

# EgoPlan-Bench: Benchmarking Egocentric Embodied Planning with Multimodal Large Language Models

Yi Chen<sup>1,2</sup>, Yuying Ge<sup>1\*</sup>, Yixiao Ge<sup>1,3</sup>, Mingyu Ding<sup>4</sup>, Bohao Li<sup>1</sup>, Rui Wang, Ruifeng Xu<sup>5\*</sup>, Ying Shan<sup>1,3</sup>, and Xihui Liu<sup>2\*</sup>

<sup>1</sup> Tencent AI Lab

<sup>2</sup> The University of Hong Kong

<sup>3</sup> ARC Lab, Tencent PCG

<sup>4</sup> University of California, Berkeley

<sup>5</sup> Peng Cheng Laboratory

yichennlp@gmail.com

[https://chenyi99.github.io/ego\\_plan/](https://chenyi99.github.io/ego_plan/)

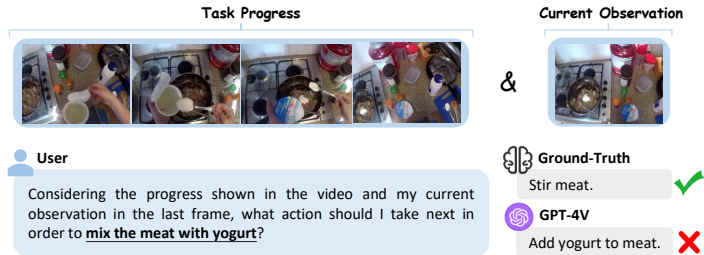
**Abstract.** Multimodal Large Language Models, combining the remarkable reasoning and generalization capabilities of Large Language Models (LLMs) with the ability to comprehend visual inputs, have opened up new avenues for embodied task planning. Given diverse environmental inputs, including real-time task progress, visual observations, and open-form language instructions, a proficient task planner is expected to predict feasible actions, which is a feat inherently achievable by Multimodal Large Language Models (MLLMs). In this paper, we aim to quantitatively investigate the potential of MLLMs as embodied task planners in real-world scenarios by introducing a benchmark with human annotations named EgoPlan-Bench. Our benchmark is distinguished by realistic tasks derived from real-world videos, a diverse set of actions involving interactions with hundreds of different objects, and complex visual observations from varied scenes. We evaluate a wide range of MLLMs, revealing that these models have not yet evolved into embodied planning generalists (even GPT-4V). We further construct an instruction-tuning dataset EgoPlan-IT from videos with human-object interactions, to facilitate the learning of high-level task planning in intricate real-world situations. The experiment results demonstrate that the model tuned on EgoPlan-IT not only significantly improves performance on our benchmark, but can also be applied as a task planner for guiding embodied agents in simulations.

**Keywords:** Multimodal Large Language Model Benchmark · Egocentric Embodied Planning

## 1 Introduction

In recent years, Large Language Models (LLMs) [4, 5, 39, 46] demonstrate strong performance in general-purpose natural language tasks. By seamlessly integrat-

\* Corresponding authors.



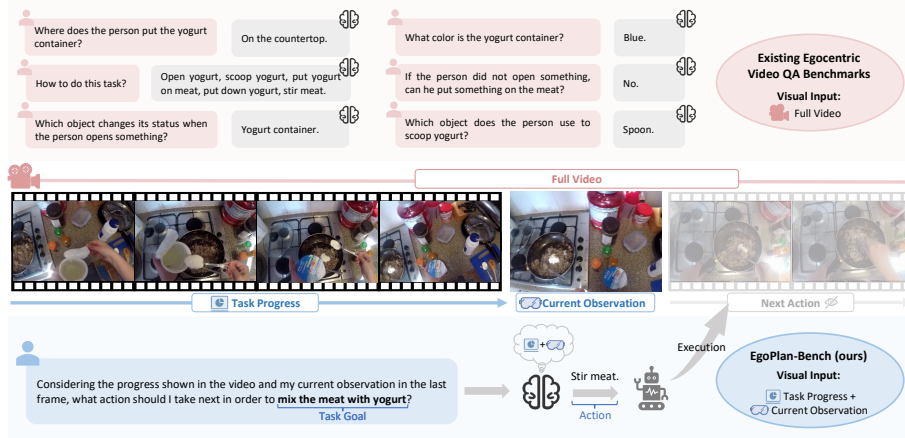
**Fig. 1:** EgoPlan-Bench presents significant challenges for embodied planning given real-time task progress and real-world visual observations (even GPT-4V makes mistakes).

ing the impressive reasoning and generalization ability of LLMs with multi-modal comprehension capabilities, Multimodal Large Language Models (MLLMs) inherently have the potential to serve as embodied task planners, which are expected to predict feasible actions given a specific task goal, real-time task progress and visual observations. Embodied task planning is crucial in enabling autonomous agents to plan for complex tasks in everyday environments, paving the way for versatile AI assistants.

However, embodied planning in real-world scenarios presents significant challenges, as it requires a comprehensive understanding of the dynamic and complicated visual environment and the identification of the key information relevant to the tasks. The model needs to comprehend various types of fine-grained visual information such as the object states and spatial relationships between objects. Moreover, for long-horizon tasks, models are expected to process a series of past visual observations to assess the task progress and make informed plans that better facilitate task completion.

As exemplified in Figure 1, in order to decide the next appropriate action for the task goal “*mixing the meat with yogurt*”, the model must be able to understand and track several state changes of the yogurt container and the spoon, and clearly identify the white paste of yogurt on the meat that requires further stirring in the current observation.

So we ask **how far are the current MLLMs from becoming an embodied planning generalist in real-world scenarios?** Despite the emergence of various benchmarks [11, 23, 33, 51] specially designed to evaluate MLLMs across a wide range of dimensions, there is a notable absence of a benchmark designed for objectively evaluating the capabilities of MLLMs in embodied planning, especially for practical real-world tasks that involve complex visual inputs. Large-scale egocentric videos [8, 16] capturing daily human activities from the first-person perspective serve as a valuable resource for research in this area. However, as illustrated in Figure 2, existing egocentric video question-answering (QA) benchmarks [10, 20] mainly focus on analyzing the **comprehension** ability of models based on a complete video demonstration. In contrast, an embodied task planning benchmark emphasizes on evaluating the **planning** capability of



**Fig. 2:** Our EgoPlan-Bench evaluates Egocentric Embodied **Planning**, where a model predicts the next feasible action by taking a video showing task progress, current visual observation and language instruction indicating the task goal as inputs. In contrast, existing egocentric video QA benchmarks mainly evaluate **Comprehension**, where a model answers questions based on the spatial and temporal understanding of the entire video. The key difference is that the setting of EgoPlan-Bench better aligns with real-world embodied AI applications, as the model’s output can directly serve as plans for agents to execute tasks in real environments.

a model to predict feasible actions for a specific task goal, given the real-time task progress shown in partial video input and current visual observation.

To mitigate the aforementioned gap, this paper establishes a rigorous Egocentric Embodied Planning Benchmark called EgoPlan-Bench, for real-world embodied planning, which is sourced from egocentric videos [8, 16] that showcase everyday human activities. All the evaluation data has undergone strict manual screening, as illustrated in Figure 3. Drawing upon the attributes of the egocentric video sources, our benchmark exhibits three main characteristics. **1) Realism of Tasks:** The tasks are extrapolated from authentic real-world videos, offering a closer reflection of daily human needs and showcasing greater variety than artificially constructed tasks. **2) Diversity of Actions:** The benchmark involves a diverse set of actions, requiring interaction with hundreds of different objects and extending beyond basic manipulation skills such as picking and placing items. **3) Intricacy of Visual Observations:** The visual observations come across various real-world scenes, where objects vary in appearance, state, and placement. Besides, the visual inputs can span extensive periods, making it difficult for models to monitor task progression and detect critical changes in object states.

We evaluate a wide range of MLLMs on our benchmark. The results indicate that our benchmark poses significant challenges for existing MLLMs, and there is still a long way to go before these models evolve into generalist embodied task

planners. In light of the empirical findings, we construct an instruction-tuning dataset EgoPlan-IT in an automatic way, which is specialized for enhancing Egocentric Embodied Planning. We further fine-tune a baseline MLLM with EgoPlan-IT, and the experiment results show that the fine-tuned model significantly improves performance in both in-domain and out-of-domain evaluations. It demonstrates the feasibility for MLLMs to learn generalizable commonsense for task planning from real-world videos. The fine-tuned model can also be applied to guide a robot in completing long-horizon tasks in a simulated environment, which showcases the potential of MLLMs as the “brain” of an embodiment in accomplishing complex tasks.

In summary, our main contributions include:

- We introduce an evaluation benchmark called EgoPlan-Bench for real-world embodied planning based on egocentric videos, which features realistic tasks, diverse actions, and intricate visual observations.
- We evaluate a wide range of MLLMs and find that our benchmark poses significant challenges for them. We further provide a detailed analysis of the possible reasons and future directions in this area.
- We construct an instruction-tuning dataset EgoPlan-IT specifically for Egocentric Embodied Planning in real-world scenarios. The model tuned on this dataset demonstrates robust performance gains on the proposed benchmark and the potential to act as a task planner for guiding embodied agents.

## 2 Related Work

**Large Foundation Models for Embodied Planning.** The growing number of studies incorporating large foundation models into embodied planning processes can be broadly categorized into two main trends. The first trend focuses on utilizing Large Language Models (LLMs) [4–6, 38, 40, 46] to facilitate reasoning and planning based on textual information. Research in this direction often depends on external modules, which convert multi-modal environment information into text [27, 45, 48] or ground LLM predictions to environment affordance [3, 18, 19, 28]. However, these approaches are constrained by the accuracy of the external modules, which may lead to error propagation. The second trend focuses on Multimodal Large Language Models (MLLMs) for end-to-end embodied planning, which extends LLMs with multi-modal perception abilities. PaLM-E [9] is the first work to demonstrate the potential of this approach. Subsequent works, such as EmbodiedGPT [37] and Otter [21, 22], construct instruction data based on egocentric videos [16] for model pretraining and show qualitative evidence of emergent embodied planning abilities. Nonetheless, these studies have not conducted a comprehensive and objective evaluation of embodied planning in real-world scenarios.

**Egocentric Video QA Benchmarks.** Egocentric videos provide a distinctive perspective for active engagement with the real world. Over the years, numerous egocentric video datasets have been developed [26, 42, 44]. Large-scale daily lifelog datasets such as Epic-Kitchens [8] and Ego4D [16] offer fine-grained information

on human-object interactions, making them valuable resources for embodied AI research. Existing egocentric video QA benchmarks primarily assess models’ understanding of activities and objects in a given egocentric video demonstration. These questions range from visual concept recognition [10] to spatial, temporal, and causal relationship reasoning [20]. However, none of these benchmarks address the specific needs of embodied planning, where a more practical question and answer would involve a task goal instruction and a prediction of the next executable action that has not yet occurred but is helpful for achieving the given task goal.

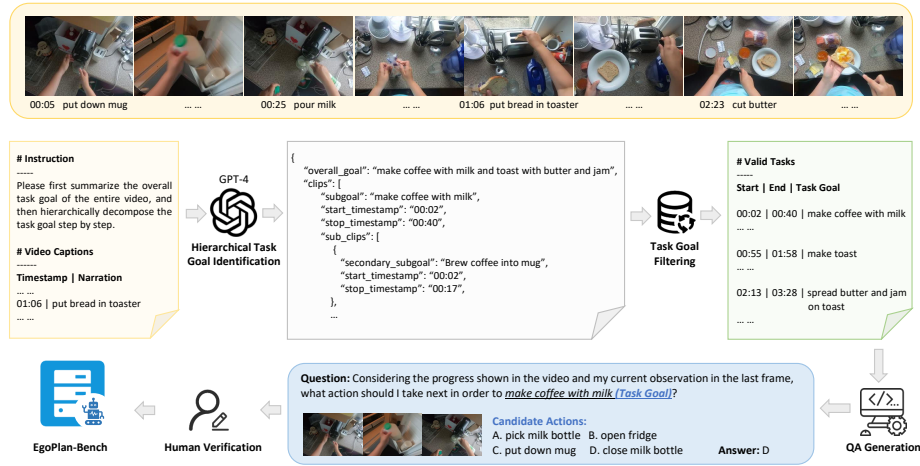
**Multimodal Large Language Models.** Recent advancements in Large Language Models (LLMs) [5, 6, 39, 46] have led to significant achievements in language understanding and generation. As a result, there is a growing interest in developing Multimodal Large Language Models (MLLMs) that combine LLMs’ impressive language capabilities with multi-modal perception abilities [2, 12, 15, 22, 36, 41, 41, 47, 49, 52, 54]. Typical research involves integrating representations from pretrained visual encoders into LLMs’ input embedding space through pre-training on image-text interleaved datasets [7, 24, 31, 32, 55]. Recent studies extend this approach to incorporate video inputs, enhancing LLMs’ video comprehension capabilities. Some works even advance MLLMs to include visual generation abilities [13, 14]. In this paper, we aim to provide a comprehensive evaluation of these models, particularly focusing on their performance in Egocentric Embodied Planning.

### 3 EgoPlan-Bench

As illustrated in Figure 2, given a language instruction demonstrating the task goal, Egocentric Embodied Planning aims to decide the next appropriate action based on the visual input. The visual input includes a video showing a sequence of previous video frames indicating the task progress, with the last frame indicating the current egocentric observation. The proposed EgoPlan-Bench contains 3.4K high-quality human-verified multiple-choice questions derived from existing egocentric video sources [8, 16]. In Sec. 3.1, we first introduce the egocentric videos we used to build the benchmark. We illustrate how to automatically extract task goals from these videos in Sec. 3.2 and construct questions based on the extracted task goals in Sec. 3.3. We then report the detailed data statistics in Sec. 3.4 and finally describe our evaluation strategy in Sec. 3.5. The entire benchmark construction pipeline is demonstrated in Figure 3.

#### 3.1 Egocentric Video Sources

Unscripted, egocentric videos encompass a wide range of daily human tasks and diverse actions, making them an ideal and challenging medium for evaluating embodied planning. We select two large-scale and well-annotated datasets, Epic-Kitchens [8] and Ego4D [16], as our video sources. In Epic-Kitchens, each video is annotated with dense narrations and start/end timestamps of actions.



**Fig. 3:** Overview of the construction pipeline for EgoPlan-Bench based on existing untrimmed egocentric videos with detailed action narrations. (1) We first leverage GPT-4 to identify task goals through hierarchical reasoning. (2) We then filter task goals based on the requisite number of actions. (3) The questions are designed in the form of multiple-choice, where the questions are automatically generated based on task goals, and the options are derived from different actions under the same task goal. (4) We employ human annotators to verify each question to ensure the benchmark quality.

The narrations are expressed as brief verb-object phrases. Ego4D offers similar annotations, except that it only provides a single timestamp for each action and the narrations are expressed as brief sentences like “*C does something*”. We convert Ego4D narrations to verb-object phrases as Epic-Kitchens with GPT-4 [38] and derive the start and end timestamps for each action segment following the EgoVLP framework [29].

### 3.2 Task Goal Extraction

**Hierarchical Task Goal Identification.** Egocentric videos, which capture tasks ranging from simple (involving fewer than three actions) to complex (entailing a long sequence of actions), present challenges in identifying task goals due to their varying lengths. To address this, we introduce a hierarchical reasoning strategy to identify task goals. As illustrated in Figure 3, given the dense narrations of actions, we utilize GPT-4 to decompose video content into three levels: overall goal → subgoal → secondary subgoal. For instance, the goal to “*make coffee with milk and toast with butter and jam*” is divided into the subgoals of “*making coffee with milk*” and “*making toast with butter and jam*,” with the former further broken down into secondary subgoals like “*brew coffee into mug*” and “*merge coffee and milk*.” This hierarchical approach allows GPT-4 to efficiently summarize tasks of varying complexity by organizing them into a

structured framework, similar to chain-of-thought reasoning. Our findings suggest that prompting GPT-4 to generate all task goals at once may lead to difficulties due to the extensive information presented in longer videos. The Detailed prompt can be found in the supplementary material.

**Task Goal Filtering.** Given the variance in video lengths, the task hierarchy across different videos might not align consistently. Specifically, the complexity of the overall goal in a shorter video could be equivalent only to a secondary subgoal in a more extended video. To address this variance, we aggregate the overall goals, subgoals, and secondary subgoals extracted from various videos under the unified term “task goals”. We then refine this set by filtering based on the requisite number of actions, retaining only those task goals that involve between 4 to 20 actions to maintain a moderate level of task complexity.


### 3.3 Multiple-Choice Question Answering

**Automatic QA Generation.** Following the mainstream MLLM benchmarks [23, 33], we evaluate Egocentric Embodied Planning with multiple-choice questions, which facilitates the convenient computation of accuracy as an objective metric. With the filtered task goals, we use templates to automatically generate questions. As shown in the blue box in Figure 3, the underlined words indicate the part to fill in with different task goals. With the same task goal, we can derive several question-answering (QA) pairs that differ in task progress and current observations. For example, given a task goal involving  $N$  actions, we can obtain  $N$  corresponding QA pairs with each action as the ground-truth answer. We align the visual input with the ground-truth answer and randomly sample three distinct actions under the same task goal as negative choices.

**Visual Input Alignment.** The visual input attached to each QA pair contains the historical observations since the whole task begins and right before the ground-truth action occurs. The last frame of the visual input indicates the current observation for planning the next action. To prevent the models from cheating with the clues from hand-object interaction (e.g., a hand touching a tap implies turning it on or off), we set the end timestamp of the visual input as 0.5 seconds preceding the actual action segment.

**Human Verification.** To ensure the reliability and objectiveness of our EgoPlan-Bench, we rigorously refined the automatically generated QA pairs with the help of 10 professional annotators, all holding at least a bachelor’s degree. Initially, annotators underwent a quality assurance phase, pre-annotating 50 samples following comprehensive guidelines. Their work was meticulously reviewed by two senior annotators. During the main annotation phase, pairs of annotators independently selected the most appropriate answers for each question, based on the task goal and visual cues, or flagged questions with unclear options for removal. We retained only those QA pairs where annotators’ choices were unanimous and aligned with the ground-truth answer. This process resulted in 3.4K high-quality multiple-choice QA pairs, closely correlated with visual observations.





**Fig. 4:** Top 20 verbs with top 8 related objects in EgoPlan-Bench candidate actions.



**Fig. 5:** Wordcloud of task goals involved in EgoPlan-Bench questions.

**Table 1:** Statistics of the evaluation data of EgoPlan-Bench.

#questions	#task goals	#actions (verb+object)	#verbs	#objects	#scenes
3,355	2,406	2,979	202	496	279

### 3.4 Data Statistics

The proposed EgoPlan-Bench, detailed in Table 1, includes 3,355 multiple-choice questions covering 2,406 task goals. Figure 5 shows the distribution of task goals. Overall, there are 2,979 actions (verb-object phrases) from the questions’ candidate options, involving 202 distinct root verbs and 496 unique objects. The most common actions are illustrated in Figure 4, demonstrating the benchmark’s variety of actions. Additionally, the visual observations of these questions originate from egocentric videos taken across 279 different scenes. This diversity in task goals, candidate actions, and visual scenes underscores the challenges presented by our benchmark.

### 3.5 Evaluation Strategy

Inspired by [3, 4, 7, 23, 30], we use the closed-set answer ranking strategy based on language completion distribution, rather than directly matching a model’s free-form generation results with candidate options. From the output distribution of MLLMs, we can derive the likelihood  $p(a|v, l)$  that an MLLM generates the content of each candidate action  $a$ , given the visual inputs  $v$  and the user instruction  $l$  as the multimodal context.

This approach avoids the influence of option sequence in the prompt on the model’s performance. More importantly, it is closer to the practical setting. An



**Table 2:** Performance of MLLMs on EgoPlan-Bench.

Model	LLM	Acc%
BLIP-2 [24]	Flan-T5-XL	25.69
InstructBLIP [7]	Flan-T5-XL	26.32
InstructBLIP Vicuna [7]	Vicuna-7B	25.37
LLaVA [32]	LLaMA-7B	25.52
MiniGPT-4 [55]	Flan-T5-XL	26.50
VPGTrans [52]	LLaMA-7B	26.61
MultiModal-GPT [15]	Vicuna-7B	27.03
Otter [22]	LLaMA-7B	26.47
OpenFlamingo [36]	LLaMA-7B	26.85
LLaMA-Adapter V2 [12]	LLaMA-7B	26.49
GVT [47]	Vicuna-7B	26.83
mPLUG-Owl [49]	LLaMA-7B	25.93
Kosmos-2 [41]	Decoder only 1.3B	26.97
Qwen-VL-Chat [2]	Qwen-7B	26.32
LLaVA-1.5 [31]	Vicuna-7B	26.80
VideoChat [25]	Vicuna-7B	26.32
Video-ChatGPT [35]	LLaMA-7B	26.35
Valley [34]	LLaMA-13B	26.17
Video-LLaMA [53]	LLaMA2-Chat-7B	28.58
SEED-LLaMA [14]	LLaMA2-Chat-13B	29.93
InternLM-Xcomposer-VL [54]	InternLM-7B	34.40
GPT-4V [1]	-	<b>37.98</b>

MLLM that achieves better performance on our benchmark indicates that it can output precise next-action plans with higher probability.

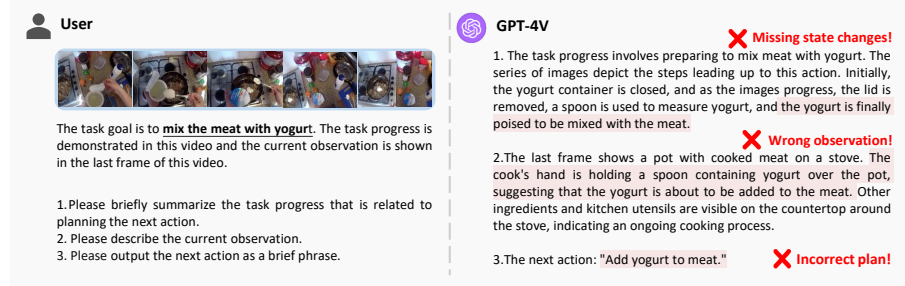
## 4 Evaluation Results

### 4.1 Models

We select 22 prominent MLLMs for a comprehensive evaluation, including BLIP-2 [24], InstructBLIP [7], InstructBLIP Vicuna [7], LLaVA [32], MiniGPT-4 [55], VPGTrans [52], MultiModal-GPT [15], Otter [22], OpenFlamingo [36], LLaMA-Adapter V2 [12], GVT [47], mPLUG-Owl [49], Kosmos-2 [41], Qwen-VL-Chat [2], LLaVA1.5 [31], VideoChat [25], Video-ChatGPT [35], Valley [34], Video-LLaMA [53], SEED-LLaMA [14], InternLM-Xcomposer-VL [54], GPT-4V [1].

### 4.2 Performance Analysis

The performance of various MLLMs on our EgoPlan-Bench is shown in Table 2. Notably, the majority of these models exhibit only marginally better results than random guessing, with the highest-performing model attaining a mere 37.98% accuracy, indicating that current MLLMs lack the ability as real-world embodied



**Fig. 6:** A case study with GPT-4V, where the model fails to detect the state changes of target object in task progress, and struggles to recognize the target object in current observation, leading to the wrong prediction of plan.

task planners. In the supplementary material, we present common failure cases of these models. Our analysis reveals that most MLLMs tend to select options that are semantically closer to the given task goal. However, they often fail to pay adequate attention to the vital aspect of visual observations. Common mistakes such as overlooking completed actions or misinterpreting key object states can significantly affect their overall performance in Egocentric Embodied Planning.

### 4.3 Case Study with GPT-4V

Since GPT-4V achieves the best performance among all the evaluated MLLMs, we use it for a qualitative case study to delve deeper into why current MLLMs may underperform on our benchmark. As depicted in Figure 6, we require GPT-4V to explicitly summarize the task progress and describe the current observation before planning the next action.

Our observations indicate that while GPT-4V successfully identifies obvious actions, like “*removing the lid of yogurt container*” or “*measuring yogurt with a spoon*,” it might overlook more subtle state changes that are critical to the task’s objective. For instance, GPT-4V fails to recognize that yogurt has been added to the meat, as indicated by the white paste visible on the meat post-measurement. Furthermore, even when the yogurt has been placed on the countertop in the final frame, GPT-4V erroneously maintains that the current observation shows “*a hand holding a spoon containing yogurt*.” This misinterpretation of the visual input leads to incorrect next plans, such as concluding that the next action is to “*Add yogurt to meat*,” when this step has already been completed. This case study highlights a key area for improvement in MLLMs’ processing of visual cues in the context of real-time task execution.

Future directions may include: 1) developing advanced algorithms or incorporating sophisticated training datasets that highlight subtle visual changes, 2) enhancing contextual understanding of MLLMs to make more informed plans based on comprehensive visual input analysis relevant to specific task goals, and 3) integrating real-time feedbacks for adaptability in dynamic environments.

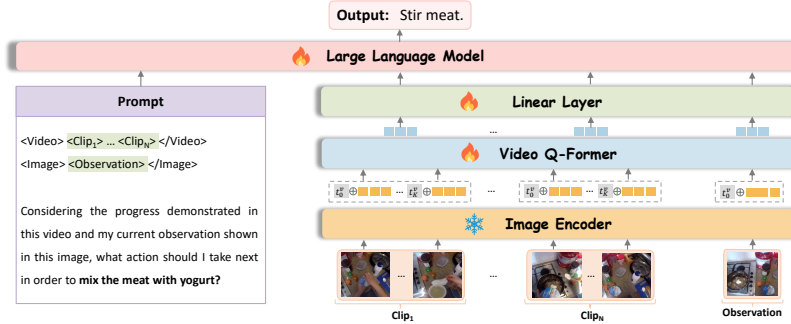


Fig. 7: The diagram of enhancing the baseline model Video-LLaMA with EgoPlan-IT.

These improvements would enhance MLLM performance in Egocentric Embodied Planning and expand their applicability in real-world scenarios.

## 5 Enhancing Egocentric Embodied Planning Capability by Instruction Tuning

Given the suboptimal performance of the evaluated MLLMs on EgoPlan-Bench, we investigate enhancing the Egocentric Embodied Planning capabilities of MLLMs through instruction-tuning. Specifically, we construct an instruction-tuning dataset, EgoPlan-IT, to enable MLLMs to effectively learn embodied planning experience from real-world videos. In this section, we first describe the data preparation process in Sec. 5.1. Then we introduce the framework of a baseline MLLM in Sec. 5.2 and the training objectives in Sec. 5.3. In Sec. 5.4, we demonstrate the effectiveness of our data and instruction-tuning strategy through experiments.

### 5.1 Data Preparation

**Construction of EgoPlan-IT.** Following the automatic data construction pipeline in Sec. 3, we additionally tailor an instruction dataset EgoPlan-IT, which contains 50K instruction-following pairs as shown in the bottom half of Figure 2, for Egocentric Embodied Planning. To assess the robustness of our instruction tuning approach, we exclusively utilize video data from Epic-Kitchens [8] for the creation of this training set. This enables us to assess the performance of the enhanced model on both the in-domain (Epic-Kitchens) and out-of-domain (Ego4D) subsets of our evaluation data, which will be discussed in detail in Section 5.4. For increasing the efficiency of the dataset construction, we do not include human annotations in this process.

**Auxiliary Data for Action Recognition.** To reduce the learning difficulty of the model, we also construct an auxiliary dataset for action recognition based on the visual inputs of the 50K training samples from EgoPlan-IT. For the example shown in the bottom half of Figure 2, the instruction of the auxiliary sample

would be a simple query for action recognition such as “*Can you enumerate the actions in the video, describing each with a short verb-noun combination?*”, and the response is the concatenation of corresponding action narrations such as “*Open yogurt, scoop yogurt, put yogurt on meat, put down yogurt.*”

**Increasing Instruction Diversity.** In order to increase the diversity of instructions, we also include an additional collection of 164K instruction data, which integrates the 150K image-instruction dataset from LLaVA [32], the 3K image-detail-description dataset from MiniGPT-4 [55] and the 11K video-instruction dataset from VideoChat [25].

## 5.2 Model Framework

We use Video-LLaMA [53] as a baseline MLLM for investigation. As illustrated in Figure 7, the vision processor of Video-LLaMA comprises four components, including an Image Encoder to extract features from video frames, a position embedding layer for adding temporal information, a Video Q-former for aggregating frame representations, and a linear layer for producing video embedding vectors that match the dimensions of LLM token embeddings.

To enhance the model’s ability to process visual inputs for Egocentric Embodied Planning, we organize historical observations based on actions and retain the most recent  $N(=4)$  clips. For each action segment, we choose  $K(=8)$  uniformly spaced keyframes, creating an aggregated video representation. To emphasize the current observation, we treat it as a separate single-frame clip, extracting a unique representation for it. The token embeddings for placeholders  $\langle \text{Clip}_i \rangle$  and  $\langle \text{Observation} \rangle$  in the textual prompt are substituted with the embeddings corresponding to the  $i$ -th action segment and the current observation, respectively, within the LLM.

## 5.3 Training Objectives

To fine-tune the model, we mix the 50K specialized EgoPlan-IT data together with the 50K auxiliary data for Action Recognition and the 164K diverse instruction data. The autoregressive loss function is defined as:

$$L_{ft} = - \sum_t \log P(a_t | \mathbf{v}, \mathbf{l}, \mathbf{a}_{<t}; \boldsymbol{\Theta}), \quad (1)$$

where  $\mathbf{v}$  and  $\mathbf{l}$  denote the given visual observations and the language instruction respectively, and  $\mathbf{a}_{<t}$  signifies the sequence of answer tokens preceding the current prediction token  $a_t$ .  $\boldsymbol{\Theta}$  is the trainable parameters including the full parameters of the Video Q-Former and the Linear Layer, and the LoRA [17] parameters of the language model. Inspired by the RRHF mechanism [50], we add a contrastive loss to the fine-tuning loss:

$$L = L_{ft} + L_{ct}, \quad (2)$$

$$L_{ct} = \max(0, s_{neg} - s_{pos}), \quad (3)$$

**Table 3:** Performance of the baseline model (Video-LLaMA) tuned on EgoPlan-IT. In-Dom. and Out-Dom. are short for in-domain and out-of-domain respectively.

Model	In-Dom.	Out-Dom.	All
GPT-4V	38.40	36.90	37.98
Video-LLaMA	27.88	30.44	28.58
Enhanced Video-LLaMA	<b>54.65</b>	<b>44.42</b>	<b>51.83</b>
w/o LoRA	53.58	42.58	50.55
w/o Contrastive Loss	52.14	40.52	48.94
w/o Diverse Instructions	53.95	43.12	50.97
w/o Action Recognition	52.34	44.31	50.13

$$s = \frac{\sum_t \log P(a_t | \mathbf{v}, \mathbf{l}, \mathbf{a}_{<t}; \Theta)}{\|\mathbf{a}\|}. \quad (4)$$


Here, for each instruction-following pair, a random action distinct from the ground-truth answer is selected under the same task goal to serve as a ‘negative’ answer. The model is then constrained to ensure that its score for the ‘positive’ answer ( $s_{pos}$ ) exceeds that for the ‘negative’ answer ( $s_{neg}$ ). This operation aims to ensure a more balanced consideration of both the language instruction and the visual context in Egocentric Embodied Planning.

#### 5.4 Experiments

**Main Results.** The results in Table 3 clearly indicate that the baseline model (Video-LLaMA) tuned on our data outperforms the previous state-of-the-art model, GPT-4V, by a considerable margin. Notably, it shows a remarkable 23.25% increase in accuracy over its vanilla version across the overall evaluation set. Furthermore, the enhanced model demonstrates robust domain transfer capabilities, outperforming the vanilla version by 13.98% on the out-of-domain evaluation subset. This indicates that the experience developed through instruction-tuning on our Ego-IT data is effectively transferable to Egocentric Embodied Planning in new environments.

**Ablations.** Our analysis of the ablated versions of the enhanced model reveals several key insights. Notably, employing the LoRA strategy, integrating contrastive loss, diversifying instructions, and fine-tuning with auxiliary action recognition data consistently enhance performance across domains. Contrastive loss contributes the most to the model’s improvement. Furthermore, fine-tuning the model with in-domain action recognition auxiliary data boosts in-domain performance but does not significantly improve domain transferability.

**Applications in Guiding Embodied Agents.** Embodied planning could facilitate various downstream applications, including personalized virtual assistants, robotics for complex tasks, adaptive gaming experiences, and so on. We demonstrate how to apply the MLLM tuned on EgoPlan-IT as a task planner



**Fig. 8:** The MLLM tuned on EgoPlan-IT can be applied as the task planner of embodied agents for completing long-horizon tasks within a simulated environment.

for guiding embodied agents, through a qualitative analysis in a simulated environment, VirtualHome [43]. The model is deployed in a zero-shot setting. As depicted in Figure 8, we use the model to predict feasible actions for an agent to execute step by step. We employ the beam search decoding algorithm with a beam width of five, choosing the top-one prediction for implementation. As shown in Figure 8, the model’s predictions closely align with the visual progress and real-time observation at each step, successfully guiding the embodied agent to complete a long-horizon task. This application case highlights the promising usage of the MLLM and instruction-tuning dataset specialized for Egocentric Embodied Planning, underscoring the significant research potential of this area. More examples can be seen in the supplementary material.

## 6 Conclusion and Discussion

In this work, we introduce EgoPlan-Bench, a sophisticated benchmark to evaluate Multimodal Large Language Models (MLLMs) in Egocentric Embodied Planning, where a model is expected to plan step-by-step executable actions by considering task progress, current visual observation, and open-form task goal. The evaluation results of various models reveal that current MLLMs have not yet developed into generalized embodied planners. We further construct EgoPlan-IT, an instruction-tuning dataset, to facilitate the learning of high-level task planning from human videos. The model tuned on EgoPlan-IT exhibits a significant performance enhancement on our benchmark. Moreover, it can be applied as an effective task planner for guiding embodied agents to complete long-horizon tasks within a simulated environment. Currently, we have not considered the situation of replanning if the agent fails to complete the given task, as the cases

are rare in existing egocentric video sources. Nevertheless, our benchmark construction approach is general, and in the future, we can expand our benchmark by recording videos in such scenarios to further investigate this area.

## References

1. Gpt-4v(ision) system card (2023), <https://api.semanticscholar.org/CorpusID:263218031> **9**
2. Bai, J., Bai, S., Yang, S., Wang, S., Tan, S., Wang, P., Lin, J., Zhou, C., Zhou, J.: Qwen-vl: A frontier large vision-language model with versatile abilities. arXiv preprint arXiv:2308.12966 (2023) **5, 9**
3. Brohan, A., Chebotar, Y., Finn, C., Hausman, K., Herzog, A., Ho, D., Ibarz, J., Irpan, A., Jang, E., Julian, R., et al.: Do as i can, not as i say: Grounding language in robotic affordances. In: Conference on Robot Learning. pp. 287–318. PMLR (2023) **4, 8**
4. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Nee-lakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* **33**, 1877–1901 (2020) **1, 4, 8**
5. Chiang, W.L., Li, Z., Lin, Z., Sheng, Y., Wu, Z., Zhang, H., Zheng, L., Zhuang, S., Zhuang, Y., Gonzalez, J.E., et al.: Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023) (2023) **1, 4, 5**
6. Chung, H.W., Hou, L., Longpre, S., Zoph, B., Tay, Y., Fedus, W., Li, E., Wang, X., Dehghani, M., Brahma, S., et al.: Scaling instruction-finetuned language models. arXiv preprint arXiv:2210.11416 (2022) **4, 5**
7. Dai, W., Li, J., Li, D., Tiong, A.M.H., Zhao, J., Wang, W., Li, B., Fung, P., Hoi, S.: Instructblip: Towards general-purpose vision-language models with instruction tuning. arXiv preprint arXiv:2305.06500 (2023) **5, 8, 9**
8. Damen, D., Doughty, H., Farinella, G.M., Furnari, A., Ma, J., Kazakos, E., Moltisanti, D., Munro, J., Perrett, T., Price, W., Wray, M.: Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision (IJCV)* **130**, 33–55 (2022), <https://doi.org/10.1007/s11263-021-01531-2> **2, 3, 4, 5, 11**
9. Driess, D., Xia, F., Sajjadi, M.S., Lynch, C., Chowdhery, A., Ichter, B., Wahid, A., Tompson, J., Vuong, Q., Yu, T., et al.: Palm-e: An embodied multimodal language model. arXiv preprint arXiv:2303.03378 (2023) **4**
10. Fan, C.: Egovqa-an egocentric video question answering benchmark dataset. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019) **2, 5**
11. Fu, C., Chen, P., Shen, Y., Qin, Y., Zhang, M., Lin, X., Yang, J., Zheng, X., Li, K., Sun, X., et al.: Mme: A comprehensive evaluation benchmark for multimodal large language models. arXiv preprint arXiv:2306.13394 (2023) **2**
12. Gao, P., Han, J., Zhang, R., Lin, Z., Geng, S., Zhou, A., Zhang, W., Lu, P., He, C., Yue, X., Li, H., Qiao, Y.: Llama-adapter v2: Parameter-efficient visual instruction model. arXiv preprint arXiv:2304.15010 (2023) **5, 9**
13. Ge, Y., Ge, Y., Zeng, Z., Wang, X., Shan, Y.: Planting a seed of vision in large language model. arXiv preprint arXiv:2307.08041 (2023) **5**



14. Ge, Y., Zhao, S., Zeng, Z., Ge, Y., Li, C., Wang, X., Shan, Y.: Making llama see and draw with seed tokenizer. arXiv preprint arXiv:2310.01218 (2023) [5](#), [9](#)
15. Gong, T., Lyu, C., Zhang, S., Wang, Y., Zheng, M., Zhao, Q., Liu, K., Zhang, W., Luo, P., Chen, K.: Multimodal-gpt: A vision and language model for dialogue with humans (2023) [5](#), [9](#)
16. Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari, A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X., et al.: Ego4d: Around the world in 3,000 hours of egocentric video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18995–19012 (2022) [2](#), [3](#), [4](#), [5](#)
17. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. arXiv preprint arXiv:2106.09685 (2021) [12](#)
18. Huang, W., Xia, F., Shah, D., Driess, D., Zeng, A., Lu, Y., Florence, P., Mordatch, I., Levine, S., Hausman, K., et al.: Grounded decoding: Guiding text generation with grounded models for robot control. arXiv preprint arXiv:2303.00855 (2023) [4](#)
19. Huang, W., Xia, F., Xiao, T., Chan, H., Liang, J., Florence, P., Zeng, A., Thompson, J., Mordatch, I., Chebotar, Y., et al.: Inner monologue: Embodied reasoning through planning with language models. In: Conference on Robot Learning. pp. 1769–1782. PMLR (2023) [4](#)
20. Jia, B., Lei, T., Zhu, S.C., Huang, S.: Egotaskqa: Understanding human tasks in egocentric videos. Advances in Neural Information Processing Systems **35**, 3343–3360 (2022) [2](#), [5](#)
21. Li, B., Zhang, Y., Chen, L., Wang, J., Pu, F., Yang, J., Li, C., Liu, Z.: Mimic-it: Multi-modal in-context instruction tuning. arXiv preprint arXiv:2306.05425 (2023) [4](#)
22. Li, B., Zhang, Y., Chen, L., Wang, J., Yang, J., Liu, Z.: Otter: A multi-modal model with in-context instruction tuning. arXiv preprint arXiv:2305.03726 (2023) [4](#), [5](#), [9](#)
23. Li, B., Ge, Y., Ge, Y., Wang, G., Wang, R., Zhang, R., Shan, Y.: Seed-bench-2: Benchmarking multimodal large language models. arXiv preprint arXiv:2311.17092 (2023) [2](#), [7](#), [8](#)
24. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. ICML (2023) [5](#), [9](#)
25. Li, K., He, Y., Wang, Y., Li, Y., Wang, W., Luo, P., Wang, Y., Wang, L., Qiao, Y.: Videochat: Chat-centric video understanding. arXiv preprint arXiv:2305.06355 (2023) [9](#), [12](#)
26. Li, Y., Ye, Z., Rehg, J.M.: Delving into egocentric actions. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 287–295 (2015) [4](#)
27. Liang, J., Huang, W., Xia, F., Xu, P., Hausman, K., Ichter, B., Florence, P., Zeng, A.: Code as policies: Language model programs for embodied control. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 9493–9500. IEEE (2023) [4](#)
28. Lin, K., Agia, C., Migimatsu, T., Pavone, M., Bohg, J.: Text2motion: From natural language instructions to feasible plans. arXiv preprint arXiv:2303.12153 (2023) [4](#)
29. Lin, K.Q., Wang, J., Soldan, M., Wray, M., Yan, R., XU, E.Z., Gao, D., Tu, R.C., Zhao, W., Kong, W., et al.: Egocentric video-language pretraining. Advances in Neural Information Processing Systems **35**, 7575–7586 (2022) [6](#)
30. Lin, S., Hilton, J., Evans, O.: Truthfulqa: Measuring how models mimic human falsehoods. arXiv preprint arXiv:2109.07958 (2021) [8](#)

31. Liu, H., Li, C., Li, Y., Lee, Y.J.: Improved baselines with visual instruction tuning. arXiv preprint arXiv:2310.03744 (2023) 5, 9
32. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. arXiv preprint arXiv:2304.08485 (2023) 5, 9, 12
33. Liu, Y., Duan, H., Zhang, Y., Li, B., Zhang, S., Zhao, W., Yuan, Y., Wang, J., He, C., Liu, Z., et al.: Mmbench: Is your multi-modal model an all-around player? arXiv preprint arXiv:2307.06281 (2023) 2, 7
34. Luo, R., Zhao, Z., Yang, M., Dong, J., Qiu, M., Lu, P., Wang, T., Wei, Z.: Valley: Video assistant with large language model enhanced ability. arXiv preprint arXiv:2306.07207 (2023) 9
35. Maaz, M., Rasheed, H., Khan, S., Khan, F.S.: Video-chatgpt: Towards detailed video understanding via large vision and language models. arXiv preprint arXiv:2306.05424 (2023) 9
36. ml\_foundations: Openflamingo. [https://github.com/mlfoundations/open\\_flamingo](https://github.com/mlfoundations/open_flamingo) (2023) 5, 9
37. Mu, Y., Zhang, Q., Hu, M., Wang, W., Ding, M., Jin, J., Wang, B., Dai, J., Qiao, Y., Luo, P.: Embodiedgpt: Vision-language pre-training via embodied chain of thought. arXiv preprint arXiv:2305.15021 (2023) 4
38. OpenAI: Gpt-4 technical report (2023) 4, 6
39. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. NeurIPS (2022) 1, 5
40. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. Advances in Neural Information Processing Systems **35**, 27730–27744 (2022) 4
41. Peng, Z., Wang, W., Dong, L., Hao, Y., Huang, S., Ma, S., Wei, F.: Kosmos-2: Grounding multimodal large language models to the world. arXiv preprint arXiv:2306.14824 (2023) 5, 9
42. Pirsiavash, H., Ramanan, D.: Detecting activities of daily living in first-person camera views. In: 2012 IEEE conference on computer vision and pattern recognition. pp. 2847–2854. IEEE (2012) 4
43. Puig, X., Ra, K., Boben, M., Li, J., Wang, T., Fidler, S., Torralba, A.: Virtual-home: Simulating household activities via programs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8494–8502 (2018) 14
44. Sigurdsson, G.A., Gupta, A., Schmid, C., Farhadi, A., Alahari, K.: Charades-ego: A large-scale dataset of paired third and first person videos. arXiv preprint arXiv:1804.09626 (2018) 4
45. Song, C.H., Wu, J., Washington, C., Sadler, B.M., Chao, W.L., Su, Y.: Llm-planner: Few-shot grounded planning for embodied agents with large language models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2998–3009 (2023) 4
46. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al.: Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971 (2023) 1, 4, 5
47. Wang, G., Ge, Y., Ding, X., Kankanhalli, M., Shan, Y.: What makes for good visual tokenizers for large language models? arXiv preprint arXiv:2305.12223 (2023) 5, 9
48. Yang, Z., Li, L., Wang, J., Lin, K., Azarnasab, E., Ahmed, F., Liu, Z., Liu, C., Zeng, M., Wang, L.: Mm-react: Prompting chatgpt for multimodal reasoning and action. arXiv preprint arXiv:2303.11381 (2023) 4

49. Ye, Q., Xu, H., Xu, G., Ye, J., Yan, M., Zhou, Y., Wang, J., Hu, A., Shi, P., Shi, Y., et al.: mplug-owl: Modularization empowers large language models with multimodality. arXiv preprint arXiv:2304.14178 (2023) 5, 9
50. Yuan, H., Yuan, Z., Tan, C., Wang, W., Huang, S., Huang, F.: Rrhf: Rank responses to align language models with human feedback. In: Thirty-seventh Conference on Neural Information Processing Systems (2023) 12
51. Yue, X., Ni, Y., Zhang, K., Zheng, T., Liu, R., Zhang, G., Stevens, S., Jiang, D., Ren, W., Sun, Y., et al.: Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. arXiv preprint arXiv:2311.16502 (2023) 2
52. Zhang, A., Fei, H., Yao, Y., Ji, W., Li, L., Liu, Z., Chua, T.S.: Transfer visual prompt generator across llms **abs/23045.01278** (2023) 5, 9
53. Zhang, H., Li, X., Bing, L.: Video-llama: An instruction-tuned audio-visual language model for video understanding. arXiv preprint arXiv:2306.02858 (2023) 9, 12
54. Zhang, P., Wang, X.D.B., Cao, Y., Xu, C., Ouyang, L., Zhao, Z., Ding, S., Zhang, S., Duan, H., Yan, H., et al.: Internlm-xcomposer: A vision-language large model for advanced text-image comprehension and composition. arXiv preprint arXiv:2309.15112 (2023) 5, 9
55. Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M.: Minigpt-4: Enhancing vision-language understanding with advanced large language models. arXiv preprint arXiv:2304.10592 (2023) 5, 9, 12