

STOP: Integrated Spatial-Temporal Dynamic Prompting for Video Understanding

Zichen Liu¹, Kunlun Xu¹, Bing Su², Xu Zou³, Yuxin Peng¹, Jiahuan Zhou^{1*}

¹ Wangxuan Institute of Computer Technology, Peking University, Beijing, China

² Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China

³ School of Artificial Intelligence and Automation, Huazhong University of Science and Technology, Wuhan, China

{lzc20180720, xkl}@stu.pku.edu.cn, subingats@gmail.com, zoux@hust.edu.cn

{pengyuxin, jiahuanzhou}@pku.edu.cn

Abstract

Pre-trained on tremendous image-text pairs, vision-language models like CLIP have demonstrated promising zero-shot generalization across numerous image-based tasks. However, extending these capabilities to video tasks remains challenging due to limited labeled video data and high training costs. Recent video prompting methods attempt to adapt CLIP for video tasks by introducing learnable prompts, but they typically rely on a single static prompt for all video sequences, overlooking the diverse temporal dynamics and spatial variations that exist across frames. This limitation significantly hinders the model's ability to capture essential temporal information for effective video understanding. To address this, we propose an integrated **S**patial-**T**emporal dynamic **P**rompting (**STOP**) model which consists of two complementary modules, the intra-frame spatial prompting and inter-frame temporal prompting. Our intra-frame spatial prompts are designed to adaptively highlight discriminative regions within each frame by leveraging intra-frame attention and temporal variation, allowing the model to focus on areas with substantial temporal dynamics and capture fine-grained spatial details. Additionally, to highlight the varying importance of frames for video understanding, we further introduce inter-frame temporal prompts, dynamically inserting prompts between frames with high temporal variance as measured by frame similarity. This enables the model to prioritize key frames and enhances its capacity to understand temporal dependencies across sequences. Extensive experiments on various video benchmarks demonstrate that **STOP** consistently achieves superior performance against state-of-the-art methods. The code is available at <https://github.com/zhoushijiahuan1991/CVPR2025-STOP>.

*Corresponding author




Figure 1. Existing video prompting methods [19, 61] typically add a static prompt across diverse videos, limiting the model's ability to capture essential temporal information. In contrast, our approach introduces dynamic intra-frame spatial prompts and inter-frame temporal prompts, guiding the model to focus on discriminative regions and key frames with significant temporal dynamics.

1. Introduction

With the advancement of deep learning [12, 20, 45, 57–60, 62], the data required for model training has been steadily increasing. Recently, vision-language models such as CLIP, which are pre-trained on large-scale image-text pairs by contrastive learning, have demonstrated remarkable zero-shot generalization capabilities in downstream tasks such

as image classification and image-text retrieval [45, 65, 66]. However, applying this pre-training and zero-shot generalization paradigm to video tasks poses unique challenges. Compared to images, labeled video data is more limited, which makes it challenging to collect large-scale video-text pairs [42, 64]. Additionally, training large-scale video-language models through contrastive learning with video-text pairs entails high computational costs, making it challenging to apply training paradigms similar to CLIP for video-language tasks [19, 61].

To address this issue, recent studies [19, 61, 64] have focused on efficiently fine-tuning large-scale pre-trained vision-language models on video data, enabling them to handle downstream video understanding tasks at a lower cost. Among these methods, video prompting, or video prompt learning, has emerged as an important category of approaches. It adapts vision-language models to video tasks by introducing learnable video prompts while keeping the model’s pre-trained parameters fixed. However, as shown in Figure. 1, existing video prompting methods [19, 61, 64] learn a single static prompt for all data, overlooking the differences in keyframes with significant temporal dynamics and the varying discriminative regions within each frame. This results in inaccuracies in the video frames and regions that the pre-trained vision-language model focuses on. As a consequence, the model’s ability to capture temporal information is limited, which ultimately hinders its capacity to effectively understand video content.

To address this issue, we propose an integrated **Spatial-Temporal dynamic Prompting** approach, named **STOP**. Regions with temporal dynamics are essential for understanding video actions, yet image-text pairs pre-trained CLIP models often struggle to attend effectively to these areas. To overcome this limitation, we introduce the Intra-frame Spatial Prompting. Our approach incorporates a lightweight 3D convolutional network to capture the temporal dynamics of different regions of the video. By combining this information with intra-frame attention weights, we identify discriminative regions within videos that include both primary objects in single frames and dynamic temporal information. Building on these insights, we design an intra-prompt prompter to generate spatial prompts for these regions, guiding the model to focus on areas with significant temporal changes and thereby enhancing its ability to capture fine-grained critical information in video data.

Furthermore, Considering that the dynamic changes vary across frames, affecting their importance for video understanding, we propose inter-frame temporal prompting to help the pre-trained model focus on keyframes. Building on the intra-frame spatial prompting that identifies discriminative regions within video frames, we further calculate the degree of change in these regions between frames. For keyframes with significant temporal dynamics, we dy-

namically generate inter-frame prompts using a lightweight prompter and insert them to provide fine-grained information between two frames. The intra-frame spatial prompts and inter-frame temporal prompts complement each other, guiding the model to focus on critical spatial and temporal locations, thereby enhancing the model’s ability to accurately understand videos, leading to improved performance.

To sum up, the main contributions of this work are: (1) To address the challenge of pre-trained vision-language models struggling to capture temporal information in videos, we propose the STOP method. First, we design intra-frame spatial prompting to highlight discriminative regions in video frames, effectively guiding the model to focus on dynamically changing areas. (2) Furthermore, we calculate the dynamic changes between video frames in the video and dynamically generate inter-frame temporal prompts. These prompts are inserted between frames with significant dynamic changes, providing fine-grained temporal information to facilitate the model focusing on and understanding the keyframes in the video. (3) Extensive experiments on various benchmarks for video action recognition and video-text retrieval demonstrate that our proposed STOP method consistently and significantly outperforms existing video prompting methods.

2. Related Work

2.1. Vision-Language Pre-training

Vision-language pre-training (VLP) aims to learn joint representations of vision and language, achieving strong performance across various downstream tasks [5, 9, 43, 45]. Recent advancements in contrastive image-text pre-training, such as CLIP [45], leveraging large-scale image-text pairs from the internet for training. Works like Flamingo [1] and ALBEF [29] further demonstrate the effectiveness of pre-trained models. For video-language pre-training, models like CLIPBERT [27] and Frozen in Time [3] have shown promise, yet challenges persist, including limited video datasets and high computational costs. To address these, methods like CLIP4Clip [42] and X-Pool [15] transfer image-text pre-trained models to video tasks, though they still require fine-tuning all model parameters, resulting in high overhead.

2.2. Video Prompting

Prompt learning, originating in NLP, aims to adapt pre-trained language models to various tasks [28, 32, 35], by introducing learnable tokens as prompts, optimizing only these tokens during training. Inspired by prompt learning’s success in NLP, this approach extends to vision-language models (VLMs) [24, 25, 41, 65, 66] and visual models [21, 31, 39, 40]. For instance, CoOp [66] uses learnable text prompts for improved classification performance,




Figure 2. The pipeline of our STOP. For each video, we begin by embedding it into image tokens. Then the intra-frame spatial prompting is introduced to locate discriminative regions and add generated prompts to these areas. Based on the degree of inter-frame variation, we dynamically generate inter-frame prompts and insert them as needed. Finally, these prompts along with the CLS token, pass through the MSA (multi-head self-attention) block to obtain a video representation, which is then used to compute similarity with the text features.

while VPT [21] introduces token-level prompts in the vision branch to capture image attributes.

Recently, prompt learning has been expanded to video tasks, like video understanding and text-video retrieval [13, 19, 61, 64]. For instance, VoP [19] innovates with visual prompts to capture spatiotemporal video characteristics, and DGL [61] enhances video-text interaction through global-local prompt coordination. Similarly, MPT [64] refines prompts to mine more detailed modality-specific information, boosting retrieval performance. However, these methods apply static prompts to all videos, neglecting video-specific frame details and dynamic changes, which limits the model’s ability to capture temporal nuances and constrains overall performance.

2.3. Video Understanding

Video action recognition and video-text retrieval are two main research areas related to video understanding. For video action recognition, researchers have focused on advancing temporal modeling efficiency. Optical flow methods, often used in two-stream fusion, are effective but computationally demanding [46, 52]. An alternative approach, 3D convolutions, extends 2D convolutions into the spatiotemporal domain but also suffers from high complex-

ity [8, 48, 49]. Other strategies embed plug-and-play temporal modules into 2D networks or adapt LSTMs for sequence analysis [33, 37]. Recently, Vision Transformer-based architectures have shown promising results in temporal modeling and feature extraction efficiency [4, 30, 38]. In addition to single-modality methods, multi-modal techniques leverage CLIP’s capabilities for tasks like video classification and retrieval [22, 42, 53].

For video-text retrieval task, early multi-modal researches [15, 36] focused on feature fusion for retrieval enhancement, while more recent studies [10, 55] emphasize cross-modal alignment at levels such as events, actions, and entities. Inspired by large-scale pre-trained models like BERT [23] and CLIP [45], contemporary methods [34] increasingly employ contrastive learning or masked language modeling to develop a shared video-text representation space. In our work, we adapt the frozen CLIP model with learnable prompts, effectively transferring its strengths to video retrieval, as validated by our empirical results.

3. Method

In this section, we illustrate the proposed STOP in detail, and the overall pipeline is depicted in Figure 2.

3.1. Notations

The backbone of our method is CLIP [45], with CLIP4clip [42] as the baseline. It mainly consists of two branches: the text encoder $\mathbf{E}_T$ and the video encoder $\mathbf{E}_V$. The pre-trained text encoder $\mathbf{E}_T$ is a transformer model [50]. For the video-text retrieval task, the input of the text encoder $\mathbf{E}_T$ is a natural language sentence $\mathbf{S}$. For the action recognition task, given a video's category text description $C = \{c_1, c_2, \dots, c_K\}$, where K denotes the number of categories, we construct a sentence using a hand-crafted prompt following [51], such as $\mathbf{S} = \text{"A video of the action [CLS]"}.$ Then, we use the text encoder $\mathbf{E}_T$ to obtain the representation $\mathbf{s} = \mathbf{E}_T(\mathbf{S})$ of the sentence $\mathbf{S}$, where $\mathbf{s} \in \mathbb{R}^d$ and d is the dimension of representation.

The image encoder in CLIP and CLIP4clip is based on the Vision Transformer (ViT) [14]. The input is a video $\mathbf{V} \in \mathbb{R}^{N_F \times 3 \times H \times W}$, where N_F represents the number of frames and $H \times W$ is the spatial size. Each video frame $\{\mathbf{F}_i\}_{i=1}^{N_F}$ is split into $N_p = \frac{H \times W}{h \times w}$ fixed-size patches of size $h \times w$, and these patches are flattened into a set of vectors $\mathbf{x}_i = \{\mathbf{x}_{i,j} \in \mathbb{R}^{h \times w}\}_{j=1}^{N_p}$, where i denotes the frame index and j denotes the patch index. These vectors are then projected into patch embeddings $\mathbf{h}_i = \{\mathbf{h}_{i,j}\}_{j=1}^{N_p}$, where $\mathbf{h}_{i,j} \in \mathbb{R}^{d_v}$ and d_v is the embedding dimension. Then all embeddings $\{\mathbf{h}_i\}_{i=1}^{N_F}$ and the CLS token $\mathbf{h}_{cls} \in \mathbb{R}^{d_v}$ are then input to the video encoder $\mathbf{E}_V$, getting the representation $\mathbf{v} \in \mathbb{R}^d$. The video representation $\mathbf{v}$ and text representation $\mathbf{s}$ are then used to calculate cosine similarity, producing the retrieval or action recognition results.

3.2. Integrated Spatial-Temporal Dynamic Prompting

Our integrated spatial-temporal dynamic prompting method includes two modules, intra-frame spatial prompting and inter-frame temporal prompting. They are functionally complementary, guiding the pre-trained vision-language model to accurately focus on discriminative regions of the video in the spatial and temporal dimensions, respectively. The intra-frame spatial prompting module first identifies the location of the discriminative region $\mathbf{r}_i \in \mathbb{R}^{N_p}$ and then generates intra-frame spatial prompts $\mathbf{p}_i^s$ for frame $\mathbf{F}_i$ through a lightweight prompter $\mathcal{P}^s$:

$$\mathbf{p}_i^s = \mathcal{P}^s(\mathbf{h}_{i-1}, \mathbf{h}_i, \mathbf{h}_{i+1}), \quad (1)$$

Then, the intra-frame spatial prompt $\mathbf{p}_i^s$ is overlaid onto the discriminative region $\mathbf{r}_i$ of $\mathbf{h}_i$, resulting in $\mathbf{h}_i^s$.

The inter-frame video prompting module first utilizes the discriminative regions obtained by the intra-frame video prompting module and a 3D convolutional network to compute the temporal variation $\{w_i\}_{i=1}^{N_F-1}$ between adjacent frames in the video. $w_t \in \mathbb{R}$ represents the degree of tem-

poral variation between frames $\mathbf{F}_i$ and $\mathbf{F}_{i+1}$. Then, we use an inter-frame prompter $\mathcal{P}^t$ to generate inter-frame prompts $\{\mathbf{p}_i^t\}_{i=1}^{N_F-1}$ with $\{\mathbf{h}_i^s\}_{i=1}^{N_F}$ and $\{\mathbf{r}_i\}_{i=1}^{N_F}$. Then, we concatenate the inter-frame prompts $\{\mathbf{p}_i^t\}_{i=1}^{N_F-1}$ with the image token $\{\mathbf{h}_i^s\}_{i=1}^{N_F}$ and use them as the input to the pre-trained model's MSA block.

3.3. Intra-frame Spatial Prompting

As mentioned in Sec 3.2, we first identify the discriminative region $\mathbf{r}_i$ of each frame $\mathbf{F}_i$. Specifically, we compute the attention map A_i of different patches $\mathbf{h}_i = \{\mathbf{h}_{i,j} \in \mathbb{R}^{d_v}\}_{j=1}^{N_p}$ of a single frame $\mathbf{F}_i$ using the self-attention module $\text{Attn}(\cdot)$ of the pre-trained model:

$$A_i = \text{Attn}(\mathbf{h}_{cls}, \mathbf{h}_i), \quad (2)$$

where $A_i \in \mathbb{R}^{N_p}$. Additionally, a 3D convolutional layer $\mathcal{N}^s$ is utilized to calculate the temporal dynamics along the temporal dimension:

$$[\tilde{\mathbf{h}}_1, \tilde{\mathbf{h}}_2, \dots, \tilde{\mathbf{h}}_{N_F}] = \mathcal{N}^s([\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{N_F}]) \quad (3)$$

$$M_{i,j} = \frac{1}{d_v} \sum_k \tilde{\mathbf{h}}_{i,j,k}^2, \quad (4)$$

where $\tilde{\mathbf{h}}_i \in \mathbb{R}^{N_p \times d_v}$ and $M_{i,j} \in \mathbb{R}$ represents the temporal dynamic of the j -th token of the i -th frame. $[\]$ denotes the concatenation operation.

Then, we compute the discriminative regions $\mathbf{r}_i$ of each frame by using the attention map A_i and the temporal dynamic variation M_i :

$$\mathbf{W}_i^s = \alpha A_i + (1 - \alpha) M_i, \quad (5)$$

where α is a weight hyper-parameter. Then, we get the discriminative regions of each video frame with N_s patches:

$$\mathbf{r}_{i,j} = \begin{cases} 1, & \text{if } W_{i,j}^s \text{ is the top } N_s \text{ largest values of } W_i^s, \\ 0, & \text{otherwise} \end{cases}, \quad (6)$$

where N_s is a hyper-parameter.

Through the above process, we comprehensively consider the degree of temporal variation and the importance of understanding individual frames, thereby obtaining the discriminative regions of the video. Next, as introduced in Eq. 1 of Sec. 3.2, we use a lightweight prompter $\mathcal{P}^s$ to generate intra-frame spatial prompts $\mathbf{p}_i^s$, which we then overlay onto the tokens corresponding to the discriminative regions:

$$\mathbf{h}_{i,j}^s = \mathbf{h}_{i,j} + \mathbf{r}_{i,j} \cdot \mathbf{p}_{i,j}^s. \quad (7)$$

With the intra-frame spatial prompting, we identify the discriminative regions of each video frame and add prompts to the corresponding tokens, guiding the pre-trained model to focus accurately.

3.4. Inter-frame Temporal Prompting

Building on the intra-frame spatial prompting that highlights the discriminative regions within frames, we further introduce inter-frame temporal prompting to identify key frames along the temporal dimension. As mentioned in Sec. 3.2, we first use a 3D convolutional layer $\mathcal{N}^t$ to obtain the dynamic variation between adjacent frames:

$$\Delta \mathbf{h}_i^s = \mathbf{h}_i^s - \mathbf{h}_{i-1}^s, \quad (8)$$

$$[\tilde{\mathbf{h}}_1^s, \dots, \tilde{\mathbf{h}}_{N_F-1}^s] = \mathcal{N}^t([\Delta \mathbf{h}_1^s, \dots, \Delta \mathbf{h}_{N_F-1}^s]), \quad (9)$$

$$W_i^t = \frac{1}{N_p \cdot d_v} \sum_j ((1 + \beta \cdot \mathbf{r}_{i,j}) \sum_k (\tilde{\mathbf{h}}_{i,j,k}^s)^2), \quad (10)$$

where $\tilde{\mathbf{h}}_i^s \in \mathbb{R}^{N_p \times d_v}$, $W_i^t \in \mathbb{R}$ and β is a weight hyper-parameter. Through the above process, when calculating frame variation, we assign higher weights to the discriminative regions $\mathbf{r}_i$ identified by intra-frame spatial prompting, enabling a more accurate assessment of changes of the main object rather than the background.

Then, based on the magnitude of the frame temporal variation W_i^t , we determine the number of prompt tokens $N_i^t \in \mathbb{R}$ to be inserted between frames:

$$N_i^t = \lceil \eta \cdot W_i^t \rceil, \quad (11)$$

where $\lceil \cdot \rceil$ denotes the ceiling (rounding up to the nearest integer) and η is a scaling factor. Next, we introduce an inter-frame prompter $\mathcal{P}^t$ to generate the inter-frame prompts $\mathbf{p}_i^t$:

$$[\mathbf{p}_1^t, \dots, \mathbf{p}_{N_F}^t] = \mathcal{P}^t([\Delta \mathbf{h}_1^s, \dots, \Delta \mathbf{h}_{N_F-1}^s]), \quad (12)$$

where $\mathbf{p}_i^t \in \mathbb{R}^{N_i^t \times d_v}$. For different prompt token numbers N_i^t , the final linear layer of the prompter corresponds to selecting a net with an output dimension of N_i^t . Then, we concatenate the inter-frame temporal prompts $\mathbf{p}_i^t$ with the image tokens $\tilde{\mathbf{h}}_i^t$ and the CLS token $\mathbf{p}_{cls}$, and input them into the multi-head self-attention (MSA) block of the pre-trained model to obtain the final video representation $\mathbf{v} \in \mathbb{R}^d$. Then, we calculate the semantic relevance between $\mathbf{v}$ and the text representation $\mathbf{s}$ using cosine similarity:

$$c(\mathbf{s}, \mathbf{v}) = \frac{\mathbf{s} \cdot \mathbf{v}}{\|\mathbf{s}\| \|\mathbf{v}\|}, \quad (13)$$

where $\|\cdot\|$ denotes the ℓ_2 norm.

3.5. Overall Optimization

For the video action recognition task, following [51, 65, 66], we use the cross-entropy loss function for training:

$$\mathcal{L}_{act} = -\frac{1}{B} \sum_{i=1}^B \log \frac{e^{c(\mathbf{v}_i, \mathbf{s}_{y_i})/\tau}}{\sum_{j=1}^K e^{c(\mathbf{v}_i, \mathbf{s}_j)/\tau}}, \quad (14)$$

where y_i is the label of video $\mathbf{v}_i$, B denotes the batch size and τ is a temperature parameter.

For the video-text retrieval task, following [61, 64], we use contrastive loss for training. We treat paired text-video data as positive samples, while other data in the batch are negative samples:

$$\mathcal{L}_{vt} = \frac{1}{2B} \sum_{i=1}^B \left(\log \frac{e^{c(\mathbf{s}_i, \mathbf{v}_i)/\tau}}{\sum_j e^{c(\mathbf{s}_j, \mathbf{v}_i)/\tau}} + \log \frac{e^{c(\mathbf{s}_i, \mathbf{v}_i)/\tau}}{\sum_j e^{c(\mathbf{s}_i, \mathbf{v}_j)/\tau}} \right), \quad (15)$$

In our proposed STOP, the parameters to be optimized include two 3D convolutional layers $\mathcal{N}^s$ and $\mathcal{N}^t$, and two prompters $\mathcal{P}^s$ and $\mathcal{P}^t$.

4. Experiment

4.1. Datasets

Action Recognition. HMDB51 [26] includes around 7,000 clips across 51 action classes, featuring diverse real-world actions with variations in background and camera angles. UCF101 [47] offers over 13,000 clips spanning 101 classes of activities, covering sports, daily actions, and social interactions, adding scene diversity. SS-V2 (Something-Something V2) [16] has approximately 220,000 clips in 174 classes, emphasizing temporally dependent actions suitable for evaluating temporal reasoning.

Text-Video Retrieval. We evaluate on four widely used text-video retrieval datasets: MSR-VTT [56], ActivityNet [6], DiDeMo [2], and VATEX [54]. MSR-VTT contains 10,000 clips with around 20 descriptions each. ActivityNet includes 20,000 YouTube videos across 200 activities, with video descriptions combined into a single paragraph for video-text retrieval. DiDeMo consists of 10,000

	Methods	HMDB51	UCF101	SS-V2
	CLIP4Clip [42]	75.2	94.1	69.4
Adapter	Bias [7]	60.1	85.6	13.6
	Adapter ^{ATTN} [17]	60.3	85.7	14.2
	Adapter ^{FFN} [11]	60.6	86.2	13.8
	Visual-Text Adapter [18]	63.4	88.3	14.6
	Video-Text Adapter [44]	64.7	89.2	15.7
Prompt	VoP ^{F+C} [19]	65.2	91.3	16.7
	DGL-Linear [61]	67.2	92.5	18.3
	DGL-Transformer [61]	69.8	93.6	18.1
	UniPT [13]	65.2	90.6	15.6
	STOP (Ours)	72.0	95.3	21.4

- The best and second best results are marked in **RED** and **BLUE**, respectively.

Table 1. Comparison with state-of-the-art on the HMDB51, UCF101 and SS-V2. Here, we report the action recognition classification accuracy (ACC@1). For a fair comparison, all methods use CLIP-ViT-B/32 [45] as the backbone.

	Methods		Params (MB) ↓	Text → Video				Video → Text			
				R@1 ↑	R@5 ↑	R@10 ↑	MnR ↓	R@1 ↑	R@5 ↑	R@10 ↑	MnR ↓
	CLIP4Clip [42]	<i>Neurocomputing'22</i>	123.54	43.1	70.4	80.8	16.2	43.1	70.5	81.2	12.4
Adapter	Bias [7]	<i>NeurIPS'20</i>	0.1	39.7	66.5	77.3	17.3	41.1	68.4	79.2	13.6
	Adapter ^{ATTN} [17]	<i>ICLR'22</i>	2.0	37.6	63.2	75.8	18.7	39.6	66.5	76.8	14.7
	Adapter ^{FFN} [11]	<i>NeurIPS'22</i>	2.0	38.2	63.5	76.4	17.9	39.9	66.8	77.7	14.2
	Visual-Text Adapter [18]	<i>ICML'19</i>	11.82	39.2	65.7	76.1	17.6	40.7	68.8	77.6	13.7
	Video-Text Adapter [44]	<i>NeurIPS'22</i>	11.94	41.1	67.0	77.1	17.4	42.6	68.4	78.4	13.8
Prompt	Efficient Prompt [22]	<i>ECCV'22</i>	6.35	36.7	64.6	-	-	-	-	-	-
	VPT [21]	<i>ECCV'22</i>	0.18	42.0	66.6	77.3	19.2	39.4	66.8	77.2	16.2
	UPT [63]	<i>arXiv'22</i>	9.57	42.1	67.7	78.2	16.5	42.6	70.3	79.3	12.3
	VoP ^{F+C} [19]	<i>CVPR'23</i>	14.10	44.6	69.9	80.3	16.3	44.5	70.7	80.6	11.5
	DGL-Linear [61]	<i>AAAI'24</i>	0.83	44.7	70.5	79.2	16.2	42.1	70.0	80.6	13.4
	DGL-Transformer [61]	<i>AAAI'24</i>	9.57	45.8	69.3	79.4	16.3	43.5	70.5	80.7	13.1
	UniPT [13]	<i>CVPR'24</i>	9.60	38.9	60.2	71.4	18.5	39.3	58.6	70.4	16.4
	MPT-Linear [64]	<i>ACMMM'24</i>	0.87	45.0	70.8	79.6	16.2	42.8	70.6	81.1	12.9
	MPT-Transformer [64]	<i>ACMMM'24</i>	9.61	46.3	70.9	80.7	15.6	45.0	70.9	80.6	12.7
	STOP (Ours)	<i>This Paper</i>	7.53	47.7	71.4	81.1	15.2	46.1	71.6	81.0	12.2

- “Params” represents the number of trainable parameters of each method.
- The best and second best results are marked in **RED** and **BLUE**, respectively.

Table 2. Comparison with state-of-the-art on the MSR-VTT dataset. For a fair comparison, all methods use CLIP-ViT-B/32 [45] as the backbone.

Flickr videos paired with 40,000 sentences, using combined descriptions for model evaluation. The VATEX dataset is a large-scale video-text dataset containing 41,250 videos with bilingual English-Chinese captions, widely used for video captioning and cross-modal learning.

Evaluation Metrics. For video action recognition, we use classification accuracy (ACC@1) as the evaluation metric. For video-text retrieval, following previous work [61, 64], we use standard retrieval metrics to assess model performance, including R@K (Recall at Rank K, higher is better ↑) and MnR (Mean Rank, lower is better ↓).

4.2. Comparison Methods

We compare our STOP with both adapter-based and prompt-based parameter-efficient finetuning methods. We also report the fully tuning results as a baseline, i.e., CLIP4Clip [42]. For adapter-based methods, we report the results of Bias [7], Adapter^{ATTN} [17], Adapter^{FFN} [11], Visual-Text Adapter [18] and Video-Text Adapter [44]. For prompt-based methods, we report the results of Efficient Prompt [22], VPT [21], UPT [63], VoP^{F+C} [19], DGL-Linear [61], DGL-Transformer [61], UniPT [13], MPT-Linear [64] and MPT-Transformer [64].

4.3. Implementation Details

Following [61, 64], the video and text encoders are initialized with pre-trained CLIP (ViT-B/32), with pre-trained weights kept frozen during training. All video frames are

resized to 224×224 and divided into 49 non-overlapping patches. For video action recognition, we use a handcrafted prompt template “a video of a person doing [CLASS]” to construct the text input. For the MSR-VTT and VATEX datasets, the maximum length of captions is set to 32, and each video is uniformly sampled to 12 frames. For other datasets, the maximum numbers of sentences and frames are both set to 64. In our method, the default number of intra-frame prompts N_s is set to 6, the scaling factor of inter-frame prompts η is set to 12, and the weight hyperparameters α and β are set to 0.4 and 4, respectively. During training, the model is trained with the AdamW optimizer, with a cosine decay scheduler.

4.4. Comparison with State-of-the-arts

4.4.1. Action Recognition

We first validate the effectiveness of our method on the video action recognition task, conducting experiments on HMDB51, UCF101, and SS-V2. As shown in Table 1, compared to other state-of-the-art parameter-efficient finetuning and video prompting methods, our STOP achieves improvements of **2.2%**, **1.7%**, and **3.1%** on the three datasets, respectively. This is because our STOP leverages integrated spatial-temporal dynamic prompting, which guides the pre-trained model to focus on discriminative regions with temporal dynamics. This approach enhances the pre-trained CLIP model’s ability to capture fine-grained temporal information, leading to more accurate action

	Methods	ActivityNet				DiDeMo				VATEX			
		R@1 ↑	R@5 ↑	R@10 ↑	MnR ↓	R@1 ↑	R@5 ↑	R@10 ↑	MnR ↓	R@1 ↑	R@5 ↑	R@10 ↑	MnR ↓
	CLIP4Clip [42]	40.5	72.4	98.1	7.5	43.4	70.2	80.6	17.5	55.9	89.2	95.0	3.9
Adapter	Bias [7]	31.3	60.3	74.2	13.4	36.5	63.4	75.2	24.8	52.2	83.1	91.3	5.2
	Adapter ^{ATTN} [17]	31.6	60.5	74.4	13.1	36.4	62.8	73.9	23.5	52.6	83.4	91.8	5.0
	Adapter ^{FFN} [11]	31.8	61.0	75.0	12.8	36.3	63.4	75.4	22.9	52.3	83.3	91.5	5.2
	Visual-Text Adapter [18]	33.5	64.8	77.5	10.9	-	-	-	-	53.1	85.0	92.3	4.9
	Video-Text Adapter [44]	36.4	66.1	79.6	10.0	-	-	-	-	53.5	85.0	92.4	4.7
Prompt	VoP ^{F+C} [19]	36.1	65.5	78.5	10.9	45.3	72.3	80.4	13.8	54.2	85.2	93.6	4.7
	DGL-Linear [61]	38.3	68.4	79.2	10.3	44.2	70.6	80.2	15.8	56.2	87.1	93.5	4.1
	DGL-Transformer [61]	40.1	69.5	80.9	9.1	45.6	71.7	81.1	14.6	54.3	85.5	92.3	4.9
	UniPT [13]	34.6	65.7	75.2	15.5	40.1	64.2	74.7	18.7	53.4	85.2	92.3	5.0
	MPT-Transformer [64]	41.4	70.9	82.9	7.8	46.4	72.2	81.4	13.4	-	-	-	-
	STOP (Ours)	43.1	71.4	83.7	6.9	47.5	73.5	82.0	12.5	57.5	88.4	93.2	4.0

- The best and second best results are marked in **RED** and **BLUE**, respectively.

Table 3. Comparison with state-of-the-art on the ActivityNet, DiDeMo, and VATEX. Briefly, we only report text-to-video retrieval (t2v) results. For a fair comparison, all methods use CLIP-ViT-B/32 [45] as the backbone.

recognition results.

4.4.2. Video-Text Retrieval

Next, to further validate the effectiveness of our method, we conducted experiments on the video-text retrieval task. As shown in Table 2, on the MSR-VTT dataset, compared to the second-best method MPT-Transformer, our method achieves a reduction of **2.08 MB** in parameters while improving the mean rank (MnR) by **0.4%** for text-to-video and **0.5%** for video-to-text, with R@1 increasing by **1.4%** and **1.1%**, respectively. Additionally, our method consistently shows improvements in the R@5 and R@10 metrics. This is because our proposed intra-frame spatial prompts and inter-frame prompts complement each other, highlighting the locations of discriminative regions and keyframes in both spatial and temporal dimensions, which facilitates the model in accurately extracting video representations.

Furthermore, as shown in Table 3, we also conducted experiments on datasets ActivityNet, DiDeMo, and VATEX. Compared to the MPT-Transformer, our STOP method on dataset ActivityNet improves R@1, R@5, and R@10 by **1.7%**, **1.5%**, and **0.8%**, respectively, while decreasing the MnR by **0.9%**. On the other two datasets, DiDeMo and VATEX, our method also achieves consistent performance improvements. This is because the positions of our intra-frame prompts and the inter-frame prompts are dynamically adjusted, allowing our method to generalize well across different datasets.

4.5. Ablation

4.5.1. Influence of Different Components

To verify the effectiveness of the intra-frame spatial prompting and inter-frame temporal prompting, we conducted ablation experiments on three datasets: HMDB51, UCF101,

Intra	Inter	HMDB51	UCF101	SS-V2
-	-	50.2	70.2	10.2
✓	-	67.1	83.1	16.5
-	✓	69.4	87.4	15.9
✓	✓	72.0	95.3	21.4

Table 4. Ablation study about the influence of components in STOP. “-” and “✓” represent without or with this component. “Intra” represents the intra-frame spatial prompting and “Inter” represents the inter-frame temporal prompting.

and SS-V2, as shown in Table 4. As demonstrated, when neither component is used, STOP degrades to a pre-trained CLIP model. When the intra-frame spatial prompting is used alone, the model’s performance improves by 16.9% on HMDB51. This is because the intra-frame spatial prompts highlight key regions with significant dynamic changes for video understanding, facilitating the pre-trained CLIP model to accurately capture fine-grained temporal information. When the inter-frame temporal prompting is used alone, the model’s performance improves by 19.2%. This is due to the inter-frame prompt capturing temporal dynamics between frames and guiding the pre-trained model to focus on keyframes in the video. When both components are used together, the model’s performance improves by an additional 2.6%, as they complement each other by guiding the pre-trained model to focus on discriminative regions in both the spatial and temporal dimensions, enhancing the model’s ability to understand temporal dependencies.

4.5.2. The Visualization Results of Attention Map

To further explore the impact of our intra-frame spatial prompting and inter-frame temporal prompting, we visu-




Figure 3. Attention map visualization results of existing method DGL and STOP (Ours). Visualization of more cases is included in the supplementary material.

alize the attention maps of each video frame. As shown in Figure 3, existing video prompting methods like DGL-Transformer use the same static prompt for all videos. This causes the pre-trained CLIP model to focus on static objects in the video, such as chairs, basketball hoops, and poles. As a result, the model struggles to accurately understand the actions performed by people in the video. In contrast, our intra-frame spatial prompting and inter-frame temporal prompting highlight the key regions with dynamic changes in the video. This enables the pre-trained model to focus on the person and his actions in the video, leading to a more accurate understanding.

4.5.3. Influence of Hyper-parameters

We conduct ablation experiments to investigate the impact of several hyper-parameters in our method, including the intra-frame prompt number N_s , the scaling factor of inter-frame prompts η , and the weight hyper-parameters α and




Figure 4. Influence of hyper-parameters of STOP in HDMB51.

β . The results, as illustrated in Figure 4, indicate that the model performs best when N is set to 6. This is because, when N_s is too small, the intra-frame spatial prompts fail to cover the discriminative regions, while if N_s is too large, it leads to an overemphasis on the background. For the scaling factor η , when it is too small, the number of inter-frame temporal prompt tokens is insufficient, limiting the model’s learning capacity. On the other hand, when η is too large, the excessive number of inter-frame temporal prompts interferes with feature extraction. The optimal performance is achieved when η is set to 12.

The weight hyper-parameter α controls the importance of the intra-frame attention map and the temporal dynamics in the inter-frame prompt when computing the discriminative regions. When α is set to 0.4, it balances both aspects, leading to more accurate identification of discriminative regions and achieving the best performance. Similarly, the weight hyper-parameter β affects the importance of discriminative regions when computing the temporal dynamics across frames. When β is set to 4, it strikes a balance and results in optimal model performance.

5. Conclusion

In this paper, we propose a novel integrated spatial-temporal dynamic prompting (STOP) approach. Specifically, we design intra-frame spatial prompts with dynamically changing positions, as well as inter-frame temporal prompts with dynamically changing quantities based on the temporal variations of the video. This approach guides the pre-trained vision-language model to focus on discriminative spatiotemporal regions, enhancing the model’s ability to accurately understand temporal dependencies within the video. The effectiveness of our proposed STOP has been validated on multiple large-scale benchmarks for video action recognition and video-text retrieval.

Acknowledgments

This work was supported by the grants from the National Natural Science Foundation of China (62376011, 61925201, 62132001, 62432001) and Beijing Natural Science Foundation (L247006).

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *NeurIPS*, 35: 23716–23736, 2022. [2](#)
- [2] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *ICCV*, pages 5803–5812, 2017. [5](#)
- [3] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, pages 1728–1738, 2021. [2](#)
- [4] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, page 4, 2021. [3](#)
- [5] Yi Bin, Wenhao Shi, Yujuan Ding, Zhiqiang Hu, Zheng Wang, Yang Yang, See-Kiong Ng, and Heng Tao Shen. Gallerygpt: Analyzing paintings with large multimodal models. In *ACM MM*, pages 7734–7743, 2024. [2](#)
- [6] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015. [5](#)
- [7] Han Cai, Chuang Gan, Ligeng Zhu, and Song Han. Tinytl: Reduce memory, not parameters for efficient on-device learning. *NeurIPS*, 33:11285–11297, 2020. [5](#), [6](#), [7](#)
- [8] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 6299–6308, 2017. [3](#)
- [9] Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*, 15(3):1–45, 2024. [2](#)
- [10] Shizhe Chen, Yida Zhao, Qin Jin, and Qi Wu. Fine-grained video-text retrieval with hierarchical graph reasoning. In *CVPR*, pages 10638–10647, 2020. [3](#)
- [11] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition. *NeurIPS*, 35:16664–16678, 2022. [5](#), [6](#), [7](#)
- [12] Leichao Dai, Lin Feng, Xinglin Shang, and Han Su. Cross modal adaptive few-shot learning based on task dependence. *Chinese Journal of Electronics*, 32(1):85–96, 2023. [1](#)
- [13] Haiwen Diao, Bo Wan, Ying Zhang, Xu Jia, Huchuan Lu, and Long Chen. Unipt: Universal parallel tuning for transfer learning with efficient parameter and memory. In *CVPR*, pages 28729–28740, 2024. [3](#), [5](#), [6](#), [7](#)
- [14] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [4](#)
- [15] Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *ECCV*, pages 214–229. Springer, 2020. [2](#), [3](#)
- [16] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The” something something” video database for learning and evaluating visual common sense. In *ICCV*, pages 5842–5850, 2017. [5](#)
- [17] Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. In *ICLR*, 2022. [5](#), [6](#), [7](#)
- [18] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morroni, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *ICML*, pages 2790–2799. PMLR, 2019. [5](#), [6](#), [7](#)
- [19] Siteng Huang, Biao Gong, Yulin Pan, Jianwen Jiang, Yiliang Lv, Yuyuan Li, and Donglin Wang. Vop: Text-video co-operative prompt tuning for cross-modal retrieval. In *CVPR*, pages 6565–6574, 2023. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [20] Zhao Huijuan, YE Ning, and Wang Ruchuan. Improved cross-corpus speech emotion recognition using deep local domain adaptation. *Chinese Journal of Electronics*, 32(3): 640–646, 2023. [1](#)
- [21] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *ECCV*, pages 709–727. Springer, 2022. [2](#), [3](#), [6](#)
- [22] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *ECCV*, pages 105–124. Springer, 2022. [3](#), [6](#)
- [23] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. page 2. Minneapolis, Minnesota, 2019. [3](#)
- [24] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *CVPR*, pages 19113–19122, 2023. [2](#)
- [25] Muhammad Uzair Khattak, Syed Talal Wasim, Muzammal Naseer, Salman Khan, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Self-regulating prompts: Foundational model adaptation without forgetting. In *ICCV*, pages 15190–15200, 2023. [2](#)
- [26] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *ICCV*, pages 2556–2563. IEEE, 2011. [5](#)
- [27] 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 *CVPR*, pages 7331–7341, 2021. [2](#)

- [28] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. 2
- [29] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *NeurIPS*, 34:9694–9705, 2021. 2
- [30] Kunchang Li, Yali Wang, Peng Gao, Guanglu Song, Yu Liu, Hongsheng Li, and Yu Qiao. Uniformer: Unified transformer for efficient spatiotemporal representation learning. *arXiv preprint arXiv:2201.04676*, 2022. 3
- [31] Qiwei Li, Kunlun Xu, Yuxin Peng, and Jiahuan Zhou. Exemplar-free lifelong person re-identification via prompt-guided adaptive knowledge consolidation. *IJCV*, pages 1–16, 2024. 2
- [32] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021. 2
- [33] Ji Lin, Chuhan Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, pages 7083–7093, 2019. 3
- [34] Fang Liu, Xiaoxue Qian, Licheng Jiao, Xiangrong Zhang, Lingling Li, and Yuanhao Cui. Contrastive learning-based dual dynamic gcnn for sar image scene classification. *IEEE Transactions on Neural Networks and Learning Systems*, 35(1):390–404, 2022. 3
- [35] Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Lam Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks. *arXiv preprint arXiv:2110.07602*, 2021. 2
- [36] Yang Liu, Samuel Albanie, Arsha Nagrani, and Andrew Zisserman. Use what you have: Video retrieval using representations from collaborative experts. *arXiv preprint arXiv:1907.13487*, 2019. 3
- [37] Zhaoyang Liu, Limin Wang, Wayne Wu, Chen Qian, and Tong Lu. Tam: Temporal adaptive module for video recognition. In *ICCV*, pages 13708–13718, 2021. 3
- [38] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *CVPR*, pages 3202–3211, 2022. 3
- [39] Zichen Liu, Yuxin Peng, and Jiahuan Zhou. Compositional prompting for anti-forgetting in domain incremental learning. *IJCV*, pages 1–18, 2024. 2
- [40] Zichen Liu, Yuxin Peng, and Jiahuan Zhou. Insyp: Efficient instance visual prompting from image itself. In *ACM MM*, pages 6443–6452, 2024. 2
- [41] Zichen Liu, Hongbo Sun, Yuxin Peng, and Jiahuan Zhou. Dart: Dual-modal adaptive online prompting and knowledge retention for test-time adaptation. In *AAAI*, pages 14106–14114, 2024. 2
- [42] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508:293–304, 2022. 2, 3, 4, 5, 6, 7
- [43] Chancharik Mitra, Brandon Huang, Trevor Darrell, and Roei Herzig. Compositional chain-of-thought prompting for large multimodal models. In *CVPR*, pages 14420–14431, 2024. 2
- [44] Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. St-adapter: Parameter-efficient image-to-video transfer learning. *NeurIPS*, 35:26462–26477, 2022. 5, 6, 7
- [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 1, 2, 3, 4, 5, 6, 7
- [46] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *NeurIPS*, 27, 2014. 3
- [47] K Soomro. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 5
- [48] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497, 2015. 3
- [49] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, pages 6450–6459, 2018. 3
- [50] A Vaswani. Attention is all you need. *NeurIPS*, 2017. 4
- [51] Hao Wang, Fang Liu, Licheng Jiao, Jiahao Wang, Zehua Hao, Shuo Li, Lingling Li, Puhua Chen, and Xu Liu. Vilt-clip: Video and language tuning clip with multimodal prompt learning and scenario-guided optimization. In *AAAI*, pages 5390–5400, 2024. 4, 5
- [52] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, pages 20–36. Springer, 2016. 3
- [53] Mengmeng Wang, Jiazhen Xing, and Yong Liu. Actionclip: A new paradigm for video action recognition. *arXiv preprint arXiv:2109.08472*, 2021. 3
- [54] Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuan-Fang Wang, and William Yang Wang. Vatex: A large-scale, high-quality multilingual dataset for video-and-language research. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4581–4591, 2019. 5
- [55] Peng Wu, Xiangteng He, Mingqian Tang, Yiliang Lv, and Jing Liu. Hanet: Hierarchical alignment networks for video-text retrieval. In *ACM MM*, pages 3518–3527, 2021. 3
- [56] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, pages 5288–5296, 2016. 5
- [57] Kunlun Xu, Chenghao Jiang, Peixi Xiong, Yuxin Peng, and Jiahuan Zhou. Dask: Distribution rehearsing via adaptive style kernel learning for exemplar-free lifelong person re-identification. *arXiv preprint arXiv:2412.09224*, 2024. 1
- [58] Kunlun Xu, Haozhuo Zhang, Yu Li, Yuxin Peng, and Jiahuan Zhou. Mitigate catastrophic remembering via continual knowledge purification for noisy lifelong person re-identification. In *ACM MM*, pages 5790–5799, 2024.

- [59] Kunlun Xu, Xu Zou, Yuxin Peng, and Jiahuan Zhou. Distribution-aware knowledge prototyping for non-exemplar lifelong person re-identification. In *CVPR*, pages 16604–16613, 2024.
- [60] Kunlun Xu, Xu Zou, and Jiahuan Zhou. Lstkc: Long short-term knowledge consolidation for lifelong person re-identification. In *AAAI*, pages 16202–16210, 2024. [1](#)
- [61] Xiangpeng Yang, Linchao Zhu, Xiaohan Wang, and Yi Yang. Dgl: Dynamic global-local prompt tuning for text-video retrieval. In *AAAI*, pages 6540–6548, 2024. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [62] Zhaoda Ye, Xiangteng He, and Yuxin Peng. Unsupervised cross-media hashing learning via knowledge graph. *Chinese Journal of Electronics*, 31(6):1081–1091, 2022. [1](#)
- [63] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Unified vision and language prompt learning. *arXiv preprint arXiv:2210.07225*, 2022. [6](#)
- [64] Haonan Zhang, Pengpeng Zeng, Lianli Gao, Jingkuan Song, and Heng Tao Shen. Mpt: Multi-grained prompt tuning for text-video retrieval. In *ACM MM*, 2024. [2](#), [3](#), [5](#), [6](#), [7](#)
- [65] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *CVPR*, pages 16816–16825, 2022. [2](#), [5](#)
- [66] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *IJCV*, 130(9):2337–2348, 2022. [2](#), [5](#)