

STEP: Spatio-Temporal Progressive Learning for Video Action Detection

Xitong Yang^{1*} Xiaodong Yang² Ming-Yu Liu²
Fanyi Xiao^{3*} Larry Davis¹ Jan Kautz²

¹University of Maryland, College Park ²NVIDIA ³University of California, Davis

Abstract

In this paper, we propose *Spatio-TEmporal Progressive (STEP) action detector*—a progressive learning framework for spatio-temporal action detection in videos. Starting from a handful of coarse-scale proposal cuboids, our approach progressively refines the proposals towards actions over a few steps. In this way, high-quality proposals (i.e., adhere to action movements) can be gradually obtained at later steps by leveraging the regression outputs from previous steps. At each step, we adaptively extend the proposals in time to incorporate more related temporal context. Compared to the prior work that performs action detection in one run, our progressive learning framework is able to naturally handle the spatial displacement within action tubes and therefore provides a more effective way for spatio-temporal modeling. We extensively evaluate our approach on UCF101 and AVA, and demonstrate superior detection results. Remarkably, we achieve mAP of 75.0% and 18.6% on the two datasets with 3 progressive steps and using respectively only 11 and 34 initial proposals.

1. Introduction

Spatio-temporal action detection aims to recognize the actions of interest that present in a video and localize them in both space and time. Inspired by the advances in the field of object detection in images [8, 21], most recent work approaches this task based on the standard two-stage framework: in the first stage action proposals are produced by a region proposal algorithm or densely sampled anchors, and in the second stage the proposals are used for action classification and localization refinement.

Compared to object detection in images, spatio-temporal action detection in videos is however a more challenging problem. New challenges arise from both of the above two stages when the temporal characteristic of videos is taken into account. First, an action tube (i.e., a sequence of bounding boxes of action) usually involves spatial displace-




Figure 1: A schematic overview of spatio-temporal progressive learning for action detection. Starting with a coarse-scale proposal cuboid, it progressively refines the proposal towards the action, and adaptively extends the proposal to incorporate more related temporal context at each step.

ment over time, which introduces extra complexity for proposal generation and refinement. Second, effective temporal modeling becomes imperative for accurate action classification, as a number of actions are only identifiable when temporal context information is available.

Previous work usually exploits temporal information by performing action detection at the clip (i.e., a short video snippet) level. For instance, [12, 17] take as input a sequence of frames and output the action categories and regressed tubelets of each clip. In order to generate action proposals, they extend 2D region proposals to 3D by replicating them over time, assuming that the spatial extent is fixed within a clip. However, this assumption would be violated for the action tubes with large spatial displacement, in particular when the clip is long or involves rapid movement of actors or camera. Thus, using long cuboids directly as action proposals is not optimal, since they introduce extra noise for action classification and make action localization more challenging, if not hopeless. Recently, there are some attempts to use adaptive proposals for action detection [16, 20]. However, these methods require an offline linking process to generate the proposals.

*Work done during an internship at NVIDIA Research.

In this paper, we present a novel learning framework, **Spatio-TEmporal Progressive (STEP)** action detector, for video action detection. As illustrated in Figure 1, unlike existing methods that directly perform action detection in one run, our framework involves a multi-step optimization process that progressively refines the initial proposals towards the final solution. Specifically, STEP consists of two components: **spatial refinement** and **temporal extension**. Spatial refinement starts with a small number of coarse-scale proposals and updates them iteratively to better classify and localize action regions. We carry out the multiple steps in a sequential order, where the outputs of one step are used as the proposals for next step. This is motivated by the fact that the regression outputs can better follow actors and adapt to action tubes than the input proposals. Temporal extension focuses on improving classification accuracy by incorporating longer-range temporal information. However, simply taking a longer clip as input is inefficient and also ineffective since a longer sequence tends to have larger spatial displacement, as shown in Figure 1. Instead, we progressively process longer sequences at each step and adaptively extend proposals to follow action movement. In this manner, STEP can naturally handle the spatial displacement problem and therefore provide more efficient and effective spatio-temporal modeling. Moreover, STEP achieves superior performance by using only a handful (e.g., 11) of proposals, obviating the need to generate and process large numbers (e.g., >1K) of proposals due to the tremendous spatial and temporal search space.

To our knowledge, this work provides the first end-to-end progressive optimization framework for video action detection. We bring up the spatial displacement problem in action tubes and show that our method can naturally handle the problem in an efficient and effective way. Extensive evaluations find our approach to produce superior detection results while only using a small number of proposals.

2. Related Work

Action Recognition. A large family of the research in video action recognition is about action classification, which provides fundamental tools for action detection, such as two-stream networks on multiple modalities [28, 35], 3D-CNN for simultaneous spatial and temporal feature learning [4, 13], and RNNs to capture temporal context and handle variable-length video sequences [25, 36]. Another active research line is the temporal action detection, which focuses on localizing the temporal extent of each action. Many methods have been proposed, from fast temporal action proposals [15], region convolutional 3D network [34], to budget-aware recurrent policy network [23].

Spatio-Temporal Action Detection. Inspired by the recent advances in image object detection, a number of efforts have been made to extend image object detectors (e.g., R-

CNN, Fast R-CNN and SSD) to the task as frame-level action detectors [10, 26, 27, 30, 33, 37, 38]. The extensions mainly include: first, optical flow is used to capture motion cues, and second, linking algorithms are developed to connect frame-level detection results as action tubes. Although these methods have achieved promising results, the temporal property of videos is not explicitly or fully exploited as the detection is performed on each frame independently. To better leverage the temporal cues, several recent work has been proposed to perform action detection at clip level. For instance, ACT [17] takes as input a short sequence of frames (e.g., 6 frames) and outputs the regressed tubelets, which are then linked by a tubelet linking algorithm to construct action tubes. Gu et al. [12] further demonstrate the importance of temporal information by using longer clips (e.g., 40 frames) and taking advantage of I3D pre-trained on the large-scale video dataset [4]. Rather than linking the frame or clip level detection results, there are also some methods that are developed to link the proposals before classification to generate action tube proposals [16, 20].

Progressive Optimization. This technique has been explored in a range of vision tasks from pose estimation [3], image generation [11] to object detection [2, 6, 7, 24]. Specifically, the multi-region detector [6] introduces iterative bounding box regression with R-CNN to produce better regression results. AttractionNet in [7] employs a multi-stage procedure to generate accurate object proposals that are then input to Fast R-CNN. G-CNN [24] trains a regressor to iteratively move a grid of bounding boxes towards objects. Cascade R-CNN [2] proposes a cascade framework for high-quality object detection, where a sequence of R-CNN detectors are trained with increasing IoU thresholds to iteratively suppress close false positives.

3. Method

In this section, we introduce the proposed progressive learning framework STEP for video action detection. We first formulate the problem and provide an overview of our approach. We then describe in details the two primary components of STEP including spatial refinement and temporal extension. Finally, the training algorithm and implementation details are presented.

3.1. Framework Overview

Proceeding with the recent work [12, 17], our approach performs action detection at clip level, i.e., detection results are first obtained from each clip and then linked to build action tubes across a whole video. We assume that each action tubelet of a clip has a constant action label, considering the short duration of a clip, e.g., within one second.

Our target is to tackle the action detection problem through a few progressive steps, rather than directly detecting actions all at one run. In order to detect the actions in a




Figure 2: Example of the 11 initial proposals: 2D boxes are replicated across time to obtain cuboids.

clip I_t with K frames, according to the maximum progressive steps S_{max} , we first extract the convolutional features for a set of clips $\mathcal{I} = \{I_{t-S_{max}+1}, \dots, I_t, \dots, I_{t+S_{max}-1}\}$ using a backbone network such as VGG16 [29] or I3D [4]. The progressive learning starts with M pre-defined proposal cuboids $B^0 = \{b_i^0\}_{i=1}^M$ and $b_i^0 \in \mathcal{R}^{K \times 4}$, which are sparsely sampled from a coarse-scale grid of boxes and replicated across time to form the initial proposals. An example of the 11 initial proposals used in our experiments is illustrated in Figure 2. These initial proposals are then progressively updated to better classify and localize the actions. At each step s , we update the proposals by performing the following processes in order:

- **Extend:** the proposals are temporally extended to the adjacent clips to include longer-range temporal context, and the temporal extension is adaptive to the movement of actions, as described in Section 3.3.
- **Refine:** the extended proposals are forwarded to the spatial refinement, which outputs the classification and regression results, as presented in Section 3.2.
- **Update:** all proposals are updated using a simple greedy algorithm, i.e., each proposal is replaced by the regression output with the highest classification score:

$$b_i^s \doteq l_i^s(c^*), \quad c^* = \arg \max_c p_i^s(c), \quad (1)$$

where c is an action class, $p_i^s \in \mathcal{R}^{(C+1)}$ is the probability distribution of the i th proposal over C action classes plus background, $l_i^s \in \mathcal{R}^{K \times 4 \times C}$ denotes its parameterized coordinates (for computing the localization loss in Eq. 3) at each frame for each class, and \doteq indicates decoding the parameterized coordinates. We summarize the outline of our detection algorithm in Algorithm 1.

3.2. Spatial Refinement

At each step s , the spatial refinement solves a multi-task learning problem that involves action classification and localization regression. Accordingly, we design a two-branch architecture, which learns separate features for the two tasks, as illustrated in Figure 3. Our motivation is that the two tasks have substantially different objectives and re-

Algorithm 1: STEP Action Detection for Clip I_t

Input : video clips \mathcal{I} , initial proposals B^0 , and maximum steps S_{max}

Output: detection results $\{(p_i^{S_{max}}, l_i^{S_{max}})\}_{i=1}^M$

- 1 extract convolutional features for video clips \mathcal{I}
- 2 **for** $s \leftarrow 1$ **to** S_{max} **do**
- 3 **if** $s == 1$ **then**
- 4 // initial proposals
- 5 $\tilde{B}^{s-1} \leftarrow B^0$
- 6 **else**
- 7 // temporal extension (Sec.3.3)
- 8 $\tilde{B}^{s-1} \leftarrow \text{Extend}(B^{s-1})$
- 9 **end**
- 10 // spatial refinement (Sec.3.2)
- 11 $\{(p_i^s, l_i^s)\}_{i=1}^M \leftarrow \text{Refine}(\tilde{B}^{s-1})$
- 12 // update proposals (Eq.1)
- 13 $B^s \leftarrow \text{Update}(\{(p_i^s, l_i^s)\}_{i=1}^M)$
- 14 **end**

quire different types of information. For accurate action classification, it demands context features in both space and time, while for robust localization regression, it needs more precise spatial cues at frame level. As a result, our two-branch network consists of a *global branch* that performs spatio-temporal modeling on the entire input sequence for action classification, as well as a *local branch* that performs bounding box regression at each frame.

Given the frame-level convolutional features and the tubelet proposals for the current step, we first extract regional features through an ROI pooling [8]. Then we take the regional features to the global branch for spatio-temporal modeling and produce the global feature. Each global feature encodes the context information of a whole tubelet and is further used to predict the classification output p_i^s . Moreover, the global feature is concatenated with the corresponding regional features at each frame to form the local feature, which is used to generate the class-specific regression output l_i^s . Our local feature not only captures the spatio-temporal context of a tubelet but also extracts the local details of each frame. By jointly training the two branches, the network learns the two separate features that are informative and adaptable for their own tasks.

Training Loss. We enforce a multi-task loss to jointly train for action classification and tubelet regression. Let \mathcal{P}^s denote the set of selected positive samples and \mathcal{N}^s the set of negative samples at step s (the sampling strategy is described in Section 3.4). We define the training loss \mathcal{L}^s as:

$$\mathcal{L}^s = \sum_{i \in \{\mathcal{P}^s, \mathcal{N}^s\}} \mathcal{L}_{cls}(p_i^s, u_i) + \lambda \sum_{i \in \mathcal{P}^s} \mathcal{L}_{loc}(l_i^s(u_i), v_i), \quad (2)$$




Figure 3: Left: the architecture of our two-branch network. Right: the illustration of our progressive learning framework, where “S” indicates spatial refinement, “T” temporal extension, “P” classification, and “L” localization, the numbers correspond to the steps, and “L0” denotes the initial proposals.

where u_i and v_i are the ground truth class label and localization target for the i th sample, and λ is the weight to control the importance of the two loss terms. We employ the multi-class cross-entropy loss as the classification loss $\mathcal{L}_{cls}(p_i^s, u_i) = -\log p_i^s(u_i)$ in Eq. 2. We define the localization loss using the averaged $\ell_{1,smooth}$ between predicted and ground truth bounding boxes over the frames of a clip:

$$\mathcal{L}_{loc}(l_i^s(u_i), v_i) = \frac{1}{K} \sum_{k=1}^K \ell_{1,smooth}(l_{i,k}^s(u_i) - v_{i,k}). \quad (3)$$

We apply the same parameterization for $v_{i,k}$ as in [9] by using a scale-invariant center translation and a log-space height/width shift relative to the bounding box.

3.3. Temporal Extension

Video temporal information, especially the long-term temporal dependency, is critical for accurate action classification [4, 36]. In order to leverage longer range of temporal context, we extend the proposals to include in more frames as input. However, the extension is not trivial since the spatial displacement problem becomes even more severe for longer sequences, as illustrated in Figure 1. Recently, some negative impacts caused by the spatial displacement problem for action detection have also been observed by [12, 17], which simply replicate 2D proposals across time to increase longer temporal length.

With the intention to alleviate the spatial displacement problem, we perform temporal extension *progressively* and *adaptively*. From the second step, we extend the tubelet proposals to the two adjacent clips at a time. In other words, at each step $1 \leq s < S_{max}$, the proposals B^s with length K^s are extended to $\tilde{B}^s = B_{-1}^s \circ B^s \circ B_{+1}^s$ with length $K^s + 2K$, where \circ denotes concatenation. Additionally, the temporal extension is adaptive to action movement by taking advantage of the regressed tubelets from the previous step. We

introduce two methods to enable the temporal extension to be adaptive as described in the following.

Extrapolation. By assuming that the spatial movement of an action satisfies a linear function approximately within a short temporal range, such as a 6-frame clip, we can extend the tubelet proposals by using a simple linear extrapolation function:

$$B_{+1,k}^s = B_{K^s}^s + \frac{k}{K-1}(B_{K^s}^s - B_{K^s-K+1}^s). \quad (4)$$

A similar function can be applied to B_{-1}^s to adapt to the movement trend, but the assumption would be violated for long sequences and therefore results in drifted estimations.

Anticipation. We can also achieve the adaptive temporal extension by location anticipation, i.e., training an extra regression branch to conjecture the tubelet locations in adjacent clips based on the current clip. Intuitively, the anticipation requires the network to infer the movement trend in adjacent clips by action modeling in the current clip. A similar idea is explored in [37], where location anticipation is used at the region proposal stage.

We formulate our location anticipation as a residual learning problem [14, 22] based on the assumption that the tubelets of two adjacent clips differ from each other by a small residual. Let x indicate the features forwarded to the output layer f of the location regressor $L^s = f(x)$ at step s . So the anticipated locations can be obtained as:

$$L_{-1}^s = L^s + f_{-1}(x), \quad L_{+1}^s = L^s + f_{+1}(x), \quad (5)$$

where f_{-1} and f_{+1} are the anticipation regressors, which are lightweight and introduce negligible computational overhead. L_{-1}^s and L_{+1}^s are then decoded to the proposals B_{-1}^s and B_{+1}^s . The loss function of location anticipation is defined in a similar way as Eq. 3, and combined with \mathcal{L}_{cls} and \mathcal{L}_{loc} with a coefficient γ to form the overall loss.

3.4. Network Training

Although STEP involves multiple progressive steps, the whole framework can be trained end-to-end to optimize the models at different steps jointly. Compared against the step-wise training scheme used in [24], our joint training is simpler to implement, runs more efficiently, and achieves better performance in our experiments.

Given a mini-batch of training data, we first perform an $(S_{max} - 1)$ -step inference pass, as illustrated in the right of Figure 3, to obtain the inputs needed for all progressive steps. In practice, the detection outputs $\{(p_i^s, l_i^s)\}_{i=1}^M$ at each step are collected and used to select the positive and negative samples \mathcal{P}^s and \mathcal{N}^s for training. We accumulate the losses of all steps and back-propagate to update the whole model at the same time.

Distribution Change. Compared to the prior work that performs detection in one run, our training could be more challenging as the input/output distributions change over steps. As shown in Figure 4, the input distribution is right-skewed or centered in a low-IoU level at early steps, and reverses at later steps. This is because our approach starts from a coarse-scale grid (see Figure 2) and progressively refines them towards generating high-quality proposals. Accordingly, the range of output distribution (i.e., the scale of offset vectors) decreases over steps.

Inspired by [2], we tackle the distribution change in three ways. First, separate headers are used at different steps to adapt to the different input/output distributions. Second, we increase IoU thresholds over the multiple steps. Intuitively, a lower IoU threshold at early steps tolerates the initial proposals to include sufficient positive samples and a higher IoU threshold at late steps encourages high-quality detection. Third, a hard-aware sampling strategy is employed to select more informative samples during training.

Hard-Aware Sampling. We design the sampling strategy based on two principles: (i) the numbers of positive and negative samples should be roughly balanced, and (ii) the harder negatives should be selected more often. To measure the “hardness” of a negative sample, we use the classification scores from the previous step. The tubelet with a high confidence but a low overlap to any ground truth is viewed as a hard sample. We calculate the overlap of two tubelets by averaging the IoU of bounding boxes over K frames of the target clip. So the negative samples with higher classification scores will be sampled with a higher chance.

Formally, given a set of proposals and the overlap threshold τ^s at step s , we first assign positive labels to the candidates with the highest overlap with ground truth. This is to ensure that each ground truth tube has at least one positive sample. After that, the proposals having an overlap higher than τ^s with any ground truth tube are added to the positive pool and the rest to the negative pool. We then sample $|\mathcal{P}^s|$ positives and $|\mathcal{N}^s|$ negatives from the two pools,




Figure 4: Change of input distribution (IoU between input proposals and ground truth) over steps on UCF101.

respectively, with the sampling probability proportional to the classification score. For the first step, the highest overlap with ground truth tubes is used as the score for sampling. Each selected positive in \mathcal{P}^s is assigned to the ground truth tube with which it has the highest overlap. Note that a single proposal can be assigned to only one ground truth tube.

3.5. Full Model

We can also integrate our model with the common practices for video action detection [12, 17, 30], such as two-stream fusion and tubelet linking.

Scene Context. It has been proven to be beneficial to object and action detection [20, 32]. Intuitively, some action-related semantic clues from scene context can be utilized to improve action classification, for example, the scene of a basketball court for recognizing “basketball dunk”. We incorporate scene context by concatenating extended features to original regional features in the global branch. The extended features can be obtained by RoI pooling of the whole image. So the global features encode both spatial and temporal context useful for action classification.

Two-Stream Fusion. Most previous methods use late fusion to combine the results at test time, i.e., the detections are obtained independently from the two streams and then fused using either mean fusion [17] or union fusion [30]. In this work, we also investigate early fusion for two-stream fusion, which concatenates RGB frames and optical flow maps in channel and input to the network as a whole. Intuitively, early fusion can model the low-level interactions between the two modalities and also obviates the need for training two separate networks. In addition, a hybrid fusion can be further performed to combine detection results from the early fusion and the two streams. Our experiment shows that early fusion outperforms late fusion, and hybrid fusion achieves the best performance.

Tubelet Linking. Given the clip-level detection results, we link them in space and time to construct the final action tubes. We follow the same linking algorithm as described in [17], apart from that we do not apply global non-maximum suppression across classes but perform temporal trimming over the linked paths as commonly used in [20, 27]. The temporal trimming enforces consecutive boxes to have smooth classification scores by solving an energy maximization problem via dynamic programming.

4. Experiments

In this section, we describe the experiments to evaluate STEP and compare against the recent competing algorithms. We start by performing a variety of ablation studies to better understand the contributions of each individual component in our approach. We then report comparisons to the state-of-the-art methods, provide in-depth analysis, and present the qualitative detection results.

4.1. Experimental Setup

Datasets. We evaluate our approach on the two benchmarks: UCF101 [31] and AVA [12]. In comparison with other action detection datasets, such as J-HMDB and UCF-Sports, the two benchmarks are much larger and more challenging, and more importantly, they are temporally untrimmed, which fits better to the spatio-temporal action detection task. UCF101 is originally an action classification dataset collected from online videos, and a subset of 24 classes with 3,207 videos are provided with the spatio-temporal annotations for action detection. Following the standard evaluation protocol [17], we report results on the first split of the dataset. AVA contains complex actions and scenes sourced from movies. We use the version 2.1 of AVA, which consists of the annotations at 1 fps over 80 action classes. Following the standard setup in [12], we report results on the most frequent 60 classes that have at least 25 validation examples per class.

Evaluation Metrics. We report the frame-level mean average precision (frame-mAP) with an IoU threshold of 0.5 for both datasets. This metric allows us to evaluate the quality of the detection results independently of the linking algorithm. We also use the video-mAP on UCF101 to compare with the state-of-the-art results.

Implementation Details. For the experiments on UCF101, we use VGG16 [29] pre-trained on ImageNet [5] as the backbone network. Although more advanced models are available, we choose the same backbone as [17] for fair comparisons. For the temporal modeling in global branch, we use three 3D convolutional layers with adaptive max pooling along the temporal dimension. All frames are resized to 400×400 and the clip length is set to $K = 6$. Similar to [17], 5 consecutive optical flow maps are stacked as a whole for the optical flow input. We train our models for 35 epochs using Adam [19] with a batch size of 4. We set the initial learning rate to 5×10^{-5} and perform step decay after 20 and 30 epochs with the decay rate 0.1.

For the experiments on AVA, we adopt I3D [4] (up to Mixed_4f) pre-trained on Kinetics-400 [18] as the backbone network. We take the two layers Mixed_5b and Mixed_5c of I3D for temporal modeling in our global branch. All frames are resized to 400×400 and the clip length is set to $K = 12$. We use 34 initial proposals and perform temporal extension only at the third step. As the

S_{max}	s				Mode	f-mAP
	1	2	3	4		
1	51.5	-	-	-	RGB	66.7
2	56.6	60.7	-	-	Flow	63.5
3	57.1	61.8	62.6	-	Late	70.7
4	58.2	62.1	62.8	62.7	Early	74.3
					Hybrid	75.0

Table 1: Comparisons of frame-mAP (%) of our models trained with different numbers of steps (left), and different input modalities and fusion methods (right).

classification is more challenging on AVA, we first pre-train our model for an action classification task using the spatial ground truth of training set. We then train the model for action detection with a batch size of 4 for 10 epochs. We do not use optical flow on this dataset due to the heavy computation and instead combine results of two RGB models. Our initial learning rate is 5×10^{-6} for the backbone network and 5×10^{-5} for the two-branch networks, and step decay is performed after 6 epochs with the decay rate 0.1.

For all experiments, we extract optical flow (if used) with Brox [1], and perform data augmentation to the whole sequence of frames during training, including random flipping and cropping. More architecture and implementation details are available in the supplementary material.

4.2. Ablation Study

We perform various ablation experiments on UCF101 to evaluate the impacts of different design choices in our framework. For all experiments in this section, we employ the 11 initial proposals as shown in Figure 2 and RGB only, unless explicitly mentioned otherwise, and frame-mAP is used as the evaluation metric.

Effectiveness of Spatial Refinement. Our primary design of STEP is to progressively tackle the action detection problem through a few steps. We thus first verify the effectiveness of progressive learning by comparing the detection results at different steps with the spatial refinement. No temporal extension is applied in this comparison. Table 1(a) demonstrates the step-wise performance under different maximum steps S_{max} . Since our approach starts from the coarse-scale proposals, performing spatial refinement once is insufficient to achieve good results. We observe that the second step improves results consistently and substantially, indicating that the updated proposals have higher quality and provide more precise information for classification and localization. Further improvement can be obtained by additional steps, suggesting the effectiveness of our progressive spatial refinement. We use 3 steps for most of our experiments as the performance saturates after that. Note that using more steps also improves the results of early steps, due to the benefits of our multi-step joint training.




Figure 5: Comparison of frame-mAP (%) of our models trained with and without temporal extension.

Effectiveness of Temporal Extension. In addition to the spatial refinement, our progressive learning contains the temporal extension to progressively process a longer sequence at each step. We compare the detection results with and without temporal extension in Figure 5. We show the results of the models taking $K = 6$ and $K = 30$ frames as inputs directly without temporal extension, and the results of the extrapolation and anticipation methods. Note that the models with temporal extension also deal with 30 frames at the third step (extension process: $6 \rightarrow 18 \rightarrow 30$).

Both of the temporal extension methods outperform the baseline ($K = 6$) by a large margin, which clearly shows the benefit of incorporating longer-range temporal context for action classification. More remarkably, simply taking $K = 30$ frames as input without temporal extension results in inferior performance, validating the importance of adaptively extending the temporal scale in the progressive manner. Furthermore, we observe that anticipation performs better than extrapolation for longer sequences, indicating that anticipation can better capture nonlinear movement trends and therefore generate better extensions.

Fusion Comparison. Table 1(b) presents the detection results of different fusions: late, early and hybrid fusion. In all cases, using both modalities improves the performance compared to individual ones. We find that early fusion outperforms late fusion, and attribute the improvement to modeling between the two modalities at the early stage. Hybrid fusion achieves the best result by further utilizing the complementary information of different methods.

Miscellaneous. We describe several techniques to improve the training in Section 3, including incorporating scene context, hard-award sampling and increasing IoU threshold. To validate the contributions of the three techniques, we conduct ablation experiments by removing one at a time, which correspondingly results in a performance drop of 2.5%, 1.5% and 1%. In addition, we observe that incorporating scene context provides more gains for later steps, suggesting that scene context is more important for action classification when bounding boxes become tight.




Figure 6: Analysis of runtime of our approach under various settings: (a) the inference speeds using different step numbers with and without temporal extension, and (b) the detection results (green dots) and speeds (blue bars) using different numbers of initial proposals.

4.3. Runtime Analysis

Although STEP involves a multi-step optimization, our model is efficient since we only process a small number of proposals. STEP runs at 21 fps using early fusion with 11 initial proposals and 3 steps on a single GPU, which is comparable with the clip based approach (23 fps) [17] and much faster than the frame based method (4 fps) [26]. Figure 6(a) demonstrates the speeds of our approach with increasing number of steps under the settings with and without temporal extension. We also report the running time and detection performance of our approach (w/o temporal extension for 3 steps) with increasing number of initial proposals in Figure 6(b). We observe substantial gains in detection accuracy by increasing the number of initial proposals, but it also results in slowed inference speed. This trade-off between accuracy and speed can be controlled according to a specified time budget.

4.4. Comparison with State-of-the-Art Results

We compare our approach with the state-of-the-art methods on UCF101 and AVA in Tables 2 and 3. Following the standard settings, we report the frame-mAP at IoU threshold 0.5 on both datasets and the video-mAP at various IoU thresholds on UCF101. STEP consistently performs better than the state-of-the-art methods on UCF101, and brings a clear gain in frame-mAP, producing 5.5% improvement over the second best result. Our approach also achieves superior result on AVA, outperforming the recently proposed ACRN by 1.2%. Notably, STEP performs detection simply from a handful of initial proposals, while other competing algorithms rely on a great amount of densely sampled anchors or an extra person detector trained with external large-scale image object detection datasets.

4.5. Qualitative Results

We visualize the detection results of our approach at different steps in Figure 7. Each row indicates the detection outputs at a certain step. A bounding box is labeled in red if the detection result is correct, otherwise it is labeled in blue.




Figure 7: Examples of the detection results on UCF101. Red boxes indicate correct detection and blue ones misclassification. (a) illustrates the effect of progressive learning to improve action classification over steps. (b) demonstrates the regression outputs by spatial refinement at each step.

Method	frame-mAP	video-mAP		
	0.5	0.05	0.1	0.2
MR-TS [26]	65.7	78.8	77.3	72.9
ROAD [30]	-	-	-	73.5
CPLA [37]	-	79.0	77.3	73.5
RTPR [20]	-	81.5	80.7	76.3
PntMatch [38]	67.0	79.4	77.7	76.2
T-CNN [16]	67.3	78.2	77.9	73.1
ACT [17]	69.5	-	-	76.5
Ours	75.0	84.6	83.1	76.6

Table 2: Comparison with the state-of-the-art methods on UCF101 by frame-mAP and video-mAP under different IoU thresholds.

Method	frame-mAP
Single Frame* [12]	14.2
I3D [12]	14.7
I3D* [12]	15.6
ACRN* [32]	17.4
Ours	18.6

Table 3: Comparison with the state-of-the-art methods on AVA by frame-mAP under $\text{IoU} = 0.5$. “*” means the results obtained by incorporating optical flow.

Figure 7(a) demonstrates the effect of progressive learning for more accurate action classification. It can be observed by the fact that the blue boxes are eliminated at later steps. In Figure 7(b), the first row corresponds to the initial proposals and the next two rows show the effect of spatial refinement of the proposals over steps. It is clear that the proposals progressively move towards the persons perform-



Figure 8: Examples of the small scale action detection by our approach. Red boxes indicate the initial proposals and orange ones the detection outputs.

ing the actions and better localization results are obtained at later steps. Although starting from coarse-scale proposals, our approach is robust to various action scales thanks to the progressive spatial refinement, as illustrated in Figure 8.

5. Conclusion

In this paper, we have proposed the spatio-temporal progressive learning framework STEP for video action detection. STEP involves spatial refinement and temporal extension, where the former starts from sparse initial proposals and iteratively updates bounding boxes, and the latter gradually and adaptively increases sequence length to incorporate more related temporal context. STEP is found to be able to more effectively make use of longer temporal information by handling the spatial displacement problem in action tubes. Extensive experiments on two benchmarks show that STEP consistently brings performance gains by using only a handful of proposals and a few updating steps.

Acknowledgement. Davis acknowledges the support from IARPA via Department of Interior/Interior Business Center (DOI/IBC) under contract number D17PC00345.

References

- [1] Thomas Brox, Andrés Bruhn, Nils Papenberg, and Joachim Weickert. High accuracy optical flow estimation based on a theory for warping. In *ECCV*, 2004. 6
- [2] Zhaowei Cai and Nuno Vasconcelos. Cascade R-CNN: Delving into high quality object detection. In *CVPR*, 2018. 2, 5
- [3] Joao Carreira, Pulkit Agrawal, Katerina Fragkiadaki, and Jitendra Malik. Human pose estimation with iterative error feedback. In *CVPR*, 2016. 2
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *CVPR*, 2017. 2, 3, 4, 6
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 6
- [6] Spyros Gidaris and Nikos Komodakis. Object detection via a multi-region and semantic segmentation-aware CNN model. In *ICCV*. 2
- [7] Spyros Gidaris and Nikos Komodakis. Attend refine repeat: Active box proposal generation via in-out localization. In *BMVC*, 2016. 2
- [8] Ross Girshick. Fast R-CNN. In *ICCV*, 2015. 1, 3
- [9] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 4
- [10] Georgia Gkioxari and Jitendra Malik. Finding action tubes. In *CVPR*, 2015. 2
- [11] Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Jimenez Rezende, and Daan Wierstra. DRAW: A recurrent neural network for image generation. In *ICML*, 2015. 2
- [12] Chunhui Gu, Chen Sun, Sudheendra Vijayanarasimhan, Caroline Pantofaru, David A Ross, George Toderici, Yeqing Li, Susanna Ricco, Rahul Sukthankar, and Cordelia Schmid. AVA: A video dataset of spatio-temporally localized atomic visual actions. In *CVPR*, 2018. 1, 2, 4, 5, 6, 8
- [13] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3D CNNs retrace the history of 2D CNNs and ImageNet? In *CVPR*, 2018. 2
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4
- [15] Fabian Heilbron, Juan Niebles, and Bernard Ghanem. Fast temporal activity proposals for efficient detection of human actions in untrimmed videos. In *CVPR*, 2016. 2
- [16] Rui Hou, Chen Chen, and Mubarak Shah. Tube convolutional neural network (T-CNN) for action detection in videos. In *ICCV*, 2017. 1, 2, 8
- [17] Vicky Kalogeiton, Philippe Weinzaepfel, Vittorio Ferrari, and Cordelia Schmid. Action tubelet detector for spatio-temporal action localization. In *ICCV*, 2017. 1, 2, 4, 5, 6, 7, 8
- [18] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, and Paul Natsev. The Kinetics human action video dataset. *arXiv:1705.06950*, 2017. 6
- [19] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. 6
- [20] Dong Li, Zhaofan Qiu, Qi Dai, Ting Yao, and Tao Mei. Recurrent tubelet proposal and recognition networks for action detection. In *ECCV*, 2018. 1, 2, 5, 8
- [21] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander Berg. SSD: Single shot multibox detector. In *ECCV*, 2016. 1
- [22] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael Jordan. Unsupervised domain adaptation with residual transfer networks. In *NeurIPS*, 2016. 4
- [23] Behrooz Mahasseni, Xiaodong Yang, Pavlo Molchanov, and Jan Kautz. Budget-aware activity detection with a recurrent policy network. In *BMVC*, 2018. 2
- [24] Mahyar Najibi, Mohammad Rastegari, and Larry Davis. G-CNN: An iterative grid based object detector. In *CVPR*, 2016. 2, 5
- [25] Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, 2015. 2
- [26] Xiaojiang Peng and Cordelia Schmid. Multi-region two-stream R-CNN for action detection. In *ECCV*, 2016. 2, 7, 8
- [27] Suman Saha, Gurbir Singh, Michael Sapienza, Philip Torr, and Fabio Cuzzolin. Deep learning for detecting multiple space-time action tubes in videos. In *BMVC*, 2016. 2, 5
- [28] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, 2014. 2
- [29] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2015. 3, 6
- [30] Gurbir Singh, Suman Saha, Michael Sapienza, Philip Torr, and Fabio Cuzzolin. Online real-time multiple spatiotemporal action localisation and prediction. In *ICCV*, 2017. 2, 5, 8
- [31] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv:1212.0402*, 2012. 6
- [32] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Kevin Murphy, Rahul Sukthankar, and Cordelia Schmid. Actor-centric relation network. In *ECCV*, 2018. 5, 8
- [33] Philippe Weinzaepfel, Zaid Harchaoui, and Cordelia Schmid. Learning to track for spatio-temporal action localization. In *ICCV*, 2015. 2
- [34] Huijuan Xu, Abir Das, and Kate Saenko. R-C3D: Region convolutional 3D network for temporal activity detection. In *ICCV*, 2017. 2
- [35] Xiaodong Yang, Pavlo Molchanov, and Jan Kautz. Multilayer and multimodal fusion of deep neural networks for video classification. In *ACM MM*, 2016. 2
- [36] Xiaodong Yang, Pavlo Molchanov, and Jan Kautz. Making convolutional networks recurrent for visual sequence learning. In *CVPR*, 2018. 2, 4
- [37] Zhenheng Yang, Jiyang Gao, and Ram Nevatia. Spatio-temporal action detection with cascade proposal and location anticipation. In *BMVC*, 2017. 2, 4, 8
- [38] Yuancheng Ye, Xiaodong Yang, and Yingli Tian. Discovering spatio-temporal action tubes. *JVCI*, 2019. 2, 8