

An Online System for Real-Time Activity Detection in Untrimmed Surveillance Videos

Aayush Jung Rana

aayushjr@knights.ucf.edu

Praveen Tirupattur

praveentirupattur@knights.ucf.edu

Mamshad Nayeem Rizve

nayeemrizve@knights.ucf.edu

Kevin Duarte

kevin_duarte@knights.ucf.edu

Ugur Demir

ugur@knights.ucf.edu

Yogesh Rawat

yogesh@crcv.ucf.edu

Mubarak Shah

shah@crcv.ucf.edu

Abstract

Activity detection in surveillance videos is a challenging problem due to multiple factors such as large field of view, presence of multiple activities, varying scales and viewpoints, and its untrimmed nature. The requirement of processing the surveillance videos in real-time makes this more challenging. In this work, we propose a real-time online system to perform activity detection on untrimmed surveillance videos. The proposed system consists of three stages: first we detect tubelets with activities, then classify them, and finally merge them to generate spatio-temporal activity detections. We propose a localization network which takes a video clip as input and makes use of feature pyramid, multi-layer loss, and atrous convolutions to address the issue of multiple scales and detect small activities in terms of tubelets. The online processing of videos at a clip level drastically reduces the computation time in detecting activities. The detected tubelets are assigned activity class scores and merged together using our proposed Tubelet-Merge Action-Split (TMAS) algorithm to form action tubes. The TMAS algorithm efficiently connects the tubelets in an online fashion to generate spatio-temporal detections which are robust against varying length activities. We perform our experiments on the DIVA (Deep Intermodal Video Analytics) dataset and demonstrate the effectiveness of the proposed approach in terms of speed (~ 100 fps) and performance with state-of-the-art results. The code and models will be made publicly available.

1. Introduction

Deep convolutional neural networks have achieved impressive action classification results in recent years [27, 4, 28]. Similar advancements have been made for the tasks of

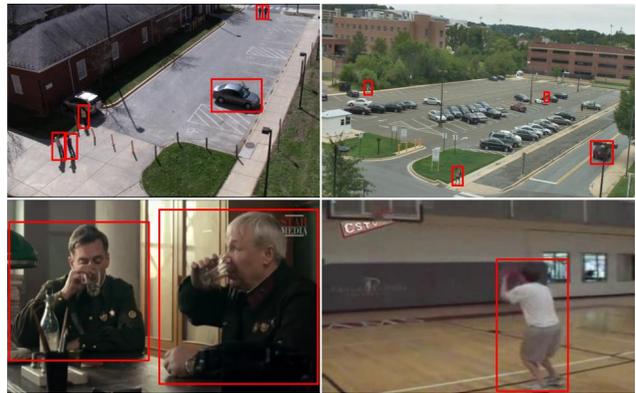


Figure 1. **Top:** Two Sample frames from different scenes of the DIVA dataset showing variation in perspective, scale and field-of-view. **Bottom:** Sample frames from the AVA dataset [10] (left) and from the THUMOS'14 dataset [13] (right). The DIVA dataset contains a greater number of concurrent actions as well as a greater variety of action scales (both spatially and temporally).

action detection in trimmed videos [14, 26, 7] and temporal action localization in untrimmed videos [32, 21]. However, these improvements have not been transferred to spatio-temporal action detection in untrimmed videos; current computer vision systems have yet to achieve high performance on this difficult task.

Action detection in untrimmed security videos poses multiple challenges. Surveillance videos comprise of multiple viewpoints and contain several actors performing multiple actions concurrently. These actors have varying scales and tend to be extremely small relative to the video frame, which makes detection of small activities extremely challenging. These challenges make it difficult to extend existing methods to detect actions in the untrimmed security videos found in the DIVA (Deep Intermodal Video Analyt-

ics) dataset [19]. Current methods are trained and evaluated on datasets which contain some, but not all of these challenges. For example, THUMOS'14 [13] is comprised of untrimmed videos, but each video contains only one or two actors performing the same action. The AVA dataset [10] contains multiple actors and actions, but each video is trimmed. Figure 1 shows sample frame from the DIVA dataset and compares them with frames from other action detection datasets.

In this work, we focus on untrimmed surveillance videos and propose an online real-time system for spatio-temporal action detection. Since activities in untrimmed videos can vary in length, it is necessary to handle both short, atomic activities, like opening a door or exiting a vehicle, as well as long, repetitive actions like walking or riding. To this end, our pipeline processes videos in an online fashion: it localizes and classifies short action tubelets. Our Tubelet-Merge Action-Split (TMAS) algorithm, merges these tubelets into action-agnostic tubes of varying length and splits the tubes into a set of final spatio-temporal action predictions. By classifying short tubelets and merging them into action tubes, our system is able to detect both atomic and repetitive actions.

Our action localization module proposes pixel-level action localizations for a short video clip. This allows our system to generate action tubelet proposals without the use of frame-based object detectors. Frame-based object detection systems [9] have two main issues: 1) processing each frame independently requires large amounts of computation, which reduces the overall speed of the system and leads to temporally inconsistent detections between adjacent frames, and 2) all objects within the frame are detected, even those which are not performing actions. Our action localization module processes multiple frames simultaneously and only produces tubelets which correspond to possible actions within the video. This results in temporally consistent localizations as well as a reduction in the number of proposals which drastically increases the speed of the overall system. Due to our localization network and the overall system's online processing of videos, inference is performed at ~ 100 frames per second, greatly exceeding the speed of contemporary action detection systems.

Our contributions are as follows:

- We propose a real-time online system that performs spatio-temporal action detection in untrimmed surveillance videos at ~ 100 frames per second.
- We propose a novel action localization network to detect activity agnostic tubelets which significantly reduces the processing time of the system.
- The proposed TMAS tubelet merging algorithm efficiently connects the tubelets in an online fashion and

produces spatio-temporal detections which are consistent across time as well as robust against varying length activities.

We evaluate the proposed approach on the DIVA (Deep Intermodal Video Analytics) dataset and obtain state-of-the-art results in terms of both speed and performance.

2. Literature Review

Convolutional Neural Networks (CNN) have been studied for video analysis and applied successfully for the action recognition problem [27, 4]. Earlier approaches fuse 2D frame features to extract temporal information [15], while recent works mostly apply 3D convolution to extract spatio-temporal features simultaneously [27, 4, 8, 29]. The work in [8, 25] use two stream network architectures to further exploit temporal dependencies.

Most action classifiers expect short trimmed videos but this is unrealistic for action recognition in real world surveillance videos. Predicting the temporal extents of actions is necessary for reliable recognition. In [21], a new layer is proposed to temporally localize activities in videos of MultiThumos dataset [32]. They represent actions as combinations of Gaussian distributions, which are predicted by their temporal convolution layer. Another major property of real world videos is that multiple actions can occur within the same scene. Most works on spatial action detection rely on a region proposal network [23] to detect multiple objects in each frame and combine them temporally to generate action tubelets [31, 20]. However, this approach becomes computationally inefficient as the number of proposals grows larger, making it unsuitable for real time methods. In [12], a 3D CNN network efficiently predicts frame-wise background-foreground segmentation map and extrapolates the action tubes. Duarte et al.[7] proposes a capsule based action detection network which segments and recognize actions jointly.

Spatio-temporal action detection in untrimmed videos requires more complex systems than aforementioned approaches. In [9], a frame level object detection and optical flow based model is proposed to solve spatio-temporal action detection for DIVA [19] dataset. They apply hierarchical clustering on all detected object regions in a video to obtain tube proposals and use optical flow to perform action classification. This approach is computationally expensive and not suitable for online processing. Our framework uses a 3D CNN network for spatio-temporal action segmentation which produces temporally consistent predictions and a fewer number of proposals. Additionally, our system processes videos in an online fashion without using computationally expensive methods (region proposal network, optical flow) and achieves better performance in real-time.

3. Methodology

3.1. Overview

The proposed system takes in a video clip as input and detects all activities in the form of tubelets. The system first operates on entire clip to spatio-temporally localize action tubelets. Once we extract potential tubelets, our classification system identifies all possible activities occurring within each tubelet. These action predictions are then fed into our TMAS system, which simultaneously filters and combines them into accurate and consistent action tubes. As an end result, we obtain spatio-temporal action detections over long untrimmed videos in an online real-time process. The following sections describe the different components of our system.

3.2. Localization Network

The tube extraction process is the first step in the pipeline, responsible for extracting all action tubes from the untrimmed video input. Localizing action regions both temporally and spatially is vital for the classification task as the length and location of action is unknown beforehand. Furthermore, each action tube could comprise of multiple actors performing multiple actions concurrently. This requires actor and action agnostic tube extraction.

As shown in Figure 2, first we divide the untrimmed videos into smaller clips, which are then forwarded into the localization network. Following an encoder-decoder approach, the network produces segmentation masks for action regions, each of which represents an action tubelet. These tubelets are then individually processed by subsequent components of our system. Since this bottom-up segmentation is performed for multiple frames simultaneously, we reduce the time taken to localize multiple actions within a video clip.

The proposed localization network uses an encoder-decoder structure, and extracts class-agnostic action features which can be used to generate segmentation masks; this requires an effective feature extractor as an encoder. To this end, we utilize a 3D convolution based encoder, I3D [4], to learn spatio-temporal features required for activity localization.

In the decoder, we use extracted action features to segment regions from the original input which contain activities. Depending on the task, this segmentation can be coarse (patches of regions) to fine grained (pixel level). We opt to produce relatively fine grained segmentations, keeping a balance between separable action tubes and a reasonable memory utilization. Following recent works in image segmentation [5, 6, 18] and video segmentation [7, 12], we use a decoder structure which combines transpose convolutions and upsampling. Stacking many transpose convolution layers is memory intensive and adds many param-

eters to the decoder, so we interleave upsampling operations to interpolate features. This results in a shallow decoder network, which prevents over-parameterization and avoids overfitting. Further architectural details of the localization network are presented in the supplementary materials.

Skip Connections: It is difficult to produce fine grained pixel level segmentations using features which have been spatially down-sampled by the encoder. To obtain these fine grained segmentations one would need a deep, memory intensive decoder. To circumvent this we pass low level features from various layers of encoder directly to corresponding layers of the decoder network, which has been found effective for image segmentation [18, 6]. This allows the decoder network to reincorporate essential features that were lost during encoding process, which results in improved pixel level segmentations.

Feature Pyramids: The objects present in surveillance videos have varying scale. Therefore, using information from layers with different feature resolutions helps in segmenting objects of different sizes. The authors in [6, 18] have recently shown that building feature pyramids from various layers, at different scales, aids in learning contextual information. Motivated by this, we stack features from various decoder layers (through upsampling) to obtain feature representations at different scales.

Atrous Convolution: Having contextual information for each pixel aids in the process of learning richer features for fine grained segmentation. Since conventional convolutions have a limited receptive field, they are unable to learn this contextual information without incurring a heavy memory overhead. To address this, we apply atrous convolution, with different receptive windows, on our final feature representation. Atrous convolution applies a convolution operation with dilated kernels centered around a pixel, which effectively increases the receptive field. We apply dilation at multiple rates to infer contextual relations among pixels with varying distances. The authors in [6] utilize atrous convolution for encoder-decoder based segmentation and observed improved performance.

Multi Layer Loss: The final output of our localization network is a segmentation mask where each pixel is assigned a probability of being a part of an action. To train the network, we compute a Binary Cross Entropy (BCE) loss, where each pixel can have a probability between zero (no activity) and one (activity). As our network is deep, only calculating the loss at the final layer affects the convergence time and the gradient update values for earlier layers. Hence, we apply the segmentation loss at multiple decoder layers, which improves the features learned in multiple layers and allows for more distributed backpropagation.

For a layer m with N total pixels, the loss is given as:

$$\mathcal{L}_m(\hat{y}, y) = -\frac{1}{N} \sum_{i=0}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (1)$$

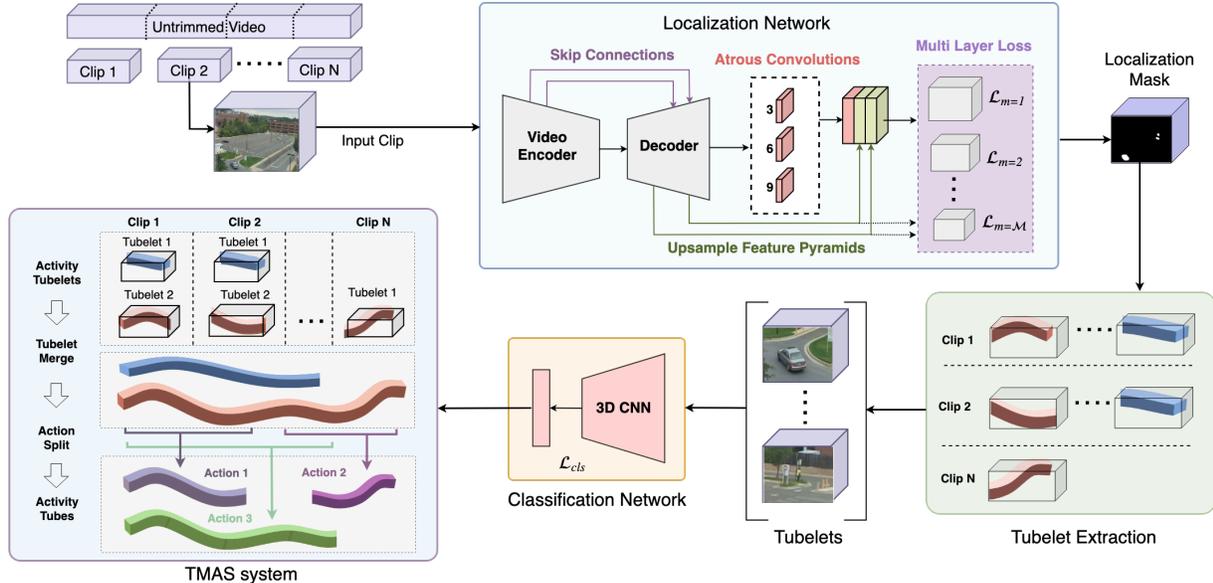


Figure 2. Overview of the proposed system for action detection in untrimmed videos. An untrimmed video is processed clip by clip and fed into the localization network, producing localization masks. Tubelet extraction produces tubelets for each clip, which are then classified and passed to our TMAS system. The classified tubelets are merged to create action-agnostic tubes, from which individual spatio-temporal action-specific detections are obtained.

where \hat{y}_i are predicted probability score of the pixel belonging to an activity and y_i are the ground truth labels of the pixels.

The combined loss for all M layers in the multi layer loss is given by:

$$\mathcal{L}_{loc} = \sum_{m=1}^M \mathcal{L}_m \quad (2)$$

Tubelet Extraction: The segmentation output for each clip is a probability mask which isolates potential action tubes. To obtain individual tubelets from this segmentation output, we threshold the output to create a binary mask followed by spatio-temporal connected component extraction. The connected component process will generate tubelets for all spatially and temporally linked pixels.

3.3. Tubelet Classification

The next step in our proposed system is tubelet classification. Our action classification network is a multi-label prediction network, which classifies the actions present within each tubelet. We treat this as a multi-label classification problem because actors can perform multiple activities simultaneously. For example, an actor can perform the actions *Riding* and *activity_carrying* at the same time. We use a 3D-Convolution based deep learning model [11] initialized with pre-trained weights on Kinetics [16] dataset for action classification. We modify the final layer of the model to have a $C + 1$ dimensional output, where C is the

number of action classes and the additional output is for the background class. A sigmoid activation is used in the final layer in place of a softmax as this is a multi-label classifier. We use BCE loss to train the classifier which is defined as,

$$\mathcal{L}_{cls}(\hat{y}, y) = -\frac{1}{C+1} \sum_{i=0}^C [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (3)$$

where \hat{y}_i is the prediction and y_i is the ground truth label.

3.4. TMAS Algorithm

To merge the tubelets and obtain the final action tubes, we propose the Tubelet-Merge Action-Split algorithm (TMAS). Each tubelet t_i is described as follows: $(f_1^i, f_2^i, \mathbf{b}^i, \mathbf{a}_c^i)$ where f_1^i is the start time, f_2^i is the end time, \mathbf{b}^i are the bounding boxes for each frame of the tubelet, and \mathbf{a}_c^i are the frame-level action probability scores for each action class $c \in \{0, 1, \dots, C\}$, where 0 is background. First, we merge the tubelets into action-agnostic tubes of varying length; then we split these action-agnostic tubes into a set of action-specific tubes which contain the spatio-temporal localizations for the various activities in the video.

Tubelet-Merge The procedure to merge tubelets into action-agnostic tubes is described in Algorithm 1. The temporally sequential stream of tubelets coming from the classification network are passed to the Tubelet-Merge procedure as input. The set of candidate tubes is initialized with

Algorithm 1 The Tubelet-Merge algorithm which merges tubelets into action-agnostic tubes. The CHECKEND function determines if a candidate tube becomes a final tube or is merged with another candidate.

Input: A stream of tubelets, \mathbf{S} , from the classifier
Output: A set of action-agnostic spatio-temporal tubes, T_{done}
Notation: $\text{Inter}_{\text{temp}}$ calculates temporal overlap between tubelets. $|\mathbf{M}[(t_c, *)]|$ returns the cardinality of the set $\{t : \mathbf{M}[(t_c, t)] > 0\}$.

```

1: procedure TUBELET-MERGE( $\mathbf{S}$ )
2:    $T_{prev}, T_{done} \leftarrow \{\}$   $\triangleright$  Initialize candidate and final tubes
3:    $\mathbf{M} \leftarrow$  initialize hash table
4:   while  $t_c$  in  $\mathbf{S}$  do  $\triangleright$  Continue until the stream of tubelets ends
5:     for all  $t_p$  in  $T_{prev}$  do
6:       if  $\text{Inter}_{\text{temp}}(t_p, t_c) > 0$  then
7:          $\mathbf{M}[(t_p, t_c)] \leftarrow \text{IoU}(t_p, t_c)$ 
8:       else
9:         CHECKEND( $t_p, T_{prev}, \mathbf{M}$ )
10:      append  $t_c$  to  $T_{prev}$   $\triangleright$  Tubelet becomes a candidate tube
11:   while  $T_{prev}$  is not empty do  $\triangleright$  Deals with remaining candidates
12:      $t_p \leftarrow T_{prev}[0]$ 
13:     CHECKEND( $t_p, T_{prev}, \mathbf{M}$ )
14:   return  $T_{done}$ 

1: function CHECKEND( $t_p, T_{prev}, \mathbf{M}$ )
2:   if  $|\mathbf{M}[(t_p, *)]| == 0$  then
3:     MOVE( $t_p, T_{prev}, T_{done}$ )  $\triangleright$  Moves  $t_p$  from  $T_{prev}$  to  $T_{done}$ 
4:   else if  $|\mathbf{M}[(t_p, *)]| == 1$  then
5:      $t_i \leftarrow \max_{t_i} \mathbf{M}[(t_p, t_i)]$ 
6:     if  $|\mathbf{M}[(*, t_i)]| == 1$  then
7:       MERGE( $t_p, t_i, T_{prev}, \mathbf{M}$ )
8:     else
9:       MOVE( $t_p, T_{prev}, T_{done}$ )
10:   else
11:      $t_i \leftarrow \max_{t_i} \mathbf{M}[(t_p, t_i)]$ 
12:     MERGE( $t_p, t_i, T_{prev}, \mathbf{M}$ )

1: function MERGE( $t_1, t_2, T_{prev}, \mathbf{M}$ )  $\triangleright$  Merges two candidate tubes
2:    $t_1 \leftarrow (f_1^1, f_2^1, \{\mathbf{b}^1, \mathbf{b}^2\}, \{\mathbf{a}^1, \mathbf{a}^2\})$   $\triangleright \{\}$  is concatenation
3:   remove  $t_2$  from  $T_{prev}$ 
4:    $\mathbf{M}[t_1, t_i] \leftarrow \mathbf{M}[t_2, t_i]$   $\triangleright$  Done for all  $t_i$  where  $\mathbf{M}[t_2, t_i] \geq 0$ 

```

the first tubelet. For each subsequent tubelet, we look for spatio-temporal overlap with the existing candidate tubes. This results in four possible outcomes: 1) If there is no overlap, the tubelet itself becomes a new candidate tube, 2) If there is a unique match found between a candidate tube and the tubelet, they are merged and become a new candidate tube, 3) if the tubelet has an overlap with multiple candidates, then the tubelet becomes a new candidate, 4) if multiple tubelets have an overlap with a single candidate tube, then the tubelet with the highest overlap is merged with that candidate and the other tubelets become separate candidate tubes. Once all tubelets are checked, the candidate tubes become the final action-agnostic tubes.

Action-Split From the action-agnostic tubes we obtain action-specific spatio-temporal localizations using the Action-Split procedure described in Algorithm 2. We start by smoothing out per-frame action confidence scores;

Algorithm 2 The Action-Split algorithm which converts the action-agnostic tubes into action-specific predictions.

Input: A set of action-agnostic tubes, T , and a set of actions, C
Output: A set of spatio-temporal action-specific tubes, A_G
Notation: The hyperparameters τ, α, β , and γ are described in the supplementary materials. $a_c^i[f]$ and $t_i[f, c]$ contain the action prediction scores and tube information at frame f , respectively.

```

1: procedure ACTION-SPLIT( $T$ )
2:    $A_G \leftarrow \{\}$   $\triangleright$  Initializes the action-specific tubes
3:   for all  $t_i$  in  $T$  do
4:      $t_{smooth} \leftarrow \text{SMOOTH}(t_i)$ 
5:     for all  $c$  in  $1 : C$  do  $\triangleright$  Loop through each action class
6:        $a_L \leftarrow \text{EXTRACT}(t_{smooth}, c)$ 
7:       append  $a_L$  to  $A_G$ 
8:   return  $A_G$ 

1: function SMOOTH( $t_i$ )
2:   for all  $f$  in  $f_1^i : f_2^i$  do
3:      $a_c^i[f] \leftarrow \frac{1}{2\tau+1} \sum_{k=-\tau}^{\tau} a_c^i[f+k]$ 
4:   return  $t_i$ 

1: function EXTRACT( $t_i, c$ )  $\triangleright$  Extracts tubes of a specific class
2:    $A_L, a_l \leftarrow \{\}$   $\triangleright$  Initialize extracted action tubes and a placeholder
3:    $count \leftarrow 0$ 
4:   for all  $f$  in  $f_1^i : f_2^i$  do
5:     if  $a_c^i[f] > \alpha$  then  $\triangleright$  Continue current action tube
6:       append  $t_i[f]$  to  $a_l$ 
7:        $count \leftarrow 0$ 
8:     else
9:        $count \leftarrow count + 1$ 
10:    if  $count > \beta$  then  $\triangleright$  Current action tube is finished
11:      append  $a_l$  to  $A_L$ 
12:       $a_l \leftarrow \{\}, count \leftarrow 0$ 
13:   remove tubes shorter than  $\gamma$  from  $A_L$ 
14:   return  $A_L$ 

```

which accounts for fragmentation caused by action misclassifications. Then we build the action-specific tubes by checking for continuous occurrences of each action class; this allows several occurrences of the same activity to occur within a single tube. For instance, a person *walking* might stop and *stand* for several seconds and start walking again; this entire sequence will be contained in a single spatio-temporal tube, but the Action-Split procedure will correctly generate two separate instances of *activity_walking* and one instance of *activity_standing*. To be robust to classification errors, action tubes with the same action label that are within a limited temporal neighborhood are combined together to form a single continuous action prediction.

Runtime Complexity The worst-case runtime of our TMAS algorithm is $\mathcal{O}(n^2)$, where n is the total number of candidate tubes at any given time. However, we sequentially process our tubelets and constantly shift the candidate tubes which can not have any possible future match to the set of final tubes. Therefore, the set of candidate tubes at any particular time is reasonably small and our TMAS algorithm contributes negligible overhead to our system's overall computation time.

4. Experimental Setup

4.1. DIVA Dataset

The DIVA dataset is a large-scale spatio-temporal action detection dataset with untrimmed surveillance videos. It consists of videos from the VIRAT [19] dataset with added annotations for action detection. There are 64 videos (2.47 hours) in training and 54 videos (1.93 hours) in the validation set, with annotations for all the activities. There are 246 videos (10.11 hours) in the held out test set for which the annotations are not made public. There are 40 different types of activities in the dataset which can be broadly categorized into activities involving only people, activities involving only vehicles, and activities involving both people and vehicles. Of these 40 classes only 18 classes are used for evaluation of the system: *activity_carrying*, *vehicle_turning_right*, *vehicle_turning_left*, *Closing*, *Opening*, *Exiting*, *Entering*, *Talking*, *Transport_HeavyCarry*, *Unloading*, *Pull*, *Loading*, *Open_Trunk*, *Closing_Trunk*, *Riding*, *specialized_texting_phone*, *specialized_talking_phone*, and *vehicle_u_turn*. Remaining classes are considered *hard_negative*, or background. Videos in this dataset are captured from surveillance cameras mounted on buildings, mostly overlooking parking lots or streets, where the activities take place. All videos are high resolution, either 1920×1080 or 1280×720 , with variation in scale, orientation of objects, and camera viewpoints.

4.2. Training Details

Network Training DIVA videos are high resolution and it is not feasible to train a localization model at that scale. Therefore, we rescale the videos to a lower resolution of 800×448 , while maintaining the aspect ratio. A stack of 16 frames is fed to the localization network to obtain the binary segmentation masks. The network is trained using SGD [24], with a learning rate of $1e - 3$ for about 100000 iterations. Our classification model is trained with a clip length of 16 frames (with a skip rate of 1 to obtain a second long clip) and a spatial-resolution of 112×112 . We use the ADAM optimizer [17] with a learning rate of $1e-4$ to train the classifier for 75000 iterations on a single NVIDIA GeForce Titan X GPU.

Data Pre-processing For the data used in training the localization network, we apply random frame jitter and crop to simulate shaking cameras and changes in scale. These augmentations add variation to the training data, and greatly reduces overfitting.

The spatial resolution of activity tubelets can be of any size, but the input to the classification model is fixed. To address this issue we crop a square region encompassing the activity within the tubelet and rescale it to the size expected by the classification model. This approach helps us

in maintaining the aspect ratio of the objects in the input clips. In temporal domain, the activity tubelets can be of arbitrary length, so we extract multiple clips by applying a sliding window over the length of the entire tubelet. To train the classification model, clips are generated using both the ground-truth annotations as well as the tubelets from the localization network. This increases the amount of training data seen by the classifier, allowing for improved performance on the test set. Training clips from the localization model that do not contain any of the 18 activity classes used for evaluation are labelled as *hard_negative*.

One of the challenging issues with the DIVA dataset is data imbalance. Refer to Figure 3 for distribution of per-class action instances/frames from DIVA ground-truth training set. Since clips are extracted using a sliding window, the class imbalance is further exasperated for longer, repetitive activities like *activity_carrying* and *Riding* as well as shorter, atomic activities like *Loading*, *Unloading*, *Entering*, and *Exiting*. To overcome this issue, we use augmentation techniques such as multi-scale cropping, frame reversal, spatial-jittering, and flipping of frames along the vertical axis. In multi-scale cropping we randomly crop a region around the center of a clip at different scales to simulate zoom-in and zoom-out effect to have actors/objects at different scales. We reverse the order of frames in a clip to generate new clips for specific pairs of classes such as (*Opening*, *Closing*), (*Loading*, *Unloading*), (*Entering*, *Exiting*), and (*Open_Trunk*, *Close_Trunk*) to increase the number of samples for these classes. With spatial-jittering we simulate camera shake due to wind, which occasionally occurs in test videos.

5. Results

In this section we present evaluation results of individual components in the proposed architecture and discuss the overall system performance.

5.1. Localization

We run several ablations to evaluate the effect of the different components - feature pyramid, atrous convolution, and multi-layer loss - of our localization network. Table 1 shows the performance of different variations of our action localization network on the DIVA validation set in terms of IoU scores. The baseline network is an I3D based encoder-decoder architecture with skip connections. The addition of these components improves our IoU scores considerably.

5.2. Classification

We experiment with multiple classification models to determine the optimal network architecture for our system. A comparison of their performance on the validation set is shown in Table 2. We use average F1-Score as a metric to compare the different models. For a fair comparison, all

Model	IoU
Baseline	0.540
Baseline+FPN	0.572
Baseline+FPN+MLL	0.613
Baseline+FPN+MLL+Atrous	0.628

Table 1. Ablation experiments to study the effect of different components of the localization network.

Architecture	Precision	Recall	F1-Score
I3D [4]	0.36	0.31	0.33
R(2+1)D [28]	0.43	0.39	0.40
P3D [22]	0.43	0.41	0.41
3D-ResNet [11]	0.46	0.43	0.44

Table 2. Average Precision, Recall and F1-score on DIVA validation set with different network architectures for classification. We use pre-trained weights on Kinetics dataset [16] to initialize all these models before training.

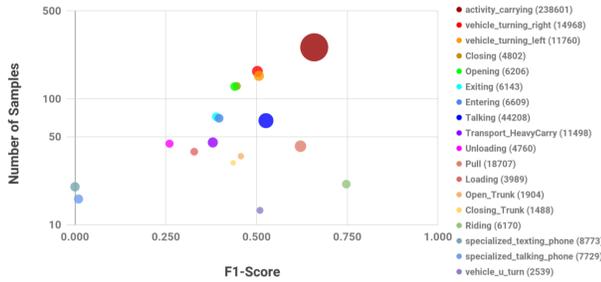


Figure 3. Number of samples per class and F1-Score achieved by the trained classifier on the DIVA validation set. The size of each bubble indicate the number of frames for an activity class. The exact numbers are indicated next to the activity name in the legend on the right.

models are initialized with pre-trained weights on the Kinetics dataset [16] and are trained with the same settings. We observe that the 3D-ResNet based model, WideResNet-50 [11], outperforms the other architectures.

We also analyze the class level performance of the classification model on the validation set. The results using our final classification model are shown in Figure 3. We can observe that our network performs better on longer activity classes such as *Riding* and *activity_carrying* with higher number of training samples when compared to shorter activity classes like *Unloading* and *Loading* which have fewer samples. Activity classes *specialized_texting_phone*, *specialized_talking_phone* have the lowest F1-Score as these have considerably fewer samples and appear very similar to *activity_standing* which is a background activity.

5.3. System Evaluation

Metrics We evaluate the performance of our system using several metrics: probability of missed detection at fixed

$P_{miss}@R_{FA}$	0.01	0.03	0.1	0.15	0.2	1.0
[9]	0.87	0.80	0.61	0.56	0.54	0.36
Ours ₁₂	0.87	0.73	0.55	0.50	0.46	0.34
Ours ₁₈	0.87	0.77	0.61	0.54	0.51	0.39

Table 3. Comparison of our temporal localization results on DIVA(AD) validation with Gleason et al.[9] on 12 classes from TRECVID-2018 [1]. The third row presents our results on 18 classes from TRECVID-2019 [2] [3] evaluation.

Team	$P_{miss}@0.15R_{FA}$	$P_{miss}@1.0R_{FA}$
P4	0.872	0.704
Xu et al. [30]	0.863	0.720
P3	0.759	0.624
P1	0.710	0.603
P2	0.624	0.621
Gleason et al. [9]	0.618	0.441
Ours	0.599	0.436

Table 4. P_{miss} scores on the DIVA test set for 12 classes from TRECVID-2018 (AD). The best performance is indicated in bold.

rate of false alarm per minute ($P_{miss}@R_{FA}$), probability of missed detection at fixed time-based false alarm per minute ($P_{miss}@T_{FA}$), and partial area under the Detection Error Tradeoff curve (AUDC). These measure the quality of action detections both temporally, for the action detection (AD) task, and spatio-temporally, for the action-object detection (AOD) task. To calculate these metrics, a one-to-one correspondence is found between the ground-truth actions and the detected actions; ground-truth actions without a corresponding detection are missed detections, while detections without corresponding ground-truth actions are false alarms. For more detailed explanations of the different evaluation metrics as well as evaluation code, we refer to TRECVID 2018 [1] and TRECVID 2019 [2][3].

Comparison For fair comparison with an earlier work [9], we compute $P_{miss}@R_{FA}$ on the validation set for the same set of 12 classes which are reported in their results. The comparison is presented in Table 3. In Table 4, we compare our results with other participants from TRECVID-2018 [1] challenge. The metric used here are P_{miss} score at R_{FA} thresholds of 0.15 and 1.0.

In Table 5, we present the results from the latest TRECVID-2019 [2] [3] leaderboard which show that our approach achieves best results on $P_{miss}@0.15T_{FA}$, $P_{miss}@0.15R_{FA}$, and comparable results on the AUDC score. As these results are on the held out test set, we only present scores shared by the independent evaluator.

The results presented so far show the performance of the system only on temporal detection of activities. We evaluate our model on the spatio-temporal detection task (AOD), results of which are presented in Table 6. These results are on the DIVA validation set. To the best of our knowledge,



Figure 4. Qualitative results of our system on some sample validation videos (shown only one frame, full videos are provided in the supplementary material). This demonstrates the ability of our system to handle varying scenes, object scales, and action types.

Team	$P_{miss@0.15} TFA$	$P_{miss@0.15} RFA$	AUDC
Fraunhofer	0.7747	0.8474	0.8270
vireoJD-MM	0.5482	0.7284	0.6012
NTT_CQUPT	0.5112	0.8725	0.6005
Hitachi	0.5099	0.8240	0.5988
BUPT-MCPRL	0.4328	0.7491	0.5240
MUDSML	0.3915	0.7979	0.4840
Ours	0.3858	0.7022	0.4909

Table 5. Temporal localization results (AD) on DIVA test set from TRECVID-2019 leaderboard. All the metrics relate to the miss-rate, so lower values indicate better performance of the system.

$P_{miss}@RFA$	0.01	0.03	0.1	0.15	0.2	1.0
Ours ₁₈	0.939	0.895	0.751	0.773	0.731	0.598

Table 6. The spatio-temporal action localization results (AOD) on DIVA validation set on 18 classes from TRECVID-2019.

there are no published works which have reported results on the AOD task to compare with.

6. Discussion

Online action detection Online action detection is different from traditional action detection as its goal is to detect an action as it happens. Our proposed system is an online action detection system as it can process frames as they come in and can perform localization, classification and temporal segmentation of activities with little or no delay. Other systems such as [9] are restricted to offline detection as they rely on object detection for every frame in the video, requiring access to future frames to generate tube proposals. This is a major advantage of our system as it can be readily used in real-world surveillance applications.

Qualitative Analysis We present some qualitative results of our system in Figure 4. Our system performs well on different viewpoints, action scales, and action types. We are able to detect activities involving multiple actors (*Talking*), as seen in the leftmost image, as well as activities involving interactions between a person and a vehicle (*Closing*), as seen in the rightmost image. Since we do not rely on

frame-based object detection, we produce fewer detections for object which are not involved in an activity, as is evident from the stationary vehicles in each of the images.

7. Conclusion

In this work, we propose a real-time online system to detect activities in untrimmed surveillance videos. The proposed system consists of three main components which include tubelet generation, action classification, and tubelet merging. The proposed approach process short video clips at a time which helps in real-time online processing. The efficient merging of tubelets using TMAS algorithm makes the spatio-temporal detections robust to varying length activities. In contrast to existing approaches, it does not require frame level object detection or optical flow extraction which are computationally expensive and need externally trained models. The proposed method provides state-of-the-art results on the DIVA dataset with a processing speed of over 100 fps.

References

- [1] Trecvid 2018 actev: Activities in extended video 1b eval august 27, 2018 - december 04, 2018.
- [2] Trecvid 2019 actev: Activities in extended video.
- [3] G. Awad, A. Butt, K. Curtis, Y. Lee, J. Fiscus, A. Godil, A. Delgado, A. F. Smeaton, Y. Graham, W. Kraaij, and G. Quénot. Trecvid 2019: An evaluation campaign to benchmark video activity detection, video captioning and matching, and video search & retrieval. In *Proceedings of TRECVID 2019*. NIST, USA, 2019.
- [4] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [5] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [6] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018.

- [7] K. Duarte, Y. Rawat, and M. Shah. Videocapsulenet: A simplified network for action detection. In *Advances in Neural Information Processing Systems*, pages 7610–7619, 2018.
- [8] C. Feichtenhofer, H. Fan, J. Malik, and K. He. Slowfast networks for video recognition, 2018.
- [9] J. Gleason, R. Ranjan, S. Schwarcz, C. Castillo, J.-C. Chen, and R. Chellappa. A proposal-based solution to spatio-temporal action detection in untrimmed videos. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 141–150. IEEE, 2019.
- [10] C. Gu, C. Sun, D. A. Ross, C. Vondrick, C. Pantofaru, Y. Li, S. Vijayanarasimhan, G. Toderici, S. Ricco, R. Sukthankar, et al. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6047–6056, 2018.
- [11] K. Hara, H. Kataoka, and Y. Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6546–6555, 2018.
- [12] R. Hou, C. Chen, and M. Shah. An end-to-end 3d convolutional neural network for action detection and segmentation in videos. *arXiv preprint arXiv:1712.01111*, 2017.
- [13] H. Idrees, A. R. Zamir, Y.-G. Jiang, A. Ghorban, I. Laptev, R. Sukthankar, and M. Shah. The thumos challenge on action recognition for videos “in the wild”. *Computer Vision and Image Understanding*, 155:1–23, 2017.
- [14] V. Kalogeiton, P. Weinzaepfel, V. Ferrari, and C. Schmid. Action tubelet detector for spatio-temporal action localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4405–4413, 2017.
- [15] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR ’14*, pages 1725–1732, Washington, DC, USA, 2014. IEEE Computer Society.
- [16] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [17] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [18] A. Kirillov, R. Girshick, K. He, and P. Dollár. Panoptic feature pyramid networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6399–6408, 2019.
- [19] S. Oh, A. Hoogs, A. Perera, N. Cuntoor, C.-C. Chen, J. T. Lee, S. Mukherjee, J. Aggarwal, H. Lee, L. Davis, et al. A large-scale benchmark dataset for event recognition in surveillance video. In *CVPR 2011*, pages 3153–3160. IEEE, 2011.
- [20] X. Peng and C. Schmid. Multi-region two-stream r-cnn for action detection. In *European conference on computer vision*, pages 744–759. Springer, 2016.
- [21] A. J. Piergiovanni and M. S. Ryoo. Temporal gaussian mixture layer for videos. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pages 5152–5161, 2019.
- [22] Z. Qiu, T. Yao, and T. Mei. Learning spatio-temporal representation with pseudo-3d residual networks. In *proceedings of the IEEE International Conference on Computer Vision*, pages 5533–5541, 2017.
- [23] 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*, pages 91–99, 2015.
- [24] H. Robbins and S. Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- [25] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 568–576. Curran Associates, Inc., 2014.
- [26] G. Singh, S. Saha, M. Sapienza, P. H. Torr, and F. Cuzzolin. Online real-time multiple spatiotemporal action localisation and prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3637–3646, 2017.
- [27] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [28] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018.
- [29] X. Wang, R. Girshick, A. Gupta, and K. He. Non-local neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [30] H. Xu, A. Das, and K. Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *Proceedings of the IEEE international conference on computer vision*, pages 5783–5792, 2017.
- [31] Z. Yang, J. Gao, and R. Nevatia. Spatio-temporal action detection with cascade proposal and location anticipation. *arXiv preprint arXiv:1708.00042*, 2017.
- [32] S. Yeung, O. Russakovsky, N. Jin, M. Andriluka, G. Mori, and L. Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision*, 2017.