YoTube : searching action proposal via recurrent and static regression networks

Zhu, Hongyuan; Vial, Romain; Lu, Shijian; Peng, Xi; Fu, Huazhu; Tian, Yonghong; Cao, Xianbin

2018

Zhu, H., Vial, R., Lu, S., Peng, X., Fu, H., Tian, Y., & Cao, X. (2018). YoTube : searching action proposal via recurrent and static regression networks. IEEE Transactions on Image Processing, 27(6), 2609-2622. doi:10.1109/TIP.2018.2806279

<https://hdl.handle.net/10356/80911>

<https://doi.org/10.1109/TIP.2018.2806279>

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at: <https://doi.org/10.1109/TIP.2018.2806279>.

Downloaded on 26 Apr 2023 22:49:57 SGT

YoTube: Searching Action Proposal via Recurrent and Static Regression Networks

Hongyuan Zhu*, Romain Vial*, Shijian Lu, Xi Peng, Huazhu Fu,
Yonghong Tian, Xianbin Cao

Abstract—In this paper, we propose *YoTube*—a novel deep learning framework for generating action proposals in untrimmed videos, where each action proposal corresponds to a spatial-temporal tube that potentially locates one human action. Most of the existing works generate proposals by clustering low-level features or linking image proposals, which ignore the interplay between long-term temporal context and short-term cues. Different from these works, our method considers the interplay by designing a new recurrent *YoTube* detector and static *YoTube* detector. The recurrent *YoTube* detector sequentially regresses candidate bounding boxes using RNN learned long-term temporal contexts. The static *YoTube* detector produces bounding boxes using rich appearance cues in every single frame. To fully exploit the complementary appearance, motion, and temporal context, we train the recurrent and static detector using RGB and Flow information. Moreover, we fuse the corresponding outputs of the detectors to produce accurate and robust proposal boxes and obtain the final action proposals by linking the proposal boxes using dynamic programming with a novel path trimming method. Benefiting from the pipeline of our method, the untrimmed video could be effectively and efficiently handled. Extensive experiments on the challenging UCF-101, UCF-Sports, JHMDB datasets show superior performance of the proposed method compared with the state of the arts.

I. INTRODUCTION

Action proposal aims to extract a small number of spatial-temporal paths to cover all potential regions corresponding to human actions. As it could significantly reduce the size of search space, there is increasing attention in video analytics tasks [1]–[12]. Different from action recognition [13]–[19], action proposal mines all potential action regions from one video rather than classifying the whole video into existing

categories. Comparing with action detection [20]–[22], action proposal generates generic action regions instead of focusing on the specific action defined by training data. A conceptual illustration of the difference among action proposal, action recognition, and action detection could be found in Fig. 1.

Fig. 1. Conceptual illustration of the difference among action proposal, action recognition, and action detection. Action proposal tries to extract all possible human actions which are highlighted using dash-lines. Action recognition aims to classify the whole video into specific categories, e.g., ‘Jump’. Action detection aims at detecting when and where a specific action appears (highlighted using solid-lines).

Despite the success of object proposals in images [23], [24], it is extremely challenging to generate action proposals in videos due to following reasons. First, the extensively investigated image object proposal only relies on appearance and spatial cues, whereas action proposal takes appearance, motion and temporal information into consideration. What is more challenging is learning effective actionness cues to differentiate human actions from commonly occurred background clutters and other dynamic motion, given the diversity and variations of human actions. Second, the search space of action proposals is exponentially larger than image object proposal since the former is with the additional temporal dimension. In practice, it is infeasible to enumerate all possible candidates to pick action proposals. Third, the raw video content is generally untrimmed, which brings in temporal noises and needs further elaborate post-processing to trim the action paths.

To holistically address the above challenges, we propose a novel deep learning framework called *YoTube* that generates spatially compact and temporally smooth action proposals for untrimmed videos by simultaneously considering the appearance, motion and temporal information. In details, our framework consists of a novel recurrent *YoTube* detector and a static *YoTube* detector. The recurrent *YoTube* detector is based on a novel recurrent regression network that sequentially predicts the bounding boxes using adjacent frame temporal contexts in one-shot. The static *YoTube* detector is designed for better exploiting the rich global appearance and motion cues within each individual frame. The outputs of these two networks are further fused to exploit the complementary

H. Zhu and H. Fu are with Institute for Infocomm Research, A*Star, Singapore (email: {zhuh,fuh}@i2r.a-star.edu.sg).

R. Vial is with the Mines ParisTech, France (e-mail:romain.vial@mines-paristech.fr).

S. Lu is with Nanyang Technological University, Singapore (email:Shijian.Lu@ntu.edu.sg)

X. Peng is with College of Computer Science, Sichuan University (e-mail:pangsaai@gmail.com)

Y. Tian is with National Engineering Laboratory for Video Technology (NELVT), School of EECS, Peking University, Beijing, China (email:yhtian@pku.edu.cn)

X. Cao is with School of Automation Science and Electrical Engineering, Beihang University, Beijing, China (email:xbcao@buaa.edu.cn)

This work was mainly done when R. Vial was interned with the Institute for Infocomm Research. H. Zhu and R. Vial are equally contributed (*Corresponding authors: H. Zhu, R. Vial.*)

Xi Peng is funded by National Nature Science Foundation of China under grant No.61432012 and No.U1435213, and the Fundamental Research Funds for the Central Universities under grant No.YJ201748. Yonghong Tian is supported by the National Natural Science Foundation of China under contract No. U1611461, No. 61390515, and No. 61425025.

information between short-term and long-term contexts. After that, our method gives initial action proposals by linking the candidate action boxes in terms of their actionness score and overlap in the spatial-temporal domain. Furthermore, we propose a novel temporal path trimming method to handle the untrimmed videos by utilizing the actionness and background score transition pattern.

The contributions of our paper lie in following aspects:

- A novel deep learning framework for action proposal is proposed, which learns discriminative actionness cues from video by considering the short-term appearance, motion information, and the long-term temporal information.
- A novel recurrent regression network is introduced to capture the long-term temporal information for action proposal, which is ignored in recent works.
- An efficient and accurate path trimming technique is proposed to deal with untrimmed videos.
- Extensive experiments are carried out on UCF-101, UCF-Sports and JHMDB dataset, which demonstrate the superior performance of our method.

A preliminary conference version of our work appeared in [25], which only adopts the RGB information and applies a less efficient path trimming method. Moreover, the experiment was conducted only on the UCF-101 dataset. This paper extends the conference work by additionally considering the motion information from flow images and proposing a more efficient and accurate path trimming method to trim the action proposal. Furthermore, this work performs more detailed analysis and extensive experiments involving more state-of-the-art methods on three challenging datasets (UCF-101, UCF-Sports, and JHMDB dataset).

II. RELATED WORK

Recurrent Neural Network:

Recurrent Neural Network, especially Long-Short Term Memory (LSTM) [26] has become popular for sequence generation and prediction tasks. A detailed survey of recent RNN models and applications could be found in [27]. In this work, we mainly discuss RNNs based action recognition and detection.

Veeriah *et al.* [28] proposed a differential gating scheme to capture the changes between two successive frames. Donahue *et al.* [29] developed a novel recurrent convolutional architecture and successfully applied it to video recognition, image description and video narration. Wu *et al.* [30], [31] and Ng *et al.* [19] demonstrated that an average fusion of LSTMs with appearance and flow boosts the prediction performance. Although our work also employs LSTM to learn long-term temporal contexts, it is different from existing works since we employ the LSTM to predict bounding boxes instead of class labels. Moreover, we trained a CNN using a larger resolution images for feature extraction, which facilitates the action proposal task.

Recently, RNN has been applied to refine tracking-by-detection result. For example, Ni *et al.* [32] applied an LSTM to refine the tracked objects or human parts captured in each image frame. Stewart and Andriluka [33] applied an LSTM

to aggregate contexts from adjacent detections so that the detected face in the image is progressively refined. Comparing with these methods, the proposed recurrent regression network is designed for video action proposal, which predicts bounding boxes in one-shot without a usage of other detectors and multiple pass regime, thus embracing computational efficiency.

Regression based Object Detection: Most recent deep learning detection methods perform detection by classifying object proposals (e.g. selective search [23], EdgeBox [34]) or directly regressing the coordinates of bounding boxes based on local features. The typical methods include but not limited to Region Proposal Network [35] and SSD [36]. Recently, [37] proposes YOLO to perform inference using global image feature and achieves impressive results, which exploits the context information of the whole image to avoid the influence from the background.

To reduce the influence from the background, our work also exploits the global image features for bounding box regression. The architecture of our CNN is different from that of YOLO. More specifically, we replace the last two fully connected layers of YOLO with a locally connected dense layer, thus reducing the computational overhead. Experimental result verifies that such a difference improves the accuracy of our method by nearly 5%. Moreover, YOLO and other detectors are mainly designed for image object detection, which neglects the useful temporal context information. In contrast, our work explores the regression capability of RNN for video action proposal. The extensive experimental study reveals that the combination of RNN and CNN yields superior performance over either one of them alone.

Action Recognition Following the impressive performance of CNNs in image recognition, deep learning approaches were applied to action recognition. A thorough survey on recent deep action recognition methods could be found in [38] and [39]. Here we summarize the influential works according to different pipelines.

Existing deep action recognition architectures could be divided into three groups according to [39]. To be exact, 1) image based action recognition methods directly extract the off-the-shelf CNN features pre-trained on the ImageNet and then pass the features through a learned SVM classifier [13], [14]; 2) end-to-end snippet learning methods learn video features using the appearance cues of short video snippets. For example, Ji *et al.* [40] introduced the 3D-CNN that operates on the stacked video frames. Karpathy *et al.* [15] compared several similar 3D-CNN architectures on the large-scale video classification task. Tran *et al.* [16] proposed the C3D model which has inspired numerous works, such as R-C3D [41] and Segment-CNN [42]. Recently, Simonyan and Zisserman [17] propose a two-stream approach to break down the video feature learning into the learning of separate spatial and temporal clues, thus greatly reducing the learning overload in C3D and improving the recognition performance. In [18], the two-stream approach has been extended with dense trajectories. 3) long-term temporal modeling overcomes the limitations of short-snippet learning method by using RNN/LSTM to capture the long range temporal contexts. For example, [29] and [30] trained an LSTM on the top of CNN for video recognition.

Fig. 2. Conceptual illustration of our method: we utilize the regression capability of RNN and CNN to directly regress the sequences of bounding boxes and then seam the boxes into longer action proposals using path linking and trimming.

Ng *et al.* [19] further stacked multiple layers of LSTM and compared different pooling strategies. The technical difference with the long-term temporal modeling is discussed in the section of Recurrent Neural Network.

Our work explores to combine the advantages of short-snippet modeling and long-term temporal modeling in action recognition for action proposal. We show that both networks are complementary with each other and could deliver significant performance improvement.

Action Proposal: To reduce the search space, action proposal generates sequences of bounding boxes with good localization of candidate human actions in the spatio-temporal domain, which avoids the rigid structure requirement of early works [43]–[46].

Unsupervised image proposals [23], [24] have been extended for video action proposal. Jain *et al.* [47] extend Selective Search [23] by clustering the videos into voxels and then hierarchically merging the voxels to produce action proposals. Similarly, Oneata *et al.* [48] extend [24] by introducing a randomized supervoxel segmentation method for proposal generation. Inspired by the video segmentation method in [49], Jain *et al.* [50] propose to generate action proposals by clustering long term point trajectories with improved speed and accuracy. However, they are based on low-level features which make difficulties in handling videos with rich motion.

Supervised detectors have also been introduced to the action proposal by learning actionness cues using labels. Yu *et al.* [51] use human detector and generate the action proposal by using max sub-path search. Inspired by the success of deep learning, Gkioxari and Malik [52] propose to train two stream R-CNN networks [53] which learns actionness cues with selective search to detect action regions. They link the high scored action boxes to form action tubes. In class-specific action detection, Saha *et al.* [21] and Peng *et al.* [22] propose to use region proposal networks (RPN) to generate frame proposals and then classify these regions by Fast-RCNN. Detection-and-tracking methods have also been used for action localization and action proposal. Weinzaepfel [20] train a two-stream R-CNN to detect action regions and an additional instance-level detector to track the regions with Spatio-Temporal Motion Histogram. Inspired by these works, Li *et al.* [54] also trains RPN [35] to replace R-CNN in [20]

for generating proposal boxes. Moreover, their method uses an improved method of [51] to generate action proposals. The missed detections are remedied by tracking-by-detection and their method has achieved the state-of-the-art performance.

Most of these works [47], [48], [50]–[52] generate proposals for each frame individually without considering the temporal contexts or just considering the contexts in very short snippets [54]. Moreover, they generally work on trimmed videos [47], [48], [50], [52]. To handle untrimmed video, extra detectors need to be trained using low-level features, thus leading to errors accumulation and higher time consumption [20]–[22], [54].

Different from the above works, we propose a 4-way network fusion scheme to combine the short-term and long-term information given by the recurrent and static regression networks respectively. Our method uses a novel recurrent regression network to capture the long-term temporal context which is largely neglected in recent works of action proposal. In addition, we use the global features instead of local features to perform inference, thus reducing the interference from background clutter. Finally, we design an efficient path trimming technique which is capable of handling untrimmed videos directly without requiring time-consuming techniques of existing methods [20]–[22], [54].

Information Fusion: Fusing information from multiple sources has shown effectiveness in various tasks. A comprehensive survey on this topic could be found in [55] and [56]. Here, we mainly focus on the works related to object detection, action/event recognition and action detection based on neural networks.

A popular approach is to perform fusion at the input level. Some static object detection methods stack images generated from multiple-sensors such as multi-spectral camera [57] and RGB-D camera [58]. To achieve action recognition, C3D [16], [40] applies the 3D-Convolution on stacked image frames. Another popular approach is to perform fusion at the feature level. In RGB-D object/scene-recognition, Wang *et al.* [59], [60] and Zhu *et al.* [61] proposed to use auto-encoder and correlation analysis to fuse the features from RGB and Depth, respectively. Wu *et al.* [30] and Feichtenhofer *et al.* [62] proposed to directly concatenate the RGB and Flow features for action recognition. The third approach is to perform fusion at

the output level. Recently, Zhang *et al.* [63] propose using an auto-encoder to perform late fusion on the results of dense trajectories generation, scene classification and object detection. Zhang *et al.* [64] also proposed dynamically fusing the motion and image cues for video description. Simonyan and Zisserman [17] proposed a late fusion scheme to train a two-stream network for the purpose of extracting features from RGB and Flow frames.

Our model adopts the late fusion since [17], [21], [22], [30] has validated its effectiveness in class-specific action recognition and action detection. The major difference between our work and existing ones is that we perform 4-way fusion by considering the recurrent and static regression networks trained on RGB and Flow frames for action proposal. Another advantage of our model is that it does not require any external multi-channel sensors such as [57]–[61]. As a result, our method is easier to deploy. It should be pointed out that [65] is proposed for object tracking, in which the “fusion” concept is different from the score/modality fusion concept discussed in our work. The fusion concept here is more about “model updating”, i.e., updating the old model by using newly classified examples.

III. METHODOLOGY

The proposed method inputs an untrimmed video and outputs the action proposal accordingly. Fig. 2 illustrates the flowchart of our framework which consists of two steps: 1) using recurrent and static *YoTube* to predict sequential candidate action bounding boxes; 2) linking and trimming action path.

Fig. 3. Using a video snippet as the input, the proposed recurrent *YoTube* detector first extracts discriminative features from each frame and then applies the LSTM to regress the coordinates of the bounding boxes. The bounding boxes of each frame is estimated by considering the rich spatial and temporal context in the forward direction.

A. *YoTube* for Action Candidate Boxes Generation

One limitation of existing deep action proposal methods is that they either process a video frame-by-frame [52] or separate spatial and temporal information learning into two isolated processes [20], [54]. Moreover, temporal dynamics and contexts between two adjacent frames have been proven to be useful in action classification and video description [29]. In this paper, we design a neural network based fusion framework to incorporate the appearance, motion and temporal context learning in an end-to-end optimizable manner.

We first describe the **recurrent *YoTube*** — a recurrent regression network, which is used to predict the bounding

boxes for each frame by leveraging the temporal contexts from adjacent frames. Fig. 3 depicts the architecture of recurrent *YoTube* detector which passes the frame f_t at time t into a CNN to produce a fixed-length feature x_t . Then a recurrent LSTM maps x_t and the previous time step hidden state h_{t-1} into a new hidden state h_t and bounding boxes o_t . The inference is sequentially conducted from top to bottom as illustrated in Fig. 3. Benefit from the used network, the context in earlier frames t_l ($t_l < t$) can be propagated into the current frame t .

The output o_t for the frame t is a $K \times K \times (B \times 5 + |S|)$ tensor which encodes the output bounding boxes information. Specifically, the image is divided into $K \times K$ grids and each grid cell will predict B bounding boxes parameterized by (x, y, w, h, c) , where (x, y) represents the center of the box relative to the bounds of the cell. The width w (height h) is normalized with respect to the image width (height). The confidence c predicts the IoU between the predicted box and the ground-truth. Moreover, each cell will also predict a score tuple $S = (s_{ac}, s_{bg})$, where s_{ac} and s_{bg} is an actionness score and a background score for the given cell, respectively.

The loss function is defined as a sum-squared error between the prediction o_t and the ground-truth $\hat{o}_t$ for the simplicity in optimization [37]:

$$\begin{aligned}
 & \lambda_{coord} \sum_{i=0}^{K^2} \sum_{j=0}^B 1_{ij}^{obj} \| (x_i, y_i) - (\hat{x}_i, \hat{y}_i) \|^2 \\
 & + \lambda_{coord} \sum_{i=0}^{K^2} \sum_{j=0}^B 1_{ij}^{obj} \| (\sqrt{h_i}, \sqrt{w_i}) - (\sqrt{\hat{h}_i}, \sqrt{\hat{w}_i}) \|^2 \\
 & + \sum_{i=0}^{K^2} \sum_{j=0}^B 1_{ij}^{obj} (c_i - \hat{c}_i)^2 \\
 & + \lambda_{noobj} \sum_{i=0}^{K^2} \sum_{j=0}^B 1_{ij}^{noobj} (c_i - \hat{c}_i)^2 \\
 & + \sum_{i=0}^{K^2} 1_i^{obj} \sum_{k \in \{ac, bg\}} (s_k^i - \hat{s}_k^i)^2
 \end{aligned} \tag{1}$$

where $\hat{o}_t^i = (\hat{x}_i, \hat{y}_i, \hat{h}_i, \hat{w}_i, \hat{c}_i, \hat{s}_{ac}^i, \hat{s}_{bg}^i)$ denotes the cell i of the ground-truth $\hat{o}_t$. 1_i^{obj} indicates that the object appears in cell i , and 1_{ij}^{obj} indicates that the j^{th} bounding box predictor in cell i is responsible for the prediction (i.e., has a higher IoU with the ground truth between the B boxes). In contrast, 1_{ij}^{noobj} denotes that the j^{th} bounding box predictor in cell i is not responsible for the prediction or that there is no ground truth boxes in cell i .

The first two terms penalize coordinates error only when the prediction is responsible for the ground truth box. Since the deviation in the predicted coordinates more better for smaller boxes than large ones, we take the square root of width and height. The third and fourth terms penalize confidence score error, reduced by a factor λ_{noobj} when the prediction is not responsible for the ground truth box. Since most of the grid cells do not contain objects, it pushes confidence score towards zero. The final term penalizes classification as “action” or

“background” error only when there is an object in the cell. In this work, we empirically set $\lambda_{coord} = 5$ and $\lambda_{noobj} = 0.5$.

Recurrent *YoTube* is doubly deep in spatial-temporal domain, which can learn the temporal action dynamics. To further exploit the RGB and Flow cues in each single frame, we propose *static YoTube* which shares the same architecture as the recurrent *YoTube* with only one difference, *i.e.*, the static *YoTube* replace the last LSTM layer of the recurrent *YoTube* with a fully-connected layer with the same number of neurons. These two networks complement each other and thus combining their outputs could further improve the performance.

B. Path Linking and Trimming

Once the detectors (Sec. III-A) output a set of bounding boxes for each frame (denoted by $\mathbf{B} = \{\{b_i^{(j)}, j \in [1 \dots N_{b_i}]\}, i \in [1 \dots T]\}$), we compute the confidence score $s_c(b_i^{(j)})$, actionness score $s_{ac}(b_i^{(j)})$ and background score $s_{bg}(b_i^{(j)})$ for each box $b_i^{(j)}$, where T is the length of the video and N_{b_i} is the number of predicted boxes in frame i . After that, we could create a set of proposal paths $\mathbf{P} = \{p_i = \{b_{m_i}, b_{m_i+1} \dots b_{n_i}\}, i \in [1 \dots |\mathbf{P}|]\}$, where m_i and n_i are the starting and ending frame of path p_i , respectively. The details are as follows.

1) *Action Path Linking*: In order to link frame-level boxes into the coherent path, we firstly define a score for each path given its confidence scores s_c of each box and the IoU of successive boxes:

$$S(p) = \underbrace{\sum_{i=1}^T s_c(b_i)}_{\text{unary}} + \lambda_0 \times \underbrace{\sum_{i=2}^T \text{IoU}(b_i, b_{i-1})}_{\text{pairwise}} \quad (2)$$

$S(p)$ will be high for path if the corresponding detection box is assigned with a high confidence score and overlap. λ_0 is a trade-off factor to balance these two terms.

Maximizing Eqn. 2 helps find paths whose detection box scores are high and consecutive detection boxes significantly overlap in spatial and temporal domain.

To solve $\hat{p}_c = \underset{p_c}{\operatorname{argmax}} S(p_c)$, we employ the Viterbi algorithm [52]. Once the optimal path is calculated, we remove the bounding boxes in previous path from the frames to construct the next path until a certain frame does not contain any boxes.

2) *Action Paths Trimming*: The generated action paths described in the last subsection span the entire video since it greedily optimizes all confidence scores across the paths. On the other hand, human actions typically take up a fraction for untrimmed video. Therefore, It is necessary to perform trimming for removing those boxes that are unlikely belong to the action regions. Mathematically, each box b_t in a path p is assigned by a binary label $y_t \in \{0, 1\}$ (where “zero” and “one” represent the “background” and “action” class respectively). With such a scheme, the boxes are near to (or far from) the valid action regions which should be assigned to the “action” (or “background”) class as much as possible in the final path labeling $\hat{Y}_p = [\hat{y}_0, \hat{y}_1, \dots, \hat{y}_T]$.

Fig. 4. Illustration of the proposed peak trimming method on one UCF-101 videos: Blue and green curves represent the background and actionness scores, respectively. Blue and green crosses denote score peaks. Green patches represent the ground-truth paths and red patches represent paths that are extracted by using the proposed peak trimming method. It can be observed that the proposed method is capable of trimming the predicted patterns accurately thanks to certain action and background score transition patterns.

We also noticed that a transition in background scores typically reflects a change between action and non-action frames. In addition, the boxes within a valid action region often have high actionness scores. As a result, detecting peaks in actionness scores is helpful to find a potential action region, while finding peaks in background score could define the start and end of the action regions.

As illustrated in Fig.4, we propose a new method by looking at the transition pattern in the actionness and background scores for the path trimming. We first smooth the scores by adopting averaging approach to reduce the influence of noisy classifier scores, and then detect all the peaks in both scores. Note that, a peak is defined as a local maximum among at least n neighbors:

$$\begin{aligned} peaks_{ac} &= \{t, s_{ac}(b_t) = \max(V_n^{(ac)}(t))\} \\ peaks_{bg} &= \{t, s_{bg}(b_t) = \max(V_n^{(bg)}(t))\} \end{aligned} \quad (3)$$

where $V_n^{(k)}(t) = \{s_k(b_i), i \in [t - n \dots t + n]\}, k \in \{ac, bg\}$.

Once the peaks have been found, we can select all subsequences to generate the final action proposals by applying the following Algorithm 1.

Algorithm 1 Action paths trimming using actionness and background score peaks.

Input: actionness score peaks $peaks_{ac}$ and background score peaks $peaks_{bg}$.

Output: set $subseq$ as a set of trimmed paths.

```

subseq = ∅
for p ∈ peaksac do
    s = max(peaksbg < p)
    e = min(peaksbg > p)
    add path {bs ... be} into subseq
end for

```

IV. IMPLEMENTATION AND BENCHMARKING

In this section, we discuss the details of implementation and benchmarking, including the dataset and evaluation metrics.

A. Training

The used CNN architecture for feature extraction in *YouTube* is similar to that used in [37]. In details, it consists of 24 convolution layers and 2 fully-connected layers. We firstly replace the last two fully connected layers with a locally connected layer consisting of 256 filters with a 3×3 kernel. On the top of the locally connected layer, we train an LSTM layer with 588 neurons to directly regress the coordinates of the bounding boxes. We employ a locally connected layer to stabilize the training process and improve the convergence. The number of neurons in the last layer means dividing the image into 7×7 grids of which each predicts 2 bounding boxes.

For the RGB stream, the convolutional part of our model is pretrained on the ImageNet 1000-class dataset [66]. For the Flow stream, the convolutional part is pretrained with the weights of the RGB stream. The top layers are initialized using the method in [67]. We found no problem in convergence by initializing the weights of our Flow model with the weights of the RGB model despite of the notable difference between the images distribution.

We perform data augmentation to prevent over-fitting. This part is non-negligible due to important correlation among frames of the same video. Besides mirroring, we use corner cropping and center cropping. In details, we take a 224×224 crop from the 320×240 frame in each corner and the center. Then we resize this crop to the input size of 448×448 . Such an operation permits increasing the size of the dataset by a factor of 12.

We use the Adam [68] optimizer for training. When training the static *YouTube* detector, we use a batch size of 32 frames from different videos during 100 epochs with an initial learning rate of 10^{-4} . After 20 epochs, the learning rate decay by 10^{-5} . When training the recurrent *YouTube*, we freeze the weights of the convolutional layers to avoid a catastrophic forgetting. We set the batch size to 10 sequences and each sequence consists of 10 frames from different videos during 50 epochs. The same learning rate scheduling is used.

B. Datasets

UCF-101. The UCF-101 dataset is a large action recognition dataset which contains 101 action categories with more than 13,000 videos and each video contains about 180 frames. In our experiments, we use a subset which consists of 3,204 videos over 24 categories for the localization task. About 25% videos have been untrimmed, which permits to validate the efficiency of our methods of trimming videos. Each video contains one or more instances of the same action class. It has large variations in terms of appearance, scale, motion, etc with much diversity in terms of actions. Three default training/testing splits are provided with the dataset, and we perform experiments on the first split which consists of 2,290 training videos and 914 testing videos.

UCF-Sports. The dataset contains 150 sport broadcast videos with realistic actions captured under dynamic and cluttered environments. The dataset considers many actions with large displacement and intra-class variation. These videos have been trimmed to contain a single action instance without interruption. There are ten categories in the dataset, e.g. “diving”, “swinging bench”, “horse riding”, etc. We used the training-testing split suggested in [50], where the training and testing partition consist of 103 and 47 videos, respectively. The ground truth is provided as the sequences of bounding boxes enclose the actions.

JHMDB. This dataset consists of 928 videos for 21 different actions such as brush hair, swing baseball or jump. Video clips are trimmed to the duration of the action. Each clip contains between 15 and 40 frames. There are 3 training/testing splits and evaluation averages the results over the three splits.

C. Evaluation metrics

ABO, MABO: We use two popular metrics as in [50] to report the overall performance, namely Average Best Overlap (ABO) and Mean ABO (MABO). The overlap (OV) between a path $\mathbf{d} = \{d_s \dots d_e\}$ and a ground truth path $\mathbf{g} = \{g_s \dots g_e\}$ is defined as follows:

$$OV(\mathbf{d}, \mathbf{g}) = \frac{1}{|\mathbf{d} \cup \mathbf{g}|} \times \sum_{i \in \mathbf{d} \cap \mathbf{g}} \frac{d_i \cap g_i}{d_i \cup g_i}$$

$$|\mathbf{d} \cup \mathbf{g}| = \max(d_e, g_e) - \min(d_s, g_s)$$

$$\mathbf{d} \cap \mathbf{g} = [\max(d_s, g_s) \dots \min(d_e, g_e)]$$

where d_s and d_e are the detected bounding boxes in the starting and ending frame of a path, g_s and g_e are the bounding boxes in the starting and ending frame of the ground-truth path.

ABO measures the best localization from the set of action proposals $D = \{d_j | j = 1 \dots m\}$ for the ground-truth G , where $ABO(c)$ is the ABO computed for the ground-truth G_c of class c . The mean ABO (MABO) measures the average performance across all classes.

$$ABO = \frac{1}{|\mathbf{G}|} \sum_{\mathbf{g} \in \mathbf{G}} \max_{\mathbf{d} \in \mathbf{D}} OV(\mathbf{d}, \mathbf{g})$$

$$ABO(c) = \frac{1}{|\mathbf{G}^c|} \sum_{\mathbf{g} \in \mathbf{G}^c} \max_{\mathbf{d} \in \mathbf{D}} OV(\mathbf{d}, \mathbf{g})$$

$$MABO = \frac{1}{|\mathbf{C}|} \sum_{c \in \mathbf{C}} ABO(c)$$

where $\mathbf{C}$ is the set of action classes, $\mathbf{G}$ is the set of ground truth paths and $\mathbf{G}^c$ is the set of ground truth paths for action class c .

Recall vs. IoU: Another popular metric is the Recall vs. IoU [34], which measures the fraction of ground-truths detected in a set of overlap threshold. An instance of action, g_i is correctly detected by an action proposal d_j if the overlap score is higher than the threshold η i.e., $OV(d_j, g_i) \geq \eta$ and $\eta \in [0, 1]$. In our work, we aim to maximize the recall at the threshold of 0.5 like [50], [54].

Precision vs Recall: The precision-recall is a metric commonly used for action detection. In experiments, we also

evaluate the precision-recall curve for action proposal to reflect a detector's tradeoff between precision and recall. The recall is similarly defined as in Recall vs IoU. The Precision describes how many detected actions are matched with respect to the total number of detected tubes.

V. EXPERIMENTAL RESULTS

Our experiments could be divided into following parts: component analysis, generalization analysis, parameter sensitivity analysis, run-time analysis and comparison with other methods.

A. Component Analysis

1) RGB vs. Flow vs. Frame Difference: In recent action analysis studies, RGB and Flow have shown their complementary role in achieving state-of-the-art performance [17]. Their success could attribute to that RGB conveys rich object information and scene context and the Flow images capture salient object motions. Furthermore, in surveillance, frame difference has also been used as a kind of inputs to speed-up analysis thanks to its computation efficiency. Fig. 5 demonstrates the performance of our method with different inputs. One can see that on UCF-101 and JHMDB, our method using flow image achieves better performance than the case of RGB inputs. One underlying reason is that Flow image could eliminate the influence of the background clutter in RGB. Moreover, we found in UCF-Sports, using RGB images achieves a better result than using Flow images, which is resulted from the contents of dataset. In words, UCF-Sports contains many videos with salient actors, hence the information from RGB images is discriminative enough.

The performance of using Frame difference is relatively lower than that of using RGB and Flow, a possible explanation is that both UCF-Sports and UCF-101 contain background clutters/motion, simply calculating the difference between frames will result in noisy responses from background.

2) AlexNet vs GoogleNet vs C3D: Our method is compatible with recent popular Convolutional Neural Networks such as AlexNet [69], VGG [70], GoogleNet [71] and C3D [16] for feature extraction. In LRCN [29], the authors apply CaffeNet (a variant of AlexNet) for feature extraction. In our experiment, we choose GoogleNet [71] as our base feature extractor to balance speed and accuracy since it gives comparable performance with VGG [70], while using a smaller number of parameters. On the counter-part, the C3D structure [16] has shown good performance in action recognition and temporal action detection by learning 3D convolutional filter. Hence, we conduct experiments using AlexNet, GoogleNet and C3D for feature extraction.

The comparison result is shown in Fig.6. The performance of GoogleNet is about 3% and 6% higher than C3D and AlexNet respectively. There are two possible reasons for the performance gain: 1) the GoogleNet is deeper than the C3D and AlexNet, which can help extract more discriminative features; 2) our GoogleNet is trained with an input image of size 448×448 , while the input to C3D and AlexNet is only 160×160 and 224×224 respectively. The low-resolution input

UCF101	ABO	MABO	Recall	#Prop.
RGB Stream				
Static (RGB)	44.94	45.42	46.13	10
Recurrent (RGB)	45.85	45.83	47.05	21
YoTube (RGB)	47.60	47.78	50.61	35
Flow Stream				
Static (FLOW)	46.87	47.02	47.63	33
Recurrent (FLOW)	46.09	46.45	46.3	9
YoTube (FLOW)	48.61	48.83	51.29	42
Ensemble				
YoTube-RGB+FLOW (NO TRIM)	45.36	46.42	52.03	80
YoTube-RGB+FLOW (Trim with [25])	47.02	46.89	54.30	75
YoTube-RGB+FLOW	52.45	52.92	59.19	73

TABLE I
COMPONENT ANALYSIS ON THE UCF-101 DATASET.

UCF-Sports	ABO	MABO	Recall	#Prop.
RGB Stream				
Static (RGB)	71.64	72.54	97.87	20
Recurrent (RGB)	70.08	71.5	93.62	20
YoTube (RGB)	72.45	73.54	97.87	30
Flow Stream				
Static (FLOW)	68.08	68.98	93.62	20
Recurrent (FLOW)	66.22	66.91	91.49	20
YoTube (FLOW)	69.08	69.95	95.74	30
Ensemble				
YoTube-RGB+FLOW (NO TRIM)	74.44	75.31	97.87	30
YoTube-RGB+FLOW (Trim with [25])	74.44	75.31	97.87	30
YoTube-RGB+FLOW	74.44	75.31	97.87	30

TABLE II
COMPONENT ANALYSIS ON THE UCF-SPORTS DATASET.

to C3D and AlexNet is less desirable for detection related tasks. Actually, we would point out that our architecture could further enjoy the progress of deep neural networks.

3) Recurrent YoTube vs Static YoTube: The component analysis between recurrent and static *YoTube* for two streams in UCF-101, UCF-Sports and JHMDB are shown in Fig. 7, Table I, II and III. In UCF-101, the performance of recurrent version is slightly better than the static version in RGB stream with around 1% improvements in recall and the ensemble model (YoTube(RGB)) achieves about 3.56% improvement as shown in Table I. These results prove that the recurrent model and static model are complementary. We conjecture this since the RNN could capture the temporal dynamics among adjacent frames. For the results of Flow stream in Table I, the recall of static version is 1.3% better than the recurrent

JHMDB	ABO	MABO	Recall	#Prop.
RGB Stream				
Static (RGB)	68.36	67.55	92.34	20
Recurrent (RGB)	56.43	56.07	91.07	20
YoTube (RGB)	70.52	69.72	94.03	30
Flow Stream				
Static (FLOW)	70.46	70.08	96.81	20
Recurrent (FLOW)	71.76	71.36	96.17	20
YoTube (FLOW)	72.55	72.17	97.65	30
Ensemble				
YoTube-RGB+FLOW (NO TRIM)	74.72	74.21	99.31	30
YoTube-RGB+FLOW (Trim with [25])	73.41	72.65	97.08	35
YoTube-RGB+FLOW	74.72	74.21	99.31	30

TABLE III
COMPONENT ANALYSIS ON THE JHMDB DATASET.

Fig. 5. Comparison on using RGB, Flow, Frame Difference and their combinations as input. From left to right column, the results are on the UCF-101, UCF-Sports, and the JHMDB dataset.

Fig. 6. Comparison by using AlexNet, GoogleNet and C3D for feature extraction.

version and the ensemble model (YoTube(FLOW)) achieves another 4% improvement than the static version. The result further confirms the complementariness of two methods. The slightly inferior result of recurrent version is probably caused by lacking training data. In addition, the ensemble flow stream (YoTube(FLOW)) performs slightly better than the ensemble rgb stream (YoTube(RGB)). This is probably because the flow field eliminates the interference from the background. The proposed model (YoTube (RGB+FLOW)) combines two streams, which achieves 8% improvement than the flow stream in recall. The result demonstrates that the RGB and Flow streams also complement each other.

For UCF-Sports dataset, the static version outperforms the recurrent version by 4% (see Table II), whereas the ensemble model performs similar to the static version in terms of recall, but achieves higher ABO and MABO for better localization. For the flow stream, the static version gives a performance gain of 2.2% over the recurrent version in terms of recall, and the ensemble model is 2% better than the static version.

For JHMDB dataset, the static version is 12%, 11% and 1.3% better than the recurrent version in RGB stream in terms of the ABO, MABO and Recall metrics. The combination of

static and recurrent version yields a further 2% performance improvement. In Flow stream, the recurrent version is about 1.3% better than the static version in terms of ABO and MABO, their combination yields around 1.5% improvement. Further ensemble yields a recall of 99.31%. This confirms the complementariness of recurrent and static networks using RGB and Flow information.

4) Comparison between different path trimming methods: We also compare our path trimming method with the method proposed in our conference work [25], one can observe that the new method is 5% and 2% better than [25] on UCF-101 and JHMDB with less amount of tubes in (see Tables I–III). Our method and [25] have the identical result in UCF-Sports in Table II, because the classifier scores in UCF-Sports is strong enough to indicate the start and end of the video.

Moreover, we investigate the performance of the model without path trimming (NO TRIM) in Table I, II and III. As UCF-101 dataset contains un-trimmed video, the performance of our method without path trimming is nearly 7% lower in recall and also contains more noisy paths. This shows that the proposed path trimming is effective for the untrimmed video. For the UCF-Sports and JHMDB dataset which consists of trimmed videos, the method with and without path trimming achieve the same result. This proves that our method is adaptive to the video content.

5) Run Time Analysis Finally, we compared the running speed of our method with the APT [50] and Gkioxari [52]. Our method runs at 20 FPS, which is 15x faster than APT and 60x faster than Gkioxari. Our path linking and trimming method runs at 0.01s and is 20x faster than [25], which is quite computational efficient.

B. Generalization Analysis

We test the generalization ability of our method by pre-training the models on UCF-101 and transferring on JHMDB. Note that, there are little overlap between these two

Fig. 7. Comparison between two stream's static YoTube and recurrent YoTube on the UCF-101 (top-row), UCF-Sports (middle row) and JHMDB (bottom row); left column shows RGB stream, middle column shows the Flow stream and right column shows their ensemble.

Fig. 8. Generalization Analysis: the models trained on UCF-101 is tested on JHMDB to demonstrate generalization capability. The dashed lines represent the results using the models trained on UCF-101. The solid lines represent the results trained in JHMDB.

Fig. 9. Parameter sensitivity test: the magnitude of the contour map is the recall rate in UCF-101, the x-axis is the test value of λ_{obj} and y-axis is the test range of λ_{noobj} .

datasets. The corresponding result is shown in Fig.8. One can observe that our model and [52] pre-trained on UCF-101 experience nearly 15% performance degradation when the JHMDB dataset is used for evaluation. Further analysis reveals that the loss mainly results from the flow stream. One possible reason is that the appearance of flow images in UCF-101 is corrupted by noises, which is difficult to generalize to JHMDB as it consists of small motions.

C. Parameter Sensitivity Test

The hyper-parameters λ_{obj} and λ_{noobj} in Eqn.1 balance the role of objects and no-objects, whose values are empirically chosen. In experiment, we investigate the influence of one

of parameters by fixing the other parameters. The sensitivity testing result is shown in Fig.8. One can observe that our method has a relative stable performance when $\lambda_{coord} \in \{5, 9\}$ and $\lambda_{noobj} \in \{0.1, 0.5\}$. One explanation for such parameter selection is that as one image typically contains a few action regions, hence the bounding boxes from background should be with small weights.

D. Comparison to state-of-the-arts

We compare our method with state of the arts on UCF-101, UCF-Sports and JHMDB datasets. The recall-vs-IoU and recall-per-class curves for the testing datasets are shown in Fig. 10 and Fig. 11, respectively.

Fig. 10. Comparison with some state of the arts on UCF-101, UCF-Sports and JHMDB dataset in terms of recall with different IoU thresholds.

Fig. 11. Comparison with some state-of-the-art methods on UCF-101, UCF-Sports and JHMDB dataset. The performance is measured by the recall on each action class.

Fig. 12. Comparison with some state-of-the-arts methods on UCF-101, UCF-Sports and JHMDB dataset. The performance is measured by the precision vs. recall.

For the UCF-101 dataset, our method outperforms the state of the art [54] by at least 20% in all the range of IoU. Although Li *et al.* proposed using RPN [35], their model only involves one stream and the achieved performance is only 4% better than the unsupervised method (APT [50]) in terms of the recall as shown in Table IV. Notwithstanding, the recall of our single stream based method in RGB and Flow remarkably outperforms [54] by 11% and 12%, respectively. The result shows the superiority of the proposed *YoTube* which employs spatial-temporal modeling and two-stream design. [51] is a low-level features based human detector and the experimental result shows that it cannot effectively handle large dynamic changes in the scene. The work in [49] is also a low-level features based method, which is designed for non-overlap segmentation. These two characteristics make it sub-optimal for the task and achieving the lowest recall. According to the per-class recall curve in Fig. 11, our method is better than Li *et al.* [54] in many cases, especially when the data set contains large motion changes (e.g. “skijet”, “floor gymnastics” and

“long jump”).

For the UCF-Sports dataset, our method also outperforms Li *et al.* [54] by nearly 6% in terms of recall in Table V. Moreover, the deep learning based approaches (our method and [54]) also remarkably outperform the unsupervised methods, which shows the effectiveness of the deep networks based methods. The per-class recall curve for all methods is also provided in Fig. 11.

For the JHMDB dataset, Table VI shows that our method outperforms [52] and [50] by nearly 13% and 40% in terms of recall. For a comprehensive evaluation, we also report the Recall vs. Precision Curve in Fig.12. The unsupervised method [50] achieves a relative low precision because they generate many tubes by using low-level cues. Our method performs better than Gkioxari and Malik [52] since we simultaneously consider the appearance and temporal contexts. Actually, Precision and Recall may not be a good metric to evaluate action proposal, as proposals aims to have as high coverage as possible to cover all the objects, hence the recall

UCF101	ABO	MABO	Recall	#Prop.
Brox & Malik [49]	13.28	12.82	1.40	3
Yu <i>et al.</i> [51]	n.a	n.a	0.0	10,000
APT [50]	40.77	39.97	35.45	2299
Li <i>et al.</i> [54]	63.76	40.84	39.64	18
YoTube-RGB	47.60	47.78	50.61	35
YoTube-FLOW	48.61	48.83	51.29	42
YoTube-RGB+FLOW	52.45	52.92	59.19	73

TABLE IV

QUANTITATIVE COMPARISON ON THE UCF-101 DATASET. RECALL IS COMPUTED AT AN IOU THRESHOLD OF 0.5.

UCF Sports	ABO	MABO	Recall	#Prop.
Brox & Malik [49]	29.84	30.90	17.02	4
Jain <i>et al.</i> [47]	63.41	62.71	78.72	1642
Oneata <i>et al.</i> [48]	56.49	55.58	68.09	3000
Gkioxari <i>et al.</i> [52]	63.07	62.09	87.23	100
APT [50]	65.73	64.21	89.36	1449
Li <i>et al.</i> [54]	89.64	74.19	91.49	12
YoTube-RGB	72.45	73.54	97.87	30
YoTube-FLOW	69.08	69.95	95.74	30
YoTube-RGB+FLOW	74.44	75.31	97.87	30

TABLE V

QUANTITATIVE COMPARISON ON THE UCF-SPORTS DATASET. RECALL IS COMPUTED AT AN IOU THRESHOLD OF 0.5.

is more important than precision to reflect the action proposal's performance. Moreover, action proposal is still relatively new problem and lacking sufficient methods to benchmark.

We also evaluate our methods in terms of other metrics, *i.e.*, ABO, MABO, and number of proposals. The results are shown in Table IV, V and VI. From the results, our method produces the highest MABO. Note that, although [54] achieves the highest ABO, there are big differences between ABO and MABO. Actually, these two measurements should be in the similar scale according to the formula in Sec. IV-C, *i.e.* MABO is the mean of ABO for all classes.

Some visual examples are shown in Fig.13 and 14. The illustrations demonstrate that our method can produce candidate paths which produce good localization of the human actions.

E. Discussion

The proposed method has some failed cases as follows. First, it may loss frame-level localization for medium- and small-scale objects, especially in cluttered environments and the video with severe motion blurs (*e.g.* row 1 and 2 of Fig. 15). One possible reason is that our method performs inference on a dense fully connected layer which cannot capture the small scale changes due to its relatively coarse receptive field.

JHMDB	ABO	MABO	Recall	#Prop.
Gkioxari <i>et al.</i> [52]	63.07	62.09	86.34	125
APT [50]	54.16	53.37	59.55	2400
YoTube-RGB	70.52	69.72	94.03	30
YoTube-FLOW	72.55	72.17	97.65	30
YoTube-RGB+FLOW	74.72	74.21	99.31	30

TABLE VI

QUANTITATIVE COMPARISON ON THE JHMDB DATASET. RECALL IS COMPUTED AT AN IOU THRESHOLD OF 0.5.

Another failed case is that the localization of the generated action proposal might be interfered when the detection quality is unsatisfactory or objects have large overlap (see Fig. 15, row 2). The reason is that the path is generated by the greedy dynamic programming which accumulates errors during optimization. How to improve the robustness against these failed cases will be explored in our future work.

VI. CONCLUSION

We propose a novel framework for video action proposal. Given an untrimmed video as input, our method produces a small number of spatially compact and temporally smooth action proposals. The proposed approach enjoys the regression capability of RNN and representation learning capability of CNN, thus producing frame-level candidate action boxes jointly using RGB, Flow and temporal contexts among frames. The action proposals are constructed using dynamic programming with a novel path trimming method. Incorporating the long-term temporal context from LSTM helps reduce the ambiguities in each single frame. The proposed path trimming method can help trim the path for untrimmed videos. The superior results on UCF-101, UCF-Sports and JHMDB datasets highlight the effectiveness of our framework.

Moreover, we recently extend our method to fully convolutional approach [72] and achieved much improved performance. In the future, we plan to investigate how to use network compression technique to prune the redundant parameters so that the method could run on an on-board device such as drone. In addition, we will also explore how to use a deeper network such as ResNet [73] for higher accuracy.

REFERENCES

- [1] G. Yu, J. Yuan, and Z. Liu, "Action search by example using randomized visual vocabularies," *IEEE Trans. Image Processing*, vol. 22, no. 1, pp. 377–390, 2013.
- [2] K. R. Jeripothula, J. Cai, and J. Yuan, "CATS: co-saliency activated tracklet selection for video co-localization," in *ECCV*, 2016.
- [3] H. Fu, D. Xu, and S. Lin, "Object-based multiple foreground segmentation in RGBD video," *IEEE Trans. Image Processing*, vol. 26, no. 3, pp. 1418–1427, 2017.
- [4] H. Fu, D. Xu, B. Zhang, S. Lin, and R. K. Ward, "Object-based multiple foreground video co-segmentation via multi-state selection graph," *IEEE Trans. Image Processing*, vol. 24, no. 11, pp. 3415–3424, 2015.
- [5] F. Meng, H. Li, Q. Wu, B. Luo, and K. N. Ngan, "Weakly supervised part proposal segmentation from multiple images," *IEEE Trans. Image Processing*, vol. 26, no. 8, pp. 4019–4031, 2017.
- [6] X. Peng, C. Lu, Y. Zhang, and H. Tang, "Connections between nuclear norm and frobenius norm based representation," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 1, pp. 218–224, Jan. 2018.
- [7] X. Peng, H. Tang, L. Zhang, Z. Yi, and S. Xiao, "A unified framework for representation-based subspace clustering of out-of-sample and large-scale data," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 27, no. 12, pp. 2499–2512, 2016.
- [8] Y. Yang, C. Deng, S. Gao, W. Liu, D. Tao, and X. Gao, "Discriminative multi-instance multitask learning for 3d action recognition," *IEEE Transactions on Multimedia*, vol. 19, no. 3, pp. 519–529, 2017.
- [9] X. Lan, A. J. Ma, P. C. Yuen, and R. Chellappa, "Joint sparse representation and robust feature-level fusion for multi-cue visual tracking," *IEEE Transactions on Image Processing*, vol. 24, no. 12, pp. 5826–5841, Dec 2015.
- [10] X. Lan, S. Zhang, P. C. Yuen, and R. Chellappa, "Learning common and feature-specific patterns: A novel multiple-sparse-representation-based tracker," *IEEE Transactions on Image Processing*, 2017.

Fig. 13. Examples of our method on 4 videos from UCF101. Green boxes are ground truth and red boxes are from the best predicted path. Best viewed in colors.

Fig. 14. Examples of our method on 3 videos from UCF Sports. Green boxes are ground truth and red boxes are from the best predicted path. Best viewed in colors.

Fig. 15. Mistakes of our method on 2 videos. Green boxes are ground truth and red boxes are from the best predicted path. Best viewed in colors.

- [11] H. Zhu, F. Meng, J. Cai, and S. Lu, "Beyond pixels: A comprehensive survey from bottom-up to semantic image segmentation and cosegmentation," *J. Visual Communication and Image Representation*, vol. 34, pp. 12–27, 2016. [Online]. Available: <https://doi.org/10.1016/j.jvcir.2015.10.012>
- [12] H. Zhu, J. Lu, J. Cai, J. Zheng, S. Lu, and N. Magnenat-Thalmann, "Multiple human identification and cosegmentation: A human-oriented CRF approach with poselets," *IEEE Trans. Multimedia*, vol. 18, no. 8, pp. 1516–1530, 2016. [Online]. Available: <https://doi.org/10.1109/TMM.2016.2571629>
- [13] S. Zha, F. Luisier, W. Andrews, N. Srivastava, and R. Salakhutdinov, "Exploiting image-trained CNN architectures for unconstrained video classification," in *BMVC*, 2015.
- [14] Z. Xu, Y. Yang, and A. G. Hauptmann, "A discriminative CNN video representation for event detection," in *CVPR*, 2015.
- [15] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and F. Li, "Large-scale video classification with convolutional neural networks," in *CVPR*, 2014.
- [16] D. Tran, L. D. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3d convolutional networks," in *ICCV*, 2015.
- [17] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *NIPS*, 2014.
- [18] L. Wang, Y. Qiao, and X. Tang, "Action recognition with trajectory-pooled deep-convolutional descriptors," in *CVPR*, 2015.
- [19] J. Y. Ng, M. J. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici, "Beyond short snippets: Deep networks for video classification," in *CVPR*, 2015.
- [20] P. Weinzaepfel, Z. Harchaoui, and C. Schmid, "Learning to track for spatio-temporal action localization," in *ICCV*, 2015.
- [21] S. Saha, G. Singh, M. Sapienza, P. H. S. Torr, and F. Cuzzolin, "Deep learning for detecting multiple space-time action tubes in videos," in *BMVC*, 2016.
- [22] X. Peng and C. Schmid, "Multi-region two-stream R-CNN for action detection," in *ECCV*, 2016.
- [23] J. R. R. Uijlings, K. E. A. van de Sande, T. Gevers, and A. W. M.

- Smeulders, "Selective search for object recognition," *IJCV*, vol. 104, no. 2, pp. 154–171, 2013.
- [24] S. Manen, M. Guillaumin, and L. J. V. Gool, "Prime object proposals with randomized prim's algorithm," in *ICCV*, 2013.
- [25] R. Vial, H. Zhu, Y. Tian, and S. Lu, "Search video action proposal with recurrent and static yolo," *ICIP*, 2017.
- [26] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [27] F. Wang and D. M. J. Tax, "Survey on the attention based RNN model and its applications in computer vision," *CoRR*, vol. abs/1601.06823, 2016. [Online]. Available: <http://arxiv.org/abs/1601.06823>
- [28] V. Veeriah, N. Zhuang, and G. Qi, "Differential recurrent neural networks for action recognition," in *ICCV*, 2015.
- [29] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, T. Darrell, and K. Saenko, "Long-term recurrent convolutional networks for visual recognition and description," in *CVPR*, 2015.
- [30] Z. Wu, X. Wang, Y. Jiang, H. Ye, and X. Xue, "Modeling spatial-temporal clues in a hybrid deep learning framework for video classification," in *ACM MM*, 2015.
- [31] Z. Wu, Y. Jiang, X. Wang, H. Ye, and X. Xue, "Multi-stream multi-class fusion of deep networks for video classification," in *ACM MM*, 2016.
- [32] B. Ni, X. Yang, and S. Gao, "Progressively parsing interactional objects for fine grained action detection," in *CVPR*, 2016.
- [33] R. Stewart, M. Andriluka, and A. Y. Ng, "End-to-end people detection in crowded scenes," in *CVPR*, 2016.
- [34] C. L. Zitnick and P. Dollár, "Edge boxes: Locating object proposals from edges," in *ECCV*, 2014.
- [35] S. Ren, K. He, R. B. Girshick, and J. Sun, "Faster R-CNN: towards real-time object detection with region proposal networks," in *NIPS*, 2015.
- [36] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. E. Reed, C. Fu, and A. C. Berg, "SSD: single shot multibox detector," in *ECCV*, 2016.
- [37] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *CVPR*, 2016.
- [38] S. Herath, M. T. Harandi, and F. Porikli, "Going deeper into action recognition: A survey," *CoRR*, vol. abs/1605.04988, 2016. [Online]. Available: <http://arxiv.org/abs/1605.04988>
- [39] Z. Wu, T. Yao, Y. Fu, and Y. Jiang, "Deep learning for video classification and captioning," *CoRR*, vol. abs/1609.06782, 2016. [Online]. Available: <http://arxiv.org/abs/1609.06782>
- [40] S. Ji, W. Xu, M. Yang, and K. Yu, "3d convolutional neural networks for human action recognition," in *ICML*, 2010.
- [41] H. Xu, A. Das, and K. Saenko, "R-C3D: region convolutional 3d network for temporal activity detection," *CoRR*, vol. abs/1703.07814, 2017. [Online]. Available: <http://arxiv.org/abs/1703.07814>
- [42] Z. Shou, D. Wang, and S. Chang, "Temporal action localization in untrimmed videos via multi-stage cnns," in *CVPR*, 2016.
- [43] Y. Ke, R. Sukthankar, and M. Hebert, "Efficient visual event detection using volumetric features," in *ICCV*, 2005.
- [44] T. Lan, Y. Wang, and G. Mori, "Discriminative figure-centric models for joint action localization and recognition," in *ICCV*, 2011.
- [45] Y. Tian, R. Sukthankar, and M. Shah, "Spatiotemporal deformable part models for action detection," in *CVPR*, 2013.
- [46] D. Tran and J. Yuan, "Max-margin structured output regression for spatio-temporal action localization," in *NIPS*, 2012.
- [47] M. Jain, J. C. van Gemert, H. Jégou, P. Bouthemy, and C. G. M. Snoek, "Action localization with tubelets from motion," in *CVPR*, 2014.
- [48] D. Oneata, J. Revaud, J. J. Verbeek, and C. Schmid, "Spatio-temporal object detection proposals," in *ECCV*, 2014.
- [49] T. Brox and J. Malik, "Object segmentation by long term analysis of point trajectories," in *ECCV*, 2010.
- [50] J. C. van Gemert, M. Jain, E. Gati, and C. G. M. Snoek, "APT: action localization proposals from dense trajectories," in *BMVC*, 2015.
- [51] G. Yu and J. Yuan, "Fast action proposals for human action detection and search," in *CVPR*, 2015.
- [52] G. Gkioxari and J. Malik, "Finding action tubes," in *CVPR*, 2015.
- [53] R. B. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *CVPR*, 2014.
- [54] N. Li, D. Xu, Z. Ying, and G. L. Zhihao Li, "Search action proposals via spatial actionness estimation and temporal path inference and tracking," in *ACCV*, 2016.
- [55] T. Baltrusaitis, C. Ahuja, and L. Morency, "Multimodal machine learning: A survey and taxonomy," *CoRR*, vol. abs/1705.09406, 2017. [Online]. Available: <http://arxiv.org/abs/1705.09406>
- [56] Y. Li, M. Yang, and Z. Zhang, "Multi-view representation learning: A survey from shallow methods to deep methods," *CoRR*, vol. abs/1610.01206, 2016. [Online]. Available: <http://arxiv.org/abs/1610.01206>
- [57] J. Wagner, V. Fischer, M. Herman, and S. Behnke, "Multispectral pedestrian detection using deep fusion convolutional neural networks," *ESANN*, 2016.
- [58] L. Qu, S. He, J. Zhang, J. Tian, Y. Tang, and Q. Yang, "Rgb-d salient object detection via deep fusion," *IEEE Transactions on Image Processing*, vol. 26, no. 5, pp. 2274–2285, May 2017.
- [59] A. Wang, J. Lu, J. Cai, G. Wang, and T. Cham, "Unsupervised joint feature learning and encoding for RGB-D scene labeling," *IEEE Trans. Image Processing*, vol. 24, no. 11, pp. 4459–4473, 2015. [Online]. Available: <https://doi.org/10.1109/TIP.2015.2465133>
- [60] A. Wang, J. Lu, J. Cai, T. Cham, and G. Wang, "Large-margin multi-modal deep learning for RGB-D object recognition," *IEEE Trans. Multimedia*, vol. 17, no. 11, pp. 1887–1898, 2015. [Online]. Available: <https://doi.org/10.1109/TMM.2015.2476655>
- [61] H. Zhu, J. Weibel, and S. Lu, "Discriminative multi-modal feature fusion for RGBD indoor scene recognition," in *CVPR*, 2016.
- [62] C. Feichtenhofer, A. Pinz, and R. P. Wildes, "Spatiotemporal residual networks for video action recognition," in *NIPS*, 2016.
- [63] X. Zhang, H. Zhang, Y. Zhang, Y. Yang, M. Wang, H. Luan, J. Li, and T. Chua, "Deep fusion of multiple semantic cues for complex event recognition," *IEEE Trans. Image Processing*, vol. 25, no. 3, pp. 1033–1046, 2016. [Online]. Available: <https://doi.org/10.1109/TIP.2015.2511585>
- [64] X. Zhang, K. Gao, Y. Zhang, D. Zhang, J. Li, and Q. Tian, "Task-driven dynamic fusion: Reducing ambiguity in video description," in *CVPR*, 2017.
- [65] N. D. Doulamis and A. D. Doulamis, "Fast and adaptive deep fusion learning for detecting visual objects," in *ECCV*, 2012.
- [66] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.
- [67] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *ICCV*, 2015.
- [68] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [69] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [70] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014. [Online]. Available: <http://arxiv.org/abs/1409.1556>
- [71] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. E. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *CVPR*, 2015.
- [72] H. Zhu, R. Vial, and S. Lu, "TORNADO: A spatio-temporal convolutional regression network for video action proposal," in *ICCV*, 2017.
- [73] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016.