

SkateFormer: Skeletal-Temporal Transformer for Human Action Recognition

Jeonghyeok Do[✉] and Munchurl Kim[✉]

Korea Advanced Institute of Science and Technology, South Korea
{ehwjdgur0913, mkimee}@kaist.ac.kr
https://kaist-viclab.github.io/SkateFormer_site/

Abstract. Skeleton-based action recognition, which classifies human actions based on the coordinates of joints and their connectivity within skeleton data, is widely utilized in various scenarios. While Graph Convolutional Networks (GCNs) have been proposed for skeleton data represented as graphs, they suffer from limited receptive fields constrained by joint connectivity. To address this limitation, recent advancements have introduced transformer-based methods. However, capturing correlations between all joints in all frames requires substantial memory resources. To alleviate this, we propose a novel approach called Skeletal-Temporal Transformer (SkateFormer) that partitions joints and frames based on different types of skeletal-temporal relation (Skate-Type) and performs skeletal-temporal self-attention (Skate-MSA) within each partition. We categorize the key skeletal-temporal relations for action recognition into a total of four distinct types. These types combine (i) two skeletal relation types based on physically neighboring and distant joints, and (ii) two temporal relation types based on neighboring and distant frames. Through this partition-specific attention strategy, our SkateFormer can selectively focus on key joints and frames crucial for action recognition in an action-adaptive manner with efficient computation. Extensive experiments on various benchmark datasets validate that our SkateFormer outperforms recent state-of-the-art methods.

Keywords: Skeleton-based Action Recognition · Transformer · Partition-Specific Attention

1 Introduction

In recent years, human action recognition (HAR) [39, 51, 57] has gained widespread applications in real-life scenarios, involving the classification of actions based on human movements. A diverse range of data sources such as videos captured from RGB cameras, optical flow generated through post-processing, 2D/3D skeletons estimated from RGB videos [2, 50], and skeletons acquired from sensors [16], contain information about human movements that can be leveraged for action recognition. However, using RGB videos as input is challenging. They are sensitive to external factors such as lighting conditions, camera distance and

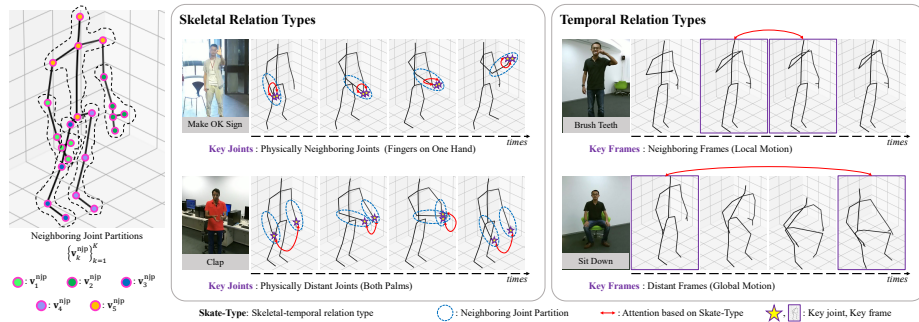


Fig. 1: SkateFormer’s partition-specific attention strategy for skeleton-based action recognition, leveraging two skeletal relation types – physically neighboring and distant joints – and two temporal relation types for local and global motion.

angles, and background variations [12]. Moreover, RGB videos require large storage capacity due to their detailed background visuals. Conversely, 3D skeletons obtained through sensors can offer a compact representation with robustness to external environmental changes, and do not require additional post-processing modules such as pose and optical flow estimations [2, 15, 49, 50].

The joints and their connections (or bones) in skeleton input correspond to the vertices and edges, respectively, in a graph structure. Consequently, many methods based on Graph Convolutional Networks (GCNs) [18] suitable for processing graph inputs have been extensively proposed for human action recognition. Most of the GCN-based methods [4–6, 10, 19–22, 24, 36, 43, 46, 47, 61, 63, 67, 68, 75] exchange information between different joints within a single frame using graph convolutions, and capture the temporal dynamics for each joint using 1D temporal convolutions. However, they struggle to capture the relation of physically distant joints (e.g., between two hands with open arms) effectively due to the direct propagation of information between physically connected joints.

To mitigate this limitation, transformer-based methods utilizing self-attention to capture relations across all joint pairs have been proposed [40, 58, 60, 64]. Yet, considering every joint in every frame is inefficient as specific joints in certain frames are more critical for particular action recognition. Some transformer-based methods have tried to reduce computational complexity by squeezing features along the joint or frame dimension before self-attention [11], only using either skeletal or temporal relations [13, 38, 44, 73, 76], or by tokenizing physically similar skeletal information [34, 62]. Separating joints and frames during self-attention can obscure the critical skeletal-temporal relations for action recognition, and tokenization may dilute vital information, potentially degrading performance.

To overcome these issues, we propose an efficient transformer-based approach, called Skeletal-Temporal Transformer (SkateFormer) that introduces joint and frame partition strategies and partition-specific self-attention based on types of skeletal-temporal relations (Skate-Type). Fig. 1 illustrates the joint and frame

partition strategies and partition-specific self-attention of our SkateFormer. For example, in the ‘make ok sign’ action class, relations between physically neighboring joints (e.g., joints on the same hand) are crucial, whereas in the ‘clap’ action class, relations between physically distant joints (e.g., between the palms of both hands) are more critical. Regarding temporal relations, for repetitive local motions like in the ‘brush teeth’ class, neighboring frame relations are essential, whereas for actions with global motion like ‘sit down’, distant frame relations become critical. Additionally, the speed at which actions are performed can vary significantly depending on the actor. To make our SkateFormer computationally efficient, we introduce a novel partition-specific attention (Skate-MSA). For this, we divide the skeletal-temporal relations into four partition types: (i) neighboring joints and local motion – Skate-Type-1, (ii) distant joints and local motion – Skate-Type-2, (iii) neighboring joints and global motion – Skate-Type-3 and (iv) distant joints and global motion – Skate-Type-4. Therefore, Skate-MSA is designed to efficiently capture skeletal-temporal relations at the joint-element-level, eliminating the need for tokenization, which typically involves joint-group-level attention [34, 62]. In summary, our contributions are as follows:

- We propose a Skeletal-Temporal Transformer (SkateFormer), a partition-specific attention strategy (Skate-MSA) for skeleton-based action recognition that captures skeletal-temporal relations and reduces computational complexity.
- We introduce a range of augmentation techniques and an effective positional embedding method, named Skate-Embedding, which combines skeletal and temporal features. This method significantly enhances action recognition performance by forming an outer product between learnable skeletal features and fixed temporal index features.
- Our SkateFormer sets a new state-of-the-art for action recognition performance across multiple modalities (4-ensemble condition) and single modalities (joint, bone, joint motion, bone motion), showing notable improvement over the most recent state-of-the-art methods. Additionally, it concurrently establishes a new state-of-the-art in interaction recognition, a sub-field of action recognition.

2 Related Works

2.1 RNN/CNN-based approaches

Recently, skeleton-based action recognition has made significant progress. Early research efforts focused on utilizing Recurrent Neural Networks (RNNs), including LSTM and GRU, to handle skeleton data due to its sequential and continuous nature [25, 28, 29, 37, 45, 71, 77]. Furthermore, some studies explored the conversion of skeleton data into pseudo-images to leverage Convolutional Neural Networks (CNNs) [8, 17, 23, 59, 65, 66]. Very recently, PoseC3D [12] employed a 2D pose estimation network [50] on RGB video to generate image-like skeleton heatmaps as input for CNNs, and Ske2Grid [1] proposed a method to transform 2D/3D skeleton data nodes into image-format grid patches.

2.2 GCN-based approaches

Skeleton data is composed of joints and bones, which correspond to vertices and edges in a graph. Consequently, research in this domain has predominantly gravitated towards Graph Neural Networks (GNNs) [42, 69] and Graph Convolutional Networks (GCNs) [3–6, 10, 19–21, 24, 36, 43, 46, 47, 61, 67, 68]. InfoGCN [7] introduced additional loss terms to form a compact latent feature space, resulting in clear decision boundaries. Similarly, FR-Head [75] incorporated contrastive learning to separate feature representations of ambiguous classes. LST [63] introduced language supervision to train GCNs effectively, while HD-GCN [22] introduced rooted trees as a data structure to create diverse connections between joints. Despite their merits, GCNs have limitations in effectively capturing topology variations depending on the input data by using the fixed-sized kernels.

2.3 Transformer-based approaches

Recent research explores transformer-based approaches for skeleton-based action recognition, which excel in capturing data-adaptive joint connections [40, 58, 60, 64]. However, these methods often require high computational costs due to the usage of their large-sized attention maps. Early methods have tried to reduce attention map sizes by employing feature pooling [11] or reshaping [44] techniques. Some studies have attempted to mitigate computational load by utilizing separate skeletal and temporal attention modules in a parallel [38, 73] or serial [13, 44, 76] configurations. Nevertheless, they face challenges in simultaneously capturing skeletal-temporal relations, which are critical aspects of skeleton-based action recognition. IGFormer [34] and ISTA-Net [62] tackle the challenge by embedding joint sets within the same partition into a unified token before attention modules, relying solely on partition strategies for tokenization without introducing partition-specific attention mechanisms. Despite incorporating skeletal-temporal attention at the joint-group-level, they encounter limitations in local action recognition due to the tokenization process leading to the loss of physically similar skeletal information. In our work, the proposed SkateFormer employs partition-specific skeletal-temporal attention modules to effectively capture skeletal-temporal relations at the joint-element-level without tokenization.

Complexity of transformer-based vision tasks. Vision Transformers (ViT) [9] revolutionized the field of computer vision but face challenges due to high computational costs. To mitigate these challenges, various studies have been conducted to reduce the size of attention maps while efficiently propagating local and global information [30, 54]. In the context of vision tasks, 2D feature maps are organized along horizontal and vertical spatial axes. In contrast, our task involves skeletal and temporal axes, making the nature of our data fundamentally different. However, the elements within skeleton data tend to scale mostly with the factors such as the number of recognized joints, the number of individuals performing actions, and the number of frames, necessitating the reduction of computational complexity. Inspired by the works [53, 54], we propose a transformer structure that efficiently captures skeletal-temporal relations at the joint-element-level, reducing the computational cost.

3 Methodology

3.1 Overview of SkateFormer

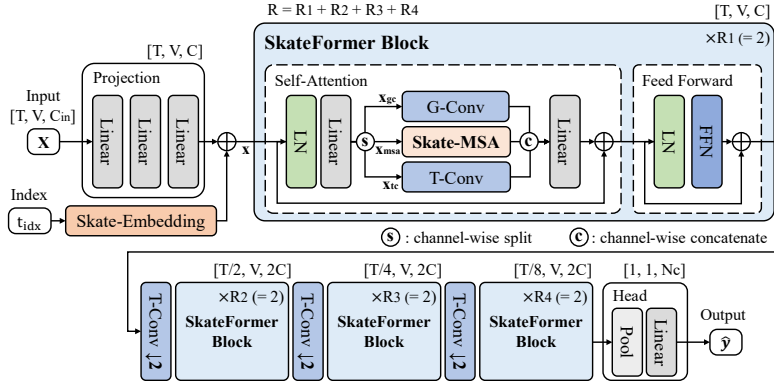


Fig. 2: The overall framework of our proposed SkateFormer.

The overall architecture of our proposed SkateFormer is depicted in Fig. 2. A skeleton sequence, corresponding to a single action, is sampled to maintain a consistent frame count of T , resulting in $\mathbf{X} \in \mathbb{R}^{T \times V \times M \times C_{in}}$, where V represents the number of joints per frame, M denotes the number of individuals involved in the action, and C_{in} is the dimensionality of data representing a single joint ($C_{in} = 3$ for 3D skeletons in our cases). Initially, we reshape \mathbf{X} from (T, V, M, C_{in}) to $(T, V \cdot M, C_{in})$ to treat joints from different individuals separately. For simplicity, we redefine V as $V \cdot M$ so that we have $\mathbf{X} \in \mathbb{R}^{T \times V \times C_{in}}$ in subsequent text. Next, the three linear layers in the SkateFormer map the low-dimensional raw skeleton data into a higher-dimensional feature space. We then perform the skeletal-temporal positional embedding by adding learnable skeletal features and fixed temporal index features to this mapped feature. The embedded features pass through R SkateFormer blocks, each comprising a self-attention layer that propagates features via skeletal-temporal relations, followed by a feed-forward layer that refines the features. The final features, after R SkateFormer Blocks, undergo skeletal-temporal pooling to produce the outcome $\hat{y} \in \mathbb{R}^{N_c}$, which is trained via backpropagation through a loss function by comparing \hat{y} with the true label y , where N_c represents the number of classes.

3.2 SkateFormer Block

Each of the SkateFormer Blocks in Fig. 2 is structured similarly as the traditional transformer blocks [30, 54], incorporating both a self-attention layer and a feed-forward network (FFN). The self-attention component of the SkateFormer Block can be expressed as follows:

$$\begin{aligned}
[\mathbf{x}_{gc}, \mathbf{x}_{tc}, \mathbf{x}_{msa}] &= \text{Split}(\text{Linear}(\text{LN}(\mathbf{x}))) \\
\mathbf{x}_{gc} &\leftarrow \text{G-Conv}(\mathbf{x}_{gc}) \\
\mathbf{x}_{tc} &\leftarrow \text{T-Conv}(\mathbf{x}_{tc}) \\
\mathbf{x}_{msa} &\leftarrow \text{Skate-MSA}(\mathbf{x}_{msa}) \\
\mathbf{x} &\leftarrow \mathbf{x} + \text{Linear}(\text{Concat}(\mathbf{x}_{gc}, \mathbf{x}_{tc}, \mathbf{x}_{msa})),
\end{aligned} \tag{1}$$

where the LN represents the Layer Normalization layer, the Linear denotes a fully-connected layer, and the Split indicates a channel splitting function. When input \mathbf{x} to the self-attention component has C channels, we split \mathbf{x} into \mathbf{x}_{gc} of $C/4$ channels, \mathbf{x}_{tc} of $C/4$ channels, and \mathbf{x}_{msa} of $C/2$ channels. To obtain inductive biases for both skeletal and temporal aspects, we utilize a one-layer GCN (G-Conv) and a single temporal convolution layer (T-Conv), respectively. With a total of H heads, we employ a learnable matrix of shape $(H/4, V, V)$ instead of a predefined adjacency matrix to perform G-Conv operations, enabling us to capture diverse connectivity patterns between joints. T-Conv is a 1D convolution layer with a kernel size of k , and performs 1D ($H/4$ -group) convolutions for \mathbf{x}_{tc} to capture temporal dynamics. The FFN of the SkateFormer Block is expressed as follows: $\mathbf{x} \leftarrow \mathbf{x} + \text{Linear}(\text{Act}(\text{Linear}(\text{LN}(\mathbf{x}))))$, where Act represents the activation layer. When downsampling is needed for the input to a SkateFormer Block, a 1D convolution with a stride of 2 and Batch Normalization layer are applied.

3.3 Skate-MSA




Fig. 3: The Skate-MSA of our proposed SkateFormer.

As shown in Fig. 3, the feature map input \mathbf{x}_{msa} to the Skate-MSA is first split channel-wise into four equal-sized features $\mathbf{x}_{msa}^1, \mathbf{x}_{msa}^2, \mathbf{x}_{msa}^3, \mathbf{x}_{msa}^4$, each of which has $C/8$ channels. For each \mathbf{x}_{msa}^i , self-attention operation is applied to discern the correlations between joints corresponding to specific relation types as:

$$\begin{aligned}
\mathbf{x}_{msa}^i &\leftarrow \mathcal{R}_i(\text{MSA}(\mathcal{P}_i(\mathbf{x}_{msa}^i))) \\
\mathbf{x}_{msa} &\leftarrow \text{Concat}(\mathbf{x}_{msa}^1, \mathbf{x}_{msa}^2, \mathbf{x}_{msa}^3, \mathbf{x}_{msa}^4),
\end{aligned} \tag{2}$$

where \mathcal{P}_i and \mathcal{R}_i represent the i -th Skate-Type partition and reverse operations, respectively. This approach enables the Skate-MSA to effectively analyze and model various joint relations in a specialized manner, contributing to its overall performance improvement.

Partition and Reverse. Let a frame be $\mathbf{f}^t \in \mathbb{R}^{V \times C_{in}}$ at time t . So, we have $\mathbf{X} = \{\mathbf{f}^t\}_{t=1}^T \in \mathbb{R}^{T \times V \times C_{in}}$ where \mathbf{f}^t consists of the 3D ($C_{in} = 3$) positions of V joints at time t . The frames of \mathbf{X} are often temporally highly correlated for most of the action types. However, the joint positions of skeletons tend to be less correlated among some joints, depending on different action types. Nevertheless, some other joints such as the joints of the head, neck, shoulders, abdomen and pelvises are relatively not highly variant among themselves for various action types. That is, they may move together in a group. So, it is

efficient to partition the whole set of joints into smaller-sized joint partitions that can be useful to distinguish different actions. Furthermore, it is also worthwhile to consider categorizing the movements of various actions into two partitions, such as local and global motions: local motion, which only changes the positions of a small number of joints with partial movements such as teeth brushing and clapping, and global motion, which appears in the entire frames and includes actions like sitting down and standing up. Motivated from this, we partition the action into four skeletal-temporal relation types. For this, we first partition the entire set of joints into total non-overlapping K subsets as neighboring joint partitions \mathbf{v}_k^{njp} . For example, when $K = 5$, we may have \mathbf{v}_1^{njp} , \mathbf{v}_2^{njp} , \mathbf{v}_3^{njp} , \mathbf{v}_4^{njp} and \mathbf{v}_5^{njp} as the right arm, left arm, right leg, left leg and torso, respectively. The elements within each subset are ordered in a manner that extends outward from the body’s central region (e.g., for $k = 3$, $v_{3,1}$, $v_{3,2}$ and $v_{3,3}$ are the 3D coordinates of the pelvis, right knee and right foot positions, respectively):

$$\mathbf{v}_k^{njp} = [v_{k,1}, v_{k,2}, \dots, v_{k,L}], \quad (3)$$

where $k = 1, 2, \dots, K$, L represents the total number of its elements, and $\mathbf{v}_i^{njp} \cap \mathbf{v}_j^{njp} = \emptyset$ for $i \neq j$. We stack \mathbf{v}_k^{njp} ’s, thus creating the skeletal axis as:

$$\mathbf{v} = [\mathbf{v}_1^{njp} | \mathbf{v}_2^{njp} | \dots | \mathbf{v}_K^{njp}], \quad (4)$$

where $|\mathbf{v}| = KL = V$. Also, to consider the relations between the joints that can be physically separated in a distance each other, we group the same-positioned elements of \mathbf{v}_k^{njp} to create possibly distant joint partitions \mathbf{v}_l^{djp} as:

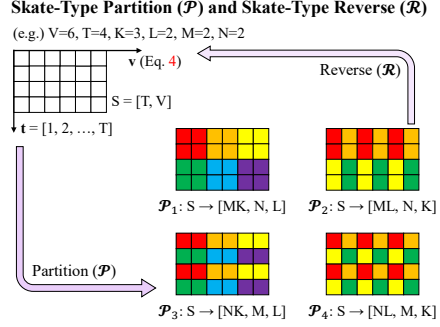


Fig. 4: Skate-Type partition and reverse.

$$\mathbf{v}_l^{\text{djp}} = [v_{1,l}, v_{2,l}, \dots, v_{K,l}], \quad (5)$$

where $l = 1, 2, \dots, L$. Similarly, we define $\mathbf{t} = [1, 2, \dots, T]$ as a time axis. We further define two semantic time axes as $\mathbf{t}_m^{\text{local}}$ for local motion comprehension and $\mathbf{t}_n^{\text{global}}$ for global motion understanding as:

$$\begin{aligned} \mathbf{t}_m^{\text{local}} &= [(m-1)N+1, (m-1)N+2, \dots, mN] \\ \mathbf{t}_n^{\text{global}} &= [n, n+N, \dots, n+(M-1)N], \end{aligned} \quad (6)$$

where $m = 1, 2, \dots, M$ with the total M number of $\mathbf{t}_n^{\text{global}}$ elements, and $n = 1, 2, \dots, N$ with the total N number of $\mathbf{t}_m^{\text{local}}$ elements. Note that $\mathbf{t}_m^{\text{local}}$ is a segment of consecutive time indices to represent a local motion within the time segment while $\mathbf{t}_n^{\text{global}}$ is an N -strided sparse time axis to capture the global motion over \mathbf{t} . So, we have $\mathbf{t} = \{\mathbf{t}_m^{\text{local}}\}_{m=1}^M = \{\mathbf{t}_n^{\text{global}}\}_{n=1}^N$ and $|\mathbf{t}| = MN = T$. Based on these time axes, our skeletal-temporal partitions of joints and frames are explained in the followings.

We partition the joints and frames together into four types in the context of skeletal-temporal relation for the Skate-MSA – Skate-Type-1, -2, -3 and -4: (i) The Skate-Type-1 partition, denoted as \mathcal{P}_1 , pertains to a self-attention branch targeting neighboring joints and local motion, based on $\mathbf{v}_k^{\text{njp}}$ and $\mathbf{t}_m^{\text{local}}$; (ii) the Skate-Type-2 partition (\mathcal{P}_2) represents a branch for distant joints and local motion, based on $\mathbf{v}_l^{\text{djp}}$ and $\mathbf{t}_m^{\text{local}}$; (iii) the Skate-Type-3 partition (\mathcal{P}_3) signifies a branch for neighboring joints and global motion, based on $\mathbf{v}_k^{\text{njp}}$ and $\mathbf{t}_n^{\text{global}}$; (iv) Lastly, the Skate-Type-4 partition (\mathcal{P}_4) corresponds to a branch targeting distant joints and global motion, based on $\mathbf{v}_l^{\text{djp}}$ and $\mathbf{t}_n^{\text{global}}$. The Skate-Type partition operations transform the shape ($S = (T, V, c)$) of $\mathbf{x}_{\text{msa}}^i$ into:

$$\begin{aligned} \mathcal{P}_1 : S &\rightarrow (MK, N, L, c) & \mathcal{P}_2 : S &\rightarrow (ML, N, K, c) \\ \mathcal{P}_3 : S &\rightarrow (NK, M, L, c) & \mathcal{P}_4 : S &\rightarrow (NL, M, K, c), \end{aligned} \quad (7)$$

where $c = C/8$. The partitioned feature map $\mathbf{x}_{\text{msa}}^{i, \mathcal{P}} = \mathcal{P}_i(\mathbf{x}_{\text{msa}}^i)$ undergoes multi-head self-attention (MSA) and is then reshaped back to its original size of (T, V, c) through a Skate-Type reverse operation \mathcal{R}_i according to Eq. 2.

Multi-head self-attention. In order to collectively consider the skeletal-temporal relation types, we can generalize feature maps $\mathbf{x}_{\text{msa}}^{i, \mathcal{P}}$ to have a shape of (B, T', V', c) . The feature maps are first reshaped into $(B, T' \cdot V', c)$. Through linear mappings of $\mathbf{x}_{\text{msa}}^{i, \mathcal{P}}$ of shape (c, c) by \mathbf{W}_Q , \mathbf{W}_K , and \mathbf{W}_V , we obtain the query (\mathbf{Q}), key (\mathbf{K}), and value (\mathbf{V}) tensors as $[\mathbf{Q}, \mathbf{K}, \mathbf{V}] = \mathbf{x}_{\text{msa}}^{i, \mathcal{P}} \cdot [\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V]$.

Since we utilized half of the total H heads for G-Conv and T-Conv operations, the remaining $H/2$ heads are divided into quarters with $H' = H/8$. They are then assigned differently with the four Skate-Types, each of which has a MSA. We reshaped $\mathbf{Q}, \mathbf{K}, \mathbf{V}$ to have a shape of $(B, H', T' \cdot V', c/H')$. This refers to the self-attention (SA) for each individual head h , given as:

$$\begin{aligned} \text{SA}_h(\mathbf{x}_{\text{msa}}^{i,\mathcal{P}}) &= \text{SoftMax}(\mathbf{Q}_h \mathbf{K}_h^T / \sqrt{c/H'} + \mathbf{B}_h) \mathbf{V}_h \\ \text{MSA}(\mathbf{x}_{\text{msa}}^{i,\mathcal{P}}) &= \text{Concat}(\text{SA}_1(\mathbf{x}_{\text{msa}}^{i,\mathcal{P}}), \dots, \text{SA}_{H'}(\mathbf{x}_{\text{msa}}^{i,\mathcal{P}})), \end{aligned} \quad (8)$$

where $\mathbf{Q}_h \mathbf{K}_h^T$ is a data-dependent term that varies based on the input, and \mathbf{B}_h represents the skeletal-temporal positional bias.

Complexity analysis. The computational complexity of a naive self-attention layer [9, 55] for feature map \mathbf{x}_{msa} with a shape of $(T, V, C/2)$ is $2(VT)^2(C/2)$. In contrast, the computation complexity of our Skate-MSA is:

$$2(VT)^2 \frac{C}{2} \left[\frac{1}{4} \left(\frac{1}{MK} + \frac{1}{ML} + \frac{1}{NK} + \frac{1}{NL} \right) \right]. \quad (9)$$

In our experimental settings, this results in approximately a $48\times$ reduction in computational complexity compared to the naive self-attention layer.

3.4 Skeletal-Temporal Positional Embedding

Our novel skeletal-temporal positional embedding method, called Skate-Embedding, is tightly associated with temporal data augmentation. So, we first explain data augmentation with our contribution and then describe the Skate-Embedding.

Intra-instance augmentation. Previous skeleton-based action recognition methods often used data augmentation by temporally sampling input frames with fixed strides or randomly for the whole input, which we refer to as *temporal augmentation*. Also, as *skeletal augmentation*, various transformations such as the actor order permutation, random shear, random rotation, random scaling, random coordinate drop, and random joint dropout have been applied [14, 26].

We propose trimmed-uniform random sampling of frames with p portion. This sampling cuts out the first and last parts of the total input sequence and performs uniform random sampling of frames. From our trimmed-uniform random sampling of frames with a p portion, the masking effect of skeleton sequences in their front and back portions, as well as a more dense sampling effect in the middle, are expected. This can lead to a concentration in the middle portion, resulting in better data augmentation.

Inter-instance augmentation. We define inter-instance augmentation as data augmentation by exchanging the bone lengths of different subjects across different frame sequences (not within each frame sequence). By doing so, the resulting data augmentation can provide the diversity of subjects with different body sizes, which can help generalization learning.

Skate-Embedding. We propose a novel skeletal-temporal positional embedding method, called Skate-Embedding, that utilizes fixed (not learnable) temporal index features and learnable (not fixed) skeletal features. The temporal index features are suitable for conveying to the first SkateFormer Block the temporal positional information of the sampled frames from sequences of various lengths.

The sampled temporal indices are designated as $t_{\text{idx}} = [t_1, t_2, \dots, t_T]$. These temporal indices are then normalized to the range $[-1, 1]$, and are used for the fixed temporal index features as done in the temporal positional embedding. The fixed temporal index features, denoted as $\text{TE} \in \mathbb{R}^{T \times C}$, are constructed for t_{idx} by using the sinusoidal positional embedding [55]. On the other hand, as skeletal joint positional (which are not the 3D coordinates of joints but their indices) embeddings, the learnable skeletal features, denoted as $\text{SE} \in \mathbb{R}^{V \times C}$, are learned within the Skate-Embedding as shown in Fig. 2. Finally, the skeletal-temporal positional embedding $\text{STE} \in \mathbb{R}^{T \times V \times C}$ is done by taking the outer product of SE and TE as $\text{STE}[i, j, d] = \text{SE}[j, d] \cdot \text{TE}[i, d]$ at the i -th time, the j -th joint and the d -th channel.

4 Experimental Results

4.1 Datasets

NTU RGB+D. This dataset [41] offers 60 action classes and includes diverse activities like drinking water, eating, brushing teeth, dropping objects, and more. It comprises total 56,880 videos captured from 40 subjects across 155 camera viewpoints. The dataset utilizes Kinect v2, encompassing RGB, IR, depth, and 3D skeleton data, and supports cross-subject (*X-Sub60*) and cross-view (*X-View60*) evaluation. Out of the 60 action classes, only 11 are related to human interaction, specifically when two individuals are present. We denote this subset as *NTU-Inter* [11, 34].

NTU RGB+D 120. This dataset [27] extends the NTU RGB+D dataset [41] with 120 action classes, covering actions such as putting on headphones, basketball shooting, juggling table tennis balls, and more. It contains 114,480 videos from 106 subjects across 155 camera viewpoints. It enables cross-subject (*X-Sub120*) and cross-setup (*X-Set120*) evaluation based on different subject and camera setups. Out of the 120 action classes, only 26 are related to human interaction. We denote this subset as *NTU-Inter 120* [11, 34, 62].

NW-UCLA. This dataset [56] includes 10 action classes, such as picking up objects, dropping trash, walking, sitting, standing, and more. It comprises total 1,475 videos from 10 subjects across three camera views, collected using Kinect v1 [74] and featuring RGB, IR, depth, and 3D skeleton data. Evaluation follows the cross-view approach, utilizing two training views and one test view.

4.2 Experiment Details

All experiments were performed using the PyTorch framework [35], running on a single NVIDIA DGX A100 GPU. Each model was trained with a total of 500 epochs. We employed a linear warm-up strategy for the learning rate, gradually increasing it from 10^{-7} to 10^{-3} during the first 25 epochs. Subsequently, a cosine-annealing scheduler [31] was used to update the learning rate at each iteration for the remaining epochs. We employed the AdamW optimizer [32] with a betas

Table 1: Top-1 accuracy of different skeleton-based action recognition methods on the NTU RGB+D, NTU RGB+D 120 and NW-UCLA datasets. The methods of RNN, CNN, GCN, and Transformer types are evaluated based on the number of input frames and the use of ensemble strategies. The **best performances** are highlighted in **bold**.

Types	Methods	Frames	NTU RGB+D (%)						NTU RGB+D 120 (%)						NW-UCLA (%)
			X-Sub60			X-View60			X-Sub120			X-Set120			
			\mathbb{E}_1	\mathbb{E}_2	\mathbb{E}_4	\mathbb{E}_1	\mathbb{E}_2	\mathbb{E}_4	\mathbb{E}_1	\mathbb{E}_2	\mathbb{E}_4	\mathbb{E}_1	\mathbb{E}_2	\mathbb{E}_4	
RNN	AGC-LSTM [45]	100	87.5	89.2	-	93.5	95.0	-	-	-	-	-	-	-	93.3
CNN	TA-CNN [65]	64	88.8	-	90.4	93.6	-	94.8	82.4	-	85.4	84.0	-	86.8	96.1
	Ske2Grid [1]	100	88.3	-	-	95.7	-	-	82.7	-	-	85.1	-	-	-
GCN	SGN [72]	20	-	89.0	-	-	94.5	-	-	79.2	-	-	81.5	-	-
	CTR-GCN [3]	64	89.9	-	92.4	-	-	96.8	84.9	88.7	88.9	-	90.1	90.6	96.5
	ST-GCN++ [10]	100	89.3	91.4	92.1	95.6	96.7	97.0	83.2	87.0	87.5	85.6	87.5	89.8	-
	InfoGCN [7]	64	-	-	92.7	-	-	96.9	85.1	88.5	89.4	86.3	89.7	90.7	96.6
	FR-Head [75]	64	90.3	92.3	92.8	95.3	96.4	96.8	85.5	-	89.5	87.3	-	90.9	96.8
	Koopman [61]	64	90.2	-	92.9	95.2	-	96.8	85.7	-	90.0	87.4	-	91.3	97.0
	LST [63]	64	90.2	-	92.9	95.6	-	97.0	85.5	-	89.9	87.0	-	91.1	97.2
	HD-GCN [22]	64	90.6	92.4	93.0	95.7	96.6	97.0	85.7	89.1	89.8	87.3	90.6	91.2	96.9
	STC-Net [21]	64	-	92.5	93.0	-	96.7	97.1	-	89.3	89.9	-	90.7	91.3	97.2
Transformer	DSTA-Net [44]	128	-	-	91.5	-	-	96.4	-	-	86.6	-	-	89.0	-
	STST [73]	128	-	-	91.9	-	-	96.8	-	-	-	-	-	-	-
	FG-STFormer [13]	128	-	-	92.6	-	-	96.7	-	-	89.0	-	-	90.6	97.0
	Hyperformer [76]	64	90.7	-	92.9	95.1	-	96.5	86.6	-	89.9	88.0	-	91.3	96.7
	SkateFormer	64	92.6	93.0	93.5	97.0	97.4	97.8	87.7	89.4	89.8	89.3	91.0	91.4	98.3

of (0.9, 0.999), a weight decay of 5×10^{-4} . Additionally, gradient clipping [70] was applied for the loss values with gradients exceeding 1. The batch size was set to 128, the random seed was fixed at 1, and we adopted the label-smoothed cross-entropy loss [33, 52] as our loss function with parameter $\alpha = 0.1$.

From various experiments, we found the empirical values as follows. For our experiments with the NTU RGB+D and NTU RGB+D 120 datasets, we used the following configuration: $V = 48$ (excluding the center of the body, resulting in 24 joints for each of the two individuals), $T = 64$, $L = 4$, $K = 12$, $M = 8$, $N = 8$, and $C = 96$. For the experiments with the NW-UCLA dataset, $V = 20$, $T = 64$, $L = 4$, $K = 5$, $M = 8$, $N = 8$, and $C = 96$.

4.3 Performance Comparison

We present a comprehensive performance comparison for our SkateFormer against recent state-of-the-art skeleton-based action recognition methods. The performance comparison is made for three ensembles of different modalities (\mathbb{E}_1 , \mathbb{E}_2 , \mathbb{E}_4): (i) \mathbb{E}_1 - *joint* modality only; (ii) \mathbb{E}_2 - *joint + bone* modalities; and (iii) \mathbb{E}_4 - *joint + bone + joint motion + bone motion* modalities. As in [3, 7, 75, 76], we train separate networks for each modality and ensemble their outputs. Table 1 shows the overall performance comparisons of various methods for the NTU RGB+D, NTU RGB+D 120 and NW-UCLA datasets under the three ensembles. Notably, several works address human interaction recognition [11, 34, 62], a sub-part of skeleton-based action recognition, specifically focusing on scenarios where two or more individuals coexist within a single action. Accordingly, we

Table 2: Comparison with human interaction recognition methods on the NTU-Inter and NTU-Inter 120 datasets. The **best performances** are highlighted in **bold**.

Types	Methods	NTU-Inter (\mathbb{E}_1 , %)		NTU-Inter 120 (\mathbb{E}_1 , %)		Params. (M)	FLOPs (G)	Time (ms)
		X-Sub60	X-View60	X-Sub120	X-Set120			
Transformer	IGFormer [34]	93.6	96.5	85.4	86.5	-	-	-
	SkeleTR [11]	94.9	97.7	87.8	88.3	3.82	7.30	-
	ISTA-Net [62]	-	-	90.6	91.7	6.22	68.18	21.71
	SkateFormer	97.1	99.3	92.3	93.2	2.02	3.62	11.25

Table 3: Comparative analysis of SkateFormer with other methods by parameters, FLOPs, inference time, and average top-1 accuracy for joint modality.

Types	Methods	Params. (M)	FLOPs (G)	Time (ms)	NTU RGB+D (\mathbb{E}_1 , %)	NTU RGB+D 120 (\mathbb{E}_1 , %)
GCN	InfoGCN [7]	1.56	3.34	12.97	-	85.7
	FR-Head [75]	1.45	3.60	18.49	92.8	86.4
	Koopman [61]	5.38	8.76	17.86	92.7	86.6
	LST [63]	2.10	3.60	18.85	92.9	86.3
	HD-GCN [22]	1.66	3.44	72.81	93.2	86.5
Transformer	DSTA-Net [44]	3.45	16.18	13.80	-	-
	Hyperformer [76]	2.71	9.64	18.07	92.9	87.3
	SkateFormer	2.03	3.62	11.46	94.8	88.5

additionally present the performance of human interaction recognition methods on the NTU Inter and NTU Inter 120 datasets in Table 2.

$\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_4$ **performance.** As shown in Tables 1, 2, the extensive experiments demonstrate that our SkateFormer outperforms all the state-of-the-art (SoTA) methods except the \mathbb{E}_4 of X-Sub120. It should be noted that the performance of our SkateFormer is relatively higher for \mathbb{E}_1 and \mathbb{E}_2 than \mathbb{E}_4 . This is indicative of the model’s efficient attention strategy, which is based on partition-specific processing, allowing for the simultaneous handling of joint (input), bone (skeletal relations), joint motion (temporal relations), and bone motion (skeletal-temporal relations) inputs within a singular network framework. As the ensemble modality increases, the potential for information redundancy also rises, which may diminish the ensemble synergy. Therefore, our SkateFormer is already somewhat a strong learner in the form of a single modality compared to others.

Ensemble methods improve performance, but their effectiveness depends on the computational complexity, scaling proportionally with the number of models in the ensemble. Also, optimizing ensemble coefficients is dataset and model-specific [22], posing challenges in real-world applications. Consequently, it is crucial to leverage diverse modalities inherently within a single model to achieve both generalization and efficiency. Therefore, highlighting the importance of \mathbb{E}_1 performance is essential as a key indicator of the model’s ability to generalize across diverse inputs.

Computational complexity analysis. Table 3 shows the complexity comparisons [48] for various skeleton-based action recognition methods. Our SkateFormer exhibits a competitive balance between model complexity and computational efficiency, compared to other methods. It maintains a comparable number of parameters and FLOPs with the GCN-based methods, while substantially

reducing these metrics in comparison to the transformer-based counterparts. Table 3 has been compiled using publicly available official code.

4.4 Ablation Studies

To see the efficacy of the key components in our SkateFormer, the ablation experiments were carried for the NTU RGB+D dataset. A detailed analysis on the effectiveness of the components is provided in Tables 4, 5, 6 to quantify the impact of different design choices on the model’s performance in terms of accuracy (%) across two benchmarks: X-Sub60 and X-View60.

Skate-Types of Skate-MSA. Table 4 presents the impact of different Skate-Types on accuracy. As shown, incorporating the skeletal relation types only improves the action classification performance over the baseline, and so does the temporal relation types only. The full model that utilizes the Skate-Types (skeletal-temporal relation types) achieves the highest accuracy, indicating that both skeletal and temporal splits are crucial in distinguishing complex actions.

Table 4: Ablation study on the influence of skeletal and temporal relation types (Skate-Types) of Skate-MSA in our SkateFormer. (*: The batch size is set to 8 due to limited GPU memory.)

Attention Types	NTU RGB+D (%)		Params. (M)	FLOPs (G)	Time (ms)	Memory (MB)
	X-Sub60	X-View60				
Baseline (no attention)	90.7	95.7	2.03	3.59	10.68	80.6
+ Naive self-attention*	90.6 (↓0.1)	95.0 (↓0.7)	2.03	4.39	17.55	1612.1
+ $\mathbf{v}_k^{\text{nip}}$ + $\mathbf{v}_l^{\text{djp}}$	91.8 (↑1.1)	96.4 (↑0.7)	2.03	3.59	11.10	86.9
+ $\mathbf{t}_{ms}^{\text{local}}$ + $\mathbf{t}_n^{\text{global}}$	91.9 (↑1.2)	96.6 (↑0.9)	2.03	3.59	11.09	85.5
+ $\mathbf{v}_k^{\text{nip}}$ + $\mathbf{v}_l^{\text{djp}}$ + $\mathbf{t}_{ms}^{\text{local}}$ + $\mathbf{t}_n^{\text{global}}$	92.6 (↑1.9)	97.0 (↑1.3)	2.03	3.62	11.46	137.6

Exploration of Skate-Embedding. Table 5 shows the ablation study on the STE to assess the impact of different skeletal and temporal embedding methods. We explore various combinations of skeletal and temporal embedding methods. The learnable skeletal embedding (SE) paired with fixed temporal embedding (TE) achieves superior performance, suggesting an optimal balance between adaptability and stability in embeddings. The Skate-Embedding allows for a more tailored feature representation, leading to improved recognition accuracy.

Evaluation of frame sampling methods. In Table 6, we compare three frame sampling strategies with (i) fix-strided (*Fixed*), (ii) uniform random (*Uniform*) and (iii) our trimmed-uniform random (*Trimmed*) sampling methods. Due to the masking effect of skeleton sequences in their front and the back portions as well and more dense sampling effect in the middles during training, the trimmed-uniform random sampling is more effective in perspectives of generalization learning and data augmentation. It surpasses traditional fix-strided and uniform random sampling approaches with 1.2 (0.6)% and 1.0 (0.2)% margins for the X-Sub60 (X-View60), respectively.

Table 5: Ablation study results on the Skate-Embedding to assess skeletal and temporal embedding impacts.

Embedding Methods		NTU RGB+D (%)	
Skeletal	Temporal	X-Sub60	X-View60
\mathbf{x}	\mathbf{x}	91.9 ($\downarrow 0.7$)	96.6 ($\downarrow 0.4$)
	Learnable	91.8 ($\downarrow 0.8$)	96.5 ($\downarrow 0.5$)
	Fixed (TE)	91.6 ($\downarrow 1.0$)	96.6 ($\downarrow 0.4$)
Learnable (SE)	\mathbf{x}	92.1 ($\downarrow 0.5$)	96.5 ($\downarrow 0.5$)
	Learnable	91.4 ($\downarrow 1.2$)	96.2 ($\downarrow 0.8$)
	Fixed (TE)	92.6	97.0

Table 6: Ablation study results on the effect of intra-instance and inter-instance augmentations.

Intra-instance		Inter-instance	NTU RGB+D (%)	
Temporal	Skeletal		X-Sub60	X-View60
Trimmed			89.8	94.6
Trimmed		✓	90.3 ($\uparrow 0.5$)	94.3 ($\downarrow 0.3$)
Trimmed	✓		92.2 ($\uparrow 2.4$)	96.6 ($\uparrow 2.0$)
Fixed	✓	✓	91.4 ($\uparrow 1.6$)	96.4 ($\uparrow 1.8$)
Uniform	✓	✓	91.6 ($\uparrow 1.8$)	96.8 ($\uparrow 2.2$)
Trimmed	✓	✓	92.6 ($\uparrow 2.8$)	97.0 ($\uparrow 2.4$)

Impact of data augmentations. Table 6 presents the effectiveness of intra-instance (traditional) and inter-instance (additional) data augmentations. The intra-instance augmentation alone improves the performance while performance drop is observed with the inter-instance augmentation alone for the X-View60. This performance drop with 0.3%-point in accuracy is due to the inherent characteristics of the cross-view setting such that the variability in bone lengths across different subjects within the X-View60 is limited. This implies that the model may already encapsulate a comprehensive representation of these attributes. Consequently, the introduction of inter-instance data augmentation does not contribute additional discriminative information and may instead introduce redundancy, leading to a performance decline. However, the amalgamation of both intra- and inter-instance augmentations provides a synergistic enhancement, especially by random scaling and shearing on bone lengths in the skeletal augmentation, achieving peak accuracy of 92.6% on X-Sub60 and 97.0% on X-View60.

5 Conclusions

In this paper, we presented SkateFormer – a novel skeletal-temporal transformer tailored for action recognition tasks. For this, we propose an effective partition-specific attention strategy in skeleton-based action recognition, both to capture essential features and to reduce computational complexity. For efficient training, our novel Skate-Embedding that combines skeletal and temporal features is presented, significantly enhancing action recognition performance by forming an outer product between learnable skeletal features and fixed temporal index features. Our SkateFormer sets a new state-of-the-art for action recognition performance across multiple modalities (4-ensemble condition) and single modalities (joints, bones, joint motions, bone motions), showing notable improvement over the most recent state-of-the-art methods.

Acknowledgements. This work was supported by IITP grant funded by the Korea government(MSIT) (No.RS2022-00144444, Deep Learning Based Visual Representational Learning and Rendering of Static and Dynamic Scenes).

References

1. Cai, D., Kang, Y., Yao, A., Chen, Y.: Ske2grid: Skeleton-to-grid representation learning for action recognition. In: International Conference on Machine Learning (2023)
2. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7291–7299 (2017)
3. Chen, Y., Zhang, Z., Yuan, C., Li, B., Deng, Y., Hu, W.: Channel-wise topology refinement graph convolution for skeleton-based action recognition. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 13359–13368 (2021)
4. Chen, Z., Li, S., Yang, B., Li, Q., Liu, H.: Multi-scale spatial temporal graph convolutional network for skeleton-based action recognition. In: Proceedings of the AAAI conference on artificial intelligence. vol. 35, pp. 1113–1122 (2021)
5. Cheng, K., Zhang, Y., Cao, C., Shi, L., Cheng, J., Lu, H.: Decoupling gcn with drop-graph module for skeleton-based action recognition. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIV 16. pp. 536–553. Springer (2020)
6. Cheng, K., Zhang, Y., He, X., Chen, W., Cheng, J., Lu, H.: Skeleton-based action recognition with shift graph convolutional network. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 183–192 (2020)
7. Chi, H.g., Ha, M.H., Chi, S., Lee, S.W., Huang, Q., Ramani, K.: Infogcn: Representation learning for human skeleton-based action recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20186–20196 (2022)
8. Ding, Z., Wang, P., Ogunbona, P.O., Li, W.: Investigation of different skeleton features for cnn-based 3d action recognition. In: 2017 IEEE International conference on multimedia & expo workshops (ICMEW). pp. 617–622. IEEE (2017)
9. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
10. Duan, H., Wang, J., Chen, K., Lin, D.: Pyskl: Towards good practices for skeleton action recognition. In: Proceedings of the 30th ACM International Conference on Multimedia. pp. 7351–7354 (2022)
11. Duan, H., Xu, M., Shuai, B., Modolo, D., Tu, Z., Tighe, J., Bergamo, A.: Skeletr: Towards skeleton-based action recognition in the wild. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13634–13644 (2023)
12. Duan, H., Zhao, Y., Chen, K., Lin, D., Dai, B.: Revisiting skeleton-based action recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2969–2978 (2022)
13. Gao, Z., Wang, P., Lv, P., Jiang, X., Liu, Q., Wang, P., Xu, M., Li, W.: Focal and global spatial-temporal transformer for skeleton-based action recognition. In: Proceedings of the Asian Conference on Computer Vision. pp. 382–398 (2022)
14. Guo, T., Liu, H., Chen, Z., Liu, M., Wang, T., Ding, R.: Contrastive learning from extremely augmented skeleton sequences for self-supervised action recognition. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 762–770 (2022)

15. Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: Flownet 2.0: Evolution of optical flow estimation with deep networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2462–2470 (2017)
16. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017)
17. Ke, Q., Bennamoun, M., An, S., Sohel, F., Boussaid, F.: Learning clip representations for skeleton-based 3d action recognition. *IEEE Transactions on Image Processing* **27**(6), 2842–2855 (2018)
18. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)
19. Korban, M., Li, X.: Ddgcnn: A dynamic directed graph convolutional network for action recognition. In: *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*. pp. 761–776. Springer (2020)
20. Kwon, T., Tekin, B., Stühmer, J., Bogo, F., Pollefeys, M.: H2o: Two hands manipulating objects for first person interaction recognition. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 10138–10148 (2021)
21. Lee, J., Lee, M., Cho, S., Woo, S., Jang, S., Lee, S.: Leveraging spatio-temporal dependency for skeleton-based action recognition. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 10255–10264 (2023)
22. Lee, J., Lee, M., Lee, D., Lee, S.: Hierarchically decomposed graph convolutional networks for skeleton-based action recognition. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 10444–10453 (2023)
23. Li, B., Dai, Y., Cheng, X., Chen, H., Lin, Y., He, M.: Skeleton based action recognition using translation-scale invariant image mapping and multi-scale deep cnn. In: *2017 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*. pp. 601–604. IEEE (2017)
24. Li, M., Chen, S., Chen, X., Zhang, Y., Wang, Y., Tian, Q.: Actional-structural graph convolutional networks for skeleton-based action recognition. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 3595–3603 (2019)
25. Li, S., Li, W., Cook, C., Zhu, C., Gao, Y.: Independently recurrent neural network (indrnn): Building a longer and deeper rnn. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 5457–5466 (2018)
26. Lin, L., Zhang, J., Liu, J.: Actionlet-dependent contrastive learning for unsupervised skeleton-based action recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2363–2372 (2023)
27. Liu, J., Shahroudy, A., Perez, M., Wang, G., Duan, L.Y., Kot, A.C.: Ntu rgb+ d 120: A large-scale benchmark for 3d human activity understanding. *IEEE transactions on pattern analysis and machine intelligence* **42**(10), 2684–2701 (2019)
28. Liu, J., Shahroudy, A., Xu, D., Wang, G.: Spatio-temporal lstm with trust gates for 3d human action recognition. In: *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*. pp. 816–833. Springer (2016)
29. Liu, J., Wang, G., Duan, L.Y., Abdiyeva, K., Kot, A.C.: Skeleton-based human action recognition with global context-aware attention lstm networks. *IEEE Transactions on Image Processing* **27**(4), 1586–1599 (2017)

30. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
31. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)
32. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
33. Müller, R., Kornblith, S., Hinton, G.E.: When does label smoothing help? *Advances in neural information processing systems* **32** (2019)
34. Pang, Y., Ke, Q., Rahmani, H., Bailey, J., Liu, J.: Igformer: Interaction graph transformer for skeleton-based human interaction recognition. In: *European Conference on Computer Vision*. pp. 605–622. Springer (2022)
35. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
36. Peng, W., Hong, X., Chen, H., Zhao, G.: Learning graph convolutional network for skeleton-based human action recognition by neural searching. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 34, pp. 2669–2676 (2020)
37. Perez, M., Liu, J., Kot, A.C.: Interaction relational network for mutual action recognition. *IEEE Transactions on Multimedia* **24**, 366–376 (2021)
38. Plizzari, C., Cannici, M., Matteucci, M.: Skeleton-based action recognition via spatial and temporal transformer networks. *Computer Vision and Image Understanding* **208**, 103219 (2021)
39. Presti, L.L., La Cascia, M.: 3d skeleton-based human action classification: A survey. *Pattern Recognition* **53**, 130–147 (2016)
40. Qiu, H., Hou, B., Ren, B., Zhang, X.: Spatio-temporal segments attention for skeleton-based action recognition. *Neurocomputing* **518**, 30–38 (2023)
41. Shahroudy, A., Liu, J., Ng, T.T., Wang, G.: Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1010–1019 (2016)
42. Shi, L., Zhang, Y., Cheng, J., Lu, H.: Skeleton-based action recognition with directed graph neural networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 7912–7921 (2019)
43. Shi, L., Zhang, Y., Cheng, J., Lu, H.: Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 12026–12035 (2019)
44. Shi, L., Zhang, Y., Cheng, J., Lu, H.: Decoupled spatial-temporal attention network for skeleton-based action-gesture recognition. In: *Proceedings of the Asian Conference on Computer Vision* (2020)
45. Si, C., Chen, W., Wang, W., Wang, L., Tan, T.: An attention enhanced graph convolutional lstm network for skeleton-based action recognition. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 1227–1236 (2019)
46. Song, Y.F., Zhang, Z., Shan, C., Wang, L.: Stronger, faster and more explainable: A graph convolutional baseline for skeleton-based action recognition. In: *proceedings of the 28th ACM international conference on multimedia*. pp. 1625–1633 (2020)
47. Song, Y.F., Zhang, Z., Shan, C., Wang, L.: Constructing stronger and faster baselines for skeleton-based action recognition. *IEEE transactions on pattern analysis and machine intelligence* **45**(2), 1474–1488 (2022)

48. Sovrasov, V.: ptflops: a flops counting tool for neural networks in pytorch framework (2018–2023), <https://github.com/sovrasov/flops-counter.pytorch>
49. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8934–8943 (2018)
50. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5693–5703 (2019)
51. Sun, Z., Ke, Q., Rahmani, H., Bennamoun, M., Wang, G., Liu, J.: Human action recognition from various data modalities: A review. *IEEE transactions on pattern analysis and machine intelligence* (2022)
52. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2818–2826 (2016)
53. Tu, Z., Talebi, H., Zhang, H., Yang, F., Milanfar, P., Bovik, A., Li, Y.: Maxim: Multi-axis mlp for image processing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5769–5780 (2022)
54. Tu, Z., Talebi, H., Zhang, H., Yang, F., Milanfar, P., Bovik, A., Li, Y.: Maxvit: Multi-axis vision transformer. In: European conference on computer vision. pp. 459–479. Springer (2022)
55. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
56. Wang, J., Nie, X., Xia, Y., Wu, Y., Zhu, S.C.: Cross-view action modeling, learning and recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2649–2656 (2014)
57. Wang, L., Huynh, D.Q., Koniusz, P.: A comparative review of recent kinect-based action recognition algorithms. *IEEE Transactions on Image Processing* **29**, 15–28 (2019)
58. Wang, L., Koniusz, P.: 3mformer: Multi-order multi-mode transformer for skeletal action recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5620–5631 (2023)
59. Wang, P., Li, Z., Hou, Y., Li, W.: Action recognition based on joint trajectory maps using convolutional neural networks. In: Proceedings of the 24th ACM international conference on Multimedia. pp. 102–106 (2016)
60. Wang, Q., Shi, S., He, J., Peng, J., Liu, T., Weng, R.: Iip-transformer: Intra-inter-part transformer for skeleton-based action recognition. In: 2023 IEEE International Conference on Big Data (BigData). pp. 936–945. IEEE (2023)
61. Wang, X., Xu, X., Mu, Y.: Neural koopman pooling: Control-inspired temporal dynamics encoding for skeleton-based action recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10597–10607 (2023)
62. Wen, Y., Tang, Z., Pang, Y., Ding, B., Liu, M.: Interactive spatiotemporal token attention network for skeleton-based general interactive action recognition. In: 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 7886–7892. IEEE (2023)
63. Xiang, W., Li, C., Zhou, Y., Wang, B., Zhang, L.: Generative action description prompts for skeleton-based action recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10276–10285 (2023)

64. Xin, W., Miao, Q., Liu, Y., Liu, R., Pun, C.M., Shi, C.: Skeleton mixformer: Multivariate topology representation for skeleton-based action recognition. In: Proceedings of the 31st ACM International Conference on Multimedia. pp. 2211–2220 (2023)
65. Xu, K., Ye, F., Zhong, Q., Xie, D.: Topology-aware convolutional neural network for efficient skeleton-based action recognition. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 2866–2874 (2022)
66. Xu, Y., Cheng, J., Wang, L., Xia, H., Liu, F., Tao, D.: Ensemble one-dimensional convolution neural networks for skeleton-based action recognition. *IEEE Signal Processing Letters* **25**(7), 1044–1048 (2018)
67. Yan, S., Xiong, Y., Lin, D.: Spatial temporal graph convolutional networks for skeleton-based action recognition. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
68. Ye, F., Pu, S., Zhong, Q., Li, C., Xie, D., Tang, H.: Dynamic gcn: Context-enriched topology learning for skeleton-based action recognition. In: Proceedings of the 28th ACM international conference on multimedia. pp. 55–63 (2020)
69. Zeng, A., Sun, X., Yang, L., Zhao, N., Liu, M., Xu, Q.: Learning skeletal graph neural networks for hard 3d pose estimation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 11436–11445 (2021)
70. Zhang, J., He, T., Sra, S., Jadbabaie, A.: Why gradient clipping accelerates training: A theoretical justification for adaptivity. *arXiv preprint arXiv:1905.11881* (2019)
71. Zhang, P., Lan, C., Xing, J., Zeng, W., Xue, J., Zheng, N.: View adaptive recurrent neural networks for high performance human action recognition from skeleton data. In: Proceedings of the IEEE international conference on computer vision. pp. 2117–2126 (2017)
72. Zhang, P., Lan, C., Zeng, W., Xing, J., Xue, J., Zheng, N.: Semantics-guided neural networks for efficient skeleton-based human action recognition. In: proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1112–1121 (2020)
73. Zhang, Y., Wu, B., Li, W., Duan, L., Gan, C.: Stst: Spatial-temporal specialized transformer for skeleton-based action recognition. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 3229–3237 (2021)
74. Zhang, Z.: Microsoft kinect sensor and its effect. *IEEE multimedia* **19**(2), 4–10 (2012)
75. Zhou, H., Liu, Q., Wang, Y.: Learning discriminative representations for skeleton based action recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10608–10617 (2023)
76. Zhou, Y., Li, C., Cheng, Z.Q., Geng, Y., Xie, X., Keuper, M.: Hypergraph transformer for skeleton-based action recognition. *arXiv preprint arXiv:2211.09590* (2022)
77. Zhu, W., Lan, C., Xing, J., Zeng, W., Li, Y., Shen, L., Xie, X.: Co-occurrence feature learning for skeleton based action recognition using regularized deep lstm networks. In: Proceedings of the AAAI conference on artificial intelligence. vol. 30 (2016)