

Tube Convolutional Neural Network (T-CNN) for Action Detection in Videos

Rui Hou, Chen Chen, Mubarak Shah

Center for Research in Computer Vision (CRCV), University of Central Florida (UCF)

houray@gmail.com, chenchen870713@gmail.com, shah@crcv.ucf.edu

Abstract

Deep learning has been demonstrated to achieve excellent results for image classification and object detection. However, the impact of deep learning on video analysis has been limited due to complexity of video data and lack of annotations. Previous convolutional neural networks (CNN) based video action detection approaches usually consist of two major steps: frame-level action proposal generation and association of proposals across frames. Also, most of these methods employ two-stream CNN framework to handle spatial and temporal feature separately. In this paper, we propose an end-to-end deep network called Tube Convolutional Neural Network (T-CNN) for action detection in videos. The proposed architecture is a unified deep network that is able to recognize and localize action based on 3D convolution features. A video is first divided into equal length clips and next for each clip a set of tube proposals are generated based on 3D Convolutional Network (ConvNet) features. Finally, the tube proposals of different clips are linked together employing network flow and spatio-temporal action detection is performed using these linked video proposals. Extensive experiments on several video datasets demonstrate the superior performance of T-CNN for classifying and localizing actions in both trimmed and untrimmed videos compared to state-of-the-arts.

1. Introduction

The goal of action detection is to detect every occurrence of a given action within a long video, and to localize each detection both in space and time. Deep learning learning based approaches have significantly improved video action recognition performance. Compared to action *recognition*, action *detection* is a more challenging task due to flexible volume shape and large spatio-temporal search space.

Previous deep learning based action detection approaches first detect frame-level action proposals by popular proposal algorithms [5, 30] or by training proposal networks [19]. Then the frame-level action proposals are associated across frames to form final action detection through



Figure 1: Overview of the proposed Tube Convolutional Neural Network (T-CNN).

tracking based approaches [32]. Moreover, in order to capture both spatial and temporal information of an action, twostream networks (a spatial CNN and a motion CNN) are used. In this manner, the spatial and motion information are processed separately.

Region Convolution Neural Network (R-CNN) for object detection in images was proposed by Girshick *et al.* [4]. It was followed by a fast R-CNN proposed in [3], which includes the classifier as well. Later, faster R-CNN [20] was developed by introducing a region proposal network. It has been extensively used to produce excellent results for object detection in images. A natural generalization of the R-CNN from 2D images to 3D spatio-temporal volumes is to study their effectiveness for the problem of action detection in videos. A straightforward spatio-temporal generalization of the R-CNN approach would be to treat action detection in videos as a set of 2D image detections using faster R-CNN. However, unfortunately, this approach does not take the temporal information into account and is not sufficiently expressive to distinguish between actions.

Inspired by the pioneering work of faster R-CNN, we

propose Tube Convolutional Neural Network (T-CNN) for action detection. To better capture the spatio-temporal information of video, we exploit 3D ConvNet for action detection, since it is able to capture motion characteristics in videos and shows promising result on video action recognition. We propose a novel framework by leveraging the descriptive power of 3D ConvNet. In our approach, an input video is divided into equal length clips first. Then, the clips are fed into Tube Proposal Network (TPN) and a set of tube proposals are obtained. Next, tube proposals from each video clip are linked according to their actionness scores and overlap between adjacent proposals to form a complete tube proposal for spatio-temporal action localization in the video. Finally, the Tube-of-Interest (ToI) pooling is applied to the linked action tube proposal to generate a fixed length feature vector for action label prediction.

Our work makes the following contributions:

• We propose an end-to-end deep learning based approach for action detection in videos. It directly operates on the original videos and captures spatio-temporal information using a single 3D network to perform action localization and recognition based on 3D convolution features. To the best of our knowledge, it is the first work to exploit 3D ConvNet for action detection.

• We introduce a Tube Proposal Network, which leverages skip pooling in temporal domain to preserve temporal information for action localization in 3D volumes.

• We propose a new pooling layer – **Tube-of-Interest** (ToI) pooling layer in T-CNN. The ToI pooling layer is a 3D generalization of Region-of-Interest (RoI) pooling layer of R-CNN. It effectively alleviates the problem with variable spatial and temporal sizes of tube proposals. We show that ToI pooling can greatly improve the recognition results.

• We extensively evaluate our T-CNN for action detection in both trimmed videos from UCF-Sports, J-HMDB and UCF-101 datasets and untrimmed videos from THU-MOS'14 dataset and achieve state-of-the-art performance.

2. Related Work

Convolutional Neural Networks (CNN) have been demonstrated to achieve excellent results for action recognition [17, 18]. Karpathy *et al.* [14] explore various frame-level fusion methods over time. Ng *et al.* [31] use recurrent neural network employing the CNN feature. Since these approaches only use frame based CNN features, the temporal information is neglected. Simonyan *et al.* [22] propose the two-stream CNN approach for action recognition. Besides a classic CNN which takes images as an input, it has a separate network for optical flow. Moreover, Wang *et al.* fuse the trajectories and CNN features. Although these methods, which take hand-crafted temporal feature as a separate stream, show promising performance on action recognition, however, they do not employ end to end deep network and

require separate computation of optical flow and optimization of the parameters. 3D CNN is a logical solution to this issue. Ji *et al.* [9] propose a 3D CNN based human detector and head tracker to segment human subjects in videos. Tran *et al.* [27] leverage 3D CNN for large scale action recognition problem. Sun *et al.* [25] propose a factorization of 3D CNN and exploit multiple ways to decompose convolutional kernels. However, to the best of our knowledge, we are the first ones to exploit 3D CNN for action *detection*.

Compared to action recognition, action detection is a more challenging problem [2, 7, 29], which has been an active area of research. Ke et al. [15] present an approach for event detection in crowded videos. Tian et al. [26] develop Spatio-temporal Deformable Parts Model [1] to detect actions in videos. Jain et al. [6] and Soomro et al. [23] use supervoxel and selective search to localize the action boundaries. Recently, researchers have leveraged the power of deep learning for action detection. Authors in [5] extract frame-level action proposals using selective search and link them using Viterbi algorithm. While in [30] frame-level action proposals are obtained by EdgeBox and linked by a tracking algorithm. Two-stream R-CNNs for action detection is proposed in [19], where a spatial Region Proposal Network (RPN) and a motion RPN are used to generate frame-level action proposals. However, these deep learning based approaches detect actions by linking frame-level action proposals and treat the spatial and temporal features of a video separately by training two-stream CNN. Therefore, the temporal consistency in videos is not well explored in the network. In contrast, we determine action tube proposals directly from input videos and extract compact and more effective spatio-temporal features using 3D CNN.

For object detection in images, Girshick et al. propose Region CNN (R-CNN) [4]. In their approach region proposals are extracted using selective search. Then the candidate regions are warped to a fixed size and fed into ConvNet to extract CNN features. Finally, SVM model is trained for object classification. A fast version of R-CNN, Fast R-CNN, is presented in [3]. Compared to the multi-stage pipeline of R-CNN, fast R-CNN incorporates object classifier in the network and trains object classifier and bounding box regressor simultaneously. Region of interest (RoI) pooling layer is introduced to extract fixed-length feature vectors for bounding boxes with different sizes. Recently, faster R-CNN is proposed in [20]. It introduces a RPN (Region Proposal Network) to replace selective search for proposal generation. RPN shares full image convolutional features with the detection network, thus the proposal generation is almost cost-free. Faster R-CNN achieves state-ofthe-art object detection performance while being efficient during testing. Motivated by its high performance, in this paper we explore generalizing faster R-CNN from 2D image regions to 3D video volumes for action detection.

name	kernel dims	output dims
	$(d \times h \times w)$	$(C \times D \times H \times W)$
conv1	$3 \times 3 \times 3$	$64 \times 8 \times 300 \times 400$
max-pool1	$1 \times 2 \times 2$	$64\times8\times150\times200$
conv2	$3 \times 3 \times 3$	$128\times8\times150\times200$
max-pool2	$2 \times 2 \times 2$	$128 \times 4 \times 75 \times 100$
conv3a	$3 \times 3 \times 3$	$256\times 4\times 75\times 100$
conv3b	$3 \times 3 \times 3$	$256\times 4\times 75\times 100$
max-pool3	$2 \times 2 \times 2$	$256\times2\times38\times50$
conv4a	$3 \times 3 \times 3$	$512\times2\times38\times50$
conv4b	$3 \times 3 \times 3$	$512 \times 2 \times 38 \times 50$
max-pool4	$2 \times 2 \times 2$	$512 \times 1 \times 19 \times 25$
conv5a	$3 \times 3 \times 3$	$512 \times 1 \times 19 \times 25$
conv5b	$3 \times 3 \times 3$	$512 \times 1 \times 19 \times 25$
toi-pool2*	-	$128 \times 8 \times 8 \times 8$
toi-pool5	-	$512 \times 1 \times 4 \times 4$
1x1 conv	-	8192
fc6	-	4096
fc7	-	4096

Table 1: T-CNN architecture. We refer kernel with shape $d \times h \times w$ where d is the kernel depth, h and w are height and width, and output matrix with shape $C \times D \times H \times W$ where C is number of channels, D is the number of frames, H and W are the height and width. toi-pool2 exists in TPN.

3. Generalizing R-CNN from 2D to 3D

Generalizing R-CNN from 2D image regions to 3D video tubes is challenging due to the asymmetry between space and time. Different from images which can be cropped and reshaped into a fixed size, videos vary widely in temporal dimension. Therefore, we divide input videos into fixed length (8 frames) clips, so that video clips can be processed with a fixed-size ConvNet architecture. Also, clip based processing mitigates the cost of GPU memory.

To better capture the spatio-temporal information in video, we exploit 3D CNN for action proposal generation and action recognition. One advantage of 3D CNN over 2D CNN is that it captures motion information by applying convolution in both time and space. Since 3D convolution and 3D max pooling are utilized in our approach, not only in the spatial dimension but also in the temporal dimension, the size of video clip is reduced while distinguishable information is concentrated. As demonstrated in [27], the temporal pooling is important for recognition task since it better models the spatio-temporal information of video and reduces some background noise. However, the temporal order information is lost. That means if we arbitrarily change the order of the frames in a video clip, the resulting 3D max pooled feature cube will be the same. This is problematic in action detection, since it relies on the feature cube to get bounding boxes for the original frames. To this end, we incorporate temporal skip pooling to retain temporal order information residing in the original frames. More details are provided in the next section.

Since a video is processed clip by clip, action tube proposals with various spatial and temporal sizes are generated for different clips. These clip proposals need to be linked into a tube proposal sequence, which is used for action label prediction and localization. To produce a fixed length feature vector, we propose a new pooling layer - Tube-of-Interest (ToI) pooling layer. The ToI pooling layer is a 3D generalization of Region-of-Interest (RoI) pooling layer of R-CNN. The classic max pooling layer defines the kernel size, stride and padding which determines the shape of the output. In contrast, for RoI pooling layer, the output shape is fixed first, then the kernel size and stride are determined accordingly. Compared to RoI pooling which takes 2D feature map and 2D regions as input, ToI pooling deals with feature cube and 3D tubes. Denote the size of a feature cube as $d \times h \times w$, where d, h and w respectively represent the depth, height and width of the feature cube. A ToI in the feature cube is defined by a *d*-by-4 matrix, which is composed of d boxes distributed in all the frames. The boxes are defined by a four-tuple $(x_1^i, y_1^i, x_2^i, y_2^i)$ that specifies the top-left and bottom-right corners in the *i*-th feature map. Since the *d* bounding boxes may have different sizes, aspect ratios and positions, in order to apply spatio-temporal pooling, pooling in spatial and temporal domains are performed separately. First, the $h \times w$ feature maps are divided into $H \times W$ bins, where each bin corresponds to a cell with size of approximately $h/H \times w/W$. In each cell, max pooling is applied to select the maximum value. Second, the spatially pooled d feature maps are temporally divided into Dbins. Similar to the first step, d/D adjacent feature maps are grouped together to perform the standard temporal max pooling. As a result the fixed output size of ToI pooling layer is $D \times H \times W$. A graphical illustration of ToI pooling is presented in Figure 2.



Figure 2: Tube of interest pooling.

Back-propagation of ToI pooling layer routes the derivatives from output back to the input. Assume x_i is the *i*-th activation to the ToI pooling layer, and y_j is the *j*-th output. Then the partial derivative of the loss function (*L*) with respect to each input variable x_i can be expressed as:

$$\frac{\partial L}{\partial x_i} = \frac{\partial y_j}{\partial x_i} \frac{\partial L}{\partial y_j} \to \frac{\partial L}{\partial x_i} = \sum_j [i = f(j)] \frac{\partial L}{\partial y_j}.$$
 (1)

Each pooling output y_j has a corresponding input position i. We use a function $f(\cdot)$ to represent the argmax selection from the ToI. Thus, the gradient from the next layer $\partial L/\partial y_j$ is passed back to only that neuron which achieved the max $\partial L/\partial x_i$. Since one input may correspond to multiple outputs, the partial derivatives are the accumulation of multiple sources.

4. T-CNN Pipeline

As shown in Figure 1, our T-CNN is an end-to-end deep learning framework that takes video clips as input. The core component is the Tube Proposal Network (TPN) (see Figure 3) to produce tube proposals for each clip. Linked tube proposal sequence represents spatio-temporal action detection in the video and is also used for action recognition.

4.1. Tube Proposal Network

For a 8-frame video clip, 3D convolution and 3D pooling are used to extract spatio-temporal feature cube. In 3D ConvNet, convolution and pooling are performed spatiotemporally. Therefore, the temporal information of the input video is preserved. Our 3D ConvNet consists of seven 3D convolution layers and four 3D max-pooling layers. We denote the kernel shape of 3D convolution/pooling by $d \times h \times w$, where d, h, w are depth, height and width, respectively. In all convolution layers, the kernel sizes are $3 \times 3 \times 3$, padding and stride remain as 1. The numbers of filters are 64, 128 and 256 respectively in the first 3 convolution layers and 512 in the remaining convolution layers. The kernel size is set to $1 \times 2 \times 2$ for the first 3D max-pooling layer, and $2 \times 2 \times 2$ for the remaining 3D max-pooling layers. The details of network architecture are presented in Table 1. We use the C3D model [27] as the pre-trained model and fine tune it on each dataset in our experiments.

After conv5, the temporal size is reduced to 1 frame (*i.e.* feature cube with depth D = 1). In the feature tube, each frame/slice consists of a number of channels specified in Table 1. Here, we drop the number of channels for ease of explanation. Following faster R-CNN, we generate bounding box proposals based on the conv5 feature cube.

Anchor bounding boxes selection. In faster R-CNN, the bounding box dimensions are hand picked, *i.e.* 9 anchor boxes with 3 scales and 3 aspect ratios. We can directly adopt the same anchor boxes in our T-CNN framework. However, it has been shown in [13] that if we choose better priors as initializations for the network, it will help the network learn better for predicting good detections. Therefore, instead of choosing hand-picked anchor boxes, we apply

k-means clustering on the training set bounding boxes to learn 12 anchor boxes (*i.e.* clustering centroids). This data driven anchor box selection approach is adaptive to different datasets.

Each bounding box is associated with an "actionness" score, which measures the probability that the content in the box corresponds to a valid action. We assign a binary class label (of being an action or not) to each bounding box. Bounding boxes with actionness scores smaller than a threshold are discarded. In the training phase, the bounding box which has an IoU overlap higher than 0.7 with any ground-truth box or has the highest Intersection-over-Union (IoU) overlap with a ground-truth box (the later condition is considered in case the former condition may find no positive sample) is considered as a positive bounding box proposal.

Temporal skip pooling. Bounding box proposals generated from conv5 feature tube can be used for frame-level action detection by bounding box regression. However, due to temporal concentration (8 frames to 1 frame) of temporal max pooling, the temporal order of the original 8 frames is lost. Therefore, we use temporal skip pooling to inject the temporal order for frame-level detection. Specifically, we map each positive bounding box generated from conv5 feature tube to conv2 feature tube which has 8 feature frames/slices. Since these 8 feature slices correspond to the original 8 frames in the video clip, the temporal order information is preserved. As a result, if there are 5 bounding boxes in conv5 feature tube for example, 5 scaled bounding boxes are mapped in each conv2 feature slice at the corresponding locations. This creates 5 tube proposals as illustrated in Figure 3, which are paired with the corresponding 5 bounding box proposals for frame-level action detection. To form a fixed feature shape, ToI pooling is applied to the variable size tube proposals as well as the bounding box proposals. Since a tube proposal covers 8 frames, the ToI pooled bounding box is duplicated 8 times to form a tube. We then L2 normalize the paired two tubes and perform vectorization. For each frame, features are concatenated. Since we use the C3D model [27] as the pre-trained model, we connect a 1x1 convolution to match the input dimension of fc6. Three fully-connected layers process each descriptor and produce the output: displacement of height, width and center coordinate of each bounding box ("bbox") in each frame. Finally, a set of refined tube proposals are generated as an output from the TPN representing spatio-temporal action localization of the input video clip.

4.2. Linking Tube Proposals

We obtain a set of tube proposals for each video clip after the TPN. We then link these tube proposals to form a proposal sequence for spatio-temporal action localization of the entire video. Each tube proposal from different clips can be linked in a tube proposal sequence (*i.e.* video tube



Figure 3: Tube proposal network.

proposal) for action detection. However, not all combinations of tube proposals can correctly capture the complete action. For example, a tube proposal in one clip may contain the action and a tube proposal in the following clip may only capture the background. Intuitively, the content within the selected tube proposals should capture an action and connected tube proposals in any two consecutive clips should have a large temporal overlap. Therefore, two criteria are considered when linking tube proposals: actionness and overlap scores. Each video proposal is then assigned a score defined as follows:

$$S = \frac{1}{m} \sum_{i=1}^{m} Actionness_i + \frac{1}{m-1} \sum_{j=1}^{m-1} Overlap_{j,j+1}$$
(2)

where $Actionness_i$ denotes the actionness score of the tube proposal from the *i*-th clip, $Overlap_{j,j+1}$ measures the overlap between the linked two proposals respectively from the j-th and (j + 1)-th clips, and m is the total number of video clips. As shown in Figure 3, each bounding box proposal from conv5 feature tube is associated with an actionness score. The actionness scores are inherited by the corresponding tube proposals. The overlap between two tube proposals is calculated based on the IoU (Intersection Over Union) of the last frame of the *j*-th tube proposal and the first frame of the (j+1)-th tube proposal. The first term of S computes the average actionness score of all tube proposals in a video proposal and the second term computes the average overlap between the tube proposals in every two consecutive video clips. Therefore, we ensure the linked tube proposals can encapsulate the action and at the same time have temporal consistency. An example of linking tube proposals and computing scores is illustrated in Figure 4. We choose a number of linked proposal sequences with highest scores in a video (see more details in Sec. 5.1).



Figure 4: An example of linking tube proposals in each video clips using network flow. In this example, there are three video clips and each has two tube proposals, resulting in 8 video proposals. Each video proposal has a score, *e.g.* S1, S2, ..., S8, which is computed according to Eq. (2).

4.3. Action Detection

After linking tube proposals, we get a set of linked tube proposal sequences, which represent potential action instances. The next step is to classify these linked tube proposal sequences. The tube proposals in the linked sequences may have different sizes. In order to extract a fixed length feature vector from each of the linked proposal sequence, our proposed ToI pooling is utilized. Then the ToI pooling layer is followed by two fully-connected layers and a dropout layer. The dimension of the last fully-connected layer is N + 1 (N action classes and 1 background class).

5. Experiments

To verify the effectiveness of the proposed T-CNN for action detection, we evaluate T-CNN on three trimmed video datasets including UCF-Sports [21], J-HMDB [8], UCF-101 [11] and one un-trimmed video dataset – THU-MOS'14 [12].

5.1. Implementation Details

We implement our method based on the Caffe toolbox [10]. The TPN and recognition network share weights in their common layers. Due to memory limitation, in training phase, each video is divided into overlapping 8-frame clips with resolution 300×400 and temporal stride 1. When training the TPN network, each anchor box is assigned a binary label. Either the anchor box which has the highest IoU overlap with a ground-truth box, or an anchor box that has an IoU overlap higher than 0.7 with any ground-truth box is assigned a positive label, the rest are assigned negative label. In each iteration, 4 clips are fed into the network. Since the number of background boxes is much more than that of action boxes, to well model the action, we randomly select some of the negative boxes to balance the number of positive and negative samples in a batch. For recognition network training, we choose 40 linked proposal sequences with highest scores in a video as Tubes of Interest.

Our model is trained in an alternative manner. First, **Initialize TPN** based on the pre-trained model in [27], then using the generated proposals to **initialize recognition networks**. Next, the weights tuned by recognition network are used to **update TPN**. Finally, the tuned weights and proposals from TPN are used for **finalizing recognition network**. For all the networks for UCF-Sports and J-HMDB, the learning rate is initialized as 10^{-3} and decreased to 10^{-4} after 30k batches. Training terminates after 50k batches. For UCF-101 and THUMOS'14, the learning rate is initialized as 10^{-3} and decreased to 10^{-4} after 60k batches. Training terminates after 100k batches.

During testing, each video is divided into nonoverlapping 8-frame clips. If the number of frames in video cannot be divided by 8, we pad zeros after the last frame to make it dividable. 40 tube proposals with highest actionness confidence through TPN are chosen for the linking process. Non-maximum suppression (NMS) is applied to linked proposals to get the final action detection results.

5.2. Datasets and Experimental Results

UCF-Sports. This dataset contains 150 short videos of 10 different sport classes. Videos are trimmed and bounding boxes annotations are provided for all frames. We follow the training and test split defined in [16]. We use the IoU criterion and generate ROC curve in Figure 5(a) when overlap criterion equals to $\alpha = 0.2$. Figure 5(b) illustrates AUC (Area-Under-Curve) measured with different overlap criterion. In direct comparison, our T-CNN clearly outperforms all the competing methods shown in the plot. We are unable to directly compare the detection accuracy against Peng *et al.* [19] in the plot, since they do not provide the ROC and AUC curves. As shown in Table 2, the frame level mAP of our approach outperforms theirs in 8 actions out of 10. Moreover, by using the same metric, the video mAP of our

approach reaches 95.2 ($\alpha = 0.2$ and 0.5), while they report 94.8 ($\alpha = 0.2$) and 94.7 ($\alpha = 0.5$).

J-HMDB. This dataset consists of 928 videos with 21 different actions. All the video clips are well trimmed. There are three train-test splits and the evaluation is done on the average results over the three splits. The experiment results comparison is shown in Table 3. We report our results with 3 metrics: frame-mAP, the average precision of detection at frame level as in [5]; video-mAP, the average precision at video level as in [5] with IoU threshold $\alpha = 0.2$ and $\alpha = 0.5$. It is evident that our T-CNN consistently outperforms the state-of-the-art approaches in terms of all three evaluation metrics.

UCF101. This dataset has 101 actions. For action detection task, a subset of 24 action classes and 3, 207 videos have spatio-temporal annotations. Similar to other methods, we perform the experiments on the first train/test split only. We report our results in Table 4 with 3 metrics: frame-mAP, video-mAP ($\alpha = 0.2$) and video-mAP ($\alpha = 0.5$). Our approach again yields the best performance. Moreover, we also report the action recognition results of T-CNN on the above three datasets in Table 5.

THUMOS'14. To further validate the effectiveness of our proposed T-CNN approach for action detection, we evaluate it using the untrimmed videos from the THU-MOS'14 dataset [12]. The THUMOS'14 spatio-temporal localization task consists of 4 classes of sports actions: BaseballPitch, golfSwing, TennisSwing and ThrowDiscus. There are about 20 videos per class and each video contains 500 to 3, 000 frames. The videos are divided into validation set and test set, but only video in the test set have spatial annotations provided by [24]. Therefore, we use samples corresponding to those 4 actions in UCF-101 with spatial annotations to train our model.

In untrimmed videos, there often exist other unrelated actions besides the action of interests. For example, "walking" and "picking up a golf ball" are considered as unrelated actions when detecting "GolfSwing" in video. We denote clips which have positive ground truth annotation as positive clips, and the other clips as negative clips (*i.e.* clips contain only unrelated actions). If we randomly select negative samples for training, the number of boxes on unrelated actions is much smaller than that of background boxes (*i.e.* boxes capturing only image background). Thus the trained model will have no capability to distinguish action of interest and unrelated actions.

To this end, we introduce a so called **negative sample mining** process. Specifically, when initializing the TPN, we only use positive clips. Then we apply the model on the whole training video (both positive clips and negative clips). Most false positives in negative clips should include unrelated actions to help our model learn the correlation between action of interest and unrelated actions. Therefore we



Figure 5: The ROC and AUC curves for UCF-Sports Dataset [21] are shown in (a) and (b), respectively. The results are shown for Jain *et al.* [6] (green), Tian *et al.* [26] (purple), Soomro *et al.* [23] (blue), Wang *et al.* [28] (yellow), Gkioxari *et al.* [5] (cyan) and Proposed Method (red). (c) shows the mean ROC curves for four actions of THUMOS'14. The results are shown for Sultani *et al.* [24] (green), proposed method (red) and proposed method without negative mining (blue).

	Diving	Golf	Kicking	Lifting	Riding	Run	SkateB.	Swing	SwingB.	Walk	mAP
Gkioxari et al. [5]	75.8	69.3	54.6	99.1	89.6	54.9	29.8	88.7	74.5	44.7	68.1
Weinzaepfel et al. [30]	60.71	77.55	65.26	100.00	99.53	52.60	47.14	88.88	62.86	64.44	71.9
Peng et al. [19]	96.12	80.47	73.78	99.17	97.56	82.37	57.43	83.64	98.54	75.99	84.51
Ours	84.38	90.79	86.48	99.77	100.00	83.65	68.72	65.75	99.62	87.79	86.7

Table 2: mAP for each class of UCF-Sports. The IoU threshold α for frame m-AP is fixed to 0.5.

	fmAP	vmAP	vmAP
	$(\alpha = 0.5)$	$(\alpha = 0.2)$	$(\alpha = 0.5)$
Gkioxari et al. [5]	36.2	-	53.3
Weinzaepfel et al. [30]	45.8	63.1	60.7
Peng et al. [19]	58.5	74.3	73.1
Ours w/o skip pooling	47.9	66.9	58.6
Ours	61.3	78.4	76.9

Table 3: Comparison to the state-of-the-art on J-HMDB. The IoU threshold α for frame m-AP is fixed to 0.5.

	fmAP	video-mAP			
IoU th.		0.05	0.1	0.2	0.3
Weinzaepfel et al. [30]	35.84	54.3	51.7	46.8	37.8
Peng et al. [19]	39.63	54.5	50.4	42.3	32.7
Ours	41.37	54.7	51.3	47.1	39.2

Table 4: Comparison to the state-of-the-art on UCF-101 (24 actions). The IoU threshold α for frame m-AP is fixed to 0.5.

select boxes in negative clips with highest scores as **hard negatives** because low scores probably infer image background. In updating TPN procedure, we choose 32 boxes which have IoU with any ground truth greater than 0.7 as positive samples and randomly pick another 16 samples as negative. We also select 16 samples from hard negative pool as negative. Therefore, we efficiently train a model, which is able to distinguish not only action of interest from background, but also action of interest from unrelated actions.

The mean ROC curves of different methods on THU-MOS'14 action detection are plotted in Figure 5(c). Our method without negative mining performs better than the baseline method Sultani *et al.* [24]. Additionally, with negative mining, the performance is further boosted.

For qualitative results, we shows examples of detected action tubes in videos from UCF-Sports, JHMDB, UCF-101 (24 actions) and THUMOS'14 datasets (see Figure 6). Each block corresponds to a different video that is selected from the test set. We show the highest scoring action tube for each video.

6. Discussion

Tol Pooling. To evaluate the effectiveness of ToI pooling, we compare action recognition performance on UCF-101 dataset (101 actions) using C3D [27] and our approach. For the C3D network, we use C3D pre-train model from [27] and fine tune the weights on UCF-101 dataset. In the C3D fine tuning process, a video is divided into 16 frames clips first. Then the C3D network is fed by clips and outputs a feature vector for each clip. Finally, a SVM classifier is trained to predict the labels of all clips, and the video label is determined by the predictions of all clips belonging to the video. Compared to the original C3D network, we use the ToI pooling layer to replace the 5-th 3d-max-pooling layer in C3D pipeline. Similar to C3D network, our approach takes clips from a video as input. The ToI pooling layer takes the whole clip as tube of interest and the pooled depth



Figure 6: Action detection results by T-CNN on UCF-Sports, JHMDB, UCF-101 and THUMOS'14. Red boxes indicate the detections in the corresponding frames, and green boxes denote ground truth. The predicted label is overlaid.

Dataset	Accuracy (%)
UCF-Sports	95.7
J-HMDB	67.2
UCF-101 (24 actions)	94.4

Table 5: Action recognition results of T-CNN on various datasets.

is set to 1. As a result, each video will output one feature vector. Therefore, it is an end-to-end deep learning based video recognition approach. Video level accuracy is used as the metric. The results are shown in Table 6. For a direct comparison, we only use the result from deep network without fusion with other features. Our approach shows a 5.2% accuracy improvement compared to the original C3D. Our ToI pooling based pipeline optimizes the weights for the whole video directly, while C3D performs clip-based optimization. Therefore, our approach can better capture the spatio-temporal information of the entire video. Furthermore, our ToI pooling can be combined with other deep learning based pipelines, such as two-stream CNN [22].

	C3D [27]	Ours
Accuracy (%)	82.3	87.5

Table 6: Video action recognition results on UCF-101.

Temporal skip connection. Since we use overlapping clips with temporal stride of 1 in training, a particular frame is included in multiple training clips at different temporal positions. The actual temporal information of that particular frame is lost if we only use the conv5 feature cube to infer action bounding boxes. Especially when the action happens periodically (*i.e.* SwingBench), it always fails to locate a phase of spinning. On the contrary, by combining conv5

with conv2 through temporal skip pooling, temporal order is preserved to localize actions more accurately. To verify the effectiveness of temporal skip pooling in our proposed TPN, we conduct an experiment using our method without skip connection. In other words, we perform bounding box regression to estimate bounding boxes in 8 frames simultaneously using only the conv5 feature cube. As shown in Table 3, without skip connection, the performance decreases a lot, demonstrating the advantage of skip connection for extracting temporal order information and detailed motion in original frames.

Computational cost. We conduct experiments on a workstation with one GPU (Nvidia GTX Titan X). For a 40-frames video, it takes 1.1 seconds to generate tube proposals, 0.03 seconds to link tube proposals and 0.9 seconds to predict the action label.

7. Conclusion

In this paper we propose an end-to-end Tube Convolutional Neural Network (T-CNN) for action detection in videos. It exploits 3D convolutional network to extract effective spatio-temporal features and perform action localization and recognition in a unified framework. Coarse proposal boxes are densely sampled based on the 3D convolutional feature cube and linked for action recognition and localization. Extensive experiments on several benchmark datasets demonstrate the strength of T-CNN for spatiotemporal localizing actions, even in untrimmed videos.

Acknowledgement. The project was supported by Award No. 2015-R2-CX-K025, awarded by the National Institute of Justice, Office of Justice Programs, U.S. Department of Justice. The opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect those of the Department of Justice.

References

- P. Felzenszwalb, D. McAllester, and D. Ramanan. A discriminatively trained, multiscale, deformable part model. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2008.
- [2] A. Gaidon, Z. Harchaoui, and C. Schmid. Temporal localization of actions with actoms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(11):2782–2795, 2013.
- [3] R. Girshick. Fast r-cnn. In IEEE International Conference on Computer Vision (ICCV), December 2015.
- [4] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014.
- [5] G. Gkioxari and J. Malik. Finding action tubes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 759–768, 2015.
- [6] M. Jain, J. Van Gemert, H. Jégou, P. Bouthemy, and C. G. Snoek. Action localization with tubelets from motion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 740–747, 2014.
- [7] M. Jain, J. C. van Gemert, and C. G. Snoek. What do 15,000 object categories tell us about classifying and localizing actions? In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 46–55, 2015.
- [8] H. Jhuang, J. Gall, S. Zuffi, C. Schmid, and M. J. Black. Towards understanding action recognition. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3192– 3199, 2013.
- [9] S. Ji, W. Xu, M. Yang, and K. Yu. 3d convolutional neural networks for human action recognition. *IEEE Transactions* on Pattern Analysis and Machine Intelligence, 35(1):221– 231, 2013.
- [10] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [11] Y.-G. Jiang, J. Liu, A. Roshan Zamir, I. Laptev, M. Piccardi, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. /ICCV13-Action-Workshop/, 2013.
- [12] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. http://crcv.ucf.edu/THUMOS14/, 2014.
- [13] R. Joseph and F. Ali. Yolo9000: Better, faster, stronger. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.
- [14] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1725–1732, 2014.
- [15] Y. Ke, R. Sukthankar, and M. Hebert. Event detection in crowded videos. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1–8, 2007.

- [16] T. Lan, Y. Wang, and G. Mori. Discriminative figure-centric models for joint action localization and recognition. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2003–2010, 2011.
- [17] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [18] F. Negin and F. Bremond. Human action recognition in videos: A survey. *INRIA Technical Report*, 2016.
- [19] X. Peng and C. Schmid. Multi-region two-stream r-cnn for action detection. In *European Conference on Computer Vi*sion (ECCV), pages 744–759, 2016.
- [20] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 91–99. 2015.
- [21] M. Rodriguez, A. Javed, and M. Shah. Action mach: a spatio-temporal maximum average correlation height filter for action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2008.
- [22] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In Advances in Neural Information Processing Systems (NIPS), pages 568– 576, 2014.
- [23] K. Soomro, H. Idrees, and M. Shah. Action localization in videos through context walk. In *IEEE International Confer*ence on Computer Vision (CVPR), pages 3280–3288, 2015.
- [24] W. Sultani and M. Shah. What if we do not have multiple videos of the same action? – video action localization using web images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [25] L. Sun, K. Jia, D.-Y. Yeung, and B. E. Shi. Human action recognition using factorized spatio-temporal convolutional networks. In *IEEE International Conference on Computer Vision (ICCV)*, pages 4597–4605, 2015.
- [26] Y. Tian, R. Sukthankar, and M. Shah. Spatiotemporal deformable part models for action detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2642–2649, 2013.
- [27] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *IEEE International Conference on Computer Vi*sion (ICCV), pages 4489–4497, 2015.
- [28] L. Wang, Y. Qiao, and X. Tang. Video action detection with relational dynamic-poselets. In *European Conference* on Computer Vision (ECCV), pages 565–580, 2014.
- [29] L. Wang, Y. Qiao, X. Tang, and L. V. Gool. Actionness estimation using hybrid fully convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2708–2717, June 2016.
- [30] P. Weinzaepfel, Z. Harchaoui, and C. Schmid. Learning to track for spatio-temporal action localization. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3164–3172, 2015.
- [31] J. Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 4694–4702, 2015.

[32] B. Zhang, Z. Li, A. Perina, A. D. Bue, V. Murino, and J. Liu. Adaptive local movement modeling for robust object tracking. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(7):1515–1526, 2017.