

# InteLiPlan: An Interactive Lightweight LLM-Based Planner for Domestic Robot Autonomy

Kim Tien Ly<sup>1</sup>, Kai Lu<sup>2</sup>, Ioannis Havoutis<sup>1</sup>

**Abstract**—We introduce an interactive LLM-based framework designed to enhance the autonomy and robustness of domestic robots, targeting embodied intelligence. Our approach reduces reliance on large-scale data and incorporates a robot-agnostic pipeline that embodies an LLM. Our framework, *InteLiPlan*, ensures that the LLM’s decision-making capabilities are effectively aligned with robotic functions, enhancing operational robustness and adaptability, while our human-in-the-loop mechanism allows for real-time human intervention when user instruction is required. We evaluate our method in both simulation and on the real Toyota Human Support Robot (HSR). Our method achieves a 93% success rate in the ‘fetch me’ task completion with failure recovery, highlighting its capability in both failure reasoning and task planning. *InteLiPlan* achieves comparable performance to state-of-the-art large-scale LLM-based robotics planners, while using only real-time onboard computing. More information about *InteLiPlan* can be found at the project website: <https://kimtienly.github.io/InteLiPlan>.

## I. INTRODUCTION

In recent years, the integration of artificial intelligence (AI) into robotics has led to significant advancements in automation and autonomous systems [1]–[5]. A key development in this field is the application of large language models (LLMs), such as GPT [6], [7], LLaMA [8], [9], and Mistral [10], which have proven to be powerful tools for understanding and generating human-like text. These models offer novel approaches for enhancing human-robot interaction and enabling more intuitive decision-making processes in robotics.

Despite their impressive capabilities, applying LLMs in robotics presents unique challenges [11]. These challenges stem primarily from the constraints of robot kinematics and the dynamic nature of the environments in which robots operate. Additionally, integrating these generative models within the existing robotics pipelines is non-trivial, as these are usually tightly coupled systems of perception, planning, and actuation. Having such systems running in real-time robotic deployment also poses a significant challenge due to the computational overhead. Moreover, robotic motion safety and robustness are critical concerns in human-robot environments. Therefore, developing a general-purpose autonomous robotic system for domestic scenarios remains a challenging open problem.

In this work, we propose a lightweight LLM-based framework designed for domestic robots, aiming to address the aforementioned challenges. Our approach reduces reliance on large-scale data and develops a robot-agnostic pipeline.

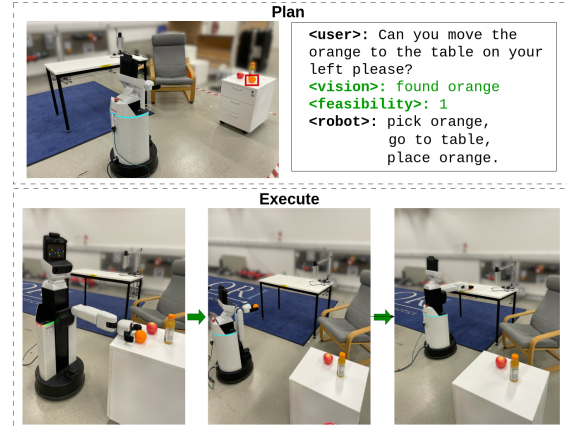


Fig. 1: Step-by-step execution of *InteLiPlan* result on the physical HSR. In our system, the robot receives requests and guidance from humans, e.g., pick-and-place of an orange. Our lightweight onboard planner will generate the robot actions for the task as shown in the figure.

Furthermore, we detail strategies to mitigate kinematic constraints through real-time feasibility checks, ensuring that decision-making processes facilitated by the LLM align effectively with the physical capabilities of robots.

An integral part of our framework is the human-in-the-loop (HITL) mechanism, which allows the system to handle system failures (e.g., partially observable environment, vision failures due to lightning, etc.) through human intervention. This feature is critical for maintaining autonomy, especially in complex environments where decisions by the LLM may fail to execute. The system allows human operators to input corrective actions directly, facilitating on-the-fly adaptation to new or unanticipated situations. This real-time interaction not only enhances the robustness and reliability of robotic operations but also enables fine-tuning of the LLM’s responses in accordance with human guidance.

By tackling these critical issues, our framework not only enhances the practical applicability of LLMs in robotics but also paves the way for more intuitive and responsive robotic systems capable of complex interactions and tasks in varied real-world scenarios. In summary, our work investigates the following questions:

- How can we achieve embodied AI with a lightweight planner that considers different components, including user commands, visual perception, and action feasibility?
- How can robots make use of LLMs to translate human-language inputs to executable action directly?

<sup>1</sup>: K. T. Ly and I. Havoutis are with the Oxford Robotics Institute, University of Oxford. Email: {ktien, ioannis}@robots.ox.ac.uk.

<sup>2</sup>: K. Lu is with the Department of Computer Science, University of Oxford. Email: kai.lu@cs.ox.ac.uk.

- How can uncertainties in robotic planning be solved through human intervention?

## II. RELATED WORK

### A. LLMs for Robotic Systems

The application of LLM models in robotics, driven by recent advancements in natural language processing, has become a focal point of current research [12]. Silver et al. [13] study how a pretrained LLM can solve planning domain definition language (PDDL) problems. Similarly, Plansformer [4] requires PDDL domain inputs and fine-tunes the transformer model on planning problems. Although the PDDL interface is easy to use with the current task planner, it is difficult for end-users to add new tasks. In this work, we instead explore how to bridge the gap between human language and formal planning language task descriptions, reducing the need for engineer-handcrafted planning domains.

Liang et al. [5] use an LLM for generating code policies that can be executed on the robot. However, they do not account for the robot’s kinematics or geometry in the planning process, which is critical in real-world applications. Conversely, SayCan [1] and Text2Motion [14] are notable for considering geometric feasibility when planning the action sequence. PIGINET [15] instead introduces a transformer-based plan feasibility predictor to be integrated in a TAMP planner. The method fuses 6 cameras, making it harder to deploy in domestic settings. Our research diverges from these models by a multimodal structure that processes user input, visual and geometric information to output a sequence of actions that directly interfaces with the motion APIs.

Recently, applying vision-language models (VLMs) [16]–[19] in robotics emerged as another promising direction. Robotics transformer (RT)-2 [3] employs vision to enable a more generalizable approach to its former version RT-1 [2]. Although these models report robust embodied intelligence, collecting large-scale robot-specific data is required, making it less applicable and practical. Our work also develops a data-based solution, but from a robot-agnostic perspective using text-only data.

### B. Human-Robot Interaction for Autonomy

Recent research indicates that LLMs may not fully replicate human reasoning capabilities [20]. Consequently, integrating HITL is essential for maintaining the reliability and safety of robotic systems [21]. Vemprala et al. [7] implement ChatGPT with HITL, emphasizing the dual benefits of enhanced model training through human feedback and increased safety during operations. Ren et al. [22] focus on human-robot interaction, where the robots decide whether or not to ask for human intervention. The framework, however, remains a sole task planner that evaluates task reasoning without considering the actual robot constraints. The decision to ask the human or not depends on the planner’s confidence in generating a logic solution rather than an explicit failure. Moreover, the developed algorithm also relies on a scene description from the prompt and does not have any perception input. Our approach underlines the importance of human oversight in critical decision-making problems, and

facilitates a clear and intuitive interface where the human is informed during robot operation.

### C. Replanning for Failure Recovery

Failure recovery has been widely addressed in the literature for robotics planners. When it comes to LLM-based solutions, researchers typically address replanning from a second-level action check when the plan is generated [23], [24]. Vision-language models are also implemented to describe the error from the scene [25]. REFLECT [26] incorporates visual and audio sensory data to explain failures from the failed actions. This showcases the success of failure reasoning with a multimodal structure, however does not consider robotic constraints and remains a pure task planner. CAPE [27] instead proposes re-prompting to replan upon failure with preconditions similar to classical planning methods (e.g. PDDL, STRIPS). Considering the context of task and motion planning, this method works in the form of hierarchical layers and replans upon execution failure. In addition to recovering from execution, we designed a multimodal LLM-based planner that incorporates action feasibility during generating plans, which helps to reduce the time taken for re-querying the planner.

### D. Data-efficiency in LLM-based robotics planner

LLM-based models are usually data-intensive and require substantial training datasets and computing resources to generate robust solutions. In robotics, this demand poses a significant challenge, as collecting large, diverse, and representative datasets from physical robots is often prohibitively time-consuming, labor-intensive, and costly. For instance, RT series [2], [3] and  $\pi_0$  [19] introduces an end-to-end approach to applying transformers to robotics, requiring a large amount of data collected in months to map vision and language to actions. Similarly, Lynch et al. [28] proposes an interactive method for language-conditioned robot skills, compiling a dataset through 2.7k hours of collection on real robots. Prompt engineering, on the other hand, has shown promise in applying LLMs to robotic tasks in a zero-shot manner [29]–[31]. However, this is often not robust for real-world execution without considering and understanding the operating environment. Parameter-efficient fine-tuning (PEFT) [32] techniques have become a promising approach to minimize fine-tuning costs while enabling LLMs to acquire task-specific parameters. Our work employs this method to incorporate robotics knowledge to LLMs while introducing a widely applicable LLM-based framework that is easy to deploy on different robot platforms, as it does not require robot-specific data.

## III. METHODOLOGY

### A. Problem Statement

Our work involves a mobile manipulator interacting with non-expert users to perform a human request. *InteLiPlan* shares the same set of skills with the motion planner, which includes one-step actions like *go()*, *pick()*, *place()*, *open()*, *close()*, *search()*, *turn()*. During operation, the robot receives input from the user  $U$ , which is a request for the first input, or

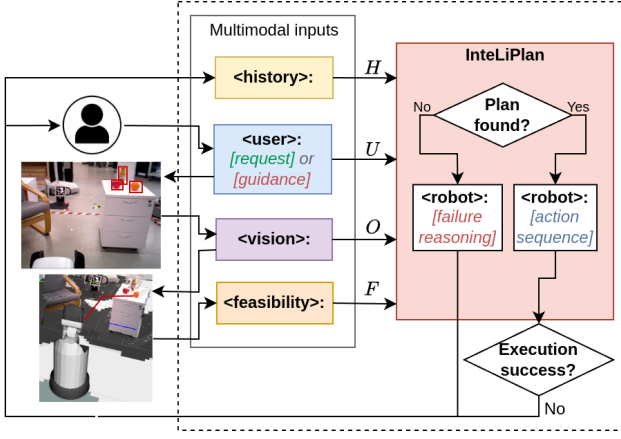


Fig. 2: System overview. Our multimodal planner integrates user input, visual perception, and action feasibility score as inputs to a fine-tuned lightweight LLM. The LLM will then generate an action sequence for a real robot to perform mobile manipulation tasks.

guidance for recovery. The centralized planner, *IntelLiPlan*, takes in the user input and the internal vision and motion verification functions, in order to generate a feasible plan. If a plan is found, the sequence of actions will be executed from a predefined skill library. If a failure is detected, the robot will notify the reason and ask for human instruction.

The robot obtains environmental information from a vision module, a state-of-the-art object detection model, processing input from the onboard RGB-D camera. The visual detection is represented as  $O$ . The motion verification module assesses the feasibility score  $F$  of the action during planning.

The result is *IntelLiPlan*, a real-time onboard LLM-based planner, with the ability to interact with humans in natural language and replan upon guidance. We evaluate *IntelLiPlan* on domestic tasks, aiming at a human-centric and interpretable framework.

### B. Interactive Planner with Multimodal Perception

Our planner uses a fine-tuned LLM model, which takes multimodal inputs from the user ( $U$ ), visual observation ( $O$ ), and motion feasibility ( $F$ ), as shown in Fig. 2. Before feeding into the central planner, the system evaluates the inputs in the following order:  $U$ ,  $O$ ,  $F$ . Algorithm 1 explains the internal interaction between the modules. When the user sends a command to the central planner, the object list is extracted and sent to the vision module. The object detector runs as a service that detects objects in real-time, and stores a list of detected objects with their corresponding positions. When a task is sent to the robot, the system checks if the mentioned items are found in the current observation. If the object is found, the motion verification module gets the object position and checks if the target object is reachable. For example, in our implementation, we use the reachability graph designed in R-LGP [33], which is used as a feasibility check service for LGP framework [34]. This reachability graph takes in the desired end-effector position (e.g. object location –  $object\_loc$ ) and current end-effector position and determines whether or not there is a whole-body collision-free trajectory

between the two end points with  $F.get\_score(object\_loc)$ . In *IntelLiPlan*, we incorporate a binary score for feasibility check, which applies, but not limited to the reachability graph in R-LGP. Last, we keep track of the conversation by inserting a history  $H$  to each planner call.  $H$  is initialized empty and records the human-robot interaction over time.

### Algorithm 1 *IntelLiPlan* pipeline

**Require:** user  $U$ , robot  $R$ , interaction history  $H$ , vision module  $O$ , feasibility module  $F$

```

1:  $inputs \leftarrow U.input()$ 
2:  $objects\_list \leftarrow R.extract(inputs)$ 
3:  $H \leftarrow None$ 
4: while not ( $R.task\_is\_complete()$  or  $time\_out()$ ) do
5:    $objects\_loc \leftarrow O.get\_position(objects\_list)$ 
6:    $result \leftarrow R.plan(H, inputs, O.is\_seen(objects\_list),$ 
      $F.get\_score(objects\_loc))$ 
7:   if  $R.get\_failure(result)$  then
8:      $H \leftarrow H + inputs + R.get\_failure(result)$ 
9:      $inputs \leftarrow U.get\_guidance()$ 
10:  else
11:     $R.execute(R.get\_plan(result))$ 

```

*IntelLiPlan* is formulated as  $A = \phi(H, U, O, F)$ , where  $A = \{a_t, \dots, a_{t+T}\}$  represents the output sequence of actions. While  $U$  is the direct input from the user,  $H$ ,  $O$  and  $F$  are the internal conversations of the system.  $H$  is a quoted string that tracks the previous conversation for contextual information. We design  $O$  as a text-based feedback from the visual observation, and  $F \in \{0, 1\}$  as a binary signal indicating the feasibility of the action. The multimodal inputs to the planner are represented in specific tokens:  $\langle history \rangle$ ,  $\langle user \rangle$ ,  $\langle vision \rangle$ ,  $\langle feasibility \rangle$ , where the token for the output is  $\langle robot \rangle$ . One example of the multimodal conversation is shown in Fig. 1.

Our framework features a centralized LLM-based planner to utilize the integration and coordination between submodules. We provide a plug-and-play modular design to leverage state-of-the-art models across various research domains (e.g. object recognition) in a zero-shot way.

### C. Failure Recovery

The ability to replan task and motion upon a failure is a critical feature towards real-world applications. This can be accomplished by active exploration, or with guidance from humans. A recent paper from [35] solves a self-recovery problem with an LLM-based planner. Our work instead addresses cases where human intervention is required. For example, when the task is not clear to the robot; when the user has specific preference towards choices; or when the robot is completely blocked by its partial observation of the environment. We blend our method with the human-in-the-loop mechanism, taking advantage of human input for corner cases. By integrating human insights, *IntelLiPlan* can navigate complex environments more effectively and perform tasks with greater precision, thereby increasing the system’s overall autonomy and operational safety.

We consider three main categories of failures: planning failure, task confusion and execution failure, as outlined in Table I. The multimodal structure allows the planner to reason around failures through feedback from the vision and motion verification modules, extracting the missing information to complete the task. When a failure occurs, our system can report to the human with a contextual description, interpreting the current observation and failure that occurred. By fine-tuning the LLM, the model can intelligently process human inputs and replan based on the human guidance. For instance, vision failures due to partial observation or adverse lighting conditions can be mitigated with human insights, leveraging the planner’s capability to maintain a record of initial commands through the internal dialogue  $H$  and to replan upon system failure as necessary, as demonstrated in Algorithm 1. This capability significantly enhances transparency, robustness and reliability. By injecting fine-tuning data, the output is guaranteed to consider necessary observations before executing any actions.

TABLE I: Cases that require human interaction.

Category	Case	Examples
Planning failure	Vision failure	Failed to find object
	Feasibility failure	Unable to reach object
Task confusion	Ambiguous reference	Found more than 1 specified object
	Ambiguous task	The user asks for "something to drink"
Execution failure	Action failure	Ask for recovery

#### D. Text-only Fine-Tuning Data for Robotics Planner

Recently, prompt engineering has become popular for implementing ready-to-use LLM-based planners, but poses a risk of plan reliability when executing on embodied agents. In addition, outputs from such planners often need translation into robotic control commands. Our method implements few-shot fine-tuning to inject the multimodal structure to the LLM, as well as inform the planner of the desired output format that enables direct calls to the action API. This approach also lays out the context and reduces the need for additional conversion, which uses computing resources that can prevent real-time planning.

To fine-tune the LLM model, we collected a pure-text customized dataset that includes scenarios both with and without failure recovery, where the interactions in failure cases are limited to a maximum of two rounds. The dataset is formulated as  $D = \{H_i, U_i, O_i, F_i, R_i | i = 0, 1, 2, \dots, N\}$ . We loop through a list of objects and commands to generate the command and expected sequence of actions. Visual observation  $O_i$  and feasibility score  $F_i$  is provided in the form of text. This text-based-only fine-tuning approach enables the model to easily adapt to different robots without further modifications or re-parameterization.

Our methodology employs Low-Rank Adaptation (LoRA) [36] to fine-tune LLM efficiently. Instead of updating all model parameters, LoRA introduces trainable low-rank matrices into key projection layers, significantly reducing com-

putational and memory overhead. LoRA is one of the PEFT [32] techniques that has also been proven to be superior over full fine-tuning as a parameter-efficient fine-tuning method in [37]. We implement supervised fine-tuning with the collected dataset, where  $\{H_i, U_i, O_i, F_i\}$  is the input, and  $R_i$  is the desired output. This approach enables efficient adaptation of the model to domain-specific tasks while maintaining scalability.

## IV. EXPERIMENTAL RESULTS

**Experimental Setup.** We evaluate our approach using the Toyota Human Support Robot (HSR) in a domestic environment. The first set of experiments is a ‘fetch me’ task, in an environment that includes 10 seen objects and 20 unseen objects, and the planner is evaluated on 4 seen request commands, 23 unseen request commands, and replan with 9 seen guidance commands, and 18 unseen guidance commands. The unseen commands represent the variants in human-like conversation. We then scale up the model to test with the 101 task requests from SayCan [1]. The vision module receives inputs from the RGB-D camera of the HSR, then processes object recognition with YOLO [38] and uses pose estimation to obtain object locations. In this way, the robot can automatically verify the presence and precise location of objects relative to the robot’s map frame. For feasibility verification, we use the reachability graph from [33], a sampling-based approach that efficiently checks for paths from a starting point to a goal while ensuring feasible configurations for the robot reaching towards the target object. The LLM model for our approach is fine-tuned from Mistral 7B [10] and Deepseek 8B [39], state-of-the-art lightweight LLM models, which is suitable for onboard deployment. The fine-tuning data includes all failure categories outlined in Table I. Before the execution, the robot is provided with a map of the room, e.g. large furniture like tables or drawers, and a list of motion APIs that are shared between the planner and the controller.

**Metrics.** To assess the performance of the proposed method we use five metrics. 1) *Task planning success rate*: the ratio of the no-failure trials where the robot successfully generates the action sequence by understanding the assigned task. 2) *Failure explanation success rate*: the ratio of the trials where the robot successfully generates the reason for not being able to find the action sequence. 3) *Failure recovery success rate*: the ratio of the trials where the robot successfully generates the recovery plan in response to human guidance. 4) *Execution success rate*: the ratio of the trials where the robot successfully executes the given request. 5) *Inference time*: the time taken to proceed the formulated multimodal inputs, showcasing the real-time capability of the model.

**Baselines.** We compare *Mistral-Ours* and *Deepseek-Ours* with few-shot prompting on Mistral - *Mistral-Prompt* and Deepseek - *Deepseek-Prompt*. The samples are in the same form as the dataset for fine-tuning, as discussed in Section III-D. The second comparison is with *SayCan* [1]. Given that this method does not have failure recovery ability, we only report the result during the scalability experiments.



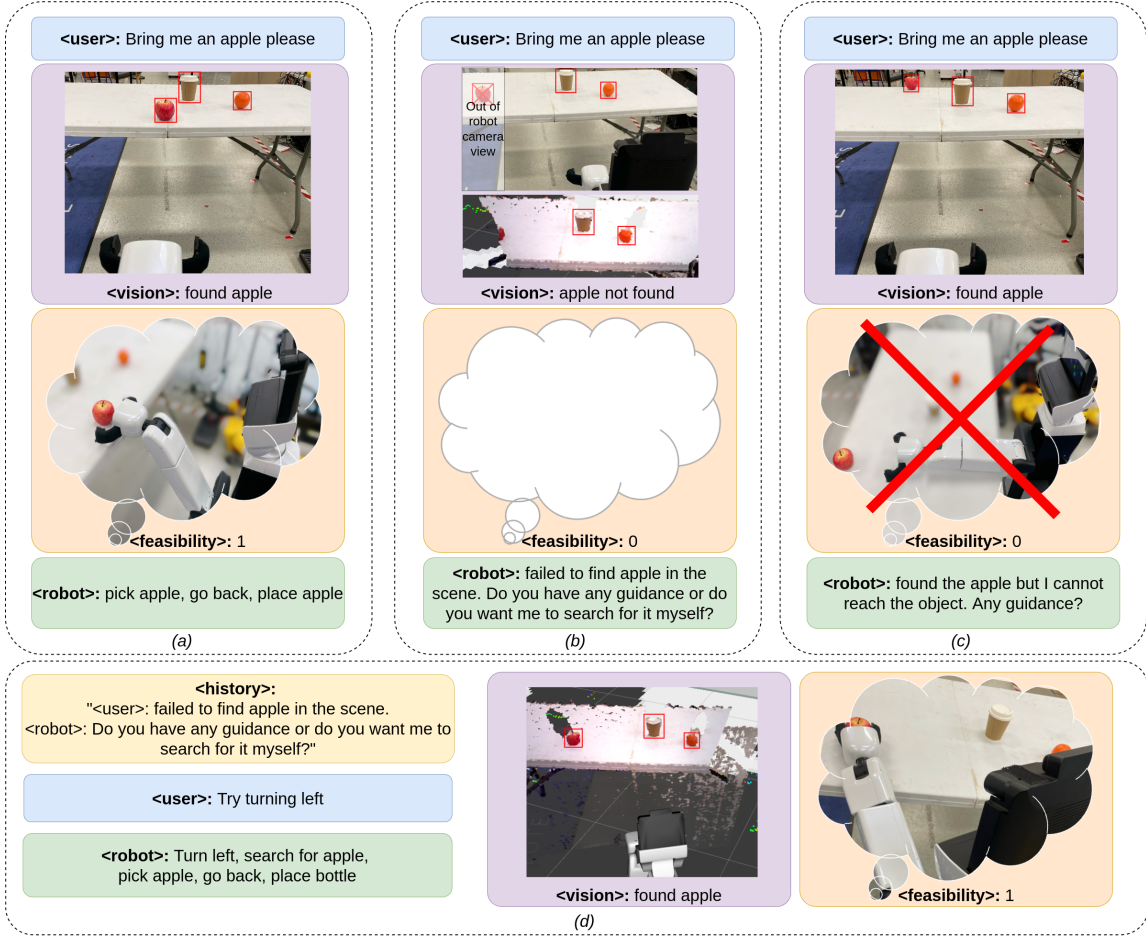


Fig. 3: Examples of the multimodal LLM-based planner. (a) presents a no-failure case for the given command. (b) and (c) depict the failure reasoning ability of the model considering the inputs, for vision and feasibility failures respectively. (d) showcases the ability to recover from the failure in (b) with human instruction.

Generative result from *PaLM 8B* [40] is also reported as a lightweight LLM baseline.

#### A. Comparison of Different Methods in Feasible Tasks

In our first experiment, we assume successful outcomes in both object detection and robotic motion feasibility. Here, the planner is solely required to generate the appropriate action sequences based on human inputs. An example of such conversation is provided in Fig. 3a.

Despite being told to list a sequence of actions only, *Mistral-Prompt* and *Deepseek-Prompt* outputs include redundant text. For a fairer comparison, we only extracted the outputted action sequence when determining the success rate, and ignored the lengthy explanation and slightly misused words (e.g. 'pick up' vs 'pick') for action API interface. As presented in the *w/o failure* column of Table II, the results demonstrate a superior performance of our models, achieving 100% success in scenarios with seen commands and objects, and maintaining high effectiveness even with unseen objects or commands. In contrast, *Mistral-Prompt* and *Deepseek-Prompt* experienced performance drops in scenarios involving unseen elements, with more sensitivity over unseen textual commands. It is also observed that *Deepseek-Prompt* behaves better than *Mistral-Prompt*, while *Deepseek-*

*Ours* has slightly lower success rates than *Mistral-Prompt*.

#### B. Evaluation of Robustness in Failure Explanation

In this test, we break down the evaluation of failure explanation in the failure recovery process, investigating the ability of the model to explicitly describe what hinders the process of planning actions. The failure experiments include all categories outlined in Table I, with examples of 'planning failure' detection in Fig. 3b,c. We only include 'execution failure' in failure recovery experiments (Section IV-C), as it is a failure description that is triggered from the controller level.

Figure 4 shows that our method achieved near-perfect success rates in scenarios with seen commands and seen/unseen objects, underscoring its reliability in familiar settings. It is observed that 'Task confusion' failure detection obtains lower success rates than multimodal failure. This is mainly due to the 'ambiguous task' experiments, where the robot fails to distinguish object categories (e.g. fruits, drinks) from the list of detected objects. Our method increases the success rates of failure explanation by more than twice in these cases, proving that fine-tuning triggers the domain-specific understanding of the robots for which factor (e.g. user preference) to consider during planning.

TABLE II: Success Rates (%) of Task Planning.

Scenario	Method	w/o failure	w/ failure
Seen cmd + Seen obj	Deepseek-Prompt	100	100
	Deepseek-Ours	100	100
	Mistral-Prompt	100	52
	Mistral-Ours	100	100
Seen cmd + Unseen obj	Deepseek-Prompt	90	35
	Deepseek-Ours	100	86
	Mistral-Prompt	90	29
	Mistral-Ours	100	90
Unseen cmd + Seen obj	Deepseek-Prompt	98.95	54
	Deepseek-Ours	100	93
	Mistral-Prompt	86.32	49
	Mistral-Ours	100	94
Unseen cmd + Unseen obj	Deepseek-Prompt	95.79	39
	Deepseek-Ours	97.89	93
	Mistral-Prompt	78.95	36
	Mistral-Ours	100	95




Fig. 4: Success Rates of Failure Explanation.

### C. Evaluation of Robustness in Failure Recovery

This section evaluates the ability of our framework’s capacity to effectively replan from human instructions. For each trial, we assume a detected failure in `<history>`, and expect the models to output corresponding action to the user inputs.

Figure. 5 demonstrate our framework’s performance in replanning in various scenarios. While ‘task confusion’ mainly processes user preference (e.g. user prefers the robot to get a coke to a 7up), ‘planning failure’ and ‘execution failure’ cases assess the ability of mapping language to the list of pre-defined motion (e.g. output ‘open cupboard’ in response to ‘the object is in the cupboard’). The lower success rates in recovering from ‘execution failure’ are due to the longer sequence of recovery than what was provided in the fine-tuning data. In general, prompting methods fails to embody LLM with sufficient robot-domain understanding. Our method significantly increases the performance of LLM models in robotic task failure recovery, indicating robustness in real-world conditions where adaptability and responsiveness to failures are crucial for operational success.

### D. Evaluation of Scalability

In this experiment, we evaluate the scalability of the approach by the 101 task instructions from SayCan [1]. To accommodate the task variants, we fine-tuned Mistral with the tasks outlined in Table III, which covers the expected actions required for the dataset solutions.




Fig. 5: Success Rates of Failure Recovery.

TABLE III: Tasks with Corresponding Expected Outcome.

Case	Expected plan
Pick object	Pick object
Go to destination	Go to destination
Fetch me	Pick object, Go back, Place object
Put away	Pick object, Go to destination, Place object
Put in drawer	Open drawer, Pick object, Place in drawer, Close drawer

Table IV presents the tested result of the frameworks with the dataset from Saycan. As a 540B model, the SayCan result explains its excellent decision-making capability with throughout understanding of the world. Its trained affordance value provides sufficient embodied knowledge, ensuring high success rate in execution. Besides, we note that prompting with *PaLM 8B* only successfully plans 38% cases. *Mistral-Prompt* with our modular structure helps the LLM model to gain embodied intelligence, with the planning success rate to 59%. It is observed that *Mistral-Prompt*’s failure cases come from its lack of sense of the operating environment. For example, some of the results tell the robot to ‘go to store’ to pick something up, despite being told that it is working in a domestic environment. It is proven that with 7B parameters, our *InteLiPlan* obtains similar results in comparison to the state-of-the-art 504B SayCan model with 83% planning success rate.

The execution success rate remains relative to the planning success rate in *Mistral-Prompt* and *Mistral-Ours*, given the same vision and reachability modules check for the system. Having the reachability graph checking whole-body feasibility explicitly instead of using skill probabilities like SayCan, our approach also increases the chances of successfully executing the plan once generated. This is shown by the reduced difference in plan and execute success rates between SayCan and our pipeline. On the other hand, as we use a sampling-based approach, there is a risk of not finding a path even if there is one exists, leading to failures from the planning level. This is where the failure recovery comes in. Another solution is to use more samples in the graph such that they sufficiently provide the reachability knowledge. Since *InteLiPlan* is a plug-and-play system, it can seamlessly integrate future state-of-the-art feasibility modules.

TABLE IV: Plan and Execution Success Rates (%) of the methods for 101 SayCan task descriptions.

Methods	Plan	Execute
PaLM-8B	38	n/a
PaLM-504B SayCan	84	74
Mistral-7B Prompt	59	58
Mistral-7B Ours	83	82

TABLE V: Breakdown of onboard planning time (s).

Vision query	Feasibility query	Planner query
6e-6	5	1.5

### E. Real-robot experiments

We validate the efficiency of the system on the physical HSR. We implemented 3 sets of experiments:

- *No-failure task*: the robot is asked to complete a task by processing a high-level command. In this task, we guarantee that the object is within the observation and is reachable.
- *1-step failure task*: we implement the planning failure from vision, where the targeted object is not within its observation, the robot is expected to report the failure to the user and replan upon guidance.
- *Multi-step failure task*: the robot must detect failures and replan twice to be able to complete the task. Specifically, we implement the planning failure from vision first, then after the targeted object is found, it is out of reach and the robot should ask again to complete the task.

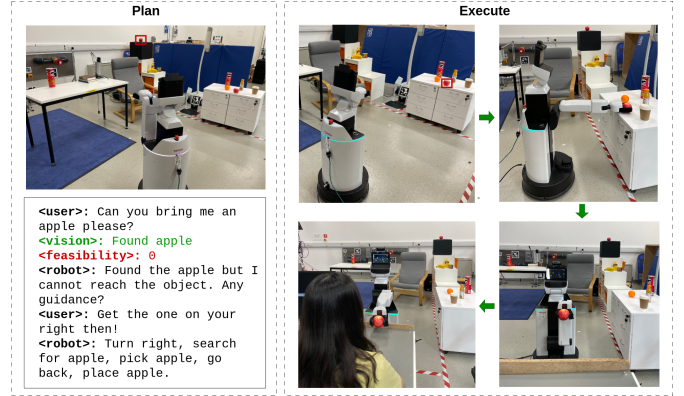
The result demonstrates that by processing both vision and reachability capabilities, all planned action sequences can be executed with collision-free trajectories. Fig. 1 and Fig. 6 showcase *InteLiPlan* results in real-world settings, which includes tasks with and without failure recovery. The multi-step failure verifies that our method enables autonomy with seamless end-user intervention.

Table V records the averaged planning time of the system over the trials, which is divided into vision, feasibility and planner queries processing time. In total, the model can plan an executable action sequence in less than 7s. The feasibility check module, which takes up the longest processing time, can be relaxed with fewer nodes in the reachability graph. Notably, using YOLO instead of VLM for vision input significantly contributes to real-time and onboard processing capability of the pipeline.

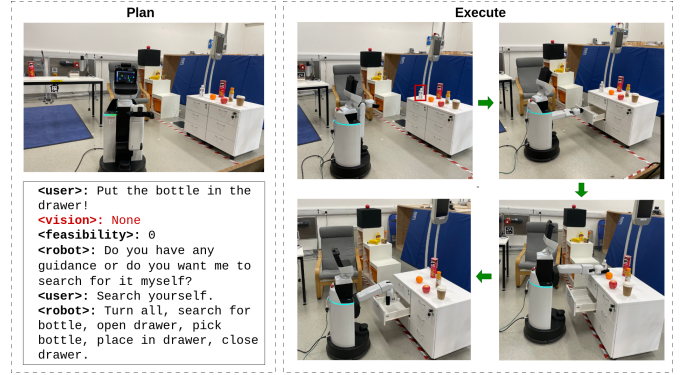
Videos of our demonstrations are available at our project page: <https://kimtienly.github.io/InteLiPlan>.

## V. CONCLUSION

We presented *InteLiPlan*, an interactive lightweight LLM-based robotics planner for reliable and robust autonomy in domestic environments. Our framework employs a conversation-like format between internal modules and facilitates human-robot interaction, allowing the system to reason about failures through a multimodal input formulation. By incorporating a human in the loop, the robot gains the ability



(a) Failure from feasibility



(b) Failure from vision

Fig. 6: Demonstration of the interactive failure recovery capability on the physical HSR.

to replan based on instructions, effectively utilizing human input for dynamic adjustments. We performed extensive evaluations to investigate common failures that happen in real scenarios, where human instruction is required. The interactive behaviors guarantee interpretability and reliability for robots operating in domestic environments.

With roughly 300 data samples for fine-tuning, our approach effectively handles the targeted tasks by leveraging the human-like text understanding capabilities of LLMs. This reduces the effort of defining traditional task planning domains in robotics, provides an intuitive, robot-independent data structure and boosts applicability in resource-constrained applications. Notably, deploying our approach with the lightweight Mistral 7B model achieves both comparable results with the SOTA baseline and – notably to our approach– real-time onboard computing. Our system was validated on the physical Toyota HSR robot.

*Limitations and Future Work* - Our framework is structured as a multimodal system with a feasibility check, however, it does not have a motion-level replanning feature. Incorporating low-level reactions together with high-level replanning capabilities in the form of a dual process can potentially improve failure recovery speed. Furthermore, the sampled path formed by the reachability graph can be further developed into a motion planning subpart of our approach, resulting into a full task and motion planning stack.

## REFERENCES

- [1] A. Brohan, Y. Chebotar, C. Finn, K. Hausman, A. Herzog, D. Ho, J. Ibarz, A. Irpan, E. Jang, R. Julian *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” in *Conference on robot learning*. PMLR, 2023, pp. 287–318.
- [2] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu *et al.*, “Rt-1: Robotics transformer for real-world control at scale,” *arXiv preprint arXiv:2212.06817*, 2022.
- [3] B. Zitkovich, T. Yu, S. Xu, P. Xu, T. Xiao, F. Xia, J. Wu, P. Wohlhart, S. Welker, A. Wahid *et al.*, “Rt-2: Vision-language-action models transfer web knowledge to robotic control,” in *Conference on Robot Learning*. PMLR, 2023, pp. 2165–2183.
- [4] V. Pallagani, B. Muppasani, K. Murugesan, F. Rossi, L. Horeh, B. Srivastava, F. Fabiano, and A. Loreggia, “Plansformer: Generating symbolic plans using transformers,” in *NeurIPS 2023 Workshop on Generalization in Planning*.
- [5] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng, “Code as policies: Language model programs for embodied control,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9493–9500.
- [6] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, “Language models are unsupervised multitask learners,” 2019.
- [7] S. Vemprala, R. Bonatti, A. Buckner, and A. Kapoor, “Chatgpt for robotics: Design principles and model abilities,” Microsoft, Tech. Rep. MSR-TR-2023-8, February 2023.
- [8] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample, “Llama: Open and efficient foundation language models,” 2023.
- [9] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale *et al.*, “Llama 2: Open foundation and fine-tuned chat models,” *arXiv preprint arXiv:2307.09288*, 2023.
- [10] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. d. l. Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier *et al.*, “Mistral 7b,” *arXiv preprint arXiv:2310.06825*, 2023.
- [11] J. Wang, Z. Wu, Y. Li, H. Jiang, P. Shu, E. Shi, H. Hu, C. Ma, Y. Liu, X. Wang *et al.*, “Large language models for robotics: Opportunities, challenges, and perspectives,” *arXiv preprint arXiv:2401.04334*, 2024.
- [12] F. Zeng, W. Gan, Y. Wang, N. Liu, and P. S. Yu, “Large language models for robotics: A survey,” *arXiv preprint arXiv:2311.07226*, 2023.
- [13] T. Silver, V. Hariprasad, R. S. Shuttlesworth, N. Kumar, T. Lozano-Pérez, and L. P. Kaelbling, “PDDL Planning with Pretrained Large Language Models,” in *NeurIPS 2022 Foundation Models for Decision Making Workshop*, 2022.
- [14] K. Lin, C. Agia, T. Migimatsu, M. Pavone, and J. Bohg, “Text2motion: From natural language instructions to feasible plans,” *Autonomous Robots*, vol. 47, no. 8, pp. 1345–1365, 2023.
- [15] Z. Yang, C. Garrett, T. Lozano-Perez, L. Kaelbling, and D. Fox, “Sequence-based plan feasibility prediction for efficient task and motion planning,” in *Robotics science and systems*, 2023.
- [16] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, “Learning transferable visual models from natural language supervision,” in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [17] J. Li, D. Li, C. Xiong, and S. Hoi, “Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation,” in *International conference on machine learning*. PMLR, 2022, pp. 12 888–12 900.
- [18] M. Shridhar, L. Manuelli, and D. Fox, “Cliport: What and where pathways for robotic manipulation,” in *Conference on robot learning*. PMLR, 2022, pp. 894–906.
- [19] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter *et al.*, “ $\pi_0$ : A vision-language-action flow model for general robot control,” *arXiv preprint arXiv:2410.24164*, 2024.
- [20] J. Kaddour, J. Harris, M. Mozes, H. Bradley, R. Raileanu, and R. McHardy, “Challenges and applications of large language models,” *arXiv preprint arXiv:2307.10169*, 2023.
- [21] C. Zhang, J. Chen, J. Li, Y. Peng, and Z. Mao, “Large language models for human-robot interaction: A review,” *Biomimetic Intelligence and Robotics*, p. 100131, 2023.
- [22] A. Z. Ren, A. Dixit, A. Bodrova, S. Singh, S. Tu, N. Brown, P. Xu, L. Takayama, F. Xia, J. Varley *et al.*, “Robots that ask for help: Uncertainty alignment for large language model planners,” *Proceedings of Machine Learning Research*, vol. 229, 2023.
- [23] M. Ahn, M. G. Arenas, M. Bennice, N. Brown, C. Chan, B. David, A. Francis, G. Gonzalez, R. Hessmer, T. Jackson *et al.*, “Vader: Visual affordance detection and error recovery for multi robot human collaboration,” *arXiv preprint arXiv:2405.16021*, 2024.
- [24] F. Joubin, A. Ceravola, P. Smirnov, F. Ocker, J. Deigmoeller, A. Belardinelli, C. Wang, S. Hasler, D. Tanneberg, and M. Gienger, “Copal: corrective planning of robot actions with large language models,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 8664–8670.
- [25] Z. Wang, S. Cai, G. Chen, A. Liu, X. Ma, Y. Liang, and T. CraftJarvis, “Describe, explain, plan and select: interactive planning with large language models enables open-world multi-task agents,” in *Proceedings of the 37th International Conference on Neural Information Processing Systems*, 2023, pp. 34 153–34 189.
- [26] Z. Liu, A. Bahety, and S. Song, “Reflect: Summarizing robot experiences for failure explanation and correction,” in *Conference on Robot Learning*. PMLR, 2023, pp. 3468–3484.
- [27] S. S. Raman, V. Cohen, I. Idrees, E. Rosen, R. Mooney, S. Tellex, and D. Paulius, “Cape: Corrective actions from precondition errors using large language models,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 14 070–14 077.
- [28] C. Lynch, A. Wahid, J. Tompson, T. Ding, J. Betker, R. Baruch, T. Armstrong, and P. Florence, “Interactive language: Talking to robots in real time,” *IEEE Robotics and Automation Letters*, 2023.
- [29] W. Huang, P. Abbeel, D. Pathak, and I. Mordatch, “Language models as zero-shot planners: Extracting actionable knowledge for embodied agents,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 9118–9147.
- [30] W. Huang, F. Xia, T. Xiao, H. Chan, J. Liang, P. Florence, A. Zeng, J. Tompson, I. Mordatch, Y. Chebotar *et al.*, “Inner monologue: Embodied reasoning through planning with language models,” in *Conference on Robot Learning*. PMLR, 2023, pp. 1769–1782.
- [31] 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.
- [32] N. Ding, Y. Qin, G. Yang, F. Wei, Z. Yang, Y. Su, S. Hu, Y. Chen, C.-M. Chan, W. Chen *et al.*, “Parameter-efficient fine-tuning of large-scale pre-trained language models,” *Nature Machine Intelligence*, vol. 5, no. 3, pp. 220–235, 2023.
- [33] K. T. Ly, V. Semenov, M. Risiglione, W. Merkt, and I. Havoutis, “Rlgp: A reachability-guided logic-geometric programming framework for optimal task and motion planning on mobile manipulators,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 14 917–14 923.
- [34] M. Toussaint, “Logic-geometric programming: an optimization-based approach to combined task and motion planning,” in *Proceedings of the 24th International Conference on Artificial Intelligence*, 2015.
- [35] M. Shirasaka, T. Matsushima, S. Tsunashima, Y. Ikeda, A. Horo, S. Ikoma, C. Tsuji, H. Wada, T. Omija, D. Komukai *et al.*, “Self-recovery prompting: Promptable general purpose service robot system with foundation models and self-recovery,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 17 395–17 402.
- [36] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen *et al.*, “Lora: Low-rank adaptation of large language models,” *International Conference on Learning Representations*, vol. 1, no. 2, p. 3, 2022.
- [37] M. Weysow, X. Zhou, K. Kim, D. Lo, and H. Sahraoui, “Exploring parameter-efficient fine-tuning techniques for code generation with large language models,” *ACM Transactions on Software Engineering and Methodology*, 2023.
- [38] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [39] D. Guo, D. Yang, H. Zhang, J. Song, R. Zhang, R. Xu, Q. Zhu, S. Ma, P. Wang, X. Bi *et al.*, “Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning,” *arXiv preprint arXiv:2501.12948*, 2025.
- [40] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann *et al.*, “Palm: Scaling language modeling with pathways,” *Journal of Machine Learning Research*, vol. 24, no. 240, pp. 1–113, 2023.