

Point3D: tracking actions as moving points with 3D CNNs

Shentong Mo*¹
shentonm@andrew.cmu.edu

Jingfei Xia*¹
jingfeix@andrew.cmu.edu

Xiaoqing Tan^{1,2}
xit31@pitt.edu

Bhiksha Raj^{†1}
bhiksha@cs.cmu.edu

¹ Carnegie Mellon University
Pittsburgh, PA 15213, United States

² University of Pittsburgh
Pittsburgh, PA 15260, United States

Abstract

Spatio-temporal action recognition has been a challenging task that involves detecting where and when actions occur. Current state-of-the-art action detectors are mostly anchor-based, requiring sensitive anchor designs and huge computations due to calculating large numbers of anchor boxes. Motivated by nascent anchor-free approaches, we propose **Point3D**, a flexible and computationally efficient network with high precision for spatio-temporal action recognition. Our Point3D consists of a *Point Head* for action localization and a *3D Head* for action classification. Firstly, Point Head is used to track center points and knot key points of humans to localize the bounding box of an action. These location features are then piped into a time-wise attention to learn long-range dependencies across frames. The 3D Head is later deployed for the final action classification. Our Point3D achieves state-of-the-art performance on the JHMDB, UCF101-24, and AVA benchmarks in terms of frame-mAP and video-mAP. Comprehensive ablation studies also demonstrate the effectiveness of each module proposed in our Point3D.

1 Introduction

Spatio-temporal action recognition has attracted much attention in the field of computer vision. It aims to locate and detect action instances of interest in a video in both space and time. The applications can range from human-computer interaction [23, 28], video surveillance [21, 26], to social activity recognition [9, 4]. Videos, unlike images, are time series of images that consist of both spatial components and temporal components. Hence, spatio-temporal action recognition requires both spatial information from key frames as well as temporal information from their previous frames [19].

Early attempts for spatio-temporal action recognition typically apply an object detector at each frame of a video clip independently and then link those frame-wise detection results to generate action tubes [12, 27, 34, 41]. However, those methods fail to consider temporal features across frames and are cumbersome to deploy in a real-time setting. Later,

*These authors contributed equally to this work. [†]Corresponding author.




Figure 1: **Motivation illustration.** Our proposed anchor-free network Point3D achieves high action detection precision by tracking the center point (red dot) and key points (green dots) with 3D-CNNs. The example figures are from the JHMDB dataset.

approaches [15, 17, 19, 33, 35, 45, 48] aim to capture temporal information in videos by conducting clip-level action recognition. These tubelet detection methods, however, are anchor-based detectors that heavily rely on a large number of pre-defined anchor boxes as candidates for calculating bounding boxes. They can be highly sensitive and computationally inefficient as they exhaust all potential actor locations in an image and perform classification for each location [22].

Recently, anchor-free approaches have gained increasing popularity in the computer vision community for their relatively simple network structures and computational efficiency [6, 20, 22, 39, 49, 50]. By eliminating the anchor boxes, anchor-free methods avoid the intense computation as well as hyper-parameters tuning related to anchor boxes, resulting in a more robust performance [22]. However, there has been little work on spatio-temporal action recognition under this framework.

Motivated by the nascent anchor-free framework, we propose a flexible and computationally efficient end-to-end detector called Point3D for the task of spatio-temporal action recognition. Our motivation is illustrated in Figure 1. Specifically, our Point3D performs the task of spatio-temporal action recognition by solving the localization and classification problems separately. Point3D consists of a *Point Head* for action localization and a *3D Head* for action classification. Specifically, the Point Head is used to track center points and knot key points of humans to localize the bounding box of an action. These location features are then piped into a time-wise attention mechanism to learn the dependencies between frames in a clip. The 3D Head is later deployed for the final action classification. We conduct extensive experiments on the JHMDB and UCF101-24 datasets as well as ablation studies to demonstrate the effectiveness of each part of our network.

The major contributions of this paper are summarized as follows: **1.** We present Point3D, an anchor-free end-to-end architecture with high precision for spatio-temporal action recognition. **2.** Our Point3D is the first of its kind to incorporate center points and knot key points together to localize the bounding box of an action without any anchors. **3.** We apply a time-wise attention mechanism to learn the temporal dependencies of each frame in an input clip to boost the performance. **4.** Our Point3D achieves competitive performance on the JHMDB, UCF101-24 and AVA datasets and extensive ablation studies show each part of our Point3D.

2 Related Works

2.1 Anchor-based Object Detection

R-CNN [11], Fast R-CNN, [10] and Faster R-CNN [52] are three of the pioneering works to introduce CNNs into the field of object detection. These three algorithms all fall into the

category of two-stage object detectors that first generate regions of interest and then pipe those region proposals for object classification and bounding-box regression. The need for a two-stage network was eliminated by the YOLO architecture [50] where the task of object detection is regarded as a simple regression problem by taking an input image and learning its class probabilities and bounding box coordinates. Later improvements [0, 29, 50] in YOLO involve optimizing small object localization, anchor boxes, and network structure.

2.2 Anchor-free Object Detection

Recent advances in anchor-free object detection show surprisingly competitive performance without the need of computing expensive anchors [20, 46, 49, 50]. CornerNet [20] detects an object bounding box as a pair of key points using a single convolution neural network. CenterNet [49] detects an object as a point in an image with a focus on modeling the center of a box as an object and uses this predicted center to get the bounding box coordinates. RepPoints [46] instead uses multiple points to automatically learn the spatial extent of an object. CenterTrack [50] tracks objects from frame to frame by using the displacement of a point between adjacent frames.

2.3 Spatio-temporal Action Recognition

Most spatio-temporal action detectors build their work in videos based on object detectors [9, 13, 15, 17, 22, 37, 40, 42]. Traditional anchor-based approaches conduct frame-level action detection by combining independent frame-level detection with linking algorithms to generate final tubes [12, 27, 34, 40, 41]. More recent works place more focus on incorporating temporal information to better the recognition of actions in videos [15, 17, 19, 24, 33, 35, 45, 48]. Typically, 3D CNN [6] and LSTM [14] are used to extract temporal features. AIA [38] also applies memory mechanism to model long-term interaction. However, many of these works apply a two-stage architecture, where proposals are first generated by RPN and then the tasks of action classification and localization are performed accordingly [15, 27, 43]. Because the proposal network and the classification network are trained separately, the results may not be optimal. Plus, as the RPN only focuses on spatial information of actors in each frame, relationships between frames are largely neglected, failing to incorporate important temporal information for action classification and localization.

More recently, an end-to-end detector called YOWO [19] has been proposed which uses a 2D-CNN branch and a 3D-CNN [18] branch to extract spatial and temporal features respectively on video data. YOWO is flexible in that its 2D-CNN and 3D-CNN branches can be replaced by any CNN architectures of interest [19]. Inspired by anchor-free object detectors, MOC Detector [22] incorporates movement information along with center detection and spatial extent detection by treating an action instance as a trajectory of moving points. However, limited works have been proposed in the field of spatio-temporal action localization with anchor-free architectures. In this work, we propose an anchor-free end-to-end network, termed as Point3D, motivated jointly by the conceptually simple CenterNet [49], the recent excellent performance of 3D-CNNs [18], as well as the flexible two-branch YOWO [19]. Besides, our Point3D can be trained in an end-to-end manner, compactly incorporating both spatial and temporal information of a video clip.

3 Method

3.1 Overview

In this work, we aim to solve the problems of localization of action bounding boxes in a video clip together with classification of the action categories. To address this problem,




Figure 2: **Point3D network overview.** An input clip is fed into Feature Extractor to obtain features of each frame. These features will then be used to solve the task of action location and classification separately. Point Head is applied to these features on each frame to track center points and knot key points of humans for action localization, where those features are piped into Time-wise Attention and a 3D Head for action classification.

we propose an anchor-free architecture called Point3D, that operates by tracking actions as moving points and achieves better performance with the help of 3D-CNNs.

As illustrated in Figure 2, Point3D consists of four components: (1) *Feature Extractor*: an input clip is fed into Feature Extractor to obtain spatial features of each frame. (2) *Point Head*: Point Head is applied for the localization task to track center points and knot key points of actors on each frame in a video clip. (3) *Time-wise Attention*: Time-wise Attention is designed to aggregate the correlation and continuity across frames for incorporating more context information. (4) *3D Head*: 3D Head with a 3D-CNN backbone is applied to improve the performance of action classification. Our Point3D is trained in an end-to-end manner, such that these components collaboratively work together to generate a more stable and accurate recognition in a video. We describe the technical details in the following sections.

3.2 Feature Extractor

Assume an input frame as $I \in \mathbb{R}^{3 \times H \times W}$, where H and W are the height and width of the frame. There are N actors in each frame and there are T frames in a clip. We implement a widely-used network in pose estimation, the Hourglass architecture [25] as our backbone for feature extraction. Following [49], the input frames are first downsampled by four times, and then they are passed through a two-stacked Hourglass structure to extract features for future localization and classification tasks. We denote the generated features of each frame to be of $\mathbf{f} \in \mathbb{R}^{C \times \frac{H}{R} \times \frac{W}{R}}$ where C denotes number of channels and is set as 128, and R denotes stride factor and is set as 4 following [22, 49].

3.3 Point Head

Point Head is applied for the localization task to track center points and knot key points of actors on each frame in a video clip. Specifically, our Point Head consists of two main detectors, a Center-Point (CP) and a Knot-Point (KP) detector, as shown in Figure 3.

Center-Point (CP) detector. In the CP detector, there are three types of outputs. The first type is a *center point heatmap*, denoted as $\mathbf{O}_h^{\text{CP}} \in \mathbb{R}^{1 \times \frac{H}{R} \times \frac{W}{R}}$, for localizing the center point of an actor in the current frame. Following [2, 20, 49], we apply a Gaussian kernel on the ground-truths of generated heatmaps. We define the pixel-wise logistic regression loss $\mathcal{L}_h^{\text{CP}}$

between center point heatmaps as:

$$\mathcal{L}_h^{\text{CP}} = -\frac{1}{N} \sum_{j=1}^{1 \times \frac{H}{R} \times \frac{W}{R}} \begin{cases} (1 - \hat{g}_j)^\alpha \log(\hat{g}_j) & \text{if } g_j = 1 \\ (1 - g_j)^\beta (\hat{g}_j)^\alpha \\ \log(1 - \hat{g}_j) & \text{otherwise} \end{cases} \quad (1)$$

where N denotes the number of actors in the frame. g_j and \hat{g}_j are the ground truth and prediction of the center point for pixel j on heatmap \mathbf{O}_h^{CP} . Following [20, 49], we set $\alpha = 2$ and $\beta = 4$ in our experiments. The second type of output is the *shape* prediction of a detected actor, denoted as $\mathbf{O}_s^{\text{CP}} \in \mathbb{R}^{2 \times \frac{H}{R} \times \frac{W}{R}}$, which determines the height and width of the detected actor. The shape prediction loss $\mathcal{L}_s^{\text{CP}}$ is computed as

$$\mathcal{L}_s^{\text{CP}} = \frac{1}{N} \sum_{i=1}^N |\hat{s}_i - s_i| \quad (2)$$

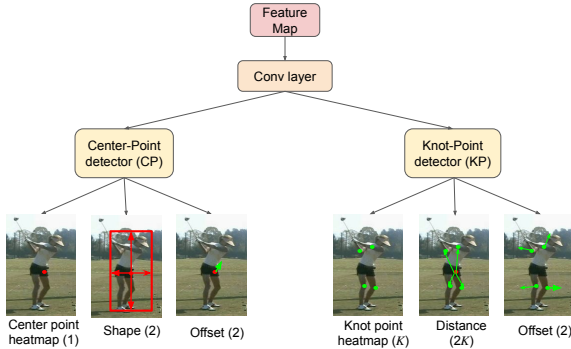


Figure 3: **Point Head structure** including CP detector (Left) and KP detector (Right).

where $s_i = (x_1^{(i)} - x_2^{(i)}, y_1^{(i)} - y_2^{(i)})$ and $(x_1^{(i)}, y_1^{(i)}, x_2^{(i)}, y_2^{(i)})$ is the bounding box for i th actor. We denote its prediction as \hat{s}_i . The third type of output is an *offset* refinement along the height and width of the actor, denoted as $\mathbf{O}_o^{\text{CP}} \in \mathbb{R}^{2 \times \frac{H}{R} \times \frac{W}{R}}$, to minimize the effect caused by strides. When information of the detected actor projects back to the original frame, we apply this offset along x and y axes to make our localization result more accurate. We define the offset loss $\mathcal{L}_o^{\text{CP}}$ inside each actor as

$$\mathcal{L}_o^{\text{CP}} = \frac{1}{N} \sum_i^N \left| \hat{o}_i - \left(\frac{p_i}{R} - \hat{p}_i \right) \right| \quad (3)$$

where \hat{o}_i and \hat{p}_i are the predictions of the offset and the center point for i th actor. We denote p_i as the ground-truth of the center point and R as the stride size.

As a result, the CP detector loss $\mathcal{L}_{loc}^{\text{CP}}$ is given by

$$\mathcal{L}_{loc}^{\text{CP}} = \lambda_h^{\text{CP}} \mathcal{L}_h^{\text{CP}} + \lambda_s^{\text{CP}} \mathcal{L}_s^{\text{CP}} + \lambda_o^{\text{CP}} \mathcal{L}_o^{\text{CP}}, \quad (4)$$

where $\lambda_h^{\text{CP}}, \lambda_s^{\text{CP}}, \lambda_o^{\text{CP}}$ denote the weight hyper-parameter for the heatmap, shape and offset loss. We set $\lambda_h^{\text{CP}} = 1, \lambda_s^{\text{CP}} = 0.1, \lambda_o^{\text{CP}} = 1$ in our all experiments.

Knot-Point (KP) detector. To make the action localization more accurate, we generate three types of outputs from the KP detector. The first output includes K -dimensional *knot point heatmap*, denoted as $\mathbf{O}_h^{\text{KP}} \in \mathbb{R}^{K \times \frac{H}{R} \times \frac{W}{R}}$, which localizes K knot points of the detected actor. The second output is K *distances* of knot points from the center point along the x and y axes, denoted as $\mathbf{O}_d^{\text{KP}} \in \mathbb{R}^{2K \times \frac{H}{R} \times \frac{W}{R}}$. The third output is an *offset* refinement along the height and width of each key point, denoted as $\mathbf{O}_o^{\text{KP}} \in \mathbb{R}^{2 \times \frac{H}{R} \times \frac{W}{R}}$. For the KP detector losses $\mathcal{L}_h^{\text{KP}}$, we have similar loss calculation as the CP detector and the only difference is that we need to add

the loss of K points together. The the KP detector loss \mathcal{L}_{loc}^{KP} is computed as

$$\mathcal{L}_{loc}^{KP} = \lambda_h^{KP} L_h^{KP} + \lambda_d^{KP} L_d^{KP} + \lambda_o^{KP} L_o^{KP} \quad (5)$$

where λ_h^{KP} , λ_d^{KP} and λ_o^{KP} denote the weight hyper-parameter for the knot point heatmap, distance and offset loss. We set $\lambda_h^{KP} = 1, \lambda_d^{KP} = 1, \lambda_o^{KP} = 1$ in our all experiments.




Figure 4: Time-wise Attention (TWA) mechanism for modeling dependencies across frames.

3.4 Time-wise Attention (TWA)

In order to capture the long-range dependencies across frames, we introduce a Time-wise Attention module in our Point3D, as shown in Figure 4. Specifically, we first stack features $\mathbf{f} \in \mathbb{R}^{C \times \frac{H}{R} \times \frac{W}{R}}$ of T timestamps generated from Feature Extractor together as $\mathbf{F} \in \mathbb{R}^{T \times C \times \frac{H}{R} \times \frac{W}{R}}$ before feeding them into TWA module. Then, we reshape \mathbf{F} into $\tilde{\mathbf{F}} \in \mathbb{R}^{T \times (C \times \frac{H}{R} \times \frac{W}{R})}$. We take the matrix multiplication between this reshaped matrix $\tilde{\mathbf{F}}$ and its transpose to get the Gram matrix \mathbf{G} , which indicates the correlations among each time step. The Gram matrix is then passed through a softmax layer and get the time-wise attention M , where each entry inside M measures how one frame is related to another. To avoid gradient vanishing on the original feature map, $\tilde{\mathbf{F}}$ is multiplied by the Gram matrix \mathbf{G} to generate matrix $\hat{\mathbf{S}}$. Finally, we reshape $\hat{\mathbf{S}}$ back to the shape of original matrix as \mathbf{S} and add the original matrix \mathbf{F} to get output \mathbf{Y} . In this way, temporal information between frames has been enhanced to help the 3D Head for later action classification.




Figure 5: **Pipeline of 3D head.** The output from TWA are fed into a 3D CNN backbone to extract the features, then we input these features to two convolutional layers and one fully-connected layer to do the class prediction. Finally, a cross entropy loss is applied to the output for better classification.

3.5 3D Head

To improve the performance of action classification, we introduce 3D CNNs in our 3D Head in Figure 5. Concretely, we first feed the output \mathbf{Y} from TWA into a 3D backbone as the feature extractor for its robust performance [44]. Then, we pipe these generated 3D features into two convolutional layers followed by a fully-connected layer and obtain action label index. During the training process, cross entropy loss is used to train the 3D Head. Assume D is the one-hot encoding vector of ground truth class label. We define the classification loss \mathcal{L}_{cls} on the output of 3D Head as

$$\mathcal{L}_{cls} = -\log(\hat{d}) \quad (6)$$

where \hat{d} be the softmax output for the correct class entry in D .

Putting localization and classification loss together, the overall loss $\mathcal{L}_{overall}$ of our Point3D is simply computed as

$$\mathcal{L}_{overall} = \lambda_{loc} \mathcal{L}_{loc} + \lambda_{cls} \mathcal{L}_{cls}, \quad (7)$$

where $\lambda_{loc}, \lambda_{cls}$ denote the weight hyper-parameter of localization and classification loss. In order to explore how much each loss affects the final performance of our Point3D, we perform extensive experiments on each parameter in the supplementary.

Tubelet linking. After obtaining the bounding box on each frame, we need to link each frame together into the video clip. We use the linking algorithm from [12] and compute scores with the definition in [19]. After all the linking scores are computed, the Viterbi algorithm [8] is applied to find the optimal action tube.

4 Experiments

4.1 Experiments Settings

Datasets and Metrics. We perform experiments on the UCF101-24 [36], JHMDB [16], and AVA [13] datasets. UCF101-24 consists of 3207 temporally untrimmed videos from 24 sports classes [36]. We report the action detection performance for the first split only following the common settings as in [17, 19, 22, 27]. JHMDB is a subset of the HMDB-51 dataset and consists of 928 short videos with 21 action categories in daily life [16]. Each video of JHMDB is well trimmed and has a single action instance across all the frames. We report results averaged over three splits following the common settings as in [17, 19, 22, 27]. Following state-of-the-art methods [12, 17, 19, 22, 41], we adopt frame mAP and video mAP to evaluate action detection accuracy. AVA is a more challenging dataset consisting of 211k training and 57k validation video segments, where each actor in the key frame is labeled at 1 FPS and with one bounding box and multiple classes from 80 atomic action categories. Following the common protocol [13, 19, 38], we report frame mAP performance with an IoU threshold of 0.5 over the top 60 action classes on AVA v2.1 and v2.2 benchmarks.

Implementation Details. We employ the Hourglass-104 [25] as the feature extractor network, closely following the same setting as the current state-of-the-art methods [19, 22] for a fair comparison. The frame is resized to 256×256 for effective computations. The spatial down-sampling ratio R is set to 4 and the resulted feature map size is 64×64 . The number of input frames T is 16. We use Adam with a learning rate of $5e-4$ to optimize the feature extractor and the Point Head. The 3D Head is optimized by using a SGD with a learning rate of $1e-3$. The overall network is trained in an end-to-end manner. The learning rate adjusts to convergence on the validation set and decreases by a factor of 10 when performance saturates. For a fair comparison, we closely follow previous work [22, 38, 43] to train 12 epochs on UCF101-24 [36], 20 epochs on JHMDB [16], and 15 epochs on AVA [13]. The total training time on 8 Tesla V100 GPUs is 153 hours.

4.2 Comparison with State of the Arts

JHMDB. We compare our method to the current state-of-the-art methods, including 2D backbone methods (SSD+ACT [17], TACNet [35], Dance+Flow [43], MOC [22]) and 3D backbone methods (C3D [19], I3D [13], S3D-G [37]). Table 1 shows the results on the JHMDB benchmark. Our method outperforms the state-of-the-art methods by a large margin in terms of both frame-mAP and three video-mAP metrics. Except for video-mAP with a threshold of 0.75 where we achieve a comparable result as MOC [22], our Point3D outperforms all other methods.

UCF101-24. We compare our method with the state-of-the-art methods in Table 1. Our Point3D outperforms the current state-of-the-art methods by a significant margin over all metrics. This direct comparisons demonstrate the improvement using Point3D. Introducing the 3D Head in Point3D helps to outperform the 2D detectors [17, 22, 27, 34, 35, 43] by a

Table 1: Comparison with state-of-the-art methods on the JHMDB and UCF101-24 benchmark. Bold and underlined numbers denote the first and second place.

Method	JHMDB					UCF101-24					
	Frame-mAP 0.5(%)	0.2	0.5	0.75	0.5:0.95	Frame-mAP 0.5(%)	0.1	0.2	0.5	0.75	0.5:0.95
<i>2D backbone:</i>											
MR-TS R-CNN [14]	58.5	74.3	73.1	-	-	39.9	50.4	42.3	-	-	-
A+AF [15]	-	73.8	72	44.5	41.6	-	-	73.5	46.3	15.0	20.4
SSD+ACT [16]	65.7	74.2	73.7	52.1	44.8	69.5	-	76.5	49.2	19.7	23.4
TACNet [17]	65.5	74.1	73.4	52.5	44.8	72.1	-	77.5	52.9	21.8	24.1
Dance+Flow [18]	-	-	74.7	53.3	45.0	-	-	78.5	50.3	22.2	24.5
MOC [19]	70.8	77.3	77.2	71.7	59.1	78.0	-	<u>82.8</u>	53.8	<u>29.6</u>	<u>28.3</u>
<i>3D backbone:</i>											
C3D [20]	61.3	78.4	76.9	-	-	41.4	-	47.1	-	-	-
I3D [21]	73.3	-	78.6	-	-	76.3	-	-	59.9	-	-
S3D-G [22]	<u>77.9</u>	-	80.1	-	-	-	-	-	-	-	-
<i>2D+3D backbone:</i>											
YOWO [23]	74.4	<u>87.8</u>	<u>85.7</u>	58.1	-	<u>80.4</u>	82.5	75.8	48.8	-	-
Point3D (ours)	79.2	89.1	86.1	<u>71.5</u>	60.9	83.5	85.4	84.5	<u>55.1</u>	33.4	31.8

large margin. We also achieve a better performance than YOWO [23] with the 3D detector with the help of our Point Head that tracks actions as moving points.

AVA. Following previous work [19, 58, 43], we report comparison results on the AVA v2.1 and v2.2 benchmarks in Table 2. As we can see, our Point3D pre-trained on Kinetics-400 [18] outperforms YOWO [23] by 10.3% mAP, which verifies the effectiveness of the proposed anchor-free architecture in spatio-temporal action localization. Typically, our Point3D pre-trained on Kinetics-700 [18] achieves a new state-of-the-art performance on both benchmarks. We also achieve comparable results as AIA [58], although we did not leverage different interactions to improve action detection performance.

Table 2: Comparison with state-of-the-art methods on the AVA v2.1 and v2.2 benchmark.

Method	Pretrain	v2.1	v2.2
<i>2D backbone:</i>			
Relation Graph [14]	Kinetics-400	22.2	-
LFB [15]	Kinetics-400	27.7	-
SlowFast [9]	Kinetics-400	28.2	29.1
<i>3D backbone:</i>			
I3D [21]	Kinetics-400	15.6	-
S3D [22]	Kinetics-400	17.4	-
VAT [9]	Kinetics-400	25.0	-
C-A RCNN [16]	Kinetics-400	28.0	-
AIA [58]	Kinetics-700	31.2	32.3
<i>2D+3D backbone:</i>			
YOWO [23]	Kinetics-400	19.2	20.2
Point3D (ours)	Kinetics-400	29.5	30.6
Point3D (ours)	Kinetics-700	31.6	32.8

Visualizations. In Figure 6, we visualize some qualitative examples of action detection on JHMDB and UCF101-24 datasets. In general, our Point3D architecture exhibits a decent job at localizing actions in videos. As can be seen in the heatmaps on the top row, our Point3D get rid of the anchor to localize and classify the action by tracking the spatial and temporal changes of key points in the frames. We defer more visualization examples to the supplementary materials. We also visualize failure cases in the last row in Figure 6. Our Point3D sometimes makes some false positive classifications at initial frames since it is hard to recognize the action before it happens.

4.3 Ablation Study

In this section, we explore extensive ablation studies on each part of our Point3D, including the Point Head, the backbone and the input of the 3D Head, the temporal length of input




Figure 6: Visualization results. The **first row** denotes the heatmaps generated from our KP detector. The **blue** bounding boxes are ground truth while **red** and **green** are true and false positive detections, respectively. Zoom in for a better view.

clips, and the hyper-parameters for weighting the overall loss. Unless specified, we conduct all ablation studies on the JHMDB benchmark.

Each component of Point3D. In Table 3, we explore each part of our Point3D, including CP detector, KP detector, 3D Head, and TWA. Introducing KP in the CP and CP+3D Head increase the frame-mAP by 0.4% and 1.3%, respectively. Adding 3D Head to CP and CP+KP improve the frame-mAP by 6.7% and 7.6%, respectively. This demonstrates the advantage of 3D Head in action classification. Adding TWA to CP+KP and CP+KP+3D Head further increases the frame-mAP by 2.0% and 1.1%, respectively, which shows the validity of the proposed Time-wise Attention in our Point3D. Similar improvements are also shown in the four video-mAP metrics.

Table 3: Exploration study on each part of our Point3D.

CP	KP	3D Head	TWA	Frame-mAP(%)		Video-mAP(%)		
				0.5	0.2	0.5	0.75	0.5:0.95
✓				70.1	76.5	77.2	69.1	58.3
✓	✓			70.5	77.8	78.3	69.9	58.5
✓	✓		✓	72.5	79.3	80.1	70.4	59.3
✓		✓		76.8	85.6	82.7	70.7	59.8
✓	✓	✓		78.1	88.2	85.6	70.8	60.4
✓	✓	✓	✓	79.2	89.1	86.1	71.5	60.9

Table 4: Exploration study on the temporal length T of input clips.

T	Speed(fps)	Frame-mAP(%)		Video-mAP(%)		
		0.5	0.2	0.5	0.75	0.5:0.95
1	38	69.8	78.5	77.4	63.2	53.8
3	32	72.1	81.4	80.8	66.4	56.7
5	27	76.3	85.2	83.6	69.5	59.6
7	23	78.6	88.5	85.4	70.8	60.3
9	19	78.8	88.7	85.8	71.1	60.5
13	15	78.9	88.9	85.9	71.2	60.7
16	13	79.2	89.1	86.1	71.5	60.9
32	7	79.8	89.6	86.5	71.9	61.4

Temporal Length of Input Clip. In this study, we explore the temporal length T of the input clip in our Point3D. We vary T from 1 to 32 and report the results in Table 4. First, we observe that when $T = 7$, our Point3D outperforms the single-frame detector by a large

margin in terms of frame-mAP and video-mAP metrics. This confirms the importance of temporal information for action recognition, which agrees with the motivation of adding a 3D Head to our Point3D. Second, we note that the detection performance will increase as we vary T from 1 to 32 and the performance increase scope becomes smaller. Third, when $T = 32$, we follow [19] and introduce the long-term feature bank [22] in the input of 3D Head, which brings further improvements on the performance.

5 Conclusion

In this paper, we propose Point3D, a novel anchor-free architecture that is flexible in structure and enjoy good computation efficiency for the task of spatio-temporal action localization. Point3D consists of a *Point Head* for action localization and a *3D Head* for action classification. We conduct extensive experiments on the JHMDB, UCF101-24 and AVA benchmarks where our method achieves the state-of-the-art results in terms of frame-mAP and video-mAP. Comprehensive ablation studies also demonstrate the effectiveness of the Point Head, the 3D Head, and TWA proposed in our Point3D.

Acknowledgement

We thank Pinxu Ren for insightful discussions.

References

- [1] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. YOLOv4: optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020.
- [2] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7291–7299, 2016.
- [3] Claudio Coppola, Diego R Faria, Urbano Nunes, and Nicola Bellotto. Social activity recognition based on probabilistic merging of skeleton features with proximity priors from rgb-d data. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5055–5061, 2016.
- [4] Claudio Coppola, Serhan Cosar, Diego R Faria, and Nicola Bellotto. Social activity recognition on continuous rgb-d video sequences. *International Journal of Social Robotics*, pages 1–15, 2019.
- [5] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6569–6578, 2019.
- [6] Kevin Duarte, Yogesh Rawat, and Mubarak Shah. Videocapsulenet: A simplified network for action detection. In *Proceedings of the Advances in Neural Information Processing Systems*, 2018.

- [7] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6202–6211, 2019.
- [8] G David Forney. The viterbi algorithm. *Proceedings of the the IEEE*, 61(3):268–278, 1973.
- [9] Rohit Girdhar, Joao Carreira, Carl Doersch, and Andrew Zisserman. Video action transformer network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 244–253, 2019.
- [10] Ross Girshick. Fast R-CNN. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1440–1448, 2015.
- [11] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 580–587, 2014.
- [12] Georgia Gkioxari and Jitendra Malik. Finding action tubes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 759–768, 2015.
- [13] Chunhui Gu, Chen Sun, David A. Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, and Jitendra Malik. AVA: a video dataset of spatio-temporally localized atomic visual actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6047–6056, 2018.
- [14] Jiawei He, Zhiwei Deng, Mostafa S. Ibrahim, and Greg Mori. Generic tubelet proposals for action localization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 343–351, 2018.
- [15] Rui Hou, Chen Chen, and Mubarak Shah. Tube convolutional neural network (t-cnn) for action detection in videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5822–5831, 2017.
- [16] Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J. Black. Towards understanding action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3192–3199, 2013.
- [17] Vicky Kalogeiton, Philippe Weinzaepfel, Vittorio Ferrari, and Cordelia Schmid. Action tubelet detector for spatio-temporal action localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4405–4413, 2017.
- [18] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [19] Okan Köpüklü, Xiangyu Wei, and Gerhard Rigoll. You only watch once: A unified cnn architecture for real-time spatiotemporal action localization. *arXiv preprint arXiv:1911.06644*, 2019.

- [20] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European Conference on Computer Vision*, pages 734–750, 2018.
- [21] Marco Leo, Tiziana D’Orazio, and Paolo Spagnolo. Human activity recognition for automatic visual surveillance of wide areas. In *Proceedings of the ACM 2nd international workshop on Video surveillance & sensor networks*, pages 124–130, 2004.
- [22] Yixuan Li, Zixu Wang, Limin Wang, and Gangshan Wu. Actions as moving points. In *Proceedings of the European Conference on Computer Vision*, pages 68–84, 2020.
- [23] Sushmita Mitra and Tinku Acharya. Gesture recognition: A survey. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 37(3):311–324, 2007.
- [24] Shentong Mo, Xiaoqing Tan, Jingfei Xia, and Pinxu Ren. Towards improving spatiotemporal action recognition in videos. *arXiv preprint arXiv:2012.08097*, 2020.
- [25] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Proceedings of the European Conference on Computer Vision*, pages 483–499, 2016.
- [26] Sangmin Oh, Anthony Hoogs, Amitha Perera, Naresh Cuntoor, Chia-Chih Chen, Jong Taek Lee, Saurajit Mukherjee, JK Aggarwal, Hyungtae Lee, Larry Davis, et al. A large-scale benchmark dataset for event recognition in surveillance video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3153–3160, 2011.
- [27] Xiaojiang Peng and Cordelia Schmid. Multi-region two-stream r-cnn for action detection. In *Proceedings of the European Conference on Computer Vision*, pages 744–759, 2016.
- [28] Siddharth S Rautaray and Anupam Agrawal. Vision based hand gesture recognition for human computer interaction: a survey. *Artificial Intelligence Review*, 43(1):1–54, 2015.
- [29] Joseph Redmon and Ali Farhadi. YOLO9000: better, faster, stronger. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6517–6525, 2017.
- [30] Joseph Redmon and Ali Farhadi. YOLOv3: an incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- [31] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016.
- [32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 91–99, 2015.
- [33] Suman Saha, Gurkirt Singh, and Fabio Cuzzolin. Amtnet: Action-micro-tube regression by end-to-end trainable deep architecture. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4414–4423, 2017.

- [34] Gurkirt Singh, Suman Saha, Michael Sapienza, Philip HS Torr, and Fabio Cuzzolin. Online real-time multiple spatiotemporal action localisation and prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3637–3646, 2017.
- [35] Lin Song, Shiwei Zhang, Gang Yu, and Hongbin Sun. TACNet: transition-aware context network for spatio-temporal action detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11987–11995, 2019.
- [36] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv: 1212.0402*, 2012.
- [37] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Kevin Murphy, Rahul Sukthankar, and Cordelia Schmid. Actor-centric relation network. In *Proceedings of the European Conference on Computer Vision*, pages 318–334, 2018.
- [38] Jiajun Tang, Jin Xia, Xinzhi Mu, Bo Pang, and Cewu Lu. Asynchronous interaction aggregation for action detection. In *Proceedings of the European Conference on Computer Vision*, pages 71–87, 2020.
- [39] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9627–9636, 2019.
- [40] Limin Wang, Yu Qiao, Xiaoou Tang, and Luc Van Gool. Actionness estimation using hybrid fully convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2708–2717, 2016.
- [41] Philippe Weinzaepfel, Zaid Harchaoui, and Cordelia Schmid. Learning to track for spatio-temporal action localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3164–3172, 2015.
- [42] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 284–293, 2019.
- [43] Jianchao Wu, Zhanghui Kuang, Limin Wang, Wayne Zhang, and Gangshan Wu. Context-Aware RCNN: a baseline for action detection in videos. In *Proceedings of the European Conference on Computer Vision*, pages 440–456, 2020.
- [44] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5987–5995, 2017.
- [45] Xitong Yang, Xiaodong Yang, Ming-Yu Liu, Fanyi Xiao, Larry S Davis, and Jan Kautz. Step: Spatio-temporal progressive learning for video action detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 264–272, 2019.
- [46] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. RepPoints: Point set representation for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9657–9666, 2019.

- [47] Yubo Zhang, Pavel Tokmakov, Martial Hebert, and Cordelia Schmid. A structured model for action detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9975–9984, 2019.
- [48] Jiaojiao Zhao and Cees GM Snoek. Dance with flow: Two-in-one stream action detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9935–9944, 2019.
- [49] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv: 1904.07850*, 2019.
- [50] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. In *Proceedings of the European Conference on Computer Vision*, pages 474–490, 2020.