

Tripping through time: Efficient Localization of Activities in Videos

Meera Hahn¹
meerahahn@gatech.edu

Asim Kadav²
asim@nec-labs.com

James M. Rehg¹
rehg@gatech.edu

Hans Peter Graf²
hpg@nec-labs.com

¹ College of Computing
Georgia Institute of Technology
Atlanta, GA

² Machine Learning Department
NEC Labs America
Princeton, NJ

Abstract

Localizing moments in untrimmed videos via language queries is a new and interesting task that requires the ability to accurately ground language into video. Previous works have approached this task by processing the entire video, often more than once, to localize relevant activities. In the real world applications of this approach, such as video surveillance, efficiency is a key system requirement. In this paper, we present TripNet, an end-to-end system that uses a gated attention architecture to model fine-grained textual and visual representations in order to align text and video content. Furthermore, TripNet uses reinforcement learning to efficiently localize relevant activity clips in long videos, by learning how to intelligently skip around the video. It extracts visual features for few frames to perform activity classification. In our evaluation over Charades-STA [14], ActivityNet Captions [15] and the TACoS dataset [16], we find that TripNet achieves high accuracy and saves processing time by only looking at 32-41% of the entire video.

1 Introduction

The increasing availability of videos and their importance in application domains, such as social media and surveillance, has created a pressing need for automated video analysis methods. A particular challenge arises in long video clips which have noisy or nonexistent labels, which are common in many settings including surveillance and instructional videos. While classical video retrieval works at the level of entire clips, a more challenging and important task is to efficiently sift through large amounts of unorganized video content and retrieve specific moments of interest.

A promising formulation of this task is the paradigm of temporal activity localization via language query (TALL), which was introduced by Gao et al. [14] and Hendricks et al. [2]. The TALL task is illustrated in Fig. 1 and shows a few frames from a climbing video along with a language query describing the event of interest “Climber adjusts his feet for the first time.” In the figure the green frame denotes the desired output, the first frame in which the feet are being adjusted. Note that the solution to this problem requires local and global

Input Query: "Climber adjusts his feet for the first time"




Figure 1: Figure demonstrating video search method where given an input query, a human would normally forward the video until relevant objects or objects arrive and then proceed slowly to obtain the appropriate temporal boundaries.

temporal information: local cues are needed to identify the specific frame in which the feet are adjusted, but global cues are also needed to identify the first time this occurs, since the climber adjusts his feet throughout the clip. The MCN [22] and CTRL [24] architectures, along with more recent work by Yuan et al. [61], have obtained encouraging results for the TALL task, however, they all suffer from a significant limitation which we address in this paper: Existing methods construct temporally-dense representations of the entire video clip which are then analyzed to identify the target events. This can be very inefficient in long video clips, where the target might be a single short moment. This approach starkly contrasts to how humans search for events of interest, as illustrated schematically in Fig. 1. A human would fast-forward through the video from the beginning, effectively sampling sparse sets of frames until they got closer to the region of interest. Then, they would look frame-by-frame until the start and end points are localized. At that point the search would terminate, leaving the vast majority of frames unexamined. Note that efficient solutions must go well beyond simple heuristics, since events of interest can occur anywhere within a target clip and the global position may not be obvious from the query.

In order to develop an efficient solution to the TALL task it is necessary to address two main challenges: 1) Devising an effective joint representation (embedding) of video and language features to support localization; and 2) Learning an efficient search strategy that can mimic the human ability to sample a video clip intelligently to localize an event of interest. Most prior works on TALL address the issue of constructing a joint embedding space through a sliding window approach that pools at different temporal scales. This strategy has achieved some success, but it is not well suited for efficient search because it provides only a coarse control over which frames are evaluated. In contrast, our usage of gated-attention architecture is effective in aligning the text queries, which often consist of an object and its attributes or an action, with the video features that consist of convolutional filters that can identify these elements. Two prior works [48, 51] also utilize an attention model, but they do not address its use in efficient search.

We address the challenge of efficient search through a combination of reinforcement learning (RL) and fine-grained video analysis. Our approach is inspired in part by strategies used by human video annotators. In addition, we make an analogy to the success of RL in language-guided navigation tasks in 3D environments [2, 2, 11, 13]. Specifically, we make an analogy between temporally localizing events in a video through playback controls and the actions an agent would take to navigate around an environment looking for a specific object of interest. We share with the navigation task, the fact that labeled data is not available for explicitly modeling the relationship between actions and rewards, but it is possible to learn a model through simulation. Our approach to temporal localization uses a novel architecture for combining the multi-modal video and text features with a policy learning module that learns to step forward and rewind the video and receives awards for accurate temporal

localization.

In summary this paper makes two contributions: First, we present a novel end-to-end reinforcement learning framework called TripNet that addresses the problem of temporal activity localization via a language query. TripNet uses gated-attention to align text and visual features, leading to improved accuracy. Second, we present experimental results on the datasets Charades-STA [14], ActivityNet Captions [26] and TACoS [56]. These results demonstrate that TripNet achieves state of the art results in accuracy while significantly increasing efficiency, by evaluating only 32-41% of the total video.

2 Related Work

Querying using Natural Language This paper is most-closely related to prior works on the TALL problem, beginning with the two works [14, 27] that introduced it. These works additionally introduced a dataset for the task, Charades-STA[14] and DiDeMo[27]. Each dataset contains untrimmed videos with multiple sentence queries and the corresponding start and end timestamp of the clip within the video. The two papers adopt a supervised cross modal embedding approach, in which sentences and videos are projected into a embedding space, optimized so that queries and their corresponding video clips will lie close together, while non-corresponding clips will be far apart. At test time both approaches run a sliding window across the video and compute an alignment score between the candidate window and the language query. The window with the highest score is then selected as the localized query. Follow-up works on the TALL task have differed in the design of the embedding process [8]. [28, 47, 51] modified the original approach by adding self-attention and co-attention to the embedding process. [48] introduces early fusion of the text queries and video features rather than using an attention mechanism. Additionally, [48] uses the text to produce activity segment proposals as their candidate windows instead of using a fixed sliding window approach. Pre-processing of the sentence queries has not been a primary focus of previous works, with most methods using a simple LSTM based architecture for sentence embedding, with Glove [63] or Skip-Thought [25] vectors to represent the sentences. Other work in video-based localization which predated TALL either uses a limited set of text vocabulary to search for specific events or uses structured videos [9, 19, 53, 43]. Other earlier works address the task of retrieving objects using natural language questions, which is less complex than than localizing actions in videos [23, 24, 30].

Temporal Localization Temporal action localization refers to localizing activities over a known set of action labels in untrimmed videos. Some existing work in this area has found success by extracting CNN features from the video frames, pooling the features and feeding them into either single or multi-stage classifiers to obtain action predictions, along with temporal labels [6, 11, 47, 50]. These methods use a sliding window approach. Often, multiple window sizes are used and candidate windows are densely sampled, meaning that the candidates overlap with each other. This method has high accuracy, but it is an exhaustive search method that leads to high computational costs. There are also TALL methods that forgo the end-to-end approach and instead use a two-stage approach: first generating temporal proposals, and second, classifying the proposals into action categories [6, 15, 16, 17, 21, 55, 59, 40]. Unlike both sets of previous methods, our proposed model TripNet, uses reinforcement learning to perform temporal localization using natural language queries. In action-recognition methods [8, 13, 27, 29, 31, 41, 42], Frame glimpse [49] and Action Search [0] use reinforce-

ment learning to identify an action while looking at the smallest possible number of frames. These works focus only on classifying actions (as opposed to parsing natural language activities) in trimmed videos and the former does not perform any type of temporal boundary localization. A recent work [20] also uses RL to learn the temporal video boundaries. However, it analyzes the entire video and does not focus on the efficiency of localizing the action. Another work [46], uses RL to estimate the activity location using fixed temporal boundaries. Furthermore, it uses Faster R-CNN trained on the Visual Genome dataset to provide semantic concepts to the model.

3D navigation and Reinforcement Learning Locating a specific activity within a video using a reinforcement learning agent is analogous to the task of navigating through a three dimensional world. How can we efficiently learn to navigate through the temporal landscape of the video? The parallels between our approach and navigation lead us to additionally review the related work in the areas of language-based navigation and embodied perception [4, 18]. Recently, the task of Embodied Question Answering [10] was introduced. In this task, an agent is asked questions about the environment, such as “what color is the bathtub,” and the agent must navigate to the bathtub and answer the question. The method introduced in [10] focuses on grounding the language of the question not directly into the pixels of the scene but into actions for navigating around the scene. We seek to ground our text query into actions for skipping around the video to narrow down on the correct clip. Another recent work [9] explores giving visually grounded instructions such as “go to the green red torch” to an agent and having them learn to navigate their environment to find the object and complete the instruction.

3 Methods

We now describe TripNet, an end to end reinforcement learning method for localizing and retrieving temporal activities in videos given a natural language query as shown in Figure 2.

Problem Formulation: The localization problem that we solve is defined as follows: Given an untrimmed video V and a language query L , our goal is to temporally localize the specific clip W in V which is described by L . In other words, let us denote the untrimmed video as $V = \{f_n\}_{n=1}^N$ where N is the number of frames in the video, we want to find the clip $W = \{f_n, \dots, f_{n+k}\}$ that corresponds best to L . It is possible to solve this problem efficiently because videos have an inherent temporal structure, such that an observation made at frame n conveys information about frames both in the past and in the future. Some challenges of the problem are, how to encode the uncertainty in the location of the target event in a video, and how to update the uncertainty from successive observations. While a Bayesian formulation could be employed, the measurement and update model would need to be learned and supervision for this is not available.

Since it is computationally feasible to simulate the search process (in fact it is only a one-dimensional space, in contrast to standard navigation tasks) we adopt an reinforcement learning (RL) approach. We are motivated by human annotators who observe a short clip and make a decision to skip forward or backward in the video by some number of frames, until they can narrow down to the target clip. We emulate this sequential decision process using RL. Using RL we train an agent that can steer a fixed sized window around the video to find W without looking at all frames of V . We employ the actor-critic method A3C [32] to learn the policy π that maps $(V, L) \rightarrow W$. The intuition is that the agent will take large

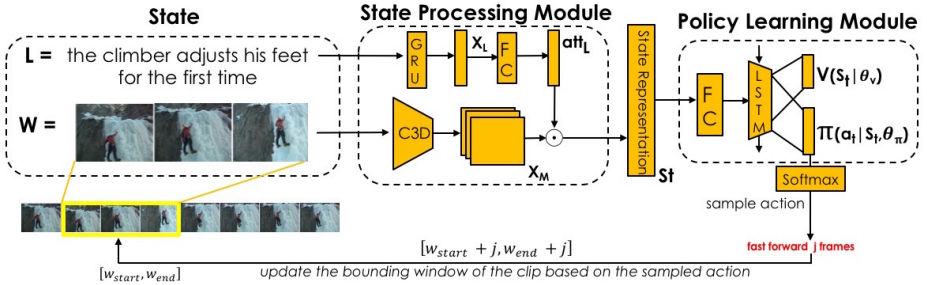


Figure 2: TripNet localizes specific moments in videos based on a natural language query. There are two main components of TripNet: the state processing module and the policy learning module. The state processing module encodes the state into a joint visual and linguistic representation which is then fed to the policy learning module on which it generates the action policy. The sampled action in this figure is skipping forward by j frames and the state is updated accordingly.

jumps around the video until it finds visual features that identify proximity to L , and then it will start to take smaller steps as it narrows in on the target clip.

3.1 State and Action Space

At each time step, the agent observes the current state, which consists of the sentence L and a candidate clip of the video. The clip is defined by a bounding window $[W_{start}, W_{end}]$ where start and end are frame numbers. At time step $t = 0$, the bounding window is set to $[0, X]$, where X is the average length of annotated clips within the dataset.¹ This window size is fixed and does not change. At each time step the state processing module creates a state representation vector which is fed into the policy learning module, on which it generates an action policy. This policy is a distribution over all the possible actions. An action is then sampled according to the policy.

Our action space consists of 7 predefined actions: move the entire bounding window W_{start}, W_{end} forward or backward by h frames, j frames, or 1 second of frames or TERMINATE. Where $h = N/10$ and $j = N/5$. These navigation steps makes aligning the visual and text features easier. Making the RL agent explicitly learn the window size significantly increases the state space and leads to a drop in accuracy and efficiency with our current framework. However, an additional function to adjust the box width over the fixed window size could be learned, analogous to the refinement of the anchor boxes used in object detectors. These actions were chosen so that amount of movement is proportional to the length of the video. If the bounding window is at the start or end of the full video and an action is chosen that would push the bounding window outside the video’s length, the bounding window remains the same as the previous time step. The action TERMINATE ends the search and returns the clip of the current state as the video clip predicted to be best matched to L .

3.2 TripNet architecture

We now describe the architecture of TripNet, which is illustrated in Figure 2.

¹Note that this means that the search process always starts at the beginning of the video. We also explored starting at the middle and end of the clip and found empirically that starting at the beginning performed the best.

State Processing Module At each time step, the state-processing module takes the current state as an input and outputs a joint-representation of the input video clip and the sentence query L . The joint representation is used by the policy learner to create an action policy from which the optimal action to take is sampled. The clip is fed into C3D [45] to extract the spatio-temporal features from the fifth convolutional layer. We mean-pool the C3D features across frames and denote the result as x_M .

To encode the sentence query, we first pass L through a Gated Recurrent Unit (GRU) [9] which outputs a vector x_L . We then transform the query embedding into an attention vector that we can apply to the video embedding. To do so, the sentence query embedding x_L is passed through a fully connected linear layer with a sigmoid activation function. The output of this layer is expanded to be the same dimension of x_M . We call the output of the linear layer the attention vector, att_L . We perform a Hadamard multiplication between att_L and x_M and the result is output as the state representation s_t . This attention unit is our gated attention architecture for activity localization. Hence, the gated attention unit is designed to gate specific filters based on the attention vector from the language query [17]. The gating mechanism allows the model to focus on specific filters that can attend to specific objects, and their attributes, from the query description.

In our experiments, we implement an additional baseline using a simple concatenation operation between the video and text representations to demonstrate the effectiveness of our gated-attention architecture. Here, we only perform self-attention over the mean pooled C3D features. Then, we take a Skip-Thought [25] encoding of the sentence query and concatenate it with the features of the video frames to produce the state representation. In the experiments, we denote this method as TripNet-Concat.

Policy Learning Module We use an actor-critic method to model the sequential decision process of grounding the language query to a temporal video location. The module employs a deep neural network to learn the policy and value functions. The network consists of a fully connected linear layer followed by an LSTM, which is followed by a fully connected layer to output the value function $v(s_t|\theta_v)$ and fully connected layer to output the policy $\pi(a_t|s_t, \theta_\pi)$, for state, s_t and action, a_t at time t , θ_v is the critic branch parameters and θ_π is actor branch parameters. The policy, π , is a probabilistic distribution over all possible actions given the current state. Since we are trying to model a sequential problem we use an LSTM so that the system can have memory of the previous states which will inevitably positively impact the future actions. Specifically, we use the asynchronous actor-critic method known as A3C [22] with Generalized Advantage Estimation [37] that reduces policy gradient variance. The method runs multiple parallel threads that each run their own episodes and update global network parameters at the end of the episode.

Since the goal is to learn a policy that returns the best matching clip, we want to reward actions that bring the bounding windows $[W_{start}, W_{end}]$ closer to the bounds of the ground truth clip. Hence, the action to take, should return a state that has a clip with more overlap with the ground-truth than the previous state. Therefore, we use reward shaping by having our reward be the difference of potentials between the previous state and current state. However, we want to ensure the agent is taking an efficient number of jumps and not excessively sampling the clip. In order to encourage this behavior, we give a small negative reward in proportion with the total number of steps thus far. As a result, the agent is encouraged to find the clip window as quickly as possible. We experiment to find the optimal negative reward factor β . We found using a negative reward factor results in the agent taking more actions

with larger frame jump. Hence, our reward at any time step t is calculated as follows:

$$\text{reward}_t = (\text{IOU}_t(s_t) - \text{IOU}_t(s_{t-1})) - \beta * t \quad (1)$$

where we set β to .01. We calculate the IOU between the clip of the state at time t , $[W_{start}^t, W_{end}^t]$, and the ground truth clip for sentence L , $[G_{start}, G_{end}]$ as follows:

$$\text{IOU}_t = \frac{\min(W_{end}^t, G_{end}) - \max(W_{start}^t, G_{start})}{\max(W_{end}^t, G_{end}) - \min(W_{start}^t, G_{start})} \quad (2)$$

We use the common loss functions for A3C for the value and policy loss. For training the value function, we set the value loss to the mean squared loss between the discounted reward sum and the estimated value:

$$\text{Loss}_{\text{value}} = \sum_t (R_t - v(s_t | \theta_v))^2 * \gamma_1, \quad (3)$$

where we set γ_1 to .5 and R_t is the accumulated reward. For training the policy function, we use the policy gradient loss:

$$\begin{aligned} \text{Loss}_{\text{policy}} = & - \sum_t \log(\pi(a_t | s_t, \theta_\pi)) * \text{GAE}(s_t) \\ & - \gamma_0 * H(\pi(a_t | s_t, \theta_\pi)), \end{aligned} \quad (4)$$

where GAE is the generalized advantage estimation function, H is the calculation of entropy and γ_0 is set to 0.5. Therefore, the total loss for our policy learning module is:

$$\text{Loss} = \text{Loss}_\pi + \gamma_1 + \text{Loss}_v. \quad (5)$$

4 Evaluation

We evaluate the TripNet architecture on three video datasets, Charades-STA [4], ActivityNet Captions [7] and TACoS [3]. Charades-STA was created for the moment retrieval task and the other datasets were created for video captioning, but are often used to evaluate the moment retrieval task. Note that we chose not to include the DiDeMo [2] dataset because the evaluation is based on splitting the video into 21 pre-defined segments, instead of utilizing continuously-variable start and end times. This would require a change in the set of actions for our agent. We do, however, compare our approach against the method from [2] on other datasets. Please refer to supplementary materials for the details on the datasets as well as the implementation details of our system.

4.1 Experiments

Evaluation Metric. We use Intersection over Union (IoU) at different alpha thresholds to measure the difference between the ground truth clip and the clip that TripNet temporally aligns to the query. If a predicted window and the ground truth window have an IoU that is above the set alpha threshold, we classify the prediction as correct, otherwise we classify it as incorrect. See Equation 2 for how the IoU is calculated. The results over the different datasets and methods are shown in Table 1. Most previous work uses a R@k-IoU, which is a IoU score for the top k returned clips. These works used a sliding window approach allowing their alignment system to return k-top candidate windows based on confidence scores. In contrast, TripNet searches the video until it finds the best aligned clip and returns only that clip. As a result, all reported IoU scores are measured at R@1.

Method	ActivityNet			TACoS			Charades		
	$\alpha@.3$	$\alpha@.5$	$\alpha@.7$	$\alpha@.3$	$\alpha@.5$	$\alpha@.7$	$\alpha@.3$	$\alpha@.5$	$\alpha@.7$
MCN [22]	21.37	9.58	-	1.64	1.25	-	-	-	-
CTRL [12]	28.70	14.00	-	18.32	13.3	-	-	23.63	8.89
TGN [8]	45.51	28.47	-	21.77	18.9	-	-	-	-
ABLR [51]	55.67	36.79	-	-	-	-	-	-	-
ACRN [28]	31.29	16.17	-	19.52	14.62	-	-	-	-
MLVI [48]	45.30	27.70	13.60	-	-	-	54.70	35.60	15.80
VAL [42]	-	-	-	19.76	14.74	-	-	23.12	9.16
SM-RL [46]	-	-	-	20.25	15.95	-	24.36	11.17	-
TripNet-Concat	36.75	25.64	10.25	18.24	14.16	6.47	41.84	27.23	12.62
TripNet-GA	48.42	32.19	13.93	23.95	19.17	9.52	54.64	38.29	16.07

Table 1: The accuracy of each method on ActivityNet, TACoS and Charades-STA measured by IoU at multiple α values.

Dataset	% of Frames Used	Avg. # Actions	Bounding Window Size
ActivityNet	41.65	5.56	35.7s
Charades	33.11	4.16	8.3s
TACoS	32.7	9.84	8.0s

Table 2: The efficiency of the TripNet architecture on different datasets. Number of actions includes the TERMINATE action. Size of the bounding window is in seconds. The methods were tested on the same Titan Xp GPU.

4.1.1 Comparison and Baseline Methods

We compare against other methods both from prior work and a baseline version of the TripNet architecture. The prior work we compare against is as follows: MCN [22], CTRL [12], TGN [8], ABLR [51], ACRN [28], MLVI [48], VAL [42] and SM-RL [46]. Besides SM-RL, all prior works tackle the task by learning to jointly represent the ground truth moment clip and the moment description query. To generate candidate windows during testing, these methods go over the whole video using a sliding window and then, choose the candidate window that best corresponds to the query encoding. This methodology relies on seeing all frames of the video at least once, if not more, during test time.

We provide accuracy results in Table 1 for 3 datasets, compiling the performance numbers reported in the prior works for comparison to the performance of TripNet. We can see that, in terms of accuracy, TripNet outperforms all other methods on the Charades-STA and TACoS datasets, and that it performs comparably to the state of the art on ActivityNet Captions. Using a state processing module for TripNet that does not use attention (TripNet-Concat) performs consistently worse than the state processing module for TripNet that uses the gated attention architecture (TripNet-GA), since multi-modal fusion between vision and language shows improvement with different attention mechanisms.

More specifically, in Table 1, ABLR processes all frames and encodes the intermediate representations of each frame and word using Bi-LSTM, followed by two levels of attention over it. It then regresses the coordinates over this comprehensive intermediate representa-

Method	Charades	ActivityNet	TACoS
CTRL [14]	44.19ms	218.95ms	342.12ms
MCN [22]	78.16ms	499.32ms	674.01ms
TGN [8]	18.2ms	90.2ms	144.7ms
TripNet-GA	5.13ms	6.23ms	11.27ms

Table 3: The average time in milliseconds that it takes to localize a moment on different datasets.

tion. Hence, this regression-based approach over all frames is beneficial for the TALL task. However, it is computationally inefficient to perform this processing over all frames. On the other hand, in our method feature extraction is only performed if requested by the RL step and is suitable for long surveillance videos where ABLR may not be practical. We also see an overall drop in performance of TripNet on TACoS, despite outperforming other recent work. We don’t believe this is due to the length of the videos but instead due to the nature of the content in the video. TACoS contains long cooking videos all taken in the same kitchen making the dataset more difficult. While ActivityNet videos are up to four times longer than Charades-STA videos, TripNet is nonetheless able to maintain high accuracy on ActivityNet, illustrating TripNet’s ability to scale to videos of different lengths.

In the analysis of our results, we found that one source of IoU inaccuracy was the size of the bounding window. TripNet currently uses a fixed size bounding window that the agent moves around the video until it returns a predicted clip. The size of the fixed window is equal to the mean length of the ground truth annotated clips. A possible direction for future work would be to add actions that expand and contract the size of the bounding window.

4.1.2 Efficiency

We now describe how the different competing methods construct their candidate windows, in order to provide an understanding of the computational cost of each approach. MCN [22]: creates 21 candidate windows across the video. CTRL [14], VAL [22], ACRN [28]: use a sliding windows of multiple sizes to generate candidate windows. TGN [8]: does not use a sliding window but generates candidate windows across the video at different scales. ABLR[50]: encodes the entire video and then uses attention via the language query to localize the clip. The method still requires two passes over the entire video. MLVI [48]: trains a separate network to generate candidate window proposals. **TripNet**: TripNet is the only method that does not need to watch the entire video to temporally localize a described moment. Instead, our trained agent efficiently moves a candidate window around the video until it localizes the described clip. Measurements of efficiency are described in Table 2. In order to get a better understanding of efficiency, we run the CTRL [14], TGN [8], MCN [22], and TripNet-GA over the dataset test splits and compute the average time it takes to localize a moment in the video. This is shown in Table 3.

4.1.3 Qualitative Results

In Figure 3, we show the qualitative results of TripNet-GA on the Charades-STA dataset. In the figure, we show the sequential list of actions the agent takes in order to temporally localize a moment in the video. The green boxes represent the bounding window of the state at time t and the yellow box represents the ground truth bounding window. The first video is

33 seconds long (792 frames) and the second video is 20 seconds long (480 frames). In both examples the agent skips both backwards and forwards. In the first example, TripNet sees 408 of the frames of the video, which is 51% of the video. In the second example TripNet sees 384 frames of the video, which is 80% of the frames.

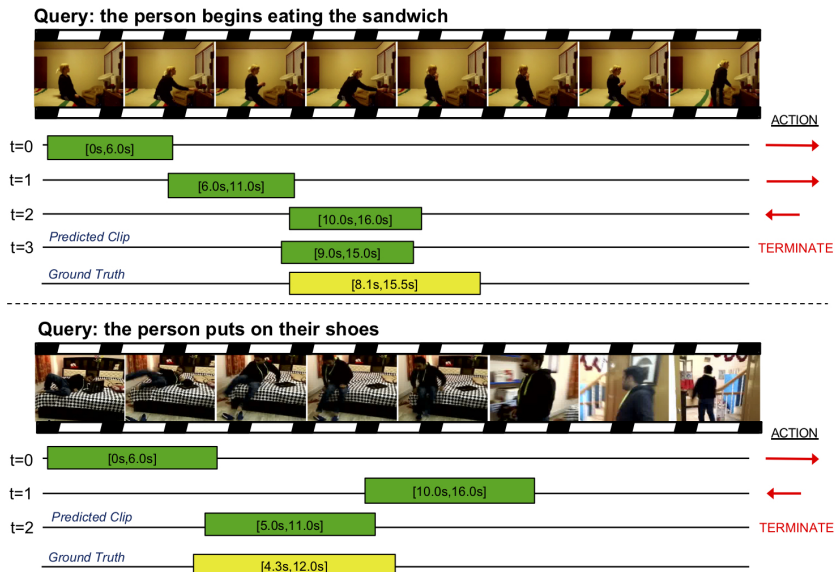


Figure 3: Qualitative performance of TripNet-GA: We show two examples where the TripNet agent skips through the video looking at different candidate windows before terminating the search. Both these videos are from the Charades-STA dataset.

5 Conclusion

Localizing moments in long, untrimmed videos using natural language queries is a useful and challenging task for fine-grained video retrieval. Most previous works are designed to analyze 100% of the video frames. In this paper, we have introduced a system that uses a gated-attention mechanism over cross-modal features to automatically localize a moment in time given a natural language text query with high accuracy. Our model incorporates a policy network which is trained for efficient search performance, resulting in a system that on average examines less than 50% of the video frames in order to localize a clip of interest, while achieving a high level of accuracy.

References

- [1] Humam Alwassel, Fabian Caba Heilbron, and Bernard Ghanem. Action search: Spotting actions in videos and its application to temporal action localization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 251–266, 2018.
- [2] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. Vision-and-language navigation:

- Interpreting visually-grounded navigation instructions in real environments. In *CVPR*, pages 3674–3683, 2018.
- [3] Yunlong Bian, Chuang Gan, Xiao Liu, Fu Li, Xiang Long, Yandong Li, Heng Qi, Jie Zhou, Shilei Wen, and Yuanqing Lin. Revisiting the effectiveness of off-the-shelf temporal modeling approaches for large-scale video classification. *arXiv preprint arXiv:1708.03805*, 2017.
- [4] Piotr Bojanowski, Rémi Lajugie, Edouard Grave, Francis Bach, Ivan Laptev, Jean Ponce, and Cordelia Schmid. Weakly-supervised alignment of video with text. In *ICCV*, pages 4462–4470, 2015.
- [5] Shyamal Buch, Victor Escorcia, Chuanqi Shen, Bernard Ghanem, and Juan Carlos Nieves. Sst: Single-stream temporal action proposals. In *CVPR*, pages 2911–2920, 2017.
- [6] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A Ross, Jia Deng, and Rahul Sukthankar. Rethinking the faster r-cnn architecture for temporal action localization. In *CVPR*, pages 1130–1139, 2018.
- [7] Devendra Singh Chaplot, Kanthashree Mysore Sathyendra, Rama Kumar Pasumarthi, Dheeraj Rajagopal, and Ruslan Salakhutdinov. Gated-attention architectures for task-oriented language grounding. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [8] Jingyuan Chen, Xinpeng Chen, Lin Ma, Zequn Jie, and Tat-Seng Chua. Temporally grounding natural sentence in video. In *EMNLP*, pages 162–171, 2018.
- [9] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [10] Xiyang Dai, Bharat Singh, Guyue Zhang, Larry S Davis, and Yan Qiu Chen. Temporal context network for activity localization in videos. In *ICCV*, pages 5793–5802, 2017.
- [11] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied question answering. In *CVPR*, volume 5, page 6, 2018.
- [12] Bhuwan Dhingra, Hanxiao Liu, Zhilin Yang, William W Cohen, and Ruslan Salakhutdinov. Gated-attention readers for text comprehension. *arXiv preprint arXiv:1606.01549*, 2016.
- [13] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, pages 2625–2634, 2015.
- [14] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. Tall: Temporal activity localization via language query. *arXiv preprint arXiv:1705.02101*, 2017.
- [15] Jiyang Gao, Zhenheng Yang, and Ram Nevatia. Cascaded boundary regression for temporal action detection. *arXiv preprint arXiv:1705.01180*, 2017.

- [16] Runzhou Ge, Jiyang Gao, Kan Chen, and Ram Nevatia. Mac: Mining activity concepts for language-based temporal localization. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 245–253. IEEE, 2019.
- [17] Soham Ghosh, Anuva Agarwal, Zarana Parekh, and Alexander Hauptmann. Excl: Extractive clip localization using natural language descriptions. *arXiv preprint arXiv:1904.02755*, 2019.
- [18] Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, and Ali Farhadi. Iqa: Visual question answering in interactive environments. In *CVPR*, pages 4089–4098, 2018.
- [19] Meera Hahn, Nataniel Ruiz, Jean-Baptiste Alayrac, Ivan Laptev, and James M Rehg. Learning to localize and align fine-grained actions to sparse instructions. *arXiv preprint arXiv:1809.08381*, 2018.
- [20] Dongliang He, Xiang Zhao, Jizhou Huang, Fu Li, Xiao Liu, and Shilei Wen. Read, watch, and move: Reinforcement learning for temporally grounding natural language descriptions in videos. *arXiv preprint arXiv:1901.06829*, 2019.
- [21] Fabian Caba Heilbron, Wayner Barrios, Victor Escorcia, and Bernard Ghanem. Scc: Semantic context cascade for efficient action detection. In *CVPR*, pages 3175–3184. IEEE, 2017.
- [22] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *ICCV*, pages 5803–5812, 2017.
- [23] Ronghang Hu, Huazhe Xu, Marcus Rohrbach, Jiashi Feng, Kate Saenko, and Trevor Darrell. Natural language object retrieval. In *CVPR*, pages 4555–4564, 2016.
- [24] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, pages 2901–2910, 2017.
- [25] Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Skip-thought vectors. In *NIPS*, pages 3294–3302, 2015.
- [26] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *ICCV*, 2017.
- [27] Jingen Liu, Jiebo Luo, and Mubarak Shah. Recognizing realistic actions from videos “in the wild”. In *CVPR*, pages 1996–2003. IEEE, 2009.
- [28] Meng Liu, Xiang Wang, Liqiang Nie, Xiangnan He, Baoquan Chen, and Tat-Seng Chua. Attentive moment retrieval in videos. In *Conference on Research & Development in Information Retrieval*, pages 15–24. ACM, 2018.
- [29] Chih-Yao Ma, Asim Kadav, Iain Melvin, Zsolt Kira, Ghassan AlRegib, and Hans Peter Graf. Attend and interact: Higher-order object interactions for video understanding. In *CVPR*, pages 6790–6800, 2018.

- [30] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *CVPR*, pages 11–20, 2016.
- [31] Antoine Miech, Ivan Laptev, and Josef Sivic. Learnable pooling with context gating for video classification. *arXiv preprint arXiv:1706.06905*, 2017.
- [32] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *ICML*, pages 1928–1937, 2016.
- [33] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543, 2014.
- [34] Michaela Regneri, Marcus Rohrbach, Dominikus Wetzel, Stefan Thater, Bernt Schiele, and Manfred Pinkal. Grounding action descriptions in videos. *ACL*, 1:25–36, 2013.
- [35] Cristian Rodriguez-Opazo, Edison Marrese-Taylor, Fatemeh Sadat Saleh, Hongdong Li, and Stephen Gould. Proposal-free temporal moment localization of a natural-language query in video using guided attention. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2453–2462. IEEE, 2020.
- [36] Marcus Rohrbach, Sikandar Amin, Mykhaylo Andriluka, and Bernt Schiele. A database for fine grained activity detection of cooking activities. In *CVPR*, pages 1194–1201. IEEE, 2012.
- [37] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.
- [38] Ozan Sener, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. Unsupervised semantic parsing of video collections. In *ICCV*, pages 4480–4488, 2015.
- [39] Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *CVPR*, pages 1049–1058, 2016.
- [40] Zheng Shou, Jonathan Chan, Alireza Zareian, Kazuyuki Miyazawa, and Shih-Fu Chang. Cdc: Convolutional-de-convolutional networks for precise temporal action localization in untrimmed videos. In *CVPR*, pages 5734–5743, 2017.
- [41] Gunnar A Sigurdsson, Santosh Divvala, Ali Farhadi, and Abhinav Gupta. Asynchronous temporal fields for action recognition. In *CVPR*, 2016.
- [42] Xiaomeng Song and Yahong Han. Val: Visual-attention action localizer. In *Pacific Rim Conference on Multimedia*, pages 340–350. Springer, 2018.
- [43] Stefanie Tellex and Deb Roy. Towards surveillance video search by natural language query. In *International Conference on Image and Video Retrieval*, page 38. ACM, 2009.
- [44] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. C3d: generic features for video analysis. *CoRR, abs/1412.0767*, 2:7, 2014.

-
- [45] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497. IEEE, 2015.
- [46] Weining Wang, Yan Huang, Liang Wang, and . Language-driven temporal activity localization: A semantic matching reinforcement learning model. In *CVPR*, pages 334–343, 2019.
- [47] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. *ICCV*, 2017.
- [48] Huijuan Xu, Kun He, Leonid Sigal, Stan Sclaroff, and Kate Saenko. Text-to-clip video retrieval with early fusion and re-captioning. *arXiv preprint arXiv:1804.05113*, 2018.
- [49] Serena Yeung, Olga Russakovsky, Greg Mori, and Li Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *CVPR*, pages 2678–2687, 2016.
- [50] Jun Yuan, Bingbing Ni, Xiaokang Yang, and Ashraf A Kassim. Temporal action localization with pyramid of score distribution features. In *CVPR*, pages 3093–3102, 2016.
- [51] Yitian Yuan, Tao Mei, Wenwu Zhu, and . To find where you talk: Temporal sentence localization in video with attention based location regression. *arXiv preprint arXiv:1804.07014*, 2018.