#### **ORIGINAL PAPER**



# Pose is all you need: the pose only group activity recognition system (POGARS)

Haritha Thilakarathne<sup>1</sup> · Aiden Nibali<sup>1</sup> · Zhen He<sup>1</sup> · Stuart Morgan<sup>1,2</sup>

Received: 11 February 2021 / Revised: 29 June 2022 / Accepted: 16 September 2022 / Published online: 10 October 2022 © The Author(s) 2022

#### Abstract

We introduce a novel deep learning-based group activity recognition approach called the Pose Only Group Activity Recognition System (POGARS), designed to use only tracked poses of people to predict the performed group activity. In contrast to existing approaches for group activity recognition, POGARS uses 1D CNNs to learn spatiotemporal dynamics of individuals involved in a group activity and forgo learning features from pixel data. The proposed model uses a spatial and temporal attention mechanism to infer person-wise importance and multi-task learning for simultaneously performing group and individual action classification. Experimental results confirm that POGARS achieves highly competitive results compared to state-of-the-art methods on a widely used public volleyball dataset despite only using tracked pose as input. Further, our experiments show by using pose only as input, POGARS has better generalization capabilities compared to methods that use RGB as input.

Keywords Group activity recognition · Human pose analysis · Self-attention · Deep learning

# 1 Introduction

Developing methods for recognizing activities performed by groups of people is an important computer vision research area applicable for many domains, with two of the most prominent tasks being analysis of sports videos and monitoring of surveillance video footage. Examples in the sports domain include identifying events such as shots on goal in soccer matches, and pick and roll events in basketball. Examples in surveillance videos include people fighting, walking together, people being followed and other similar

 Haritha Thilakarathne p.thilakarathne@latrobe.edu.au
 Aiden Nibali a.nibali@latrobe.edu.au
 Zhen He z.he@latrobe.edu.au
 Stuart Morgan stuart.morgan@sportaus.gov.au

<sup>1</sup> Department of Computer Science and Information Technology, La Trobe University, Plenty Road and Kingsbury Dr, Bundoora, VIC 3086, Australia

<sup>2</sup> Sport Strategy and Investment, Australian Institute of Sport, Leverrier St, Bruce, ACT 2617, Australia human activities. The major difference between group activity recognition and individual action classification is the need to simultaneously reason about multiple people. Thus, group activity recognition methods should be specifically designed to understand spatial and temporal evolution of pose for groups of people.

It is common for existing work on group activity recognition domain to operate on video pixel data directly, but thus far it has remained an open question as to how well group activity can be predicted from extracted human poses. The objective of our research is to explore the predictive power of pose data by developing a deep learning model that predicts a volleyball group activity label only from tracked pose, without directly using RGB pixel data.

Majority of the existing group activity recognition models utilize person appearance features extracted via Convolutional Neural Network (CNN)-based models to learn their spatiotemporal dynamics throughout consecutive frames in the video instance [1–3]. Variants of Recurrent Neural Networks (RNNs) such as Long Short-Term Memory (LSTM) networks or Gated Recurrent Units (GRUs) are most widely used to model temporal relationships of activities. To better map the relationships between individuals in the scene, some such existing works have introduced pooling strategies and other interpersonal processing techniques [4,5]. In contrast to most existing methods which are focused on taking RGB features as input [1,2] or RGB with additions like optical flow [5,6] and/or pose [3,7], proposed method uses only pose to predict the group activity label.

Extracting useful high level information such as pose and tracking data from video is a common practice in sports domain [8-11]. Using pose to represent human subjects has several advantages. Pose representations ignore less relevant factors in the input data such as person appearance features and background which allows the models trained using pose to generalize better to new situations like sporting events involving previously unseen teams and venues. Ignoring person level identifiable appearance features allows pose data to preserve privacy of the people whose data are being captured. In addition, recent commercial computer vision systems such as Microsoft Azure Kinect are capable of directly producing pose information as a better representation of human dynamics in video streams. As more researchers realize the usefulness of pose information and methods for capturing such data become more readily available, we anticipate that pose information will be collected for a range of use cases such as analyzing basketball player movements, motion tracking and other similar human activities. Our method can then take advantage of the already available pose information to more efficiently perform group action recognition. This is due to the much lower dimensionality of pose information (usually just 16 key points per person) compared to RGB images.

In addition to pose information, Perez et al. [12] recently explored the usage of related object (i.e., ball) trajectories for identifying group activities in sports videos, finding that the additional input information can substantially improve classification accuracy. Our experiments partially support this finding, with the inclusion of ball tracklets exhibiting slightly improved empirical results. However, our approach is not reliant on ball trajectories to achieve high classification accuracy.

In this study, we propose Pose Only Group Activity Recognition System (POGARS), a novel system for performing group activity recognition by utilizing extracted tracked pose as input. POGARS uses 1D CNNs to learn the spatiotemporal dynamics of individuals using their pose keypoint estimations and position tracklets. We found 1D CNNs were able to model the temporal dynamics in the video more effectively than the RNNs used by most existing methods for group activity recognition.

In most group activity settings, there are certain people who play more central roles than others and are hence more predictive of the group activity recognition task. For example, when recognizing a spike in volleyball, we expect more attention should be placed on the players directly involved with spiking and blocking the ball. POGARS uses a spatial self-attention mechanism for identifying the importance of each individual for the particular group activity. In addition, temporal attention is also important. For instance, video frames at which the volleyball is passed can be considered as an important temporal instance in the video clip. Proposed approach uses a temporal self-attention mechanism to assign different importance weights to the different video frames.

While recognizing group activity of the video instance, we simultaneously perform individual action recognition as an additional task, thus making our model perform multitask learning. This was found to result in a noticeable accuracy improvement with respect to the primary task of group activity recognition.

Though group activity recognition is an important research area in computer vision, there are few video datasets available for model evaluation. The volleyball dataset introduced by Ibrahim et al. [13] is one of the most popular group activity recognition datasets. Existing works such as [1,3,5,6] have used the volleyball dataset to evaluate group activity recognition models. In order to evaluate and compare our proposed approach with the past related work, we have conducted our experiments on the volleyball dataset.

POGARS achieves 93.2% accuracy for group activity recognition by only using tracked pose information on the volleyball dataset. To contextualize this result, the state-ofthe-art actor-transformer model proposed by Guvrilyuk et al. [7] achieves 92.3% when using pose only. (Their stateof-the-art accuracy of 94.4% requires optical flow data as an additional input modality.) Another very important benefit of only using tracked pose is that the trained model can generalize better to testing data with different characteristics by ignoring such factors as lighting conditions, team uniform color, the color of the court, audience appearance, etc. To test this hypothesis, we created a skewed train/test split of the volleyball data where the training data consisted of only games played in a single venue in the London 2012 Olympics and the test set consisted of matches from all other venues. The results show POGARS's accuracy only dropped slightly from 93.2 to 89.7% when using the skewed data. In contrast, a competitive I3D model [9] trained using RGB images as input showed a very large drop in accuracy, from 84.6 to 73.9%.

Our contributions can be summarized as follows:

- We propose POGARS, the first group activity recognition system that solely utilizes tracked poses as input. We showed a POGARS model trained on a single venue generalizes significantly better to different test venues compared to a competitive model trained on RGB.
- We performed a series of experiments and found a 1D CNN-based approach that uses spatial and temporal attention outperformed the traditional approach of using RNNs for group activity recognition.

• We conducted experiments on a popular group activity recognition dataset to show that POGARS achieves near state-of-the- art accuracy by only using tracked pose for model input.

# 2 Related work

In this section, we review existing work in the area of group activity recognition and the use of deep learning techniques to solve the problem.

Group activity recognition considers the multi-person behavioral and interaction dynamics to interpret the collective activity instances. Most of the early work on group activity recognition is based on models that learn information in video frames using hand-crafted features. Hierarchical graphical models [14,15] and dynamic Bayesian networks [16] are some of the popular methods used for interpreting group activity in video. Choi et al. [17] introduced a spatio-temporal feature descriptor based on shape context [18] for interpreting collective activities of people. Hierarchical models such as [14] interpreted the group activity by representing the actions of each individual person as well as person to person interactions using a tree-like graphical descriptor. Handcrafted feature engineering approaches for collective activity recognition come with various limitations such as domain dependency, large computational costs and poorer resultant accuracy.

The focus of recent collective activity recognition models has shifted towards using deep neural networks [19]. Deng et al. [20] introduced Structure Inference Machine, which has a RNN-based backbone for analyzing interactions between people in the scene in order to recognize group activity. Ibrahim et al. [1] proposed a group activity recognition model that consists of two LSTM-based sub-models where the first one encodes the individual player level actions and their temporal dynamics, while a second model fuses outputs from the first model for providing temporal dynamics of the group activity. Inspired by this work, two-level RNN-based hierarchical methods have become the most popular approach in the domain of collective activity recognition [1,2,21,22]. Though these models learn the spatial features of each person in the scene, most of them do not consider the spatial position of the individuals. In contrast to the above existing work, our proposed model uses person level pose and spatial location information (position tracklets) for predicting the collective activity since combination of pose and location information is a dense information representation.

In order to learn relations between individuals, Wu et al. [23] employed graph convolutional networks in their Actor Relation Graph (ARG) designed for identifying group activities. 2D CNNs are used to extract features from person bounding boxes which are used to capture the appearance and position relations between actors by ARG. In contrast to ARG, our proposed model is able to learn the temporal relationship of the individuals in the group activity through the inclusion of 1D convolutional layers.

Azar et al. [5] proposed a multi-stream CNN framework for group activity recognition which uses optical flow, warped optical flow, pose heatmaps and the RGB frame as different input modalities for the predictive model. Features extracted from CNNs are first applied to each input stream and then fused to produce the group activity prediction for video instances. Azar et al. [6] recently introduced Convolutional Relational Machine (CRM). Emphasizing the importance of spatial position information of individuals in group activity classification, the authors have proposed to learn the spatial relations between the activities of people using activity maps in CRM. Similar to [5], CRM also uses optical flow information to enhance prediction accuracy.

Following a similar approach of using CNNs to predict group activities in videos, Gavrilyuk et al. [7] introduced an actor-transformer model which uses optical flow for representing temporal dynamics, while pose information has been used for interpreting spatial information of the people. Our method does not use optical flow information as input due to the high computational costs.

# 2.1 Pose-based action recognition

Pose data are lightweight and more robust to visual clutter and noise compared to RGB data. Hence, many individual action recognition models use different forms of pose data representations as input feature embeddings [8,11,24,25]. Zhu et al. [11] introduced "Action Machine" which first feeds cropped images of people as input into an I3D model [9]. The model has a RGB stream and also a pose stream. These two streams are fused together to predict action classes. In contrast to these works (which perform individual action recognition), we perform group activity recognition using tracked pose data.

Inspired by the success of pose-based individual action recognition, Lu et al. [3] proposed a group activity recognition model built from pose skeleton data and deep RGB features of individuals in each video frame. Temporal dynamics of group activity were learnt using GRUs.

Perez et al. [12] recently proposed a skeleton-based relational reasoning model for group activity analysis which uses pose skeleton data and ball tracklets for learning interactions between individuals in related objects.

When compared to all the models in this section, we differ in the following ways: (1) We use only pose data instead of both RGB and pose data; and (2) we are the only solution that uses 1D CNNs to model temporal patterns.

#### 2.2 Attention models

Attention has been successfully applied to both computer vision and natural language understanding domains [26,27]. In the most recent studies, collective activity recognition models have used attention mechanisms either as spatial attention or temporal attention. Ramanathan et al. [2] introduced the concept of key actors to understand collective activity. A bidirectional LSTM-based model was used to represent the tracks of people in the scene, and the model is capable of attending for a subset of the people actively engaged in collective activity. The latent embedding model introduced by Tang et al. [28] consisted of an attention mechanism to encode the importance of each individual person in the collective activity scenario.

In order to identify video frames which make more contribution towards the prediction, different temporal attention approaches have been used in group activity recognition problem domain [29,30]. Attention mechanisms employed in these models were designed to extract a small subset of predictive frames from the whole temporal span of the video. Instead of making a hard decision to include frames as these works do (so-called hard attention), we use a soft temporal attention mechanism in POGARS which weights the contribution of features from each temporal instant.

Lu et al. [3] introduced an attention-based collective activity recognition model which uses pose and RGB features of each person in the scene to predict the group activity label. It combines both spatial and temporal attention mechanisms. While RGB data were the primary input for their model, the authors used pose features for determining the distribution of attention among the people. Our work differs from [3] in the following ways: 1) We use 1D CNNs instead of RNNs to extract temporal patterns; and 2) we only use pose as input instead of both pose and RGB as attention inputs.

# 3 Pose only group activity recognition system (POGARS)

In comparison with raw image frames, tracked poses of the individuals from the video contain a much denser set of important information allowing the deep learning model to focus only on the most important features. Such as a person who is about to perform a spike in volleyball has a vastly different pose from someone who is running for the ball. Pose keypoints can be represented as numerical coordinates, as opposed to the high-dimensional pixel data of an RGB image.

One of the biggest challenges for modeling data in video is finding temporal coherence (tracking where one point in one frame moves to the next frame). Modeling temporal coherence in sequences of poses is more natural than in video data since each keypoint is assigned to a different dimension in the pose vector. Hence, 1D CNNs can effectively model the changes in position of each keypoint across time. Further highlighting the most important data for activity recognition, we also use a self-attention mechanism to assign a higher weight to more important people and frames for accurate group activity classification. Using extracted pose features instead of RGB features can produce a model that is better able to generalize across datasets since pose excludes factors from the input, such as lighting conditions, team uniform color, the color of the court, audience appearance and other nuisance factors.

Figure 1 shows an overview of the proposed POGARS approach. First, position coordinates and associated bounding boxes of each person across time in the input video are acquired from manual annotations used in [31]. Next, we generate 16 2D keypoint estimations for each individual by feeding the bounding box tracklets to stacked hourglass human pose estimation algorithm [32]. Position coordinate tracklets are fed into a fully connected layer to create a descriptive feature embedding for the position of each individual in the frame. Position embeddings and keypoint pose representations are concatenated together to produce a composite feature embedding for each individual that contains pose and position information.

Feature representations of consecutive frames in the video clip contribute in different degrees to class prediction. The concatenated feature embeddings of each person in each frame are fed into a temporal attention mechanism to calculate frame-wise attention weights. An element-wise multiplication between the feature embedding and attention weights across the temporal dimension is performed before feeding the features into a set of 1D CNN blocks. Subsequent CNN blocks with residual connections are able to learn temporal dynamics of the individuals using the pose and position embeddings.

POGARS performs multi-task learning by predicting both group activity label  $p_G$  and individual action labels  $p_I$ . This is achieved by attaching two separate heads (one for group activity prediction and the other for individual activity prediction) to the output feature set from 1D CNN blocks.

To implement spatial attention, attention weights of each individual are multiplied element-wise with the corresponding feature embedding to ensure the feature embeddings of key people in the video get more attention in the process of predicting the group activity. Detailed descriptions of each component of our method are presented in the following subsections.

#### 3.1 Tracking and pose estimation

Since our collective activity recognition model is based on position and pose feature representations of individuals, we **Fig. 1** Overview of pose only group activity recognition system (POGARS)



employ manually annotated bounding boxes of the people from [31] which are in the image space. Our approach decouples tracking and pose estimation from the model used to classify the group activities. In the event that manually annotated tracks are not available, any multi-object detection and tracking algorithm can be used, with the final classification accuracy expected to correlate positively with the performance of these algorithms.

After object detection, we extract the human pose of each person. There are many existing 2D human pose estimation

models which are appropriate for inferring such pose keypoints [32–34]. Pose estimation methods which first detect body joints and then group them to form individual poses per person are called bottom-up approaches. In contrast, top-down approaches first detect people in the scene and then predict keypoints of each detected person. Top-down approaches have been found to be more accurate than bottomup approaches for performing pose estimation [35]. Hence, we take a top-down approach by first using [36] to find bounding boxes around the people and then employing the pre-trained stacked hourglass model [32] to estimate joint coordinates of each person. Due to the modular nature of our group activity recognition method, any human tracking algorithm and 2D human pose estimation algorithm can be used with our proposed approach. This includes bottom-up pose tracking approaches, which may be suitable candidates for situations where computational speed is of key importance.

Given an input video clip *V* containing *N* people and spanning *T* frames, manual annotations have been used to generate *N* bounding box tracklets  $B \in \mathbb{R}^{N \times 4 \times T}$  and position coordinate tracklets of each bounding box  $C \in \mathbb{R}^{N \times 2 \times T}$ .

For each individual, 16 2D pose keypoints  $K \in \mathbb{R}^{N \times 32 \times T}$  are estimated using the stacked hourglass pose estimation algorithm. A linear layer is added on top of the bounding box position coordinates *C* in order to generate a descriptive feature embedding  $\hat{C} \in \mathbb{R}^{N \times 32 \times T}$  to represent the spatial position of the people in each frame.

Pose feature representation *K* and position feature representation  $\hat{C}$  of each corresponding individual are concatenated, so  $(K \mid \hat{C}) = F \in \mathbb{R}^{N \times 64 \times T}$  feature embedding of the tracked pose is generated.

Since the temporal evolution block is responsible for learning each person level feature separately, the order of people is important. Thus, we ensured that the pose feature representations of each individual in each frame are ordered consistently by stacking them in a left to right in terms of x-position in the video frame.

## 3.2 POGARS with different person level fusion and temporal modeling approaches

POGARS captures the spatial and temporal evolution of individuals in order to predict the group activity label. Person level spatiotemporal features are fused together to generate a person level feature representation. Learning group activities in the video instances is done by modeling temporal evolution of the fused feature representations. We explored two different person level fusion strategies and two temporal modeling variations within POGARS.

#### 3.2.1 Person-level fusion

Figure 2 illustrates the two different person level fusion approaches we explored. These two approaches are early person fusion and late person fusion. In **early person fusion**, person-level feature independence is removed before modeling the temporal evolution by fusing the tracked pose features of all people before input into the module that captures temporal dynamics. Alternatively, in **late person fusion** the temporal evolution of each person is first modeled individually before the learned features are fused together at the end.

In the late person fusion method, the temporal evolution block is responsible for learning the tracked pose of each person separately while the early person fusion method combines person level feature embeddings of each frame before combining the temporal dimension using the temporal modeling block. The early person fusion approach eliminates the need for person level tracking by combining the features of each person before temporal modeling. Although not needing to perform tracking eliminates one potential source of error, our experiments show person-wise spatiotemporal feature learning (as provided by late person fusion) is essential for achieving higher classification accuracy.

#### 3.2.2 Temporal evolution modeling

Most of the existing group activity recognition mechanisms [1,3,20,21] use LSTM- or GRU-based networks for learning temporal dynamics of the video. In our experiments, we evaluated two temporal evolution modeling approaches: bidirectional LSTMs and 1D CNN, and found that POGARS networks which utilize 1D convolutions outperform their LSTM counterparts. We give a detailed description of our 1D CNN-based temporal modeling block in Sect. 3.4.

#### 3.3 Temporal attention mechanism

Group activities have different important stages in their temporal span. For example, in a volleyball spiking activity, the frames in close temporal proximity to the instant that the player strikes the ball are typically more important than rest of the frames. Directing the model to attend to features from the most important frames prevents the model from overfitting to less important frames which may contain irrelevant information. Hard attention-based methods such as [29,37] detect a specific set of video frames that maximally contribute to the final prediction. Rather than selecting a small number of frames to keep for further analysis and discarding the rest, we propose a soft self-attention mechanism which assigns an importance weight to every frame.

The attention mechanism we have employed in POGARS contains a set of linear layers and a softmax layer. Given the feature embedding of individuals F, temporal attention is computed as:

$$F' = \operatorname{softmax} \left(\phi(F)\right) \odot F \tag{1}$$

where  $\phi$  is the temporal attention function (modeled using linear layers). Computed temporal attention weights are multiplied element-wise across the time dimension with the feature embedding *F*. This produces  $F' \in \mathbb{R}^{N \times 64 \times T}$  which is the feature embedding of the video clip weighted according to frame-wise temporal importance.



Fig. 2 Different person level fusion approaches of POGARS. Each approach was tested with two temporal evolution modeling variations: bidirectional LSTMs and 1D CNNs

# 3.4 Temporal evolutions of individuals based on tracked pose

In the proposed method, people in the video clip are represented by feature embeddings of their respective tracked poses. Learning temporal evolution of the pose and position dynamics of individuals is critical for predicting the collective activity performed by the group. We employed 1D CNNs to model the temporal evolutions of individuals instead of RNNs, based on the study we did with different POGARS architectures. In contrast to RNNs, the ability of 1D CNNs to learn deep related temporal features and translational equivariance makes them a better fit for representing fine grained pose and position evolutions of individuals across time. Figure 3 shows the composition of convolutional blocks used in POGARS. Each convolutional block consists of 3 1D CNN layers with a skip connection, and POGARS contains 4 such blocks stacked together. The stack takes the post-attention



Fig. 3 Composition of a single convolutional block of POGARS

feature embeddings F' as input and analyzes the motion of individuals by producing a temporally convolved feature representation  $E \in \mathbb{R}^{N \times 1024 \times T}$ .

#### 3.5 Spatial attention mechanism

In POGARS, we employ a spatial attention mechanism that uses temporally convolved feature extractions of each individual to identify each person's importance score for the particular group activity. Our spatial attention block contains a set of linear layers and a softmax layer. Temporally convolved feature embeddings E are average pooled in the time dimension to produce a feature embedding for each person in the video clip  $\hat{E} \in \mathbb{R}^{N \times 1024 \times 1}$  where N is the number of individuals in the video. The spatial attention is computed as:

$$\hat{E}' = \operatorname{softmax}(\psi(\hat{E})) \odot \hat{E}$$
 (2)

where  $\psi$  is the spatial attention function. Computed spatial attention weights are multiplied element-wise across the person dimension with the feature embedding  $\hat{E}$ . This produces  $\hat{E}' \in \mathbb{R}^{N \times 1024 \times 1}$  which is the feature embedding of all individuals in the video clip weighted accordingly to each person's importance.

# 3.6 Multi-task learning for group activity recognition and individual action recognition

In parallel to learning the collective activity label, POGARS uses parameter sharing to also predict the individual actions of each person in the video clip. The average pooled feature representation is shared between two task-specific linear layer stacks. One of these layers predicts the group activity label, while the other predicts individual action labels for each person. Feature embeddings used for individual activity recognition are not weighted with the spatial attention mechanism.

In order to train our proposed multi-task model, the losses of group activity recognition task and individual action recognition are optimized simultaneously. Considering  $\mathcal{L}_{IA}$  as the individual action prediction loss and  $\mathcal{L}_{GA}$  as the group activity recognition loss, the multi-task loss function  $\mathcal{L}_{MT}$  of the model is defined as follows:

$$\mathcal{L}_{MT} = \alpha \mathcal{L}_{IA} + \mathcal{L}_{GA} \tag{3}$$

We used multi-class cross-entropy loss as the loss function for both individual and group activity prediction in our experiments. To achieve the optimum accuracy for group activity recognition, we found that  $\alpha$  set to 2.0 gave the best results.





Fig. 4 Adding ball as an input modality for POGARS

# 3.7 Relationship between individuals and relevant objects for group activity recognition

Using specific activity-related objects such as a ball, bat or a puck is a common scenario in group sports. Perez et al. [12] explored the effectiveness of learning ball tracklets for predicting group activity instances in volleyball videos using their GIRN model. In order to benchmark POGARS with GIRN, we experimented with adding position coordinates of the volleyball as an input modality for our network (see Fig. 4).

Given the ball position coordinates in the image space  $O \in$  $\mathbb{R}^{2 \times T}$ , sinusoidal positional encoding [38] is applied in order to generate a descriptive feature embedding  $\hat{O} \in \mathbb{R}^{64 \times T}$ . The feature embeddings are fed into a stack of 1D CNN blocks (similar to the architecture described in Sect. 3.4) to produce a temporally convolved feature representation  $G \in \mathbb{R}^{1024 \times T}$ , followed by global average pooling to remove the temporal dimension. The resulting ball features are concatenated to  $\hat{E}'$  for generating  $H \in \mathbb{R}^{(N+1) \times 1024}$ , which is employed for predicting group activities by using a linear layer stack.

Page 9 of 15 95

Method	Accuracy
POGARS - early person fusion setting with LSTM-based network	81.4
POGARS - late person fusion setting with LSTM-based network	85.9
POGARS - early person fusion setting with 1D CNN-based network	83.0
POGARS - late person fusion setting with 1D CNN-based network	88.3
POGARS - with only spatial attention	90.8
POGARS - with only temporal attention	91.7
POGARS - with spatial and temporal attention	92.1
POGARS ALL - with spatial, temporal attention and multitask learning	93.2

The entry with the bold letters (POGARS ALL) is gives the highest accuracy on Volleyball dataset

# **4 Experiments**

Table 1Comparison ofdifferent settings of thePOGARS evaluated onVolleyball dataset

#### 4.1 Experimental setup

POGARS is evaluated by performing experiments on the popular volleyball dataset [13]. Manually annotated bounding box tracklets and position coordinate tracklets of each player in the volleyball videos are obtained from [31]. The annotations are in image space. We utilized stacked hourglass pose keypoint estimation algorithm [32] for predicting 16 keypoints (*ankles, knees, hips, pelvis, spine, neck, head, wrists, elbows and shoulders*) of each player in the video frames. The stacked hourglass network has been pretrained using the MPII dataset [39]. Coordinates of the extracted keypoints are expressed relative to the pelvis joint location, whereas the position coordinate tracklets are expressed in absolute frame space.

Coordinate tracklets of the people are scaled to values between 0 and 1. From each video clip, 36 frames are utilized in model training and evaluation. We used stochastic gradient descent with ADAM [40] with initial learning rate set to  $10^{-3}$ , step-wise learning rate decay with a step size of 7 and fixed moving average decay rates  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ . Models are implemented using the PyTorch deep learning framework [41] and trained on an NVIDIA GeForce RTX 2080 Ti GPU.

#### 4.2 Volleyball dataset

The volleyball dataset introduced by Ibrahim et al. [13] contains 55 videos collected from publicly available YouTube volleyball matches. It includes 4830 trimmed group activity instances which belong to 8 activity classes. Each activity instance contains 41 frames where the middle frame is labeled with the group activity label and individual action labels. The 9 individual action labels are *spiking*, *blocking*, *setting*, *jumping*, *digging*, *standing*, *falling*, *waiting*, and *moving*. The 8 group activity labels are *right spike*, *left spike*, *right set*, *left set*, *right pass*, *left pass*, *right winpoint*, and *left winpoint*. For our experiments, we use the same manual bounding box annotations and player positions as [31]. We have also followed the train/test splits suggested by [13]. In order to reduce model overfitting, we perform data augmentation by also training on horizontally flipped versions of examples from the training set and flipping the activity labels accordingly.

#### 4.3 Experiment results

#### 4.3.1 Ablation study

POGARS consists of two different attention mechanisms and a fusion mechanism that promotes learning spatiotemporal dynamics of the individuals involved with the activity. We performed an ablation study to assess the impact of these key features on overall performance. Moreover, we performed experiments on two different temporal modeling methods for POGARS.

As the results in Table 1 show, both spatial and temporal attentions are beneficial and complementary. Furthermore, the results show that the late person fusion approach achieves higher accuracy than early fusion. Consequently, we use late person fusion for the remainder of our experiments.

Instead of RNNs, we used 1D CNNs for the backbone of our model. Results show that POGARS using 1D Convolutions (late person fusion without attention) outperforms an LSTM-based counterpart model with the same basic architecture by 2.4%. This demonstrates the ability of 1D CNNs to learn lower-level fine grained feature dynamics of individuals. 1D CNNs are less sensitive to temporal shift due to implicit inductive bias in the model itself. This could, for example, allow the model to more easily learn feature maps that represent certain movements regardless of when they happen in the clip. In contrast, RNNs rely on hidden state to learn time-varying patterns, and this prevents RNNs from being temporally equivariant, making it harder to learn patterns which are applicable at any time offset. Moreover, we theorize that the translational equivariance of CNNs is useful for understanding repetitive movement dynamics of group activities.

While predicting the group activity label, our model is capable of predicting the individual action labels of people in the scene with the aid of a multi-task learning approach. We used a combined loss function as stated in Eq. 3 for model training. Superior accuracy of the multi-task model clearly indicates that multi-task learning leads to the implicit development of shared features, and that these features are more predictive of group activity than features learned through group activity recognition alone. Though individual action recognition is not the focus of our work, we achieved 79.5% accuracy for that particular task using pose information only.

#### 4.3.2 Generalization capability of POGARS

To test the generalization capability of POGARS, we created two experiments. In the first experiment in which we split the volleyball dataset according to the match venues, 29 videos out of the 55 videos in the dataset were recorded at the same 2012 London Olympics venue. We designated that particular set of videos as the train set, while the rest were used for model testing. We named this special data split as the "Olympic split".

Since most of the group activity recognition approaches [3,6,42] have used 3D convolution based networks, we used a I3D-based network [9] for the comparison with POGARS. The I3D network is a well-recognized and strong baseline algorithm for action recognition. The I3D network we used was pretrained on the Imagenet [43] and Kinetics datasets [9]. Table 2 summarizes the validation accuracies of I3D model and POGARS for the different data splits of volleyball dataset.

As shown in Table 2, POGARS exhibits only a 3.5 percentage point decrease in test accuracy when training and evaluating on the skewed Olympics data split, as opposed to the much larger drop of 10.7 percentage points exhibited by the I3D model. This observation implies training a model only using tracked pose (like for POGARS) is better able to generalize across different playing environments and venues compared to using RGB image input to train a model (like I3D model) for group activity recognition.

In the second experiment, we tested the generalizability of POGARS by testing the accuracy of our model trained on the public volleyball dataset [13] on a new real-world volleyball game we downloaded separately from YouTube. To increase the challenge in the task, we did not do any manual person detection, ball position annotation, tracking and pose estimation. Instead, we used the OpenPifPaf pose detection and tracking algorithm [44] to automatically generate player pose estimations and tracklets. This contrasts from our previous experiments with the volleyball dataset where ground truth detection and tracking information was used. Therefore, for this task our model needs to deal with erroneous tracks which were not present in the previous experiments.

In this task, we used the Brazil Vs. USA women's volleyball gold medal match of the 2020 Tokyo Olympics for our experiment. For testing POGARS, 8 group activity classes (same classes used in rest of the experiments) from the first set of the match were manually annotated. We randomly selected 6 activity instances from each class which were then trimmed into 36 frames long video instances. After using OpenPifPaf to extract the pose tracklets, we manually remove any tracks that do not belong to volleyball players.

The experimental results showed that POGARS was able to achieve 75% test accuracy on this more challenging problem setup. This is a good result since no ground truth tracking/ pose information was used; instead, the model needed to overcome any erroneous tracks. This result shows POGARS can cope with erroneous tracks and be able to generalize to different camera angles and playing environments.

# 4.3.3 Visualization of the spatial and temporal attention mechanisms

In order to qualitatively assess the effectiveness of the spatial and temporal attention mechanisms, we visualized both spatial and temporal attention outputs of the best performing POGARS.

Since group activities have different important stages in their temporal span, our temporal attention mechanism described in Sect. 3.3 is capable of assigning attention weights to each frame in accordance with its perceived importance. Each activity clip of the volleyball dataset we used in our experiments consists of 36 frames where the 16th frame is the central keyframe of the group activity. Therefore, we expect that frames neighboring the 16th frame (t = 15)should generally receive high attention when predicting the group activity label. Figure 5 illustrates the predicted mean temporal attention weights for each frame index as assigned by POGARS. The graph indicates a clear peak in attention weights when t = 15 (key frame) and comparatively high attention weights for the neighboring frames. This observation indicates that our temporal attention mechanism can identify the predictive power of frames from group activity video instances.

The spatial self-attention mechanism discussed in Sect. 3.5 attends to the key people involved in the group activity. Though group activities are labeled according to the side of the volleyball court they were performed from (such as *left spike* and *right spike*), players on both sides of the court contribute to a group activity. For example, the players on the left-hand side of the court may form a defensive formation when a right spike activity is taking place.

Figure 6 visualizes the spatial attention weights from six correct volleyball group activity predictions done by POG-

 Table 2
 Test accuracy for

 volleyball dataset with different
 train/test splits

Method	Default data split	Olympic data split	Difference	
I3D [9]	84.6	73.9	10.7	
POGARS ALL	93.2	89.7	3.5	



Fig. 6 Spatial attention assigned for the players by POGARS. Player bounding boxes are colored according to the spatial attention score. Key player of each group activity is annotated with dotted bounding boxes

ARS. The first row of the figure represents successful spatial attention cases, while the second row represents three failure spatial attention cases. The key player of each group activity instance was manually determined and denoted by dotted red bounding boxes.

Instance (a) is a *left spike* and instance (b) is a *right spike* group activity instance of volleyball data. In both cases, the player with highest attention is the one spiking the ball. In instance (c), which is a *left set* group activity, the player who's having the control of the ball is having the highest spatial attention. It is clear that POGARS has been able to identify the key player correctly in those group activity instances.

In (d) *left pass* and (e) *right pass* activity instances, players who have control of the ball are in a falling position which is different from the generic pass posture of a player. These different posture patterns in such cases can be interpreted as the cause for the failure of spatial attention assignment.

Instance (f) is a *right spike* group activity of volleyball data where spatial attention assigned by POGARS is incorrect. In this particular instance, two players have spike jumping postures which may cause incorrect identification of the prominent player of the activity. These instances can be defined as edge cases where players are having drastic differences in their posture as well as their position on the volleyball court compared to generic group activity instances.



Fig. 7 Spatial attention assigned for the players involved in right spike activity by POGARS. Key frame of the activity is visualized with player bounding boxes colored according to the spatial attention score

For more clarity, Fig. 7 shows a more zoomed in view of the algorithm assigning attention weights to players in a *right spike* group activity, along with several frames cropped around the player with the highest attention. In this case, the player with highest attention is the one spiking the ball, so it is clear that POGARS has correctly assigned attention to the keyplayer in this case.

# 4.3.4 Performance comparison of POGARS against rival methods

We compare the performance of POGARS with state-ofthe-art methods and 3 baselines. Table 3 reports the group activity recognition accuracy of the stated models on the volleyball dataset [13].

The baselines are:

- **B1** Activity recognition with keyframe spatial features: This baseline uses the middle frame of the activity video clip (21st frame among 41 frames) as the input to a ImageNet pre-trained ResNet34 model [48].
- **B2 Activity recognition with person level spatial and position features:** Similar to B1, except that we introduce a second stream for processing keyframe player bounding boxes represented as heatmaps. Learned spatial features from the two streams are fused together late in the architecture to predict the group activity.
- **B3** Activity recognition with spatiotemporal RGB features: This baseline takes 16 consecutive frames of the collective activity as input to a network based on

ResNet18 architecture with 3D CNN layers. B3 can be thought of as a temporally extended version of B1.

Both B1 and B2 use keyframe-based spatial feature extractions for learning the group activity. Though these models do not consider the temporal aspect of the activities, they performed reasonably well in collective activity recognition by leveraging state-of-the-art image classification network architectures and transfer learning [49]. In B2 we use a second input stream consisting of heatmap representations of player bounding boxes in the keyframe. The slight increase in the accuracy in B2 compared to B1 implies that using player positions as additional input is beneficial in predicting group activity labels.

Compared to B1 and B2, B3 consumes more memory and computational power due to the 3D convolutional layers used in the network. The low accuracy of B3 suggests that the inclusion of deep 3D CNN networks is not a favorable tradeoff in terms of computation versus accuracy. All baselines (B1, B2 and B3) use RGB-based feature extraction to learn the spatiotemporal dynamics of the group activities. In our proposed group activity recognition model, we utilize pose keypoint estimations of the people and their position coordinates to create a descriptive feature embedding for each individual. Compared with the baseline approaches, POG-ARS reports a significantly higher achieved accuracy as well as faster model convergence.

POGARS performs better than the majority of existing group activity recognition algorithms on the Volleyball dataset (see Table 3). This includes methods which use RGB image-based feature representations ([1,45]), variants

Table 3         Comparison of           different baselines and         Image: Comparison of the second	Method	Feature extraction method	Accuracy	
state-of-the-art methods with POGARS evaluated on volleyball dataset	B1 - keyframe-based model	RGB	73.3	
	B2 - person level feature-based model	RGB	75.1	
	B3 - Spatio Temporal feature-based model	RGB	63.0	
	Two-stage Hierarchical model [1]	RGB	81.9	
	CERN [45]	RGB	83.3	
	I3D [9]	RGB	84.6	
	Social Scene [4]	RGB	89.9	
	Action Relation Graph [46]	RGB	92.6	
	Joint learning with social groups [47]	RGB	93.1	
	Multi-stream CNN [5]	RGB + Optical flow	90.5	
	Spatio-temporal attention-based model [3]	Pose + RGB	91.7	
	Convolutional Relational Machine [6]	RGB + Optical flow	93.0	
	Actor Transformers [7]	Pose + Optical flow	94.4	
	Group Interaction Relational Network [12]	Pose	88.4	
	Actor Transformers [7]	Pose	92.3	
	POGARS ALL	Pose	93.2	
	Group Interaction Relational Network [12]	Pose + Ball tracklets	92.2	
	POGARS with ball	Pose + Ball tracklets	93.9	

The entry with the bold letters (Actor Transformers) is the current state-of-art accuracy on the Volleyball dataset. The entry with the bold italic letters (POGARS All and POGARS with ball) is accuracy on volleyball dataset using our approach

of optical flow ([5,6]), and pose information-based networks ([12]).

To our knowledge, the Actor Transformer model [7] is the only group activity recognition model that achieves higher accuracy (1.2% improvement) than POGARS on the volleyball dataset. [7] makes use of a computationally expensive 3D CNN-based architecture which uses optical flow and pose information as input modalities. Although we recognize that utilizing different input modalities can improve benchmark accuracy results, they require more resource-intensive models to process.

We ran an additional experiment to evaluate the effectiveness of employing tracklets of related objects for group activity recognition by using ball tracklets as an additional input modality for POGARS. Our results in Table 3 show that POGARS is able to take advantage of the ball tracklets to improve classification accuracy by 0.7 percentage points. POGARS with ball tracking outperforms Group Interaction Relational Network [12], the only other existing approach which makes use of ball trajectories.

Figure 8 contains a confusion matrix summarizing correct and incorrect group activity label predictions. For the most part, POGARS is capable of providing accurate predictions for all 8 classes in the volleyball dataset. It can be seen that the proposed model will occasionally confuse "*pass*" and "*set*" activities, which can be explained by the fact that both of these activities are defined by the ball being passed

l-winpoint	97	3	0	0	0	2	0	0
r-winpoint	4	82	0	0	0	1	0	0
r-set	0	0	159	1	2	18	10	2
l-pass	0	0	0	207	1	3	3	12
l-spike	0	0	0	2	168	1	1	7
r-pass	1	0	8	2	4	192	3	0
r-spike	0	0	2	4	2	4	161	0
l-set	1	0	1	11	9	1	1	144
	l-winpoint	r-winpoint	r-set	l-pass	l-spike	r-pass	r-spike	l-set

Fig.8 Confusion matrix obtained using best performing POGARS for the volleyball dataset

between teammates and hence often involve similar player configurations.

# **5** Conclusions

We propose POGARS, a novel architecture for identifying group activities. It does this by learning spatial and temporal dynamics of individuals and fusing learned features to classify the group activity. Using only tracked pose as input, POGARS is able to achieve highly competitive results on the Volleyball dataset with a very important benefit of generalizing the model to testing conditions that differ significantly from training conditions. We also demonstrated that our method's temporal attention mechanism is capable of identifying the specific importance of each frame in the video clip, while the spatial attention mechanism provides a person-wise importance score based on their involvement for the particular collective activity. For future work, we plan to explore solving the problem of temporal action localization in untrimmed video using only tracked pose.

**Funding** Open Access funding enabled and organized by CAUL and its Member Institutions.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecomm ons.org/licenses/by/4.0/.

# References

- Ibrahim, M.S., Muralidharan, S., Deng, Z., Vahdat, A., Mori, G.: A hierarchical deep temporal model for group activity recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016). https://doi.org/10.1109/cvpr.2016.217
- Ramanathan, V., Huang, J., Abu-El-Haija, S., Gorban, A., Murphy, K., Fei-Fei, L.: Detecting events and key actors in multi-person videos. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016). https://doi.org/10.1109/cvpr.2016. 332
- Lu, L., Di, H., Lu, Y., Zhang, L., Wang, S.: Spatio-temporal attention mechanisms based model for collective activity recognition. Signal Process. Image Commun. 74, 162–174 (2019)
- Bagautdinov, T., Alahi, A., Fleuret, F., Fua, P., Savarese, S.: Social scene understanding: End-to-end multi-person action localization and collective activity recognition. 2017 IEEE Conference on computer vision and pattern recognition (CVPR) (2017). https://doi. org/10.1109/cvpr.2017.365
- Azar, S.M., Atigh, M.G., Nickabadi, A.: A multi-stream convolutional neural network framework for group activity recognition. arXiv:1812.10328 (2018)
- Azar, S.M., Atigh, M.G., Nickabadi, A., Alahi, A.: Convolutional relational machine for group activity recognition. arXiv:1904.03308 (2019)
- Gavrilyuk, K., Sanford, R., Javan, M., Snoek, C.G.: Actortransformers for group activity recognition. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition, pp. 839–848 (2020)
- Baradel, F., Wolf, C., Mille, J.: Pose-conditioned spatio-temporal attention for human action recognition. arXiv:1703.10106 (2017)

- Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. 2017 IEEE Conference on computer vision and pattern recognition (CVPR) (2017). https://doi.org/10. 1109/cvpr.2017.502
- Mehrasa, N., Zhong, Y., Tung, F., Bornn, L., Mori, G.: Learning person trajectory representations for team activity analysis. arXiv:1706.00893 (2017)
- Zhu, J., Zou, W., Xu, L., Hu, Y., Zhu, Z., Chang, M., Huang, J., Huang, G., Du, D.: Action machine: rethinking action recognition in trimmed videos. arXiv:1812.05770 (2018)
- 12. Perez, M., Liu, J., Kot, A.C.: Skeleton-based relational reasoning for group activity analysis (2020)
- Ibrahim, M.S., Muralidharan, S., Deng, Z., Vahdat, A., Mori, G.: A hierarchical deep temporal model for group activity recognition. arXiv:1511.06040 (2015)
- Lan, T., Wang, Y., Yang, W., Robinovitch, S.N., Mori, G.: Discriminative latent models for recognizing contextual group activities. IEEE Trans. Pattern Anal. Mach. Intell. 34(8), 1549–1562 (2012). https://doi.org/10.1109/TPAMI.2011.228
- Amer, M.R., Lei, P., Todorovic, S.: Hirf: Hierarchical random field for collective activity recognition in videos. In: European Conference on computer vision, pp. 572–585 (2014). Springer
- Zhu, Y., Nayak, N.M., Roy-Chowdhury, A.K.: Context-aware modeling and recognition of activities in video. In: 2013 IEEE conference on computer vision and pattern recognition, pp. 2491– 2498 (2013). https://doi.org/10.1109/CVPR.2013.322
- Choi, W., Shahid, K., Savarese, S.: What are they doing? : Collective activity classification using spatio-temporal relationship among people. 2009 IEEE 12th International Conference on computer vision workshops, ICCV Workshops, 1282–1289 (2009)
- Belongie, S., Malik, J., Puzicha, J.: Shape matching and object recognition using shape contexts. IEEE Trans. Pattern Anal. Mach. Intell. (2002). https://doi.org/10.1109/34.993558
- Wu, D., Sharma, N., Blumenstein, M.: Recent advances in video-based human action recognition using deep learning: A review. In: 2017 International joint conference on neural networks (IJCNN), pp. 2865–2872 (2017). https://doi.org/10.1109/IJCNN. 2017.7966210
- Deng, Z., Vahdat, A., Hu, H., Mori, G.: Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition. 2016 IEEE Conference on computer vision and pattern recognition (CVPR) (2016). https://doi.org/10. 1109/cvpr.2016.516
- Tsunoda, T., Komori, Y., Matsugu, M., Harada, T.: Football action recognition using hierarchical lstm, pp. 155–163 (2017). https:// doi.org/10.1109/CVPRW.2017.25
- Qi, M., Wang, Y., Qin, J., Li, A., Luo, J., Van Gool, L.: stagnet: an attentive semantic rnn for group activity and individual action recognition. IEEE Trans. Circ. Syst. Video Technol. PP, 1–1 (2019). https://doi.org/10.1109/TCSVT.2019.2894161
- Wu, J., Wang, L., Wang, L., Guo, J., Wu, G.: Learning actor relation graphs for group activity recognition. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition, pp. 9964–9974 (2019)
- Liu, J., Shahroudy, A., Xu, D., Wang, G.: Spatio-temporal lstm with trust gates for 3d human action recognition. arXiv:1607.07043 (2016)
- Wang, H., Wang, L.: Modeling temporal dynamics and spatial configurations of actions using two-stream recurrent neural networks. arXiv:1704.02581 (2017)
- Mnih, V., Heess, N., Graves, A., Kavukcuoglu, K.: Recurrent models of visual attention. arXiv:1406.6247 (2014)
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A.C., Salakhutdinov, R., Zemel, R.S., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. arXiv:1502.03044 (2015)

- Tang, Y., Zhang, P., Hu, J., Zheng, W.: Latent embeddings for collective activity recognition. In: 2017 14th IEEE International Conference on advanced video and signal based surveillance (AVSS), pp. 1–6 (2017). https://doi.org/10.1109/AVSS.2017.8078522
- Varona, X., Gonzàlez, J., Roca, F.X., Villanueva, J.J.: Automatic selection of keyframes for activity recognition. In: International Conference on articulated motion and deformable objects, pp. 173– 181 (2000). Springer
- Raptis, M., Sigal, L.: Poselet key-framing: A model for human activity recognition. In: 2013 IEEE Conference on computer vision and pattern recognition, pp. 2650–2657 (2013). https://doi.org/10. 1109/CVPR.2013.342
- Sendo, K., Ukita, N.: Heatmapping of people involved in group activities. In: 2019 16th International Conference on machine vision applications (MVA), pp. 1–6 (2019). https://doi.org/10. 23919/MVA.2019.8757971
- Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. Lecture Notes in Computer Science, 483– 499 (2016)
- Shotton, J., Fitzgibbon, A., Cook, M., Sharp, T., Finocchio, M., Moore, R., Kipman, A., Blake, A.: Real-time human pose recognition in parts from single depth images. In: CVPR 2011, pp. 1297–1304 (2011). https://doi.org/10.1109/CVPR.2011.5995316
- Cao, Z., Simon, T., Wei, S., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. arXiv:1611.08050 (2016)
- He, K., Gkioxari, G., Dollar, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE International Conference on computer vision (ICCV) (2017)
- Danelljan, M., Häger, G., Khan, F.: Accurate scale estimation for robust visual tracking. British machine vision conference, 1–11 (2014)
- Yan, X., Gilani, S.Z., Qin, H., Feng, M., Zhang, L., Mian, A.S.: Deep keyframe detection in human action videos. arXiv:1804.10021 (2018)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. arXiv:1706.03762 (2017)
- Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2d human pose estimation: New benchmark and state of the art analysis. In: 2014 IEEE Conference on computer vision and pattern recognition, pp. 3686–3693 (2014). https://doi.org/10.1109/CVPR.2014.471

- Kingma, D., Ba, J.: Adam: A method for stochastic optimization. International Conference on learning representations (2014)
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. In: Advances in Neural Information Processing Systems, pp. 8024– 8035 (2019)
- Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. 2015 IEEE International Conference on Computer Vision (ICCV) (2015). https://doi.org/10.1109/iccv.2015.510
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
- Kreiss, S., Bertoni, L., Alahi, A.: Openpifpaf: composite fields for semantic keypoint detection and spatio-temporal association. IEEE Trans. Intell. Transp. Syst. (2021). https://doi.org/10.1109/ TITS.2021.3124981
- Shu, T., Todorovic, S., Zhu, S.-C.: Cern: Confidence-energy recurrent network for group activity recognition. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017). https://doi.org/10.1109/cvpr.2017.453
- Wu, J., Wang, L., Wang, L., Guo, J., Wu, G.: Learning actor relation graphs for group activity recognition. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition, pp. 9964–9974 (2019)
- 47. Ehsanpour, M., Abedin, A., Saleh, F., Shi, J.: Joint learning of social groups, individuals action and sub-group activities in videos
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. 2016 IEEE Conference on computer vision and pattern recognition (CVPR) (2016). https://doi.org/10.1109/cvpr. 2016.90
- Yosinski, J., Clune, J., Bengio, Y., Lipson, H.: How transferable are features in deep neural networks? arXiv:1411.1792 (2014)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.