

Cascaded Boundary Regression for Temporal Action Detection

Jiyang Gao
jiyangga@usc.edu
Zhenheng Yang
zhenheny@usc.edu
Ram Nevatia
nevatia@usc.edu

Institute for Robotics and Intelligent
Systems
University of Southern California
Los Angeles, CA, USA

Abstract

Temporal action detection in long videos is an important problem. State-of-the-art methods address this problem by applying action classifiers on sliding windows. Although sliding windows may contain an identifiable portion of the actions, they may not necessarily cover the entire action instance, which would lead to inferior performance. We adapt a two-stage temporal action detection pipeline with Cascaded Boundary Regression (CBR) model. Class-agnostic proposals and specific actions are detected respectively in the first and the second stage. CBR uses temporal coordinate regression to refine the temporal boundaries of the sliding windows. The salient aspect of the refinement process is that, inside each stage, the temporal boundaries are adjusted in a cascaded way by feeding the refined windows back to the system for further boundary refinement. We test CBR on THUMOS-14 and TVSeries, and achieve state-of-the-art performance on both datasets. The performance gain is especially remarkable under high IoU thresholds, *e.g.* $\text{map@tIoU}=0.5$ on THUMOS-14 is improved from 19.0% to 31.0%.

1 Introduction

Temporal action detection in long videos is an important and challenging problem, which has been receiving increasing attention recently. Given a long video, the task of action detection is to localize intervals where actions of interest take place and also predict the action categories.

Good progress has been achieved in action classification [21, 24], where the task is to predict action classes in "trimmed" videos. Current state-of-the-art methods [15, 21, 23] on action detection extend classification methods to detection by applying action classifiers on dense sliding windows. However, while sliding windows may contain an identifiable portion of the action, they do not necessarily cover the entire action instance or they could contain extraneous background frames, which may lead to inferior performance. Similar observations have also been made for use of sliding windows in object detection [19]. Inspired by object detection, Shou *et al.* [20] proposed a two-stage pipeline for action detection, called SCNN. In the first stage, it produces actionness scores for multi-scale sliding windows and outputs the windows with high scores as class-agnostic temporal proposals; in the second

stage, SCNN categorizes the proposals to specific actions. However, SCNN still suffers from the imprecision of sliding window intervals.

To improve temporal localization accuracy, recently a method called TURN [15] proposed to use temporal boundary regression. TURN takes sliding windows and their surrounding context as input and refines their temporal boundaries by learning a boundary regressor. We propose that the process of boundary estimation can be improved by deploying a cascade, where a regressed clip is fed back to the system for further refinement. As the system could observe different content in each round of refinement, the boundary can be refined gradually.

We adapt a two-stage action detection model with temporal coordinate regression. In the first stage, our model takes sliding windows as input, and outputs class-agnostic temporal proposals. In the second stage, our model detects actions based on the proposals. The salient aspect in our model is that, inside each stage, we propose to use Cascaded Boundary Regression (CBR) to adjust temporal boundaries in a regression cascade, where regressed clips are fed back to the system for further boundary refinement. The main differences compared with TURN are (1) the use of cascaded boundary regression, (2) a full pipeline for action detection (TURN only focus on proposal generation). We evaluate CBR on two challenging datasets: THUMOS-14 and TVSeries [16]. CBR outperforms state-of-the-art methods on both temporal action proposal generation and action detection tasks by a large margin. The performance gain is especially remarkable under high IoU thresholds, e.g $\text{map@tIoU}=0.5$ on THUMOS-14 is improved from 19.0% to 31.0%.

Our contributions are two-fold:

(1) We propose a Cascaded Boundary Regression method for temporal boundary estimation, which is shown to be effective on both proposal generation and action detection.

(2) We evaluate CBR on both proposal generation and action detection, and achieve state-of-the-art performance on both THUMOS-14 and TVSeries [16].

2 Related Work

Temporal action detection, temporal proposal generation and object detection are related to our work, we will introduce these three topics in this section.

Temporal Action Detection Temporal action localization has been received much attentions recently. S-CNN [17] presented action detection framework which involves two stages: the first stage uses proposal network to generate temporal action proposals; the second stage classifies the proposals with localization network, which is trained using classification and localization loss. Singh *et al.* [18] extended two-stream [17] framework to multi-stream network and use bi-directional LSTM networks to encode temporal information, they achieved state-of-the-art performance on MPII-Cooking dataset [18]. Ma *et al.* [19] addressed the problem of early action detection. They proposed to train a LSTM network with ranking loss and merge the detection spans based on the frame-wise prediction scores generated by the LSTM. Sun *et al.* [20] proposed to transfer knowledge from web images to address temporal detection in untrimmed web videos.

Temporal Action Proposal. Similar to object proposal generation, temporal proposal generation aims to produce class-agnostic proposals efficiently and accurately. Sparse-prop [21] presented a method that use STIPs [22] and dictionary learning for class-independent proposal generation. SCNN-prop [23] presented a method that fine-tunes 3D convolutional network [24] for binary proposal classification. DAPs [6] used LSTM networks to encode a video stream and produce proposals inside the video stream. Gao *et al.* [5] proposed a

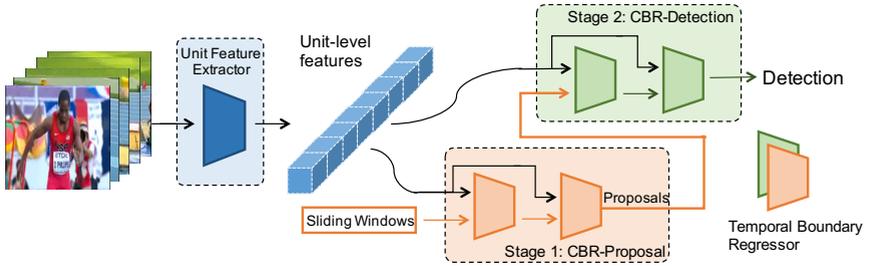


Figure 1: Architecture of two-stage action detection pipeline with Cascaded Boundary Regression (CBR)

method, called TURN, to use unit-level temporal coordinate regression to refine the temporal action boundary.

Object Proposals and Object Detection. Recent successful object detection frameworks [6, 7, 12] are built on high quality object proposals. SelectiveSearch [25] and Edgebox [24] rely on hand-crafted low-level visual features. R-CNN [7] and Fast R-CNN [6] use this type of object proposals as input. RPNs [12] proposed to use anchor boxes and spatial regression for object proposal generation, which is based on ConvNet’s conv-5 featmap. YOLO [16] proposed to divide the input image into grid cells and estimate the object bounding box by coordinate regression. SSD [13] further adopted multi-scale grid cells to predict bounding boxes.

3 Methods

In this section, we describe the two-stage Cascaded Boundary Regression (CBR) network and the training procedure, its architecture is shown in Figure 1. This work is inspired by [6], so we use similar notations in the following.

3.1 Video Unit Feature Extraction

A video V containing T frames, $V = \{f_i\}_1^T$, is divided into T/u_f consecutive video units, where u_f is the number of frames in a unit. A video unit can be represented as $u = \{t_i\}_{f_s}^{f_s+u_f-1}$, where f_s is the starting frame, $f_s + u_f - 1$ is the ending frame f_e . Units do not overlap with each other. Each unit is processed by a visual encoder E_v to get a unit-level representation $f_u = E_v(u)$. In our experiments, C3D [24] and two-stream CNN models [21] are investigated. Details are given in Section 4.

3.2 Video Clip Modeling

A clip c is composed of units, $c = \{u_j\}_{u_s}^{u_s+c_u-1}$, where u_s is the index of starting unit and c_u is the number of units inside c . $u_e = u_s + c_u - 1$ is the index of ending unit u_e , and $\{u_j\}_{u_s}^{u_e}$ are called *internal units* of c . Besides the internal units, *surrounding units* for c are also modeled. $\{u_j\}_{u_s-n_{ctx}}^{u_s-1}$ and $\{u_j\}_{u_e+1}^{u_e+n_{ctx}}$ are the surrounding units before and after c respectively, n_{ctx} is the number of units we consider. The surrounding units provide temporal *context* for clips, which are important for temporal boundary inferring. Internal feature and context features

are pooled from unit-level features separately by mean pooling operation P . The final feature f_c for a clip is the concatenation of context features and the internal feature.

$$f_c = P(\{u_j\}_{s_u-n_{ctx}}^{s_u}) \parallel P(\{u_j\}_{s_u}^{e_u}) \parallel P(\{u_j\}_{e_u+n_{ctx}}^{e_u}) \quad (1)$$

where \parallel represents vector concatenation. We scan a video by multi-scale temporal sliding windows. The temporal sliding windows are modeled by two parameters: window length l_i and window overlap o_i . Note that, although multi-scale clips would have temporal overlaps, the clip-level features are computed from unit-level features, which are only calculated once.

3.3 Temporal Coordinate Regression

We first introduce temporal coordinate regression and then introduce the two-stage proposal and detection pipeline. Our goal is to design a method which is able to estimate the temporal boundaries of actions. For spatial boundary regression, previous works [6, 17] use parameterized coordinate offsets, that is, the boundary coordinates are first parameterized by the central coordinates and the size (*i.e.* length and width) of the bounding box, the offsets are calculated based on these parameterized coordinates. In the temporal coordinate settings, the parameterization offsets could be represented as,

$$o_x = (x^{gt} - x^{clip})/l_{clip}, \quad o_l = \log(l^{gt}/l^{clip}) \quad (2)$$

where x and l denote the clip’s center coordinate and clip length respectively. Variables x^{gt}, x^{clip} are for ground truth clip and test clip (likewise for l).

Instead of using parameterization, non-parameterized offset is to use the start and end coordinates directly. Specifically, there are two levels of coordinates: frame-level and unit-level. The frame-level coordinate is the index of the frame f_i ; the unit-level coordinate is the index of the unit u_j . For an action instance, the ground truth start and end coordinates t_s^{gt} and t_e^{gt} are usually annotated in seconds, which could be always transferred to frame-level (multiplied by FPS) f_s^{gt} and f_e^{gt} , the unit-level ground truth coordinates are calculated by rounding:

$$u_s^{gt} = \langle f_s^{gt}/u_f \rangle, \quad u_e^{gt} = \langle f_e^{gt}/u_f \rangle \quad (3)$$

where $\langle \cdot \rangle$ represents rounding, u_f is the frame number in a unit. The non-parameterized regression offsets are

$$o_s = s^{clip} - s^{gt}, \quad o_e = e^{clip} - e^{gt} \quad (4)$$

where s^{clip}, e^{clip} are the start and end coordinates of the input clip, which could be at frame-level or unit-level. s^{gt}, e^{gt} are the coordinates for the matched ground truth action instance. The intuition behind unit-level coordinate regression is that, as the basic unit-level features are extracted to encode n_u frames, the feature may not be discriminative enough to regress the coordinates at frame-level. Comparing with frame-level regression, unit-level coordinate regression is easier to learn, though with coarser boundaries.

3.4 Two-Stage Proposal and Detection Pipeline

Inspired by the proposal and detection pipeline in object detection, we design a two-stage pipeline for temporal action detection, in which class-agnostic proposals and class-specific detections are generated respectively, shown in Figure 1. In both stages, temporal coordinate regression is used to infer temporal action boundaries. Specifically, given a clip $c = \langle s, e \rangle$,

it is first processed by the proposal network, which outputs two boundary regression offsets $\langle o_s, o_e \rangle$ and an actionness score p indicating whether c is an action instance. If the output p is higher than a threshold θ , the detection network takes $c' = \langle s', e' \rangle$ (new temporal boundaries) as input and generates $n + 1$ softmax scores $p^z, x \in [1, n + 1]$ and n pairs of boundary offsets $\langle o_s^z, o_e^z \rangle, x \in [1, n]$, where n is the number of action categories.

3.5 Cascaded Boundary Regression

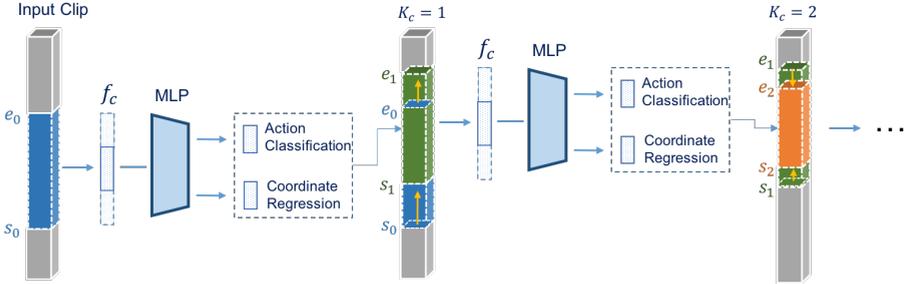


Figure 2: Unrolled model of Cascaded Boundary Regression (CBR), the parameters of the MLPs are shared.

In each stage (*i.e.* the proposal and detection stages), boundary regression is applied in a cascaded manner—the output boundaries are fed back as input to the network for further refinement, as shown in Figure 2. For proposal network, given a input clip $c = \langle s, e \rangle$, the output clip $c^1 = \langle s^1, e^1 \rangle$ is fed back as input to do a second round of refinement, and the second output is $c^2 = \langle s^2, e^2 \rangle$. The iteration process takes K_c^p steps, the final boundaries and the actionness score for c are

$$c_{K_c^p} = \langle s_{K_c^p}, e_{K_c^p} \rangle, \quad p = \prod_{i=1}^{K_c^p} p_i \quad (5)$$

The cascade process of detection network is similar to that of proposal network. For detection network, it outputs n pairs of temporal boundary offsets and $n + 1$ category scores. Among the n non-background categories, we take the category with the highest score as the prediction x , and pick the corresponding boundary offsets $\langle o_{s,x}, o_{e,x} \rangle$. The refined clip $c^1 = \langle s_x^1, e_x^1 \rangle$ is fed back into the network. After K_c^d steps, the final boundaries and score for the predicted category x are

$$c_{K_c^d} = \langle s_{K_c^d}^z, e_{K_c^d}^z \rangle, \quad p = \prod_{i=1}^{K_c^d} p_i^z \quad (6)$$

The proposal network and the detection network are trained separately, details could be found in the next section. In either stage (proposal or detection), each cascade step could be trained separately, but here we have chosen to use the same network parameters in each step for simplicity.

3.6 Loss Function

To train CBR, we collect training samples from sliding windows, whose parameters (overlap and window length) will be introduced in Section 4. A class label is assigned to a sliding window if: (1) the window clip with the highest temporal Intersection over Union (tIoU) overlaps with a ground truth clip; or (2) the window clip has tIoU larger than 0.5 with any of the ground truth clips. Note that, a single ground truth clip may assign its label to multiple window clips. Negative labels are assigned to non-positive clips whose tIoU is equal to 0.0 (*i.e.* no overlap) for all ground truth clips. We design a multi-task loss L to jointly train classification and coordinate regression.

$$L = L_{cls} + \lambda L_{reg} \quad (7)$$

where L_{cls} is the loss for classification, which is a standard cross-entropy loss. For proposal network, L_{cls} is a binary classification cross-entropy loss; for detection network, L_{cls} is a standard multi-class cross-entropy loss. L_{reg} is for temporal coordinate regression and λ is a hyper-parameter, which is set empirically. The regression loss is

$$L_{reg} = \frac{1}{N} \sum_{i=1}^N \sum_{z=1}^n l_i^z [R(\hat{o}_{s,i}^z - o_{s,i}^z) + R(\hat{o}_{e,i}^z - o_{e,i}^z)] \quad (8)$$

where R is $L1$ distance, N is batch size and n is the total number of categories, l_i^z is the label, when the i th sample is from category z , $l_i^z = 1$, otherwise, $l_i^z = 0$. \hat{o} is the regression estimate offset, and o is the ground truth offset. For parameterized offsets, $o_{s,i}^z$ and $o_{e,i}^z$ are replaced by $o_{s,i}^z$ and $o_{l,i}^z$.

The learning rate and batch size are set as 0.005 and 128 respectively. We use the Adam [14] optimizer to train CBR. The ratio of sample numbers of background to non-background in a mini-batch is set to be 10 for training proposal network. For training detection network, the number background samples are equal to the average sample numbers of all categories. λ is set to 2 for both proposal and detection network.

4 Evaluation

We evaluate the effectiveness of the proposed Cascaded Boundary Regression (CBR) on standard benchmarks THUMOS-14 and TVSeries for both temporal action proposal generation and action detection.

Unit-level Feature Extraction. C3D unit-level features: The C3D model is pre-trained on Sports1M [14], we uniformly sample 16 frames in a unit and then input them into C3D; the output of *fc6* layer is used as unit-level feature. Two-stream features: We use the two-stream model [26] that is pre-trained on ActivityNet v1.3 training set. In each unit, the central frame is sampled to calculate the appearance CNN feature, which is the output of "Flatten_673" layer in ResNet [8]. For the motion feature, we sample 6 consecutive frames at the center of a unit and calculate optical flows [9] between them; these flows are then fed into the pretrained BN-Inception model [9, 26] and the output of "global pool" layer is extracted. The motion features and the appearance features are concatenated into 4096-dimensional vectors, which are used as unit-level features.

4.1 Experiments on THUMOS-14

We first introduce the datasets and the experiment setup, then discuss the experimental results on THUMOS-14.

Dataset. The temporal action localization part of THUMOS-14 contains over 20 hours of videos from 20 sport classes. There are 200 untrimmed videos in validation set and 213 untrimmed videos in test set. The training set of THUMOS-14 contains only trimmed videos. We train our model on the validation set and test it on the test set.

Experimental setup. We perform the following experiments on THUMOS-14: (1) explore components in the proposed framework: (a) parametrized offsets vs non-parameterized unit-level regression vs non-parameterized frame-level regression, (b) cascaded steps for boundary regression; (2) comparison with state-of-the-art approaches. The unit size u_f is 16, the surrounding unit number n_{ctx} is set to 4. The sliding window lengths and overlaps are $\{16(16), 32(16), 64(16), 128(32), 256(64), 512(128)\}$, where the numbers out of brackets are lengths of sliding windows, and the numbers in brackets are the corresponding overlaps of the sliding windows.

Temporal coordinate regression. To explore which type of coordinate offsets is most effective for boundary regression in temporal action detection, we test three types: (a) parameterized coordinate offsets, which are similar to the ones in object detection [14], (b) non-parameterized frame-level coordinate offsets and (c) non-parameterized unit-level coordinate offsets. The results of temporal action detection are listed in Table 1. The cascade step K_c is set to be 1 for both proposal stage and detection stage. Both C3D feature and Two-stream CNN feature are tested.

Table 1: Comparison of different coordinate offsets on action localization (% mAP@tIoU=0.5): parameterized, non-parameterized frame-level, non-parameterized unit-level. The performance with no boundary regression is also listed.

	no regression	parameterized	non-para, frame-level	non-para, unit-level
CBR-C3D	16.7	19.4	18.8	20.5
CBR-TS	22.3	26.1	25.3	27.7

We can see that all three regression offsets provide improvement over "no regression". Unit-level offsets are more effective than frame-level offsets; we think the reason is that, the features are extracted at unit level, frame-level coordinates contain redundant information, which may make the regression task more difficult. The performance of parameterized coordinate offsets is lower than that of non-parameterized unit-level offsets. We think that the reason is that unlike objects which can be re-scaled in images with camera projection, actions' time spans can not be easily re-scaled in videos, although the time spans of the same action can be varied in different videos. Therefore, "time" itself work as a standard scale for action instances.

Cascaded boundary regression. We explore the effects of boundary regression cascade. Cascade step K_c^p and K_c^d are the number of boundary regression conducted in proposal stage and detection stage respectively. The results are shown in Table 2 and Table 3. We investigate the cascade step with C3D feature and two-stream CNN feature. Non-parameterized unit-level coordinate offset is adopted.

For the proposal network (shown in Table 2), we can see that cascaded boundary regression increase the performance from 42.7 to 45.0 for two-stream features, and from 38.6 to 39.6 for C3D features. When $K_c^p = 3$, two-stream CBR achieves the best performance, and after the performance peak, the performance drops slightly.

Table 2: Comparison of cascaded step $K_c^p = 1, 2, 3, 4$ for temporal proposal generation (% AR@F=1.0) on THUMOS-14.

	$K_c^p = 1$	$K_c^p = 2$	$K_c^p = 3$	$K_c^p = 4$
CBR-C3D	38.6	39.6	39.4	37.8
CBR-TS	42.7	44.5	45.2	44.8

Table 3: Comparison of cascaded step $K_c^d = 1, 2, 3, 4$ for temporal action detection (% mAP@tIoU=0.5) on THUMOS-14.

	$K_c^d = 1$	$K_c^d = 2$	$K_c^d = 3$	$K_c^d = 4$
CBR-C3D	21.5	22.7	22.4	22.2
CBR-TS	28.4	31.0	30.5	30.2

To test the effects of cascaded boundary regression for action detection, we fix $K_c^p = 3$. As shown in Table 3, we observe a similar trend as in proposal generation: when $K_c^d = 2$, CBR increases the performance from 28.4 to 31.0 for two-stream features, and from 21.5 to 22.7 for C3D features. After $K_c^d = 2$, the performance becomes saturated.

Table 4: Comparison with state-of-the-art on temporal action proposal generation. Average Recall at Proposal Frequency (AR@F=1.0) performance are reported.

Method	Sparse-prop[10]	DAPs[9]	SCNN-prop[12]	TURN-FL [8]	CBR-FL	CBR-TS
AR@AN=200	32.3	34.1	37.2	42.8	43.5	44.2
AR@F=1.0	33.3	35.7	38.3	43.5	44.4	45.2

Comparison with state-of-the-art on temporal proposal generation. We compare CBR-P with state-of-the-art methods on temporal action proposal generation, including SCNN-prop [[12](#)], DAPs[[9](#)] and Sparse-prop [[10](#)] and TURN [[8](#)]. The results are shown in Table 4. To fairly compare with TURN, we also provide the performance using only optical flow CNN features, which is the same for TURN-FL. We can see that CBR-FL outperforms state-of-the-art (TURN-FL) and CBR-TS provides further improvement over CBF-FL.

Comparison with state-of-the-art on temporal action detection. We compare our method with other state-of-the-art temporal action localization methods on THUMOS-14, the results are shown in Table 5. We compare with the challenge results [[15](#)], and recent methods including based on segment window C3D [[12](#)], score pyramids [[18](#)] and deep recurrent reinforcement learning [[17](#)]. Both SCNN and CBR-C3D are based on C3D features, we can see that CBR-C3D outperforms SCNN at all tIoU thresholds, especially at high tIoU, which shows the effectiveness of CBR. If two-stream features are adopted, CBR outperforms state-of-the-art methods by 12% at tIoU=0.5.

4.2 Experiments on TVSeries

We first introduce the datasets and the experiment setup, then discuss the experimental results on TVSeries [[9](#)].

Dataset. The TVSeries Dataset [[9](#)] is a realistic, large-scale dataset for temporal action detection, which contains 16 hours of videos (27 episodes) from six recent popular TV series. 30 daily life action categories are defined in TVSeries, such as "close door", "drive car", "wave". There are totally 6231 action instances annotated with start and end times and action categories in the dataset. The train/validation/test sets contain 13/7/7 episodes respectively.

Experimental setup. We test the cascaded steps of boundary regression for both proposal generation and action detection, and then compare with state-of-the-art performance on

Table 5: Temporal action detection performance (mAP %) comparison at different tIoU thresholds on THUMOS-14.

tIoU	0.1	0.2	0.3	0.4	0.5	0.6	0.7
Oneata <i>et al.</i> [13]	36.6	33.6	27.0	20.8	14.4	8.5	3.2
Yeung <i>et al.</i> [27]	48.9	44.0	36.0	26.4	17.1	-	-
Yuan <i>et al.</i> [28]	51.4	42.6	33.6	26.1	18.8	-	-
SCNN [22]	47.7	43.5	36.3	28.7	19.0	10.3	5.3
CBR-C3D	48.2	44.3	37.7	30.1	22.7	13.8	7.9
CBR-TS	60.1	56.7	50.1	41.3	31.0	19.1	9.9

action detection. The unit size u_f is 6, the surrounding unit number n_{ctx} is set to 4. The sliding window lengths and overlaps are $\{12(6), 24(6), 48(12), 72(18), 96(24), 192(48), 384(96)\}$, where the numbers out of brackets are lengths of sliding window, and the numbers in brackets are the corresponding overlaps of the sliding window.

Cascaded boundary regression. We explore the effects of cascaded boundary regression on TVSeries. Note that K_c is the cascaded step. We investigate the cascaded step with two-stream CNN features.

Table 6: Comparison of cascaded step $K_c = 0, 1, 2, 3, 4$ for temporal action detection (% mAP@tIoU=0.5) on TVSeries. $K_c = 0$ means that the system only do classification, no boundary regression.

	$K_c = 0$	$K_c = 1$	$K_c = 2$	$K_c = 3$	$K_c = 4$
Proposal (AR@F=1.0)	20.4	24.3	25.6	26.1	25.9
Detection (mAP@IoU=0.2)	6.2	8.8	9.5	9.2	9.0

As shown in Table 6, comparing with $K_c = 0$ and $K_c = 1$, we can see that temporal coordinate regression brings a big improvement, which shows its effectiveness. We can also see that when $K_c = 3$, CBR achieves the best performance for proposal network. To test detection network, we fix $K_c^p = 3$, the results show that when $K_c^d = 2$ CBR achieves the best performance for action detection. After the peak, the performance starts to decrease. The performance distribution of cascaded step is consistent with THUMOS-14.

Comparison with state-of-the-art on action detection. We compare CBR with state-of-the-art performance on TVSeries in Table 7. Overall, we can see that TVSeries is a more challenging dataset than THUMOS-14. To provide another comparison, we train SVM classifiers based on two-stream features, which is shown in Table 7 as **SVM-TS**. The SVM classifiers are trained and tested using the same samples as CBR, which are described in Section 3.6; clip-level features are mean-pooled from unit-level features. We can see that with the same features, CBR outperforms the SVM-based classifiers by 3.5% at tIoU=0.2. At tIoU = 0.2, CBR achieves 9.5, while the state-of-the-art method FV [2] only achieves 4.9. We also report mAP performance at tIoU = 0.1 and tIoU = 0.3, which are 11.0 and 7.9 respectively.

Table 7: Temporal action detection performance (mAP %) comparison at different tIoU thresholds on TVSeries.

tIoU	CNN [2]	LSTM [2]	FV [2]	SVM-TS	CBR-TS
0.1	-	-	-	7.3	11.0
0.2	1.1	2.7	4.9	6.0	9.5
0.3	-	-	-	4.6	7.9

5 Conclusion

We present a novel two-stage action detection pipeline with Cascaded Boundary Regression (CBR), which achieves state-of-the-art performance on standard benchmarks. In the first stage, temporal proposals are generated; based on the proposals, actions are detected in the second stage. Cascaded boundary regression are conducted in both stages. Detailed experiments and analysis on cascaded steps are conducted, which show the effectiveness of CBR for both temporal proposal generation and action detection. Different temporal regression offset settings are also investigated and discussed. State-of-the-art performance has been achieved on both THUMOS-14 and TVSeires dataset.

References

- [1] Fabian Caba Heilbron, Juan Carlos Niebles, and Bernard Ghanem. Fast temporal activity proposals for efficient detection of human actions in untrimmed videos. In *CVPR*, 2016.
- [2] Roeland De Geest, Efstratios Gavves, Amir Ghodrati, Zhenyang Li, Cees Snoek, and Tinne Tuytelaars. Online action detection. In *ECCV*, 2016.
- [3] Victor Escorcia, Fabian Caba Heilbron, Juan Carlos Niebles, and Bernard Ghanem. Daps: Deep action proposals for action understanding. In *ECCV*, 2016.
- [4] Gunnar Farnebäck. Two-frame motion estimation based on polynomial expansion. In *Scandinavian conference on Image analysis*. Springer, 2003.
- [5] Jiyang Gao, Zhenheng Yang, Chen Sun, Kan Chen, and Ram Nevatia. Turn tap: Temporal unit regression network for temporal action proposals. *arXiv preprint arXiv:1703.06189*, 2017.
- [6] Ross Girshick. Fast r-cnn. In *ICCV*, 2015.
- [7] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [9] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. 2015.
- [10] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014.
- [11] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [12] Ivan Laptev. On space-time interest points. *International Journal of Computer Vision*, 64(2-3):107–123, 2005.

- [13] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, and Scott Reed. Ssd: Single shot multibox detector. In *ECCV*, 2016.
- [14] Shugao Ma, Leonid Sigal, and Stan Sclaroff. Learning activity progression in lstms for activity detection and early detection. In *CVPR*, 2016.
- [15] Dan Oneata, Jakob Verbeek, and Cordelia Schmid. The lear submission at thumos 2014. In *ECCV Workshop*, 2014.
- [16] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016.
- [17] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.
- [18] Marcus Rohrbach, Sikandar Amin, Mykhaylo Andriluka, and Bernt Schiele. A database for fine grained activity detection of cooking activities. In *CVPR*, 2012.
- [19] Pierre Sermanet, David Eigen, Xiang Zhang, Michaël Mathieu, Rob Fergus, and Yann LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. In *ICLR*, 2014.
- [20] Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *CVPR*, 2016.
- [21] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014.
- [22] Bharat Singh, Tim K. Marks, Michael Jones, Oncel Tuzel, and Ming Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *CVPR*, 2016.
- [23] Chen Sun, Sanketh Shetty, Rahul Sukthankar, and Ram Nevatia. Temporal localization of fine-grained actions in videos by domain transfer from web images. In *ACM MM*, 2015.
- [24] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.
- [25] Jasper RR Uijlings, Koen EA van de Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *International journal of computer vision*, 104(2):154–171, 2013.
- [26] Yuanjun Xiong, Limin Wang, Zhe Wang, Bowen Zhang, Hang Song, Wei Li, Dahua Lin, Yu Qiao, Luc Van Gool, and Xiaoou Tang. Cuhk & ethz & siat submission to activitynet challenge 2016. *arXiv preprint arXiv:1608.00797*, 2016.
- [27] Serena Yeung, Olga Russakovsky, Greg Mori, and Li Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *CVPR*, 2016.
- [28] Jun Yuan, Bingbing Ni, Xiaokang Yang, and Ashraf A. Kassim. Temporal action localization with pyramid of score distribution features. In *CVPR*, 2016.
- [29] C Lawrence Zitnick and Piotr Dollár. Edge boxes: Locating object proposals from edges. In *ECCV*, 2014.