

EPAM-Net: An efficient pose-driven attention-guided multimodal network for video action recognition

Ahmed Abdelkawy*, Asem Ali¹, Aly Farag

Computer Vision and Image Processing Laboratory (CVIP), University of Louisville, Louisville, KY, USA

ARTICLE INFO

Communicated by J. Liu

Keywords:

Human action recognition
Multimodal learning
X3D network
X-shiftNet
Spatial-temporal attention
Activities of daily living

ABSTRACT

Existing multimodal-based human action recognition approaches are computationally intensive, limiting their deployment in real-time applications. In this work, we present a novel and efficient pose-driven attention-guided multimodal network (EPAM-Net) for action recognition in videos. Specifically, we propose eXpand temporal Shift (X-ShiftNet) convolutional architectures for RGB and pose streams to capture spatio-temporal features from RGB videos and their skeleton sequences. The X-ShiftNet tackles the high computational cost of the 3D CNNs by integrating the Temporal Shift Module (TSM) into an efficient 2D CNN, enabling efficient spatiotemporal learning. Then skeleton features are utilized to guide the visual network stream, focusing on keyframes and their salient spatial regions using the proposed spatial-temporal attention block. Finally, the predictions of the two streams are fused for final classification. The experimental results show that our method, with a significant reduction in floating-point operations (FLOPs), outperforms and competes with the state-of-the-art methods on NTU RGB-D 60, NTU RGB-D 120, PKU-MMD, and Toyota SmartHome datasets. The proposed EPAM-Net provides up to a 72.8x reduction in FLOPs and up to a 48.6x reduction in the number of network parameters. The code will be available at <https://github.com/ahmed-nady/Multimodal-Action-Recognition>.

1. Introduction

Human action recognition (HAR), which assigns an action class label to an input video segment, has been an active research area in computer vision. It plays a key role in several real-world applications such as video indexing and retrieval, visual surveillance, healthcare, patient rehabilitation, sports analysis, and measuring the behavioral engagement of students in classrooms [1–5]. Recently, unimodal methods of HAR, such as skeleton-based or RGB video-based methods, have witnessed remarkable improvements.

The RGB video-based methods model the spatial-temporal representation from video data and its corresponding estimated optical flow using network architectures such as two-stream network [6], CNN-LSTM network [7], and 3D convolutional neural networks (3D CNN) [8–12]. Among these architectures, X3D is an efficient 3D CNN architecture, achieving competitive performance for action recognition. In this network architecture, a tiny 2D mobile CNN architecture, which utilizes channel-wise separable convolutions instead of standard ones, is progressively expanded along several axes: time, space, depth, and network width. Despite the efficiency of the X3D network, it still has higher computational demands compared to its 2D CNN counterpart.

On the other hand, the skeleton-based methods represent human action through the trajectories of body keypoints [13]. Skeleton data can be obtained either from RGB videos using pose estimation algorithms or from motion capture systems, e.g., Kinect. Skeleton-based approaches can be grouped into four categories based on the used network architecture: Convolutional Neural Network (2D-CNN) [14, 15], Recurrent Neural Network (RNN), Graph Convolutional Network (GCN) [16,17], and 3D CNN [18]. In 2D CNN-based methods [14,15], manually designed transformations are utilized to model the skeleton sequence as a pseudo image, while RNN-based methods model temporal context within the skeleton sequence. Such input representations limit the exploitation of structural information of the skeleton sequence. GCN-based approaches [16,19] represent pose sequences as spatiotemporal graphs. However, the limitations of these approaches are their non-robustness to noises in pose estimation and the necessity of careful design for integrating the skeleton with other modalities. In contrast, in the 3D-CNN based method [18], the input is represented as a volume of heatmaps, which capture the structure of skeleton joints and their dynamics over time. In this paper, we adopt the representation of the human skeleton sequence using a 3D pseudo-heatmap volume,

* Corresponding author.

E-mail addresses: a0nady01@louisville.edu (A. Abdelkawy), asem.ali@louisville.edu (A. Ali), aly.farag@louisville.edu (A. Farag).

<https://doi.org/10.1016/j.neucom.2025.129781>

Received 20 December 2024; Received in revised form 7 February 2025; Accepted 15 February 2025

Available online 25 February 2025

0925-2312/© 2025 Elsevier B.V. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

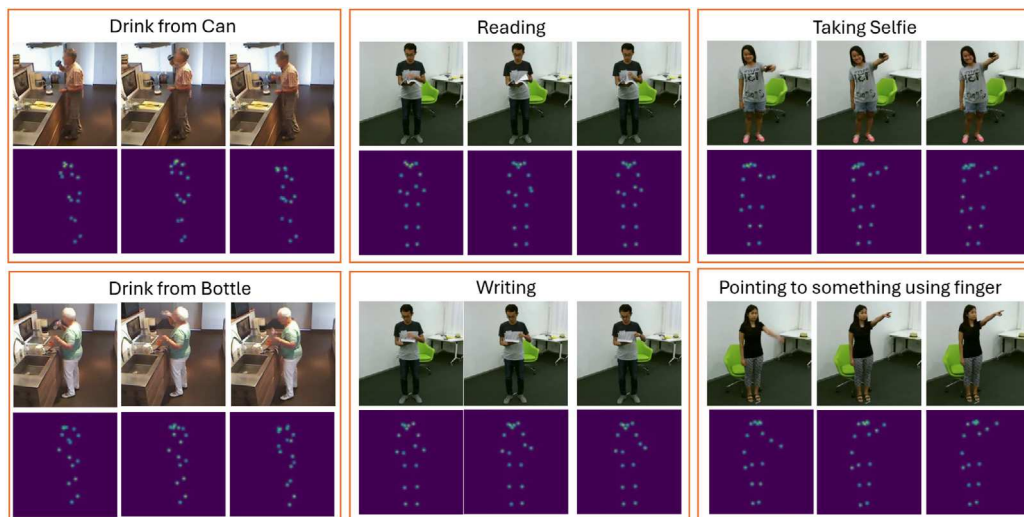


Fig. 1. An example of action pairs that are challenging to differentiate using the skeleton modality in Toyota-Smarthome (left) and in NTU-RGB+D dataset (middle and right), where the skeleton heatmaps of each pair of actions (e.g., reading–writing) are similar.

and then the proposed X-ShiftNet is utilized to learn spatial–temporal representation from such pseudo heatmap volume.

RGB video and skeleton modalities offer distinct perspectives on human actions. The RGB modality provides detailed appearance information, including scene context and object interactions. However, RGB video-based approaches are vulnerable to changes in viewpoint, background, and illumination conditions. In contrast, human skeleton data represents actions as a sequence of moving skeleton joints, making it robust against the challenges of RGB-based approaches. Nevertheless, certain actions appear ambiguous when viewed solely through skeleton data because of lacking appearance details, such as interacting objects. Fig. 1 shows an example of action pairs that have similar skeleton movements, such as drinking from a can and drinking from a bottle, reading and writing, as well as pointing to something with a finger and taking a selfie.

Therefore, multimodal HAR methods, which take advantage of the complementarity of RGB and skeleton modalities to improve the performance of HAR, have recently gained attention [20]. Previous fusion-based multimodal HAR methods can be grouped into three categories: score-level, feature-level, and model-level fusion. Score fusion-based models handle the skeleton and RGB information separately and then aggregate their scores from Softmax layers. However, RGB and skeleton modalities fail to boost each other for feature representation. On the other hand, feature fusion-based approaches concatenate modalities' features at the fully connected layers of modality-specific models. These methods slightly improved the HAR performance due to not considering the alignment of RGB data and its corresponding human body poses. Das et al. [21] addressed such an alignment problem through proposing spatial embedding, which projects visual features and 3D skeletons in the same referential. Moreover, the temporal alignment is performed by assuming the existence of a 3D pose for each frame. However, this approach is computationally expensive. In contrast, model-level fusion-based methods employ the knowledge from one data modality to facilitate modeling in other data modalities [22]. Bruce et al. [22] used skeleton modality to learn spatial attention and then weight the spatiotemporal region of interest (ST-ROI) map, which is constructed from video input, accordingly. Although this approach uses a 2D CNN to learn visual features from the ST-ROI map, it is still computationally intensive.

To tackle the mentioned limitations, we propose a novel Efficient Pose-driven Attention-guided Multimodal Network (EPAM-Net) that exploits the complementarity of RGB and skeleton modalities by aligning skeleton sequences and RGB videos in spatial and temporal dimensions. EPAM-Net consists of two eXpand temporal Shift networks

(X-ShiftNets), which integrate the Temporal Shift Module (TSM) into an efficient 2D CNN, enabling efficient spatiotemporal learning, and a lightweight spatiotemporal attention block. In summary, our main contributions are as follows:

1. We introduce two X-ShiftNet models, which achieve and surpass efficient 3D CNNs (e.g., X3D network) performance while requiring 1.1x and 1.4x fewer FLOPs and network parameters.
2. We introduce a pose-driven spatiotemporal attention block to guide the visual stream focusing on discriminative frames and their salient human body regions.
3. Using the aforementioned two components, we construct an efficient multimodal architecture (EPAM-Net) that exploits the complementarity of the RGB and skeleton modalities in an end-to-end manner.
4. Our EPAM-Net outperforms and competes with state-of-the-art methods on four benchmarks while reducing FLOPs and network parameters up to 72.8x and 48.6x, respectively.

2. Related work

In this section, we review the HAR literature according to the model's type: unimodal HAR (RGB video-based or skeleton-based HAR) and multimodal HAR.

2.1. RGB video-based action recognition approaches

Due to the easy collection of RGB data, RGB video-based methods have rapidly developed and obtained impressive results. RGB video-based methods can be divided into three categories [13]: Two-stream 2D CNN-based, RNN-based, and 3D CNN-based methods. Two-stream 2D CNN methods [23] comprise two 2D CNN streams to learn the appearance and motion features from RGB video and its corresponding estimated optical flow. RNN-based methods [7] extract frame-level visual features using 2D CNN, then utilize gated RNN architectures, e.g., Long-Short Term Memory (LSTM) to capture the long-term temporal dynamic in a video sequence. To concurrently learn the spatial and temporal information from RGB video, 3D CNN-based methods are introduced. Two-stream Inflated 3D CNN [10] was introduced through temporally extending the convolutional and pooling kernels of a 2D CNN. Feichtenhofer et al. [11] proposed a SlowFast network with two paths that operate on video frames but at two different speeds. The slow pathway, which works at a low frame rate, models spatial semantics, while the fast pathway models fine motion by working at a high

frame rate. The lateral connection is used to fuse the two pathways. Despite the impressive results of 3D CNN-based methods, they require heavy computation to extract spatio-temporal features from videos. As a result, Temporal Shift Module (TSM) [24] is proposed to enable 2D CNN networks to achieve 3D CNN performance without adding computational overhead. TSM facilitates temporal interactions among features of neighboring frames through shifting a part of the channels along the temporal dimension. Feichtenhofer [12] introduced an efficient 3D CNN architecture (X3D) for action recognition by progressively expanding a tiny 2D mobile image classification architecture into a spatiotemporal one along several possible axes: temporal duration γ_t , spatial resolution γ_s , network depth γ_d , network width γ_w , bottleneck width γ_b , and frame rate γ_r . This progressive expansion starts by expanding the bottleneck width of the model γ_b , followed by frame rate γ_r , spatial resolution γ_s , depth of the network γ_d , temporal/duration γ_t , and finally, global width of the layers γ_w . Due to the efficiency and competitive performance of the X3D network, we built upon it to develop our X-ShiftNet model.

2.2. Skeleton-based action recognition approaches

Skeleton-based approaches can be grouped into four classes based on architectures: 2D CNN, RNN, GNN, and 3D CNN. For CNN, Choutas et al. [14] introduced video clip-level human pose-based representation that encodes body joints' motion during the entire clip. The pose motion representation (PoTion) is constructed by temporally aggregating the heatmaps of each joint by colorizing each of them based on their order in the video clip. After that, PoTion is used as input for the proposed 2D-CNN network architecture to predict the action category. Liu et al. [25] observed that using pose estimation maps that maintain more details of human body shape is more beneficial for action recognition than depending solely on the inaccurate 2D coordinates of body joints. So, they generated two evolution maps: a body shape evolution map from a sequence of averaged body joints heatmaps and a body joints evolution map from a sequence of pseudo-heatmaps of estimated body joints 2D coordinates. The two evolution maps are used as input for the proposed two-stream 2D CNN architecture, and then the prediction label is obtained by averaging the two CNN scores.

For GCN-based approaches, Yan et al. [16] proposed the Spatial-Temporal Graph Convolutional Network (ST-GCN) to model the dynamics of skeleton sequence and the spatial arrangements of its joints. They created a spatial graph by utilizing the inherent connections between joints in the human body, and they also introduced temporal edges between corresponding joints in consecutive frames to extend the spatial graph into the spatiotemporal domain. The limitation of this work is that the skeleton GCN's graph is heuristically preset and it reflects only the human body's physical structure. It also remains fixed across all layers and input samples. Lie et al. [17] addressed this limitation by introducing an adaptive graph convolutional network that is capable of adaptively learning the graph topology for different GCN layers and skeleton samples. To improve the accuracy of classification, the joints and second-order information of the skeleton represented in the bone lengths and their directions were modeled using a two-stream adaptive graph convolutional network (2s-AGCN).

For 3D CNN approaches, Duan et al. [18] represented the sequence of the human skeleton using a 3D pseudo-heatmap volume. Then, a 3D CNN network called slowOnly is used to classify the 3D heatmap volumes into one of the action categories. Compared to the SlowOnly model [18], we develop X-ShiftNet model to learn spatiotemporal information from the pseudo-heatmap volume of the skeleton sequence.

2.3. Multimodal action recognition approaches

Multimodal HAR approaches can be grouped into two categories: fusion and distillation-based approaches. Fusion-based approaches can be further classified into three groups: score-level fusion, feature-level

fusion, and model-level fusion. In score fusion, the skeleton and RGB information are modeled separately, and then the classification scores of both network streams are fused to obtain the final prediction. On the other hand, feature fusion-based methods concatenate modality-specific features either at the fully connected layers of modality-specific models or at several layers using lateral connections [18]. Zolfaghari et al. [26] introduced a deep network architecture to utilize the three visual cues (pose, motion, and appearance) and fuse them sequentially using a Markov chain model. Each modality has its 3D-CNN architecture (C3D). In the chained architecture, the prediction of each stream relies not only on the stream input but also on all predictions of previous streams. As a result, the network in that stream refines class labels from prior streams by learning complementing features. Li et al. [27] introduced a skeleton-guided multimodal network (SGM-Net) to exploit the complementarity of RGB and skeleton modalities at the feature level. The network is composed of three components: ST-GCN [16] to extract pose features, the R(2+1)D network [28] to extract visual features, and a guided block to pay attention to action-related features in RGB videos. Unlike SGM-Net, where visual features are extracted from the whole spatial resolution of video frames, we extract these features from cropped human regions to reduce the interference of background in action classification. Das et al. [21] introduced video-pose network (VPN), whose input is 64 video frames and their corresponding 3D poses. They used the I3D network to extract spatiotemporal features from these 64 frames, while the GCN model was employed to learn pose features from 3D poses. The pose features are used to learn spatial-temporal attention, and then spatial embedding is proposed to find the correspondence between human joints and their relevant image regions to modulate visual features, which are used for classification. Unlike VPN [21], our proposed approach represents a skeleton sequence using a pseudo heatmap volume, which is spatially aligned with a cropped RGB human region.

For model-based fusion, Bruce et al. [22] proposed a model-based multimodal network (MMNet) that learns spatial attention from the skeleton modality using a GCN network and then weights the spatiotemporal region of interest (ST-ROI) map accordingly. An ST-ROI map is constructed by cropping the frame's body area of the actor(s) head, hands, and feet from five sampled frames of a video input. After that, such ST-ROI is weighted and fed to 2D CNN to learn the appearance feature. Unlike MMNet [22], which uses spatial attention on ST-ROI, not spatial-temporal attention that may not assign different weights for each body part through frames, we use person-centered modeling and spatial-temporal attention to modulate visual features accordingly.

Conversely, distillation-based approaches incorporate pose information into a network with RGB input, eliminating the need for poses at inference. Das et al. [29] proposed VPN++ network to augment the RGB representation with 3D pose information through features and attention-level distillation. However, it neglected the alignment between 2D skeleton joints and their corresponding spatial regions in RGB frames. To address the alignment problem, Reilly and Das [30] proposed 2D and 3D pose induction modules to integrate 2D and 3D pose information into TimeSFormer with RGB input. However, their method is computation-intensive with 590.0 GFLOPs.

In this work, we propose a multimodal action recognition approach that addresses these limitations and reduces the computation complexity as well as enhances the performance by developing X-ShiftNet to learn spatial-temporal features and modulating visual features using spatial-temporal attention. Table 1 shows a comparison of our EPAM-Net with existing person-centric multimodal HAR approaches in terms of modality interactions, contributions, and limitations.

3. The proposed EPAM-Net

Our proposed approach focuses on person-centric modeling by capturing both human body movements and interacting objects. To achieve

Table 1
Comparison between person-centric multimodal HAR approaches.

Work	Modality interactions	Contributions	Limitations
VPN [21]	Feature-level fusion	◊ Learn spatial and temporal attention weights separately, then fuse them via a Hadamard product. ◊ Align 3D joints with image regions via spatial embedding.	◊ Intensive computation (107.9 GFLOPs). ◊ Performance depends on 3D pose quality.
VPN++[29]	Knowledge distillation	◊ Augment RGB representation with 3D pose information through features and attention-level distillation.	◊ Intensive computation (125.8 GFLOPs). ◊ Ignoring skeleton-image alignment. ◊ Comparable results demand 3D poses at inference.
MMNet [22]	Model-level and score-level fusion	◊ Construct ST-ROI from five-RGB frame actors' head/hands/feet crops. ◊ Utilize GCN-guided spatial attention to weigh the ST-ROI map before 2D CNN classification.	◊ Use spatial, not spatiotemporal, attention on ST-ROI, thus lacking per-part, per-frame weighting. ◊ Performance depends on 3D pose quality. ◊ High computational cost (89.2 GFLOPs). ◊ Need 2D pose to crop head/hands/feet from RGB frames and 3D pose for GCN-based pose stream input.
TCEM-MMNet [31]	Model-level and score fusion	◊ Introduce TCEM, composed of ResNet18 and a 3-layer LSTM for spatial and temporal feature learning of ST-ROI.	◊ Intensive computation (85.4 GFLOPs). ◊ Performance depends on 3D pose quality. ◊ Need 2D pose to crop head/hands/feet from RGB frames and 3D pose for GCN-based pose stream input.
π -ViT [30]	Knowledge distillation	◊ Integrate 2D/3D pose information into TimeSFormer's RGB backbone via 2D/3D pose induction module.	◊ High computation cost (590.0 GFLOPs). ◊ Comparable results demand 3D poses at inference.
Ours	Feature-level and score-level fusion	◊ Introduce X-ShiftNets to learn spatiotemporal features from aligned RGB frames and pose sequences. ◊ Introduce a lightweight nesting spatiotemporal attention block. ◊ EPAM-Net rivals SOTA methods with 8.1 GFLOPs.	◊ Performance depends on 2D pose quality.

this, we first determine the minimum bounding box that encompasses all 2D human skeletons across video frames. Then, each frame is cropped according to this bounding box and resized to the target dimensions. Fig. 2 illustrates an overview of the proposed network architecture. The input to our proposed network is cropped RGB frames and their corresponding skeleton sequence. Specifically, the pose stream input is the pseudo-heatmap volume constructed from N uniformly sampled frames from the input video clip, while the RGB stream input consists of M frames selected from these N frames by picking one frame out of every $\frac{N}{M}$ frames.

The pipeline of the proposed multimodal network includes: (1) Extracting the spatiotemporal dynamics of skeleton sequences and cropped RGB frames separately using the proposed X-ShiftNet models; (2) Guiding the visual network stream to focus on discriminative human body part(s), interacting objects, and keyframes a using spatial-temporal attention block; (3) Fusing the class scores of the two streams of the proposed network for further performance improvement. The rationale behind this pipeline is that the M -frame RGB video input (e.g., $M = 16$) might not be sufficient to capture the full temporal dynamics of an action. In contrast, the pose network, which utilizes N uniformly sampled frames (e.g., $N = 48 > M$), can capture more temporal information. Below we discuss the visual network, pose network, and spatial-temporal attention block in detail.

3.1. Visual stream

The X-ShiftNet network is proposed to extract spatiotemporal features $F_r \in \mathbb{R}^{C \times T \times H \times W}$ from RGB video frames, where C represents the number of channels, T the number of frames, and $H \times W$ the spatial resolution. The X-ShiftNet network, inspired by TSM [24] and X3D networks [12], achieves the effect of 3D convolution using 2D convolution by moving a portion of the input feature channels along the temporal dimension. The proposed X-ShiftNet combines the architectural design strength of the X3D network and the temporal modeling capability of TSM, capturing spatio-temporal features without increasing network parameters or computation overhead. Specifically, we obtain the 2D CNN counterpart of the X3D network by removing the temporal convolution in the conv1 stage and replacing all 3D convolutions with 2D convolutions. Then, TSM is added for each inverted residual block of

the network. Moreover, we employ one fully connected layer instead of two for the classification head. Table 2 shows the instantiating of the RGB-X-ShiftNet network.

3.2. Pose stream

We develop an X-ShiftNet-s network to extract spatiotemporal dynamics from skeleton sequences. The proposed RGB-X-ShiftNet network is modified for skeleton-based action recognition as follows: (1) Remove the first stage, and (2) Change the spatial stride of the first convolution layer of the stem layer from 2 to 1. This makes the spatial resolution of the final feature map match that of the visual backbone feature maps. Table 2. shows the instantiation of the Pose-X-ShiftNet network.

Following the work [18], we employ the Top-Down pose estimation approach [18] with Faster-RCNN as a detector and HRNet [32] as a pose estimation. Having skeleton coordinate triplets (x_k, y_k, c_k) for each frame, the K heatmaps are generated using a Gaussian map centered at each joint:

$$H_k(x, y) = \exp - \frac{(x - x_k)^2 + (y - y_k)^2}{2\sigma^2} \quad (1)$$

where x_k and y_k are the coordinates of k th joint and σ controls Gaussian map variance. Finally, all K joints heatmaps are stacked along the temporal dimension to form the 3D heatmap volume with size $K \times T \times H \times W$ where K is a number of human body keypoints, T is temporal length, and H and W are the height and width of such maps.

3.3. Spatial-temporal attention block

After extracting skeleton features using the Pose-X-ShiftNet network and visual features using the RGB-X-ShiftNet network, the nesting spatial-temporal attention block, as shown in Fig. 3, is proposed to learn which spatial regions in each frame and which frames are worth paying attention to using skeleton features and then weigh visual features accordingly. The nesting spatial-temporal attention block consists of a spatial attention module, followed by a nested temporal attention module, which uses the spatial attention map as an input. Our nesting spatiotemporal attention, like VPN's spatiotemporal coupler [21], modulates visual features by assigning different weights to each frame

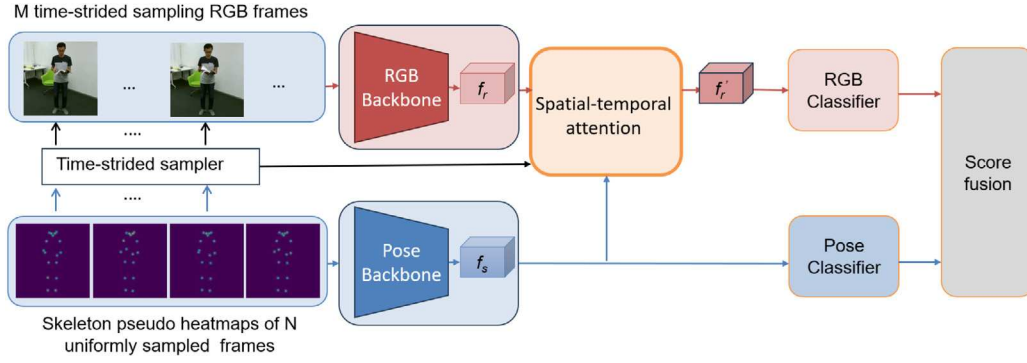


Fig. 2. The EPAM-Net architecture consists of visual and pose backbones to extract spatial–temporal features from RGB videos and pose sequences, respectively; a pose-driven spatial–temporal attention block to re-weight visual features accordingly; two classification heads; and final score fusion. The input of the pose network stream is a pseudo-heatmaps volume from N uniformly sampled frames, while the input of the visual network stream consists of M frames selected from these N frames by sampling one out of every $\frac{N}{M}$ frame. f_s and f_r represent skeleton features and visual features, respectively.

Table 2

X-ShiftNet architectures for RGB and Pose streams. The kernel dimensions are represented by $T \times S^2$, C for temporal, spatial, and channel sizes, respectively. TSM is the temporal shift module.

Stage	RGBNet	PoseNet	Output sizes $T \times H \times W$
data layer			RGB: 16x224x224 Pose: 48x56x56
conv1	$1 \times 3^2, 24$	$1 \times 3^2, 24$	RGB: 16x112x112 Pose: 48x56x56
res2	$\begin{bmatrix} TSM \\ 1 \times 1^2, 54 \\ 1 \times 3^2, 54 \\ 1 \times 1^2, 24 \end{bmatrix} \times 3$	None	RGB: 16x56x56 Pose: 48x56x56
res3	$\begin{bmatrix} TSM \\ 1 \times 1^2, 108 \\ 1 \times 3^2, 108 \\ 1 \times 1^2, 48 \end{bmatrix} \times 5$	$\begin{bmatrix} TSM \\ 1 \times 1^2, 54 \\ 1 \times 3^2, 54 \\ 1 \times 1^2, 24 \end{bmatrix} \times 5$	RGB: 16x28x28 Pose: 48x28x28
res4	$\begin{bmatrix} TSM \\ 1 \times 1^2, 216 \\ 1 \times 3^2, 216 \\ 1 \times 1^2, 96 \end{bmatrix} \times 11$	$\begin{bmatrix} TSM \\ 1 \times 1^2, 108 \\ 1 \times 3^2, 108 \\ 1 \times 1^2, 48 \end{bmatrix} \times 11$	RGB: 16x14x14 Pose: 48x14x14
res5	$\begin{bmatrix} TSM \\ 1 \times 1^2, 432 \\ 1 \times 3^2, 432 \\ 1 \times 1^2, 192 \end{bmatrix} \times 7$	$\begin{bmatrix} TSM \\ 1 \times 1^2, 216 \\ 1 \times 3^2, 216 \\ 1 \times 1^2, 96 \end{bmatrix} \times 7$	RGB: 16x7x7 Pose: 48x7x7
conv5	$1 \times 1^2, 432$	$1 \times 1^2, 216$	RGB: 16x7x7 Pose: 48x7x7
global average pooling, fc			#classes

and its spatial regions. However, rather than computing spatial and temporal attentions separately as in VPN [21], we explicitly model their interaction, gaining an advantage of attention to attention. This contrasts with MMNet [22] as well, which uses skeleton joint weights (i.e., spatial attention) to weight the ST-ROI map, the input of the 2D CNN of the RGB stream.

To properly use spatial–temporal attention, we align skeleton pseudo-heatmaps with corresponding RGB frames. For spatial alignment, video frames are cropped according to the minimum bounding box involving all 2D human skeletons across the video frames, and then the skeleton pseudo-heatmap volume is generated accordingly. Moreover, the spatial resolution of the final feature maps for visual and pose backbones is matched to ensure spatial correspondence between the two modalities. For temporal alignment, since the RGB and pose modalities have different temporal resolutions, their feature maps should be matched in time for accurate action recognition. In particular, denoting the shape of the pose feature as $\{C_p, T_N, S^2\}$, and the shape of the RGB feature as $\{C_r, T_M, S^2\}$, so, T_N is aligned with T_M through time-strided sampling. Below, we discuss the spatial and temporal attention module in detail.

3.3.1. Spatial attention module

Given the skeleton feature maps $F_s \in \mathbb{R}^{C \times T \times H \times W}$, the spatial attention map $A_S \in \mathbb{R}^{1 \times T \times H \times W}$ is obtained by compressing channel-wise features using max-pooling and average-pooling operations, followed by a $1 \times 7 \times 7$ convolution. Specifically, the process of spatial attention can be expressed as follows:

$$A_S = \phi(g^{1 \times 7 \times 7}([GAP(F_s); GMP(F_s)])) \quad (2)$$

where GAP denotes global average pooling, GMP represents the global max pooling.

The spatial attention map reveals the importance of each spatial region in each video frame, with those of larger weights representing discriminative regions for the action.

3.3.2. Temporal attention module

The temporal attention module is inspired from [33]. It has two operations: a squeeze operation, in which global average pooling is used to aggregate spatial dimensions of a spatial attention map A_S , and an excitation operation, in which a 1D convolution models temporal-wise interactions among neighboring frames. Overall, the two operations of

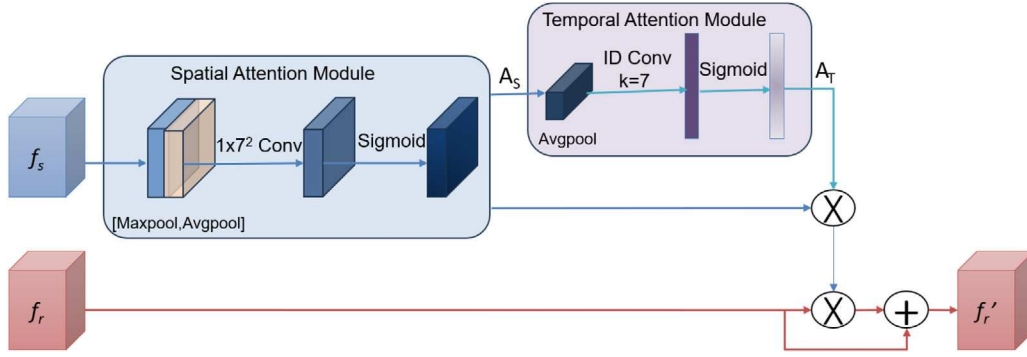


Fig. 3. Illustration of the proposed spatial-temporal attention block. A spatial attention map weights discriminative spatial regions, while a temporal attention map weights keyframes.

the temporal attention block can be formulated as:

$$Z_t = \frac{1}{W \times H} \sum_{i=1}^W \sum_{j=1}^H A_S(:, i, j, :), \quad (3)$$

$$A_T = \phi(\text{Conv1D}(Z_t)) \quad (4)$$

where Conv1D is 1D Conv with a kernel size of 7.

The temporal attention map A_T represents the importance of the T frames, with frames having larger weights in A_T expected to be keyframes. The spatiotemporal attention map A_{ST} is obtained by multiplying spatial attention map A_S and temporal attention map A_T , i.e., $A_{ST} = A_S \otimes A_T$. After that, the RGB feature is modulated according to A_{ST} as follows: $f_r' = f_r + f_r \otimes A_{ST}$. The reason for adopting a residual connection in modulating the RGB feature is the low quality of 2D pose estimation due to occlusion, low resolution, and truncation for some datasets, e.g., Toyota-Smarthome dataset.

4. Experiments

We evaluate the proposed multimodal network on NTU RGB+D 60 [34], NTU RGB+D 120 [35], PKU-MMD [36], and Toyota-Smarthome [37] datasets. We report the mean Top-1 accuracy for Toyota-Smarthome dataset [37] following [22,30] and Top-1 accuracy for other datasets using 1-clip per video.

4.1. Datasets

NTU RGB-D dataset [34,35]. The dataset is a large-scale, multi-modalities human action recognition dataset captured in a lab-controlled environment. It is available in two variants, NTU-60 and NTU-120. The NTU RGB-D 60 dataset has 56,880 video clips of 60 human actions performed by 40 volunteers, whereas the NTU RGB-D 120 dataset has 114,480 videos of 120 human actions performed by 106 volunteers. Each action is simultaneously captured from three distinct horizontal views for several camera setups. Each camera setup has a different height, and the three cameras are positioned at that height. The datasets have three settings for evaluation: cross-subject (X-Sub), cross-view (X-View for NTU 60), and cross-setup (X-Set for NTU-120). In cross-subject (X-Sub), half of the subjects are used for training and the other half for testing. For X-View, video samples are split based on camera IDs (cameras 2 and 3 for training, camera 1 for testing), while X-Set splits are based on camera setups (even setup IDs for training and odd ones for testing).

Toyota-Smarthome dataset [37]. This dataset contains activities of daily living (ADL) collected in a smart home where 18 elderly people perform daily living tasks spontaneously. This dataset comprises 16,115 video clips with 31 activity classes captured from 7 viewpoints. The dataset evaluation protocols are cross-subject (CS) and cross-view

(CV2) [29]. The cross-view (CV1) evaluation setting is ignored due to the small number of training samples.

PKU-MMD dataset [36]. The dataset comprises 21,545 action samples in 51 action categories performed by 66 participants and recorded simultaneously from left, center, and right viewpoints using three Microsoft Kinect v2 cameras. It employs two evaluation protocols: Cross-Subject (x-sub), where 57 subjects are used for training and 9 for testing, and Cross-View (x-view), where data from the center and right cameras are used for training, while the left camera is reserved for testing.

4.2. Implementation details

For skeleton modality, following the work [18], we utilize a Top-Down pose estimation approach instantiated with HRNet-W32 [32] to extract 2D poses from videos and save the coordinate triplets (x, y, score). We then generate the pseudo heatmaps volume, the input of the pose stream, as follows: (1) 48 frames are sampled uniformly from the video. (2) 17 heatmaps (one heatmap per joint) are generated for each sampled frame (see Section 3.2) and then all such heatmaps are stacked along the temporal dimension. (3) The heatmaps are cropped with the global box that envelops all persons in the video and resized to 56 x 56. For RGB video modality, we select 16 frames from these 48 sampled frames of the pose stream via time-strided sampling. We then crop such frames with the global box and resize them to a resolution of 224 x 224. The cropped RGB frames are used as input for the RGB stream. Horizontal flipping is applied as data augmentation during training for RGB and pose streams.

The training process of the proposed multimodal network involves two phases: first, pre-training RGB and skeleton networks, followed by fine-tuning the entire multimodal network for final classification. The RGB stream (RGB-X-ShiftNet) is trained with a batch size of 64 and a learning rate of 0.05 for 200 epochs, while the skeleton stream (Pose-X-ShiftNet) is trained with a batch size of 96 and a learning rate of 0.0375 for 240 epochs. The entire multimodal network is trained with a batch size of 24 and a learning rate of 0.001 for 10 epochs. The optimizer used to train all networks is the Stochastic Gradient Descent (SGD) with a momentum of 0.9. The learning rate is decreased using the cosine annealing schedule. The loss function of the proposed multimodal network is the summation of two cross-entropy losses of RGB and skeleton streams. For X-ShiftNets, we empirically shift 1/8 of feature channels forward and another 1/8 backward in the RGB stream, while in the Pose stream, we shift 1/4 forward and another 1/4 backward, as detailed in Section 4.4. On Toyota-Smarthome dataset, we pretrain the two-stream X-ShiftNets on the NTU RGB-D 120 dataset since training sets on cross-subject and cross-view2 protocols are 8761 and 7685 video clips, respectively. We adopt PyTorch for implementation and train the RGB and Pose streams on 8 NVIDIA Tesla V100 PCIe 16 GB and finetune the proposed multimodal network on 2 NVIDIA TITAN RTX GPUs 24 GB.

Table 3

Ablations of the proposed multimodal architecture components on three benchmarks: NTU RGB+D 60, NTU RGB+D 120, and Toyota SH. We report the mean Top-1 accuracy (%) for Toyota-Smarthome dataset and Top-1 accuracy (%) for other datasets using 1-clip per video. STA stands for spatial-temporal attention.

#	Method	NTU 60		NTU 120		Toyota SH	
		X-Sub	X-View	X-Sub	X-Set	X-Sub	X-View2
1	Pose-X-ShiftNet	92.72	96.39	84.0	87.26	65.74	59.30
2	RGB-X-ShiftNet	94.83	98.04	90.13	91.58	62.73	58.81
3	Score fusion (#1,#2)	96.0	98.77	91.40	93.41	70.29	64.78
4	Ours with STA (#1,#2)	96.1	99.0	92.4	94.3	71.7	67.8

Table 4

Ablation study of proportion shift's impact on X-ShiftNet performance, evaluated on NTU RGB+D 60's X-Sub and X-View protocols.

Portion shift	RGB Stream		Pose stream	
	X-Sub	X-View	X-Sub	X-View
1/4	94.83	98.04	92.6	95.7
1/2	95.03	97.98	92.7	96.4

4.3. Ablation studies

In this section, we assess the effectiveness of each component of the proposed pose-driven attention multimodal architecture. Moreover, we compare the performance of network architectures for RGB and pose streams according to the complexity/accuracy trade-off. Finally, we evaluate the effectiveness of the nesting spatial-temporal attention block.

4.3.1. Effectiveness of the proposed multimodal architecture components

From Table 3, we can notice that RGB video-based network has higher top-1 accuracies compared to its skeleton-based counterpart on the NTU RGB-D 60 and NTU RGB-D 120. This is because the NTU RGB-D 60, and NTU RGB-D 120 datasets are captured in a lab-controlled environment where there are neither illumination changes nor background variations. Conversely, on Toyota-Smarthome, the RGB video-based network performs similarly or worse than the skeleton-based network due to background clutter and viewpoint variations (see Fig. 1). Also, we can observe that the proposed spatial-temporal attention enhances the accuracy by 0.1%, 0.2%, 1.0%, and 0.9% for NTU 60 (cross-subject and cross-view evaluation setting) and NTU 120 (XSub and XSet), respectively. In addition, on Toyota-Smarthome dataset, our attention-based multimodal network surpasses its counterpart without such spatial-temporal attention with 1.4% and 3.0% in mean class accuracy on CS and CV2 protocols.

This confirms that the proposed multimodal method takes full advantage of the complementarity of RGB and skeleton modalities.

4.4. Choosing the proportion shift of the X-ShiftNet

Since X-ShiftNet utilizes the Temporal Shift Module (TSM) for temporal modeling, we need to study the impact of the different proportion shifts on the X-ShiftNet performance. From Table 4, we can observe the following: (1) The Pose-X-ShiftNet reaches its peak performance when shifting 1/2 feature channels between neighboring frames (1/4 in each direction). This confirms the intuition that the pose stream focuses on modeling the action dynamics since the included information in the skeleton sequence is the skeleton joint coordinates. (2) X-ShiftNet for RGB stream achieves similar performance when 1/2 or 1/4 of feature channels are shifted, although the proportion shift 1/2 has a higher data movement overhead than that of 1/4. Therefore, we use proportion shifted channels 1/4 for the RGB stream and 1/2 for the pose stream for the rest of the paper.

Table 5

Comparison with RGB video-based methods on NTU RGB+D 60 X-Sub Protocol.

Backbone	T_{RGB}	NTU60-XSub	FLOPs	Params
TSM [24]	8	92.3	33.0G	23.6M
TSM [24]	16	93.7	65.9G	23.6M
SlowOnly [18]	8	94.9	42.0G	31.9M
X3D [12]	16	94.0	5.0G	3.1M
RGB-X-ShiftNet (our)	16	94.8	4.5G	2.0M

Table 6

Comparison with skeleton-based methods on NTU RGB+D 60 X-Sub Protocol.

Backbone	T_{pose}	NTU60-XSub	FLOPs	Params
C3D [18]	48	92.5	16.8 G	3.4 M
SlowOnly [18]	48	93.1	15.9 G	2.0 M
X3D [12]	48	92.84	4.0 G	586.2 k
Pose-X-ShiftNet (our)	48	92.72	3.6 G	514.7 k

4.4.1. Choosing the appearance and pose network architectures

A comparison of RGB video-based methods is shown in Table 5. We can observe that the proposed RGB-X-ShiftNet outperforms TSM [24] and the X3D network [12] and obtains similar results compared to the SlowOnly network [18], while requiring 14.5x, 1.1x, and 9.3x fewer GFLOPs. For the pose network, it is noticeable from Table 6 that the proposed Pose-X-ShiftNet network performs better than C3D [18] while requiring 4.7x fewer FLOPs. Also, it achieves competitive performance with the X3D-s network [12], and SlowOnly [18] while requiring 1.1x and 4.4x less GFLOPs.

4.4.2. Choosing the spatial-temporal attention block

To evaluate the effectiveness of the proposed nesting spatial-temporal attention block, we conducted experiments with the nesting spatial-temporal attention block from [38], which consists of both spatial attention and a nested temporal attention module. The spatial attention module involves a 1x3x3 spatial convolution to compress the number of channels f_s to 1, followed by a 1x7x7 spatial convolution. This can be expressed as:

$$A_S = \phi(g^{1 \times 7 \times 7}(\delta(g^{1 \times 3 \times 3}(f_s))))), \quad (5)$$

where ϕ and δ are the Sigmoid and RELU activation functions, respectively.

The temporal attention block, which is inspired by the Squeeze and Excitation(SE) block [39], aggregates spatial information of the spatial attention map A_S using global average pooling and then models the temporal-wise dependencies by two fully connected layers with non-linear activation functions (RELU and Sigmoid). This process is formulated as:

$$A_T = \phi(W_2(\delta(W_1(GAP(A_S))))), \quad (6)$$

The weights of the fully connected layers are represented by W_1 and W_2 .

Our multimodal network with the proposed nesting spatial-temporal attention block achieves 96.14% and 99.0% on the X-Sub and X-View evaluation settings of the NTU RGB-D 60 dataset, respectively, compared to 96.20% and 98.9% with the spatial-temporal attention

Table 7

Comparison of top-1 accuracy (%) with state-of-the-art methods on NTU RGB+D 60 and NTU RGB+D 120 Toyota SmartHome.[†] indicates the results are obtained using 10 clips per video. ◦ means the skeleton is used in training but not in inference.

Method	Modality		NTU 60			NTU 120		GFLOPs	Param (M)
	Skeleton	RGB	XSub	XView	Average	XSub	XSet		
ST-GCN [16]	✓	–	81.5	88.3	84.6	79.0	81.3	3.8	3.1
2s-AGCN [17]	✓	–	88.5	95.1	91.8	82.9	84.9	8.8	3.5
MS-G3D [19]	✓	–	91.5	96.2	93.9	86.9	88.4	16.7	2.8
PoseConv3D [18]	✓	–	93.7	96.6	95.2	86.0	89.6	15.9	2.0
C3D [8]	–	✓	63.5	70.3	66.9	–	–	38.5	78.4
I3D-Resnet50 [40]	–	✓	93.2	97.7	95.3	–	–	51.7	33.0
STAR-Transformer [41]	✓	✓	92.0	96.5	94.3	–	–	–	–
TSMF [42]	✓	✓	92.5	97.4	95.0	87.0	89.1	85.4	20.8
VPN [21](I3D)	✓	✓	93.5	96.2	94.6	86.3	87.8	107.9	24.0
VPN++ [29]	◦	✓	91.9	94.9	93.4	86.7	89.3	–	–
VPN++ + 3D Poses [29]	✓	✓	94.9	98.1	96.5	90.7	92.5	125.8	15.5
TCEM-MMNet [31]	✓	✓	94.3	98.8	96.6	–	–	–	–
MMNet [22](Inception-v3)	✓	✓	95.3	98.4	96.8	92.9	94.4	89.2	34.2
π -ViT [†] [30]	◦	✓	94.0	97.9	96.0	91.9	92.9	590.0	121.4
Our proposed approach	✓	✓	96.1	99.0	97.6	92.4	94.3	8.1	2.5

Table 8

Comparison of mean top-1 accuracy (%) with state-of-the-art methods on Toyota SmartHome.[†] indicates the results are obtained using 10 clips per video. ◦ means the skeleton is used in training but not in inference.

Method	Modality		Toyota SmartHome		
	Skeleton	RGB	CS	CV2	Average
TSMF [42]	✓	✓	53.8	28.9	41.4
VPN [21](I3D)	✓	✓	65.2	54.1	59.7
VPN++ [29]	◦	✓	69	54.9	62.0
VPN++ + 3D Poses [29]	✓	✓	71.0	58.1	64.6
MMNet [22](ResNet18)	✓	✓	65.1	33.4	49.3
π -ViT [†] [30]	◦	✓	72.9	64.8	68.9
π -ViT + 3D Poses [†] [30]	✓	✓	73.1	65.0	69.1
Our proposed approach	✓	✓	71.7	67.8	69.8

Table 9

Comparison of top-1 accuracy (%) with state-of-the-art methods on the PKU-MMD dataset.

Method	Modality		PKU-MMD	
	Skeleton	RGB	XSub	XView
STA-LSTM [43]	✓	–	86.9	92.6
CNN-based [44]	✓	–	90.4	93.7
SRNet [45]	✓	–	93.1	97.0
TSMF [42]	✓	✓	95.8	97.8
MMNet [22](ResNet18)	✓	✓	96.3	98.0
MMNet [22](Inception-v3)	✓	✓	97.2	98.1
TCEM-MMNet [31]	✓	✓	96.8	98.0
Our proposed approach	✓	✓	96.2	98.4

one [38]. The similar performance of the two spatial-temporal attention blocks demonstrates the versatility of the proposed multimodal architecture across different attention block designs. Notably, our spatiotemporal attention block is significantly more parameter-efficient, requiring only 107 parameters, compared to the 2.87k parameters used by the attention block in [38]. This reduction in parameters highlights the efficiency of our approach without sacrificing the performance.

4.5. Comparison with state-of-the-art methods

In Table 7, we compare the performance of the proposed multimodal method with previous single-modal and multimodal methods on NTU-RGB+D 60 and NTU RGB+D 120 datasets. For the NTU-RGB+D 60 dataset, it is noticeable that our approach outperforms and competes with previous skeleton-based, RGB video-based, and multimodal methods. In particular, our approach boosts the average accuracies

of state-of-the-art skeleton-based, RGB video-based, and multimodal-based by 2.4%, 2.3%, and 0.8%, respectively. It also surpasses π -ViT, a distillation-based approach, by 1.2% average accuracy and achieves competitive performance on X-Sub and X-view protocols with 96.14% and 99.0% compared to 96.3% and 99.0% of π -ViT + 3D Poses. The proposed multimodal approach outperforms π -ViT and matches the performance of π -ViT + 3D Poses, while using 1-clip per video at testing compared to their 10 clips. Furthermore, our method leverages efficient CNN-based backbones for both visual and pose streams, unlike the Transformer-based architectures of π -ViT and π -ViT+ 3D Poses.

For the NTU RGB-D 120 dataset, we can notice that the proposed multimodal method achieves comparable top-1 accuracy with other methods while drastically decreasing FLOPs and the number of network parameters. Specifically, FLOPs of our method are less than those of TSMF [22], MMNet [22], and π -ViT [30] with 10.5x, 11.0x, and 72.8x, respectively. This shows that the proposed network meets the real-time requirement of a practical HAR system.

For Toyota-Smarthome dataset (Table 8), the proposed method surpasses prior works on cross-view2 (CV2) evaluation protocol with large margins (3.0%) and obtains the second-best mean top-1 accuracy on the cross-subject (CS) protocol. In particular, our approach enhanced the average accuracy of state-of-the-art Transformer-based π -ViT + 3D Poses [30] by 0.7%. The performance improvement may stem from utilizing 2D poses estimated with HRNet [32], which are of better quality than the 3D poses employed by π -ViT + 3D Poses [30], especially in the spontaneous activity-driven home environment. The lightweight and competitive performance of the proposed multimodal method makes it suitable for the recognition of activities of daily living.

For PKU-MMD [36] (Table 9), our EPAM-Net surpasses existing skeleton-based methods and achieves competitive performance against multimodal-based methods. Specifically, our EPAM-Net enhances the results of skeleton-based SRNet [45] by 3.1% and 1.4% for X-Sub and X-View evaluation protocols, respectively. Regarding multimodal-based approaches, it achieves the best performance (98.4% top-1 accuracy) under the X-View protocol and comparable performance under X-Sub protocols. The competitive performance of our EPAM-Net on the PKU-MMD dataset is consistent with that on the NTU RGB-D 60, RGB-D 120, and Toyota-Smarthome datasets, indicating the generalizability of our approach.

Fig. 4 shows the skeleton-based (Pose-X-ShiftNet) and the proposed multimodal-based action recognition accuracy per action for the difficult actions (e.g., actions shown in Fig. 1) on the X-Sub protocol of NTU RGB+D 120 and Toyota-Smarthome datasets. It is noticeable that the proposed multimodal method enhances the accuracy of such actions significantly compared to the pose-X-ShiftNet method.




Fig. 4. Classification accuracy per action for top-10 challenging actions on across-subject protocol of NTU RGB+D 120 (a) and Toyota-Smarthome (b) datasets.

5. Conclusion

Throughout this work, we addressed the extensive computation problem of existing multimodal action recognition networks by introducing a novel and efficient pose-driven attention-guided multimodal network that competes and is computationally efficient with the state of the art. The proposed network utilizes an efficient X-ShiftNet network to model the spatial and temporal information from RGB frames and their corresponding pose sequences. The spatial-temporal attention block guides the visual features to pay attention to the keyframes and their salient spatial regions using the pose features. The final classification score is obtained using the fused predictions of the RGB and the skeleton streams. NTU-60, NTU-120, PKU-MMD, and Toyota-Smarthome datasets showed competitive performance of the proposed method compared to state-of-the-art ones with up to 72.8x reduction in FLOPs and up to 48.6x reduction in the number of network parameters.

CRediT authorship contribution statement

Ahmed Abdelkawy: Writing – original draft, Software, Methodology, Investigation, Conceptualization. **Asem Ali:** Writing – review & editing, Supervision, Investigation. **Aly Farag:** Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

I have shared the link to my code at abstract.

References

- [1] M. Karim, S. Khalid, A. Aleryani, J. Khan, I. Ullah, Z. Ali, Human action recognition systems: A review of the trends and state-of-the-art, *IEEE Access* (2024).
- [2] M. Karim, S. Khalid, A. Aleryani, N. Tairan, Z. Ali, F. Ali, Hade: Exploiting human action recognition through fine-tuned deep learning methods, *IEEE Access* 12 (2024) 42769–42790.
- [3] A. Abdelkawy, A. Farag, I. Alkabbany, A. Ali, C. Foreman, T. Tretter, N. Hindy, Measuring student behavioral engagement using histogram of actions, *Pattern Recognit. Lett.* 186 (2024) 337–344.
- [4] S. Khalid, S. Wu, Supporting scholarly search by query expansion and citation analysis, *Eng. Technol. Appl. Sci. Res.* 10 (4) (2020) 6102–6108.
- [5] A. Garba, S. Khalid, I. Ullah, Understanding the impact of query expansion on federated search, *Multimedia Tools Appl.* 83 (4) (2024) 10393–10407.
- [6] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, L. Fei-Fei, Large-scale video classification with convolutional neural networks, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1725–1732.
- [7] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, T. Darrell, Long-term recurrent convolutional networks for visual recognition and description, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 2625–2634.
- [8] D. Tran, L. Bourdev, R. Fergus, L. Torresani, M. Paluri, Learning spatiotemporal features with 3d convolutional networks, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 4489–4497.
- [9] Z. Qiu, T. Yao, T. Mei, Learning spatio-temporal representation with pseudo-3d residual networks, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 5533–5541.
- [10] J. Carreira, A. Zisserman, Quo vadis, action recognition? a new model and the kinetics dataset, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6299–6308.
- [11] C. Feichtenhofer, H. Fan, J. Malik, K. He, Slowfast networks for video recognition, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 6202–6211.
- [12] C. Feichtenhofer, X3d: Expanding architectures for efficient video recognition, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 203–213.
- [13] Z. Sun, Q. Ke, H. Rahmani, M. Bennamoun, G. Wang, J. Liu, Human action recognition from various data modalities: A review, *IEEE Trans. Pattern Anal. Mach. Intell.* (2022).
- [14] V. Choutas, P. Weinzaepfel, J. Revaud, C. Schmid, Potion: Pose motion representation for action recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7024–7033.
- [15] C. Caetano, J. Sena, F. Brémond, J.A. Dos Santos, W.R. Schwartz, Skelemotion: A new representation of skeleton joint sequences based on motion information for 3d action recognition, in: *2019 16th IEEE International Conference on Advanced Video and Signal Based Surveillance, AVSS, IEEE*, 2019, pp. 1–8.
- [16] S. Yan, Y. Xiong, D. Lin, Spatial temporal graph convolutional networks for skeleton-based action recognition, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, 2018.
- [17] L. Shi, Y. Zhang, J. Cheng, H. Lu, Two-stream adaptive graph convolutional networks for skeleton-based action recognition, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12026–12035.
- [18] H. Duan, Y. Zhao, K. Chen, D. Lin, B. Dai, Revisiting skeleton-based action recognition, in: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR*, 2022, pp. 2959–2968.
- [19] Z. Liu, H. Zhang, Z. Chen, Z. Wang, W. Ouyang, Disentangling and unifying graph convolutions for skeleton-based action recognition, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 143–152.
- [20] T. Baltrušaitis, C. Ahuja, L.-P. Morency, Multimodal machine learning: A survey and taxonomy, *IEEE Trans. Pattern Anal. Mach. Intell.* 41 (2) (2018) 423–443.
- [21] S. Das, S. Sharma, R. Dai, F. Bremond, M. Thonnat, Vpn: Learning video-pose embedding for activities of daily living, in: *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, (2020) 23–28, Proceedings, Part IX 16*, Springer, 2020, pp. 72–90.
- [22] X. Bruce, Y. Liu, X. Zhang, S.-h. Zhong, K.C. Chan, Mmnet: A model-based multimodal network for human action recognition in rgb-d videos, *IEEE Trans. Pattern Anal. Mach. Intell.* 45 (3) (2022) 3522–3538.
- [23] K. Simonyan, A. Zisserman, Two-stream convolutional networks for action recognition in videos, *Adv. Neural Inf. Process. Syst.* 27 (2014).
- [24] J. Lin, C. Gan, S. Han, Tsm: Temporal shift module for efficient video understanding, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 7083–7093.
- [25] M. Liu, J. Yuan, Recognizing human actions as the evolution of pose estimation maps, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1159–1168.

- [26] M. Zolfaghari, G.L. Oliveira, N. Sedaghat, T. Brox, Chained multi-stream networks exploiting pose, motion, and appearance for action classification and detection, in: Proceedings of the IEEE International Conference on Computer Vision, 2017, pp. 2904–2913.
- [27] J. Li, X. Xie, Q. Pan, Y. Cao, Z. Zhao, G. Shi, Sgm-net: Skeleton-guided multimodal network for action recognition, *Pattern Recognit.* 104 (2020) 107356.
- [28] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, M. Paluri, A closer look at spatiotemporal convolutions for action recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 6450–6459.
- [29] S. Das, R. Dai, D. Yang, F. Bremond, Vpn++: Rethinking video-pose embeddings for understanding activities of daily living, *IEEE Trans. Pattern Anal. Mach. Intell.* 44 (12) (2021) 9703–9717.
- [30] D. Reilly, S. Das, Just add?! pose induced video transformers for understanding activities of daily living, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 18340–18350.
- [31] D. Liu, F. Meng, Q. Xia, Z. Ma, J. Mi, Y. Gan, M. Ye, J. Zhang, Temporal cues enhanced multimodal learning for action recognition in rgb-d videos, *Neurocomputing* 594 (2024) 127882.
- [32] K. Sun, B. Xiao, D. Liu, J. Wang, Deep high-resolution representation learning for human pose estimation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 5693–5703.
- [33] Q. Wang, B. Wu, P. Zhu, P. Li, W. Zuo, Q. Hu, Eca-net: Efficient channel attention for deep convolutional neural networks, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 11534–11542.
- [34] A. Shahroudy, J. Liu, T.-T. Ng, G. Wang, Ntu rgb+d: A large scale dataset for 3d human activity analysis, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 1010–1019.
- [35] J. Liu, A. Shahroudy, M. Perez, G. Wang, L.-Y. Duan, A.C. Kot, Ntu rgb+d 120: A large-scale benchmark for 3d human activity understanding, *IEEE Trans. Pattern Anal. Mach. Intell.* 42 (10) (2019) 2684–2701.
- [36] C. Liu, Y. Hu, Y. Li, S. Song, J. Liu, Pku-mmd: A large scale benchmark for continuous multi-modal human action understanding, 2017, arXiv preprint arXiv:1703.07475.
- [37] S. Das, R. Dai, M. Koperski, L. Minciullo, L. Garattoni, F. Bremond, G. Francesca, Toyota smarhome: Real-world activities of daily living, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 833–842.
- [38] J. Li, P. Wei, N. Zheng, Nesting spatiotemporal attention networks for action recognition, *Neurocomputing* 459 (2021) 338–348.
- [39] J. Hu, L. Shen, G. Sun, Squeeze-and-excitation networks, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 7132–7141.
- [40] J. Zhu, W. Zou, Z. Zhu, L. Xu, G. Huang, Action machine: Toward person-centric action recognition in videos, *IEEE Signal Process. Lett.* 26 (11) (2019) 1633–1637.
- [41] D. Ahn, S. Kim, H. Hong, B.C. Ko, Star-transformer: a spatio-temporal cross attention transformer for human action recognition, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2023, pp. 3330–3339.
- [42] X. Bruce, Y. Liu, K.C. Chan, Multimodal fusion via teacher-student network for indoor action recognition, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, 2021, pp. 3199–3207.
- [43] S. Song, C. Lan, J. Xing, W. Zeng, J. Liu, Spatio-temporal attention-based lstm networks for 3d action recognition and detection, *IEEE Trans. Image Process.* 27 (7) (2018) 3459–3471.
- [44] C. Li, Q. Zhong, D. Xie, S. Pu, Skeleton-based action recognition with convolutional neural networks, in: 2017 IEEE International Conference on Multimedia & Expo Workshops, ICMEW, IEEE, 2017, pp. 597–600.
- [45] W. Nie, W. Wang, X. Huang, Srnet: Structured relevance feature learning network from skeleton data for human action recognition, *IEEE Access* 7 (2019) 132161–132172.



Ahmed Abdelkawy received B.Sc. and M.Sc. degrees in Computer Science from Helwan University, Egypt in 2012 and 2018, respectively. He is currently pursuing a Ph. D degree in Electrical and Computer Engineering at the University of Louisville. His research interests include video analysis, action recognition, and affective computing.



Asem Ali received B.Sc. and M.Sc. degrees in Electrical Engineering from Assiut University, Egypt in 1999 and 2002, respectively. He received his Ph.D. degree from University of Louisville in ECE in 2008. He has been a post-doctoral fellow then a research scientist with the CVIP Lab, UoFL, during the period 2008-2012 and 2015-2024, respectively. Also, he spent three years with the College of Engineering at Assuit University as an assistant professor. He is an instructor at Bioengineering and ECE departments, J. B. Speed School of Engineering, UoFL. His research interest intersects Artificial Intelligence, Affective Computing, Image Analysis, and Computer Vision.



Aly A. Farag, Fellow, IEEE and IAPR: Received the bachelor's degree in electrical engineering from Cairo University, Egypt and the Ph.D. degree from Purdue University in Electrical Engineering. He also holds master's degrees in biomedical engineering from the Ohio State University and the University of Michigan, Ann Arbor. He joined the University of Louisville in August 1990, where he is currently a Professor of Computer Science and Engineering (CSE), and Professor of Electrical and Computer Engineering (ECE). At the University of Louisville, Dr. Farag founded the Computer Vision and Image Processing Laboratory (CVIP Lab) which focuses on imaging science, computer vision and biomedical imaging. Dr. Farag main research focus is scene analysis, object reconstruction from multimodality imaging, statistical and variational methods for object modeling, and facial biometrics. He has authored over 400 technical papers, edited two volumes on Deformable Models for Biomedical Applications (Springer 2007). He is the author of two textbooks: Biomedical Image Analysis: Variational and Statistical Approaches – Cambridge University Press, 2014, and Introduction to Probability Theory with Engineering Application, Cognella, 2020.