

Text-Visual Prompting for Efficient 2D Temporal Video Grounding

Yimeng Zhang^{1,2}, Xin Chen², Jinghan Jia¹, Sijia Liu¹, Ke Ding²
¹Michigan State University, ²Applied AI, Intel
 {zhan1853, jiajingh, liusiji5}@msu.edu,
 {xin.chen, ke.ding}@intel.com

Abstract

In this paper, we study the problem of temporal video grounding (TVG), which aims to predict the starting/ending time points of moments described by a text sentence within a long untrimmed video. Benefiting from fine-grained 3D visual features, the TVG techniques have achieved remarkable progress in recent years. However, the high complexity of 3D convolutional neural networks (CNNs) makes extracting dense 3D visual features time-consuming, which calls for intensive memory and computing resources. Towards efficient TVG, we propose a novel *text-visual prompting* (TVP) framework, which incorporates optimized perturbation patterns (that we call ‘prompts’) into both visual inputs and textual features of a TVG model. In sharp contrast to 3D CNNs, we show that TVP allows us to effectively co-train vision encoder and language encoder in a 2D TVG model and improves the performance of cross-modal feature fusion using only low-complexity sparse 2D visual features. Further, we propose a *Temporal-Distance IoU* (TDIoU) loss for efficient learning of TVG. Experiments on two benchmark datasets, Charades-STA and ActivityNet Captions datasets, empirically show that the proposed TVP significantly boosts the performance of 2D TVG (e.g., 9.79% improvement on Charades-STA and 30.77% improvement on ActivityNet Captions) and achieves 5× inference acceleration over TVG using 3D visual features. Codes are available at [Open.Intel](https://openintel.org).

1. Introduction

In recent years, we have witnessed great progress on temporal video grounding (TVG) [30, 74]. One key to this success comes from the fine-grained dense 3D visual features extracted by 3D convolutional neural networks (CNNs) (e.g., C3D [56] and I3D [3]) since TVG tasks demand spatial-temporal context to locate the temporal interval of the moments described by the text query. However, due to the high cost of the dense 3D feature extraction, most existing TVG models only take these 3D visual features ex-

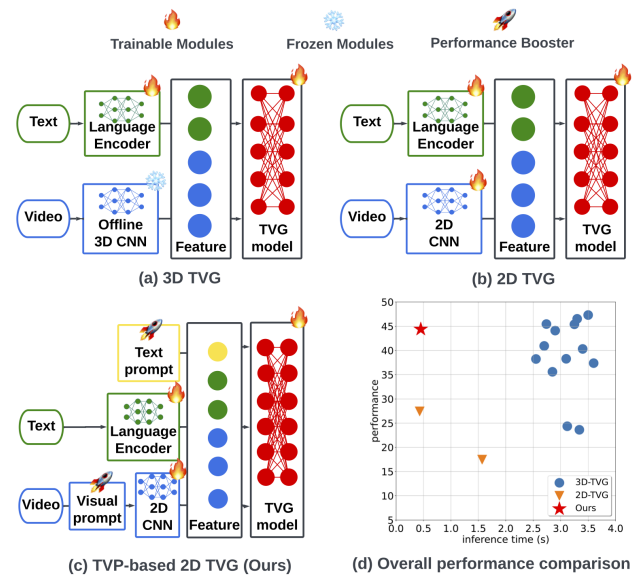


Figure 1. The architecture and performance comparison among TVG methods: **a)** 3D TVG methods [14, 16, 18, 34, 43, 60–62, 64, 67, 69, 71, 73], **b)** 2D TVG methods [1, 7], and **c)** TVP-based 2D TVG (Ours), **d)** overall performance comparison. Ours is the most efficient (least inference time) and achieves competitive performance compared to 3D TVG methods. In contrast to existing TVG methods, which utilize dense video features extracted by non-trainable *offline 3D CNNs* and textual features, our proposed framework utilizes a trainable *2D CNN* as the vision encoder to extract features from sparsely-sampled video frames with a universal set of frame-aware visual prompts and adds text prompts in textual feature space for end-to-end regression-based modeling.

tracted by offline 3D CNNs as inputs instead of co-training during TVG model training.

Although models using 3D visual features (that we call ‘**3D methods or models**’) outperform these using the 2D features (that we call ‘**2D methods or models**’), a unique advantage of 2D methods is that extracting 2D visual features can significantly reduce the cost in TVG tasks [14, 15, 30, 34, 61, 62, 69, 74, 75]. An efficient and lightweight solution with reasonable performance is also demanded in computer vision, NLP, and video-language tasks [19, 23, 38, 41, 68, 76–80]. As discussed above, the methods employing 3D video features are challenging to be

employed in practical applications. It thus has significant practical and economic value to develop compact 2D solutions for TVG tasks. In this paper, we ask:

How to advance 2D TVG methods so as to achieve comparable results to 3D TVG methods?

To address this problem, we propose a novel text-visual prompting (TVP) framework for training TVG models using 2D visual features. As shown in Fig. 1, for existing 2D TVG and 3D TVG methods, they all utilize offline pre-trained vision encoders and language encoders to perform feature extraction. In contrast, our proposed TVP framework is end-to-end trainable. Furthermore, benefiting from text-visual prompting and cross-modal pretraining on large-scale image-text datasets, our proposed framework could achieve comparable performance to 3D TVG methods with significant inference time acceleration.

Conventionally, TVG methods consist of three stages: ① extracting feature from visual and text inputs; ② multi-modal feature fusion; ③ cross-modal modelling. In contrast to conventional methods, TVP incorporates optimized input perturbation patterns (that we call ‘prompts’) into both visual inputs and textual features of a TVG model. We apply trainable parameters in the textual features as text prompts and develop a universal set of frame-aware patterns as visual prompts. Specially, we sample a fixed number of frames from a video and optimize text prompts for the input query sentence and a set of visual prompts for frames with different temporal locations during training. During testing, the same set of optimized visual prompts and textual prompts are applied to all test-time videos. We refer readers to Fig. 2 for illustrations of visual prompts and text prompts introduced. To the best of our knowledge, our work makes the first attempt to utilize prompt learning to successfully improve the performance of regression-based TVG tasks using 2D visual features.

Compared to 3D CNNs, 2D CNNs loses spatiotemporal information of the video during feature extraction. Inspired by the success of transformers on the vision-language tasks [9, 22, 35, 44, 47, 54, 55] and the recent application of prompt learning to transformers in both vision and language domains [2, 25, 27, 32, 37, 40], we choose transformer as our base TVG model and propose to utilize prompts to compensate for the lack of spatiotemporal information in 2D visual features. Furthermore, we develop a Temporal-Distance IoU (TDIoU) loss for training our proposed framework. There are two aspects that distinguish our proposed framework from existing works. First, our proposed framework is designed to boost the performance of the regression-based TVG methods utilizing 2D CNNs as the vision encoder, not for transfer learning [2, 21, 26]. Second, our proposed framework utilizes 2D CNN to extract visual features from

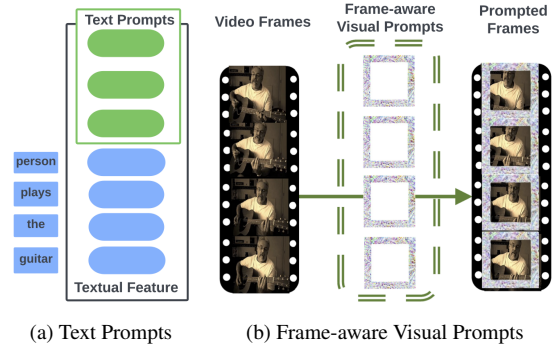


Figure 2. Text-visual prompting illustration. (a) Text prompts are directly applied in the feature space. (b) A set of visual prompts are applied to video frames in order.

sparingly-sampled video frames, which requires less memory and is easier to be applied in practical applications compared to 3D methods [34, 60–62, 69, 75], especially for long videos. Furthermore, thanks to the compact 2D CNN as the vision encoder, our proposed framework could implement the language encoder and visual encoder co-training for better multimodal feature fusion. In summary, the **contributions** of this work are unfolded below:

- We propose an effective and efficient framework to train 2D TVG models, in which we leverage TVP (text-visual prompting) to improve the utility of sparse 2D visual features without resorting to costly 3D features. To the best of our knowledge, it is the first work to expand the application of prompt learning for resolving TVG problems. Our method outperforms all of 2D methods and achieves competitive performance to 3D TVG methods.
- Technology-wise, we integrate visual prompt with text prompt to co-improve the effectiveness of 2D visual features. On top of that, we propose TDIoU (temporal-distance IoU)-based prompt-model co-training method to obtain high-accuracy 2D TVG models.
- Experiment-wise, we show the empirical success of our proposal to boost the performance of 2D TVG on Charades-STA and ActivityNet Captions datasets, e.g., 9.79% improvement in Charades-STA, and 30.77% in ActivityNet-Captions together with 5× inference time acceleration over 3D TVG methods.

2. Related Work

Video Temporal Grounding (TVG). The objective of the TVG is to predict the starting/ending time points of target moments within an untrimmed video, which is described by a text sentence. Early TVG solutions [7, 14, 20, 39, 62, 64, 70] mainly employ two-stage “propose-and-rank” pipeline: ①

Propose: utilize sliding windows or proposal network to generate proposal candidates from the input video. ② Rank: the proposed candidates would be ranked according to the text query, and then the proposal with the highest ranking would be the final prediction decision. In contrast to proposal-based methods, regression-based methods [16, 67, 69] directly predict the starting/ending time points of the target moments without ranking massive proposal candidates. Thus, regression-based methods are much faster than proposal-based methods, which is one reason why our work focuses on the regression-based TVG. Furthermore, reinforcement learning (RL)-based methods formulate the TVG task as a sequence of decisions to make [18, 60]. In particular, they train an agent to control the movement of a window by shifting or scaling. During training, the agent would be rewarded or punished based on whether the window is close to the target moment after an adjustment.

Temporal Action Detection (TAD). TAD aims to determine whether predefined actions occur in a video and to predict the corresponding time intervals during which these actions occur [12, 13, 48, 53, 56, 59, 63]. Different from TVG, the input of TAD is only a video. In other words, TAD only requires a semantic understanding of videos. Compared to TAD, TVG is more challenging since it requires a semantic understanding of both videos and natural languages. Furthermore, TVG needs to process the multimodal interaction between videos and natural languages.

Text Prompting. Prompting has recently achieved great success in the domain of natural language processing [25, 32, 37, 40, 46, 49–52, 58]. Text prompting is a process that leverages a data-agnostic perturbation operation applied to text inputs or their embeddings to improve the performance of the downstream task. The simplest way is to construct an input context template originating from human contemplation [46, 49–51]. Although the manually-crafted context templates are simple and interpretable, they are typically not the optimal input prompts. To tackle this issue, other work has focused on searching the optimal prompting in the discrete input space [25, 52, 58] or in the language model’s embedding space [32, 37, 40].

Visual Prompting. Inspired by the idea of prompt learning in NLP [37], visual prompting (VP) was first proposed by Bahng *et. al.* [2] to reprogram a source vision model (*e.g.*, ImageNet-pretrained classifier) to accomplish downstream target tasks (*e.g.*, CIFAR-10 image classification). VP shares almost the same idea with the model reprogramming technology in the vision domain [4–6, 11, 57, 65, 72, 81], which incorporates a universal input perturbation into testing data so as to improve a desired performance metric, *e.g.*, target task accuracy, robustness, and fairness.

Multi-Modal Prompting. Although visual prompting and text prompting have recently attracted much attention, they

are under-explored in the multi-modal learning, especially on the temporal video grounding task. The existing works [2, 27, 66] mainly focus on integrating text and visual prompts with the CLIP (Contrastive Language–Image Pre-trained) model to improve downstream tasks with imagery data. The problem of multi-modal prompting in the video understanding task has not been studied. In this paper, we for the first time develop the text-visual prompting technique to improve the performance of temporal video grounding using 2D visual features.

3. Methods

In this section, we begin with the problem formulation of regression-based TVG. Then we demonstrate the design of TVP (text-visual prompts) and present the overview of our proposed TVP framework.

3.1. Problem Definition

Let $\mathbf{v} \in \mathbb{R}^{N_{\text{vid}} \times C \times H \times W}$ be an untrimmed video consisting of a sequence of N_{vid} video frames, and $\mathbf{s} \in \mathbb{R}^{N_{\text{tex}}}$ be a text query consisting of a sequence of N_{tex} language tokens. Here, the video-query pair (\mathbf{v}, \mathbf{s}) belongs to a video-language dataset \mathcal{D} . Given \mathbf{v} and \mathbf{s} , TVG aims to predict the time interval $\hat{\mathbf{T}} = (\hat{t}_{\text{sta}}, \hat{t}_{\text{end}})$ of the target video moments described by the query \mathbf{s} . The TVG model that fuses the vision-language modalities can be described as:

$$\hat{\mathbf{T}} = f(g_{\text{tex}}(\mathbf{s}), g_{\text{vid}}(\mathbf{v})), \quad (1)$$

where f denotes TVG model, and g_{vid} and g_{tex} represent vision encoder and language encoder, respectively.

3.2. TDIoU Loss Function

Conventionally, the TVG model can be learned by minimizing the **temporal IoU loss** $\mathcal{L}_{\text{tIoU}}$ defined below:

$$\mathcal{L}_{\text{tIoU}} = \left(1 - \frac{\hat{\mathbf{T}}(\boldsymbol{\theta}) \cap \mathbf{T}}{\hat{\mathbf{T}}(\boldsymbol{\theta}) \cup \mathbf{T}} \right), \quad (2)$$

where for ease of notation let $\boldsymbol{\theta}$ denote all the trainable parameters involved in (1), and $\mathbf{T} = (t_{\text{sta}}, t_{\text{end}})$ is the label (*i.e.*, the ground-truth time interval) of the target moment associated with the input video-query pair (\mathbf{v}, \mathbf{s}) . The rationale behind (2) is to maximize the overlapping between the predicted time interval and its ground truth.

However, for non-overlapping cases, the temporal IoU loss $\mathcal{L}_{\text{tIoU}}$ would encounter a gradient vanishing problem. Inspired by [82], we develop a novel **TDIoU** (Temporal-Distance IoU) loss for training our proposed TVG models by incorporating the normalized central time point distance and duration difference between the predicted video clips and the target video clips. We elaborate on the proposed loss below.

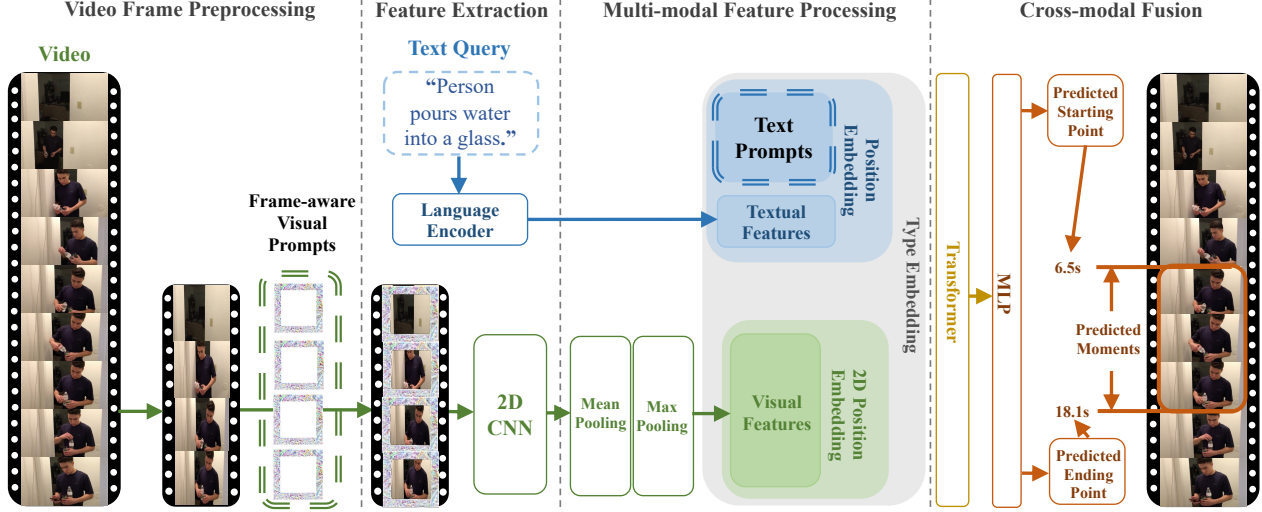


Figure 3. Overview of our proposed TVP (text-visual prompting) framework for 2D TVG (temporal video grounding). The whole process contains four phases: ❶ Video frame preprocessing: uniformly sample frames from input video and apply a set of frame-aware visual prompts to the sampled frames in order; ❷ Feature extraction: 2D CNN extracts features from sampled video frames with visual prompts, and the language encoder extracts textual features. In addition, the visual features would be spatially downsampled and temporally fused by max pooling and mean pooling, respectively. ❸ Multimodal feature processing: after spatial downsampling and temporal fusion, the 2D visual features would be integrated into the prompted textual features. ❹ Crossmodal fusion: the multimodal features would be processed by a 12-layer transformer encoder, and MLP would predict the starting/ending time points of the target moment.

Distance Loss \mathcal{L}_{dis} . To avoid the gradient vanishing problem caused by the non-overlapping case, we involve distance loss \mathcal{L}_{dis} to directly minimize the normalized central time point distance. In addition, we add a threshold α_1 to prevent oscillation in the later training phase. The distance loss is then given by:

$$\mathcal{L}_{\text{dis}} = \max \left(\frac{|(t_{\text{sta}} + t_{\text{end}})/2 - (\hat{t}_{\text{sta}} + \hat{t}_{\text{end}})/2|}{|\hat{\mathbf{T}} \cup \mathbf{T}|}, \alpha_1 \right), \quad (3)$$

where recall that $\mathbf{T} = (t_{\text{sta}}, t_{\text{end}})$, $\hat{\mathbf{T}}$ is predicted by the TVG model (1), and we choose $\alpha_1 = 0.2$ in experiments.

Duration Loss \mathcal{L}_{dur} . The introduction of distance loss \mathcal{L}_{dis} avoids the gradient vanishing problem but only considers the central time point distance. Yet, this may not be precise enough. For example, even if the central time points are completely overlapped, the duration of two video clips may not be identical. Inspired by the above, we propose the duration loss:

$$\mathcal{L}_{\text{dur}} = \max \left(\frac{|\mathbf{T} - \hat{\mathbf{T}}(\theta)|}{|\mathbf{T}|}, \alpha_2 \right), \quad (4)$$

where α_2 is the precision tolerance threshold and set by 0.4 in our experiments.

Finally, the proposed Temporal-Distance IoU (TDIoU) loss is given by

$$\mathcal{L} = \mathcal{L}_{\text{IoU}} + \beta_1 \mathcal{L}_{\text{dis}} + \beta_2 \mathcal{L}_{\text{dur}}, \quad (5)$$

where $\beta_1 > 0$ and $\beta_2 > 0$ are regularization parameters.

3.3. Text-Visual Prompt Design

Inspired by the application of prompts on transformers [2, 21, 36, 37], we propose jointly text-visual prompting to boost the performance of our models, in which prompts are optimized perturbation patterns. To improve data processing efficiency, we uniformly sample video frames from the untrimmed video \mathbf{v} to obtain $\mathbf{v}_{\text{sam}} \in \mathbb{R}^{N_{\text{sam}} \times C \times H \times W}$, where N_{sam} is the number of sampled video frames. In addition, we introduce a set of frame-aware visual prompts $\delta_{\text{vp}} \in \mathbb{R}^{N_{\text{sam}} \times d_{\text{vp}}}$ in the pixel space of sampled video frames \mathbf{v}_{sam} , and introduce text prompts $\delta_{\text{tp}} \in \mathbb{R}^{N_{\text{tp}} \times d_{\text{tp}}}$ in the textual feature space. By incorporating video frame sampling and text-visual prompts into the TVG model (1), we obtain:

$$(\hat{t}_{\text{sta}}, \hat{t}_{\text{end}}) = f(\delta_{\text{tp}}, g_{\text{tex}}(\mathbf{s}), g_{\text{vid}}(\mathbf{v}_{\text{sam}} + \delta_{\text{vp}})). \quad (6)$$

Given a pre-trained 2D TVG model f , the objective of text-visual prompting (TVP) is to learn a universal set of visual prompts δ_{vp} and text prompts δ_{tp} to be integrated into sampled video frames and textual features, respectively. Specially, a set of different visual prompts are applied to uniformly-sampled frames of one untrimmed video in order. During training, only the set of visual prompts and text prompts are updated through backpropagation. During fine-tuning, prompts are frozen, and the parameters of the TVG model and encoders are updated. During testing, the set of optimized visual prompts and the optimized text prompts are applied to all test-time video-query pairs.

3.4. Framework

Inspired by the success of transformers in vision-language tasks, we choose ClipBERT [31] as the base model for 2D TVG. Extended from ClipBERT, the input of our regression-based TVG model would be describable sentences and uniformly sampled frames of one untrimmed video as shown in Fig. 3. Then, the predicted starting and ending time points of the target video clip would be model outputs. As described in Algorithm 1, there are four phases of our proposed TVP framework: ① **Video frame preprocessing**: We obtain sparsely-sampled frames \mathbf{v}_{sam} from one input untrimmed video \mathbf{v} , and apply universal frame-aware visual prompts δ_{vp} on top of frames at the padding location. ② **Feature extraction**: 2D vision encoder (first 5 ConvBlock of ResNet-50) g_{vid} and language encoder (a trainable word embedding layer) g_{tex} would extract features from the prompted frames \mathbf{v}'_{sam} and textual inputs s , respectively. ③ **Multimodal feature processing**: Following the setting of Pixel-BERT [22], the 2D visual features \mathbf{Q}_{vid} are downsampled spatially by a 2×2 max-pooling layer and fused temporally by a mean-pooling layer. Then, text prompts δ_{tp} are integrated into textual features \mathbf{Q}_{tex} . In addition, trainable 2D visual position embeddings $\mathbf{M}_{2\text{D}}$ and textual position embeddings \mathbf{M}_{pos} are applied to the processed 2D visual features \mathbf{Q}'_{vid} and prompted textual features \mathbf{Q}'_{tex} , respectively [10, 31]. Afterwards, the processed and position-encoded 2D visual features $\mathbf{Q}''_{\text{vid}}$ are flattened and integrated into prompted and position-encoded textual features $\mathbf{Q}''_{\text{tex}}$. Moreover, type embeddings \mathbf{M}_{type} would be added to the integrated multimodal features \mathbf{Q}_{all} to indicate the source type of features. ④ **Crossmodal fusion**: A 12-layer transformer [10] is utilized for crossmodal fusion on \mathbf{Q}_{all} , and then multilayer perceptron (MLP) ending with sigmoid function is used as the prediction head to process the last-layer crossmodal representation \mathbf{Q}_{CM} of the transformer for generating the predicted starting/ending time points $(\hat{t}_{\text{sta}}, \hat{t}_{\text{enda}})$ of the target moments described by the text query input.

4. Experiments

In this section, we demonstrate the effectiveness of our proposed TVP framework on Charades-STA and ActivityNet Captions datasets.

4.1. Experiment Setup

Datasets. The evaluations are implemented on two standard benchmark datasets for TVG task, Charades-STA [14] and ActivityNet Captions [28]. Tab. 1 summarizes the details of both datasets. **Charades-STA** dataset contains 6,672 videos and 16,124 text queries in total. The average length of videos is 30.6s, and the average length of text query is 7.2 words. The average length of moments cor-

Algorithm 1 Overview of TVP framework

Input: vision encoder g_{vid} , language encoder g_{tex} , position embeddings \mathbf{M}_{pos} , 2D position embeddings $\mathbf{M}_{2\text{D}}$, type embeddings \mathbf{M}_{type} , transformer f , prediction head MLP , visual prompts δ_{vp} , text prompts δ_{tp}

Output: Predicted time interval $\hat{\mathbf{T}} = (\hat{t}_{\text{sta}}, \hat{t}_{\text{enda}})$

Phase ①: Video frame preprocessing

- 1: $\mathbf{v}_{\text{sam}} \leftarrow$ uniformly sample video frames from an untrimmed video \mathbf{v}
- 2: $\mathbf{v}'_{\text{sam}} \leftarrow$ apply visual prompts δ_{vp} to the sampled video frames \mathbf{v}_{sam}

Phase ②: Feature Extraction

- 3: $\mathbf{Q}_{\text{vid}} = g_{\text{vid}}(\mathbf{v}'_{\text{sam}}) \leftarrow$ extracting 2D visual features
- 4: $\mathbf{Q}_{\text{tex}} = g_{\text{tex}}(s) \leftarrow$ extracting textual features

Phase ③: Multimodal feature processing

- 5: $\mathbf{Q}'_{\text{vid}} \leftarrow$ apply spatial downsampling and temporal fusion to 2D visual features \mathbf{Q}_{vid}
- 6: $\mathbf{Q}'_{\text{tex}} \leftarrow$ apply text prompts δ_{tp} to textual features \mathbf{Q}_{tex}
- 7: $\mathbf{Q}''_{\text{vid}} \leftarrow$ add 2D visual position embeddings $\mathbf{M}_{2\text{D}}$ on the processed 2D visual features \mathbf{Q}'_{vid}
- 8: $\mathbf{Q}''_{\text{tex}} \leftarrow$ add position embeddings \mathbf{M}_{pos} to prompted textual features \mathbf{Q}'_{tex}
- 9: $\mathbf{Q}_{\text{all}} \leftarrow$ integrate the processed and position-encoded textual features $\mathbf{Q}''_{\text{tex}}$ and the processed and position-encoded 2D visual features $\mathbf{Q}''_{\text{vid}}$
- 10: $\mathbf{Q}_{\text{all}} + \mathbf{M}_{\text{type}} \leftarrow$ add type embeddings \mathbf{M}_{type} to the integrated multimodal features \mathbf{Q}_{all}

Phase ④: Crossmodal fusion

- 11: $\mathbf{Q}_{\text{CM}} = f(\mathbf{Q}_{\text{all}} + \mathbf{M}_{\text{type}}) \leftarrow$ implement crossmodal fusion through transformer f
 - 12: $(\hat{t}_{\text{sta}}, \hat{t}_{\text{enda}}) = MLP(\mathbf{Q}_{\text{CM}}) \leftarrow$ prediction head generates the predicted time interval according to crossmodal representation \mathbf{Q}_{CM}
-

Table 1. Statistics of TVG benchmark datasets (Charades-STA and ActivityNet Captions datasets).

Dataset	Charades-STA	ActivityNet Captions
Domain	Indoor Activity	Indoor/Outdoor Activity
# Videos	6,672	14,926
Avg. Video Length (second)	30.6	117.6
# Moments	11,767	71,953
Avg. Moment Length (second)	8.1	37.1
Vocabulary Size	1,303	15,505
# Queries	16,124	71,953
Avg. Query Length (word)	7.2	14.4

responding to the text query is 8.1s. Following the same dataset split as [14] for fair comparisons, there are 12,408 video-query pairs for training and 3,720 pairs for testing. **ActivityNet Captions** dataset contains 14,926 videos and 71,953 text queries in total. The average length of videos is 117.6s, and the average length of text query is 14.4 words.

The average length of moments corresponding to the text query is 37.1s. ActivityNet Captions dataset is split into training set, validation set, and testing set in a 2 : 1 : 1 ratio. Since the testing set is withheld for competition, only a training set and two validation sets (*vall* and *val2*) can be accessed publicly. For fair comparisons, we evaluate our proposed framework on *vall*.

Baselines. We compare our proposal with 15 baseline methods: ① **Proposal-based:** CTRL [14], MCN [1], SAP [7], BPNet [62], LPNet [61], QSPN [64], MAN [71]; ② **Proposal-free:** ABLR [67], DRN [69], CPNet [34], DEBUG [43], ExCL [16], VSLNet [73]; ③ **Reinforcement learning:** TSP-PRL [60], TripNet [18].

Evaluation metrics. Following [14], we adopt Acc(R@1, IoU=m) as the performance evaluation metric, which represents the percentage accuracy of top-1 predicted moments whose tIoU (temporal IoU) with the ground-truth moment is larger than m . By convention, we consider the following tIoU threshold values $m = \{0.3, 0.5, 0.7\}$.

Crossmodal pretraining setup. Our 2D vision encoder (ResNet-50) is initialized with the weight from grid-feat [24], which can extract effective grid features from visual inputs. In addition, both the language encoder and 12-layer transformer are initialized with the BERT-base model weight [10], which are pretrained on English Wikipedia and BookCorpus [83]. Thanks to the compact 2D vision encoder, TVP (our proposal) is able to directly utilize image-text pairs for end-to-end training. Since the benefits of cross-modal pretraining has been demonstrated by [22, 44, 55], our base model is pretrained on two large-scale image-text datasets, which are Visual Genome Captions [29] and COCO Captions [8]. To be more specific, image-text matching [44, 55] and masked language modeling [10] are employed for cross-modal pretraining.

Implementation setup. For video inputs, we uniformly sample N_{sam} frames from a video ($N_{\text{sam}} = 48$ for Charades-STA and $N_{\text{sam}} = 64$ for ActivityNet Captions). In addition, all video frames are resized to have a maximum longer side of 448 with an original aspect ratio, and then the frames are zero-padded to 448×448 . The default visual prompt sizes for both dataset are 96. The default text prompt sizes are 10 and 20 for Charades-STA and ActivityNet Captions, respectively. We utilize the first 5 ConvBlocks of ResNet-50 as the 2D vision encoder and a trainable embedding layer as the language encoder for both Charades-STA and ActivityNet Captions datasets. For text queries, all word tokens are maintained after lower-case conversion and tokenization. We use AdamW [42] for end-to-end model training, with $\beta_1 = 1.0$, $\beta_2 = 0.1$, $\alpha_1 = 0.2$, $\alpha_2 = 0.4$. Initial learning rates are $1e - 1$ and $5e - 7$ for prompt training and model finetuning, respectively. In addition, the learning rate linearly decays to 0 with the first 10% training step for warmup. Our experiments are imple-

Table 2. Performance comparison of different thresholds m on the Charades-STA dataset.

Type	Method	Visual Feature	Acc(R@1, IoU=m)		
			$m=0.3$	$m=0.5$	$m=0.7$
3D TVG	CTRL [14]	C3D	-	23.63	8.89
	ABLR [67]	C3D	-	24.36	9.01
	BPNet [62]	C3D	55.46	38.25	20.51
	LPNet [61]	C3D	59.14	40.94	21.13
	QSPN [64]	C3D	54.70	35.60	15.80
	TSP-PRL [60]	C3D	-	45.45	24.75
	TripNet [18]	C3D	54.64	38.29	16.07
	DRN [69]	C3D	-	45.40	26.40
	CPNet [34]	C3D	-	40.32	22.47
	DEBUG [43]	C3D	54.95	37.39	17.92
	ExCL [16]	I3D	61.50	44.1	22.40
	VSLNet [73]	I3D	64.30	47.31	30.19
	MAN [71]	I3D	-	46.53	22.72
2D TVG	MCN [1]	VGG	-	17.46	8.01
	SAP [7]	VGG	-	27.42	13.36
Ours					
TVP-Based 2D TVG	Base w/o prompts	ResNet	61.29	40.43	19.89
	Base + Visual Prompts		65.38	44.31	20.22
	Base + Text Prompts		65.81	43.44	20.65
	Base + Both Prompts		65.92	44.39	21.51

Table 3. Performance comparison of different thresholds m on the ActivityNet Captions dataset.

Type	Method	Visual Feature	Acc(R@1, IoU=m)		
			$m=0.3$	$m=0.5$	$m=0.7$
3D TVG	CTRL [14]	C3D	28.70	14.00	-
	BPNet [62]	C3D	59.98	42.07	24.69
	LPNet [61]	C3D	64.29	45.92	25.39
	QSPN [64]	C3D	45.30	27.70	13.60
	TSP-PRL [60]	C3D	56.02	38.83	-
	TripNet [18]	C3D	48.42	32.19	13.93
	DRN [69]	C3D	-	45.45	24.36
	CPNet [34]	C3D	-	40.56	21.63
	ABLR [67]	C3D	55.67	36.79	-
	DEBUG [43]	C3D	55.91	39.72	-
	ExCL [16]	C3D	63.00	43.60	24.10
	VSLNet [73]	C3D	63.16	43.22	26.16
	Ours				
TVP-Based 2D TVG	Base w/o prompts	ResNet	57.20	40.16	19.14
	Base + Visual Prompts		60.12	43.39	23.71
	Base + Text Prompts		60.48	42.58	24.39
	Base + Both Prompts		60.71	43.44	25.03

mented in PyTorch [45], and models and prompts are finetuned separately for 12 epochs with the mixed precision on 8 NVIDIA V100 GPUs.

4.2. Experiment Results

Effectiveness of TVP on Charades-STA. The performance comparisons with SOTA methods on the Charades-STA dataset are summarized in **Tab. 2**. Our proposed TVP framework can achieve competitive performance at all tIoU thresholds m in the case of utilizing 2D visual features extracted by ResNet-50, and reach the highest score at $m = 0.3$. Compared to the 2D TVG methods using VGG as the vision encoder, our proposed framework could achieve around $2.5\times$ and $2.7\times$ performance gain at thresholds 0.5 and 0.7, respectively. Furthermore, we can find that for our base model only one of visual prompts and text prompts can achieve up to 7.37% and 9.60% improvement

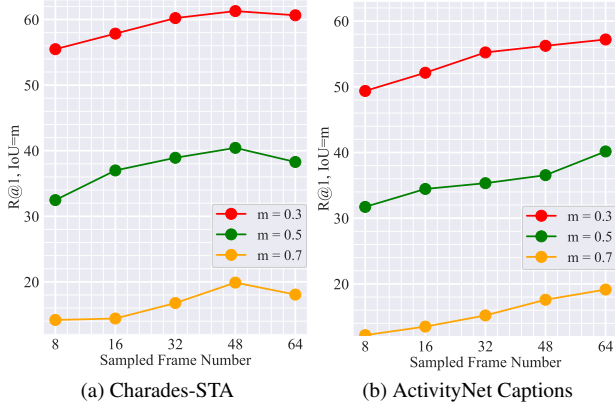


Figure 4. Impact of sampled frame numbers.

at tIoU thresholds $m = 0.3$ and $m = 0.5$. The combination of text and visual prompts can not only achieves 7.55% and 9.79% improvements at tIoU thresholds $m = 0.3$ and $m = 0.5$, but also improve the performance by 8.14% at $m = 0.7$. This demonstrates the effectiveness and necessity of the joint text-visual prompting.

Effectiveness of TVP on ActivityNet Captions. We focus on the performance comparisons with 3D TVG methods on ActivityNet since there are no results of 2D TVG method reported on ActivityNet Captions. The results of multiple methods on ActivityNet Captions datasets are reported in **Tab. 3**. Even on this more challenging dataset, our proposed method still has achieved competitive performance compared to 3D TVG methods. Different from the performance of TVP on Charades-STA dataset, text prompts or visual prompts can achieve a significant performance boost on the base model over all IoU thresholds m alone (5.73% at $m = 0.3$, 8.04% at $m = 0.5$, 27.43% at $m = 0.7$), and the text-visual prompt combination could further boost the performance (6.14% at $m = 0.3$, 8.17% at $m = 0.5$, 30.77% at $m = 0.7$). It is worth noting that the performance gap over $m = 0.7$ between 2D TVG methods and 3D TVG methods is narrowed significantly.

In summary, through the experimental results on Charades-STA and ActivityNet Captions datasets, we can find that our proposed TVP framework could achieve competitive performance overall tIoU thresholds on Charades-STA and ActivityNet Captions by improving the utility of sparse 2D visual features. Thanks to the lightweight 2D vision encoder, cotraining language encoder and vision encoder on large-scale image-text datasets can be performed, which benefits the base model to achieve good performance. Furthermore, the combination of text and visual prompts can achieve better results than any single kind of prompts on both datasets, which again proves the importance of cross-modal training.

Video frame sampling effect. **Fig. 4** demonstrates the performance of base model with different number N_{sam}

Table 4. The performance comparison of different visual prompt sizes on Charades-STA dataset.

Visual Prompt Size	Acc(R@1, IoU=m)			Prompt + Frame
	$m=0.3$	$m=0.5$	$m=0.7$	
0	61.29	40.43	19.89	
16	61.29	40.43	20.00	
32	61.94	39.78	19.35	
48	63.66	42.37	20.00	
72	63.87	43.66	19.78	
96	65.38	44.31	20.22	
128	64.73	43.66	19.78	

Table 5. The performance comparison of different text prompt sizes on Charades-STA dataset.

Text Prompt Size	Acc(R@1, IoU=m)		
	$m=0.3$	$m=0.5$	$m=0.7$
0	57.20	40.16	19.14
5	65.38	41.94	20.43
10	65.81	43.44	20.65
15	65.59	43.23	21.29
20	64.95	43.87	21.51
25	63.66	42.80	20.65
30	64.46	42.63	20.51

of sampled video frames as visual inputs. For Charades dataset, the base model performance keeps increasing before N_{sam} reaches 48, but when it exceeds 48, performance starts to degrade. This is because frequent background changes harm the performance of object re-identification in videos, which are noisy for object motion analysis [17].

For ActivityNet Caption dataset, base model performance continues to improve even when sampled frame number N_{sam} exceeds 48, due to the longer average video length in ActivityNet Captions dataset. Balancing the frame number and batch size for training, we choose $N_{\text{sam}} = 64$ for ActivityNet Captions.

TVP performance vs. prompt size. As shown in **Tab. 4**, we can find that when visual prompts are small, they cannot bring changes to the base model, and when visual prompts are too large, the performance starts to decrease. This is because key information within video frames might be removed. However, the text prompts can bring significant performance boost even when the text prompt size is small as shown in **Tab. 5**, which is because the textual features has a smaller dimension compared to visual features, and also the text prompts are directly optimized in feature space during training.

TVP performance vs. visual prompt operation. Visual prompt is first proposed by [2], where visual prompts are

Table 6. The performance comparison of different visual prompt operations (*remove*, *add*, *replace*) with fixed visual prompt size $p = 96$ on Charades-STA and ActivityNet Captions datasets.

Operation	Charades-STA			ActivityNet Captions		
	R@1, IoU= m			R@1, IoU= m		
	$m=0.3$	$m=0.5$	$m=0.7$	$m=0.3$	$m=0.5$	$m=0.7$
Original	61.29	40.43	19.89	57.20	40.16	19.14
Remove	61.29	40.43	20.0	57.20	40.16	19.14
Add	61.08	39.57	20.22	57.15	40.16	19.27
Replace	65.38	44.31	20.22	60.12	43.39	23.71




Figure 5. Inference time comparison. (a) inference time comparison between 2D vision encoder (ResNet-50) and 3D vision encoder (C3D). (b) inference time comparison between the vision encoder and the other modules of the 2D TVG model, where the sampled frame number for our TVP framework is $1.2 \times$ the length of the video in seconds.

added to the image for transfer learning on classification tasks. In contrast, our proposed prompting framework is designed to compensate for the spatiotemporal information loss in 2D visual features. Due to the differences in the task, we try two different prompt operation strategies, *replace* and *add*. *add* is to add the visual prompts to the pixel value of the video frame at the corresponding padding locations. *replace* is to replace the pixel values of video frames with visual prompts at corresponding padding locations. *remove* is in order to study the impact of removing the pixel values at the padding location. As shown in **Tab. 6**, *add* or *remove* prompt operations have limited effects on the base model. However, *replace* does boost the base model performance.

TVP achieves inference efficiency. As shown in **Fig. 5**, we can find that the inference time required for visual feature extraction accounts for more than half of the inference time of the whole model, while the inference time required for the 3D vision encoder is more than $5 \times$ compared to the 2D vision encoder, and even more than the time required for the whole TVG model using 2D vision encoder, which fully demonstrates the feasibility of accelerating the overall inference speed by reducing the complexity of the vision encoder. Need to note that if there are

Table 7. The performance comparison of different loss designs on Charades-STA dataset.

Loss Function Selection	R@1, IoU= m		
	$m=0.3$	$m=0.5$	$m=0.7$
\mathcal{L}_{tIoU}	55.05	29.89	11.82
$\mathcal{L}_{tIoU} + \mathcal{L}_{dis}$	60.64	31.18	16.77
$\mathcal{L}_{tIoU} + \mathcal{L}_{dur}$	59.78	30.97	16.34
$\mathcal{L}_{tIoU} + \mathcal{L}_{dis} + \mathcal{L}_{dur}$	61.29	40.43	19.89




Figure 6. Loss landscape visualization in 2D plane: Finetuning w/o prompts (left) and using prompts (right); see [33] for implementation.

multiple model weights for different sampled frame number settings and model weights can be adopted adaptively for different lengths of videos, the inference speed for short videos should increase, and the prediction results for long videos will be further improved.

Ablation studies. Through **Tab. 7**, we can find that the addition of either distance loss \mathcal{L}_{dis} or duration loss \mathcal{L}_{dur} will result in a performance increase, but the combination of the two will result in a significant performance increase (11.34% at $m = 0.3$, 35.26% at $m = 0.5$, 68.27% at $m = 0.7$), especially over tIoU thresholds $m = 0.5$ and $m = 0.7$. This demonstrates that distance loss \mathcal{L}_{dis} and duration loss \mathcal{L}_{dur} could provide more precise training guides compared to only using temporal IoU loss \mathcal{L}_{tIoU} . Furthermore, we posit that prompting may encode additional spatial-temporal supervision to help the model trainer to escape from bad local optima as shown in **Fig. 6**, where fine-tuning w/ prompts yields a flatter loss landscape than the one w/o prompts.

5. Conclusion

In this paper, we propose text-visual prompting to boost the performance of 2D TVG methods by compensating for the lack of spatiotemporal information in 2D visual features. In contrast to 3D TVG methods, TVP allows us to effectively co-train vision encoder and language encoder in a 2D TVG model and improves the performance of cross-modal feature fusion using only low-complexity sparse 2D visual features. The effectiveness of our proposed TVP (text-visual prompting) framework has been demonstrated on two standard datasets, Charades-STA and ActivityNet. Our models outperform all 2D models significantly, and also achieve comparable performance to 3D models. What is more, we achieve over $5 \times$ inference speedup over TVG methods of using 3D visual features.

References

- [1] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *Proceedings of the IEEE international conference on computer vision*, pages 5803–5812, 2017. 1, 6
- [2] H Bahng, A Jahanian, S Sankaranarayanan, and P Isola. Exploring visual prompts for adapting large-scale models. *arXiv preprint arXiv:2203.17274*, page 2022, 2022. 2, 3, 4, 7
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 1
- [4] Aochuan Chen, Peter Lorenz, Yuguang Yao, Pin-Yu Chen, and Sijia Liu. Visual prompting for adversarial robustness. *arXiv preprint arXiv:2210.06284*, 2022. 3
- [5] Aochuan Chen, Yuguang Yao, Pin-Yu Chen, Yihua Zhang, and Sijia Liu. Understanding and improving visual prompting: A label-mapping perspective. *arXiv preprint arXiv:2211.11635*, 2022. 3
- [6] Pin-Yu Chen. Model reprogramming: Resource-efficient cross-domain machine learning. *arXiv:2202.10629*, 2022. 3
- [7] Shaoxiang Chen and Yu-Gang Jiang. Semantic proposal for activity localization in videos via sentence query. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8199–8206, 2019. 1, 2, 6
- [8] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 6
- [9] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020. 2
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 5, 6
- [11] Gamaleldin F Elsayed, Ian Goodfellow, et al. Adversarial reprogramming of neural networks. *arXiv:1806.11146*, 2018. 3
- [12] Christoph Feichtenhofer. X3d: Expanding architectures for efficient video recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 3
- [13] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019. 3
- [14] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. Tall: Temporal activity localization via language query. In *Proceedings of the IEEE international conference on computer vision*, pages 5267–5275, 2017. 1, 2, 5, 6
- [15] Yanjun Gao, Lulu Liu, Jason Wang, Xin Chen, Huayan Wang, and Rui Zhang. Evoquer: Enhancing temporal grounding with video-pivoted backquery generation. *arXiv preprint arXiv:2109.04600*, 2021. 1
- [16] Soham Ghosh, Anuva Agarwal, Zarana Parekh, and Alexander Hauptmann. Excl: Extractive clip localization using natural language descriptions. *arXiv preprint arXiv:1904.02755*, 2019. 1, 3, 6
- [17] Xinqian Gu, Hong Chang, Bingpeng Ma, Hongkai Zhang, and Xilin Chen. Appearance-preserving 3d convolution for video-based person re-identification. In *ECCV*, 2020. 7
- [18] Meera Hahn, Asim Kadav, James M Rehg, and Hans Peter Graf. Tripping through time: Efficient localization of activities in videos. *arXiv preprint arXiv:1904.09936*, 2019. 1, 3, 6
- [19] Xiaochen Han, Bo Wu, Zheng Shou, Xiao-Yang Liu, Yimeng Zhang, and Linghe Kong. Tensor fista-net for real-time snapshot compressive imaging. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10933–10940, 2020. 1
- [20] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with temporal language. *arXiv preprint arXiv:1809.01337*, 2018. 2
- [21] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019. 2, 4
- [22] Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. *arXiv preprint arXiv:2004.00849*, 2020. 2, 5, 6
- [23] Jinghan Jia, Mingyi Hong, Yimeng Zhang, Mehmet Akçakaya, and Sijia Liu. On the robustness of deep learning-based mri reconstruction to image transformations. *arXiv preprint arXiv:2211.04930*, 2022. 1
- [24] Huaizu Jiang, Ishan Misra, Marcus Rohrbach, Erik Learned-Miller, and Xinlei Chen. In defense of grid features for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10267–10276, 2020. 6
- [25] Zhengbao Jiang, Frank F Xu, Jun Araki, and Graham Neubig. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438, 2020. 2, 3
- [26] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. *arXiv preprint arXiv:2112.04478*, 2021. 2
- [27] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. *arXiv preprint arXiv:2210.03117*, 2022. 2, 3
- [28] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *Proceedings of the IEEE international conference on computer vision*, pages 706–715, 2017. 5

- [29] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. [6](#)
- [30] Xiaohan Lan, Yitian Yuan, Xin Wang, Zhi Wang, and Wenwu Zhu. A survey on temporal sentence grounding in videos. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 2021. [1](#)
- [31] Jie Lei, Linjie Li, Luowei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learning via sparse sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7331–7341, 2021. [5](#)
- [32] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. [2](#), [3](#)
- [33] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. *Advances in neural information processing systems*, 2018. [8](#)
- [34] Kun Li, Dan Guo, and Meng Wang. Proposal-free video grounding with contextual pyramid network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1902–1910, 2021. [1](#), [2](#), [6](#)
- [35] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. [2](#)
- [36] Muheng Li, Lei Chen, Yueqi Duan, Zhilan Hu, Jianjiang Feng, Jie Zhou, and Jiwen Lu. Bridge-prompt: Towards ordinal action understanding in instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. [4](#)
- [37] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021. [2](#), [3](#), [4](#)
- [38] Yanyu Li, Pu Zhao, Geng Yuan, Xue Lin, Yanzhi Wang, and Xin Chen. Pruning-as-search: Efficient neural architecture search via channel pruning and structural reparameterization. *arXiv preprint arXiv:2206.01198*, 2022. [1](#)
- [39] Meng Liu, Xiang Wang, Liqiang Nie, Qi Tian, Baoquan Chen, and Tat-Seng Chua. Cross-modal moment localization in videos. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 843–851, 2018. [2](#)
- [40] Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. Gpt understands, too. *arXiv preprint arXiv:2103.10385*, 2021. [2](#), [3](#)
- [41] Xiao-Yang Liu, Yimeng Zhang, Yukang Liao, and Ling Jiang. Dynamic updating of the knowledge base for a large-scale question answering system. *ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP)*, 2020. [1](#)
- [42] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [6](#)
- [43] Chujie Lu, Long Chen, Chilie Tan, Xiaolin Li, and Jun Xiao. Debug: A dense bottom-up grounding approach for natural language video localization. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5144–5153, 2019. [1](#), [6](#)
- [44] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32, 2019. [2](#), [6](#)
- [45] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 2019. [6](#)
- [46] Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*, 2019. [3](#)
- [47] Di Qi, Lin Su, Jia Song, Edward Cui, Taroon Bharti, and Arun Sacheti. Imagebert: Cross-modal pre-training with large-scale weak-supervised image-text data. *arXiv preprint arXiv:2001.07966*, 2020. [2](#)
- [48] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatiotemporal representation with pseudo-3d residual networks. In *proceedings of the IEEE International Conference on Computer Vision*, 2017. [3](#)
- [49] Timo Schick and Hinrich Schütze. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676*, 2020. [3](#)
- [50] Timo Schick and Hinrich Schütze. Few-shot text generation with pattern-exploiting training. *arXiv preprint arXiv:2012.11926*, 2020. [3](#)
- [51] Timo Schick and Hinrich Schütze. It’s not just size that matters: Small language models are also few-shot learners. *arXiv preprint arXiv:2009.07118*, 2020. [3](#)
- [52] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980*, 2020. [3](#)
- [53] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *Advances in neural information processing systems*, 2014. [3](#)
- [54] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visiolinguistic representations. *arXiv preprint arXiv:1908.08530*, 2019. [2](#)
- [55] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*, 2019. [2](#), [6](#)
- [56] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. [1](#), [3](#)
- [57] Yun-Yun Tsai et al. Transfer learning without knowing: Reprogramming black-box machine learning models with scarce data and limited resources. *arXiv:2007.08714*, 2020. [3](#)

- [58] Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. Universal adversarial triggers for attacking and analyzing nlp. *arXiv preprint arXiv:1908.07125*, 2019. [3](#)
- [59] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018. [3](#)
- [60] Jie Wu, Guanbin Li, Si Liu, and Liang Lin. Tree-structured policy based progressive reinforcement learning for temporally language grounding in video. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12386–12393, 2020. [1](#), [2](#), [3](#), [6](#)
- [61] Shaoning Xiao, Long Chen, Jian Shao, Yueting Zhuang, and Jun Xiao. Natural language video localization with learnable moment proposals. *arXiv preprint arXiv:2109.10678*, 2021. [1](#), [2](#), [6](#)
- [62] Shaoning Xiao, Long Chen, Songyang Zhang, Wei Ji, Jian Shao, Lu Ye, and Jun Xiao. Boundary proposal network for two-stage natural language video localization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2986–2994, 2021. [1](#), [2](#), [6](#)
- [63] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *Proceedings of the European conference on computer vision (ECCV)*, 2018. [3](#)
- [64] Huijuan Xu, Kun He, Bryan A Plummer, Leonid Sigal, Stan Sclaroff, and Kate Saenko. Multilevel language and vision integration for text-to-clip retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9062–9069, 2019. [1](#), [2](#), [6](#)
- [65] Chao-Han Huck Yang, Yun-Yun Tsai, et al. Voice2series: Reprogramming acoustic models for time series classification. In *ICML*. PMLR, 2021. [3](#)
- [66] Yuan Yao, Ao Zhang, Zhengyan Zhang, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. Cpt: Colorful prompt tuning for pre-trained vision-language models. *arXiv preprint arXiv:2109.11797*, 2021. [3](#)
- [67] Yitian Yuan, Tao Mei, and Wenwu Zhu. 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, pages 9159–9166, 2019. [1](#), [3](#), [6](#)
- [68] Ofir Zafrir, Ariel Larey, Guy Boudoukh, Haihao Shen, and Moshe Wasserblat. Prune once for all: Sparse pre-trained language models. *arXiv preprint arXiv:2111.05754*, 2021. [1](#)
- [69] Runhao Zeng, Haoming Xu, Wenbing Huang, Peihao Chen, Mingkui Tan, and Chuang Gan. Dense regression network for video grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10287–10296, 2020. [1](#), [2](#), [3](#), [6](#)
- [70] Yawen Zeng, Da Cao, Xiaochi Wei, Meng Liu, Zhou Zhao, and Zheng Qin. Multi-modal relational graph for cross-modal video moment retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2215–2224, 2021. [2](#)
- [71] Da Zhang, Xiyang Dai, Xin Wang, Yuan-Fang Wang, and Larry S Davis. Man: Moment alignment network for natural language moment retrieval via iterative graph adjustment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1247–1257, 2019. [1](#), [6](#)
- [72] Guanhua Zhang, Yihua Zhang, Yang Zhang, Wenqi Fan, Qing Li, Sijia Liu, and Shiyu Chang. Fairness reprogramming. *arXiv preprint arXiv:2209.10222*, 2022. [3](#)
- [73] Hao Zhang, Aixin Sun, Wei Jing, and Joey Tianyi Zhou. Span-based localizing network for natural language video localization. *arXiv preprint arXiv:2004.13931*, 2020. [1](#), [6](#)
- [74] Hao Zhang, Aixin Sun, Wei Jing, and Joey Tianyi Zhou. The elements of temporal sentence grounding in videos: A survey and future directions. *arXiv preprint arXiv:2201.08071*, 2022. [1](#)
- [75] Songyang Zhang, Houwen Peng, Jianlong Fu, and Jiebo Luo. Learning 2d temporal adjacent networks for moment localization with natural language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12870–12877, 2020. [1](#), [2](#)
- [76] Yimeng Zhang, Akshay Karkal Kamath, Qiucheng Wu, Zhiwen Fan, Wuyang Chen, Zhangyang Wang, Shiyu Chang, Sijia Liu, and Cong Hao. Data-model-circuit tri-design for ultra-light video intelligence on edge devices. In *Proceedings of the 28th Asia and South Pacific Design Automation Conference*, pages 745–750, 2023. [1](#)
- [77] Yimeng Zhang, Xiao-Yang Liu, Bo Wu, and Anwar Walid. Video synthesis via transform-based tensor neural network. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2454–2462, 2020. [1](#)
- [78] Yimeng Zhang, Yuguang Yao, Jinghan Jia, Jinfeng Yi, Mingyi Hong, Shiyu Chang, and Sijia Liu. How to robustify black-box ml models? a zeroth-order optimization perspective. *arXiv preprint arXiv:2203.14195*, 2022. [1](#)
- [79] Yihua Zhang, Yuguang Yao, Parikshit Ram, Pu Zhao, Tianlong Chen, Mingyi Hong, Yanzhi Wang, and Sijia Liu. Advancing model pruning via bi-level optimization. *arXiv preprint arXiv:2210.04092*, 2022. [1](#)
- [80] Yihua Zhang, Guanhua Zhang, Prashant Khanduri, Mingyi Hong, Shiyu Chang, and Sijia Liu. Revisiting and advancing fast adversarial training through the lens of bi-level optimization. In *International Conference on Machine Learning*, pages 26693–26712. PMLR, 2022. [1](#)
- [81] Yang Zheng, Xiaoyi Feng, et al. Why adversarial reprogramming works, when it fails, and how to tell the difference. *arXiv:2108.11673*, 2021. [3](#)
- [82] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. Distance-iou loss: Faster and better learning for bounding box regression. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 12993–13000, 2020. [3](#)
- [83] Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*, pages 19–27, 2015. [6](#)