

S³D: Single Shot multi-Span Detector via Fully 3D Convolutional Network

Da Zhang¹
dazhang@cs.ucsb.edu
Xiyang Dai²
xdai@umiacs.umd.edu
Xin Wang¹
xwang@cs.ucsb.edu
Yuan-Fang Wang¹
yfwang@cs.ucsb.edu

¹ Computer Vision Lab
Dept. of Computer Science
UC Santa Barbara, USA
² Dept. of Computer Science
University of Maryland, USA

Abstract

In this paper, we present a novel Single Shot multi-Span Detector for temporal activity detection in long, untrimmed videos using a simple end-to-end fully three-dimensional convolutional (Conv3D) network. Our architecture, named S³D, encodes the entire video stream and discretizes the output space of temporal activity spans into a set of default spans over different temporal locations and scales. At prediction time, S³D predicts scores for the presence of activity categories in each default span and produces temporal adjustments relative to the span location to predict the precise activity duration. Unlike many state-of-the-art systems that require a separate proposal and classification stage, our S³D is intrinsically simple and dedicatedly designed for single-shot, end-to-end temporal activity detection. When evaluating on THUMOS'14 detection benchmark, S³D achieves state-of-the-art performance and is very efficient and can operate at 1271 FPS.

1 Introduction

Advances in deep Convolutional Neural Network (CNN) have led to significant progress in video analysis over the past few years. While the performance of activity recognition has improved a lot [10, 29, 32, 34, 35, 37], the detection performance still remains unsatisfactory [27, 36, 40]. Comparing to activity recognition, which only aims at classifying the categories of manually trimmed video clips, activity detection is for detecting and recognizing activity instances from long, untrimmed video streams. It is substantially more challenging, as it is expected to handle activities with variable lengths, predicting both the activity category and the precise temporal boundaries of each instance.

A typical framework used by many state-of-the-art systems [24, 27, 28, 36] is *detection by classification*, where temporal proposals are generated by sliding windows [24, 27] or advanced proposal methods [3, 33] and separate activity classifier is applied to predict the final detection results. However, there may be certain limitations to these frameworks: (1) Temporal proposal and classification are independent processes and optimized separately

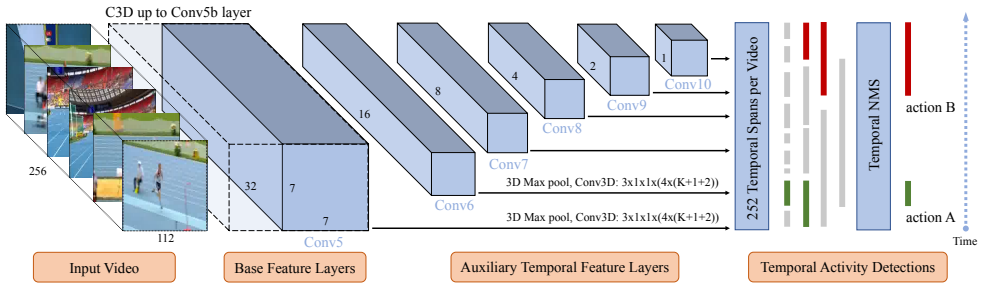


Figure 1: S^3D network architecture: Our network takes a video of 256 frames with spatial size 112×112 as input and computes base features using a standard C3D [52] network up to conv5b. We add auxiliary Conv3D layers on top of conv5 to produce a temporal feature hierarchy with multi-scale default spans at each layer. For each temporal feature map cell, we predict K class confidence scores, 1 activity confidence score and 2 location offsets with a set of Conv3D filters. Temporal NMS is applied to produce the final detection results. Refer to Figure 2 for a detailed illustration of default spans.

with different networks, resulting in sub-optimal performance, (2) the classification network only takes the proposal frames as input, thus forbidding it to see a larger temporal context which can be beneficial, and (3) this two-stage approach is usually slow due to inefficient proposal method and duplicate operations repeated in the proposal and classification stages.

We propose a *Single Shot multi-Span Detector* (S^3D), a simple yet novel fully Conv3D-based framework for activity detection in continuous untrimmed video streams. As illustrated in Figure 1, S^3D produces a fixed-size collection of temporal spans and scores for the presence of activity class instances in those spans, followed by a temporal non-maximum suppression step to generate the final detection results. S^3D is a highly-unified network by eliminating explicit temporal proposal and classification stages and solving the detection problem in one single shot. We set multi-scale default spans at feature maps with different temporal resolutions to naturally handle activities of different lengths. Furthermore, we predict the temporal offsets to adjust each default span in order to predict precise temporal boundaries. The network takes as input a whole video stream, allowing our scheme to see a larger temporal context and produce better detection results. The whole network is end-to-end trainable with a joint loss to directly maximize the detection performance.

The contributions of our paper are: (1) We introduce S^3D , a single shot end-to-end activity detection model based completely on Conv3D networks that can effectively predict both the precise temporal boundaries and confidence scores of multiple activity categories in untrimmed videos. (2) We demonstrate experimentally that our S^3D achieves state-of-the-art performance on temporal activity detection task on THUMOS’14 benchmark. (3) Besides its strong performance, the simple S^3D network is also very efficient and can run at 1271 FPS on a single GPU.

2 Related Work

Here, we review relevant works in activity recognition, object detection, and temporal activity detection. Other works on spatial-temporal activity detection and temporal video seg-

mentation are beyond the scope of this paper.

Activity Recognition. Activity recognition is an important research topic for video analysis and has been extensively studied in the past few years. Earlier methods were often based on hand-crafted visual features. Approaches include improved Dense Trajectory (iDT) [52, 53], feature encoding with Fisher Vector (FV) [23, 25], VLAD [17], etc. With the vast successes of deep learning methods [6, 9, 11, 12, 15, 30, 42], recent works, such as two-stream networks [10, 29, 57], 3D CNN architecture (C3D) [52] and I-3D [9], adopted deep neural networks and significantly improved the performance. However, most methods assume well-trimmed videos, where the activity of interests lasts for the entire clip duration. Although they do not consider the difficult task of localizing activity instances, these methods are widely used as the base network for the detection task.

Object Detection. Activity detection in untrimmed videos is closely related to object detection [13, 20, 26] in spatial images, where detection is performed by classifying region proposals into foreground classes or a background class. Earlier work [13] relied on an external region proposal method and trained a CNN classifier to classify each proposed region. Faster-RCNN [26] incorporated a region proposal network and RoI pooling to jointly generate and classify region proposals with a single network, resulting in a large improvement of the accuracy and efficiency. SSD [20] completely eliminated proposal generation and subsequent feature re-sampling stages and encapsulated all computation in a single network to directly output object locations and confidence scores. Our network is inspired by SSD [20] and adopt similar design philosophies into temporal activity detection. Like SSD [20], our S³D model is also designed for both accuracy and efficiency.

Temporal Activity Detection. Unlike activity recognition, the detection task focuses on learning how to detect activity instances in untrimmed videos with annotated temporal boundaries and instance category. The problem has recently received significant research attention due to its potential application in video data analysis.

Previous works on activity detection mainly used sliding windows as candidates and classified them with activity classifiers trained on multiple features [10, 16, 21, 23, 31]. Recently, some approaches bypassed the need for exhaustive sliding window search by proposing better temporal proposal schemes [2, 3, 9, 22, 33, 41]. Along this line of attack, some recent works incorporated deep networks into the detection framework and obtained improved performance [8, 27, 28, 39, 43]: S-CNN [27] proposed a multi-stage CNN which adopted 3D ConvNet with multi-scale sliding window to boost accuracy; CDC [28] used temporal deconvolutional network to generate per-frame classification scores for refining temporal boundaries; R-C3D [39] proposed an end-to-end trainable activity detector based on Faster-RCNN [26]; Dai *et al.* [8] explicitly modeled temporal contextual information into the proposal stage; SSN [43] utilized temporal pyramid pooling to model the complicated temporal structures and achieved state-of-the-art performance. However, all these methods require a separate temporal proposal and activity classification method.

Most recently, several attempts were made towards single shot temporal activity detection: SSAD [19] proposed to directly predict activity instances in untrimmed videos with a separate feature extraction and detection network. SS-TAD [10] have investigated the use of gated recurrent memory module in a single-stream detection framework. Our approach is one of the first within this group to propose a highly-integrated detection architecture. With a simple end-to-end Conv3D network which learns directly from raw video frames to final detection outputs, S³D is able to jointly optimize feature representation and prediction layers, resulting in a simple, fast and robust architecture achieving both state-of-the-art performance and fast runtime speed.

3 Approach

We introduce a *Single Shot multi-Span Detector* (S^3D), a simple yet novel fully Conv3D-based framework for activity detection in long untrimmed video streams. The S^3D approach, illustrated in Figure 1, is based on a feed-forward fully Conv3D network that produces a fixed-size collection of temporal spans and scores for the presence of activity class instances in those spans, followed by a temporal NMS step to generate the final detection results.

3.1 Model

Our model consists of four major components: base feature layers, auxiliary temporal feature layers, multi-scale default spans and convolutional predictors. The *base feature layers* are used to extract high-level features given an input video stream. We then add *auxiliary temporal feature layers* to generate rich spatial-temporal feature hierarchies. These layers decrease in temporal dimension progressively and allow predictions of temporal spans at different locations and scales. We associate *multi-scale default spans* with each feature map cell and the default spans tile the feature map in a convolutional manner. At each feature map cell, we predict the temporal offsets relative to the default span in the cell, as well as the confidence scores that indicate the presence of an activity instance in each of those spans. These are done by adding *convolutional predictors* on top of each cell.

Base Feature Layers. We use Conv3D filters to extract rich feature hierarchies from a given input video stream. Specifically, the input to our model is a sequence of RGB video frames which can be represented as a tensor with dimension $\mathbb{R}^{L \times H \times W \times 3}$, where L is the number of frames, H and W are the height and width of each frame. We apply the standard C3D architecture [62] as it has been proven as an effective building block in prior works [4, 19, 69]. We adopt the Conv3D layers (conv1a to conv5b) of C3D and generate a feature map $C_{conv5} \in \mathbb{R}^{\frac{L}{8} \times \frac{H}{16} \times \frac{W}{16} \times 512}$. We use C_{conv5} as our base feature since it is a rich yet compact spatial-temporal representation of the input video stream.

Auxiliary Temporal Feature Layers. To allow the model to predict variable scale temporal spans, we add temporal feature layers to the end of the base feature layers. Similar to [62], we first down sample C_{conv5} by a factor of 2 in both spatial and temporal dimension via 3D max pooling and then add auxiliary Conv3D layers to produce a sequence of feature maps that progressively decrease in temporal dimension while keeping the same spatial resolution. In more detail, we stack Conv3D layers with temporal kernel size 3 to extend the temporal receptive field and the stride is set to 2 for progressively decreasing the temporal dimension. We also add bottleneck Conv3D layers to help prevent over-fitting and improve runtime efficiency. The detailed network configurations are illustrated in Figure 1 when $L = 256$ and $H = W = 112$.

The network is intrinsically simple by only applying Conv3D filters, but builds a rich feature hierarchy by summarizing a continuous video stream in multiple temporal resolutions, allowing us to add default temporal spans at certain layers to get temporal predictions at multiple scales.

Multi-scale Default Spans. To handle different activity locations and scales, [27] suggests processing the video at different segment levels and combining the results afterward, while [4] uses a gated recurrent network to assign a number of anchors at different time steps. However, by utilizing feature maps from several different layers in a single network for prediction we can mimic the same effect, while also sharing parameters across all temporal scales. We use feature maps with different temporal resolutions for detection since earlier

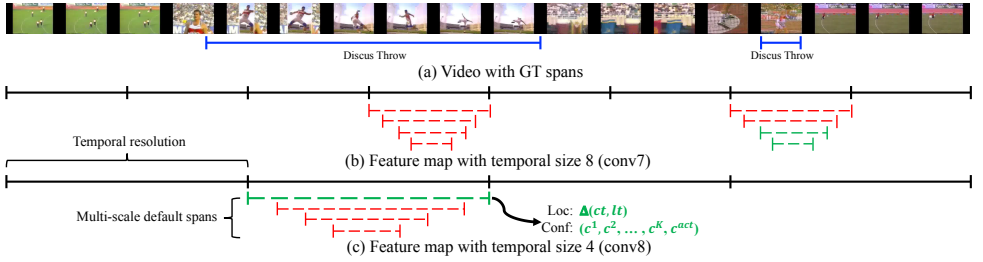


Figure 2: S³D framework. (a) Input video with temporal ground-truth annotations. We evaluate a small set (e.g. 4) of multi-scale default spans at each location in several feature maps with different temporal resolutions (e.g. conv7 in (b) and conv8 in (c)). For each default span, we predict both the temporal offsets and the confidences for presence of activity and all activity categories. At training time, we match the default spans to the ground truth spans.

feature maps have higher resolution and capture finer details of the input video, and deeper feature maps have larger receptive fields and contain more temporal contexts.

In our design, we use conv5 to conv10 as our temporal feature maps and associate a set of multi-scale default spans with each temporal feature map cell. We design the tiling of default spans so that specific feature maps learn to be responsive to particular locations and lengths of the activities. Regrading a temporal feature map f with temporal length L_f , the scale of the default spans for this feature map is set as $S_f = \frac{1}{L_f}$ (as the input video length is normalized to 1). We impose different scale ratios for the default spans, and denote them as $r \in \{0.25, 0.5, 0.75, 1.0\}$. We can compute the length ($L'_f = S_f \cdot r$) for each default span, and we set the center of each default span to $\frac{i+0.5}{L_f}$, where i indicates the i -th temporal feature cell, $i \in [0, L_f)$. So for an temporal feature map with length L_f and R different scale ratios ($R = 4$), the number of default spans is $L_f \cdot R$.

By combining predictions for all default spans with different scales from all locations of multi-scale feature maps, we have a diverse set of predictions, covering various activity locations and lengths. A concrete example is illustrated in Figure 2 where $L_{conv7} = 8$ and $L_{conv8} = 4$ for feature map conv7 and conv8 respectively.

Convolutional Predictors. Each temporal feature layer can produce a fixed set of detection predictions using a set of Conv3D filters. These are indicated on top of the feature network architecture in Figure 1. For a temporal feature map $C_f \in \mathbb{R}^{L_f \times H_f \times W_f \times d_f}$, the basic operation for predicting parameters of a potential temporal detection is a $3 \times H_f \times W_f$ kernel that produces scores for activity presence and categories, or temporal offsets relative to the default location and scale. Specifically, for each default span at a given temporal location, we compute K positive class confidence scores plus one activity confidence score and two temporal offsets. This results in a total of $(K + 1 + 2) \times R$ filters that are applied around each location in the feature map, yielding $(K + 1 + 2) \times R \times L_f$ outputs for a temporal feature map C_f . For an illustration of default spans, please refer to Figure 2. Each default span gets a prediction score vector $v_{pred} = (c^1, c^2, \dots, c^K, c^{act}, \Delta ct, \Delta lt)$ with length $K + 1 + 2$, where c^{act} is a class-agnostic confidence score to estimate the presence of activity, c^1 to c^K are used to predict default span's category and $\Delta ct, \Delta lt$ are temporal offsets relative to the locations of default spans.

3.2 Training

The key step of training S³D is that the ground truth information needs to be assigned to specific outputs in the fixed set of detector outputs. Once this assignment is determined, the loss function and back propagation are applied. We also discuss training data construction and hard negative mining strategies used in our model.

Training Data Construction and Augmentation. In theory, because S³D is a fully Conv3D network, it can be applied to an input of arbitrary size. Therefore, our S³D network can operate on videos of variable lengths. In practice, due to GPU memory limitations, we slide a temporal window of size L frames on the video and feed each windowed segment individually into the S³D network to obtain temporal detections. Although the input window size is fixed, we decode the input video stream with a small frame rate, allowing the network to encode enough temporal contexts for precisely detecting activity instances. Therefore, given a set of training videos, we obtain a training collection of windows with temporal activity annotations inside each windowed video segment. To make the model more robust to various activity locations and scales, we further improve the training dynamics by augmenting the training videos with temporal and spatial jittering [10].

Matching Strategy. During training, we need to determine which default spans correspond to a ground truth detection and train the network accordingly. Specifically, for each default span, we compute the Intersection-over-Union (IoU) score with all ground truth instances. If the highest IoU score is higher than 0.5, we match the default span with the corresponding ground truth span and regard it as positive, otherwise negative. So a ground truth instance can match multiple default spans while a default span can only match one ground truth instance at most. This simplifies the learning problem, allowing the network to predict high scores for multiple overlapping default spans.

Hard Negative Mining. After the matching step, most of the default spans are negatives. This introduces a significant imbalance between the positive and negative training examples. Instead of using all the negative examples, we sort them using the highest activity confidence loss for each default span and pick the top ones so that the ratio between the negatives and positives is nearly 1 : 1. We found that this leads to faster optimization and a more stable training.

Training Objective. The training objective of S³D is to solve a multi-task optimization problem. Let $x_{ij}^k = \{1, 0\}$ be an indicator for matching the i -th default span to the j -th ground truth span of category $k \in [1, K]$, and s_i be the highest IoU score with any ground truth spans. The overall objective loss function is a weighted sum of the localization loss (loc), class confidence loss (conf) and activity confidence loss (act):

$$Loss = L_{loc}(x, t, g) + \alpha L_{conf}(x, c) + \beta L_{act}(s, c) \quad (1)$$

where α and β are the weight terms used for balancing each part of the loss function.

The localization loss is a Smooth L1 loss [10] between the predicted temporal offsets (t) and the ground truth span parameters (g). In temporal domain, we regress to offsets for the center (ct) of the default span (d) and for its length (lt):

$$L_{loc}(x, t, g) = \frac{1}{N_{pos}} \sum_i^{N_{pos}} \sum_{m \in \{ct, lt\}} x_{ij}^k smooth_{L1}(t_i^m - \hat{g}_j^m) \quad (2)$$

where N_{pos} is the number of positive matching default spans in a batch, and the temporal offset parameters \hat{g}_j^m are defined similarly like the bounding box offset in object detection [10]:

$$\hat{g}_j^{ct} = \Delta ct_i = (g_j^{ct} - d_i^{ct})/d_i^{lt} \quad \hat{g}_j^{lt} = \Delta lt_i = \log\left(\frac{g_j^{lt}}{d_i^{lt}}\right) \quad (3)$$

where g_j^{ct} , d_i^{ct} are the centers and g_j^{lt} , d_i^{lt} are the lengths for the ground truth span and the matching default temporal span respectively.

The class confidence loss is a softmax loss over multiple class confidences (c):

$$L_{conf}(x, c) = -\frac{1}{N_{pos}} \sum_i^{N_{pos}} x_{ij}^k \log(\hat{c}_i^k) \quad (4)$$

where $\hat{c}_i^k = \frac{\exp(c_i^k)}{\sum_k \exp(c_i^k)}$ is the softmax probability for the ground truth class of this instance.

The class confidence loss is only used to distinguish between multiple positive classes not including the background. We use another activity confidence score to predict activity class agnostic scores.

The activity confidence loss is a binary classification loss using sigmoid cross-entropy. Rather than using a hard ground truth score for positive (1) and negative (0), we use the IoU score s_i as ground truth for each default span. This helps the training procedure since positive default spans are assigned different confidence levels based on its overlap with the ground truth span. We define the activity confidence loss as:

$$L_{act}(s, c) = -\frac{1}{N} \sum_i^N (s_i \log(c_i^{act}) + (1 - s_i) \log(1 - c_i^{act})) \quad (5)$$

where N is the number of total training default spans in a batch and $N = N_{pos} + N_{neg}$; c_i^{act} is the predicted activity confidence score. Note that we separate the activity confidence score and class confidence scores via two separate losses. Comparing to only having one softmax classification loss containing all positive classes and one background class, we find this configuration is more robust, leads to better validation performance and makes the network architecture more flexible.

3.3 Prediction

Activity prediction in S³D is single shot with one forward pass of the network. Given an input video stream, we generate all default spans with class confidence scores, activity confidence score and temporal location offsets. The temporal location offset is in the form of relative displacement of the center point and length of each instance as described in Equation 3, which is applied on the default span to predict accurate start time and end time. Then the default spans with low activity confidence score will be filtered out and the remaining spans are refined via NMS with threshold value 0.5. Each remaining span is considered as a positive prediction and assigned the activity label with the highest class confidence score, which we consider as the final temporal detection results of S³D.

4 Experiments

We evaluate the proposed framework on the THUMOS'14 [18] large-scale activity detection benchmark dataset. As shown in the experiments, our S³D not only achieves state-of-the-art performance but also acquires fast runtime speed at 1271 FPS.

IoU threshold	0.3	0.4	0.5	0.6	0.7
S-CNN [27]	36.3	28.7	19.0	10.3	5.3
CDC [28]	40.1	29.4	23.3	13.1	7.9
SSAD [19]	43.0	35.0	24.6	-	-
TCN [8]	-	33.3	25.6	15.9	9.0
R-C3D [69]	44.8	35.6	28.9	-	-
SSN [43]	50.6	40.8	29.1	-	-
SS-TAD [10]	40.1	-	29.2	-	9.6
S ³ D	47.9	41.2	32.6	23.3	14.3

Table 1: Temporal activity detection mAP on THUMOS’14. The top performing methods in existing papers are shown. S³D achieves state-of-the-art performance at different overlap threshold. (- indicates that results are unavailable in the corresponding papers).

4.1 Experimental Setup

Dataset [18]. The temporal activity detection task of THUMOS’14 dataset is challenging and widely used. Over 20 hours of video and 20 activity categories are involved and annotated temporally, resulting in 200 validation and 213 test untrimmed videos. Following the standard practice, we train our models on the validation set and evaluate them on the testing set. We follow the conventional metrics used in THUMOS’14, computing the Average Precision (AP) for each activity category and calculating mean AP (mAP) for evaluation.

Implementation Details. S³D takes as input $L = 256$ raw video frames with size $H = W = 112$. We decode each video at 8 FPS and produce a collection of training windows. Thus, each window contains 32 seconds of a video stream and this is motivated by the fact that more than 99% of activity instances in THUMOS’14 are less than 32 seconds. We use conv5 to conv10 as the temporal feature layers with temporal dimension $\{32, 16, 8, 4, 2, 1\}$ and associate a set of default spans at each temporal feature cell with four ratios $\{0.25, 0.5, 0.75, 1.0\}$, resulting in 252 default spans in total; the default spans correspond to spans of duration between 0.25s and 32s uniformly distributed at different temporal locations. We initialize base feature layers with C3D weights pre-trained on Sports-1M by the authors in [25], and other layers from scratch. We allow all the layers of S³D to be trained on THUMOS’14 with the end-to-end loss function. Our code is available at <https://github.com/dazhang-cv/S3D>.

4.2 Comparison with State-of-the-art

The comparison results between our S³D and other top-performing methods are summarized in Table 1, and our S³D outperforms all previous state-of-the-art methods. Furthermore, S³D improves the state-of-the-art by a large margin when the evaluation IoU thresholds are set at higher levels (0.5 to 0.7), indicating its superior ability to predict precise temporal boundaries of different activities.

In comparison with the proposed S³D model: previous systems on top of C3D networks (S-CNN [27], CDC [28]) largely relies on good temporal proposals generated by external proposal methods, restricting them from directly optimizing the detection performance. R-C3D [69] is able to process a long video stream and predict multi-scale activity instances, but it only applies anchors on a single feature map with fixed temporal dimension. With the proposed S³D framework, we jointly optimize the feature representation and detection layers at different temporal levels by processing an untrimmed input video stream with enough

temporal context.

4.3 Ablation Study

To understand S³D better, we evaluate our network with different variants on THUMOS'14 to study their effects. For all experiments, we only change the certain part of the network and use the same evaluation settings. We compare the result of different variants using the mAP at IoU threshold 0.5.

include 1.0 span	✓	✓	✓	✓
include 0.25 span		✓	✓	✓
include 0.5 span			✓	✓
include 0.75 span				✓
# Spans	63	126	189	252
mAP@0.5	27.5	29.5	31.1	32.6

Table 2: Effects of various design choices on S³D performance, the span with ratio 1.0 is included by default.

Default Span Ratio. By default, we use 4 default spans per each temporal location. If we remove the spans with ratio 0.75, the mAP drops by 1.5%. By further removing the spans with ratio 0.25 and 0.5, the mAP drops another 3.6%. By only keeping the span with ratio 1.0, our model already has a strong performance (mAP 27.5%) since it already covers most ground truth instances in the dataset. Using a variety of default ratios make the task of predicting spans easier for the network and result in better performance.

Span Regression. The default spans are defined at fixed temporal locations. In order to generate precise predictions for starting and ending time of each activity instance, we adjust each default span by applying a temporal offset described in Equation 3. This technique, which we call span regression, allows our model to predict temporal spans at much smaller granularities. As shown in Table 3, span regression improves the mAP from 28.6% to 32.6%.

Span regression	conv5	conv6	conv7	conv8	conv9	conv10	mAP@0.5	# Spans
✓	✓	✓	✓	✓	✓	✓	32.6	252
	✓	✓	✓	✓	✓	✓	28.6	252
✓	✓	✓	✓	✓	✓		31.8	248
✓	✓	✓	✓	✓			30.7	240
✓	✓	✓	✓				27.6	224

Table 3: Effects of using multiple temporal feature layers and span regression.

Multi-scale Default Spans. A major advantage of S³D is using default spans of different scales on different temporal feature layers. To measure the advantage gained, we progressively remove layers and compare results. Table 3 shows a decrease in accuracy with fewer layers, dropping monotonically from 32.6% to 27.6%. This is because that different layers are responsible for predicting temporal activities at different lengths, which reinforces the message that it is critical to spread spans of different scales over different layers.

Qualitative Results. We provide qualitative results to demonstrate the effectiveness and robustness of our proposed S³D network. As shown in Figure 3, different video streams contain very diversified background context and different activity instances vary a lot in temporal location and scale. S³D is able to predict the accurate temporal span as well as the

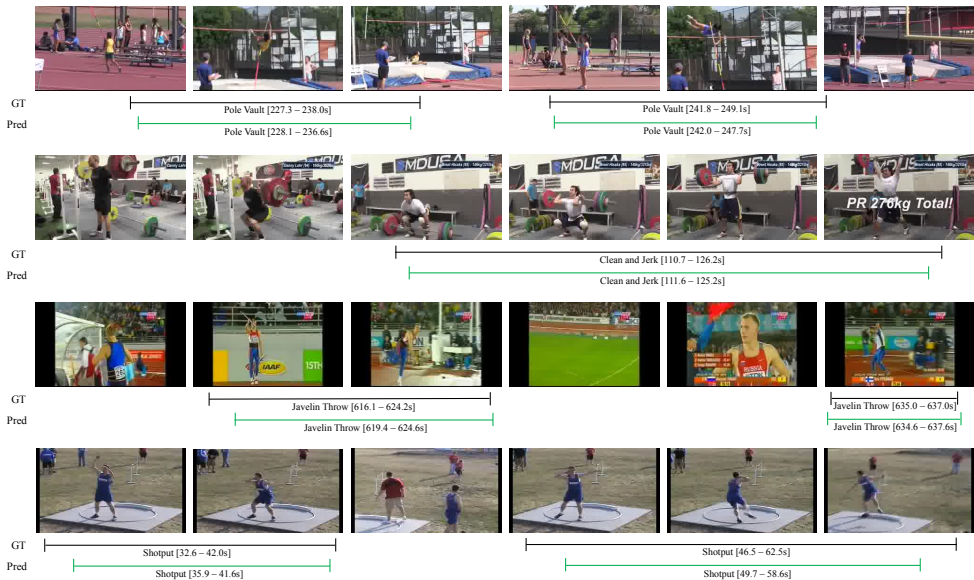


Figure 3: Qualitative visualization of the top detected activities by S³D (best viewed in color) on four different activity categories in THUMOS'14 dataset: *Pole Vault*, *Clean and Jerk*, *Javelin Throw* and *Shotput*. Ground truth activity segments are marked in black and predicted activity segments are marked in green.

correct activity category. Furthermore, S³D can distinguish activity with minor differences such as the normal weightlifting compared to *Clean and Jerk*. It is also capable of detecting the same activity sequence with different playing speed as shown in the *Shotput* example.

Activity Detection Speed. Since our model has a single-shot, end-to-end design with simple Conv3D building blocks, it is also very efficient. We benchmark our model on a GeForce GTX 1080 Ti GPU, and our S³D can run much faster than real time at 1271 FPS. For comparison, previous top performing methods [19, 27, 28] have significantly lower FPS for the whole detection pipeline. Comparing to some recent works [0, 39] providing good runtime efficiency, our S³D achieves much better accuracy.

5 Conclusion

In this paper, we introduce S³D, a Single Shot multi-Span Detector for temporal activity detection. We design a simple network architecture by using only a fully Conv3D network on top of the raw video frames to jointly predict the temporal boundaries as well as activity categories. A key feature of S³D is the use of multi-scale temporal span outputs attached to multiple temporal feature maps. With this framework, we achieved state-of-the-art performance on THUMOS'14 benchmark dataset, while being efficient to run much faster than real time on a single GPU.

References

- [1] S Buch, V Escorcia, B Ghanem, L Fei-Fei, and JC Niebles. End-to-end, single-stream temporal action detection in untrimmed videos. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2017.
- [2] Shyamal Buch, Victor Escorcia, Chuanqi Shen, Bernard Ghanem, and Juan Carlos Niebles. Sst: Single-stream temporal action proposals. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 6373–6382. IEEE, 2017.
- [3] Fabian Caba Heilbron, Juan Carlos Niebles, and Bernard Ghanem. Fast temporal activity proposals for efficient detection of human actions in untrimmed videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1914–1923, 2016.
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4733. IEEE, 2017.
- [5] Dongdong Chen, Jing Liao, Lu Yuan, Nenghai Yu, and Gang Hua. Coherent online video style transfer. In *Proc. Intl. Conf. Computer Vision (ICCV)*, 2017.
- [6] Dongdong Chen, Lu Yuan, Jing Liao, Nenghai Yu, and Gang Hua. Stylebank: An explicit representation for neural image style transfer. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, 2017.
- [7] Xiyang Dai, Joe Yue-Hei Ng, and Larry S Davis. Fason: First and second order information fusion network for texture recognition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 00, pages 6100–6108, July 2017.
- [8] Xiyang Dai, Bharat Singh, Guyue Zhang, Larry S Davis, and Yan Qiu Chen. Temporal context network for activity localization in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5793–5802, 2017.
- [9] Victor Escorcia, Fabian Caba Heilbron, Juan Carlos Niebles, and Bernard Ghanem. Daps: Deep action proposals for action understanding. In *European Conference on Computer Vision*, pages 768–784. Springer, 2016.
- [10] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1933–1941, 2016.
- [11] Adrien Gaidon, Zaid Harchaoui, and Cordelia Schmid. Temporal localization of actions with actoms. *IEEE transactions on pattern analysis and machine intelligence*, 35(11): 2782–2795, 2013.
- [12] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.
- [13] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.

- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [15] Mingming He, Dongdong Chen, Jing Liao, Pedro V Sander, and Lu Yuan. Deep exemplar-based colorization. *ACM Transactions on Graphics (Proc. of Siggraph 2018)*, 2018.
- [16] Mihir Jain, Jan Van Gemert, Hervé Jégou, Patrick Bouthemy, and Cees GM Snoek. Action localization with tubelets from motion. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 740–747, 2014.
- [17] Hervé Jégou, Matthijs Douze, Cordelia Schmid, and Patrick Pérez. Aggregating local descriptors into a compact image representation. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 3304–3311. IEEE, 2010.
- [18] YG 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, 2014.
- [19] Tianwei Lin, Xu Zhao, and Zheng Shou. Single shot temporal action detection. In *Proceedings of the 2017 ACM on Multimedia Conference*, pages 988–996. ACM, 2017.
- [20] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [21] Pascal Mettes, Jan C van Gemert, Spencer Cappallo, Thomas Mensink, and Cees GM Snoek. Bag-of-fragments: Selecting and encoding video fragments for event detection and recounting. In *Proceedings of the 5th ACM on International Conference on Multimedia Retrieval*, pages 427–434. ACM, 2015.
- [22] Pascal Mettes, Jan C van Gemert, and Cees GM Snoek. Spot on: Action localization from pointly-supervised proposals. In *European Conference on Computer Vision*, pages 437–453. Springer, 2016.
- [23] Dan Oneata, Jakob Verbeek, and Cordelia Schmid. Action and event recognition with fisher vectors on a compact feature set. In *Proceedings of the IEEE international conference on computer vision*, pages 1817–1824, 2013.
- [24] Dan Oneata, Jakob Verbeek, and Cordelia Schmid. The lear submission at thumos 2014. 2014.
- [25] Florent Perronnin, Jorge Sánchez, and Thomas Mensink. Improving the fisher kernel for large-scale image classification. *Computer Vision–ECCV 2010*, pages 143–156, 2010.
- [26] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [27] Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1049–1058, 2016.

- [28] Zheng Shou, Jonathan Chan, Alireza Zareian, Kazuyuki Miyazawa, and Shih-Fu Chang. Cdc: Convolutional-de-convolutional networks for precise temporal action localization in untrimmed videos. In *CVPR*, 2017.
- [29] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, pages 568–576, 2014.
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [31] Kevin Tang, Bangpeng Yao, Li Fei-Fei, and Daphne Koller. Combining the right features for complex event recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2696–2703, 2013.
- [32] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015.
- [33] Jan C van Gemert, Mihir Jain, Ella Gati, Cees GM Snoek, et al. Apt: Action localization proposals from dense trajectories. 2015.
- [34] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In *Proceedings of the IEEE international conference on computer vision*, pages 3551–3558, 2013.
- [35] Heng Wang, Alexander Kläser, Cordelia Schmid, and Cheng-Lin Liu. Action recognition by dense trajectories. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 3169–3176. IEEE, 2011.
- [36] Limin Wang, Yu Qiao, and Xiaoou Tang. Action recognition and detection by combining motion and appearance features. *THUMOS14 Action Recognition Challenge*, 1(2): 2, 2014.
- [37] Limin Wang, Yuanjun Xiong, Zhe Wang, and Yu Qiao. Towards good practices for very deep two-stream convnets. *arXiv preprint arXiv:1507.02159*, 2015.
- [38] Limin Wang, Yu Qiao, Xiaoou Tang, and Luc Van Gool. Actionness estimation using hybrid fully convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2708–2717, 2016.
- [39] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017.
- [40] Serena Yeung, Olga Russakovsky, Greg Mori, and Li Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2678–2687, 2016.
- [41] Gang Yu and Junsong Yuan. Fast action proposals for human action detection and search. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1302–1311, 2015.

- [42] Da Zhang, Hamid Maei, Xin Wang, and Yuan-Fang Wang. Deep reinforcement learning for visual object tracking in videos. *arXiv preprint arXiv:1701.08936*, 2017.
- [43] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaoou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *The IEEE International Conference on Computer Vision (ICCV)*, volume 8, 2017.