

# SkeleTR: Towards Skeleton-based Action Recognition in the Wild

Haodong Duan<sup>1\*</sup>    Mingze Xu<sup>2</sup>    Bing Shuai<sup>2</sup>    Davide Modolo<sup>2</sup>  
Zhuowen Tu<sup>2</sup>    Joseph Tighe<sup>2</sup>    Alessandro Bergamo<sup>2</sup>  
<sup>1</sup>The Chinese University of Hong Kong    <sup>2</sup>AWS AI Labs

dhd.efz@gmail.com, {xumingze, bshuai, dmodolo, ztu, tighej, bergamo}@amazon.com

## Abstract

We present *SkeleTR*, a new framework for skeleton-based action recognition. In contrast to prior work, which focuses mainly on controlled environments, we target more general scenarios that typically involve a variable number of people and various forms of interaction between people. *SkeleTR* works with a two-stage paradigm. It first models the intra-person skeleton dynamics for each skeleton sequence with graph convolutions, and then uses stacked Transformer encoders to capture person interactions that are important for action recognition in general scenarios. To mitigate the negative impact of inaccurate skeleton associations, *SkeleTR* takes relative short skeleton sequences as input and increases the number of sequences. As a unified solution, *SkeleTR* can be directly applied to multiple skeleton-based action tasks, including video-level action classification, instance-level action detection, and group-level activity recognition. It also enables transfer learning and joint training across different action tasks and datasets, which result in performance improvement. When evaluated on various skeleton-based action recognition benchmarks, *SkeleTR* achieves the state-of-the-art performance.

## 1. Introduction

Skeleton-based action recognition has received increasing attention in recent years due to its robustness against background and illumination changes [64, 16]. Existing methods [36, 7, 8, 15] mostly use Graph Convolutional Networks (GCNs) as backbones and focus on modeling spatio-temporal information from long skeleton sequences separately. While observing promising results in simplified and controlled environments [46, 32], they are less effective in more realistic scenarios with a variable number of people performing diverse actions. First, considering that linking skeletons (*e.g.*, using skeletons’ confidence scores [64]) in a long time span may inevitably generate more association errors [12, 48], requiring long sequences as input can bring

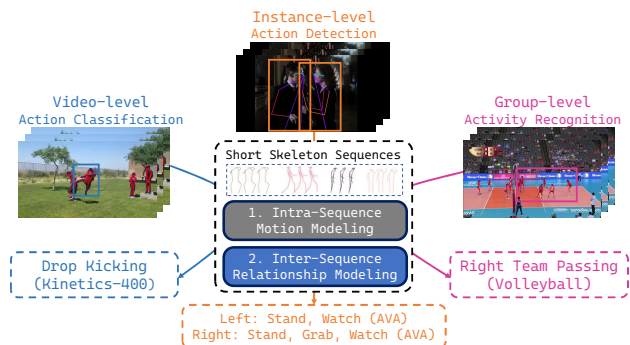


Figure 1: **SkeleTR is a general framework for skeleton action recognition which can handle different tasks**, including video-level action classification, instance-level action detection, and group-level activity recognition. It adopts a two-stage paradigm: the intra-sequence and inter-sequence modeling, respectively.

lots of noises. Second, prior works pay little attention to modeling the relations among a group of people, which plays an important role in recognizing both person interactions (*e.g.*, talking to someone [22]) and group activities (*e.g.*, pass/win-point in a volleyball game [24]).

In this paper, we propose *SkeleTR* for skeleton action recognition in real-world scenarios and handle multiple action tasks (*i.e.*, video-level, instance-level, and group-level action recognition) in a general framework. *SkeleTR* samples a large number of short skeleton sequences instead of a few long ones (practice of prior works [64, 16]) from each video. Each short sequence directly corresponds to an atomic action of a single person, thus the collection of multiple short sequences in space and time can represent hierarchical actions (with a longer time span) of a single person or a video-level action involving multiple people. The new input paradigm can be applied to instance-level and video-level recognition of actions with different temporal granularities, and it has several unique advantages. First, association errors are less likely to happen in a shorter skeleton sequence [12, 48]. We show that the improved sequence consistency leads to better intra-sequence modeling. Second, this new format is more flexible and helps to create comprehensive yet compact inputs. With a finer temporal

\*The work was done during an Amazon internship.

granularity, the same input size budget can cover more human skeletons. It reduces redundant zero-paddings for a small number of people, and mitigates information loss during input sampling for a large number of people.

SkeleTR adopts a two-stage paradigm as shown in Fig 1. Stage 1 performs intra-sequence modeling to capture the motion patterns of each individual skeleton sequence with a GCN backbone. The GCN backbone outputs a spatio-temporal feature for each sequence. Stage 2 concatenates all sequence features and feeds them into Transformer encoders to model inter-sequence relationships. At the end of stage 2, average pooling aggregates all features belonging to the same sequence into a single feature vector. By leveraging inter-sequence modeling, the output sequence feature incorporates the interactions between people and the global context of the video, which is helpful in both instance-level and video-level action recognition. To reduce the dimension of spatio-temporal skeleton features and keep the computations of Transformer tractable, we insert a novel Mix Pooling module before stage 2. It generates a compact representation with fine granularity for each sequence by applying multiple partial dimension reduction strategies in parallel.

SkeleTR easily enables transfer learning and joint learning across different datasets and tasks. With a unified input format and model architecture, one can jointly train multiple tasks with different datasets, each with its own loss function and ground truth annotations. This allows for leveraging auxiliary datasets for pre-training or joint training of a target task with limited labeled data. In experiments, we find that joint training serves as a strong regularizer and improves recognition performance, especially in scenarios with limited training data.

To evaluate the effectiveness of SkeleTR, we conduct comprehensive study on three skeleton-based action recognition tasks: action classification [6, 65, 46], spatio-temporal action detection [22], and group activity recognition [24]. Experimental results show that SkeleTR works well with different GCN backbones and significantly outperforms prior works on all three tasks. For example, with ST-GCN++ [15] backbone, SkeleTR outperforms the state-of-the-art methods by 3.8%, 3.8%, 0.7%, and 1.1% on Posetics [64], AVA [22], Volleyball [24], and NTU-Inter [40]. With Posetics joint training, the AVA performance is further improved by 4% mAP. The great capability of SkeleTR leads to another noteworthy result: we first observe the strong complementarity of skeleton action recognition. Combining SkeleTR prediction and the RGB state-of-the-art [2, 39] with simple late fusion leads to great performance boost (2.3% Posetics Top-1, 3% AVA mAP).

## 2. Related Works

**Human Action Understanding.** Action classification [6, 21] is the most fundamental task in human action under-

standing, aiming at recognizing the video-level action in a short video clip. Another task of great interest is recognizing the instance/person-level actions in a video, either at the atomic [22, 28] or hierarchical [30] level. Meanwhile, other areas of interest include group activity recognition [24, 9], action localization [17, 25, 70], or action segmentation [27, 50]. Various modalities in videos can be utilized to solve the action understanding tasks, such as RGB appearance [58, 6, 19, 59], optical flow [49], audio waves [61], and human skeletons [64, 16]. Spatio-temporal representations learned end-to-end from RGB appearance yield good performance on standard benchmarks. However, they are data intensive and less robust to context changes. In this work, we propose a unified framework to learn video representations based on human skeletons, that can be easily applied to multiple action tasks.

**Skeleton-based Action Recognition.** Human skeletons are robust against background and illumination changes [64, 47], and can be easily obtained with sensors [69] or pose estimators [4, 51]. In skeleton action recognition, the majority of research has focused on modeling the intra-person skeleton dynamics. With the rise of deep learning, manual feature engineering [23, 55, 57] has gradually been replaced by end-to-end representation learning with deep neural architectures (GCN [64], 2D-CNN [10], or 3D-CNN [16]). Among them, GCN-based approaches [64, 47, 36, 7] have been the most popular paradigm, which construct spatial-temporal graphs of human keypoint coordinates and perform spatial and temporal modeling. These methods model the intra-person skeleton dynamics for in-lab collected videos [46, 11], featuring small number of persons (1 or 2) per video and accurate person identities. For more challenging scenarios [6], the model performance decreases due to the noises produced in matching skeletons to sequences. Besides, GCN approaches primarily focus on predicting video-level actions, and do not benefit from context or interaction modeling when applied to instance-level action recognition. To this end, we propose SkeleTR, which expands existing GCN approaches and works universally well for diversified scenarios and recognition tasks.

**Human Interaction Recognition.** Human interaction recognition [43, 53, 67] is a sub-task of action recognition. It aims at identifying and categorizing activities between people. Early works [66, 26] built hand-crafted features for interaction modeling and adopted classic machine learning algorithms for activity classification. Subsequent works proposed to use two-stream RNNs to model the two-person interactions [41, 37, 56]. Recently, IGFormer [40] first proposed to use a two-stream Transformer for two-person interaction modeling and observed improved performance. However, most deep methods share the same limitations: 1. they are largely restricted to two-person scenarios and tend to suffer from exponentially growing inference

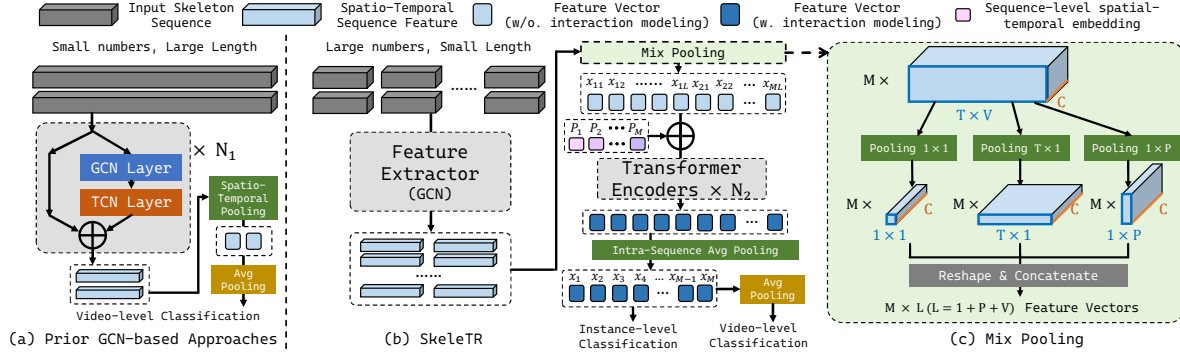


Figure 2: **The architecture of SkeleTR.** SkeleTR takes a large set of short skeleton sequences as input and performs a two-stage modeling process. It first adopts a GCN backbone to model the intra-sequence skeleton dynamics. Then it processes the sequence features with stacked Transformer encoders, which helps to model person interactions and leverage the long-range context. Between the two stages, we adopt a Mix Pooling layer to compress the spatio-temporal features and keep computations tractable.

costs with the person number; 2. they fail to leverage the strong intra-sequence modeling capabilities of GCNs. To overcome these challenges, we introduce SkeleTR, which integrates intra-sequence and inter-sequence modeling with good scalability and small computation cost.

**Visual Transformer.** With the great success in NLP [54, 13], Transformers have been increasingly applied to diversified vision tasks [14, 34, 5, 74] as well as the RGB-based action understanding [35, 18, 1, 2, 63, 71], and have achieved promising results. For skeleton action recognition, there exist works [42, 68] that directly use Transformers for intra-sequence spatio-temporal modeling. However, without exploiting the locality, the Transformer-based motion modeling only leads to inferior recognition performance compared to GCN-based ones. In SkeleTR, to enjoy the best of both worlds, we adopt GCNs for intra-sequence motion modeling, and use Transformer encoders for inter-sequence relationship modeling among multiple persons’ fine-grained skeleton features.

### 3. SkeleTR

#### 3.1. Overview

Both SkeleTR and prior GCN approaches directly take human joint coordinates as input. A person in a specific frame corresponds to a skeleton with  $V$  joints (each has  $C$  dimensions,  $C$  is 2 or 3). Skeletons in  $T$  different frames (1 skeleton per frame) can be linked in temporal order to form a skeleton sequence  $\mathbf{X}_i \in \mathbb{R}^{T \times V \times C}$  with length  $T$ . A skeleton action recognizer takes  $M$  skeleton sequences as input, and outputs a score vector  $\mathbf{S} \in \mathbb{R}^K$  for video-level action recognition, where  $K$  is the number of action classes.

Spatial-temporal GCNs are widely used to model skeleton sequences for action classification, which focus on modeling intra-sequence skeleton dynamics. Most GCN methods [64, 47, 7] share a similar high-level design (Fig 2(a)), where a backbone with  $N_1$  GCN blocks is used to extract features of each skeleton sequence. A GCN block com-

prises a GCN layer and a TCN layer. The GCN layer performs spatial modeling by fusing joint features within the same skeleton. The TCN layer employs 1D Convs to capture the motion patterns of each joint. GCNs model long skeleton sequences that span the entire video in parallel, and aggregate multiple sequence features with average pooling to predict the video-level action. On existing skeleton action recognition benchmarks, GCNs exhibit superior intra-sequence modeling capability than other architectures (*e.g.*, RNN [73, 33] or Transformers [42, 68]).

SkeleTR inherits the effective intra-sequence modeling of GCNs, while also enhancing the flexibility, generality, and inter-person modeling capability. SkeleTR takes  $M$  short skeleton sequences as input (Sec. 3.2), and outputs one feature vector per sequence. It first adopts a GCN backbone to model the motion patterns of each skeleton sequence, then uses stacked Transformer encoders to capture the interactions and leverage long-range context (Sec. 3.3). Between two stages, we use a Mix Pooling layer to compress spatio-temporal sequence features into a compact multi-view representation to keep computation tractable. Fig 2 compares the design of SkeleTR and previous GCN methods.

#### 3.2. The Input Construction of SkeleTR

For skeleton action classification in real-world scenarios, prior works [64, 47, 36] use a few long skeleton sequences as input, where skeletons at different timestamps are linked with rudimentary association methods. They use a naive confidence score based matching to link skeletons to sequences as long as the entire video (as shown in Fig 3). Specifically, they mostly keep  $\hat{M} = 2$  skeletons with the highest confidence scores per frame, and link skeletons based on confidence rank. In cases where it is applied to challenging videos involving multiple persons, this strategy can lead to noisy skeleton sequences and sub-optimal performance.

In contrast, we propose to sample a large number of short skeleton sequences as input, and link skeletons at differ-

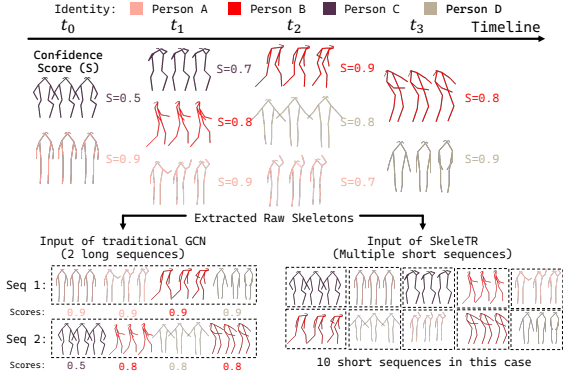


Figure 3: **The difference of input formats.** Prior works link skeletons at different timestamps based on the confidence score rank to form  $\hat{M} = 2$  long sequences, each covers the entire video. SkeleTR, instead, directly takes multiple shorter sequences as input. Each sequence only covers 1~2 seconds, thus has better consistency than long sequences.

ent timestamps into sequences based on their IOU similarity. Specifically, to sample sequences with length  $T$ , we first uniformly select a set of key frames with a stride of  $T/2$ . For each skeleton in every key frame, we create a sequence by linking skeletons within a window of  $T$  frames using IoU-based matching [3], resulting in a total of  $M'$  sequences ( $M'$  is variable for different videos). Our new input format is compatible with a variable number of input sequences. By using different number of sequences ( $M$ ) for each training batch, SkeleTR can handle inputs with different  $M$ , which helps to reduce the usages of zero-padding at test time and improves efficiency.

### 3.3. Inter-sequence Modeling Stage

After processing short skeleton sequences with the GCN backbone, SkeleTR employs stacked Transformer encoders to model inter-sequence relationships. The process starts with reducing the feature dimension using a Mix Pooling layer, resulting in  $L$  feature vectors for each sequence. Next, a sequence-level spatial-temporal embedding is combined with the feature vectors. The resulting sequence features ( $M \times L$  vectors in total) are concatenated and modeled using a standard Transformer with multiple encoder blocks [14]. This architecture is depicted in Fig 2(b). We present more details in the next paragraphs.

**Mix Pooling.** The cost of computing the attention matrices in a Transformer block grows quadratically with the input dimensionality, resulting in a computational cost of  $\Theta((M \times T \times V)^2)$  in our case, if no dimension reduction is performed ( $L = T \times V$ ). Given a video with a typical configuration ( $T = 15, V = 17$ ) and a Transformer consisting of only  $N_2 = 2$  encoder blocks, the Transformer itself consumes 52.7 GFLOPs, which is impractical to compute. Another extreme is to remove the dimension  $T$  and  $V$  with global average pooling, which reduces the computational complexity to  $\Theta(M^2)$  but also results in loss of fine-grained

features for each skeleton sequence.

Between two extremes ( $L = T \times V$  or  $L = 1 \times 1$ ), we adopt multi-stream partial dimension reductions to reduce the computation cost. We denote a partial dimension reduction stream as  $A \times B$ , indicating reducing temporal dimension ( $T$ ) to  $A$  and spatial dimension ( $V$ ) to  $B$ . Each stream provides a specific view of the spatio-temporal skeleton features. Multiple reduction streams can be combined in Mix Pooling to compress raw features into a comprehensive representation. Specifically, we adopt 3 partial dimension reduction streams:  $1 \times 1, T \times 1$ , and  $1 \times P$  ( $V$  joints aggregated to  $P = 5$  parts: head, two arms, two legs) (Fig 2(c)). Mix Pooling concatenates the output of 3 streams and keeps  $L = 1 + T + P$  features per sequence. The heterogeneous features are then processed with one single Transformer. Mix Pooling enables the cross-view feature interaction and helps keep computations tractable, yielding a complexity of  $\Theta(M^2(1 + T + P)^2)$  ( $\sim 0.6$  GFLOPs / video).

**Sequence-level Spatial Temporal Embedding.** We further combine  $L$  features per sequence with a sequence-level spatial-temporal embedding. We utilize a fully-connected layer to encode both the bounding box and the normalized timestamp (computed based on the entire video) of the center skeleton for each sequence, which then serves as the embedding. The embeddings are added with the Mix-Pooled features, and further processed by a layer normalization.

**Transformer-based Interaction Modeling.** We use a standard Transformer with  $N_2$  encoder blocks for modeling sequence features. Each block consists of a multi-head self-attention layer (MHSA) and a feed-forward network (FFN), and adopts the post-layernorm style. The multi-head design in the encoder blocks helps to capture various kinds of relationships among skeleton features, including but not limited to the inter-person interactions and intra-person long-range context. At last, we use intra-sequence average pooling to aggregate  $L = (1 + T + P)$  features per sequence into a single one. The output sequence features, which incorporate both person interactions and a comprehensive context, are then used to predict the instance-level or video-level action. We present the detailed configuration of encoder blocks in the appendix.

### 3.4. SkeleTR as a General Framework

As a general framework for skeleton action recognition, SkeleTR can be applied to various skeleton-based action recognition tasks. It also enables and benefits from transfer learning and joint learning across different tasks.

**Video-level Action Classification.** The tasks of action classification [6, 65] and group activity recognition [24, 9] both take as input a video and predict the action at video-level. Note that the input videos may contain a variable number of people. In SkeleTR, we sample numerous short skeleton sequences that last 1 ~ 2 seconds. We then average the  $M$




Figure 4: **The curve of training / validation loss and validation mAP on AVA.** Joint training with Posetics [65] significantly alleviates overfitting for AVA action detection. The jointly-trained model achieves much higher validation mAP, with the gap of training / validation losses largely shrinks.

output features, and use a fully-connected layer to predict the video action class. Thanks to the local motion modeling with GCN and interaction modeling with Transformers, SkeleTR achieves strong performance in both action classification and group activity recognition.

**Instance-level Action Recognition.** Spatio-temporal action detection [22] aims to recognize actions of each person at specific timestamps (instance-level recognition). For this task, we sample short skeleton sequences with length  $T^1$  in a longer temporal window (5 ~ 10s) as input, and utilize the sequence-level features for instance-level action recognition. SkeleTR enables the modeling of both intra-sequence motion patterns and inter-sequence interactions for skeleton-based spatio-temporal action detection.

**Transfer Learning and Joint Learning.** SkeleTR enables to: 1. transfer the representation learned on one task to another new task; 2. jointly train different datasets and tasks, which in turn alleviate the overfitting issues on small datasets. Fig 4 shows the curves of training and validation losses, and validation mAP for AVA. We see that joint training serves as a strong regularizer and alleviates the severe overfitting. With joint training, the gap between training and validation losses largely shrinks, and the recognition performance on the validation set becomes much stronger.

## 4. Experiments

### 4.1. Datasets

**Kinetics-400 (K400)** [6] contains videos with variable number of people. Some K400 actions cannot be distinguished by only using skeletons, thus we use a subset of K400, named **Posetics** [65], as the main benchmark for action classification. Posetics merges together classes featuring identical skeletons, resulting in 320 total classes (e.g., ‘eating burger’ and ‘eating hotdog’ are merged into ‘eating something’). For K400 and Posetics, we use 2D human skeletons provided by [16] (extracted with Faster-

<sup>1</sup> $T$  is restricted by annotation granularity. For example, AVA [22] annotates atomic actions at 1fps, for which  $T$  should cover at most 2 seconds.

RCNN [44] (R50) and HRNet [51]). We provide K400 results in the appendix.

**NTU-Inter** is a commonly used benchmark for two-person interaction recognition [40, 41]. It’s a subset of the in-lab collected NTURGB+D [32, 46], containing only mutual action classes (with two subjects). Consistent with NTURGB+D, NTU-Inter includes 4 data splits: NTU60-XSub, NTU60-XView, NTU120-XSub, NTU120-XSet. We use the official 3D skeletons for NTU-Inter experiments.

**AVA 2.2** [22] is a spatio-temporal action detection benchmark. It’s a multi-label dataset with 80 action classes, falling into three major categories: person movement, object manipulation, and person interaction. Following the common practice [60], we use Faster-RCNN (ResNeXt-101-FPN [31, 62]) for person detection and HRNet for pose estimation. The detector is pretrained on COCO and finetuned on AVA training bboxes. Following the benchmark guidelines, we evaluate 60 action classes. We report the overall mAP and the mAP of each major category.

**Volleyball** [24] is a group activity recognition dataset with 8 volleyball group activities. Each clip has up to 13 persons, with the center frame annotated with GT person boxes. Following common practice [72, 52], we use the tracklets from [45] and HRNet for pose extraction, train models on training and validation sets, and report the test set results.

### 4.2. Implementation Details

**Architecture.** Unless otherwise noted, we adopt ST-GCN++ [15] as the GCN backbone. Particularly, we consider two variants: **SkeleTR-S** (with fewer GCN blocks and smaller output dimension, details in appendix), and **SkeleTR-L** (original ST-GCN++). The inter-sequence modeling stage consists of a Mix Pooling module and two Transformer encoder blocks. The feature dimension in each encoder blocks equals to the output channels of the GCN backbone (128 for SkeleTR-S and 256 for SkeleTR-L).

**Training.** We train all models end-to-end with SGD and batch size 128. For all models, we set the base learning rate (LR) to 0.1 and adopt a multi-step LR scheduler with the decay factor 0.1. For Posetics and NTU-Inter, we train models for 100 epochs, and set milestones at the 70<sub>th</sub> and 85<sub>th</sub> epoch. For AVA and Volleyball, we train the models for 40 epochs, and set milestones at the 30<sub>th</sub> and 35<sub>th</sub> epoch. We use cross entropy loss ( $L_{CE}$ ) for Posetics, NTU-Inter and Volleyball, and use binary cross entropy loss ( $L_{BCE}$ ) for AVA (multi-label). We explicitly set a loss weight  $\lambda = 100$  in all AVA experiments due to the small magnitude of  $L_{BCE}$ . In Posetics-AVA joint training, we concatenate samples in two datasets and optimize the model with a weighted combination of two losses:  $L = L_{CE} + \lambda L_{BCE}$ . For AVA, we sample short skeleton sequences in a small temporal window (5s / 10s) to form a training sample.

**Inference.** We follow standard practices [15] to predict ac-

Methods	Posetics Top-1	Posetics Top-5	Methods	NTU60 XSub	NTU60 XView	NTU120 XSub	NTU120 XSet	Methods	Overall mAP	Person Movement	Object Manipulation	Person Interaction	Methods	Volleyball Top-1
ST-TR [42]	47.5	71.3	ST-LSTM [33]	83.0	87.3	63.0	66.6	ST-TR [42]	17.3	37.3	8.4	18.8	AT+ [20]	92.3
MS-G3D [36]	52.6	75.8	LSTM-IRN [41]	90.5	93.5	77.7	79.6	MS-G3D [36]	18.4	38.7	9.6	19.7	POGARS [52]	93.2
CTR-GCN [7]	54.1	76.7	GeomNet [38]	93.6	96.3	86.5	87.6	CTR-GCN [7]	18.9	39.4	9.8	20.5	COMPOSER [72]	93.7
ST-GCN++ [15]	52.5	76.0	IGFormer [40]	93.6	96.5	85.4	86.5	ST-GCN++ [15]	18.6	39.1	9.7	20.0	ST-TR [42]	93.3
UNIK [65]	52.5	75.7	AS-GCN [29]	87.1	92.0	77.8	79.3	UNIK [65]	19.0	39.7	9.9	20.4	CTR-GCN [7]	93.6
PoseC3D [16]	53.1	77.1	CTR-GCN [7]	91.6	94.3	83.2	84.4	PoseC3D [16]	18.3	39.3	9.3	19.1	ST-GCN++ [15]	93.3
SkeleTR-S	55.9	78.9	SkeleTR-S	94.6	97.1	87.4	88.0	SkeleTR-S	22.8	43.8	12.3	26.9	SkeleTR-S	<b>94.4</b>
SkeleTR-L	<b>57.9</b>	<b>80.3</b>	SkeleTR-L	94.7	97.5	<b>87.8</b>	<b>88.3</b>	SkeleTR-S*	26.8	48.2	<b>17.7</b>	27.8	SkeleTR [C]	<b>94.4</b>
SkeleTR [C]	57.6	80.1	SkeleTR [C]	<b>94.8</b>	<b>97.7</b>	<b>87.8</b>	<b>88.3</b>	SkeleTR [C]*	<b>27.3</b>	<b>50.9</b>	17.3	<b>28.3</b>		

(a) Action Classification: Posetics.

(b) Action Recognition: NTU-Inter.

(c) Spatio-temporal Action Detection: AVA v2.2.

(d) Group Activity Recognition: Volleyball.

Table 1: **Comparing SkeleTR with previous state-of-the-art on versatile skeleton action recognition tasks.** [C] denotes CTR-GCN backbone; \* denotes joint training with Posetics. For Posetics and AVA, we follow the common practice [7, 65, 16] and compare the two-stream<sup>2</sup> performance. For NTU-Inter and Volleyball, we only compare the joint modality performance as done in the prior works [52, 72]. In Table 1d, the method AT<sup>+</sup> uses groundtruth bboxes of all video frames for skeleton extraction.

tion labels for trimmed videos. For AVA untrimmed videos, in inference we use a sliding window strategy to generate predictions, with a 1-second time step. We get an action prediction for a human proposal from each window containing the proposal, and then average all predictions associated with a human proposal to get the final result.

### 4.3. Main Results

We compare SkeleTR with state-of-the-art methods on three different skeleton action understanding tasks: action classification, spatio-temporal action detection, and group activity recognition. Unless otherwise noted, all methods use the same skeleton input as SkeleTR.

**Action Classification.** Table 1a compares SkeleTR with state-of-the-art methods on Posetics. We use official training configurations when available [36, 16, 65]. For methods without official configurations [7, 15], we use the same hyper-parameter setting as SkeleTR. SkeleTR with all GCN backbones notably outperforms existing methods. Compared to CTR-GCN, SkeleTR-L achieves 3.6% higher Top-1 accuracy. Quantitative analysis shows that the improvement is more significant for multi-person scenarios: actions with  $\geq 5$  persons per frame have around 10% improvement in mean class Top-1, which is twice as much as the improvement for other actions. SkeleTR also achieves the highest accuracy on Posetics (with OpenPose skeletons) and K400 (with HRNet/OpenPose skeletons).<sup>3</sup> Table 1b shows that on the more specific task of two-person interaction recognition, SkeleTR also outperforms existing interaction recognition methods [41, 38, 40] and skeleton action classification methods by a large margin.

**Spatio-temporal Action Detection.** Skeleton-based spatio-temporal action detection is a new setting first proposed in this paper and we use AVA as the primary benchmark. For comparisons, we adapt state-of-the-art skeleton action classification algorithms by taking a single skeleton sequence (with 2 seconds length, centered at a person pro-

Methods	Posetics Top-1	Params (M)	GFLOPs	Inference Speed (2080Ti, clip/second)
ST-TR [42]	47.45	6.44	26.92	97
MS-G3D [36]	52.63	2.82	20.86	99
CTR-GCN [7]	54.10	1.36	5.36	195
ST-GCN++ [15]	52.46	1.31	5.44	235
UNIK [65]	52.52	3.28	12.15	156
PoseC3D [16]	53.09	1.90	14.69	53
SkeleTR-S	55.85	<b>0.81</b>	<b>2.74</b>	<b>336</b>
SkeleTR-L	<b>57.90</b>	3.82	7.30	200

Table 2: Efficiency statistics of models in Table 1a.

posal in the key frame) as input to predict the action for this person. We train all models with the same hyper-parameter setting as SkeleTR. Due to the lack of global context and interaction modeling, all existing methods perform poorly on the new task, while SkeleTR significantly outperforms them. Compared to the baseline ST-GCN++, SkeleTR-S improves the overall mAP by 4.1%. Among three major action categories, SkeleTR-S makes the largest improvement on person interaction (7.0%), thanks to its great inter-person modeling capability. When jointly trained with Posetics, SkeleTR-S enjoys another 4.1% overall mAP improvement. However, this time the improvement is basically for person movement and object manipulation, since AVA-style person interactions rarely appear in Kinetics videos.

**Group Activity Recognition.** Most prior works on skeleton-based group activity recognition do not utilize the effective GCN architecture to model intra-person motion patterns. Instead, they directly model the relationship among skeletons with complicated transformer architectures. For a comprehensive comparison, we also adapt the state-of-the-art skeleton action recognition algorithms and benchmark them on the Volleyball dataset. Table 1d shows that, with good capability of both intra-sequence motion modeling and inter-sequence relationship modeling, SkeleTR outperforms all previous state-of-the-art on Volleyball, including the adapted GCN-based approaches.

**Computational Efficiency.** SkeleTR is designed to be efficient and lightweight. Table 2 shows that SkeleTR-S outperforms previous state-of-the-art methods for skeleton-based action classification in terms of accuracy, while having the model with the smallest parameter size, GFLOPs, and fastest inference speed. In the appendix, we provide ad-

<sup>2</sup>Joint and Bone. Each stream takes a separate modality derived from skeleton coordinates as input. For joint stream, the input is the coordinates itself. For bone stream it is the difference of neighboring joints.

<sup>3</sup>Please see the appendix for these results.

Methods	Posetics		AVA v2.2			
	Top-1	Top-5	Overall mAP	Person Movement	Object Manipulation	Person Interaction
RGB	74.42	91.29	31.30	48.44	23.67	32.74
ST-GCN++	52.46	75.97	18.62	39.10	9.68	19.98
RGB + ST-GCN++	75.98	91.58	31.95	52.28	23.14	33.13
SkeleTR	57.90	80.33	26.83	48.19	17.72	27.75
RGB + SkeleTR	<b>76.70</b>	<b>92.02</b>	<b>34.34</b>	<b>54.78</b>	<b>25.77</b>	<b>34.93</b>

Table 3: **SkeleTR demonstrates great complementarity when fused with RGB recognizers.** We fuse predictions by state-of-the-art RGB recognizers (TimeSformer [2] for Posetics, ACAR-R101 [39] for AVA) with skeleton predictions with 2 : 1 ratio.

Interaction Modeling	Posetics Top-1	AVA Overall mAP
NA	48.21	17.56
Add Global	51.06	20.00
Concat Global	50.01	20.14
Non-Local	48.43	19.78
Transformer Encoders	<b>52.81</b>	<b>20.92</b>

Table 4: **Transformer-based interaction modeling significantly outperforms multiple baselines.**

Feature Aggregation	Posetics Top-1	AVA Overall mAP	Head Computation (GFLOPs)
$1 \times 1$	52.81	20.92	0.013
$T \times P$	52.78	21.50	5.211
$T \times 1$	53.48	20.64	0.356
$1 \times P$	53.40	21.35	0.080
$1 \times V$	52.15	21.01	0.429
$1 \times 1 + T \times 1$	53.75	21.00	0.392
$1 \times 1 + 1 \times P$	53.00	21.43	0.101
$T \times 1 + 1 \times P$	53.16	20.87	0.550
$1 \times 1 + T \times 1 + 1 \times P$	<b>53.90</b>	<b>21.65</b>	0.594

Table 5: **Comparing different feature aggregation strategies.**

ditional evidence of SkeleTR’s computational efficiency by comparing it to other methods designed for different tasks, such as IGFormer [40] and COMPOSER [72].

**Improved complementarity with RGB modality.** To demonstrate the complementarity of SkeleTR upon RGB methods, we directly fuse the prediction of SkeleTR and the state-of-the-art RGB-based action recognizers. Table 3 shows that late fusion with SkeleTR leads to a significant boost (2.3% Posetics top-1, 3% AVA mAP) compared to the RGB state-of-the-art, which notably outperforms ST-GCN++ late fusion. On AVA, the improvement is extremely large for person movement (6.3% mAP), which emphasizes the importance of the skeleton in modeling human motions.

#### 4.4. Ablation on Inter-Sequence Modeling

We conduct ablation experiments to validate the efficacy of our inter-sequence modeling stage. In experiments, we adopt the light ST-GCN++ as the backbone. We set  $M$  to 20,  $T$  to 30, and set the temporal window size to 5s for AVA. We report the recognition performance with a single joint modality on Posetics and AVA.

We first compare the Transformer-based modeling with other simple baselines for modeling global context or interaction. In all experiments, we pool both spatial and temporal dimensions and retain a single feature vector per sequence. We consider the following baselines: 1. **NA**: no




Figure 5: **Visualization of the weights in different attention heads in the last transformer encoder.**  $P_{ij}$  denotes the short skeleton sequence of person  $j$  at the timestamp  $i$ .

interaction modeling; 2. **Add Global**: process the average feature with an MLP to get a Global feature, then add it back to each sequence feature; 3. **Concat Global**: concatenate the average feature with each sequence feature, then process with an MLP; 4. **Non-Local**: process all sequence features with a standard non-local [59] block, and apply an MLP to the output. Table 4 shows that, despite its low computation cost (0.013 GFLOPs / clip), the Transformer markedly enhances recognition performance relatively to the **NA** baseline (4.6% Posetics top-1, 3.4% AVA mAP), and considerably outperforms all alternatives.

Both Non-Local and Transformer encoders adopt self-attention to model the inter-sequence relationship, while there is a significant difference in recognition performance between them. The reason is that the multi-head design in Transformers allows for modeling diverse relationships between individual sequences. Visualizing the attention weights of different attention heads, we find that some of them have clearly interpretable semantics. Fig 5 shows the attention weights of a ‘handshake’ action. Of the eight attention heads, the first head aggregates temporal features within a single person, while the fourth head models interactions between individuals. Specifically, the attention weights of the fourth head reveal that  $P_{21}$  and  $P_{22}$  attend most strongly to other skeleton features, which corresponds to the exact moment of the handshaking interaction.

We then delve deeper into the Mix Pooling design, which helps to keep fine-grained sequence features and serves as a good trade-off between efficiency and performance. Table 5 compares Mix Pooling ( $1 \times 1 + T \times 1 + 1 \times P$ ) with other feature aggregation strategies. Mix Pooling outperforms each individual pooling strategy or the combination of two strategies. It achieves the best performance on both datasets with a reasonable computation cost of 0.594 GFLOPs per clip.

#### 4.5. Ablation on Input Construction

In this section we demonstrate the advantages of our new input format, both in recognition performance and computational efficiency. We report the recognition performance with a single joint modality.

**Large Number of Short Sequences.** We use a large number of short skeleton sequences as input for skeleton action




Figure 6: **Fix the total number of skeletons and adjust the sequence number.** Posetics Top-1 accuracy is visualized.

Methods	$M$	$T$	Posetics Top-1	Avg GFLOPs for Posetics Video Clips
SkeleTR-S	20	30	53.90	2.430
SkeleTR-L	[5, 20]	30	<b>55.06</b>	4.934
SkeleTR-L	20	30	54.95	6.653
SkeleTR-S	40	15	54.17	2.789
SkeleTR-L	[10, 40]	15	55.35	3.532
SkeleTR-L	40	15	<b>55.55</b>	7.394

Table 6: **Sequence number  $M$ : fixed or adaptive?** Utilizing an adaptive sequence number  $M$  can save great amounts of computational resources and achieve comparable performance.

understanding tasks. Besides being a unified input format for different tasks, the new format also benefits standard skeleton-based action classification under real-world scenarios. To investigate the effect of sequence granularity, we fix the total number of skeletons  $M \times T$ , and adjust the sequence number  $M$ . We compare three models: ST-GCN++, SkeleTR-S, and SkeleTR-L (Fig 6). Interestingly, we find that all models obtain their highest performance when using a large number of short sequences as input, specifically,  $40 \times 15$ . Meanwhile, the improvement is especially large for SkeleTR: 2.1% top-1 for SkeleTR-S, 3% top-1 for SkeleTR-L (compared to the  $2 \times 300$  setting).

**Flexible Number of Sequences.** SkeleTR is compatible with different numbers of sequences. To train and test SkeleTR with an adaptive sequence number, we first set a dynamic range  $[M_{min}, M_{max}]$ , and for each training batch, we randomly sample  $\tilde{M}$  from the range and use it as the sequence number for all videos in the batch. For each testing video, we use the actual sequence number  $M'$  if  $M_{min} \leq M' \leq M_{max}$ , and cap it to  $M_{min}$  or  $M_{max}$  otherwise. We present an ablation study using two input formats:  $20 \times 30$  and  $40 \times 15$ . For each format, we use a dynamic range ( $[5, 20]$ ,  $[10, 40]$ ) of  $M$  to train SkeleTR-L on Posetics, and report the results in Table 6. Compared to using a fixed  $M$ , using an adaptive  $M$  can achieve comparable recognition performance with a largely reduced inference cost (26% less GFLOPs for  $T = 30$ , 52% less for  $T = 15$ ).

#### 4.6. Transfer Learning and Joint Learning

SkeleTR is a general framework for versatile skeleton action recognition tasks, which enables transfer learning and joint learning across different tasks. To demonstrate this,

Training	Temporal Window	Overall mAP	Person Movement mAP	Object Manipulation mAP	Person Interaction mAP
From Scratch	5	21.65	42.34	11.37	25.66
Finetuning	5	22.12	43.02	11.91	25.76
Joint Training	5	23.92	44.25	14.65	26.06
From Scratch	10	21.33	42.24	11.26	24.68
Finetuning	10	23.11	44.06	13.32	25.85
Joint Training	10	<b>25.69</b>	<b>47.13</b>	<b>16.53</b>	<b>26.64</b>

Table 7: **Transfer Learning and Joint Learning with SkeleTR-S.** For a target task with limited training data, an auxiliary dataset can be leveraged to improve the performance and mitigate the overfitting, either by transfer learning or joint learning.

we consider two tasks: Posetics action classification and AVA spatio-temporal action detection (with AVA being the smaller dataset). From-scratch training on AVA is easy to overfit (Fig 4). Pretraining on Posetics or joint training on AVA and Posetics both help to mitigate the overfitting, with joint training leading to the best recognition performance. Enlarging the temporal window size for AVA sampling, which leads to less diversified training samples, amplifies the gain of pretraining or joint training.

#### 4.7. Limitations

SkeleTR is currently trained with skeletons pre-extracted from RGB videos, and the adopted pose estimation solution is relatively heavy. To use SkeleTR in real-world apps with limited computation resource, additional efforts are required to switch to a light-weight pose estimation solution that can perform fast inference on CPUs or edge devices, and to retrain the models accordingly. Besides, the short sequence format proposed by SkeleTR strikes a good balance between mitigating skeleton matching noises and modeling intra-sequence dynamics in real-world scenarios. The format may not lead to extra benefits when the ground-truth skeleton sequences are available (details in appendix).

### 5. Conclusion

We presented SkeleTR, a unified framework for skeleton action recognition in general scenarios. It extends existing GCN backbones and achieves state-of-the-art performance on multiple skeleton action understanding tasks, including video-level, instance-level, and group-level action recognition. Besides, SkeleTR is compatible with transfer learning and joint learning across different datasets and tasks, which further improves the capability of learned representations.

**Broader Impact.** Using human skeletons for action recognition is more robust to environment changes and can help preserve user privacy. Real-world applications of skeleton action recognition include monitoring patient or elderly health, enabling autonomous driving, and facilitating human-computer interactions. However, there could be unintended usages and we advocate responsible usage complying with applicable laws and regulations.

## References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *ICCV*, 2021. 3
- [2] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, 2021. 2, 3, 7
- [3] Erik Bochinski, Volker Eiselein, and Thomas Sikora. High-speed tracking-by-detection without using image information. In *AVSS*, 2017. 4
- [4] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *TPAMI*, 2019. 2
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 3
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 2, 4, 5
- [7] Yuxin Chen, Ziqi Zhang, Chunfeng Yuan, Bing Li, Ying Deng, and Weiming Hu. Channel-wise topology refinement graph convolution for skeleton-based action recognition. In *ICCV*, 2021. 1, 2, 3, 6
- [8] Hyung-gun Chi, Myoung Hoon Ha, Seunggeun Chi, Sang Wan Lee, Qixing Huang, and Karthik Ramani. Infogcn: Representation learning for human skeleton-based action recognition. In *CVPR*, 2022. 1
- [9] Wongun Choi, Khuram Shahid, and Silvio Savarese. What are they doing?: Collective activity classification using spatio-temporal relationship among people. In *ICCVW*, 2009. 2, 4
- [10] Vasileios Choutas, Philippe Weinzaepfel, Jérôme Revaud, and Cordelia Schmid. Potion: Pose motion representation for action recognition. In *CVPR*, 2018. 2
- [11] Liu Chunhui, Hu Yueyu, Li Yanghao, Song Sijie, and Liu Jiaying. Pku-mmd: A large scale benchmark for continuous multi-modal human action understanding. *arXiv:1703.07475*, 2017. 2
- [12] Patrick Dendorfer, Hamid Rezaatfighi, Anton Milan, Javen Shi, Daniel Cremers, Ian Reid, Stefan Roth, Konrad Schindler, and Laura Leal-Taixé. Mot20: A benchmark for multi object tracking in crowded scenes. *arXiv:2003.09003*, 2020. 1
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv:1810.04805*, 2018. 3
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv:2010.11929*, 2020. 3, 4
- [15] Haodong Duan, Jiaqi Wang, Kai Chen, and Dahua Lin. Pyskl: Towards good practices for skeleton action recognition. *ACMMM*, 2022. 1, 2, 5, 6
- [16] Haodong Duan, Yue Zhao, Kai Chen, Dahua Lin, and Bo Dai. Revisiting skeleton-based action recognition. In *CVPR*, 2022. 1, 2, 5, 6
- [17] Bernard Ghanem Fabian Caba Heilbron, Victor Escorcia and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, 2015. 2
- [18] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *ICCV*, 2021. 3
- [19] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019. 2
- [20] Kirill Gavrilyuk, Ryan Sanford, Mehrsan Javan, and Cees GM Snoek. Actor-transformers for group activity recognition. In *CVPR*, 2020. 6
- [21] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The” something something” video database for learning and evaluating visual common sense. In *ICCV*, 2017. 2
- [22] Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *CVPR*, 2018. 1, 2, 5
- [23] Mohamed E Hussein, Marwan Torki, Mohammad A Gawayyed, and Motaz El-Saban. Human action recognition using a temporal hierarchy of covariance descriptors on 3d joint locations. In *IJCAI*, 2013. 2
- [24] Mostafa S Ibrahim, Srikanth Muralidharan, Zhiwei Deng, Arash Vahdat, and Greg Mori. A hierarchical deep temporal model for group activity recognition. In *CVPR*, 2016. 1, 2, 4, 5
- [25] Haroon Idrees, Amir R Zamir, Yu-Gang Jiang, Alex Gorban, Ivan Laptev, Rahul Sukthankar, and Mubarak Shah. The thumos challenge on action recognition for videos “in the wild”. *CVIU*, 2017. 2
- [26] Yanli Ji, Guo Ye, and Hong Cheng. Interactive body part contrast mining for human interaction recognition. In *2014 IEEE international conference on multimedia and expo workshops (ICMEW)*, pages 1–6. IEEE, 2014. 2
- [27] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *CVPR*, 2014. 2
- [28] Ang Li, Meghana Thotakuri, David A Ross, João Carreira, Alexander Vostrikov, and Andrew Zisserman. The ava-kinetics localized human actions video dataset. *arXiv:2005.00214*, 2020. 2
- [29] Maosen Li, Siheng Chen, Xu Chen, Ya Zhang, Yanfeng Wang, and Qi Tian. Actional-structural graph convolutional networks for skeleton-based action recognition. In *CVPR*, 2019. 6
- [30] Yixuan Li, Lei Chen, Runyu He, Zhenzhi Wang, Gangshan Wu, and Limin Wang. Multisports: A multi-person video dataset of spatio-temporally localized sports actions. In *ICCV*, 2021. 2

- [31] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 5
- [32] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C Kot. Ntu rgb+d 120: A large-scale benchmark for 3d human activity understanding. *TPAMI*, 2020. 1, 5
- [33] Jun Liu, Amir Shahroudy, Dong Xu, and Gang Wang. Spatio-temporal lstm with trust gates for 3d human action recognition. In *ECCV*, 2016. 3, 6
- [34] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 3
- [35] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *CVPR*, 2022. 3
- [36] Ziyu Liu, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, and Wanli Ouyang. Disentangling and unifying graph convolutions for skeleton-based action recognition. In *CVPR*, 2020. 1, 2, 3, 6
- [37] Qianhui Men, Edmond SL Ho, Hubert PH Shum, and Howard Leung. A two-stream recurrent network for skeleton-based human interaction recognition. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 2771–2778. IEEE, 2021. 2
- [38] Xuan Son Nguyen. Geomnet: A neural network based on riemannian geometries of spd matrix space and cholesky space for 3d skeleton-based interaction recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13379–13389, 2021. 6
- [39] Junting Pan, Siyu Chen, Mike Zheng Shou, Yu Liu, Jing Shao, and Hongsheng Li. Actor-context-actor relation network for spatio-temporal action localization. In *CVPR*, 2021. 2, 7
- [40] Yunsheng Pang, Qihong Ke, Hossein Rahmani, James Bailey, and Jun Liu. Igformer: Interaction graph transformer for skeleton-based human interaction recognition. In *ECCV*, 2022. 2, 5, 6, 7
- [41] Mauricio Perez, Jun Liu, and Alex C Kot. Interaction relational network for mutual action recognition. *IEEE Transactions on Multimedia*, 24:366–376, 2021. 2, 5, 6
- [42] Chiara Plizzari, Marco Cannici, and Matteo Matteucci. Skeleton-based action recognition via spatial and temporal transformer networks. *CVIU*, 2021. 3, 6
- [43] Michalis Raptis and Leonid Sigal. Poselet key-framing: A model for human activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2650–2657, 2013. 2
- [44] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS*, 2015. 5
- [45] Kohei Sendo and Norimichi Ukita. Heatmapping of people involved in group activities. In *MVA*, 2019. 5
- [46] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+d: A large scale dataset for 3d human activity analysis. In *CVPR*, 2016. 1, 2, 5
- [47] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *CVPR*, 2019. 2, 3
- [48] Bing Shuai, Alessandro Bergamo, Uta Buechler, Andrew Berneshawi, Alyssa Boden, and Joseph Tighe. Large scale real-world multi-person tracking. *ECCV*, 2022. 1
- [49] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *NeurIPS*, 2014. 2
- [50] Sebastian Stein and Stephen J McKenna. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In *UbiComp*, 2013. 2
- [51] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, 2019. 2, 5
- [52] Haritha Thilakarathne, Aiden Nibali, Zhen He, and Stuart Morgan. Pose is all you need: The pose only group activity recognition system (pogars). *Machine Vision and Applications*, 2022. 5, 6
- [53] Arash Vahdat, Bo Gao, Mani Ranjbar, and Greg Mori. A discriminative key pose sequence model for recognizing human interactions. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 1729–1736. IEEE, 2011. 2
- [54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 3
- [55] Raviteja Vemulapalli, Felipe Arrate, and Rama Chellappa. Human action recognition by representing 3d skeletons as points in a lie group. In *CVPR*, 2014. 2
- [56] Hongsong Wang and Liang Wang. Modeling temporal dynamics and spatial configurations of actions using two-stream recurrent neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 499–508, 2017. 2
- [57] Jiang Wang, Zicheng Liu, Ying Wu, and Junsong Yuan. Mining actionlet ensemble for action recognition with depth cameras. In *CVPR*. IEEE, 2012. 2
- [58] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016. 2
- [59] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018. 2, 7
- [60] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *CVPR*, 2019. 5
- [61] Fanyi Xiao, Yong Jae Lee, Kristen Grauman, Jitendra Malik, and Christoph Feichtenhofer. Audiovisual slowfast networks for video recognition. *arXiv:2001.08740*, 2020. 2
- [62] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017. 5
- [63] Mingze Xu, Yuanjun Xiong, Hao Chen, Xinyu Li, Wei Xia, Zhuowen Tu, and Stefano Soatto. Long short-term transformer for online action detection. *NeurIPS*, 2021. 3

- [64] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, 2018. 1, 2, 3
- [65] Di Yang, Yaohui Wang, Antitza Dantcheva, Lorenzo Garattoni, Gianpiero Francesca, and François Brémond. Unik: A unified framework for real-world skeleton-based action recognition. *ArXiv*, 2021. 2, 4, 5, 6
- [66] Kiwon Yun, Jean Honorio, Debaleena Chattopadhyay, Tamara L Berg, and Dimitris Samaras. Two-person interaction detection using body-pose features and multiple instance learning. In *2012 IEEE computer society conference on computer vision and pattern recognition workshops*, pages 28–35. IEEE, 2012. 2
- [67] Yimeng Zhang, Xiaoming Liu, Ming-Ching Chang, Weina Ge, and Tsuhan Chen. Spatio-temporal phrases for activity recognition. In *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part III 12*, pages 707–721. Springer, 2012. 2
- [68] Yuhan Zhang, Bo Wu, Wen Li, Lixin Duan, and Chuang Gan. Stst: Spatial-temporal specialized transformer for skeleton-based action recognition. In *ACMMM*, 2021. 3
- [69] Zhengyou Zhang. Microsoft kinect sensor and its effect. *IEEE multimedia*, 2012. 2
- [70] Hang Zhao, Antonio Torralba, Lorenzo Torresani, and Zhicheng Yan. Hacs: Human action clips and segments dataset for recognition and temporal localization. In *ICCV*, 2019. 2
- [71] Jiaojiao Zhao, Yanyi Zhang, Xinyu Li, Hao Chen, Bing Shuai, Mingze Xu, Chunhui Liu, Kaustav Kundu, Yuanjun Xiong, Davide Modolo, et al. Tuber: Tubelet transformer for video action detection. In *CVPR*, 2022. 3
- [72] Honglu Zhou, Asim Kadav, Aviv Shamsian, Shijie Geng, Farley Lai, Long Zhao, Ting Liu, Mubbasir Kapadia, and Hans Peter Graf. Composer: Compositional learning of group activity in videos. *ECCV*, 2022. 5, 6, 7
- [73] Wentao Zhu, Cuiling Lan, Junliang Xing, Wenjun Zeng, Yanghao Li, Li Shen, and Xiaohui Xie. Co-occurrence feature learning for skeleton based action recognition using regularized deep lstm networks. In *AAAI*, 2016. 3
- [74] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv:2010.04159*, 2020. 3