

PlanLLM: Video Procedure Planning with Refinable Large Language Models

Dejie Yang¹, Zijing Zhao¹, Yang Liu^{1,2*}

¹ Wangxuan Institute of Computer Technology, Peking University

² State Key Laboratory of General Artificial Intelligence, Peking University
{yjdj,zijingzhao}@stu.pku.edu.cn, yangliu@pku.edu.cn

Abstract

Video procedure planning, *i.e.*, planning a sequence of action steps given the video frames of start and goal states, is an essential ability for embodied AI. Recent works utilize Large Language Models (LLMs) to generate enriched action step description texts to guide action step decoding. Although LLMs are introduced, these methods decode the action steps into a closed-set of one-hot vectors, limiting the model’s capability of generalizing to new steps or tasks. Additionally, fixed action step descriptions based on world-level commonsense may contain noise in specific instances of visual states. In this paper, we propose PlanLLM, a cross-modal joint learning framework with LLMs for video procedure planning. We propose an LLM-Enhanced Planning module which fully uses the generalization ability of LLMs to produce free-form planning output and to enhance action step decoding. We also propose Mutual Information Maximization module to connect world-level commonsense of step descriptions and sample-specific information of visual states, enabling LLMs to employ the reasoning ability to generate step sequences. With the assistance of LLMs, our method can both closed-set and open vocabulary procedure planning tasks. Our PlanLLM achieves superior performance on three benchmarks, demonstrating the effectiveness of our designs.

Code — <https://github.com/idejie/PlanLLM>

Introduction

Mastering procedure planning is of great importance for building artificial intelligence systems capable of comprehending and imitating human actions, and eventually assisting humans in solving goal-directed problems (Tellex et al. 2011; Jansen 2020; Ahn et al. 2022; Mishra et al. 2018; Mysore et al. 2019). Previous research (Chang et al. 2020) has pointed out that instructional videos are natural resources for acquiring the skill of procedure planning, proposing the problem of video procedure planning. For example, given the visual frames of start state “raw food” and goal states “hummus”, the model is expected to produce a sequence of action steps to achieve the goal, *i.e.*, first “peel garlics”, then “add to crusher”, and finally “crush food”.

*Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

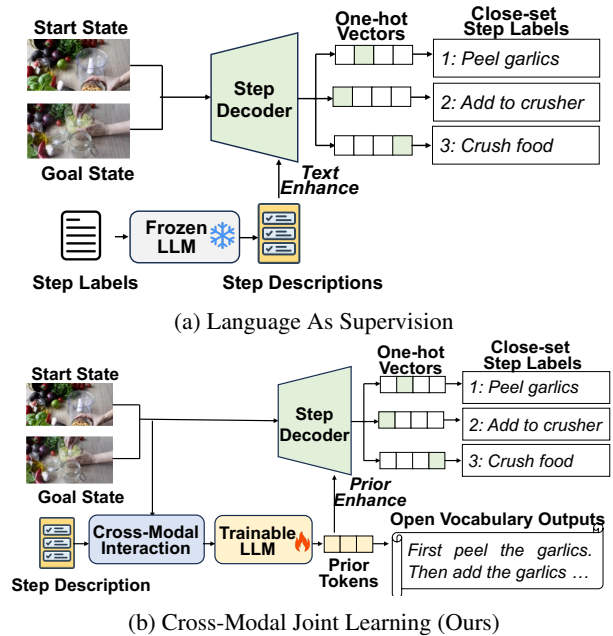


Figure 1: Language As Supervision methods use a frozen LLM to enhance the textual description of step labels, and are limited to closed-set one-hot vector predictions. In contrast, our Cross-Modal Joint Learning framework utilizes a **trainable LLM** that allows our model to output both one-hot vectors and **free-form open-vocabulary** step descriptions, providing stronger generalization for predicting new steps in new datasets.

Early works employ both visual keyframes and textual labels of action steps to train the video procedure planning model (fully-supervised), which requires labor-intensive annotations. Recent weakly-supervised methods rely only on the textual step sequences during training, and have gained more attention.

Existing weakly supervised video procedure planning methods (Chang et al. 2020; Wang et al. 2023b; Li et al. 2023c) use an encoder-decoder structure, first learning a latent visual space and then decoding action step labels from this space. These methods primarily treat video procedure planning as a visual-based action classification task on a

closed-set of labels. Recent approaches (Niu et al. 2024; Zhao et al. 2022; Wang et al. 2023a) introduce language models to take advantage of semantic information from textual action steps. For instance, as shown in Figure 1a, language as supervision method (Niu et al. 2024) uses a frozen large language model (LLM) to generate enriched descriptions of action steps based on world-level commonsense extracted from the LLM. By combining these language descriptions with visual states, the model is better trained to decode action steps, resulting in improved performance.

Though making progress, existing methods using LLMs still face issues: (1) They use a step decoder that treats each action step as a one-hot vector in a closed set, which can not generalize to open-world scenarios. Moreover, these approaches ignore the semantic relationships between steps. For example, “peel food” and “crush food” are likely to occur together, but decoding within a closed set cannot handle new steps or novel tasks. (2) Methods that employ a frozen LLM generate consistent action step descriptions based on world-level commonsense, but these may not always be accurate for specific visual start and goal states. For instance, “crush food” might generally suggest “mixed food pieces,” but if the visual end state shows hummus in a food crusher, the correct outcome would be “hummus in the food crusher.” The lack of sample-specific visual cues limits the LLM’s ability to apply contextual commonsense in step sequence reasoning.

To address the above challenges, we propose PlanLLM, a cross-modal joint learning framework with LLMs for video procedure planning. As shown in Figure 1b, we utilize LLM’s planning priors to enhance the model’s generalization ability through an LLM Enhanced Planning module, producing free-form and open-vocabulary procedure planning outputs. Notably, we are the first to use a **trainable LLM** to output action step sequences that are **not confined to a predefined set** and can generalize to new steps or novel tasks. To build cross-modal interaction between visual states and step descriptions, we introduce a Mutual Information Maximization mechanism to connect sample-specific visual information with world-level commonsense descriptions. This allows the LLM to use its generalizable reasoning ability, grounded in sample-specific visual commonsense, to enhance procedure learning. We propose a progressive alignment scheme during training. We first freeze the LLM to align the visual embeddings to textual space using a mutual information maximization loss, and then fine-tune the LLM jointly with other modules for procedure planning. Finally, the output from the LLM branch can contribute to conventional decoding by providing insights into semantic relevance of steps, thereby enabling our approach to handle both closed-set and open-vocabulary procedure planning tasks.

Overall, the main contributions of our work are as follows: (1) We propose PlanLLM, a cross-modal joint learning framework with a trainable LLM for video procedure planning, which is the first to consider both the model’s generalization ability in the open world and its planning performance on closed sets. (2) We introduce the LLM-Enhanced Planning module, which utilizes the LLM’s textual planning

priors to enhance the capability of video procedure planning models. Additionally, we propose a Mutual Information Maximization module to connect world-level commonsense in step descriptions with sample-specific visual state information, enabling LLMs to utilize reasoning abilities for generating step sequences. (3) PlanLLM achieves superior performance on three commonly used datasets: COIN (Tang et al. 2019), CrossTask (Zhukov et al. 2019), and NIV (Alayrac et al. 2016), and improves cross-dataset procedure planning performance, demonstrating the effectiveness and generalization of our method.

Related Works

Video Procedure Planning. Video procedure planning task (Chang et al. 2020) aims to generate a sequence of action steps based on visual observations of the start and goal states. Early works (Srinivas et al. 2018; Chang et al. 2020; Bi, Luo, and Xu 2021; Sun et al. 2022) train the models in a fully supervised manner, requiring annotations for textual step sequences and intermediate visual states. Recently, weakly supervised approaches, which rely only on textual step sequences during training, have gained attention due to their lower annotation costs. These methods (Chang et al. 2020; Wang et al. 2023b; Li et al. 2023c; Sun et al. 2022) typically use an encoder-decoder structure to learn a latent visual space and decode action step labels from it. PDPP (Wang et al. 2023b) introduces a diffusion probabilistic model for generating intermediate action labels, while KEPP (Nagasinghe et al. 2024) proposes a probabilistic procedural knowledge graph to improve planning. Instead of focusing on visual representations, several methods (Zhao et al. 2022; Wang et al. 2023a; Liu et al. 2023b) utilize language encoders to capture the textual information of action steps. SCHEMA (Niu et al. 2024) leverages large language models like GPT-3.5 to extend the descriptions of textual action step labels, then uses the extended descriptions to guide the decoding of action steps.

In this paper, we achieve weakly supervised video procedure planning relying only on textual action step labels during training. Unlike existing methods that decode action steps into a closed-set of one-hot vectors, we harness the generalization ability of LLMs to generate free-form procedure planning outputs. Our approach not only enhances the decoding of closed-set labels, but also allows the model to generate free-form outputs for new action steps and planning tasks. Note that the step descriptions used in both our method and the Language As Supervision approaches such as (Niu et al. 2024) are dataset-specific, aggregated from step labels and enhanced by the LLM, and are independent of the sample-specific visual states.

LLMs for Planning. Pre-trained large language models (LLMs) (Brown et al. 2020; Chowdhery et al. 2023) have demonstrated their ability to generate high-level plans (Xie et al. 2023; Wang et al. 2023c; Liu et al. 2023a; Ye et al. 2024; Yang and Liu 2024) and video understanding (Zheng et al. 2023). Some works train LLMs to generate verbal reasoning traces and text actions (Yao et al. 2022; Ting Lei and Liu 2024) in an interleaved manner, allowing the model to

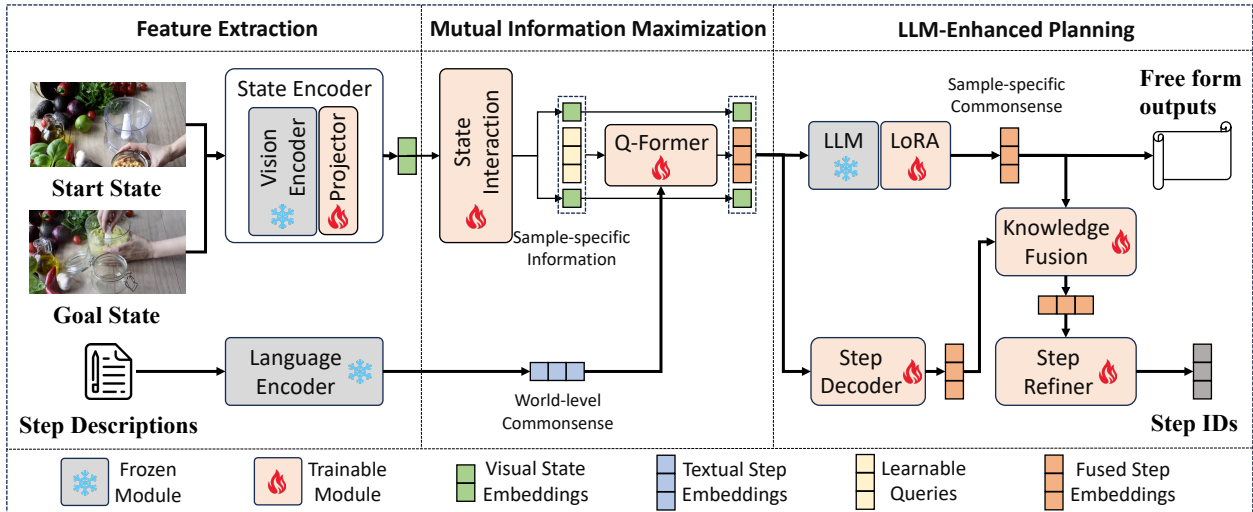


Figure 2: **The framework of our PlanLLM.** PlanLLM mainly consists of three parts: Feature Extraction, Mutual Information Maximization and LLM Enhanced Planning.

interface with external sources (e.g., knowledge bases or environments) for additional information. Other works leverages LLMs for few-shot planning (Song et al. 2023), interaction (Lei et al. 2024) and temporal understanding (Zheng et al. 2024), enhancing them with physical grounding to generate and update plans based on the current environment. These enhanced LLMs can follow natural language instructions to complete complex tasks in visually-perceived environments. Although achieving success, they treat LLMs as textual planners, relying solely on textual descriptions of the environment, which may not fully capture the complexity of visual states, thus limiting visual reasoning abilities.

Unlike existing methods that focus on textual planning with LLMs, our approach exploits the capabilities of LLMs for video procedure planning. We connect sample-specific visual state information with world-level commonsense from LLMs, training the model to learn multi-modal representations for a better understanding of the environment.

Methodology

Problem Formulation and Overview

Video procedure planning task aims to generate a sequence of action steps based on visual observations of the start and goal states. As shown in Figure 2, given the visual observations (images) of the start frame v_0 and the goal v_T , the target is to plan a procedure, represented as a sequence of action steps $\hat{\pi} = a_{1:T}$, which are sequentially performed to achieve the transition from v_0 to v_T . The video procedure planning problem can be formulated as $p(a_{1:T}|v_0, v_T)$. We follow the weakly supervised video procedure planning setting (Wang et al. 2023b; Niu et al. 2024), which rely only on the ground truth textual action sequence $a_{1:T}$ without requiring corresponding intermediate visual states (keyframes between two states). Following previous work (Niu et al. 2024), we use the textual information of action step labels,

treating video procedure planning as a multi-modal task. Notably, the textual information of step descriptions are collected from the training set and shared across different visual state inputs, making them independent of specific inputs.

As depicted in Figure 2, our PlanLLM encompasses three parts: (1) *Feature Extraction*: We use a visual encoder and a language encoder to extract features from the visual inputs $\{v_0, v_T\}$ representing the start and goal states, capturing sample-specific visual details essential for task context, and from the textual descriptions D of action steps, capturing world-level commonsense. (2) *Mutual Information Maximization*: To fuse sample-specific visual information with world-level textual commonsense, we introduce Q-Former architecture which takes visual embeddings processed by the state interaction module and step description embeddings as inputs to generate cross-modal joint step embeddings. (3) *LLM Enhanced Planning*: To address the limitations of one-hot decoding of action steps, which ignores semantic relevance between steps and cannot generalize to new steps or tasks, we propose the LLM-Enhanced Planning module. This module takes multi-modal step embeddings from the Mutual Information Maximization module as inputs, allowing the LLM to be fine-tuned via LoRA (Hu et al. 2022) to directly produce free-form procedure planning outputs, capable of handling new steps and tasks. Additionally, when combined with our proposed knowledge fusion and step refinement module, the general reasoning ability of LLMs is effectively leveraged to refine the original step decoding process, improving performance.

Feature Extraction

We first introduce a visual state encoder E_v and a language encoder E_l to extract features from sample-wise visual state inputs $\{v_0, v_T\}$ and world-level action step descriptions D , respectively.

Visual State Feature Extraction: The visual state encoder

takes the video frames of the start and goal states as input and outputs embeddings that represent sample-specific visual information. We adopt a pre-trained vision backbone f_v together with a trainable projection layer p_v as the visual state encoder E_v . The embedding x_t of visual state $v_t, t \in \{0, T\}$ is extracted by:

$$x_t = E_v(v_t) = p_v(f_v(v_t)), t \in \{0, T\} \quad (1)$$

Textual Step Feature Extraction: As the original step texts are too concise, following (Niu et al. 2024), we extend all the action step labels $A = \{a_i\}_{i=1}^N$ to enriched action step descriptions $D = \{d_i\}_{i=1}^N$ using large language models (such as GPT-3.5) with world-level commonsense knowledge:

$$d_i = \text{LLM}(\text{prompt}, a_i), i \in \{1, 2, \dots, N\} \quad (2)$$

where N denotes the number of all possible action step in the dataset. The details of description generation are provided in supplemental materials. We adopt a fixed pre-trained language encoder E_l to extract the embedding y_i of the action step descriptions d_i :

$$y_i = E_l(d_i) \quad (3)$$

Mutual Information Maximization

As the visual state embeddings contain only sample-specific information while the textual step embeddings capture only world-level commonsense knowledge from LLMs, we propose a Mutual Information Maximization module to integrate both embeddings for a better representation of action steps.

Visual State Interaction: To learn the mutual context awareness of visual states, we introduce a self-attention layer to process the original embeddings x_0^v, x_T^v to interacted embeddings x_0^v, x_T^v :

$$x_0^v, x_T^v = \text{SelfAtten}(x_0, x_T) \quad (4)$$

Interacted visual embeddings provide a better representation when integrated with world-level textual step embeddings and when aligned with LLM input space.

Mutual Information Maximization(MIM) Q-Former: To integrate sample-specific visual state embeddings and textual step embeddings with world-level commonsense, we adopt a Q-Former (Li et al. 2023a) architecture to learn fused step embeddings. Specifically, we use the integrated visual start and goal state embeddings x_0^v, x_T^v as the image input tokens for Q-Former, and the textual action step embeddings $x_{1:T}^a$ corresponding to the ground-truth steps $a_{1:T}$ as the textual input tokens. Using learnable queries $q_{1:T}$, the Q-Former outputs the integrated step embeddings $x_{1:T}^q$:

$$x_{1:T}^q = \text{QFormer}(q_{1:T}; [x_0^v, x_T^v]; a_{1:T}) \quad (5)$$

Following Q-Former, we optimize the Vision-Language Contrastive (VLC) loss and Vision-Language Matching (VLM) loss to maximize the mutual information between vision and language embeddings. Within a batch, we treat the visual state embeddings and their corresponding ground truth step embeddings as positive pairs, and unmatched vision and textual embeddings as negative pairs. The Vision-Language Contrastive (VLC) loss contrasts the embedding

similarity of a positive pair against the negative ones:

$$\mathcal{L}_{VLC} = -\log \frac{\sum_j e^{s(x^v[j], x^q[j])}}{\sum_j e^{s(x^v[j], x^q[j])} + \sum_{j \neq k} e^{s(x^v[j], x^q[k])}} \quad (6)$$

where $s(\cdot, \cdot)$ denotes the similarity function, and $x^v[j]$ and $x^q[j]$ represent the interacted visual state embeddings from the visual state interaction module and the integrated step embeddings from Q-Former, respectively, for the j th sample in a batch, with $j, k \in \{1, 2, \dots, B\}$, where B is the batch size. For the Vision-Language Matching (VLM) loss, the model is tasked with a binary classification to predict whether an image-text pair is positive (matched) or negative (unmatched):

$$\mathcal{L}_{VLM} = \sum_{j,k} \mathcal{L}_{BCE}(x^v[j], x^q[k]; \mathbf{1}(j=k)) \quad (7)$$

The total Mutual Information Maximization (MIM) loss combines VLC loss and VLM loss:

$$\mathcal{L}_{MIM} = \mathcal{L}_{VLC} + \mathcal{L}_{VLM}. \quad (8)$$

LLM Generated Free Form Planning

Different from previous methods that decode action steps into a closed set of one-hot vectors, our approach fully leverages the LLM’s generalization ability to directly generate free-form procedure planning outputs, enabling it to handle open vocabulary procedure planning and new planning tasks.

Free Form Procedure Planning Output: We introduce a generative LLM, i.e., Vicuna-7B (Chiang et al. 2023). Taking the integrated step embeddings $x_{1:T}^q$ from the MIM Q-Former along with the visual state embeddings x_0^v, x_T^v as inputs, the LLM’s encoder provides hidden state embeddings $h_{1:T}$, representing the enhanced action step embeddings with sample-specific commonsense, based on the general reasoning ability of the LLM. Its decoder then directly generates free-form output tokens O as the captioning of the action steps:

$$h_{1:T} = \text{LLM}_{enc}(x_0^v, x_T^v, x_{1:T}^q) \quad (9)$$

$$O = \text{LLM}_{dec}(h_{1:T}) \quad (10)$$

Progressive Alignment Training Scheme: To effectively train the modules to produce free-form procedure planning outputs, we propose a two-stage progressive training scheme. Specifically, in the first stage, we fix the generative LLM and train the feature extractors and MIM modules using \mathcal{L}_{MIM} as defined in Equation 8. This stage aligns the visual state embeddings and fused action step embeddings with the LLM’s input space, enabling the LLM to understand and reason with this information. In the second stage, we fine-tune the LLM using LoRA (Hu et al. 2022), along with other trainable modules, optimizing an Action Step Captioning loss \mathcal{L}_{ASC} and the original MIM loss \mathcal{L}_{MIM} . The \mathcal{L}_{ASC} loss enables the LLM to learn step caption generation, with the ground truth action step texts organized into a formatted sentence. With our progressive alignment training scheme, the LLM is effectively trained to produce free-form procedure planning outputs.

LLM Enhanced Close-set Step Decoding

Existing methods decode action steps as a closed set of one-hot vectors. With the help of LLMs which understand the semantic information of different action step labels, the conventional action step decoding process can also be enhanced to improve performance.

Action Step Decoder: Following previous work (Niu et al. 2024), we employ a cross-attention module to process the visual state embeddings, action step description embeddings, and learnable queries to produce final action step representations, which are then passed through a classifier to distinguish specific action steps. The final action step representations $r_{1:T}^{SD}$ are obtained by:

$$r_{1:T}^{SD} = \text{CrossAtten}(q_{1:T}^{SD}, [x_0^v, x_T^v, x_{1:T}^q]) \quad (11)$$

where x_0^v, x_T^v denotes the visual start and end state embeddings, $x_{1:T}^q$ denotes fused action step embeddings from MIM Q-Former, and $q_{1:T}^{SD}$ denotes learnable queries.

Knowledge Fusion: To integrate the semantic understanding knowledge of LLMs, we fuse the sample-specific commonsense of hidden states embeddings $h_{1:T}$ from the LLM encoder with the final action step representations $r_{1:T}^{SD}$ from the step decoder to produce step representations:

$$r_{1:T}^{KF} = \text{CrossAtten}(r_{1:T}^{SD}; h_{1:T}) \quad (12)$$

With knowledge fusion module, the knowledge fused step representations are aware of the semantic relevance between action step labels with the help of LLM general knowledge.

Step Refiner: Given the knowledge fused representations $r_{1:T}^{KF}$ as query, we finally refine the step representations from the embeddings of the action step descriptions with another cross-attention module :

$$r_{1:T}^{SR} = \text{CrossAtten}(r_{1:T}^{KF}; y), \quad (13)$$

where y is the embeddings of the action step descriptions via Equation (3). The refined step representations $r_{1:T}^{SR}$ are projected into a one-hot vector $\hat{a}_{1:T}$ for classification, supervised by cross-entropy loss as the Step Decoding (SD) loss:

$$\mathcal{L}_{SD} = \mathcal{L}_{CE}(\hat{a}_{1:T}, a_{1:T}) \quad (14)$$

Where \mathcal{L}_{CE} denotes cross-entropy loss, and $a_{1:T}$ denotes ground truth action step labels.

Training and Inference

Training: As has been discussed in *Progressive Alignment Training Scheme*, in the first stage, we fix the visual backbone encoder, the language encoder and the generative LLM, and only train the Mutual Information Maximization Module with \mathcal{L}_{MIM} in Equation 8 to align the visual state embeddings and the action step embeddings to the LLM input space. In the second stage, we jointly fine-tune the LLM through LoRA(Hu et al. 2022) and the other modules with the corresponding losses. The overall objective can be elaborated as:

$$\mathcal{L} = \mathcal{L}_{MIM} + \mathcal{L}_{ASC} + \mathcal{L}_{SD} \quad (15)$$

where \mathcal{L}_{ASC} is the Action Step Captioning loss, \mathcal{L}_{SD} is the State Decoding loss.

Inference: Our method can handle both closed-set action step classification tasks and free-form open-vocabulary procedure planning tasks. For conventional action step classification, our proposed LLM-Enhanced step decoding branch outputs the action step IDs following (Niu et al. 2024; Wang et al. 2023b; Zhao et al. 2022). For open-vocabulary planning, the generative LLM provides free-form procedure planning outputs and encodes the captions and new textual action step labels into vectors using a frozen language encoder. We then retrieve the top T action labels based on the similarity between captions and textual action step labels, where T is the number of action steps in a sequence.

Experiments

Evaluation

Datasets We employ three commonly used video datasets: CrossTask (Zhukov et al. 2019), NIV (Alayrac et al. 2016), and COIN (Tang et al. 2019). The CrossTask comprises 2,750 videos, illustrating 18 unique procedural tasks. NIV (Narrated Instructional Videos) contains 150 videos encompassing five procedures. COIN is the largest dataset in our evaluation, boasting 11,827 videos and 778 procedures.

Metrics We evaluate performance using three metrics as outlined in (Niu et al. 2024): (1) Mean Intersection over Union (mIoU) measures whether the model correctly identifies the set of steps required to execute the plan, regardless of the order of actions. (2) Mean Accuracy (mAcc) compares the predicted and actual action sequences element-wise, considering the correct order of actions. (3) Success Rate (SR) assesses the plan’s success only if it exactly matches the ground truth, requiring precise correspondence between the predicted and actual sequences. Note that SR is a stricter metric, offering a stringent assessment of the model’s performance and requires exact correspondence between the predicted and actual sequences to be successful.

Implementation details Following recent advancements (Niu et al. 2024; Zhao et al. 2022; Wang et al. 2023b,a), we leverage a pretrained S3D network (Miech et al. 2020) as the visual backbone and the textual encoder of the pretrained CLIP (Radford et al. 2021) as our language encoder. Additionally, we deploy the trainable Q-Former architecture initialized from BLIP2 (Li et al. 2023a). Following (Li et al. 2023b), we utilize the pretrained Vicuna (Chiang et al. 2023) as the Large Language Model (LLM). During the frozen-LLM training stage, we set the learning rate to 1×10^{-4} for the Q-Former and 1×10^{-3} for other modules, training the model with a batch size of 32 on NVIDIA A800 GPUs.

Results

Comparisons on CrossTask Table 1 shows the comparisons between our method and others on CrossTask. (1) Compared to previous fully-supervised methods (Ehsani et al. 2018; Abu Farha and Gall 2019) that require intermediate visual states as inputs, recent weakly-supervised methods (Zhao et al. 2022; Li et al. 2023c; Liu et al. 2023b; Niu et al. 2024) achieve better results due to more advanced

Method	Supervision	t=3			t=4		
		SR↑	mAcc↑	mIoU↑	SR↑	mAcc↑	mIoU↑
UPN(Srinivas et al. 2018)	V+A	2.89	24.39	31.56	1.19	21.59	27.85
DDN(Chang et al. 2020)	V+A	12.18	31.29	47.48	5.97	27.10	48.46
Ext-GAIL(Bi, Luo, and Xu 2021)	V+A	21.27	49.46	61.70	16.41	43.05	60.93
PlaTe(Sun et al. 2022)	V+A	16.00	36.17	65.91	14.00	35.29	55.36
P ³ IV(Zhao et al. 2022)	A	23.34	49.96	73.89	13.40	44.16	70.01
EGPP(Wang et al. 2023a)	A	20.14	38.36	67.29	11.32	18.85	70.53
PDPP(Wang et al. 2023b)	A	37.20	64.67	66.57	21.48	57.82	65.13
SkipPlan(Li et al. 2023c)	A	28.85	61.18	74.98	15.56	55.64	70.30
LangFirst(Liu et al. 2023b)	A	25.01	53.79	75.43	14.11	47.93	73.21
KEPP(Nagasinghe et al. 2024)	A	38.12	<u>64.74</u>	67.15	24.15	59.05	66.04
SCHEMA(Niu et al. 2024)	A	<u>38.93</u>	63.80	<u>79.82</u>	<u>24.50</u>	58.48	<u>78.42</u>
PlanLLM(Ours)	A	39.74	65.78	81.50	27.54	<u>59.01</u>	77.58

Table 1: Comparisons on CrossTask for procedure planning with prediction horizon $t \in \{3, 4\}$. *Supervision* denotes the supervision type, where *V* denotes the methods use intermediate visual states (frames between start and goal states) as supervisions, and *A* only uses the action or task category without visual states.

model designs. We also follow the weakly-supervised setting in our method. (2) Methods (Zhao et al. 2022; Li et al. 2023c; Liu et al. 2023b; Niu et al. 2024) that use embeddings of action or task categories (Supervision is A) are better than those using intermediate visual states as supervision (Supervision is V). This may be because embeddings offer more robust language representations for planning, unlike visual states, which have limited quantity and representation. (3) The performance of the models is better at $t = 3$ than 4 since the longer action sequence increases prediction difficulty. (4) At horizon $t = 3$, we surpass the previous best by +0.81% on SR (a stricter, order-sensitive metric). At $t = 4$, our method also shows gains on SR (+3.04%). These results show the effectiveness of incorporating LLMs.

Comparisons on NIV and COIN Table 2 shows the comparisons on NIV and COIN datasets. (1) The methods using the embeddings of action or task categories obtain better performance thanks to the powerful representation of language encoders. (2) On NIV, compared to previous best, our method obtains gains of +0.41% on SR (a more strict metric, order-sensitive) when prediction horizon $t = 4$. This may be because LLM can provide commonsense for long action sequences. (3) On COIN dataset, our method remains the best across metrics. Specifically, our method makes improvement on SR (+1.13%), mAcc(+1.19%) and mIoU(+0.26%), when prediction horizon $t = 3$. And our PlanLLM also has gains (+1.66% on SR, +1.67% on mAcc and +1.56% on mIoU) when $t = 4$. The results suggest that our model performs well with different scales.

Comparisons on cross-dataset generalization Unlike to previous methods that only evaluate performance within a dataset, we provide cross-dataset comparisons in Table 3, where methods are trained on one dataset and evaluated on another. (1) We generate free-form procedure planning outputs, treating them as step captions. These captions, along with new textual action step labels, are encoded into vectors using a frozen language encoder. We then retrieve the top T action labels based on the similarity between cap-

tions and action labels, where T represents the number of steps in a sequence. We applied a similar encoding approach to step IDs. (2) Our method outperforms the previous best method (Niu et al. 2024) across different datasets and metrics, likely because PlanLLM generates free-form planning descriptions by incorporating LLMs into video procedure planning. (3) Both methods perform better on “train on COIN, test on CrossTask” compared to “train on CrossTask, test on COIN.” This might be due to COIN having more samples than CrossTask, leading to more optimal models when trained on COIN, but it is more challenging vice versa.

Ablation Studies

Effectiveness of proposed components In Table 4, we evaluate the effectiveness of our proposed components by establishing a baseline method without Mutual Information Maximization (MIM) or LLM-enhanced planning (1st row, *w/o* MIM/LLM). (1) Adding the MIM module (2nd row, *w/o* LLM) improves performance across various metrics and prediction horizons (+0.96/0.56% on SR with $t = 3/4$), showing its ability to integrate visual cues with commonsense information from textual embeddings, enhancing video procedure planning. (2) Introducing the LLM component (3rd row, *w/o* MIM) further boosts performance compared to the baseline (+1.00/1.38% on SR with $t = 3/4$), indicating that the LLM provides valuable planning priors and improves action sequence accuracy. (3) Combining both MIM and LLM modules results in the greatest improvement (last row, +1.71/3.68% on SR with $t = 3/4$), highlighting their complementarity and collective effectiveness.

Effectiveness of progressive multi-modal training In Table 5, we validate the effectiveness of progressive multi-modal training. (1) We create a variant with a frozen LLM (1st row), where all LLM parameters are fixed during training. Progressive training (3rd row) outperforms the frozen LLM, indicating that fine-tuning LLM parameters is more effective for video procedure planning. (2) One-stage training (2nd row) does not achieve optimal results compared to

Horizon	Method	Supervision	NIV			COIN		
			SR \uparrow	mAcc \uparrow	mIoU \uparrow	SR \uparrow	mAcc \uparrow	mIoU \uparrow
$t = 3$	DDN (Chang et al. 2020)	V+A	18.41	32.54	56.56	13.9	20.19	64.78
	Ext-GAIL (Bi, Luo, and Xu 2021)	V+A	22.11	42.20	65.93	-	-	-
	P ³ IV (Zhao et al. 2022)	A	24.68	<u>49.01</u>	74.29	15.4	21.67	76.31
	EGPP (Wang et al. 2023a)	A	26.05	51.24	75.81	19.57	31.42	<u>84.95</u>
	PDPP(Wang et al. 2023b)	A	30.20	48.45	57.28	21.33	45.62	51.82
	SkipPlan(Li et al. 2023c)	A	-	-	-	23.65	47.12	78.44
	LangFirst(Liu et al. 2023b)	A	-	-	-	28.35	<u>53.14</u>	78.56
	SCHEMA(Niu et al. 2024)	A	<u>27.93</u>	41.64	76.77	<u>32.09</u>	49.84	83.83
	KEPP(Nagasinghe et al. 2024)	A	24.44	43.46	86.67	<u>20.25</u>	39.87	-
	PlanLLM(Ours)	A	26.74	42.97	<u>77.23</u>	33.22	54.33	85.21
$t = 4$	DDN (Chang et al. 2020)	V+A	15.97	27.09	53.84	11.13	17.71	68.06
	Ext-GAIL (Bi, Luo, and Xu 2021)	V+A	19.91	36.31	53.84	-	-	-
	P ³ IV (Zhao et al. 2022)	A	20.14	38.36	67.29	11.32	18.85	70.53
	EGPP (Wang et al. 2023a)	A	21.37	41.96	74.90	13.59	26.72	<u>84.72</u>
	PDPP(Wang et al. 2023b)	A	<u>26.67</u>	<u>46.89</u>	59.45	14.41	44.10	51.39
	SkipPlan(Li et al. 2023c)	A	-	-	-	<u>23.65</u>	<u>47.12</u>	78.44
	LangFirst(Liu et al. 2023b)	A	-	-	-	16.04	43.19	77.07
	SCHEMA(Niu et al. 2024)	A	23.26	39.93	76.75	22.02	45.33	83.47
	KEPP(Nagasinghe et al. 2024)	A	22.71	41.59	91.49	15.63	39.53	-
	PlanLLM(Ours)	A	27.08	46.96	<u>77.89</u>	25.31	48.79	86.28

Table 2: Evaluation results on NIV and COIN with prediction horizon $t \in \{3, 4\}$.

Method	Train Data	Test Data	t=3			t=4		
			SR \uparrow	mAcc \uparrow	mIoU \uparrow	SR \uparrow	mAcc \uparrow	mIoU \uparrow
SCHEMA	COIN	CrossTask	5.37	17.21	54.52	4.45	11.32	50.26
PlanLLM(Ours)	COIN	CrossTask	12.13	19.79	58.32	8.32	14.13	53.44
SCHEMA	CrossTask	COIN	3.23	10.56	49.21	1.45	9.63	46.51
PlanLLM(Ours)	CrossTask	COIN	9.97	12.43	51.17	7.18	12.21	52.21

Table 3: Performance comparisons on cross-dataset with prediction horizon $t \in \{3, 4\}$.

Method	t=3			t=4		
	SR \uparrow	mAcc \uparrow	mIoU \uparrow	SR \uparrow	mAcc \uparrow	mIoU \uparrow
w/o MIM/LLM	38.03	62.71	79.10	21.86	56.28	74.81
w/o LLM	38.99	63.23	79.76	22.42	57.25	75.17
w/o MIM	39.03	64.12	80.25	23.24	57.94	76.11
Full	39.74	65.78	81.50	27.54	59.01	77.58

Table 4: Effectiveness of proposed components

Method	t=3			t=4		
	SR \uparrow	mAcc \uparrow	mIoU \uparrow	SR \uparrow	mAcc \uparrow	mIoU \uparrow
Frozen LLM	37.68	61.26	77.48	22.92	56.19	74.34
One-Stage	38.83	62.17	77.82	23.55	57.67	75.21
Progressive	39.74	65.78	81.50	27.54	59.01	77.58

Table 5: Effectiveness of progressive multi-modal training

progressive training. Our method’s two-stage approach allows the frozen LLM in the first stage to better align embeddings with the input space, and in the second stage, training the LLM optimizes it for downstream tasks. This confirms the effectiveness of progressive training in planning.

Different planning generation strategy In Table 6, we compare the effectiveness of different planning generation strategies. We create two variants: one using only the LLM (without knowledge fusion, step decoder, and step refiner) to generate action sequences (“LLM only”, 1st row), and another using only step prediction without the LLM and

Method	t=3			t=4		
	SR \uparrow	mAcc \uparrow	mIoU \uparrow	SR \uparrow	mAcc \uparrow	mIoU \uparrow
LLM only	37.76	63.25	79.81	24.76	58.01	76.13
Step only	38.21	63.52	79.13	24.03	58.17	75.98
Fusion	39.74	65.78	81.50	27.54	59.01	77.58

Table 6: Different planning generation strategy

knowledge fusion module (“Step only”, 2nd row). (1) “Step only” outperforms “LLM only,” likely because LLM’s free-form outputs may introduce bias in a closed set, such as inconsistent captions with action labels in the test set. (2) The best result is achieved by combining information from both branches using the Knowledge Fusion module (3rd row), demonstrating that the two channels are complement.

Conclusion

In this paper, we present PlanVLM, a multi-modal model designed to enhance LLMs’ planning and visual perception. Our method integrates an LLM with a progressive training strategy to align visual and textual tokens. We also introduce a Mutual Information Maximization module to connect commonsense step descriptions with specific visual states, improving LLMs’ reasoning abilities for generating coherent step sequences. Experimental results on three datasets demonstrate the effectiveness of our approach.

Acknowledgements

This work was supported by the grants from the National Natural Science Foundation of China 62372014 and Beijing Natural Science Foundation 4252040.

References

- Abu Farha, Y.; and Gall, J. 2019. Uncertainty-aware anticipation of activities. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 0–0.
- Ahn, M.; Brohan, A.; Brown, N.; Chebotar, Y.; Cortes, O.; David, B.; Finn, C.; Fu, C.; Gopalakrishnan, K.; Hausman, K.; et al. 2022. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*.
- Alayrac, J.-B.; Bojanowski, P.; Agrawal, N.; Sivic, J.; Laptev, I.; and Lacoste-Julien, S. 2016. Unsupervised learning from narrated instruction videos. In *CVPR*, 4575–4583.
- Bi, J.; Luo, J.; and Xu, C. 2021. Procedure planning in instructional videos via contextual modeling and model-based policy learning. In *ICCV*, 15611–15620.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chang, C.-Y.; Huang, D.-A.; Xu, D.; Adeli, E.; Fei-Fei, L.; and Niebles, J. C. 2020. Procedure planning in instructional videos. In *ECCV*, 334–350.
- Chiang, W.-L.; Li, Z.; Lin, Z.; Sheng, Y.; Wu, Z.; Zhang, H.; Zheng, L.; Zhuang, S.; Zhuang, Y.; Gonzalez, J. E.; Stoica, I.; and Xing, E. P. 2023. Vicuna: An Open-Source Chatbot Impressing GPT-4 with 90%* ChatGPT Quality.
- Chowdhery, A.; Narang, S.; Devlin, J.; Bosma, M.; Mishra, G.; Roberts, A.; Barham, P.; Chung, H. W.; Sutton, C.; Gehrmann, S.; et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240): 1–113.
- Ehsani, K.; Bagherinezhad, H.; Redmon, J.; Mottaghi, R.; and Farhadi, A. 2018. Who let the dogs out? modeling dog behavior from visual data. In *CVPR*, 4051–4060.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- Jansen, P. A. 2020. Visually-grounded planning without vision: Language models infer detailed plans from high-level instructions. *arXiv preprint arXiv:2009.14259*.
- Lei, T.; Yin, S.; Peng, Y.; and Liu, Y. 2024. Exploring Conditional Multi-Modal Prompts for Zero-shot HOI Detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*.
- Li, K.; Wang, Y.; He, Y.; Li, Y.; Wang, Y.; Liu, Y.; Wang, Z.; Xu, J.; Chen, G.; Luo, P.; et al. 2023b. Mvbench: A comprehensive multi-modal video understanding benchmark. *arXiv preprint arXiv:2311.17005*.
- Li, Z.; Geng, W.; Li, M.; Chen, L.; Tang, Y.; Lu, J.; and Zhou, J. 2023c. Skip-Plan: Procedure planning in instructional videos via condensed action space learning. In *ICCV*, 10297–10306.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023a. Visual Instruction Tuning. In *NeurIPS*.
- Liu, J.; Li, S.; Wang, Z.; Li, M.; and Ji, H. 2023b. A language-first approach for procedure planning. In *Findings of the Association for Computational Linguistics: ACL 2023*, 1941–1954.
- Miech, A.; Alayrac, J.-B.; Smaira, L.; Laptev, I.; Sivic, J.; and Zisserman, A. 2020. End-to-end learning of visual representations from uncurated instructional videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9879–9889.
- Mishra, B. D.; Huang, L.; Tandon, N.; Yih, W.-t.; and Clark, P. 2018. Tracking state changes in procedural text: a challenge dataset and models for process paragraph comprehension. *arXiv preprint arXiv:1805.06975*.
- Mysore, S.; Jensen, Z.; Kim, E.; Huang, K.; Chang, H.-S.; Strubell, E.; Flanigan, J.; McCallum, A.; and Olivetti, E. 2019. The materials science procedural text corpus: Annotating materials synthesis procedures with shallow semantic structures. *arXiv preprint arXiv:1905.06939*.
- Nagasinghe, K. R. Y.; Zhou, H.; Gunawardhana, M.; Min, M. R.; Harari, D.; and Khan, M. H. 2024. Why Not Use Your Textbook? Knowledge-Enhanced Procedure Planning of Instructional Videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18816–18826.
- Niu, Y.; Guo, W.; Chen, L.; Lin, X.; and Chang, S.-F. 2024. SCHEMA: State CHangEs MATter for Procedure Planning in Instructional Videos. In *ICLR*. Vienna, Austria.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Song, C. H.; Wu, J.; Washington, C.; Sadler, B. M.; Chao, W.-L.; and Su, Y. 2023. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In *ICCV*, 2998–3009.
- Srinivas, A.; Jabri, A.; Abbeel, P.; Levine, S.; and Finn, C. 2018. Universal planning networks: Learning generalizable representations for visuomotor control. In *International Conference on Machine Learning*, 4732–4741. PMLR.
- Sun, J.; Huang, D.-A.; Lu, B.; Liu, Y.-H.; Zhou, B.; and Garg, A. 2022. PlaTe: Visually-grounded planning with transformers in procedural tasks. *IEEE Robotics and Automation Letters*, 7(2): 4924–4930.
- Tang, Y.; Ding, D.; Rao, Y.; Zheng, Y.; Zhang, D.; Zhao, L.; Lu, J.; and Zhou, J. 2019. Coin: A large-scale dataset

for comprehensive instructional video analysis. In *CVPR*, 1207–1216.

Tellex, S.; Kollar, T.; Dickerson, S.; Walter, M.; Banerjee, A.; Teller, S.; and Roy, N. 2011. Understanding natural language commands for robotic navigation and mobile manipulation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1507–1514.

Ting Lei, S. Y.; and Liu, Y. 2024. Exploring the Potential of Large Foundation Models for Open-Vocabulary HOI Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16657–16667.

Wang, A.-L.; Lin, K.-Y.; Du, J.-R.; Meng, J.; and Zheng, W.-S. 2023a. Event-Guided Procedure Planning from Instructional Videos with Text Supervision. In *ICCV*, 13565–13575.

Wang, H.; Wu, Y.; Guo, S.; and Wang, L. 2023b. Pdpp: Projected diffusion for procedure planning in instructional videos. In *CVPR*, 14836–14845.

Wang, Z.; Cai, S.; Chen, G.; Liu, A.; Ma, X.; and Liang, Y. 2023c. Describe, explain, plan and select: Interactive planning with large language models enables open-world multi-task agents. *arXiv preprint arXiv:2302.01560*.

Xie, Y.; Yu, C.; Zhu, T.; Bai, J.; Gong, Z.; and Soh, H. 2023. Translating natural language to planning goals with large-language models. *arXiv preprint arXiv:2302.05128*.

Yang, D.; and Liu, Y. 2024. Active Object Detection with Knowledge Aggregation and Distillation from Large Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16624–16633.

Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K.; and Cao, Y. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.

Ye, Q.; Yu, Z.; Shao, R.; Xie, X.; Torr, P.; and Cao, X. 2024. CAT: Enhancing Multimodal Large Language Model to Answer Questions in Dynamic Audio-Visual Scenarios. *arXiv preprint arXiv:2403.04640*.

Zhao, H.; Hadji, I.; Dvornik, N.; Derpanis, K. G.; Wildes, R. P.; and Jepson, A. D. 2022. P3iv: Probabilistic procedure planning from instructional videos with weak supervision. In *CVPR*, 2938–2948.

Zheng, M.; Cai, X.; Chen, Q.; Peng, Y.; and Liu, Y. 2024. Training Free Video Temporal Grounding using Large-scale Pre-trained Models. In *Proceedings of the European Conference on Computer Vision (ECCV)*.

Zheng, M.; Gong, S.; Jin, H.; Peng, Y.; and Liu, Y. 2023. Generating Structured Pseudo Labels for Noise-resistant Zero-shot Video Sentence Localization. In *Annual Meeting of the Association for Computational Linguistics*.

Zhukov, D.; Alayrac, J.-B.; Cinbis, R. G.; Fouhey, D.; Laptev, I.; and Sivic, J. 2019. Cross-task weakly supervised learning from instructional videos. In *CVPR*, 3537–3545.