

# Hypotheses Tree Building for One-Shot Temporal Sentence Localization

Daizong Liu<sup>1,2</sup>, Xiang Fang<sup>3</sup>, Pan Zhou<sup>1\*</sup>, Xing Di<sup>4</sup>, Weining Lu<sup>5</sup>, Yu Cheng<sup>6</sup>

<sup>1</sup>School of Cyber Science and Engineering, Huazhong University of Science and Technology

<sup>2</sup>Wangxuan Institute of Computer Technology, Peking University

<sup>3</sup>Nanyang Technological University

<sup>4</sup>ProtagoLabs Inc

<sup>5</sup>Tsinghua University

<sup>6</sup>Microsoft Research

dzliu@hust.edu.cn, xfang9508@gmail.com, panzhou@hust.edu.cn, xing.di@protagolabs.com, luwn@tsinghua.edu.cn, yu.cheng@microsoft.com

## Abstract

Given an untrimmed video, temporal sentence localization (TSL) aims to localize a specific segment according to a given sentence query. Though respectable works have made decent achievements in this task, they severely rely on dense video frame annotations, which require a tremendous amount of human effort to collect. In this paper, we target another more practical and challenging setting: one-shot temporal sentence localization (one-shot TSL), which learns to retrieve the query information among the entire video with only one annotated frame. Particularly, we propose an effective and novel tree-structure baseline for one-shot TSL, called Multiple Hypotheses Segment Tree (MHST), to capture the query-aware discriminative frame-wise information under the insufficient annotations. Each video frame is taken as the leaf-node, and the adjacent frames sharing the same visual-linguistic semantics will be merged into the upper non-leaf node for tree building. At last, each root node is an individual segment hypothesis containing the consecutive frames of its leaf-nodes. During the tree construction, we also introduce a pruning strategy to eliminate the interference of query-irrelevant nodes. With our designed self-supervised loss functions, our MHST is able to generate high-quality segment hypotheses for ranking and selection with the query. Experiments on two challenging datasets demonstrate that MHST achieves competitive performance compared to existing methods.

## Introduction

Temporal sentence localization (TSL) (Anne Hendricks et al. 2017; Gao et al. 2017) has drawn increasing attention in recent years, which aims to retrieve a temporal video segment that semantically corresponds to a given sentence query. An illustrative example of TSL is shown in Figure 1 (a). Clearly, this task involves both computer vision and natural language processing techniques for multi-modal encoding and cross-modal reasoning, to correctly locate the segment boundaries from an untrimmed video.

Most previous TSL works (Liu et al. 2018b; Ge et al. 2019; Yuan et al. 2019; Zhang et al. 2019a, 2020a; Chen et al. 2020a; Yuan, Mei, and Zhu 2019; Liu et al. 2020b,a,

\*Corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

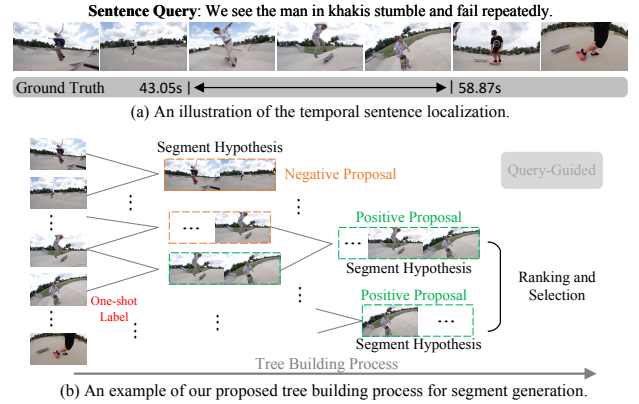


Figure 1: (a) The example of the temporal sentence localization (TSL) task. (b) Illustration of the one-shot TSL setting and our proposed tree-structure method.

2021a; Liu, Qu, and Zhou 2021; Liu et al. 2022d,b) proposed for this task are under fully-supervised setting, where each frame is manually labeled as the query-relevant or query-irrelevant frame. In spite of their great advances, these methods severely rely on abundant video-query annotations, which is labor-intensive and time-consuming to collect in real-world scenarios. To alleviate this problem, some recent works explore a weakly-supervised setting (Mithun, Paul, and Roy-Chowdhury 2019; Lin et al. 2020; Zhang et al. 2020c; Liu et al. 2022c; Chen et al. 2020b; Song et al. 2020) with only the video-query correspondence rather than the dense frame annotations. However, their performances are validated to be less satisfied with such weak supervision.

In this paper, we aim to tackle a more practical and challenging scenario for TSL task, *i.e.*, one-shot TSL, which only has the label of one query-relevant frame in each video. This one-shot TSL setting is first proposed by ViGA (Cui et al. 2022), however, the performance of this model is severely underestimated by current limited designs. We find that previous supervised TSL methods also can not be directly applicable to this challenging setting due to the following limitations: 1) Firstly, most existing works generally pre-define multiple segment proposals and utilize the dense annotations to score these proposals for ranking and selec-

tion. However, one-shot TSL has only one frame label in each video, lacking sufficient knowledge to build and score the segment proposals. 2) Secondly, there are many unlabeled frames that may be either irrelevant to the query or the labeled frame within each video. These frames may bring confounding between the frame label and irrelevant visual features. Therefore, how to discriminate their frame-wise representations for precise segment boundary estimation is also a challenging issue. 3) Thirdly, with limited supervision signals, traditional supervised loss functions are ineffective enough to train the one-shot TSL models. Thus, an appropriate framework and the training strategy should be well-designed for this one-shot setting.

To this end, we propose a novel and effective baseline model for one-shot TSL task, called Multiple Hypotheses Segment Tree (MHST), which adopts a tree structure to generate learnable segment hypothesis by merging the adjacent frames sharing the same visual-linguistic semantics with the query and the labeled frame. As shown in Figure 1 (b), the tree building process is taken as the segment hypothesis construction process. Specifically, we start from treating video frames as the initial leaf-nodes, and then merge the adjacent nodes to the upper non-leaf nodes based on the visual relevance of themselves and their linguistic relevance with the query. During the tree building, we further introduce a tree pruning strategy to selectively weaken the impact of the query-irrelevant nodes. The final hypotheses tree will have several root nodes, which represent corresponding segment hypotheses constructed by the frames of its contained leaf-nodes. At last, we take the segments containing the labeled frame as positive samples and the others as the negative ones. We devise several self-supervised training losses to jointly learn discriminative frame-wise representations for accurately segment hypotheses scoring and selection. During the inference stage, we directly choose the segment hypothesis with the highest confidence score as the final prediction. To sum up, our main contributions are as follows:

- In this paper, we investigate the structure limitations of previous TSL methods on the challenging one-shot setting. Specifically, we devise a new tree-structure framework for one-shot TSL, called MHST, to construct query-related segment hypothesis based on the solely one labeled frame.
- As for the hypotheses tree building, we propose to merge the adjacent nodes into their upper node based on both visual and linguistic relevances. We also introduce a tree pruning strategy to filter out the query-irrelevant nodes. During the training, we apply the self-supervised losses to learn the model under limited annotations.
- We conduct comprehensive experiments on two challenging datasets (ActivityNet Captions and CharadesSTA). The results demonstrate the effectiveness of our proposed method, where MHST achieves decent results and outperforms most fully-supervised methods.

## Related Work

**Temporal sentence localization.** Most of the existing TSL methods refer to fully-supervised setting where all video-

query pairs are annotated in details, including corresponding segment boundaries. Therefore, the main challenge in such setting is how to align multi-modal features well to predict precise boundary. Some works (Liu, Qu, and Hu 2022; Liu and Hu 2022b,a; Chen et al. 2018; Liu et al. 2018b; Ge et al. 2019; Guo, Liu, and Zhou 2022; Zhang et al. 2019a; Liu et al. 2021b, 2022a; Fang et al. 2021, 2022b,a; Liu et al. 2022e) integrate sentence information with each fine-grained video clip unit, and predict the scores of candidate segments by gradually merging the fusion feature sequence over time. Although these methods achieve good performances, they severely rely on the quality of the segment proposals and are time-consuming. Without using proposals, some latest methods (Nan et al. 2021; Zhang et al. 2020a; Chen et al. 2020a; Yuan, Mei, and Zhu 2019) are proposed to leverage the interaction between video and sentence to directly predict the starting and ending frames. However, the above methods heavily rely on the datasets that require numerous manually labelled annotations for training.

To ease the human labelling efforts, several recent works (Mithun, Paul, and Roy-Chowdhury 2019; Chen et al. 2020b; Song et al. 2020; Lin et al. 2020; Zhang et al. 2020c) consider a weakly-supervised setting which only access the information of matched video-query pairs without accurate segment boundaries. However, their performance is less satisfied with such weak supervision. Considering that we are more likely to have a limited annotation budget rather than full annotation or no annotation in practice, (Cui et al. 2022) introduce a new practical setting for TSL task, *i.e.*, solely labeling one query-relevant frame in each video. However, the performance of this model is severely underestimated by current limited designs.

**Multiple hypotheses construction.** Multiple hypotheses strategies are firstly used in the field of object tracking (Blackman and Popoli 1999; Blackman 2004) in videos. Their widely used hypotheses tracking (Cox and Hingorani 1996) algorithm originally evaluates its usefulness in the context of visual tracking and motion correspondence. To further improve the tracking quality, multiple hypotheses tracking in (Kim et al. 2015) propose a tracking tree with a scoring function to prune the hypothesis space efficiently and accurately which is suited to current visual tracking context. In our method, we adapt the multiple hypotheses strategy to the video segment construction scenario, where propagation between the adjacent frames is newly determined by their visual-linguistic relevance instead of the unreliable appearance similarity. We build such multiple hypotheses tree to generate query-related segment proposals within each video by grouping the query-aware semantically closed adjacent frames, and measure their scores for selection. Different from (Li et al. 2022), we propose newly designed self-supervised losses for TSL task.

## The Proposed Model

### Overview

**Problem definition.** Given an untrimmed video  $V = \{v_i\}_{i=1}^{N_v}$  and the natural language query  $Q = \{q_i\}_{i=1}^{N_q}$ , where  $N_v$  and  $N_q$  are the number of frames and words, the tem-

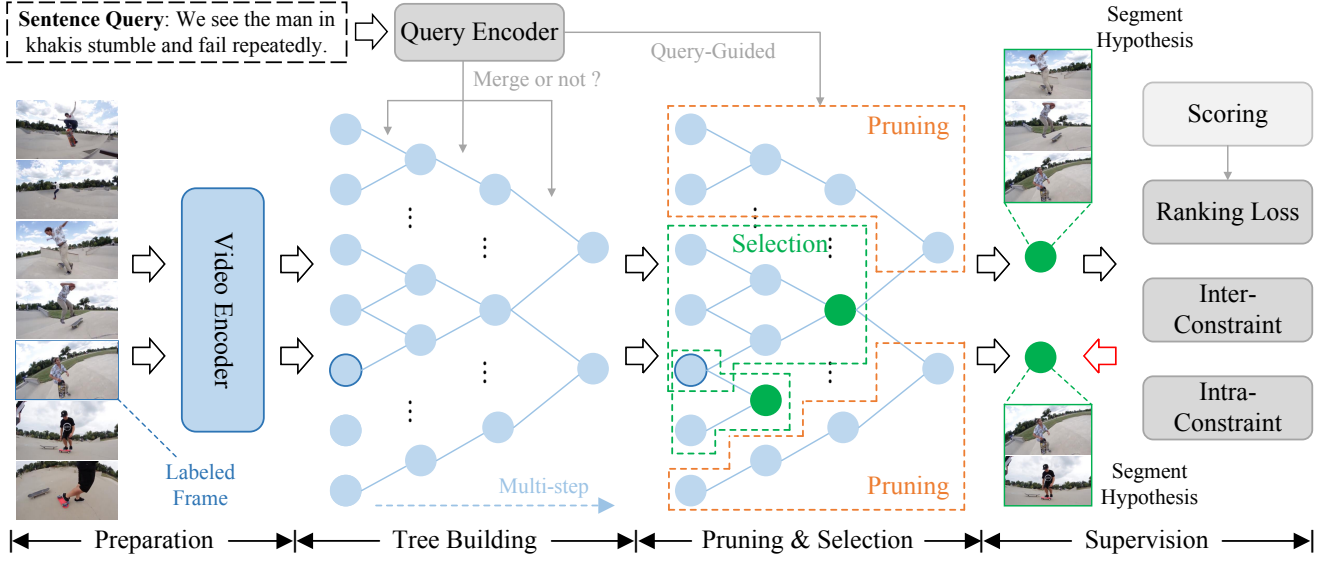


Figure 2: An overview of the proposed architecture for the one-shot TSL task. Given the video-query input, we first take each video frame as leaf-node, and then iteratively build the hypotheses segment tree by merging the nodes sharing the same visual-linguistic semantics. During the tree building, we also employ a tree pruning strategy to remove and down-weight the query-irrelevant nodes. After that, we take the final tree root nodes as the segment hypotheses. At last, we utilize several self-supervised losses to score and rank these segments for learning more discriminative representations.

poral sentence localization (TSL) task aims to localize the query-described activity segment from the video. Particularly, in one-shot TSL setting, there is solely one matched frame annotated with the query in video  $V$ . Therefore, it is more challenging than previous TSL works to predict the accurate segment with the limited supervision signals.

**Pipeline.** In this section, considering the few supervision in one-shot TSL task, we propose a novel pipeline named Multiple Hypotheses Segment Tree (MHST) to construct multiple segment hypotheses that are visual-related and linguistic-related to the labeled frame and the query. Specifically, as shown in Figure 2, MHST contains four main steps: *Firstly*, given a video-query pair, we extract frame-wise features via a video encoder and extract sentence-level features via a query encoder for preparation. *Secondly*, we build the hypotheses tree to discriminate frame-wise information for segment construction. At the beginning, each video frame is taken as a leaf-node. Then, the adjacent nodes will be iteratively merged into an upper node based on their visual-relevance and the query-relevance. *Thirdly*, to avoid the tree explosion and the negative impact of unnecessary nodes, we further introduce a tree pruning strategy to remove and down-weight the nodes conditioned on the query. *At last*, the final generated root nodes are the segment hypotheses constructed by the frames of their contained leaf-nodes. We develop a scoring head to rank these segments and design two self-supervised losses to train the whole hypotheses tree for discriminating the frame-wise representation.

### Preparation

Before building the hypotheses tree, we first extract both video and query features for preparation.

**Video encoding.** Following previous works (Zhang et al.

2019b, 2020a), given the video  $V$ , we first extract its frame-wise features by a pre-trained C3D network (Tran et al. 2015), and then employ a multi-head self-attention (Vaswani et al. 2017) module to capture the long-range dependencies among video frames. We denote the extracted video features as  $V = \{v_i\}_{i=1}^{N_v} \in \mathbb{R}^{N_v \times d}$ , where  $d$  is feature dimension.

**Query encoding.** Given the query  $Q$ , we first utilize the Glove (Pennington, Socher, and Manning 2014) model to embed each word into dense vector, and then employ Bi-GRU (Chung et al. 2014) to encode its sequential information. The final sentence-level feature  $q \in \mathbb{R}^d$  can be obtained by concatenating the last hidden unit outputs in Bi-GRU.

### Hypotheses Tree Building

Here, we will illustrate how to construct the hypotheses tree for generating the possible query-related segments. Given the frame features  $\{v_i\}_{i=1}^{N_v}$ , we treat each frame as the initial leaf-node, and iteratively merge the adjacent frames as the upper non-leaf node if they share the same visual-linguistic semantics. We repeat this merging process until there is none matched adjacent nodes and take the final node as the root node. Therefore, the leaf-nodes contained in the root node are the consecutive frames for constructing the possible query-related segment. To generate multiple root nodes in once tree building process, we try to measure and merge each pair of adjacent leaf-nodes at the beginning.

Specifically, we take two leaf-nodes  $v_i, v_j$  as an example to illustrate the node-wise merging process. To determine whether the node pair  $v_i, v_j$  should be merged into the same segment (or upper node), we calculate both the linguistic relevance between each node pair with the language query and the visual relevance between the nodes as the judgments.

For the linguistic relevance calculation, we first compute the query-guided information relation between the query feature and each frame feature as follows:

$$r_{v_i}^{qv} = \text{sigmoid}((\mathbf{W}_1 \mathbf{v}_i)(\mathbf{W}_2 \mathbf{q})^\top), \quad (1)$$

$$r_{v_j}^{qv} = \text{sigmoid}((\mathbf{W}_1 \mathbf{v}_j)(\mathbf{W}_2 \mathbf{q})^\top), \quad (2)$$

where  $\mathbf{W}_1, \mathbf{W}_2$  are the projection matrices. Then, we measure the query-aware difference between the node pair via  $r^{qv} = |r_{v_i}^{qv} - r_{v_j}^{qv}|$ . For the visual relevance calculation, we directly compute the appearance relation between two frame features via a cosine similarity function as follows:

$$r^{vv} = \frac{\mathbf{v}_i \mathbf{v}_j^\top}{\|\mathbf{v}_i\|_2 \|\mathbf{v}_j\|_2}. \quad (3)$$

Therefore, we can utilize both linguistic and visual relevances as scores to evaluate the semantic difference between the adjacent node pair  $\mathbf{v}_i, \mathbf{v}_j$  by:

$$r = \lambda_1 r^{qv} + \lambda_2 r^{vv}, \quad (4)$$

where  $\lambda_1, \lambda_2$  are used to control the balance.

With the relevant difference value  $r$ , we can rank all node pairs in each step and pick out the top  $\alpha$ , which is a hyperparameter to be set as a percentage. For the top  $\alpha$  node pair  $\mathbf{v}_i, \mathbf{v}_j$ , we merge them into a new ancestor non-leaf node:

$$\mathbf{v}_{i,j}^{new} = \mathbf{W}_3 \mathbf{v}_i + \mathbf{W}_3 \mathbf{v}_j + \mathbf{b}, \quad (5)$$

where  $\mathbf{W}_3, \mathbf{b}$  are the learnable weights. During the hypotheses tree construction, we repeat this node-merging process until there is no-relevant node pairs.

## Hypotheses Tree Pruning

During the building process of the hypotheses tree, not all the non-leaf nodes in the tree branches are closely related to the sentence query. Therefore, we have to take a pruning step to remove these non-leaf nodes and their descendant non-leaf nodes. In other words, we need to determine the mostly likely query-relevant non-leaf nodes for accurately constructing most query-relevant video segment. To this end, we first calculate the semantic relevance  $r$  for each non-leaf node and the query following the Equation (4) and then remove the nodes that are with semantic relevance less than a hyperparameter  $\tau$ . To reserve enough frame nodes (leaf-nodes) for other hypotheses construction, we do not remove their leaf-nodes. Instead, we down-weight the contribution of these lead-nodes to the final segment as follows:

$$(\mathbf{v}_i)' = \lambda_\tau r_{v_i}^{qv} \mathbf{v}_i, \quad (6)$$

where  $r_{v_i}^{qv}$  is the linguistic relevance in Equation (1) to measure the similarity between frame node and query. In particular, we apply this down-weighting process to all leaf-nodes, where  $\lambda_\tau$  of the lead-nodes in the removed non-leaf nodes are set to 0.5,  $\lambda_\tau$  of other leaf-nodes are set to 1.0.

We take a  $L$ -scan pruning method to prune the disturbing non-leaf nodes gradually instead of pruning the whole tree. In particular, for every  $L$  step, we apply the pruning process to remove the query-irrelevant non-leaf nodes in current step. Then, we track and remove their descendant non-leaf nodes in previous  $L - 1$  step. Finally, we down-weight their leaf-nodes via Equation (6).

## Segment Hypothesis Selection and Training

After the hypotheses tree building and pruning, we can get the final segment hypotheses tree in which each root node represents a segment hypothesis containing the consecutive frames of its descendant leaf-nodes. We directly select the root nodes that contain the ground-truth annotated leaf-node as the positive hypotheses, and denote their corresponding segment proposals as  $P = \{p_i\}_{i=1}^{N_p}$  where  $N_p$  is the proposal number. To generate confidence scores  $S = \{s_i\}_{i=1}^{N_p}$  for these segment proposals, we feed their merged node-wise features (obtained by Equation (5)) into a fully connected layer with further sigmoid function. Then, we apply the selection algorithm that considers the confidence score to select the top- $K$  proposals  $P^K = \{p_i^K\}_{i=1}^K$  and give their corresponding confidence scores as  $S^K = \{s_i^K\}_{i=1}^K$ .

**Training.** In order to learn and correct the above confidence scores, we apply a ranking loss based on a reward policy. Specifically, we define the reward  $R_i^K$  for the proposal  $p_i^K$  with a reward function to encourage proposals with higher linguistic relevance  $r^{qv}$ . Then the strategy of policy gradient is used to correct the scores. Note that the confidence scores are normalized by a softmax layer, which is an extremely important operation to highlight the semantically matching proposals and weaken the mismatched ones. The ranking loss can be formulated by:

$$\mathcal{L}_{rank} = \sum_{i=1}^K -R_i^K \log\left(\frac{\exp(s_i^K)}{\sum_{j=1}^K \exp(s_j^K)}\right). \quad (7)$$

Beside, to further assist the node-wise representation learning during the hypotheses tree construction, we additionally develop an *inter-constrain* and an *intra-constrain* losses to train the tree-structure framework. As for the *inter-constrain* loss, we take the matched video-sentence pairs as positive samples and the unmatched video-sentence pairs as negative samples, and use the weighted binary cross-entropy loss to supervise the query-relevance  $r^{qv}$  of the top- $K$  hypothesis segments as:

$$\mathcal{L}_{inter} = \sum_{j=1}^J \sum_{i=1}^K (-y_j \log(r_i^{qv}) - (1 - y_j) \log(1 - r_i^{qv})), \quad (8)$$

where  $J$  is the number of all video-sentence pairs,  $y_j$  the label of the  $j$ -th sample that equals 1 for the matched video-sentence pairs and 0 for the unmatched video-sentence pairs. As for the *intra-constrain* loss, within the same video, we take the top- $K$  segment proposals as positive samples and define the random segments unoverlapped with the top- $K$  segment proposals as negative samples. We adopt a hinge loss to supervise their semantics as:

$$\mathcal{L}_{intra} = \sum_{i=1}^K \max(0, \beta - r_i^{qv} + \bar{r}_i^{qv}), \quad (9)$$

where  $\bar{r}_i^{qv}$  denotes the query-relevance of negative samples. During the training, we utilize three fixed weights to balance the values of above three losses to joint learn the model.

Methods	Type	ActivityNet Captions			
		R@1	R@1	R@5	R@5
		IoU=0.3	IoU=0.5	IoU=0.3	IoU=0.5
Random	FS	18.64	7.73	52.78	29.49
VSA-RNN	FS	39.28	24.43	70.84	55.52
VSA-STV	FS	41.71	24.01	71.05	56.62
CTRL	FS	-	29.01	-	59.17
TGN	FS	43.81	27.93	54.56	44.20
2D-TAN	FS	59.45	44.51	85.53	77.13
IVG-DCL	FS	63.22	43.84	-	-
DRN	FS	-	45.45	-	<b>77.97</b>
CTF	WS	44.30	23.60	-	-
ICVC	WS	46.62	29.52	80.92	66.61
MARN	WS	47.01	29.95	72.02	57.49
SCN	WS	47.23	29.22	71.45	55.69
LCNet	WS	48.49	26.33	82.51	62.66
CCL	WS	50.12	31.07	77.36	61.29
VCA	WS	50.45	31.00	71.79	53.83
WSTAN	WS	52.45	30.01	79.38	63.42
CRM	WS	55.26	32.19	-	-
ViGA	OS	59.61	35.79	-	-
<b>MHST</b>	OS	<b>64.34</b>	<b>45.68</b>	<b>86.92</b>	<b>77.75</b>

Table 1: Performance comparisons on ActivityNet Captions dataset, where FS: fully-supervised setting, WS: weakly-supervised setting, and OS: one-shot setting.

**Testing.** During the testing, we directly choose the root node from the constructed tree with the maximum confidence score, and then generate corresponding segment as the final prediction.

## Experiments

### Dataset

**ActivityNet Captions.** This dataset is built from ActivityNet v1.3 dataset (Caba Heilbron et al. 2015) for dense video captioning. It contains 20000 YouTube videos with 100000 queries. On average, videos are about 120 seconds and queries are about 13.5 words. We follow the public split of the dataset that contains a training set and two validation sets val 1 and val 2. Following common settings, we use val 1 as our validation set and use val 2 as our testing sets.

**Charades-STA.** This dataset is built from the Charades (Sigurdsson et al. 2016) dataset and transformed into temporal sentence localization task by (Gao et al. 2017). It contains 16128 video-sentence pairs with 12408 pairs used for training and 3720 for testing. The videos are about 30 seconds on average. The annotations are generated by sentence decomposition and keyword matching with manually check.

### Experimental Settings

**Evaluation metric.** We adopt “R@n, IoU=m” as the evaluation metrics, following previous works (Gao et al. 2017; Liu et al. 2018a; Zhang et al. 2020a). The “R@n, IoU=m” denotes the percentage of language queries having at least one result whose Intersection over Union (IoU) with ground truth is larger than m in top-n retrieved segment hypotheses.

**Implementation details.** In order to make a fair comparison with previous works, we utilize the pre-trained C3D

Methods	Type	Charades-STA			
		R@1	R@1	R@5	R@5
		IoU=0.5	IoU=0.7	IoU=0.5	IoU=0.7
Random	FS	8.51	3.03	37.12	14.06
VSA-RNN	FS	10.50	4.32	48.43	20.21
VSA-STV	FS	16.91	5.81	53.89	23.58
CTRL	FS	23.62	8.89	58.92	29.52
2D-TAN	FS	39.81	23.25	79.33	52.15
DRN	FS	45.40	26.40	88.01	55.38
IVG-DCL	FS	<b>50.24</b>	32.88	-	-
SCN	WS	23.58	9.97	71.80	38.87
CTF	WS	27.30	12.90	-	-
WSTAN	WS	29.35	12.28	76.13	41.53
ICVC	WS	31.02	16.53	77.53	41.91
MARN	WS	31.94	14.18	70.00	37.40
CCL	WS	33.21	15.68	73.50	41.87
CRM	WS	34.76	16.37	-	-
VCA	WS	38.13	19.57	78.75	37.75
LCNet	WS	39.19	18.17	80.56	45.24
ViGA	OS	35.11	15.11	-	-
<b>MHST</b>	OS	49.62	<b>34.48</b>	<b>89.29</b>	<b>57.50</b>

Table 2: Performance comparisons on Charades-STA dataset, where FS: fully-supervised setting, WS: weakly-supervised setting, and OS: one-shot setting.

(Tran et al. 2015) model to extract video features and employ the Glove model (Pennington, Socher, and Manning 2014) to obtain word embeddings. As some videos are too long, we set the length of video feature sequences to 128 for Charades-STA and 256 for ActivityNet Captions, respectively. We fix the query length to 10 in Charades-STA and 20 in ActivityNet Captions. The feature dimension  $d$  is set to 512. For the hyper-parameters, we set the percentage  $\alpha$  to 60%, and set the pruning threshold  $\tau$  as 0.7. The balanced weights  $\lambda_1, \lambda_2$  are set to 1.0, 1.0. The step  $L$  in L-scan pruning is set to 3. During training, the learning rate is by default 0.00005, and decays by a factor of 10 for every 35 epochs. The batch size is 1 and the maximum training epoch is 100. For one-shot setting, we randomly selection one labeled frame from the ground-truth as the annotation for each video. All the experiments are implemented by PyTorch.

### Comparison with State-of-the-Art

**Compared methods.** We compare MHST with the state-of-the-art TSL methods grouped into three categories: 1) Fully-supervised setting: Random Selection (Gidaris, Singh, and Komodakis 2018), VSA-RNN and VSA-STV (Gao et al. 2017), CTRL (Gao et al. 2017), TGN (Chen et al. 2018), 2D-TAN (Zhang et al. 2020b), IVG-DCL (Nan et al. 2021), DRN (Zeng et al. 2020). 2) Weakly-supervised setting: CTF (Chen et al. 2020b), ICVC (Chen et al. 2022), MARN (Song et al. 2020), SCN (Lin et al. 2020), LCNet (Yang et al. 2021), CCL (Zhang et al. 2020c), VCA (Wang, Chen, and Jiang 2021), WSTAN (Wang et al. 2021), CRM (Huang et al. 2021). 3) One-shot setting: ViGA (Cui et al. 2022).

**Comparison and analysis.** As shown in Table 1 and 2, we compare our method with existing works on both ActivityNet Captions and Charades-STA datasets, respectively. From the tables, we have the following findings:

Segment Tree	Tree Pruning	Self-supervision	R@1 IoU=0.3	R@1 IoU=0.5	R@5 IoU=0.3	R@5 IoU=0.5
×	×	×	51.97	31.95	73.82	64.47
×	×	✓	54.64	34.71	76.38	66.99
✓	×	×	58.84	39.76	81.17	71.93
✓	×	✓	62.27	43.05	84.53	75.28
✓	✓	×	61.16	42.22	83.39	74.41
✓	✓	✓	<b>64.34</b>	<b>45.68</b>	<b>86.92</b>	<b>77.75</b>

Table 3: Main ablation study on ActivityNet Captions dataset, where we remove each key individual component to investigate its effectiveness.

- Firstly, compared to the weakly-supervised methods, our method outperforms them by a large margin. Specifically, on ActivityNet Captions dataset, compared to the SOTA method CRM, we bring the improvement of 9.08 and 13.49 in terms of R@1, IoU=0.3 and R@1, IoU=0.5. On Charades-STA dataset, we also outperform the SOTA method LCNet by 10.43, 16.31, 8.73 and 12.26 on all metrics. *Note that*, the one-shot setting costs similar human labors as the weakly-supervised setting since the latter also requests the annotators to determine the matched video-query pair with at least one query-relevant frame. However, our performance is much better than the weakly-supervised one, demonstrating the effectiveness of the proposed tree-structure framework.
- Secondly, compared to the fully-supervised methods, our one-shot setting has much less annotations (one labeled frame vs. dense video annotations) for model training. However, as shown in the tables, our method achieves very competitive performances, which are even better than the previous SOTA methods. Specifically, on ActivityNet Captions dataset, compared to DRN, we outperform it by 0.23 in R@1, IoU=0.5. On Charades-STA dataset, compared to IVG-DCL, we bring improvement of 1.60 in R@1, IoU=0.7. This demonstrates that our framework is robust to the weak annotations. Although we have much less annotations than the fully-supervised works, our well-designed tree-structure framework can group the query- and labeled frame-related adjacent frames into the same segment under the supervision of the proposed effective self-supervised losses.
- Thirdly, our performance outperforms ViGA method a lot, demonstrating that our tree-structure pipeline is more effective and reasonable for the one-shot TSL setting.

## Ablation Study

In this section, we conduct ablation study to validate the effectiveness of each components on ActivityNet Captions.

**Main ablation.** To analyze how each model component contributes to the task, we perform main ablation study as shown in Table 3. We start from the baseline model which does not rely on both the tree-structure framework and the self-supervised strategy to address the one-shot TSL. Specifically, this baseline model directly generates multiple coarse segment proposals like previous supervised methods and then utilizes the rank loss in Eq.(7) for training. It shows that this baseline achieves relatively worse performance com-

Methods	Type	R@1 IoU=0.3	R@1 IoU=0.5	R@5 IoU=0.3	R@5 IoU=0.5
2D-TAN	FS	59.45	44.51	85.53	77.13
2D-TAN*	OS	48.76	34.29	75.88	66.02
WSTAN	WS	52.45	30.01	79.38	63.42
WSTAN*	OS	54.69	32.17	80.93	65.58
<b>MHST</b>	<b>OS</b>	<b>64.34</b>	<b>45.68</b>	<b>86.92</b>	<b>77.75</b>

Table 4: Applying the one-shot setting to both fully- and weakly-supervised methods on ActivityNet Captions.

Visual Similarity	Linguistic Similarity	R@1 IoU=0.3	R@1 IoU=0.5	R@5 IoU=0.3	R@5 IoU=0.5
✓	×	60.96	42.17	83.58	74.42
×	✓	62.10	43.34	84.66	75.25
✓	✓	<b>64.34</b>	<b>45.68</b>	<b>86.92</b>	<b>77.75</b>

Table 5: Effect of visual- and linguistic similarities for node merging during the tree building.

Hyperparameters	R@1 IoU=0.3	R@1 IoU=0.5	R@5 IoU=0.3	R@5 IoU=0.5
$\alpha = 50\%$	63.28	44.39	85.56	76.31
$\alpha = 60\%$	<b>64.34</b>	45.68	<b>86.92</b>	<b>77.75</b>
$\alpha = 70\%$	64.32	<b>45.70</b>	86.84	77.69
$\tau = 0.6$	63.47	44.82	85.83	76.95
$\tau = 0.7$	<b>64.34</b>	<b>45.68</b>	<b>86.92</b>	<b>77.75</b>
$\tau = 0.8$	63.96	45.33	86.60	77.41
$L = 1$	63.58	45.01	86.20	77.12
$L = 3$	<b>64.34</b>	45.68	86.92	77.75
$L = 5$	64.26	<b>45.75</b>	<b>87.07</b>	<b>77.83</b>

Table 6: Effect of different hyperparameters for node merging during and tree pruning.

pared to most weakly- and fully-supervised methods. By applying the self-supervised constraint of Eq.(8)(9) to the baseline, the performance improves a lot. Applying the tree-structure framework brings the largest improvement since our well-designed merging strategy helps to distinguish the ambiguous adjacent frames for constructing more accurate segments under the limited supervision signals. Moreover, the self-supervised constraint seems to be more robust to the tree-structure framework due to the high quality of the segment hypotheses. The tree pruning strategy also helps the model to reduce the negative influence of the query-irrelevant frames. Overall, the whole tree-structure framework with both pruning strategy and self-supervised constraints achieves the best results.

**Effect of our one-shot pipeline.** To fairly compare with the existing fully- and weakly-supervised methods, we re-implement some works with our one-shot setting as shown in Table 4. Specifically, for 2D-TAN, we replace the loss of 2D map with our rank loss in Eq.(7) for handling the weak labels. For WSTAN, we also add the rank loss for enriching the model contexts. We keep all the other settings being the same as their original works. From this table, we can find that, although one-shot setting brings additional one frame label to the weakly setting, the improvement of WSTAN\* is still limited since it lacks sufficient self-supervision for discriminating the frame-wise representations and cannot gen-

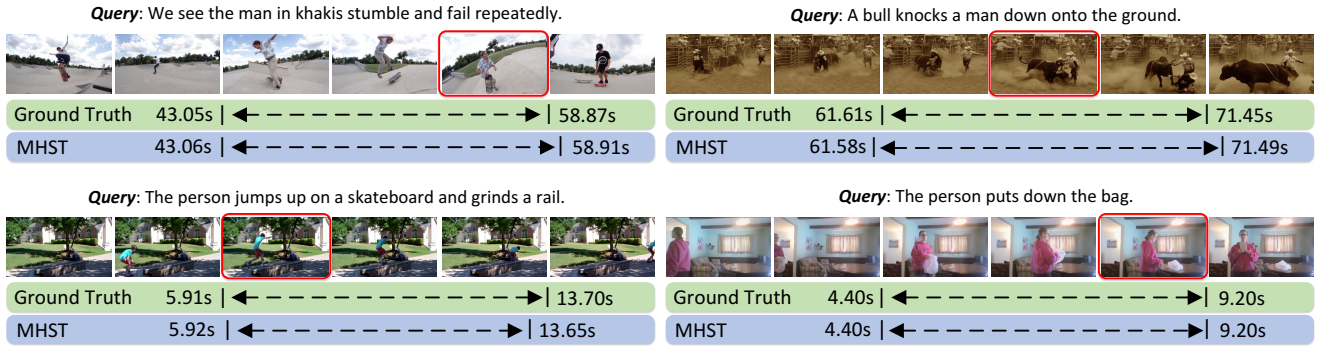


Figure 3: Qualitative results on ActivityNet Captions and Charades-STA. The red rectangle denotes the labeled frame in video.

$\mathcal{L}_{rank}$	$\mathcal{L}_{inter}$	$\mathcal{L}_{intra}$	R@1 IoU=0.3	R@1 IoU=0.5	R@5 IoU=0.3	R@5 IoU=0.5
✓	×	×	61.16	42.22	83.39	74.41
✓	✓	×	62.98	44.01	85.27	76.44
✓	×	✓	62.46	43.49	84.83	76.02
✓	✓	✓	<b>64.34</b>	<b>45.68</b>	<b>86.92</b>	<b>77.75</b>

Table 7: Ablation study on the supervision losses.

erate accurate segment by measuring both visual-linguistic relevances like us. Moreover, the performance of 2D-TAN\* degenerates a lot, demonstrating that previous fully setting is not robust to the weak label. Overall, it indicates that the one-shot setting is worth being investigated, and our proposed tree-structure framework is effectiveness.

**Effect of the visual-linguistic relevance.** As shown in Table 5, we investigate the effect of the visual- and linguistics relevances for node merging during the tree building process. It shows that both of them are crucial for segment construction, since the visual relevance helps to determine the visual similar adjacent frames near the labeled frame while the linguistic relevance helps to filter our the query-irrelevant frames. By applying both of them for node merging, our model can achieve the best performance.

**Hyper-parameters of tree building and pruning.** Moreover, we investigate the robustness of the proposed model to different hyper-parameters in Table 6. For the hyperparameter  $\alpha$  in node merging, we find that the model achieves the best result when  $\alpha$  is set to 60%. Smaller  $\alpha$  will lead negative node merging while larger  $\alpha$  will filter out some positive nodes. As for  $\tau$  in tree pruning, the model achieves the best result when  $\tau = 0.7$ . Larger  $\tau$  will filter out positive nodes. Besides, we also investigate the effect of different number  $L$  in the pruning process. It shows that the model achieves the best result when  $L = 5$ . However, lager  $L$  gets the punishment in speed. Therefore, we set  $L = 3$  in all experiments.

**Effect of the supervision losses.** At last, we investigate the effectiveness of the proposed self-supervised losses of Eq.(8) and (9) in Table 7. Here, the ranking loss  $\mathcal{L}_{rank}$  is the baseline for addressing the one-shot TSL setting. It shows that the ranking loss  $\mathcal{L}_{rank}$  is well-designed to the one-shot setting, and achieves great performance. Moreover, by applying the inter-constraint loss or the intra-constraint loss to the baseline, the model has significant performance by learn-

ing more discriminative frame-wise representations. Overall, both two self-constraint losses  $\mathcal{L}_{inter}$  and  $\mathcal{L}_{intra}$  contributes a lot to the final performance, and we can achieve the best performance by jointly utilizing them.

## Qualitative Results

As shown in Figure 3, we give the visualization of the localization results on both ActivityNet Captions and Charades-STA datasets for examples. In this figure, each video has only one frame (in red rectangle) labeled according to the query. It shows that our proposed MHST is robust to the one-shot setting and can well predict the segment boundaries.

## Conclusion

In this paper, we address the challenging one-shot setting into the temporal sentence localization task. To achieve this goal, the model needs to make full use of the solely one labeled frame in each video to retrieve the target segment according to the query semantics. Considering that the invalid frames unrelated to the query sentence or the labeled frame may bring confounding to the one-shot training process, we design a novel tree-structure framework called Multiple Hypotheses Segment Tree (MHST) to avoid this issue. Specifically, the hypotheses tree module merges adjacent frames sharing the similar visual-linguistic semantics into a new upper node. Then, by iteratively building the tree and pruning the invalid nodes, we can get the complete and query-related root nodes which represent the segment hypotheses constructed by their contained consecutive frames. Finally, we develop several self-supervised losses to train the segment tree and predict the confidence scores for each segment hypothesis. During the inference, we directly choose the segment with highest score as the final prediction. Experimental results on two challenging datasets (ActivityNet Captions and Charades-STA) demonstrate that our MHST achieves a competitive performance compared to existing fully- and weakly-supervised methods.

## Acknowledgements

This work is supported by National Natural Science Foundation of China (NSFC) under Grant No. 61972448.

## References

- Anne Hendricks, L.; Wang, O.; Shechtman, E.; Sivic, J.; Darrell, T.; and Russell, B. 2017. Localizing moments in video with natural language. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Blackman, S.; and Popoli, R. 1999. Design and analysis of modern tracking systems(Book). Norwood, MA: Artech House, 1999.
- Blackman, S. S. 2004. Multiple hypothesis tracking for multiple target tracking. *IEEE Aerospace and Electronic Systems Magazine*, 19(1): 5–18.
- Caba Heilbron, F.; Escorcia, V.; Ghanem, B.; and Carlos Niebles, J. 2015. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 961–970.
- Chen, J.; Chen, X.; Ma, L.; Jie, Z.; and Chua, T.-S. 2018. Temporally grounding natural sentence in video. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 162–171.
- Chen, J.; Luo, W.; Zhang, W.; and Ma, L. 2022. Explore Inter-contrast between Videos via Composition for Weakly Supervised Temporal Sentence Grounding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 267–275.
- Chen, L.; Lu, C.; Tang, S.; Xiao, J.; Zhang, D.; Tan, C.; and Li, X. 2020a. Rethinking the Bottom-Up Framework for Query-based Video Localization. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Chen, Z.; Ma, L.; Luo, W.; Tang, P.; and Wong, K.-Y. K. 2020b. Look closer to ground better: Weakly-supervised temporal grounding of sentence in video. *arXiv preprint arXiv:2001.09308*.
- Chung, J.; Gulcehre, C.; Cho, K.; and Bengio, Y. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *Advances in Neural Information Processing Systems (NIPS)*.
- Cox, I. J.; and Hingorani, S. L. 1996. An efficient implementation of Reid’s multiple hypothesis tracking algorithm and its evaluation for the purpose of visual tracking. *IEEE Transactions on pattern analysis and machine intelligence*, 18(2): 138–150.
- Cui, R.; Qian, T.; Peng, P.; Daskalaki, E.; Chen, J.; Guo, X.; Sun, H.; and Jiang, Y.-G. 2022. Video Moment Retrieval from Text Queries via Single Frame Annotation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*.
- Fang, X.; Hu, Y.; Zhou, P.; and Wu, D. 2021. ANIMC: A Soft Approach for Autoweighted Noisy and Incomplete Multiview Clustering. *IEEE Transactions on Artificial Intelligence*, 3(2): 192–206.
- Fang, X.; Liu, D.; Zhou, P.; and Hu, Y. 2022a. Multi-Modal Cross-Domain Alignment Network for Video Moment Retrieval. *IEEE Transactions on Multimedia*, 1–16.
- Fang, X.; Liu, D.; Zhou, P.; Xu, Z.; and Li, R. 2022b. Hierarchical Local-Global Transformer for Temporal Sentence Grounding. *arXiv preprint arXiv:2208.14882*.
- Gao, J.; Sun, C.; Yang, Z.; and Nevatia, R. 2017. Tall: Temporal activity localization via language query. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 5267–5275.
- Ge, R.; Gao, J.; Chen, K.; and Nevatia, R. 2019. Mac: Mining activity concepts for language-based temporal localization. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 245–253.
- Gidaris, S.; Singh, P.; and Komodakis, N. 2018. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*.
- Guo, C.; Liu, D.; and Zhou, P. 2022. A Hybrid Alignment Loss for Temporal Moment Localization with Natural Language. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. IEEE.
- Huang, J.; Liu, Y.; Gong, S.; and Jin, H. 2021. Cross-sentence temporal and semantic relations in video activity localisation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7199–7208.
- Kim, C.; Li, F.; Ciptadi, A.; and Rehg, J. M. 2015. Multiple hypothesis tracking revisited. In *Proceedings of the IEEE international conference on computer vision*, 4696–4704.
- Li, M.; Wang, T.; Zhang, H.; Zhang, S.; Zhao, Z.; Miao, J.; Zhang, W.; Tan, W.; Wang, J.; Wang, P.; et al. 2022. End-to-End Modeling via Information Tree for One-Shot Natural Language Spatial Video Grounding. *arXiv preprint arXiv:2203.08013*.
- Lin, Z.; Zhao, Z.; Zhang, Z.; Wang, Q.; and Liu, H. 2020. Weakly-supervised video moment retrieval via semantic completion network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 11539–11546.
- Liu, D.; Fang, X.; Hu, W.; and Zhou, P. 2022a. Exploring Optical-Flow-Guided Motion and Detection-Based Appearance for Temporal Sentence Grounding. *arXiv preprint arXiv:2203.02966*.
- Liu, D.; and Hu, W. 2022a. Learning to Focus on the Foreground for Temporal Sentence Grounding. In *Proceedings of the 29th International Conference on Computational Linguistics*, 5532–5541.
- Liu, D.; and Hu, W. 2022b. Skimming, Locating, then Perusing: A Human-Like Framework for Natural Language Video Localization. In *Proceedings of the 30th ACM International Conference on Multimedia*, 4536–4545.
- Liu, D.; Qu, X.; Di, X.; Cheng, Y.; Xu, Z. X.; and Zhou, P. 2022b. Memory-Guided Semantic Learning Network for Temporal Sentence Grounding. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Liu, D.; Qu, X.; Dong, J.; and Zhou, P. 2020a. Reasoning Step-by-Step: Temporal Sentence Localization in Videos via Deep Rectification-Modulation Network. In *Proceedings of the 28th International Conference on Computational Linguistics*, 1841–1851.
- Liu, D.; Qu, X.; Dong, J.; and Zhou, P. 2021a. Adaptive Proposal Generation Network for Temporal Sentence Localization in Videos. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 9292–9301.

- Liu, D.; Qu, X.; Dong, J.; Zhou, P.; Cheng, Y.; Wei, W.; Xu, Z.; and Xie, Y. 2021b. Context-aware Biaffine Localizing Network for Temporal Sentence Grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Liu, D.; Qu, X.; and Hu, W. 2022. Reducing the Vision and Language Bias for Temporal Sentence Grounding. In *Proceedings of the 30th ACM International Conference on Multimedia*, 4092–4101.
- Liu, D.; Qu, X.; Liu, X.-Y.; Dong, J.; Zhou, P.; and Xu, Z. 2020b. Jointly Cross-and Self-Modal Graph Attention Network for Query-Based Moment Localization. In *Proceedings of the 28th ACM International Conference on Multimedia*, 4070–4078.
- Liu, D.; Qu, X.; Wang, Y.; Di, X.; Zou, K.; Cheng, Y.; Xu, Z.; and Zhou, P. 2022c. Unsupervised Temporal Video Grounding with Deep Semantic Clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Liu, D.; Qu, X.; and Zhou, P. 2021. Progressively Guide to Attend: An Iterative Alignment Framework for Temporal Sentence Grounding. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 9302–9311.
- Liu, D.; Qu, X.; Zhou, P.; and Liu, Y. 2022d. Exploring Motion and Appearance Information for Temporal Sentence Grounding. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Liu, D.; Zhou, P.; Xu, Z.; Wang, H.; and Li, R. 2022e. Few-Shot Temporal Sentence Grounding via Memory-Guided Semantic Learning. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Liu, M.; Wang, X.; Nie, L.; He, X.; Chen, B.; and Chua, T.-S. 2018a. Attentive moment retrieval in videos. In *Proceedings of the 41st International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 15–24.
- Liu, M.; Wang, X.; Nie, L.; Tian, Q.; Chen, B.; and Chua, T.-S. 2018b. Cross-modal moment localization in videos. In *Proceedings of the 26th ACM international conference on Multimedia*, 843–851.
- Mithun, N. C.; Paul, S.; and Roy-Chowdhury, A. K. 2019. Weakly supervised video moment retrieval from text queries. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 11592–11601.
- Nan, G.; Qiao, R.; Xiao, Y.; Liu, J.; Leng, S.; Zhang, H.; and Lu, W. 2021. Interventional Video Grounding with Dual Contrastive Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543.
- Sigurdsson, G. A.; Varol, G.; Wang, X.; Farhadi, A.; Laptev, I.; and Gupta, A. 2016. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *European Conference on Computer Vision (ECCV)*, 510–526.
- Song, Y.; Wang, J.; Ma, L.; Yu, Z.; and Yu, J. 2020. Weakly-supervised multi-level attentional reconstruction network for grounding textual queries in videos. *arXiv preprint arXiv:2003.07048*.
- Tran, D.; Bourdev, L.; Fergus, R.; Torresani, L.; and Paluri, M. 2015. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 4489–4497.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems (NIPS)*, 5998–6008.
- Wang, Y.; Deng, J.; Zhou, W.; and Li, H. 2021. Weakly supervised temporal adjacent network for language grounding. *IEEE Transactions on Multimedia*.
- Wang, Z.; Chen, J.; and Jiang, Y.-G. 2021. Visual co-occurrence alignment learning for weakly-supervised video moment retrieval. In *Proceedings of the 29th ACM International Conference on Multimedia*, 1459–1468.
- Yang, W.; Zhang, T.; Zhang, Y.; and Wu, F. 2021. Local correspondence network for weakly supervised temporal sentence grounding. *IEEE Transactions on Image Processing*, 30: 3252–3262.
- Yuan, Y.; Ma, L.; Wang, J.; Liu, W.; and Zhu, W. 2019. Semantic Conditioned Dynamic Modulation for Temporal Sentence Grounding in Videos. In *Advances in Neural Information Processing Systems (NIPS)*, 534–544.
- Yuan, Y.; Mei, T.; and Zhu, W. 2019. To find where you talk: Temporal sentence localization in video with attention based location regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 9159–9166.
- Zeng, R.; Xu, H.; Huang, W.; Chen, P.; Tan, M.; and Gan, C. 2020. Dense regression network for video grounding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 10287–10296.
- Zhang, D.; Dai, X.; Wang, X.; Wang, Y.-F.; and Davis, L. S. 2019a. Man: Moment alignment network for natural language moment retrieval via iterative graph adjustment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1247–1257.
- Zhang, H.; Sun, A.; Jing, W.; and Zhou, J. T. 2020a. Span-based Localizing Network for Natural Language Video Localization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 6543–6554.
- Zhang, S.; Peng, H.; Fu, J.; and Luo, J. 2020b. Learning 2D Temporal Adjacent Networks for Moment Localization with Natural Language. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Zhang, Z.; Lin, Z.; Zhao, Z.; and Xiao, Z. 2019b. Cross-modal interaction networks for query-based moment retrieval in videos. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 655–664.
- Zhang, Z.; Zhao, Z.; Lin, Z.; He, X.; et al. 2020c. Counterfactual Contrastive Learning for Weakly-Supervised Vision-Language Grounding. *Advances in Neural Information Processing Systems (NIPS)*, 33: 18123–18134.