

# SCP: Soft Conditional Prompt Learning for Aerial Video Action Recognition

Xijun Wang<sup>\*1</sup>, Ruiqi Xian<sup>\*2</sup>, Tianrui Guan<sup>1</sup>, Fuxiao Liu<sup>1</sup> and Dinesh Manocha<sup>1</sup>

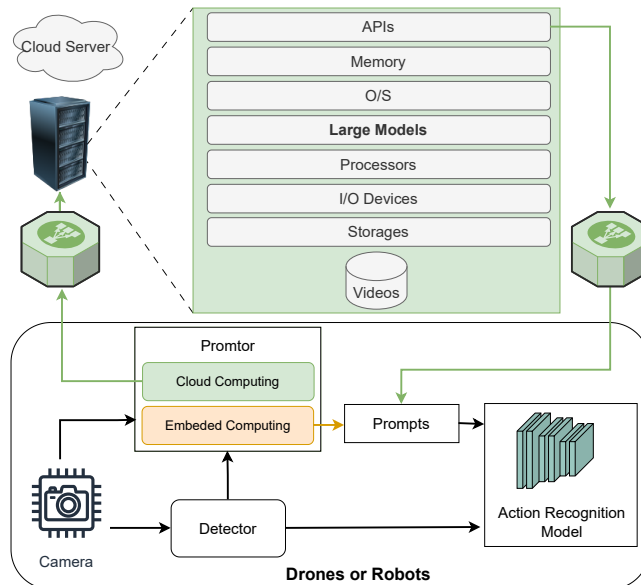
**Abstract**—We present a new learning approach, **Soft Conditional Prompt Learning (SCP)**, which leverages the strengths of prompt learning for aerial video action recognition. Our approach is designed to predict the action of each agent by helping the models focus on the descriptions or instructions associated with actions in the input videos for aerial/robot visual perception. Our formulation supports various prompts, including learnable prompts, auxiliary visual information, and large vision models to improve the recognition performance. We present a soft conditional prompt method that learns to dynamically generate prompts from a pool of prompt experts under different video inputs. By sharing the same objective with the task, our proposed SCP can optimize prompts that guide the model’s predictions while explicitly learning input-invariant (prompt experts pool) and input-specific (data-dependent) prompt knowledge. In practice, we observe a 3.17 – 10.2% accuracy improvement on the aerial video datasets (Okutama [1], NECDrone [2]), which consist of scenes with single-agent and multi-agent actions. We further evaluate our approach on ground camera videos to verify the effectiveness and generalization and achieve a 1.0 – 3.6% improvement on SSV2 [3]. We integrate our method into the ROS2 as well.

## I. INTRODUCTION

In the realm of unmanned aerial vehicles (UAVs), the ability to accurately recognize human actions from video footage is paramount for safe and effective operation. This entails extracting meaningful insights into the activities and movements of people and objects within the environment, leveraging video sequences captured by the onboard camera. This capability serves as a cornerstone technology for UAV applications, including human-UAV interaction, search and rescue, and comprehensive aerial surveillance.

Similar to ground robots, recent advancements in deep learning techniques have yielded significant strides in human action recognition for UAV videos. However, a crucial challenge persists: the majority of existing approaches rely heavily on extensive, meticulously labeled training datasets and adhere to a purely supervised learning paradigm that primarily focuses on optimizing the network architecture design. When directly applied to aerial video datasets, these methods often experience significant degradation in performance due to inherent challenges unique to this domain including small target sizes, disparate viewing angles, and camera motion dynamics.

To address these limitations, our work explores the application of prompt learning for UAV video action recognition. Prompt-based learning techniques, recently demonstrating



**Fig. 1: Overall Architecture:** Our action recognition method is designed to run one edge devices (on mobile robots) and cloud servers. This includes lightweight prompts (embedded), which can be easily embedded in any action recognition model without much extra computational cost. For large vision models, we perform these computations on cloud server and use low-latency communication with the robots.

success in natural language processing tasks [4], circumvent the requirement for extensive labeled data by leveraging pre-trained language models. In the context of UAV action recognition, prompt learning offers a promising avenue for designing more robust recognition models. By incorporating high-level texture descriptions or instructions associated with actions, prompts can effectively guide the model’s learning process. This targeted guidance allows the model to focus on discriminative spatiotemporal patterns in the aerial video data, especially when dealing with challenging visual features like small targets or unusual camera angles. Furthermore, the ease of obtaining or embedding prompt information within existing robotic systems facilitates the practical implementation of this approach.

**Main Results:** In this paper, we propose a novel prompt-learning approach to address the challenges of UAV video action recognition. Our approach integrates prompts to enhance the model’s ability to process video data effectively. These prompts can be either learnable or pre-defined templates specifically designed for action recognition tasks. By incorporating prompts, our method facilitates the model’s focus on critical regions of interest within video frames. This targeted focus enables the learning of complex visual concepts, such as recognizing interactions between multiple

<sup>\*</sup>These authors contributed equally

<sup>1</sup>Authors are with Dept. of Computer Science, University of Maryland, College Park, MD, USA. xijun@umd.edu

<sup>2</sup>Author is with the Dept. of Electrical and Computer Engineering, University of Maryland, College Park, MD, USA rxian@umd.edu

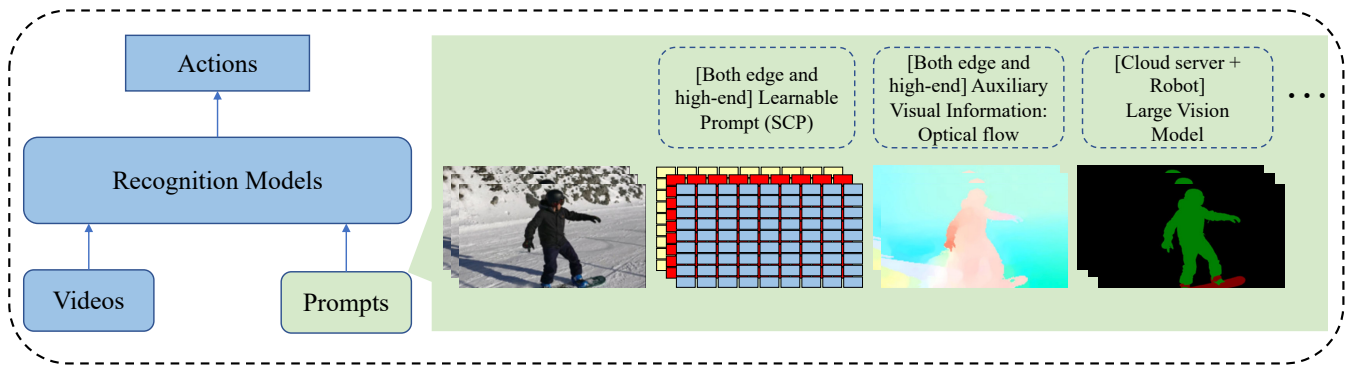


Fig. 2: **Task Overview:** We use prompt learning for action recognition. Our method leverages the strengths of prompt learning to guide the learning process by helping models better focus on the descriptions or instructions associated with actions in the input videos. We explore various prompts, including optical flow, large vision models, and proposed SCP to improve recognition performance. The recognition models can be CNNs or Transformers.

agents in aerial footage.

In our prompt learning paradigm, we explore and discuss different types of prompts, including learnable prompts, auxiliary visual information (optical flow, detection, etc.), and large vision models. For learnable prompts, our SCP dynamically generates prompts from a pool of prompt experts under different inputs. Our goal is to optimize prompts that guide the model’s predictions while explicitly learning input-invariant (prompt experts) and input-specific (data-dependent) prompt knowledge. For auxiliary visual information, we can easily obtain them from the robot’s built-in system. Our SCP can be easily embedded in any model without much extra computational cost, especially suitable for edge and mobile devices. We validate the generalization by performing evaluations on datasets comprised of aerial videos and ground camera videos on scenarios involving single-agent and multi-agent actions. We demonstrate that our technique can improve performance and enhance the generalization capabilities of video action recognition models in different scenarios. Our main contributions include:

- 1) We present a general learning approach to use prompt learning and auto-regressive techniques for aerial video action recognition.
- 2) We propose a new soft conditional learnable prompt method that can guide the model’s predictions while explicitly learning input-invariant (prompt experts) and input-specific (data-dependent) prompt knowledge.
- 3) To the best of our knowledge, ours is the first approach to explore the possibility of using large vision models as the prompt to instruct the models on aerial video action recognition tasks.
- 4) Through empirical evaluations, we demonstrate the potential and effectiveness of prompt learning techniques for aerial video action recognition tasks. Specifically, we observe a 3.17-10.2% accuracy improvement on the aerial video datasets. Moreover, we observe a 1.0-3.6% accuracy improvement on the ground camera video dataset Something Something V2.

## II. RELATED WORKS

### A. Action Recognition

Human action recognition, i.e., recognizing and understanding human actions, is crucial for a number of real-world applications. Recently, many deep learning architectures have been proposed to improve the performance. At a broad level, they can be classified into three categories: Two-stream 2D Convolutional Neural Network [5], [6], [7], [8], [9], [10], 3D CNN-based methods [11], [12], [13], [14], [15], Transformer-based approaches [16], [17], [18].

Although these methods have had good success on the ground data and YouTube videos, they cannot achieve a similar level of accuracy on videos captured using Unmanned Aerial Vehicles (UAVs) [19], [20]. Compared to ground or YouTube videos, UAV videos have unique characteristics like small resolution, scale and size variations, and moving cameras. [19] proposed auto zoom algorithms with an attention mechanism for inference on both edge devices and desktop GPUs. [21] proposed a mutual information-based feature alignment and sampling method to extract spatial-temporal features corresponding to human actors for better recognition accuracy. [22] introduced Fourier transformation into attention modules to aggregate the motion salience. [20] proposed a novel frame sampler for aerial action recognition by measuring the similarity between frame patches. Our SCP can help the above methods better focus on the target agents.

### B. Prompt Learning

The concept of prompt learning, initially introduced by [23], has garnered significant attention in the field of Natural Language Processing (NLP)[24], [25], [26], [4], [27], [28]. Prompt learning revolves around the fundamental idea of treating pre-trained language models like BERT or GPT as knowledge repositories, enabling their utilization in downstream tasks. Early studies, exemplified by[23], [29], concentrated on crafting prompts manually to enhance language model performance. Subsequently, researchers like [30], [25] aimed to automate this process using cost-effective, data-driven approaches. More recently, some works[31], [32], [33] have ventured into learning continuous prompts as an alternative to seeking discrete prompts.

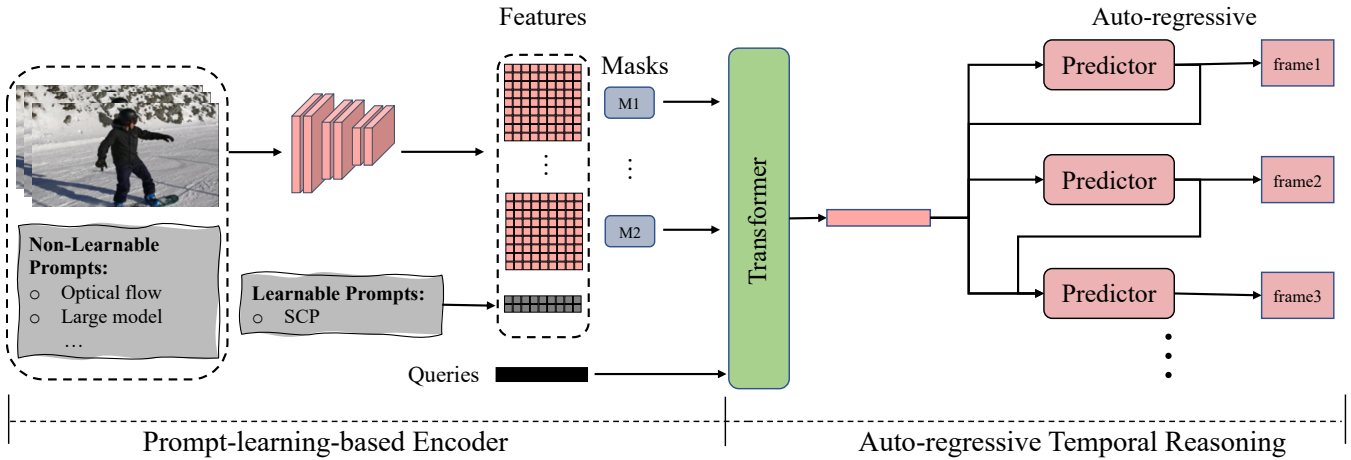


Fig. 3: **Overview of the action recognition framework:** We use transformer-based action recognition methods as an example. We designed a prompt-learning-based encoder to help better extract the feature and use our auto-regressive temporal reasoning algorithm for recognition models for enhanced inference ability.

In [34], the versatility of expressing a wide range of robot manipulation tasks through multimodal prompts is demonstrated using VIMA, a transformer-based generalist robot agent that processes prompts and generates motor actions autoregressively. [35] introduces a programmatic LLM prompt structure to facilitate plan generation adaptable to various settings, robot functionalities, and tasks. Additionally, [36] proposes a strategy combining prompt engineering principles and a high-level function library to enhance ChatGPT’s adaptability to diverse robotics tasks, simulation environments, and hardware setups. In fashion, [37] use scenes as prompts to help style-matched recommendations. In foundation model design, [38] explores different spatial information as prompts. [39], [40] explore how to find the keyframes efficiently as prompts for LLMs. Recently, more and more researchers started exploring prompt learning techniques in vision tasks [41], [42], [43], [44], [45].

While previous research has predominantly concentrated on prompt learning for ground robot tasks, the application of prompt learning to UAV tasks has received limited attention. This paper introduces a comprehensive learning framework aimed at assessing the efficacy of prompt learning in the context of UAV video comprehension, particularly in the realm of action recognition in both ground/YouTube and aerial videos. The objective is to bridge this gap and broaden the applicability of prompt learning to video understanding tasks within this domain.

### III. OUR APPROACH

We denote the input as  $X_i = \{x_1, x_2, \dots, x_m\}, i \in [1, N]$ , where  $x_j$  is the  $j$ th frame in the  $i$ th video,  $m$  is the total frame number, and  $N$  is the total number of videos. The overall approach predicts the action categories by using model  $f(X_i)$ , which can be CNNs or Transformers. As shown in Figure 3, taking transformer-based methods as an example, we follow the same scheme to extract the features, followed by using the reasoning process to predict the action labels. We also present a prompt-learning-based encoder to help better extract the feature and then propose an auto-regressive

temporal reasoning algorithm for recognition models for enhanced inference ability. Specifically, in an action model:

$$f = f_a \circ f_e([X, P]), \quad (1)$$

where  $f_e$  is the *prompt-learning-based input encoder*,  $P$  is the prompt, and  $f_a$  is the *auto-regressive-based temporal reasoning* model, which is used for the temporal dimension.

#### A. Prompt Learning-based Input Encoder

For the first part of the input encoder, inspired by these prompt-based techniques in NLP, we present a new general prompt learning-based input encoder for action recognition. Our formulation leverages the strengths of prompt learning to guide the optimization by providing high-level descriptions or instructions associated with actions in the inputs. We use this to alleviate the burden of models’ optimization by helping models better focus on the active region.

Prompts can enhance the model’s ability to process customized inputs by utilizing prompt tokens. By leveraging prompts, models can more easily focus on the interest targets, and prompt learning enables the model to learn complex visual concepts and capture discriminative spatio-temporal patterns effectively. Specifically, our prompts can be either predefined templates (non-learnable prompt: optical flow, large vision models) or learnable tokens (learnable prompt) that include task-specific information. They can be used either alone or in combination.

1) *Learnable Prompt: Soft Conditional Prompt Learning (SCP)*: To better adapt to the input data, we propose a soft conditional prompt learning (SCP), which learns to dynamically generate prompts from a pool of prompt experts under different inputs. Prompt experts are learnable parameters that can be updated from the training process. As shown in Figure 4, in our design, we use input-invariant (prompt experts) and input-specific (data dependent) prompts. The input-invariant prompts contain task information, and we use a dynamic mechanism to generate input-specific prompts for different inputs.




Fig. 4: **Soft Conditional Prompt Learning (SCP)**: Learning input-invariant (prompt experts) and input-specific (data dependent) prompt. The input-invariant prompts will be updated from all the inputs, which contain task information, and we use a dynamic mechanism to generate input-specific prompts for different inputs. Add/Mul means element-wise operations.  $B \times S \times C$  is the input features' shape, and  $l$  is the expert's number in the prompt pool.

There are different actions and domains (different video sources) for different videos, so it's challenging to learn a single general prompt for all videos. Therefore, we design an input-invariant prompt experts pool, which contains  $l$  learnable prompts. Unless otherwise specified, the default value of  $l$  is 8.

$$P = \{P_1, \dots, P_l\}, \quad (2)$$

those prompt experts are learnable and will be updated from all the inputs. For a specific input  $X^*$ ,

$$P^* = \text{Matmul}(\sigma(\text{FC}(X^*)), P), \quad (3)$$

We use an FC layer and sigmoid function  $\sigma$  to get dynamic weights. Then we apply these dynamic weights to the input-invariant prompt pool to get a customized prompt  $P^*$  for  $X^*$ .

$$x_i^p = f_e([x_i, p_i]), x_i \in X^*, p_i \in P^*, \quad (4)$$

where  $x_i^p$  is the prompt-based feature.

2) *Non-Learnable Prompt*: Non-Learnable prompts make use of statistical methods (e.g., optical flow) or existing powerful large vision models, which can offer reliable prompts without training.

a) *Optical Flow Prompt*: Optical flow is a fundamental concept in computer vision that involves estimating the motion of objects within a video sequence. It represents the apparent motion of pixels between consecutive frames, providing valuable information about the movement of objects and their relative velocities.

We divide a video into  $m$  clips. For raw frame  $x_i$  and frame  $x_j$  from the video, the optical flow is:

$$o_i = O(x_i, x_j), x_i \in \text{clip}_i, x_j \in \text{clip}_j, \quad (5)$$

where  $\text{clip}_i$  and  $\text{clip}_j$  are two adjacent clips from a video, and each clip contains several frames. When computing the optical flow, we only use one frame from each clip in a video and then apply the optical flow to this whole clip. This formulation is more efficient because it avoids many calculations for every frame. Therefore, the input with optical flow prompt becomes:

$$[X, P] = \{x_k * o_i | x_k \in \text{clip}_i, i \in [1, m]\} \quad (6)$$

where  $\text{clip}_i$  has  $k$  frames. We use  $[X, P]$  to replace the original  $X$  in video action recognition.

b) *Large Vision Model Prompt*: Recently, large models have been attracting more attention for NLP and other applications. These large models are considered powerful since they are trained on huge amounts of data and don't need to be finetuned on new tasks as an auxiliary input (i.e. prompt). Our goal is to use these large models to generate prompts (e.g. mask, bbox) for video action recognition.

One popular work is the Segment Anything Model (SAM [46]), which can segment any object in an image given only some prompts like a single click or box. SAM is trained on a dataset of 11 million images and 1.1 billion masks. SAM can segment objects with high accuracy, even when they are new or have been modified from the training data. SAM generalizes to new objects and images without the need for additional training, so we don't need to finetune the model on our dataset. For some frames in a video clip, we generate a segmentation mask using a large vision model, SAM [46]. Next, these masks are used as prompts and fused with input frames to optimize the recognition model. Specifically, for frame  $x_i$ , the output from SAM is:

$$p_i = \text{SAM}(x_i, \text{boxes}/\text{points}), x_i \in \text{clip}_i \quad (7)$$

$\text{clip}_i$  is a video clip containing a few frames,

$$[X, P] = \{x_i * p_i | i \in [1, m]\} \quad (8)$$

We use  $[X, P]$  to replace the original  $X$ .

### B. Auto-regressive Temporal Reasoning

Temporal reasoning is important for sequence data. Therefore, we propose an Auto-regressive Temporal Reasoning algorithm to better model the time-varying data. Auto-regressive models are statistical models that make predictions based on previous observations. They assume that the future values of a variable can be estimated by considering its past values. For temporal reasoning, this concept is extended to capture dependencies between different frames in a video.

After getting the prompt-based feature  $X^p = \{x_1^p, x_2^p, \dots, x_m^p\}$ , where  $x_i^p$  represents the observation at time step  $i$ , the goal is to predict the future values,

$$\hat{x}_{i+1}^p = f_a\left(\prod_{j < (i+1)} f_a(x_j^p) + x_{i+1}^p\right) \quad (9)$$

where  $f_a$  denotes the auto-regressive model that maintains an internal state and updates according to the sequential input.  $\prod$  means a series of functions here. The auto-regressive temporal reasoning model considers the past observations of the sequence and the corresponding future observations to learn the underlying temporal dependencies.

### C. Single-agent and Multi-agent Objective

The supervision formats used for single-agent and multi-agent action recognitions are different. As a result, we choose different loss functions. Specifically, we choose the classical cross-entropy loss for single-agent action recognition,

$$L_n = - \sum_{c=1}^C \log \frac{\exp(\hat{x}_{n,c}^p)}{\sum_{i=1}^C \exp(\hat{x}_{n,i}^p)} y_{n,c}, \quad (10)$$

Method	Frame size	Accuracy
AARN [47], [48]	crops	33.75%
Lite ECO [49], [48]	crops	36.25%
I3D(RGB)[15], [48]	crops	38.12%
3DCapsNet-DR[50], [48]	crops	39.37%
3DCapsNet-EM[50], [48]	crops	41.87%
DroneCaps[48]	crops	47.50%
DroneAttention without bbox[51]	720×420	61.34%
SCP without bbox (Ours)	224×224	71.54%
DroneAttention with bbox [51]	720×420	72.76%
SCP with bbox (Ours)	224×224	75.93%

TABLE I: Comparison with the state-of-the-art results on the Okutama dataset. With bbox information, we achieved 10.20% improvement over the SOTA method. Without bbox information, we outperformed the SOTA by 3.17%. crops: from detection.

where  $C$  is the class number,  $n$  is the video number, and  $\hat{x}_{n,c}^p$  is the SCP’s output feature.  $y$  is the label. For multi-agent on Okutama, we use the BCEWithLogitsLoss,

$$L_{n,c} = - [y_{n,c} \cdot \log \sigma(\hat{x}_{n,c}^p) + (1 - y_{n,c}) \cdot \log(1 - \sigma(\hat{x}_{n,c}^p))] \quad (11)$$

where  $\hat{x}_{n,c}^p$  is the SCP’s output feature.  $\sigma$  is a sigmoid function. This loss combines a sigmoid function and the BCELoss, which is more numerically stable than using a plain sigmoid followed by a BCELoss because by combining the operations into one layer, it takes advantage of the log-sum-exp for numerical stability. For both single-agent and multi-agent videos, by sharing the same objective, our learning approach can optimize prompts that guide the model’s predictions while explicitly learning input-invariant (prompt experts pool) and input-specific (data-dependent) prompt knowledge.

#### IV. DATASETS AND RESULTS

To verify the effectiveness of SCP, empirical evaluations were conducted on Okutama [1] and NEC Drone [2] comprising both single-agent and multi-agent actions. We further evaluate on Something-something V2 [3] ground camera videos to verify the effectiveness and generalization.

##### A. Datasets and Experiment Settings

a) *Okutama* [1]: The Okutama dataset consists of 43 minute-long sequences with 12 action classes, providing a challenge with dynamic action transitions, changing scales and aspect ratios, camera movement, and multi-labeled actors. All the frames extracted from the video datasets were scaled to  $224 \times 224$ . The backbone is Swin-T [52]. Following [51], the feature maps obtained were processed in the ROIAlign function (crop size of  $5 \times 5$ ) to get the desired ROIs. Other training settings follow [52].

b) *NEC Drone* [2]: features 5,250 videos depicting 16 distinct actions performed by 19 actors. The initial learning rate is set 0.05. Stochastic Gradient Descent (SGD) is used as the optimizer with 0.0005 weight decay and 0.9 momentum. We use cosine/poly annealing for learning rate decay. All the frames from the video datasets were scaled to  $224 \times 224$ .

c) *Something-something v2 (SSV2 [3])*: The SSV2 dataset is regarded as a substantial and comprehensive benchmark for action recognition, encompassing a vast collection of 220k action clips. Following [53], we train for 100 epochs using 8 GPUs with a batch size of 64 and a base learning rate of  $5e-5$  with a cosine learning rate schedule. We use Adamw and use a weight decay of  $1e-4$  and a drop path rate of 0.4. For other training and testing settings, we follow [53]. And the backbone is MVITv2-S [53].

##### B. Results on Okutama

Okutama is an aerial multi-agent action recognition dataset in which multiple actors sequentially perform a diverse set of actions, which makes it very challenging. In the real world, it’s difficult to ensure that only a single agent is in the scene for action recognition. Therefore, multi-agent action recognition is a very practical and important direction. We compare our SCP with state-of-the-art (SOTA) works.

As shown in Table I, if there is no bbox information, we achieved 10.20% improvement over the SOTA method. If there is bbox information, we outperform the SOTA by 3.17%. This demonstrates the effectiveness of our method.

##### C. Results on NECDrone

We compare our method with other existing methods on NEC-Drone. The frames are extracted from raw videos and augmented as in X3D [54]. The baseline methods use uniform and random sampling. As shown in Table II, on NEC Drone, our SCP outperforms the X3D by 4.0 - 7.4% and improves 23.1% over the K-centered.

Method	pretrain	Top-1 Accuracy
X3D-random [54]	Kinetics400	52.0%
X3D-uniform [54]	Kinetics400	55.4%
K-centered [55]	N/A	36.3%
SCP-uniform (Ours)	Kinetics400	59.4%

TABLE II: Comparison with existing methods on NEC Drone. Our SCP improves 4.0-7.4% over X3D and 23.1% over K-centered.

##### D. Results on Something-something V2

Something-something V2 is a challenging ground camera dataset for visual common sense because it requires models to understand the relationships between objects and actions. For example, to predict the category of a video, a model must understand that ”something bounces a ball” is different from ”something rolls a ball”. We evaluate our SCP’s reasoning and temporal modeling ability on Something-somethingV2.

As shown in Table III, our SCP improves 3.6% over MVITv1 and 1.0% over MVITv2, which illustrates the effectiveness of our proposed prompt learning and Auto-regressive temporal modeling.

##### E. Ablation Study

First, we conducted ablation studies on various prompts, including optical flow, large vision models, and learnable prompts (SCP), to verify their effectiveness. Then we further evaluate the effect of each component of our method.




Fig. 5: **Visualization** We first detect the interested target and generate the prompts, then predict the action.

Method	pretrain	Top-1 Acc.	Top-5 Acc.
TEA [56]	ImageNet 1k	65.1%	89.9%
MoViNet-A3 [57]	N/A	64.1%	88.8%
ViT-B-TimeSformer [17]	ImageNet 21k	62.5%	/
SlowFast R101, 8×8 [58]	Kinetics400	63.1%	87.6%
MViTv1-B, 16×4 [59]	Kinetics400	64.7%	89.2%
MViTv2-S, 16×4 [53]	Kinetics400	67.3%	91.0%
SCP (Ours)	Kinetics400	68.3%	91.4%

TABLE III: Comparison with the state-of-the-art results on the Something Something V2. Our SCP improves 3.6% over MViTv1 and 1.0% over strong SOTA MViTv2.

Component	Frame size	Accuracy
Baseline	224x224	71.54%
Baseline + ROI	224x224	73.61%
Baseline + ROI + Large Vision Model (SAM)	224x224	74.68%
Baseline + ROI + SCP	224x224	76.34%

TABLE IV: Ablation study in terms of the effect of different components in our method on the Okutama dataset. We evaluated ROI, Large Vision Model (SAM), and SCP. The experiments showed the effectiveness of our proposed methods.

**Different Prompts** To evaluate the effectiveness of different prompts, various prompts, including optical flow, large vision model (SAM [46]), and learnable prompts, are examined in this work. As shown in Table V, the large vision model and SCP achieved better accuracy.

**Effect of Each Component of Our Method** We also evaluated the effect of the components in our methods, including Region of Interest alignment (ROI), Large Vision

Method	Frame size	Accuracy
Baseline	224×224	71.54%
Baseline + Optical Flow	224×224	72.13%
Baseline + Large Vision Model (SAM)	224×224	74.68%
Baseline + SCP	224×224	75.93%

TABLE V: Ablation study in terms of different prompts on the Okutama dataset. We evaluated various prompts, including optical flow, a large vision model(SAM [46]), and SCP. From our experiment, the large vision model and SCP achieved better accuracy.

Model, and Learnable Prompt. As shown in Table IV, ROI can achieve 2.07% improvement, ROI combined with Large Vision Model can achieve 3.14% improvement, ROI combined with our SCP can achieve 4.80% improvement. The experiments showed the effectiveness of our proposed methods.

## V. CONCLUSION

We present a general prompt learning approach to alleviate the optimization burden by providing high-level texture descriptions or instructions associated with actions. These prompts enable the model to capture discriminative spatio-temporal patterns effectively. Our proposed SCP learns to dynamically generate prompts from a pool of prompt experts under different inputs. Our objective is to optimize prompts that guide the model’s predictions while explicitly learning input-invariant (prompt experts) and input-specific (data-dependent) prompt knowledge. We observe good accuracy improvements on the challenging datasets.

**Acknowledgement** This work was supported in part by ARO Grants W911NF2310046 W911NF2310352 and U.S. Army Cooperative Agreement W911NF2120076.

## REFERENCES

- [1] M. Barekatin, M. Martí, H.-F. Shih, S. Murray, K. Nakayama, Y. Matsuo, and H. Prendinger, “Okutama-action: An aerial view video dataset for concurrent human action detection,” in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2017, pp. 28–35.
- [2] J. Choi, G. Sharma, M. Chandraker, and J.-B. Huang, “Unsupervised and semi-supervised domain adaptation for action recognition from drones,” in *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2020, pp. 1717–1726.
- [3] R. Goyal, S. Ebrahimi Kahou, V. Michalski, J. Materzynska, S. Westphal, H. Kim, V. Haenel, I. Fruend, P. Yianilos, M. Mueller-Freitag *et al.*, “The” something something” video database for learning and evaluating visual common sense,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 5842–5850.
- [4] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig, “Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing,” *ACM Computing Surveys*, vol. 55, no. 9, pp. 1–35, 2023.
- [5] K. Simonyan and A. Zisserman, “Two-stream convolutional networks for action recognition in videos,” *Advances in neural information processing systems*, vol. 27, 2014.
- [6] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, “Large-scale video classification with convolutional neural networks,” in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2014, pp. 1725–1732.
- [7] L. Wang, Y. Qiao, and X. Tang, “Action recognition with trajectory-pooled deep-convolutional descriptors,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 4305–4314.
- [8] J. Sánchez, F. Perronnin, T. Mensink, and J. Verbeek, “Image classification with the fisher vector: Theory and practice,” *International journal of computer vision*, vol. 105, pp. 222–245, 2013.
- [9] G. Chéron, I. Laptev, and C. Schmid, “P-cnn: Pose-based cnn features for action recognition,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 3218–3226.
- [10] R. Girdhar, D. Ramanan, A. Gupta, J. Sivic, and B. Russell, “Actionvlad: Learning spatio-temporal aggregation for action classification,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 971–980.
- [11] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, “Learning spatiotemporal features with 3d convolutional networks,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4489–4497.
- [12] S. Ji, W. Xu, M. Yang, and K. Yu, “3d convolutional neural networks for human action recognition,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 1, pp. 221–231, 2012.
- [13] H. Zhang, L. Zhang, X. Qi, H. Li, P. H. Torr, and P. Koniusz, “Few-shot action recognition with permutation-invariant attention,” in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*. Springer, 2020, pp. 525–542.
- [14] X. Li, B. Shuai, and J. Tighe, “Directional temporal modeling for action recognition,” in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16*. Springer, 2020, pp. 275–291.
- [15] J. Carreira and A. Zisserman, “Quo vadis, action recognition? a new model and the kinetics dataset,” in *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6299–6308.
- [16] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučić, and C. Schmid, “Vivit: A video vision transformer,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 6836–6846.
- [17] G. Bertasius, H. Wang, and L. Torresani, “Is space-time attention all you need for video understanding?” in *ICML*, vol. 2, no. 3, 2021, p. 4.
- [18] X. Wang, S. Zhang, Z. Qing, Y. Shao, Z. Zuo, C. Gao, and N. Sang, “Oadtr: Online action detection with transformers,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 7565–7575.
- [19] X. Wang, R. Xian, T. Guan, C. M. de Melo, S. M. Nogar, A. Bera, and D. Manocha, “Aztr: Aerial video action recognition with auto zoom and temporal reasoning,” *arXiv preprint arXiv:2303.01589*, 2023.
- [20] R. Xian, X. Wang, D. Kothandaraman, and D. Manocha, “Pmi sampler: Patch similarity guided frame selection for aerial action recognition,” *arXiv preprint arXiv:2304.06866*, 2023.
- [21] R. Xian, X. Wang, and D. Manocha, “Mitfas: Mutual information based temporal feature alignment and sampling for aerial video action recognition,” *arXiv preprint arXiv:2303.02575*, 2023.
- [22] D. Kothandaraman, T. Guan, X. Wang, S. Hu, M.-S. Lin, and D. Manocha, “Far: Fourier aerial video recognition,” in *European Conference on Computer Vision*, 2022.
- [23] F. Petroni, T. Rocktäschel, P. Lewis, A. Bakhtin, Y. Wu, A. H. Miller, and S. Riedel, “Language models as knowledge bases?” *arXiv preprint arXiv:1909.01066*, 2019.
- [24] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, “Language models are few-shot learners,” *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [25] Z. Jiang, F. F. Xu, J. Araki, and G. Neubig, “How can we know what language models know?” *Transactions of the Association for Computational Linguistics*, vol. 8, pp. 423–438, 2020.
- [26] X. L. Li and P. Liang, “Prefix-tuning: Optimizing continuous prompts for generation,” *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, vol. abs/2101.00190, 2021.
- [27] Y. Tian, Y. Wang, D. Krishnan, J. B. Tenenbaum, and P. Isola, “Rethinking few-shot image classification: a good embedding is all you need?” in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*. Springer, 2020, pp. 266–282.
- [28] Z. Li, P. Xu, F. Liu, and H. Song, “Towards understanding in-context learning with contrastive demonstrations and saliency maps,” *arXiv preprint arXiv:2307.05052*, 2023.
- [29] N. Poerner, U. Waltinger, and H. Schütze, “E-bert: Efficient-yet-effective entity embeddings for bert,” in *Findings*, 2019.
- [30] T. Shin, Y. Razeghi, R. L. L. IV, E. Wallace, and S. Singh, “Eliciting knowledge from language models using automatically generated prompts,” *ArXiv*, vol. abs/2010.15980, 2020.
- [31] X. Han, W. Zhao, N. Ding, Z. Liu, and M. Sun, “Ptr: Prompt tuning with rules for text classification,” *AI Open*, vol. 3, pp. 182–192, 2022.
- [32] B. Lester, R. Al-Rfou, and N. Constant, “The power of scale for parameter-efficient prompt tuning,” *arXiv preprint arXiv:2104.08691*, 2021.
- [33] Z. Zhong, D. Friedman, and D. Chen, “Factual probing is [mask]: Learning vs. learning to recall,” *arXiv preprint arXiv:2104.05240*, 2021.
- [34] Y. Jiang, A. Gupta, Z. Zhang, G. Wang, Y. Dou, Y. Chen, L. Fei-Fei, A. Anandkumar, Y. Zhu, and L. Fan, “Vima: General robot manipulation with multimodal prompts,” *arXiv preprint arXiv:2210.03094*, 2022.
- [35] I. Singh, V. Blukis, A. Mousavian, A. Goyal, D. Xu, J. Tremblay, D. Fox, J. Thomason, and A. Garg, “Progprompt: Generating situated robot task plans using large language models,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 11 523–11 530.
- [36] S. Vemprala, R. Bonatti, A. Buckler, and A. Kapoor, “Chatgpt for robotics: Design principles and model abilities,” *Microsoft Auton. Syst. Robot. Res.*, vol. 2, p. 20, 2023.
- [37] X. Wang, A. Liang, J. Liang, M. Lin, Y. Lou, and S. Yang, “Icar: Image-based complementary auto reasoning,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 6, 2024, pp. 5633–5641.
- [38] X. Wang, X. Chu, C. Han, and X. Zhang, “Scsc: Spatial cross-scale convolution module to strengthen both cnns and transformers,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 731–741.
- [39] X. Wang, J. Liang, C.-K. Wang, K. Deng, Y. Lou, M. Lin, and S. Yang, “Vlap: Efficient video-language alignment via frame prompting and distilling for video question answering,” *arXiv preprint arXiv:2312.08367*, 2023.
- [40] X. Wang, J. Liang, C.-K. Wang, K. Deng, Y. M. Lou, M. Lin, and S. Yang, “Vila: Efficient video-language alignment for video question answering,” in *ECCV 2024*, 2024. [Online]. Available: <https://www.amazon.science/publications/vila-efficient-video-language-alignment-for-video-question-answering>

- [41] Y. Rao, W. Zhao, G. Chen, Y. Tang, Z. Zhu, G. Huang, J. Zhou, and J. Lu, "Denseclip: Language-guided dense prediction with context-aware prompting," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 18 082–18 091.
- [42] C. Ju, T. Han, K. Zheng, Y. Zhang, and W. Xie, "Prompting visual-language models for efficient video understanding," in *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*. Springer, 2022, pp. 105–124.
- [43] K. Zhou, J. Yang, C. C. Loy, and Z. Liu, "Learning to prompt for vision-language models," *International Journal of Computer Vision*, vol. 130, pp. 2337 – 2348, 2021.
- [44] F. Liu, K. Lin, L. Li, J. Wang, Y. Yacoob, and L. Wang, "Aligning large multi-modal model with robust instruction tuning," *arXiv preprint arXiv:2306.14565*, 2023.
- [45] F. Liu, X. Wang, W. Yao, J. Chen, K. Song, S. Cho, Y. Yacoob, and D. Yu, "Mmc: Advancing multimodal chart understanding with large-scale instruction tuning," *arXiv preprint arXiv:2311.10774*, 2023.
- [46] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo *et al.*, "Segment anything," *arXiv preprint arXiv:2304.02643*, 2023.
- [47] F. Yang, S. Sakti, Y. Wu, and S. Nakamura, "A framework for knowing who is doing what in aerial surveillance videos," *IEEE Access*, vol. 7, pp. 93 315–93 325, 2019.
- [48] A. M. Algamdi, V. Sanchez, and C.-T. Li, "Dronecaps: Recognition of human actions in drone videos using capsule networks with binary volume comparisons," in *2020 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2020, pp. 3174–3178.
- [49] M. Zolfaghari, K. Singh, and T. Brox, "Eco: Efficient convolutional network for online video understanding," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 695–712.
- [50] P. ZHANG, P. Wei, and S. Han, "Capsnets algorithm," in *Journal of Physics: Conference Series*, vol. 1544, no. 1. IOP Publishing, 2020, p. 012030.
- [51] S. K. Yadav, A. Luthra, E. Pahwa, K. Tiwari, H. Rathore, H. M. Pandey, and P. Corcoran, "Droneattention: Sparse weighted temporal attention for drone-camera based activity recognition," *Neural Networks*, vol. 159, pp. 57–69, 2023.
- [52] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10 012–10 022.
- [53] Y. Li, C.-Y. Wu, H. Fan, K. Mangalam, B. Xiong, J. Malik, and C. Feichtenhofer, "Mvitv2: Improved multiscale vision transformers for classification and detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4804–4814.
- [54] C. Feichtenhofer, "X3d: Expanding architectures for efficient video recognition," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 203–213.
- [55] S. H. Park, J. Tack, B. Heo, J.-W. Ha, and J. Shin, "K-centered patch sampling for efficient video recognition," in *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*. Springer, 2022, pp. 160–176.
- [56] Y. Li, B. Ji, X. Shi, J. Zhang, B. Kang, and L. Wang, "Tea: Temporal excitation and aggregation for action recognition," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 909–918.
- [57] D. Kondratyuk, L. Yuan, Y. Li, L. Zhang, M. Tan, M. Brown, and B. Gong, "Movidets: Mobile video networks for efficient video recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 16 020–16 030.
- [58] C. Feichtenhofer, H. Fan, J. Malik, and K. He, "Slowfast networks for video recognition," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 6202–6211.
- [59] H. Fan, B. Xiong, K. Mangalam, Y. Li, Z. Yan, J. Malik, and C. Feichtenhofer, "Multiscale vision transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6824–6835.