

CFAD: Coarse-to-Fine Action Detector for Spatiotemporal Action Localization

Yuxi Li¹, Weiyao Lin^{1,2*}, John See³, Ning Xu⁴
Shugong Xu², Ke Yan⁵, and Cong Yang⁵

¹ Department of Electronic Engineering, Shanghai Jiao Tong University, China

² Institute for Advanced Communication and Data Science, Shanghai University, China

³ Faculty of Computing and Informatics, Multimedia University, Malaysia

⁴ Adobe Research, USA

⁵ Clobotics, China

Abstract Most current pipelines for spatio-temporal action localization connect frame-wise or clip-wise detection results to generate action proposals, where only local information is exploited and the efficiency is hindered by dense per-frame localization. In this paper, we propose Coarse-to-Fine Action Detector (CFAD), an original end-to-end trainable framework for efficient spatio-temporal action localization. The CFAD introduces a new paradigm that first estimates coarse spatio-temporal action tubes from video streams, and then refines the tubes' location based on key timestamps. This concept is implemented by two key components, the Coarse and Refine Modules in our framework. The parameterized modeling of long temporal information in the Coarse Module helps obtain accurate initial tube estimation, while the Refine Module selectively adjusts the tube location under the guidance of key timestamps. Against other methods, the proposed CFAD achieves competitive results on action detection benchmarks of UCF101-24, UCFSports and JHMDB-21 with inference speed that is $3.3\times$ faster than the nearest competitor.

Keywords: Spatiotemporal action detection; Coarse-to-fine paradigm; Parameterized modeling.

1 Introduction

Spatial-temporal action detection is the task of recognizing actions from input videos and localizing them in space and time. In contrast to action recognition or temporal localization, this task is far more complex, requiring both temporal detection along the time span and spatial detection at each frame when the actions occur.

Most existing methods for spatiotemporal action detection [6,19,26,25,29,9,39,38] are implemented in two stages (illustrated in Fig. 1(a)). First, a spatial detector is applied to generate dense action box proposals on each frame. Then, these

* Correspondance author, wylin@sjtu.edu.cn




Figure 1. The comparison between pipelines of detection and linking and our coarse-to-fine framework. (a) workflow of detection and linking method in previous works. (b) Our coarse-to-fine method to detect action tubes. (Best viewed in color.)

frame-level detections are linked together by a certain heuristic algorithm to generate final output, which is a series of boxes or an *action tube*. Nevertheless, since these approaches take a single or stack of frames as input, the information utilized by the detectors is limited within a fixed time interval, hence limiting the representative capacity of the learned features for classification. The similar problem is encountered in the aspect of localization. During training phase, models could be supervised by only a temporal fragment of the tubes, which can output accurate local proposals but may fail to locate entire tubes in a consistent manner. Additionally, IOU-based linking algorithms may result in accumulative localization error when noisy bounding box proposals are produced. Since the transition within action tubes is usually smooth and gradual, we hypothesize that using lesser number of boxes could be adequate to depict the action tube shape. Current pipelines, in their present state, relies heavily on dense per-frame predictions, which are redundant and a hindrance to efficient action detection.

With these considerations, we depart from classic detect-and-link strategies by proposing a new coarse-to-fine action detector (CFAD) that can generate more accurate action tubes with higher efficiency. Unlike previous approaches that detect dense boxes at first, the CFAD (as illustrated in Fig. 1(b)) goes on a progressive approach of estimating at a rougher level before ironing out the details. This strategy first estimates coarser action tubes, and then selectively refine these tubes at key timestamps. The action tubes are generated via two important components in our pipeline: Coarse Module and Refine Module.

The Coarse Module is designed to address the lack of global information and low efficiency in previous detect-and-link paradigm. In a *global* sense, it supervises the tube regression with the full tube shape information. In addition, within this module, a parameterized modeling scheme is introduced to depict action tubes. Instead of predicting large amount of box location at each frame, Coasre Module only predict a few trajectory parameters to describe the tube of various endurance. As a result, this module learns a robust representation that accurately and efficiently characterizes action tube changes.

The Refine Module delves into the *local* context of each tube, to find precise temporal locations that are essential to further improve the estimated action tubes, which in turn, improves overall detection performance and efficiency. To properly refine the action tubes, a labelling algorithm is designed to generate labels that guide the learning of key timestamps selection. By a search scheme,

the original coarse boxes are replaced by the largest scoring box proposals at these temporal locations, which then interpolate the final tube.

In summary, our contributions are three folds. (1) We propose a novel *coarse-to-fine* framework for the task of spatial-temporal action detection, which differs from the conventional paradigm of detect-and-link. Our new pipeline achieves state-of-the-art results on standard benchmarks with inference speed of $3.3\times$ faster than the nearest competitor. (2) Under this framework, we design a novel action tube estimation method based on parametric modeling to fully exploit global supervision signal and handle time variant box coordinates by predicting limited amount of parameters. (3) We also propose a simple yet effective method of predicting an importance score for each sampled frame which is used to select key timestamps for the refinement of output action tubes.

2 Related Works

2.1 Action recognition

Deep learning techniques have shown to be effective and powerful in the classification of still images [8,11,28], and some existing works have extended such schemes to the task of human action recognition in video. Direct extensions attempt to model sequential data with serial or parallel networks. [18,33] combined 2D CNN with a RNN structure to model spatial and temporal relations separately. In [27], the authors found that the involvement of optical flow is beneficial for temporal modeling of actions and thus, proposed a two-stream framework that extracts features from RGB and optical flow data using separate parallel networks; the inference result being the combination of both modalities. In [34], the authors designed a 3D convolution architecture to automatically extract a high dimensional representation for input video. The I3D network [3] further improved the 3D convolution technique by inflating convolution kernels of networks pre-trained on ImageNet (2D) [4] into an efficient 3D form for action recognition. Although these methods achieved good results on video classification benchmarks, they can only make video level predictions and are unable to ascertain the position of actors and the duration of action instances.

2.2 Spatio-temporal action detection

The task of spatio-temporal action detection is more complex than direct classification of videos. It requires both correct categorization and accurate localization of actors during the time interval when the action happens. Gkioxari *et al.* proposed the first pipeline for this task in [6], where R-CNN [5] was applied on each frame to locate actors and classify actions, the results are then linked by viterbi algorithm. Saha *et al.* [26] designed a potential-based temporal trimming algorithm to extend general detection methods to untrimmed video datasets. Following the workflow of these two works, [19,15,9,21] tried learning more discriminative features of action instances with larger spatial or temporal




Figure 2. Overview of the proposed CFAD framework. TPN block denotes the temporal action proposal network. A 3D Conv Head block indicates a cascaded NL-3D Conv structure (“NL” represents the NonLocal Block of [35]). 2D Conv Head block denotes cascaded 2D spatial convolutions. (Best viewed in color)

context, a concept greatly enhanced by [16] through a multi-channel architecture that learns recurrently from tubelet proposals. Some works [29,12] aimed to improve heuristic linking for better localization. Recent works [39,32] proposed innovative two-stream fusion schemes for this task. [38] took a novel route to progressively regressing clip-wise action tubelets and linking them along time. Overall, all these works require temporally dense detections for each video, which is cumbersome. This inefficiency gets worse when optical flow computation is taken into account.

Among the existing works, [38] is the most similar to CFAD with its refinement process. However, our method is different from it from three aspects. Firstly, CFAD estimates coarse level tubes with parametric modeling and global supervision, while [38] relies on per-frame detection. Secondly, our approach does not require further temporal linking or trimming process. Finally, [38] refines the boxes densely for each frame, while our method only refines the box locations at selected key timestamps.

2.3 Weight prediction

Weight prediction is a meta-learning concept where machine learning models are exploited to predict parameters of other structures [1]. For example, the STN [13] utilized deep features to predict affine transformation parameters. [10] used category-specific box parameters to predict an instance weighted mask, while MetaAnchor [37] learned to predict classification and regression functions from both box parameters and data. The Coarse Module of our method is also inspired by such similar mechanisms, where the trajectory parameters are predicted by relevant spatio-temporal structures to depict the tube variation along time. To the best of our knowledge, our approach is the first attempt at exploiting parameterized modeling to handle action tube estimation.

3 Methodology

3.1 Framework

In this section, we introduce the proposed Coarse-to-Fine Action Detector (CFAD) in detail. We first formulate the problem and provide an overview of our approach. Then we discuss more elaborately on the two primary components of CFAD – the Coarse Module for tube estimation and Refine Module for final proposal.

One action tube instance in videos can be formulated as a set, $\mathcal{A} = \{(t_i, b_i) | i = 0, \dots, T_A - 1\}$, where t_i is the timestamp of a certain frame, $b_i = (x_i, y_i, w_i, h_i)$ is the corresponding actor box within this frame, and T_A denotes the total number of bounding boxes in a ground-truth tube. Each tube \mathcal{A} is accompanied with a category label c .

The workflow of CFAD is shown in Fig. 2. Firstly, the input video is resampled to a fixed length T and fed into 3D CNN for spatio-temporal feature extraction. Then the feature is processed a temporal proposal network (TPN) to obtain class-agnostic temporal proposals (t_s, t_e) . t_s is the start timestamp and t_e denotes the end timestamp. In this paper, we instantiate the temporal proposal network by implementing one that is similar to [36], readers can refer to Appendix A for architecture detail. Given the temporal proposal, we uniformly sample N 2D features $\{\mathbf{F}_i | i \in [0, N - 1]\}$ along the time axis within interval (t_s, t_e) , which are sent to Coarse and Refine Module simultaneously. In Coarse Module, these 2D sampled features are used to estimate coarse level action tubes. Next, the estimated tube and the sampled 2D features in the Refine Module are exploited for frame selection and tube refinement at identified key timestamps.

3.2 Coarse tube estimation

We design two convolutional brunches in the Coarse Module to process The sampled 2D features, one branch processes the input features directly and the other branch handles the temporal residual component $\{\mathbf{F}_{i+1} - \mathbf{F}_i | i = 0, N - 2\}$ of the input features, which is output through the “Diff Block” in Fig. 2. We add residual processing since the temporal residual component can provide more time variant information, which is beneficial to discriminate different actions and predict localization changes along time. For each branch, a Non-Local Block [35] is cascaded with a 3D convolution blocks to construct the “3D Conv Head” module in Fig. 2, which aggregates information from both spatial and temporal context. The output of the two branches are fused by element-wise summation and aggregated via temporal average pooling.

To estimate coarse-level action tubes, we adopt a *parameterized modeling* scheme, where we define a coarse-level tube estimation as a parameterized mapping $\hat{\mathcal{A}}(t; \theta) : [0, 1] \rightarrow \mathbb{R}^4$. $\hat{\mathcal{A}}$ tries to predict the coarse spatial location, i.e. $[x(t), y(t), w(t), h(t)]$, given a normalized timestamp t and trajectory parameter θ . The mapping parameters θ are predicted by the deep features from the temporal pooling block. To this end, we slide predefined anchor boxes of different sizes

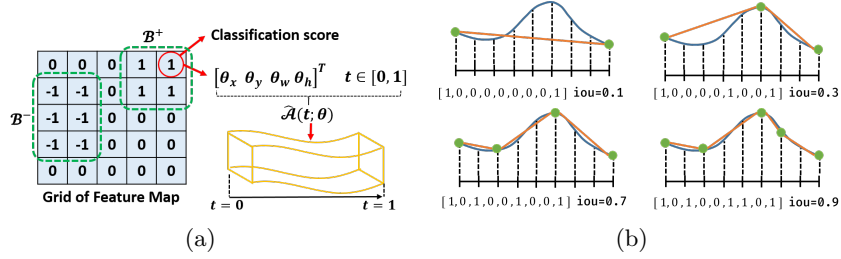


Figure 3. (a). Illustration of coarse tube estimation, where “1” and “-1” symbols denote the positive and negative samples after matching, and “0” are ignored samples. (b) Key timestamps label selection process in the refine module. For ease of simplification, this figure depicts the case of 1-dimensional linear interpolation. The blue curve is the ground-truth, the orange one is the interpolated curve and green nodes represent selected timestamps. (Best viewed in color)

on the 2D output feature map from temporal pooling block to obtain positive samples \mathcal{B}^+ and negative samples \mathcal{B}^- as according to the IOUs between anchors and tubes (illustrated in Fig. 3(a)). For each sample in \mathcal{B}^+ , the network should predict its corresponding classification score and the tube shape parameter θ through an additional 1×1 convolution layer.

Segment-wise matching. To measure the overlap between an anchor box b_a and ground-truth \mathcal{A} , an intuitive idea is to calculate the average value of IOUs between b_a and each boxes belonging to \mathcal{A} . However, since tube shapes may include motion and camera shake, such matching strategy might result in small IOU value and induce the imbalance issue of samples. Hence, we design a segment-wise matching scheme to separate positive and negative samples. To be specific, We take the boxes on first K frames in \mathcal{A} as a valid segment for matching positive anchors, where K is a predefined segment length. We take the segment from the beginning of the tube because we found the final model performance is not sensitive to the segment position, readers can refer to Appendix B for detail analysis. If the average overlap between b_a and these K boxes is larger than a threshold, it is taken as a positive sample. Further, if b_a has high overlap with multiple concurrent tubes, we choose the ground-truth with largest segment IOU as the matched tube. On the other hand, to find negative samples, we split the ground-truth tube \mathcal{A} into $\lfloor \frac{T_A}{K} \rfloor$ segments and compute IOUs between b_a and each segments as discussed above, if the maximum IOU among all segments is still less than a threshold, then it is taken as a negative sample, the intuition behind such design is that negative samples should have low overlap with any boxes in \mathcal{A} .

Parameterized modeling. Generally, any parameterized function that takes a single scalar as input and outputs a 4-dimensional vector can be used as the tube mapping. In this paper, we use the family of high order polynomial functions to model action tube variations along the timestamp. This is because action tubes typically change smoothly and gradually, while polynomial functions are capable enough of describing the patterns of tube shape changes. Therefore, the

instantiation of parameterized estimation function $\hat{\mathcal{A}}(t; \boldsymbol{\theta})$ can be formulated as:

$$\hat{\mathcal{A}}(t; \boldsymbol{\theta}) = [x(t; \boldsymbol{\theta}_x), y(t; \boldsymbol{\theta}_y), w(t; \boldsymbol{\theta}_w), h(t; \boldsymbol{\theta}_h)] = \begin{bmatrix} \boldsymbol{\theta}_x^T \mathbf{t} & \boldsymbol{\theta}_y^T \mathbf{t} & \boldsymbol{\theta}_w^T \mathbf{t} & \boldsymbol{\theta}_h^T \mathbf{t} \end{bmatrix} \quad (1)$$

where the trajectory of each coordinate is regarded as a polynomial curve of order k , the predicted parameter matrix $\boldsymbol{\theta} = [\boldsymbol{\theta}_x, \boldsymbol{\theta}_y, \boldsymbol{\theta}_w, \boldsymbol{\theta}_h]$ of size $(k+1) \times 4$ is composed of the polynomial coefficient for each bounding box coordinates. The vector $\mathbf{t} = [1, t, t^2, \dots, t^k]^T$ contains various orders of current timestamp. To learn features that are invariant to anchor transitions, we do not use $\hat{\mathcal{A}}(t; \boldsymbol{\theta})$ to directly estimate the absolute coordinates, but instead perform estimation of relative coordinates w.r.t matched bounding box b_a following the method of encoding in [23].

During training, the model learns to separate positive and negative samples, to predict the correct action classes and relative coordinates of a coarse tube under the supervision of the loss function in Eq. 2,

$$L_{coarse} = \frac{1}{|\mathcal{B}^+ \cup \mathcal{B}^-|} L_c + \frac{1}{|\mathcal{B}^+|} L_r \quad (2)$$

where $|\cdot|$ denotes the size of the set. L_c is the classification loss in [23] while L_r is the regression loss from the supervision of the whole ground-truth tube:

$$L_r = \frac{1}{T_A} \sum_{b_a \in \mathcal{B}^+} \sum_{(t_i, b_i) \in \mathcal{A}} \left\| \hat{\mathcal{A}}(\hat{t}_i; \boldsymbol{\theta}_a) - \mathbf{enc}(b_i, b_a) \right\|^2 \quad (3)$$

The function $\mathbf{enc}(\cdot, \cdot)$ in Eq. 3 is the same as the encoding function in [23] to encode the 4-dimensional relative offsets from anchor box to ground-truth box. $\boldsymbol{\theta}_a$ is the predicted tube shape parameter associated with anchor b_a . The symbol \hat{t}_i defined in Eq. 4 is the normalized timestamp of ground-truth bounding boxes in tube \mathcal{A} . We normalize the input timestamp before calculating the tube shape in order to avoid value explosion when the polynomial order increases.

$$\hat{t}_i = \frac{t_i - t_0}{t_{T_A-1} - t_0} \quad \forall (t_i, b_i) \in \mathcal{A} \quad (4)$$

3.3 Selective refinement

After the estimated coarse tube $\hat{\mathcal{A}}(t; \boldsymbol{\theta})$ has been generated by the Coarse Module, its location is further refined in the Refine Module. The Refine Module first selects the samples attached with key timestamps for action tube localization, and then refine tube boxes based on these features and guidance of coarse tube.

Key timestamp selection. One simple and intuitive refinement scheme is to observe the tube location at each sampled 2D feature map and then refine the box according to the features within that area. However, when the sample number N increases, such a scheme is costly in computation. Since changes in the action tubes are usually smooth, there is only a limited number of sparsely distributed bounding boxes that are decisive to the shape of tubes. Thus, we design

a selector network in the Refine Module to dynamically sample key timestamps that are most essential for localization.

In our implementation, we perform importance evaluation by squeezing the input 2D sampled features $\{\mathbf{F}_i | i \in [0, N - 1]\}$ with spatial pooling and applying a 1D hourglass network along the time dimension. This outputs an importance score p_i for each sample (shown in Fig. 2). During inference phase, we only take samples that satisfy $p_i \geq \alpha$ as samples of key timestamps and then proceed to refinement.

In the training phase, we heuristically define sets of labels to guide the training of the selector network. Specifically, first the ground-truth action tube \mathcal{A} is uniformly split into $N - 1$ segments along temporal axis with N endpoints. The i -th endpoint is associated with the i -th sampled feature \mathbf{F}_i , and its normalized timestamp is defined as $s_i = i/(N - 1)$. Let the timestamp set be defined as $\mathcal{U} = \{s_i | i = 0, \dots, N - 1\}$ and the key points set as \mathcal{U}_k . We start from $\mathcal{U}_k = \{s_0, s_{N-1}\}$ having the start and end points and gradually append other s_i into \mathcal{U}_k . The process can be illustrated in Fig. 3(b), whereby for each iteration, we greedily select the timestamps s^* which maximizes the overlap between the interpolated tube and ground-truth \mathcal{A} as in Eq. 5, and append this timestamp into \mathcal{U}_k . The process stops when the IOU between interpolated tube and ground-truth tube is larger than a predefined threshold ϵ .

$$s^* = \arg \max_{s_i \in \mathcal{U}/\mathcal{U}_k} IOU(\text{Interp}(\mathcal{U}_k \cup \{s_i\}), \mathcal{A}) \quad (5)$$

Here, the function **Interp**(\cdot) can be any polynomial interpolation over the input timestamp set. To avoid the large numerical oscillation around the endpoint, we choose the piece-wise cubic spline interpolation in this paper as instantiation. We assign feature samples in \mathcal{U}_k with label 1 and samples in $\mathcal{U}/\mathcal{U}_k$ with label 0. We utilize these labels to train the timestamp selector network with binary cross-entropy loss.

Sample-wise location refinement. In the Refine Module, the selected 2D features are first processed by cascaded 2D convolution blocks (shown in Fig. 2), then a class-specific regional proposal network (RPN) [23] is applied over these features to generate bounding box proposals at corresponding timestamps. With the estimated action tube function $\hat{\mathcal{A}}(t; \theta)$, we can now obtain the estimated action bounding boxes at i -th sampled timestamps s_i with Eq. 6, where **dec**(\cdot) is the inverse operation of **enc**(\cdot, \cdot) in Eq. 3.

$$\hat{x}_i, \hat{y}_i, \hat{w}_i, \hat{h}_i = \text{dec}(\hat{\mathcal{A}}(s_i; \theta)) \quad (6)$$

We design a simple local search scheme to refine the estimated bounding box at selected key timestamps. For each selected 2D sample, a searching area Ω is defined as,

$$\Omega = [\hat{x}_i - \sigma \hat{w}_i, \hat{x}_i + \sigma \hat{w}_i] \times [\hat{y}_i - \sigma \hat{h}_i, \hat{y}_i + \sigma \hat{h}_i] \quad (7)$$

where σ is a hyperparameter that controls the size of searching area. We choose the action box proposal (from RPN) with the largest score where its center is located inside Ω , as the replacement of the original coarsely estimated box.

The final output action tube is obtained via interpolation over all refined boxes and unrefined bounding boxes (localized via Eq. 6). The associated action score is the smooth average of classification score and RPN score.

4 Experiment Results

4.1 Experiment configuration

Datasets. We conduct our experiment on three common datasets for the task of action tube detection – UCF101-24, UCFSports and JHMDB-21 datasets. Although the AVA [7] dataset also includes bounding box annotations, it mainly focuses on the problem of *atomic action classification* on sparse single key frames instead of *spatiotemporal action detection* at the video level, which is the task we are focusing here. Hence, we did not conduct our experiments on the AVA dataset.

The *UCF101-24 dataset* [31] contains 3,207 untrimmed videos with frame level bounding box annotations for 24 sports categories. The dataset is challenging due to frequent camera shake, dynamic actor movements and a large variance in action duration. Following previous works [26], we report results for the first split with 2,275 videos for training and the other videos for validation. We use the corrected annotation [29] for model training and evaluation. The *JHMDB-21* is a subset of HMDB-51 dataset [14], which contains a total of 928 videos with 21 types of actions. All video sequences are temporally trimmed. The results are reported as the average performance over 3 train-val splits. The *UCFSports* dataset [24] contains 150 trimmed videos in total and we report the results on the first split. Note that although videos in JHMDB-21 and UCFSports are trimmed temporally, their samples are still suitable for our framework as they comprised mostly of cases where actions span the whole video.

Metrics. We report the video-mAP (v-mAP) [6] with different IOU thresholds as our main evaluation metric for spatial-temporal action localization on all datasets. In addition, frame-level mAP at threshold 0.5 is reported to evaluate per-frame detection performance. A proposal is regarded as positive when its overlap with the ground-truth is larger than threshold δ . We also adopt video-level mean Average Best Overlap (v-MABO) [15] in the ablation study to evaluate the localization performance of our approach. The criterion calculates the mean of largest overlap between ground-truth tubes and action proposals, averaged over all classes.

Implementation details. We use the I3D network [3] pretrained on Kinetics-600 as our 3D feature extractor, taking the feature from *mixed_5b* layer as our 3D feature. We set the video resampling length T to 96 frames for UCF101-24 and 32 frames for JHMDB-21 and UCFSports. The hyperparameter ϵ in our paper is set to 0.8 and the segment length for matching K is set to 6 frames. The number of sampling points N is set to 16 for UCF101-24, 6 for JHMDB-21 and 8 for UCFSports according to the average length of action instances. For the anchor design, we follow the strategy of [22] by clustering the bounding boxes from

	UCF101-24				JHMDB-21		
k	2	3	4	5	1	2	3
no refine	46.0	48.4	51.6	50.1	79.7	80.9	80.3
$\sigma = 0.4$	57.5	57.6	58.8	58.0	80.8	82.5	81.3
$\sigma = 0.6$	59.9	60.1	61.7	60.0	81.4	83.2	82.4
$\sigma = 0.8$	60.3	62.0	62.7	61.6	82.3	83.7	83.2

Table 1. Ablation study on the effectiveness of refinement with different hyperparameter settings.

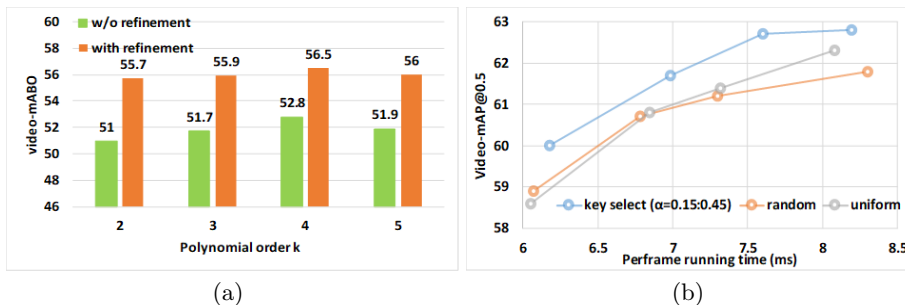


Figure 4. (a). v-MABO value of action tubes with different polynomial orders on UCF101-24. (b) Time-performance trade-off with different timestamp selection schemes on UCF101-24. (best viewed in color)

training set into 6 centers and taking their respective center coordinates as the default anchor boxes. In the training phase, we train the network in two stages. First, we train the temporal proposal network and backbone separately, and then we jointly train the entire network end-to-end to learn the final action tubes. We use the SGD solver to train CFAD with a batch size of 8. In inference stage, to handle concurrent action instances, the Coarse Module outputs at most 3 (the maximum number of instances in a video based on the datasets) estimated tubes followed by a tube-wise non-maximal suppression process with IOU threshold of 0.2 in order to avoid duplicated action tubes.

4.2 Ablation study

In this section, we report the video-mAP results with $\delta = 0.5$ of various ablation study experiments. The input modality is only RGB data unless specified.

The effectiveness of refinement. First, we analyze the effects of the refinement process on the accuracy of coarse tube estimation on UCF101-24 and JHMDB-21. The results are reported in Table 1 where “no refine” denotes the configuration without location refinement. From these results, it is obvious that the Refine Module can bring large improvements in v-mAP regardless of the polynomial order of estimated tubes; the largest performance gain can be up to +14.3% when $k = 2$ on UCF101-24. The improvement is less obvious on JHMDB-21, which we think is owing to the fact that JHMDB-21 is less dynamic and coarse-level estimations may be close to the ground-truth tubes. We




Figure 5. Statistics of true positive and various false positive proposals of CFAD on UCF101-24 and split-3 of JHMDB-21. (Best viewed in color)

also evaluate the v-MABO value on UCF101-24 as shown in Fig. 4(a), where improvements by at most +4.7% are possible by the refinement process. The results show that the Refine component is essential to better detection performance.

Meanwhile, from Table 1, we can also see that as the searching area gets larger, the mAP performance can be improved to some extent, since larger searching area can cover more centers of action proposals. We did not try larger searching area *i.e.* $\sigma > 0.8$ since we find the performance improvement is marginal (less than +0.2%) beyond $\sigma = 0.8$. This is because larger searching area also makes the refinement more vulnerable to noisy proposals.

Polynomial order selection. We also report in Table 1 the effect of different polynomial order k which decides the form of estimated tube mapping $\hat{\mathcal{A}}(t; \theta)$. Overall, we find that the performance improves along with the increase of k for both with and without refinement, since higher order polynomial functions show stronger ability in characterizing variations of action tubes.

On the other hand, we found that as the order gets larger (than $k = 5$ on UCF101-24 and $k = 3$ on JHMDB-21), the detection performance drops comparatively against the optimal value in both cases. We think the reason behind this is that although higher order polynomial functions are usually more representative, they are more complex requiring more coefficients, and the parameters predicting coefficients of higher orders are more difficult to be trained efficiently since the corresponding gradients are very small. The similar tendency is also reflected in the MABO results on UCF101-24 shown in Fig. 4(a), where the localization did improve (for both refine and no refine cases) from $k = 2$ to $k = 4$, but MABO drops after that with higher orders. Also, we observed during training that configurations with a larger polynomial order tends to slow down the convergence process and possibly result in numerical oscillation of the loss function.

Effectiveness of residual processing branch. Here, we conduct experiments on UCF101-24 and JHMDB-21 to analyze how the temporal residual information impacts the output action tube results. To test this, we remove the branch with differential module (“Diff Block” in Fig. 2) as our baseline. For a detailed comparison, we break down the final proposals into four mutually exclusive types.

- **True Positive:** the proposal classifies an action correctly and has tube-overlap with ground-truth that is larger than δ .
- **Wrong Classification Error:** a proposal with incorrectly classified action although it overlaps more than δ with ground-truth.
- **Bad Localization Error:** a proposal that has correct action class but it overlaps less than δ with ground-truth.
- **Duplicated Detection Error:** a proposal with correct action class and overlaps more than δ with a ground-truth that has been detected.

Fig. 5 illustrates the statistics of these proposals with/without the differential module. From Fig. 5, we observe that the residual processing branch is particularly important for accurate action classification. With the help of information from the temporal residual feature, wrongly-classified samples are reduced by 42% on UCF101-24 and 14% on JHMDB-21. Furthermore, models with residual processing also benefit from better tube localization while the overall recall also improves due to the increase in true positive results. These results are evidential of the effectiveness of temporal residual component in Coarse Module.

Key timestamp selection. We conduct an experiment on UCF101-24 to analyze the impact of the proposed key timestamp selection mechanism. In the experiment setting, we gradually increase the selection threshold α from 0.15 to 0.45 (in increments of 0.1) and report their respective v-mAP value and per-frame time cost. For comparisons, we design two baseline methods: (1) Random selection of samples from the input N 2D features with their corresponding timestamps taken as key timestamps, denoted as “random”. (2) Selection of timestamps across s_i based on a fixed time step, denoted as “uniform”.

The time-performance trade-off curves are shown in Fig. 4(b). We can observe that when the per-frame time costs are similar, our dynamic selection scheme is superior to the other two baseline methods. It is also worth noting that when the time cost gets smaller, the performance of “random” and “uniform” deteriorates faster than our scheme. This result indicates that the key timestamp selection process finds the important frames for location refinement and is reasonably robust to the reduction of available 2D features.

4.3 Comparison with state-of-the-art

In this section, we compare the proposed CFAD with other recent state-of-the-art approaches in the spatio-temporal action localization task on the UCF101-24, JHMDB-21 and UCFSports benchmarks. These results are listed in Table 2. We also evaluate the performance of CFAD with two-stream input, where the optical flow is extracted using the method of [2]. For simplicity, we opt for an early fusion strategy [39] to maintain efficiency of our approach.

It is worth noting that in Table 2, our method with only RGB input outperforms most other approaches that rely on two-stream features on UCF101-24 and JHMDB-21. While it is still worse than the state-of-the-art method on UCFSports, we think the reasons behind this can be that this dataset is relatively simpler and smaller in scale with less dynamic movements, thus it could be

method	JHMDB-21				UCF101-24				UCFSports	
δ	0.2	0.5	0.75	0.5:0.95	0.2	0.3	0.5	0.5:0.95	0.2	0.5
2D backbone										
Saha <i>et al.</i> [26]	72.6	71.5	-	-	66.7	54.9	35.9	14.4	-	-
Peng <i>et al.</i> [19]	74.3	73.1	-	-	72.8	65.7	30.9	7.1	94.8	94.7
Saha <i>et al.</i> [25]	57.8	55.3	-	-	63.1	51.7	33.0	10.7	-	-
Kalogeiton <i>et al.</i> [15]	74.2	73.7	52.1	44.8	76.5	-	49.2	23.4	92.7	92.7
Singh <i>et al.</i> [29]	73.8	72.0	44.5	41.6	73.5	-	46.3	20.4	-	-
Yang <i>et al.</i> [38]	-	-	-	-	76.6	-	-	-	-	-
Zhao <i>et al.</i> [39]	-	58.0	42.8	34.6	75.5	-	48.3	23.9	-	92.7
Rizard <i>et al.</i> [20]	86.0	84.0	52.8	49.5	82.3	-	51.5	24.1	-	-
Song <i>et al.</i> [30]	74.1	73.4	52.5	44.8	77.5	-	52.9	24.1	-	-
Li <i>et al.</i> [16]	82.7	81.3	-	-	76.3	71.4	-	-	97.8	97.8
Li <i>et al.</i> [17]	77.3	77.2	71.7	59.1	82.8	-	53.8	28.3	-	-
3D backbone										
Hou <i>et al.</i> [9]	78.4	76.9	-	-	73.1	69.4	-	-	95.2	95.2
Gu <i>et al.</i> [7]	-	76.3	-	-	-	-	59.9	-	-	-
Su <i>et al.</i> [32]	82.6	82.2	63.1	52.8	84.3	-	61.0	27.8	-	-
Qiu <i>et al.</i> [21]	85.7	84.9	-	-	82.2	75.6	-	-	-	-
CFAD	84.8	83.7	62.4	51.8	79.4	76.7	62.7	25.5	90.2	88.6
CFAD+OF	86.8	85.3	63.8	53.0	81.6	78.1	64.6	26.7	94.5	92.7

Table 2. Comparison with state-of-the-art methods (on video-mAP), ‘-’ denotes that the result is not available, ‘+OF’ indicates the input is combined with optical flow. All compared methods take both RGB and optical flow as input except [25,9]

more challenging to learn robust tube estimation. For fair benchmarking, we compare our method with other approaches utilizing 3D spatiotemporal features [9,7,21,32]. With RGB as input, CFAD achieves competitive performance on all datasets under different tested threshold criterion. Overall, our method achieves state-of-the-art under small threshold while there is still a margin towards the performance of [32,17] under more strict criterion. Besides, we also observe that the optical flow information is helpful for the overall detection performance.

Frame-mAP. In Table 3, we compare CFAD with other approaches on frame-level detection in UCF101-24. In our setting, we assign the video level score of a tube proposal to all boxes included by the tube to generate frame-level proposals. It can be observed that CFAD outperforms three pipelines with two-stream input. While it is still worse than some approaches [7,20,38], we think this is due to the less accurate interpolated boxes between sampled frames, which might result in many false positives with high score (which in turn lowers the overall metric). We argue that although such interpolation sacrifices frame-level accuracy, it enhances the system efficiency and video-level accuracy in return.

Efficiency. We also compare the runtime (inference) and model size of CFAD with RGB input on UCF101-24 against other approaches that also report their runtime. The speed is evaluated based on per-frame processing time, which is

method	input modal	frame-mAP@0.5
Peng <i>et al.</i> [19]	RGB+OF	65.7
Kalogeiton <i>et al.</i> [15]	RGB+OF	69.5
Yang <i>et al.</i> [38]	RGB+OF	75.0
Rizard <i>et al.</i> [20]	RGB+OF	73.7
Song <i>et al.</i> [30]	RGB+OF	72.1
Gu <i>et al.</i> [7]	RGB+OF	76.3
CFAD	RGB+OF	72.5
Hou <i>et al.</i> [9]	RGB	41.4
Yang <i>et al.</i> [38]	RGB	66.7
CFAD	RGB	69.7

Table 3. Comparison with state-of-the-art methods on frame-level mAP@0.5 on UCF101-24 dataset. ‘+OF’ indicates the input is combined with optical flow.




Figure 6. (a). Comparisons of time-performance trade-off among different state-of-the-art approaches. (b). Comparisons of trade-off between model size and performance among different state-of-the-art approaches. (Best viewed in color)

obtained by taking the runtime per video and dividing it by input length T . Since some other works only reported per-video time on JHMDB-21 [26,25], we compute their per-frame time in the same manner. The runtime comparison is illustrated in Fig. 6(a) and the model size comparison is reported in Fig. 6(b). We observe that CFAD only requires a small number of parameters (close to [29,39], and much less than others) while achieving superior running speed compared to other state-of-the-art methods. This vast improvement in processing efficiency can be attributed to the coarse-to-fine paradigm of CFAD, which does not require dense per-frame action detection followed by linking, and the RGB input of CFAD avoids the additional computation to process optical flow. Specifically, the proposed CFAD runs $\approx 3.3\times$ faster than the nearest approach [29] (7.6 ms vs. 25 ms).

4.4 Qualitative results

Fig. 7 shows some qualitative results of detected action tubes from the UCF101-24 dataset. The green boxes denote the estimated action tube output from the Coarse Module while the red boxes are the refined action tubes. We can observe that the selective refinement process has effectively corrected some poorly




Figure 7. Visualization of detected action tubes. The green boxes denote the estimated action tubes from the Coarse Module. The red boxes are the final refined action tubes. (Best viewed in color)

located action tubes, causing the bounding boxes to wrap tighter and more accurately around the actors. These visuals can evidently explain the robustness of the coarse tube estimation method, and its capability at handling a variety of dynamic actions.

5 Conclusion

In this paper, we propose a novel framework CFAD for spatio-temporal action localization. Its pipeline follows a new coarse-to-fine paradigm, which does away with the need for dense per-frame detections. The action detector comprises of two components (Coarse and Refine Modules) which play vital roles in coarsely estimating and then refining action tubes based on selected timestamps. Our CFAD achieves state-of-the-art results for a good range of thresholds on benchmark datasets and is also an efficient pipeline, running at $3.3\times$ faster than the nearest competitor.

Acknowledgement

The paper is supported in part by the following grants: China Major Project for New Generation of AI Grant (No.2018AAA0100400), National Natural Science Foundation of China (No. 61971277). The work is also supported by funding from Clobotics under the Joint Research Program of Smart Retail.

References

1. Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M.W., Pfau, D., Schaul, T., Shillingford, B., De Freitas, N.: Learning to learn by gradient descent by gradient descent. In: NeurIPS. pp. 3981–3989 (2016)
2. Brox, T., Bruhn, A., Papenberger, N., Weickert, J.: High accuracy optical flow estimation based on a theory for warping. In: ECCV. Springer (2004)

3. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: CVPR. pp. 6299–6308 (2017)
4. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR. pp. 248–255. IEEE (2009)
5. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: CVPR. pp. 580–587 (2014)
6. Gkioxari, G., Malik, J.: Finding action tubes. CVPR (Jun 2015)
7. Gu, C., Sun, C., Vijayanarasimhan, S., Pantofaru, C., Ross, D.A., Toderici, G., Li, Y., Ricco, S., Sukthankar, R., Schmid, C., Malik, J.: Ava: A video dataset of spatio-temporally localized atomic visual actions. CVPR pp. 6047–6056 (2018)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
9. Hou, R., Chen, C., Shah, M.: An end-to-end 3d convolutional neural network for action detection and segmentation in videos. ICCV (2017)
10. Hu, R., Dollár, P., He, K., Darrell, T., Girshick, R.: Learning to segment every thing. In: CVPR. pp. 4233–4241 (2018)
11. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: CVPR. pp. 4700–4708 (2017)
12. Huang, J., Li, N., Zhong, J., Li, T.H., Li, G.: Online action tube detection via resolving the spatio-temporal context pattern. In: ACM MM. pp. 993–1001. ACM (2018)
13. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: NeurIPS. pp. 2017–2025 (2015)
14. Jhuang, H., Gall, J., Zuffi, S., Schmid, C., Black, M.J.: Towards understanding action recognition. In: ICCV. pp. 3192–3199 (Dec 2013)
15. Kalogeiton, V., Weinzaepfel, P., Ferrari, V., Schmid, C.: Action tubelet detector for spatio-temporal action localization. In: ICCV. pp. 4405–4413 (2017)
16. Li, D., Qiu, Z., Dai, Q., Yao, T., Mei, T.: Recurrent tubelet proposal and recognition networks for action detection. In: ECCV. pp. 303–318 (2018)
17. Li, Y., Wang, Z., Wang, L., Wu, G.: Actions as moving points. In: arXiv preprint arXiv:2001.04608 (2020)
18. Li, Z., Gavriluyk, K., Gavves, E., Jain, M., Snoek, C.G.: Videolstm convolves, attends and flows for action recognition. Computer Vision and Image Understanding **166**, 41–50 (2018)
19. Peng, X., Schmid, C.: Multi-region two-stream r-cnn for action detection. In: ECCV. pp. 744–759. Springer (2016)
20. Pramono, R.R.A., Chen, Y.T., Fang, W.H.: Hierarchical self-attention network for action localization in videos. In: ICCV (2019)
21. Qiu, Z., Yao, T., Ngo, C.W., Tian, X., Mei, T.: Learning spatio-temporal representation with local and global diffusion. In: CVPR. pp. 12056–12065 (2019)
22. Redmon, J., Farhadi, A.: Yolo9000: better, faster, stronger. In: CVPR. pp. 7263–7271 (2017)
23. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NeurIPS. pp. 91–99 (2015)
24. Rodriguez, M.D., Ahmed, J., Shah, M.: Action mach a spatio-temporal maximum average correlation height filter for action recognition. In: CVPR. pp. 1–8 (June 2008)
25. Saha, S., Singh, G., Cuzzolin, F.: Amtnet: Action-micro-tube regression by end-to-end trainable deep architecture. In: ICCV. pp. 4414–4423 (2017)
26. Saha, S., Singh, G., Sapienza, M., Torr, P.H., Cuzzolin, F.: Deep learning for detecting multiple space-time action tubes in videos. BMVC (2016)

27. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: NeurIPS. pp. 568–576 (2014)
28. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
29. Singh, G., Saha, S., Sapienza, M., Torr, P.H., Cuzzolin, F.: Online real-time multiple spatiotemporal action localisation and prediction. In: ICCV. pp. 3637–3646 (2017)
30. Song, L., Zhang, S., Yu, G., Sun, H.: Tacnet: Transition-aware context network for spatio-temporal action detection. In: CVPR. pp. 11987–11995 (2019)
31. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild (2012)
32. Su, R., Ouyang, W., Zhou, L., Xu, D.: Improving action localization by progressive cross-stream cooperation. In: CVPR. pp. 12016–12025 (2019)
33. Sun, L., Jia, K., Chen, K., Yeung, D.Y., Shi, B.E., Savarese, S.: Lattice long short-term memory for human action recognition. In: ICCV. pp. 2147–2156 (2017)
34. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. ICCV (Dec 2015)
35. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. CVPR (Jun 2018)
36. Xu, H., Das, A., Saenko, K.: R-c3d: Region convolutional 3d network for temporal activity detection. In: ICCV. pp. 5783–5792 (2017)
37. Yang, T., Zhang, X., Li, Z., Zhang, W., Sun, J.: Metaanchor: Learning to detect objects with customized anchors. In: NeurIPS. pp. 320–330 (2018)
38. Yang, X., Yang, X., Liu, M.Y., Xiao, F., Davis, L.S., Kautz, J.: Step: Spatio-temporal progressive learning for video action detection. In: CVPR. pp. 264–272 (2019)
39. Zhao, J., Snoek, C.G.: Dance with flow: Two-in-one stream action detection. In: CVPR. pp. 9935–9944 (2019)