeleton-based human action recognition. In order to extract discriminative features in temporal dimension of data, we propose TAM which highlights the most informative skeletons in a sequence. Based on this process, we apply skeleton selection to highly reduce the overall computational cost both during training and inference.

Similar to most of the recently proposed works [30], [29], the proposed model is composed of multiple layers of spatio-temporal units. In each unit, the embedded spatial features in each skeleton are first extracted by employing a standard 2D

convolution as described in Eq. (3). Motivated by [29], the learnable spatial attention map  $M_p$  is added (and not element-wise multiplied) to the graph adjacency matrix  $A_p$ . Adding the attention map to the graph adjacency matrix has the advantage that, not only the strength of the existing graph edges can be highlighted, but relationships between graph nodes which are not connected in the adjacency matrix based on the physical connectivity of the human body skeleton joints can also be captured [29]. Such connections can be important in effectively describing actions, e.g. encoding the relationships between the joints of the human arms and the head is particularly important for describing actions like 'hand waving'. Thus, the spatial attention module is employed to selectively focus on the most informative joints in each skeleton and the spatial convolution in Eq. (3) takes the following form:

$$\mathbf{X}' = \text{ReLU} \left( \sum_p (\hat{A}_p + M_p) \mathbf{X} \mathbf{W}_p \right) \quad (6)$$

The TAM takes a tensor  $\mathbf{H}^{(l)} \in \mathbb{R}^{C_l \times T \times N}$  as input, which can be the input data, i.e.  $\mathbf{H}^{(0)} = \mathbf{X}$ , or the output of the  $l^{\text{th}}$  hidden layer of the network. First, two average pooling operations in both feature and spatial dimensions is performed to produce the  $\mathbf{h}^{(l)} \in \mathbb{R}^{1 \times T \times 1}$ , which denotes the average feature value of each skeleton in the sequence. Then, this tensor is introduced to a fully connected layer which transforms features in temporal dimension  $T$  to produce the attention tensor  $\mathbf{a} \in \mathbb{R}^{1 \times T \times 1} = \{a_0, a_1, \dots, a_T\}$  as follows:

$$\mathbf{a} = \text{Sigmoid} \left( \mathbf{h}^{(l)} \Theta \right), \quad (7)$$

where  $\Theta \in \mathbb{R}^{T \times T}$  denotes the learnable transformation matrix, and the resulted attention tensor  $\mathbf{a}$  indicates the importance of each skeleton in the sequence.

To highlight the most informative skeletons, we create the attention tensor  $\mathbf{A} \in \mathbb{R}^{C_l \times T \times N}$ , a duplicated version of attention map  $\mathbf{a}$  with  $C_l \times N$  copies, which is subsequently dot multiplied to  $\mathbf{H}^{(l)}$  as follows:

$$\hat{\mathbf{H}}^{(l)} = \text{ReLU}(\mathbf{H}^{(l)} \otimes \mathbf{A}), \quad (8)$$

where  $\otimes$  denotes element-wise multiplication. To select a subset of  $T'$  skeletons from  $\hat{\mathbf{H}}^{(l)}$ , the values in the attention map  $\mathbf{a}$  are sorted in descending order and the skeletons corresponding to the  $T'$  highest attention values are selected to be introduced into next layers of the network. To improve the computational efficiency of model in both training and inference phases, it would be preferable to use the TAM in the early layers of the network so that the next layers will process less number of skeletons.

The proposed model is composed of 6 GCN layers and 1 TAM. The first two GCN layers, map data into 64 dimensional feature space using spatial convolution followed by batch normalization layer and element-wise ReLU activation function. The mapped data is then introduced into the TAM which selects a subset of skeletons. In order to find the most discriminative skeletons using the TAM, the temporal convolution operation which smoothens the features in temporal domain is not used in the first two GCN layers. The  $3^{rd}$  and  $4^{th}$  layers change data dimensions from 64 to 128 and the last two GCN layers increase the number of channels to 256. In the last 4 GCN layers, both spatial and temporal convolutions are applied and each operation is followed by batch normalization and element-wise ReLU activation function. All the GCN layers except the first two layers utilize the ResNet [42] module to benefit from the input skeleton data too. The strides of the temporal convolution layers in  $3^{rd}$  and the  $5^{th}$  GCN layers are set to 2 as pooling layer. At the end, the refined spatio-temporal features are introduced into a global average pooling layer which produces an output feature vector of size  $256 \times 1$  for each sequence of skeletons and it is introduced into a fully connected layer which is equipped by a SoftMax classifier to classify the action. The model is trained with backpropagation in an end-to-end manner to minimize the classification error while it is learning to select the most discriminative skeletons.

Motivated by the method in [29], which utilizes both joints and bones features to enhance the classification performance, we also explore the skeleton bones' length and direction as the second-order information. Each bone is represented as a 3D vector bounded with two joints. The source joint is the one that is closer to the skeletons' center of gravity than the target joint. Therefore, each bone pointing from its source joint to the target joint holds both the length and direction information between two joints.

After extracting the bone features from the skeleton data, the two feature tensors are concatenated on their first dimension and the merged joint-bone tensor of size  $2C_{in} \times T \times N$  is introduced to the model as input. The overall architecture of the proposed method is shown in Fig. 2.

## V. EXPERIMENTS

In this section, we describe experiments evaluating the performance of the proposed TA-GCN model in skeleton based human action recognition. We conducted experiments on two widely adopted datasets for evaluating the performance of skeleton-based action recognition methods. These datasets are:

1) *The NTU-RGB+D [10]*: is the largest multi-modality indoor-captured action recognition dataset. The dataset includes RGB videos, infrared videos, 3D skeletons and depth sequences. The 3D skeleton data is captured by the Microsoft Kinect-v2 camera and is used in our experiments. This dataset consists of 56,880 video clips from 60 different human action classes which are captured from three different views. Each skeleton is represented by 25 joints which are featured by 3D coordinate values. In our experiments, we follow exactly the same data splits and benchmark evaluations as in [10]. In Cross-View (CV) benchmark, the training set contains 37,920 samples captured from cameras two and three and the test set contains 18,960 samples captured from the first camera. In Cross-Subject (CS) evaluation, 40,320 videos which represent 20 different action classes are used for training and the remaining 16,560 videos are used for testing. The number of frames for each sample is 300 and for the samples which have less than 300 frames, the frame sequence is repeated until it reaches 300 frames. In practice, the input data is a tensor of size  $(3 \times 300 \times 25)$ .

2) *The Kinetics-Skeleton [43]*: is a very large action recognition dataset that contains 300,000 video clips of 400 different human actions collected from YouTube. The 2D joints' coordinates  $(x, y)$  on every frame and their confidence score  $c$  are estimated using the public OpenPose toolbox [7]. Each skeleton in this dataset contains 18 body joints and each body joint is represented by a 3D vector  $(x, y, c)$ . The number of frames for each sample is fixed to 300 in a similar way as explained for NTU-RGB+D dataset and the input data would be a tensor of size  $(3 \times 300 \times 18)$ . We use the Kinetics skeleton data which is provided by [30]. The training and validation sets contain 240,000 and 20,000 skeleton sequences, respectively.

### A. Experimental setting

All experiments were conducted on PyTorch deep learning framework [44] with 4 GRX 1080-ti GPUs and batch size of 32 and 128 for NTU-RGB+D and Kinetics datasets, respectively. The SGD optimizer is employed to optimize the model with Cross-entropy loss function through back-propagation with weight decay set to 0.0001. We followed exactly the same setting explained by authors for the baseline method [30] and the state-of-the-art [29]. In more details, the learning rate is not fixed through all epochs. For NTU-RGB+D dataset, it starts with 0.1 and it is divided by 10 at epochs 30, 40 while the total number of training epochs is fixed to 50. For Kinetics dataset, it starts with 0.1 and it is divided by 10 at epochs 45, 55 and the total number of training epochs is fixed to 65. For the Kinetics-Skeleton dataset we don't perform data-augmentation method which is used in [30]. In more details, we utilize all the 300 skeletons in a sequence.

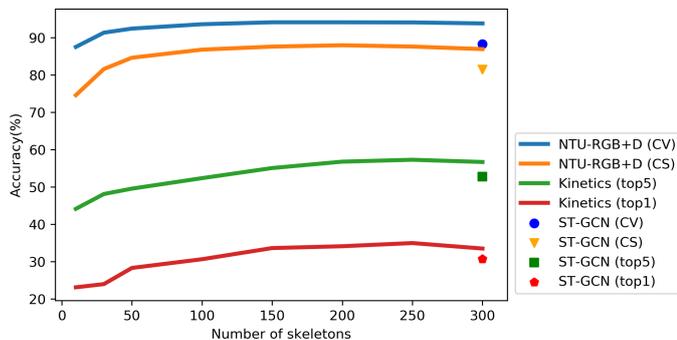


Fig. 3. The classification accuracy of the proposed TA-GCN method for a varying number of skeletons ( $T'$ ) selected by the proposed TAM. The performance of the baseline ST-GCN is also provided.

### B. Do we need all skeletons for action recognition?

The hyperparameter  $T'$  of the proposed method defines the number of skeletons which are selected by TAM. To evaluate the model's performance with different number of selected skeletons, we applied experiments with varying values of  $T' = \{10, 30, 50, 100, 150, 200, 250, 300\}$ . Fig. 3 shows the obtained performance in terms of classification accuracy on Kinetics-Skeletons dataset and NTU-RGB+D dataset with both CV and CS benchmarks. As can be seen in Fig. 3, by increasing the number of skeletons, the models' accuracy is increased gradually. For NTU-RGB+D dataset, the best performance achieved by our method is equal to 94.2% and 87.97% for CV and CS benchmarks, respectively. This performance is achieved by selecting 150 most informative skeletons and increasing the number of skeletons doesn't improve the performance necessarily. This confirms our hypothesis that the model does not need to process all available skeletons to perform action classification. Since the Kinetics-Skeleton dataset is more challenging than NTU-RGB+D, both top1 and top5 accuracies are reported. The best top1 and top5 accuracies are 34.95% and 57.28%, respectively, which are achieved by selecting 250 skeletons. Overall, we can observe that by applying the proposed TAM for skeleton selection we can achieve competitive performance even for a small number of selected skeletons. For example, the performance in NTU-RGB+D (CV) for  $T' = 50$  is equal to 92.44%, compared to 93.83% corresponding to using all  $T = 300$  skeletons. This is an advantage for application scenarios with computational restrictions.

### C. Comparison with the state-of-the-art methods

We compare the performance of the proposed method with the of state-of-the-art methods in Tables. I and II on the NTU-RGB+D and Kinetics-Skeletons datasets, respectively. Table I is divided in three blocks grouping the methods in three categories, i.e. RNN-based, CNN-based and GCN-based methods, respectively. As can be seen, CNN-based methods perform better than RNN-based methods in general, while the state-of-the-art performance is achieved by GCN-based methods.

TABLE I  
COMPARISONS OF THE CLASSIFICATION ACCURACY WITH STATE-OF-THE-ART METHODS ON THE TEST SET OF NTU-RGB+D DATASET

Method	CS(%)	CV(%)	#Streams	Skel.sel.
HBRNN [1]	59.1	64.0	5	✗
Deep LSTM [10]	60.7	67.3	1	✗
ST-LSTM [2]	69.2	77.7	1	✗
STA-LSTM [11]	73.4	81.2	1	✓
VA-LSTM [12]	79.2	87.7	1	✗
ARRN-LSTM [13]	80.7	88.8	2	✗
Two-Stream 3DCNN [14]	66.8	72.6	2	✗
TCN [15]	74.3	83.1	1	✗
Clips+CNN+MTLN [16]	79.6	84.8	1	✗
Synthesized CNN [17]	80.0	87.2	1	✗
3scale ResNet152 [18]	85.0	92.3	1	✗
CNN+Motion+Trans [19]	83.2	89.3	2	✗
ST-GCN [30]	81.5	88.3	1	✗
DPRL+GCNN [32]	83.5	89.8	1	✓
AS-GCN [33]	86.8	94.2	2	✗
2s-AGCN [29]	88.5	95.1	2	✗
GCN-NAS [31]	89.4	95.7	2	✗
DGNN [34]	89.9	96.1	4	✗
<b>TA-GCN (<math>T' = 150</math>)</b>	<b>87.97</b>	<b>94.2</b>	<b>1</b>	<b>✓</b>
<b>2s-TA-GCN (<math>T' = 150</math>)</b>	<b>88.5</b>	<b>95.1</b>	<b>2</b>	<b>✓</b>
<b>4s-TA-GCN (<math>T' = 150</math>)</b>	<b>89.91</b>	<b>95.8</b>	<b>4</b>	<b>✓</b>

TABLE II  
COMPARISONS OF THE CLASSIFICATION ACCURACY WITH STATE-OF-THE-ART METHODS ON THE TEST SET OF KINETICS-SKELETON DATASET

Method	Top1(%)	Top5(%)	#Streams	Skel.sel.
Deep LSTM [10]	16.4	35.3	1	✗
TCN [15]	20.3	40.0	1	✗
ST-GCN [30]	30.7	52.8	1	✗
AS-GCN [33]	34.8	56.5	2	✗
2s-AGCN [29]	36.1	58.7	2	✗
DGNN [34]	36.9	59.6	4	✗
GCN-NAS [31]	37.1	60.1	2	✗
1s-TA-GCN ( $T' = 250$ )	34.95	57.28	1	✓
2s-TA-GCN ( $T' = 250$ )	36.1	58.72	2	✓
4s-TA-GCN ( $T' = 250$ )	36.9	59.77	4	✓

The results for the NTU-RGB+D dataset in Table I show that the proposed method, TA-GCN, outperforms all the CNN-based and RNN-based methods with a large margin. Besides, TA-GCN outperforms ST-GCN, which is the baseline in GCN-based methods, and DPRL+GCNN by a large margin in both CV and CS benchmarks. Compared to AS-GCN, the proposed method achieves higher accuracy in CS benchmark and a similar performance in CV benchmark. Here we should also note that the only competing GCN-based method that performs skeleton selection, i.e. DPRL+GCNN [32], performs poorly compared to all variants of the proposed method. The TA-GCNs' competitive methods are 2s-AGCN and GCN-NAS and the best performing method is DGNN. For Kinetics-Skeleton dataset (Table II), the proposed method outperforms the RNN-based methods, ST-GCN and AS-GCN methods and it has competitive performance with 2s-AGCN. The best performing methods on this dataset are GCN-NAS and DGNN. When the proposed model is trained using two different data streams joints and bones, (2s-TA-GCN), it achieves similar top1 accuracy with 2s-AGCN and its top5 accuracy exceeds

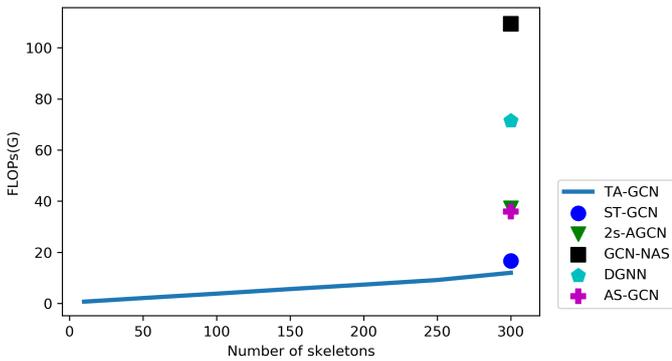


Fig. 4. Computational complexity comparison between the proposed method, when it selects different number of skeletons, and all the state-of-the-art methods which utilize all the skeletons of the input sequence.

TABLE III  
COMPARISONS OF THE COMPUTATIONAL COMPLEXITY WITH STATE-OF-THE-ART METHODS ON THE ON CV BENCHMARK OF NTU-RGB+D DATASET. THE NUMBER OF FLOPs AND PARAMS FOR 1s-TA-GCN ( $T' = 150$ ) ARE 5.64G AND 2.24M, RESPECTIVELY.

Method	FLOPs	# Params
1s-TA-GCN ( $T' = 150$ )	$\times 1$	$\times 1$
2s-TA-GCN ( $T' = 150$ )	$\times 2$	$\times 2$
4s-TA-GCN ( $T' = 150$ )	$\times 4$	$\times 4$
ST-GCN	$\times 2.9$	$\times 1.3$
AS-GCN	$\times 6.3$	$\times 3.2$
2s-AGCN	$\times 6.6$	$\times 3$
DGNN	$\times 12.6$	$\times 3.6$
GCN-NAS	$\times 19.3$	$\times 8.9$

2s-AGCN slightly. The 4s-TA-GCN which is trained using four data streams outperforms DGNN (top5) and achieves competitive performance with GCN-NAS while it has 4.8 times less computational complexity. Table. III shows the computational complexity comparison in terms of FLOPs and model parameters (Params) between the GCN-based competing methods on NTU-RGB+D (CV) dataset. Since the source code of the DPRL-GCNN method is not provided by the authors, it is not mentioned in the Table. III. Fig. 4, illustrates the computational complexity (FLOPs) of our method, when it selects different number of skeletons, and also the state-of-the-art methods which process all the skeletons. It can be seen that TA-GCN has less computational complexity compared to all the competing methods, even when it process all the skeletons in the sequence (i.e. for  $T' = 300$ ).

Columns 3 and 4 in Tables I and II indicate the number of network streams used by each GCN-based method and whether each method performs skeleton selection (or otherwise employs all  $T = 300$  skeletons). As can be seen, methods employing more than one streams generally outperform those with one stream. Although these methods achieve higher classification accuracy compared to single-stream methods, they have high computational complexity. Similar to our proposed method, the ST-GCN and DPRL-GCNN methods train only one network stream. They utilize only the joint data to train the model and we fuse both joint and bone information in

the first layer (input) which prevents increasing number of model parameters and FLOPs. The results show that not only does our method outperform the baseline ST-GCN in terms of classification accuracy, but also it has 2.9 times less FLOPs and 1.3 times less number of parameters. Therefore, TA-GCN can be a strong and efficient baseline for GCN-based human action recognition which achieves better performance than ST-GCN. The 2s-AGCN, GCN-NAS methods are built on top of ST-GCN method and train two networks using joint and bone data separately which doubles the number of model parameters and FLOPs. DGNN which has the best classification accuracy on both datasets trains 4 different networks with joint, bone, joint-motion and bone-motion data. DGNN and GCN-NAS have 12.6 and 19.3 times more FLOPs than the proposed method, respectively, and from 3.6 to 8.9 times more parameters.

To conduct a fair comparison, we also trained TA-GCN method with 4 different data streams separately. In Table. I, II, 2s-TA-GCN shows the ensembled softmax scores of the trained models with joint and bone data and 4s-TA-GCN shows the classification accuracy of the model when the softmax scores of all the 4 streams, joint, bone, joint-motion and bone-motion, are ensembled. The results indicate that training multiple models and ensembling the softmax scores can increase the accuracy by at least 1%. 2s-TA-GCN has a similar accuracy with 2s-AGCN while it has 3.3 times less FLOPs and 1.5 times less parameters. 4s-TA-GCN outperforms the competing methods 2s-AGCN and GCN-NAS with 1.65, 4.82 times less FLOPs, respectively. As can be seen in Table. III, GCN-NAS has the maximum number of FLOPs and it has competitive performance with 2s-TA-GCN which has with 9.6 times less computational complexity. In comparison with state-of-the-art, our proposed TA-GCN achieves good performance with much less computational complexity which makes TA-GCN applicable for many practical tasks with limited computational capacity.

## VI. CONCLUSION

In this paper, we proposed a temporal attention-augmented GCN to improve the computational efficiency in skeleton-based action recognition. Our method trains the attention mechanism to select the most informative skeletons for each action in an end-to-end manner. On two widely used benchmark datasets, the proposed method has competitive performance with the state-of-the-art, while being up to 10 times less computationally complex, and it outperforms the baseline in terms of both classification accuracy and computational complexity with a large margin. Therefore, it can be an efficient and strong baseline for skeleton-based action recognition.

## ACKNOWLEDGMENT

This work received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 871449 (OpenDR). This publication reflects the authors’ views only. The European Commission is not responsible for any use that may be made of the information it contains.

## REFERENCES

- [1] Y. Du, W. Wang, and L. Wang, "Hierarchical recurrent neural network for skeleton based action recognition," in *IEEE conference on computer vision and pattern recognition*, 2015, pp. 1110–1118.
- [2] J. Liu, A. Shahrourdy, D. Xu, and G. Wang, "Spatio-temporal lstm with trust gates for 3d human action recognition," in *European conference on computer vision*. Springer, 2016, pp. 816–833.
- [3] A. Iosifidis, A. Tefas, and I. Pitas, "View-invariant action recognition based on artificial neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 23, no. 3, pp. 412–424, 2012.
- [4] F. Han, B. Reily, W. Hoff, and H. Zhang, "Space-time representation of people based on 3d skeletal data: A review," *Computer Vision and Image Understanding*, vol. 158, pp. 85–105, 2017.
- [5] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake, "Real-time human pose recognition in parts from single depth images," in *IEEE Conference on Computer Vision and Pattern Recognition*. Ieee, 2011, pp. 1297–1304.
- [6] K. Sun, B. Xiao, D. Liu, and J. Wang, "Deep high-resolution representation learning for human pose estimation," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5693–5703.
- [7] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, "Realtime multi-person 2d pose estimation using part affinity fields," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7291–7299.
- [8] Z. Zhang, "Microsoft kinect sensor and its effect," *IEEE multimedia*, vol. 19, no. 2, pp. 4–10, 2012.
- [9] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, "Lstm: A search space odyssey," *IEEE transactions on neural networks and learning systems*, vol. 28, no. 10, pp. 2222–2232, 2016.
- [10] A. Shahrourdy, J. Liu, T.-T. Ng, and G. Wang, "Ntu rgb+ d: A large scale dataset for 3d human activity analysis," in *IEEE conference on computer vision and pattern recognition*, 2016, pp. 1010–1019.
- [11] 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 *AAAI conference on artificial intelligence*, 2017.
- [12] 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 *IEEE International Conference on Computer Vision*, 2017, pp. 2117–2126.
- [13] L. Li, W. Zheng, Z. Zhang, Y. Huang, and L. Wang, "Skeleton-based relational modeling for action recognition," *arXiv preprint arXiv:1805.02556*, vol. 1, no. 2, p. 3, 2018.
- [14] H. Liu, J. Tu, and M. Liu, "Two-stream 3d convolutional neural network for skeleton-based action recognition," *arXiv preprint arXiv:1705.08106*, 2017.
- [15] T. S. Kim and A. Reiter, "Interpretable 3d human action analysis with temporal convolutional networks," in *IEEE conference on computer vision and pattern recognition workshops*. IEEE, 2017, pp. 1623–1631.
- [16] Q. Ke, M. Bennamoun, S. An, F. Sohel, and F. Boussaid, "A new representation of skeleton sequences for 3d action recognition," in *IEEE conference on computer vision and pattern recognition*, 2017, pp. 3288–3297.
- [17] M. Liu, H. Liu, and C. Chen, "Enhanced skeleton visualization for view invariant human action recognition," *Pattern Recognition*, vol. 68, pp. 346–362, 2017.
- [18] B. Li, Y. Dai, X. Cheng, H. Chen, Y. Lin, and M. He, "Skeleton based action recognition using translation-scale invariant image mapping and multi-scale deep cnn," in *IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*. IEEE, 2017, pp. 601–604.
- [19] C. Li, Q. Zhong, D. Xie, and S. Pu, "Skeleton-based action recognition with convolutional neural networks," in *IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*. IEEE, 2017, pp. 597–600.
- [20] R. Pascanu, C. Gulcehre, K. Cho, and Y. Bengio, "How to construct deep recurrent neural networks," in *International Conference on Learning Representations*, 2014.
- [21] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *International Conference on Learning Representations*, 2017.
- [22] D. K. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams, "Convolutional networks on graphs for learning molecular fingerprints," in *Advances in neural information processing systems*, 2015, pp. 2224–2232.
- [23] M. Niepert, M. Ahmed, and K. Kutzkov, "Learning convolutional neural networks for graphs," in *International conference on machine learning*, 2016, pp. 2014–2023.
- [24] J. Atwood and D. Towsley, "Diffusion-convolutional neural networks," in *Advances in neural information processing systems*, 2016, pp. 1993–2001.
- [25] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Advances in neural information processing systems*, 2017, pp. 1024–1034.
- [26] F. Monti, D. Boscaini, J. Masci, E. Rodola, J. Svoboda, and M. M. Bronstein, "Geometric deep learning on graphs and manifolds using mixture model cnns," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5115–5124.
- [27] T. Kipf, E. Fetaya, K.-C. Wang, M. Welling, and R. Zemel, "Neural relational inference for interacting systems," *International conference on machine learning*, 2018.
- [28] N. Heidari and A. Iosifidis, "Progressive graph convolutional networks for semi-supervised node classification," *arXiv preprint arXiv:2003.12277*, 2020.
- [29] L. Shi, Y. Zhang, J. Cheng, and H. Lu, "Two-stream adaptive graph convolutional networks for skeleton-based action recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 026–12 035.
- [30] S. Yan, Y. Xiong, and D. Lin, "Spatial temporal graph convolutional networks for skeleton-based action recognition," in *AAAI conference on artificial intelligence*, 2018.
- [31] W. Peng, X. Hong, H. Chen, and G. Zhao, "Learning graph convolutional network for skeleton-based human action recognition by neural searching," in *AAAI conference on artificial intelligence*, 2020, pp. 2669–2676.
- [32] Y. Tang, Y. Tian, J. Lu, P. Li, and J. Zhou, "Deep progressive reinforcement learning for skeleton-based action recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5323–5332.
- [33] M. Li, S. Chen, X. Chen, Y. Zhang, Y. Wang, and Q. Tian, "Actional-structural graph convolutional networks for skeleton-based action recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 3595–3603.
- [34] L. Shi, Y. Zhang, J. Cheng, and H. Lu, "Skeleton-based action recognition with directed graph neural networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 7912–7921.
- [35] X. Gao, W. Hu, J. Tang, J. Liu, and Z. Guo, "Optimized skeleton-based action recognition via sparsified graph regression," in *ACM International Conference on Multimedia*, 2019, pp. 601–610.
- [36] B. Zhang, J. Han, Z. Huang, J. Yang, and X. Zeng, "A real-time and hardware-efficient processor for skeleton-based action recognition with lightweight convolutional neural network," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 66, no. 12, pp. 2052–2056, 2019.
- [37] M. Courbariaux, Y. Bengio, and J.-P. David, "Binaryconnect: Training deep neural networks with binary weights during propagations," in *Advances in neural information processing systems*, 2015, pp. 3123–3131.
- [38] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding," in *International Conference on Learning Representations*, 2016.
- [39] D. T. Tran, A. Iosifidis, and M. Gabbouj, "Improving efficiency in convolutional neural network with multilinear filters," *Neural Networks*, vol. 105, pp. 328–339, 2018.
- [40] Y. Chen, Z. Wang, Y. Peng, Z. Zhang, G. Yu, and J. Sun, "Cascaded pyramid network for multi-person pose estimation," in *IEEE conference on computer vision and pattern recognition*, 2018, pp. 7103–7112.
- [41] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [42] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, inception-resnet and the impact of residual connections on learning," in *AAAI conference on artificial intelligence*, 2017.
- [43] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev *et al.*, "The kinetics human action video dataset," *arXiv preprint arXiv:1705.06950*, 2017.
- [44] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," in *NIPS-W*, 2017.