

Interactive Planning Using Large Language Models for Partially Observable Robotic Tasks

Sun, Lingfeng; Jha, Devesh K.; Hori, Chiori; Jain, Siddarth; Corcodel, Radu; Zhu, Xinghao; Tomizuka, Masayoshi; Romeres, Diego

TR2024-052 May 14, 2024

Abstract

Designing robotic agents to perform open vocabulary tasks has been the long-standing goal in robotics and AI. Recently, Large Language Models (LLMs) have achieved impressive results in creating robotic agents for performing open vocabulary tasks. However, planning for these tasks in the presence of uncertainties is challenging as it requires “chain-of-thought” reasoning, aggregating information from the environment, updating state estimates, and generating actions based on the updated state estimates. In this paper, we present an interactive planning technique for partially observable tasks using LLMs. In the proposed method, an LLM is used to collect missing information from the environment using a robot, and infer the state of the underlying problem from collected observations while guiding the robot to perform the required actions. We also use a fine-tuned Llama 2 model via self-instruct and compare its performance against a pre-trained LLM like GPT-4. Results are demonstrated on several tasks in simulation as well as real-world environments.

IEEE International Conference on Robotics and Automation (ICRA) 2024

Interactive Planning Using Large Language Models for Partially Observable Robotic Tasks

Lingfeng Sun^{1,2}, Devesh K. Jha², Chiori Hori², Siddarth Jain²
 Radu Corcodel², Xinghao Zhu^{1,2}, Masayoshi Tomizuka¹, Diego Romeres²

Abstract—Designing robotic agents to perform open vocabulary tasks has been the long-standing goal in robotics and AI. Recently, Large Language Models (LLMs) have achieved impressive results in creating robotic agents for performing open vocabulary tasks. However, planning for these tasks in the presence of uncertainties is challenging as it requires “chain-of-thought” reasoning, aggregating information from the environment, updating state estimates, and generating actions based on the updated state estimates. In this paper, we present an interactive planning technique for partially observable tasks using LLMs. In the proposed method, an LLM is used to collect missing information from the environment using a robot, and infer the state of the underlying problem from collected observations while guiding the robot to perform the required actions. We also use a fine-tuned Llama 2 model via self-instruct and compare its performance against a pre-trained LLM like GPT-4. Results are demonstrated on several tasks in simulation as well as real-world environments.

I. INTRODUCTION

Designing robots that have the physical intelligence to perform open vocabulary tasks is extremely challenging. This requires that robots be able to interpret tasks from an open set of instructions and execute them robustly while performing the required reasoning. One can argue that this could be the most challenging problem facing artificial intelligence (AI). However, designing such agents can truly revolutionize the way robots would be integrated into our future society. Recently, large language models (LLMs) [1]–[3] have been shown to be very impressive at solving tasks of different complexities [4]–[8]. Large language models can help understand the tasks and decompose them into a sequence of actions, reward functions, or goals for policy given appropriate prompts and training data. Motivated by these developments, we present a problem of interactive planning in uncertain environments where a robot may not have complete information to perform the task. In these tasks, the robot needs to interact with its environment and collect additional information to complete the task.

Partial observability and uncertainty are the norm, rather than the exception, in the real world. For example, consider task T2 shown in Figure 1, where a robot needs to understand how it can gather information to identify the empty cup and then throw it in the bin. Unlike the tasks with complete information, it would be challenging to design a sequence of skills or a suitable reward function that can solve this task.

¹Mechanical Systems Control Lab, UC Berkeley, Berkeley, CA, USA {lingfengsun, zhuxh, tomizuka}@berkeley.edu
²Mitsubishi Electric Research Laboratories, Cambridge, MA, USA {jha, chori, sjain, corcodel, romeres}@merl.com

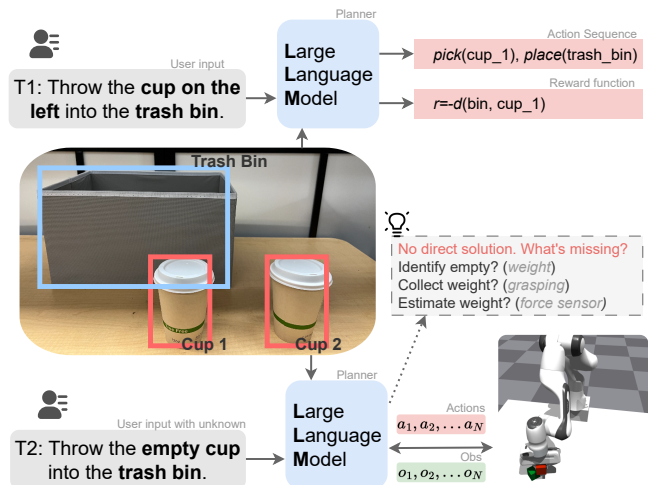


Fig. 1. An example task where the uncertainty is present in the content of the cups. For task T1, the robot is asked to throw the cup on the left into the bin. An LLM agent can generate feasible action sequences for the robot to perform the task. When asked to throw the empty cup, the agent cannot reason which cup is empty based on current information. It needs to interact with the cups and use feedback from observations (e.g., force sensor reading) to identify the empty cup.

This problem can be formulated as a Partially Observable Markov Decision Process (POMDP)[9]. However, solving POMDPs could be computationally intractable. It requires reasoning in the belief state of the problem and does not scale well with the dimensionality of the problem. Prior work on using LLMs for robotic tasks has demonstrated good reasoning capability of LLMs as well as mapping of the reasoning to robot actions [5], [10], [11]. Inspired by these advancements, we believe that we can leverage the reasoning and “chain-of-thoughts”(CoT) capability of LLMs to solve partially observable tasks while interacting with the environment. What makes this challenging for current LLMs is the requirement to understand real robot observations from different modalities and use them for task planning.

Most of the prior works using LLMs in robotics focused on step-wise scene and task understanding making full use of the current available modalities to infer the optimal action and/or reward [5], [6], [10], [11]. In this work, we focus on performing interactive planning under cases of partial observability. This requires planning to aggregate information from the environment, reasoning about the correct state of the system, and updating the state estimates based on the sensor observations collected by the robot. Furthermore, we also try to understand how well a fine-tuned smaller model like Llama2-7B [3] performs in comparison with a pre-trained LLM like GPT-4. The smaller models are generally desirable

for practical reasons but it could be challenging to distill the reasoning capability of models like GPT-4 for complex robotic tasks discussed in this paper. To understand this, we propose an instruction data generation pipeline following the self-instruction [12] scheme to understand the limitations of smaller models and potential ways to overcome them.

In summary, our work makes the following contributions:

- We introduce the Large Language Model for Partially Observable Task Planning (LLM-POP) framework to interactively plan for tasks with uncertainties. We demonstrate the framework in simulation and real-world environments.
- We compare the performance of pre-trained LLMs with fine-tuned smaller models in the proposed framework.

II. RELATED WORK

LLMs for embodied AI tasks: The recent strides in the field of LLMs have renovated the robotics community in various task domains, especially on task planning [4], [7], [13]–[15], where the robot is asked to reason about language instructions to generate robot actions, reward functions for online controllers [6], and code [8], [16]. Previous works also combine LLMs with Task and Motion Planning (TAMP) to make use of the traditional motion planning algorithms [17]–[19]. To solve robotics tasks in larger and more complex settings, researchers have utilized LLMs to process multimodal information, [5], [11], [20], [21], enable multi-robot [22] and human-robot collaboration [23]. A series of works use feedback in the planning process [5], [10], [23]–[26] to improve the LLM-generated plans, and use the natural language provided by the user to correct robot behaviors. Other works model the uncertainties in tasks [27], [28] and gather reasoning by involving human knowledge in the decision loop. In our problem setting, we focus on the partial observation environment setting where initial information is insufficient for task solving, and the robot actively seeks task-relevant information via sensory feedback.

LLMs for Data Synthesis and instruction tuning: Recent open-sourced models like LLaMa [29][3], Alpaca [30], and Gorrila [31], explore how pre-trained models can be improved for general and specific tasks [32]–[35]. Researchers have proposed strategies to prompt LLMs to perform high-quality synthetic data by in-context learning [36], task decomposition [24], [37], as well as methods to improve fine-tuning performance [38], [39], and pre-train model for specific tasks like coding [40]. Our research incorporates these methods to get synthetic data for instruction-tuning but focuses on tasks that require reasoning, but with limited labeled data.

III. INTERACTIVE LLM PLANNING WITH UNCERTAINTIES

The objective of the proposed framework is to perform long-horizon robotic tasks in the presence of various kinds of uncertainties using LLMs. These tasks require a closed-loop, interactive planning where the robot should be able to collect useful observations from the environment and then make optimal decisions. An example of such a task is illustrated in

Figure 1 where the robot’s task is to throw the empty cup into the bin. However there exists uncertainty in the contents of the cups, and therefore this information needs to be obtained by sensorimotor operations and provided as the feedback to the LLM. For clarity of presentation, this section delves into the formulation of the underlying problem using the notion of POMDPs. We then elaborate on the pivotal role of LLMs in the interactive planning framework.

A. Problem Formulation

1) **Partial Observation setting:** A POMDP is an extension of a traditional MDP that tackles decision-making scenarios where the agent lacks complete state information. A POMDP is defined by a tuple (S, A, P, R, Ω, O) , with Ω as the observation set and O as the observation function. At each time step, the environment is in state $s \in S$. The agent takes action $a \in A$ and causes the environment to transit to s' accordingly to the transition function $P(s'|s, a)$. At the same time step, the agent gets an observation $o \in \Omega$ which depends on the current state of the environment $O(o|s')$. Unlike the policy function in MDP $\pi(a|s)$, which maps the underlying states to the actions, POMDP’s policy $\pi(a|b)$ is a mapping from the belief states b to the actions. The belief state b is a probabilistic estimation of the full state s . The updated belief state b' after observing o is described by: $b'(s') = C \cdot O(o|s') \sum_{s \in S} P(s'|s, a)$ where C is a normalizing constant.

We also want the proposed framework to be generalizable to a variety of tasks. For different tasks τ , the information required to make decisions can differ. This adds additional complexity since now the LLM has to reason about a generalizable state space S . In the open-vocabulary robotics task scenarios, the robot observations are determined by on-board sensors. Not all information about the environment is relevant to the task; some of them can be directly extracted from observations, while some are unknown and require exploration. Thus, we end up getting task-dependent belief state b^τ , and the task-related states s^τ for task τ . Both finding the necessary state abstraction for different tasks and finding the optimal policy π under the task-specific MDP is important in this task-dependent POMDP setting.

2) **Action space of robots:** For long-horizon tasks, using a pre-trained set of parameterized skills as action space is a common choice. In this paper, we use a set of parameterized skills like **{pick, place, reach, reset}**. All these skills can be performed using robot observations and thus we do not consider partial observability during robot skills execution. It is noted that we do not consider continuous sensory feedback during skill execution— however, that could be incorporated by training skills using RL.

3) **Uncertainties in Tasks:** The uncertainty in decision-making in the tasks we test mainly arises from two aspects: **Environmental Uncertainty:** These uncertainties arise in the POMDP settings due to the agent’s lack of complete environmental knowledge. For example, physical properties of the objects that cannot be directly observed. The uncertainties in the belief b^τ can be reduced with certain observations. This is a major challenge we target to solve in this paper.

Skill Execution Uncertainty: Even with a well-defined plan, the actual execution of actions on robots might not always lead to the expected outcome. This can be mainly attributed to the difference between the transition functions P , P_{real} of the designed and real system as well as unexpected disturbances during execution.

With the challenges explained above, we propose a framework where LLMs are used as policy as well as for state abstraction for the underlying POMDP.

B. Language-based Planners

Based on the problems described in the previous section, we propose to use a LLM to play a multifaceted role in the interactive planning process:

1) *LLM for State Abstraction:* Given the environment description and sensor observations, LLM needs to analyze the available information and abstract the sufficient statistics (or the appropriate state) to solve the task. Furthermore, it needs to reason about what is uncertain based on the current observations. It needs to update its belief based on the observations when prompted with historical information.

2) *LLM as Policy:* Given the observation and action space, LLM needs to plan actions that gather environmental information to mitigate the uncertainty and update the agent’s belief state. The LLM-based policy is also expected to generate the optimal plan to maximize the reward based on the task description with minimal steps. Also, since we use open-loop parameterized skills for the robot, the LLM is also used to provide feedback to the robot in cases of failure in execution of these skills. This feedback needs to be provided in a way that is still executable by the robot.

We use LLM to reason about these problems during task execution. It is noted that actions in the POMDP setting is conditioned on new observations and updated beliefs. There are a few additional challenges when using LLM as closed-loop policy for tasks with uncertainties that we consider in the paper. To update the belief state of the task, the LLM must understand the robot observations from different modalities (pose detections, force sensors, etc.). These data formats might be new to the LLM model and thus, must be properly included in the prompt template to the LLM. Furthermore, the skills available to the robot are parameterized by continuous position and orientation coordinates which might be challenging to reason about while performing robotic tasks. Similarly, the output of the language model needs to be executable by the robot; the response should be written in a template that the downstream controller can understand. In the next section, we will discuss how we use the LLMs to solve the interactive planning task.

IV. LLM-POP: INTERACTIVE PLANNING FRAMEWORK

The proposed framework (LLM-POP) for interactive planning is illustrated in Figure 2. As introduced in the problem formulation, the language-based policy in our framework has multiple tasks to do in the planning loop. At each step, the input to the language model contains the **task description** from a user, **the current observation** from the robot, and the

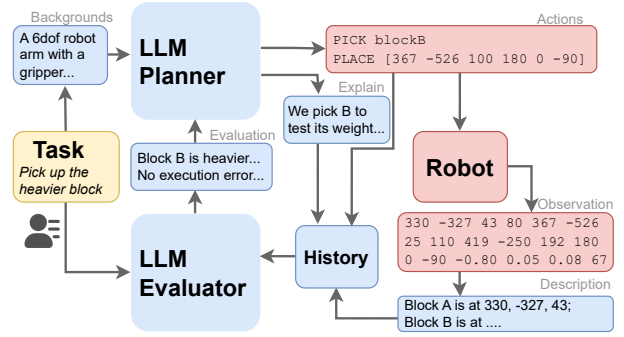


Fig. 2. An example showing how the framework works during solving the task “Pick up the heavier block”. The LLM planner outputs an executable action sequence to the robot. The robot executes the action, and the observation description and action pair are added into the history buffer. The LLM evaluator analyzes the historical information and outputs the updated information to the planner as intermediate info to generate a new plan.

historical action and observation sequence from previous steps. The model output includes an executable **sequence of actions** and the **corresponding text explanation**. The robot will execute the actions provided by the policy output and return the observations for a next-round query of the LLM. The language model must finish the reasoning task and output the policies in the designed format. The task description is the only user-provided input during the planning process. In the following sections, we show how we use a pre-trained LLM (GPT-4) as well as a fine-tuned smaller model to serve as the planner.

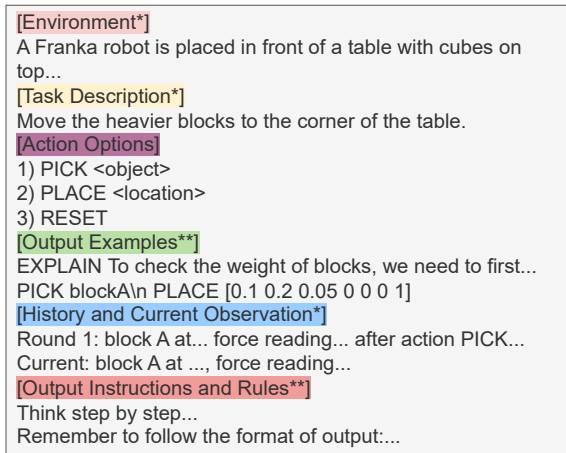
A. Prompt structure for GPT-4

Using powerful LLMs like GPT-4 as interactive planners relies on its strong chain-of-thought reasoning and in-context learning capability. Therefore, the prompt (input of a single round LLM query) to the LLM requires careful design to ensure it can generalize to robotics tasks and avoid hallucination (generating actions in wrong formats or not executable for the robot) in responses.

As shown in Figure 3, the prompt template for the *planner* consists of the following parts:

- *Environment description, action options, output rules:* Background information that help understand the task settings. This information is preset by the user and is constant throughout planning for different tasks.
- *Task description:* texts describing tasks from users.
- *Example outputs:* in-context examples for planning.
- *Current observation and historical information:* text-format descriptions of current observation and historic information. If the observation is poses and force, use vectors with explanations.

The explanation in output, together with the action sequence, will be included in historical information. This helps the LLM to understand the past actions it has performed and avoid reasoning about it again. Note that the LLM planner needs to specify the parameters in the actions based on its own understanding of the environment, task, and the action space description. For manipulation tasks, this includes location and orientation for the target pose.



* same for evaluator ** adjust for evaluator

Fig. 3. Prompt template for GPT planner and evaluator. The task description is taken from user input, the others are pre-defined according to the environment and robot. Output rules will be adjusted for the evaluator.

As shown in Figure 2, along with the LLM *planner*, we also designed an LLM *evaluator* using a similar prompt structure. The evaluator also takes in the background information, task description, and history observations after executing past actions. It evaluates the task execution status and appends it to next-round prompting. As described in Section III-B, the evaluator here will explicitly ask the LLM to finish the “state abstraction” (analyze what’s the missing information), “belief update” in policy (analyze information from historical observations), and “correct execution errors”(identify failures from the history). It’s possible to put all the requirements into the LLM planner, asking it to analyze the feedback and then do the planning, which means an evaluator isn’t necessary for the framework, especially with stronger LLMs. For current version GPT-4, we find the decomposition improving the reasoning results.

B. Fine-tuning a smaller model as planner

Fine-tuning a language model, rather than directly querying a GPT-4, not only enables offline deployment but also holds distinct advantages in the context of interactive planning. One prominent reason is the incorporation of multimodality in the data. Our system doesn’t solely rely on text descriptions but also utilizes the robot’s observations. While these observations can theoretically be converted into text form, they constitute a novel data type that GPT-4 has not been trained on, thereby resulting in limited zero-shot generalizability. For example, in experiments using GPT-4, if poses in robot observations and action parameters are in different frames of reference, the LLM will have trouble transforming them. A second reason is the requirement of large contexts in the input. A direct query to GPT would necessitate the inclusion of environment settings and generation constraints at each instance, which is inefficient and cost-intensive.

The difficulty of fine-tuning a smaller pre-trained LLM model mainly comes from two sides: 1) **Lack of data** for complex tasks. Most robotics data in the wild[41]–[43] has no partial observable tasks involved, and force-torque sensor data is usually not included since they are noisy and vary

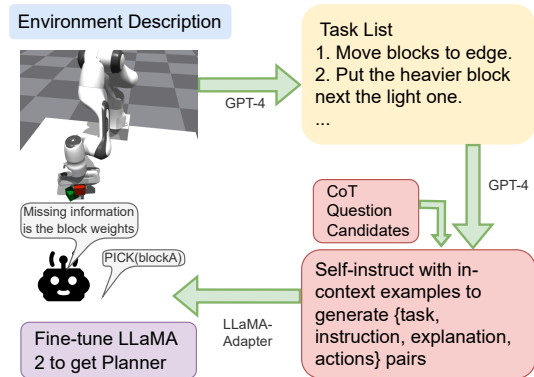


Fig. 4. The training procedure of the fine-tuned LLMs as an interactive planner as described in Sec IV-B. During inference, questions come from the pre-defined CoT question set, inputs come from robot observation.

across robots. 2) Smaller models are **worse at reasoning tasks**, CoT is fundamentally tied with larger models [44].

In order to get the required data to fine-tune a model as a planner in interactive planning under partial observation, we follow the procedure shown in Figure 4, using self-instruct[12] to generate an instruction dataset and fine-tune a LLaMA2-7B[29] model. The full pipeline includes:

Task Generation: The description of the environment, robot, potential uncertainties, action options, and example tasks are provided to GPT-4 to generate a number of tasks that are feasible to solve. We encourage GPT-4 to make the task set diverse in difficulty.

Instruction Generation: The generated tasks are used to generate pairs of instructions and responses, following the self-instruct paradigm. The instruction includes task descriptions and questions, the input encompasses the robot’s observations. The output generated by the model includes the same verbal explanations and actions as GPT-4 planners. We add format instructions to guarantee the “response” format.

CoT question designs: Finishing the state abstraction, belief update, and action planning in one query is hard for smaller models. Therefore, we create CoT questions[45] to ask *if missing information exists, how to collect information, and how to solve the task with full information*. The planner will choose questions to ask based on binary options in response.

Integrating collected robot observations: For the pre-trained actions, we collect success trajectories of the robot finish the actions and use them as in-context reference examples in the *Instruction Generation* process.

Fine-tuning: For the fine-tuning process, we adopted the LLaMA-adapter [39]. This approach allows us to enhance the model’s performance by leveraging a specifically curated dataset and fine-tuning it to our unique task generation and interactive planning scenario.

V. EXPERIMENTS

The experiments aim to validate the proposed interactive planner and to answer the following three questions:

- 1) Is the proposed framework able to solve complex tasks with uncertainties?
- 2) Does the proposed method apply to sim and real robots with different observation and action spaces?

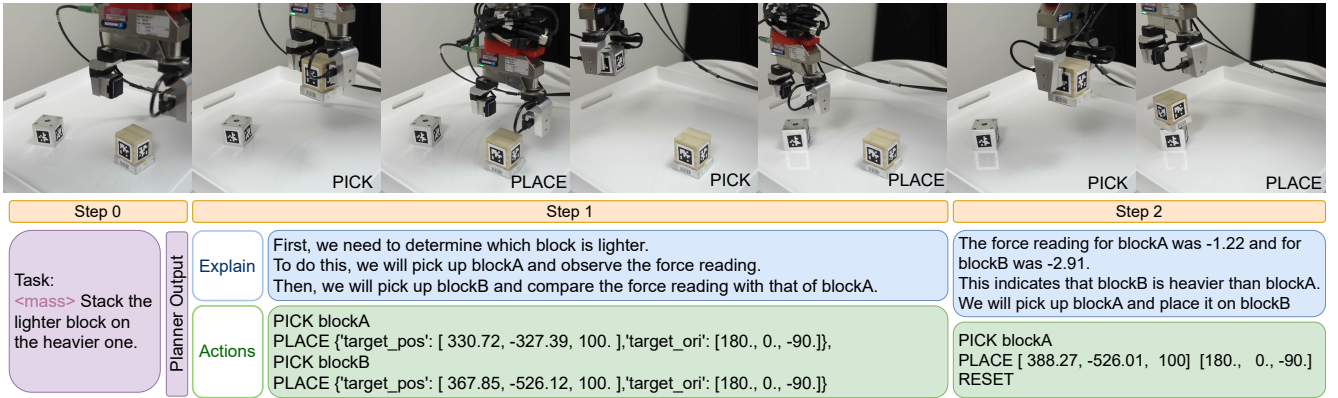


Fig. 5. An example rollout of LLM-POP solving *T4: Stack the lighter block on the heavier one*. In the first step, the planner figured out the plan to determine which block was lighter by picking up and placing both blocks down. In the second step, it figured out the heavier block. The Evaluator output is not shown here due to the space limit. In the next round(not shown in the figure), the evaluator recognized the completion of the task.

3) Can the fine-tuned LLM also solve partial observable tasks? What are the gaps between GPT-4?

A. Experimental setup

Environment: LLM-POP is evaluated on a set of manipulation tasks in a tabletop block rearrangement environment. A robot arm with a parallel gripper is equipped with a pre-trained skill set of {Pick, Place, Reach, Reset}. Each scenario is initialized with identical-size blocks with randomized positions and orientations on the table.

Uncertainties: We introduce two uncertainties in this environment 1) *mass*: Density(mass) of the blocks is randomized. 2) *fix*: blocks are randomized to be fixed/movable on the table. We design a task set containing tasks at different difficulty levels (horizon length to solve the task) under uncertainty assumptions to evaluate the planner’s performance. **Observations:** The robot observations include the pose of the robot end effector, the pose of the blocks on the table, gripper opening positions, and force-torque (F/T) readings.

Pre-trained skills and parameters: *pick(object)*: Pick up the specified *object* on the table. *place(pose)*: Move the end effector to desired *pose* and open the gripper. *reach(pose)*: Move the end effector to desired *pose*. *reset()*: Reset the arm and gripper to the initial pose. where *object* is a text name string, *pose* is the position and orientation.

Evaluation metrics: How to evaluate the task success rate is non-trivial since the desired outcome of the tasks in Table I (e.g., the lighter block is on top of the heavier block) could also be achieved through an incomplete decision-making process (e.g., stack a block on top of the other one that by chance respects the right weight relationship.) For this reason, even the LLM evaluator proposed in Section IV-A cannot confidently determine the success of the tasks, and we eventually relied on a manual check of each experiment.

For each task in the evaluation task set, we evaluate the success rate of finishing the task under ten random initializations on the positions and the uncertainties of blocks (5 for real robot settings since it’s harder to randomize the uncertainties). Table I includes our evaluation tasks. Default LLM-POP uses GPT-4 for the planner and evaluator.

TABLE I
UNCERTAINTY AND TASK DESCRIPTION TABLE

Type	Task Descriptions
/	1. Stack one block onto another.
/	2. Move the blocks to the corners of the table.
mass	3. Pick up the heavier block.
mass	4. Stack the lighter block on the heavier.
fix	5. Pick up the movable block and put it at the table corner.
fix	6. Find the movable block and put it on the fixed block.

B. Simulation: Block manipulation with Pre-trained LLM

We first evaluate our approach for solving the tasks in Table I in a simulated robotic system in IsaacGym[46]. We use the *FrankaCubeStack* task as a template environment, but change the blocks to the same size with random densities for *mass* uncertainty and randomly fix the block on the table for *fix* uncertainty. The action parameters for *place* include 3D target position and quaternion of the end effector. This is different from the block orientation quaternions in the observation, and explicitly including this (compared to using observation-action examples) in the prompt is essential for the LLM planner to get the correct action parameters and understand the observations.

We use two ablations: 1. GPT-3.5 as planner. 2. Remove the evaluator which explicitly asks for state abstraction and belief updates¹. Evaluation results are shown in Table II.

TABLE II
PLANNER SUCCESS RATE ON EVALUATION TASK SET

Model	T1	T2	T3	T4	T5	T6
GPT-3.5(w/o E)	5/10	7/10	0/10	0/10	0/10	0/10
GPT-3.5	6/10	6/10	0/10	0/10	0/10	0/10
LLM-POP(w/o E)	10/10	10/10	9/10	4/10	6/10	4/10
LLM-POP	10/10	10/10	10/10	8/10	7/10	8/10
LLM-POP*	5/5	5/5	5/5	4/5	5/5	4/5
FT-Vanilla	4/10	2/10	0/10	0/10	0/10	0/10
FT-CoT	10/10	10/10	4/10	3/10	8/10	6/10

(w/o E): No evaluator. * Real-world experiment. FT: Fine-tuned Llama2

Compared to the GPT-4-based planner, GPT-3.5 is also able to reason correct action sequences for stacking tasks, but can not always reason about the geometric (position and orientation) parameters. It can understand what’s the missing

¹To avoid the uncertainty from GPT versions, we use *gpt-4-0314* for all GPT-4 and *gpt-3.5-turbo-16k-0613* for GPT-3.5 usage.

information for tasks with uncertainty, but fails to generate multi-step plans to collect and update the information.

We observe a performance improvement for the GPT-4 based planner, especially on longer-horizon tasks with uncertainty when the evaluator is added to the pipeline. An example is shown in Figure 5. The LLM is asked explicitly what is missing and how such information can be analyzed from historical observations. This enforces the LLM to perform “state abstraction” and “belief update” before planning. In experiments, the GPT-4 planner without the evaluator sometimes makes the wrong plan by repeating the same collecting actions even if it already collected sufficient information. As the history of observations grows longer, the chance that LLM makes wrong reasoning also increases. This is partially due to the long-text handling challenge for the current GPT-4 version, and decomposing the reasoning tasks into an evaluator helps improve the stability.

C. Hardware: Real robot Block Manipulation with GPT

We implement the same version of our method on a MELFA Assista robot arm with a WSG 32 two-finger gripper. We put AprilTags [47] on the sides of the blocks to get the pose estimate of blocks. We use blocks with different materials and added weight for uncertainty in mass. We use the force-torque (F/T) sensor mounted on the robot’s wrist to get force readings. The action parameters for position controls on the real robot are in Euler angles, and the angles for the gripper and blocks are in different frames. For safety, we set the task to fail if action parameters get out of safety bounds or a collision happens. Results are in Table II.

The LLM-POP framework (GPT-4) achieves better performance in the real robot compared to the simulation domain. This is mostly because of the very accurate position controller implemented on the real robot, leading to fewer execution errors. The stiffness controller and F/T sensor used on the real robot allow us to recover accurate force readings with no noise compared to the “sensor” in the simulator, which is affected by robot movement and gravity. These experiments show that the proposed framework can solve tasks with varying levels of difficulties and uncertainties reliably in simulated as well as real systems.

D. Simulation: Block manipulation with fine-tuned model

Using the self-instruction method introduced in IV-B, we generate an Alpaca [30]-like dataset and use it to finetune a Llama2-7B [3] model for reasoning. We test it on the same IsaacGym environment. Results are also shown in Table II. *FT-Vanilla* uses directly generated instruction pairs, while *FT-CoT* uses CoT decomposition instruction pairs. The biggest challenge during data generation for both is the correspondence between imagined observation generated by GPT-4 and the ground truth. For example, after picking up the block, the gripper position in the generated data is sometimes not close to the block position (and thus incorrect). This increases the difficulty for fine-tuned models to do correct reasoning based on observations. In the experiments, the fine-tuned model is able to reason about the missing information

based on the task description and generate plans to collect the information. The results show that the fine-tuned model benefits from the CoT decomposition of instructions, and failures mostly come from wrong reasoning (wrong “heavy” block based on history). The current gap between the fine-tuned model and GPT-4 lies in the ability to analyze the historical information for updating the information and the ability to avoid and adjust wrong action parameters since it’s not included in the current data synthesis procedure. A potential improvement is to add an auxiliary task of “observation understanding” in training and use diverse environment settings to improve the reasoning capability. We leave this to our future research.

E. Common Failures in Sim and Real Experiments

1) *Execution failures*: This appears more in the simulation environment when the place action sets a target pose with less tolerance between objects and the robot moves at high speed (The control gains in the simulator are not fine-tuned for various block weights). LLM planners can also generate actions that cause collisions since there’s no online collision avoidance in skills.

To explicitly test if the evaluator can help in correcting execution errors, we did an ablation on stacking(*TI*) by adding offsets (1cm) on the grasping position of the block. The initial target placing position will fail because of this offset. With the evaluator included, which asks the LLM to analyze failure action based on history and propose correcting suggestions, the planner outputs a better target position (higher) in the next round. This shows that having an evaluator can actually help to correct execution errors. Detailed analysis is deferred to a longer draft of the paper.

2) *Belief update failures*: Incorrect plan for collecting information (e.g., trying to reach above the block to measure its weight); wrong analysis results from the observations (e.g., not comparing weight using the force sensor z-axis value but using noise in other axes). In the LLM-POP with GPT-4 case, most failures come from the wrong analysis.

VI. DISCUSSION

In this work, we proposed an interactive planning framework LLM-POP using LLM to solve tasks under partial observation. The framework is verified in simulated as well as real robot systems on various partially observable tasks. Task distribution that the current framework can solve strongly depends on the diversity and robustness of the pre-trained skills. Current skills are open-loop actions based on initial observation. If the pre-trained skills are closed-loop policies with collision avoidance and online adjustment, the framework would be able to solve more challenging tasks.

Overall, the gap between fine-tuned model and GPT-4 is clear, especially in reasoning for complex tasks. Our goal is not to replace the GPT-4 but to propose a method for generating self-instruct data for robotic tasks with limited demonstration data. We verify its usage as an interactive planner and leave the task of involving more modalities in the observation like image [48], [49] and solving complex environments like HomerRobot [50] to our future research.

REFERENCES

- [1] OpenAI, “Gpt-4 technical report,” 2023.
- [2] R. Anil, A. M. Dai, O. Firat, M. Johnson, D. Lepikhin, A. Passos, S. Shakeri, E. Taropa, P. Bailey, Z. Chen, E. Chu, J. H. Clark, L. E. Shafey, Y. Huang, K. Meier-Hellstern, G. Mishra, E. Moreira, M. Omernick, K. Robinson, S. Ruder, Y. Tay, K. Xiao, Y. Xu, Y. Zhang, G. H. Abrego, J. Ahn, J. Austin, P. Barham, J. Botha, J. Bradbury, S. Brahma, K. Brooks, M. Catasta, Y. Cheng, C. Cherry, C. A. Choquette-Choo, A. Chowdhery, C. Crepy, S. Dave, M. Dehghani, S. Dev, J. Devlin, M. Díaz, N. Du, E. Dyer, V. Feinberg, F. Feng, V. Fienber, M. Freitag, X. Garcia, S. Gehrmann, L. Gonzalez, G. Gur-Ari, S. Hand, H. Hashemi, L. Hou, J. Howland, A. Hu, J. Hui, J. Hurwitz, M. Isard, A. Ittycheriah, M. Jagielski, W. Jia, K. Kenealy, M. Krikun, S. Kudugunta, C. Lan, K. Lee, B. Lee, E. Li, M. Li, W. Li, Y. Li, J. Li, H. Lim, H. Lin, Z. Liu, F. Liu, M. Maggioni, A. Mahendru, J. Maynez, V. Misra, M. Moussalem, Z. Nado, J. Nham, E. Ni, A. Nystrom, A. Parrish, M. Pellat, M. Polacek, A. Polozov, R. Pope, S. Qiao, E. Reif, B. Richter, P. Riley, A. C. Ros, A. Roy, B. Saeta, R. Samuel, R. Shelby, A. Slone, D. Smilkov, D. R. So, D. Sohn, S. Tokumine, D. Valter, V. Vasudevan, K. Vodrahalli, X. Wang, P. Wang, Z. Wang, T. Wang, J. Wieting, Y. Wu, K. Xu, Y. Xu, L. Xue, P. Yin, J. Yu, Q. Zhang, S. Zheng, C. Zheng, W. Zhou, D. Zhou, S. Petrov, and Y. Wu, “Palm 2 technical report,” 2023.
- [3] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. C. Ferrer, M. Chen, G. Cucurull, D. Esiobu, J. Fernandes, J. Fu, W. Fu, B. Fuller, C. Gao, V. Goswami, N. Goyal, A. Hartshorn, S. Hosseini, R. Hou, H. Inan, M. Kardas, V. Kerkez, M. Khabsa, I. Kloumann, A. Korenev, P. S. Koura, M.-A. Lachaux, T. Lavril, J. Lee, D. Liskovich, Y. Lu, Y. Mao, X. Martinet, T. Mihaylov, P. Mishra, I. Molybog, Y. Nie, A. Poulton, J. Reizenstein, R. Rungta, K. Saladi, A. Schelten, R. Silva, E. M. Smith, R. Subramanian, X. E. Tan, B. Tang, R. Taylor, A. Williams, J. X. Kuan, P. Xu, Z. Yan, I. Zarov, Y. Zhang, A. Fan, M. Kambadur, S. Narang, A. Rodriguez, R. Stojnic, S. Edunov, and T. Scialom, “Llama 2: Open foundation and fine-tuned chat models,” 2023.
- [4] S. Vemprala, R. Bonatti, A. Bucker, and A. Kapoor, “Chatgpt for robotics: Design principles and model abilities,” Microsoft, Tech. Rep. MSR-TR-2023-8, 2023. [Online]. Available: <https://www.microsoft.com/en-us/research/publication/chatgpt-for-robotics-design-principles-and-model-abilities/>
- [5] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, A. Herzog, D. Ho, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, E. Jang, R. J. Ruano, K. Jeffrey, S. Jesmonth, N. J. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, K.-H. Lee, S. Levine, Y. Lu, L. Luu, C. Parada, P. Pastor, J. Quiambao, K. Rao, J. Rettinghouse, D. Reyes, P. Sermanet, N. Sievers, C. Tan, A. Toshev, V. Vanhoucke, F. Xia, T. Xiao, P. Xu, S. Xu, M. Yan, and A. Zeng, “Do as i can, not as i say: Grounding language in robotic affordances,” 2022.
- [6] W. Yu, N. Gileadi, C. Fu, S. Kirmani, K.-H. Lee, M. G. Arenas, H.-T. L. Chiang, T. Erez, L. Hasenclever, J. Humplik, B. Ichter, T. Xiao, P. Xu, A. Zeng, T. Zhang, N. Heess, D. Sadigh, J. Tan, Y. Tassa, and F. Xia, “Language to rewards for robotic skill synthesis,” 2023.
- [7] S. Li, X. Puig, C. Paxton, Y. Du, C. Wang, L. Fan, T. Chen, D.-A. Huang, E. Akyurek, A. Anandkumar, J. Andreas, I. Mordatch, A. Torralba, and Y. Zhu, “Pre-trained language models for interactive decision-making,” 2022.
- [8] 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,” 2023.
- [9] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, “Planning and acting in partially observable stochastic domains,” *Artificial Intelligence*, vol. 101, no. 1, pp. 99–134, 1998. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S000437029800023X>
- [10] W. Huang, F. Xia, T. Xiao, H. Chan, J. Liang, P. Florence, A. Zeng, J. Tompson, I. Mordatch, Y. Chebotar, P. Sermanet, N. Brown, T. Jackson, L. Luu, S. Levine, K. Hausman, and B. Ichter, “Inner monologue: Embodied reasoning through planning with language models,” 2022.
- [11] D. Driess, F. Xia, M. S. M. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu, W. Huang, Y. Chebotar, P. Sermanet, D. Duckworth, S. Levine, V. Vanhoucke, K. Hausman, M. Toussaint, K. Greff, A. Zeng, I. Mordatch, and P. Florence, “Palm-e: An embodied multimodal language model,” 2023.
- [12] Y. Wang, Y. Kordi, S. Mishra, A. Liu, N. A. Smith, D. Khashabi, and H. Hajishirzi, “Self-instruct: Aligning language model with self generated instructions,” 2022. [Online]. Available: <https://arxiv.org/abs/2212.10560>
- [13] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. J. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich, “Rt-1: Robotics transformer for real-world control at scale,” 2023.
- [14] G. Wang, Y. Xie, Y. Jiang, A. Mandlekar, C. Xiao, Y. Zhu, L. Fan, and A. Anandkumar, “Voyager: An open-ended embodied agent with large language models,” 2023.
- [15] P. Sharma, A. Torralba, and J. Andreas, “Skill induction and planning with latent language,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Dublin, Ireland: Association for Computational Linguistics, 2022, pp. 1713–1726. [Online]. Available: <https://aclanthology.org/2022.acl-long.120>
- [16] I. Singh, V. Blukis, A. Mousavian, A. Goyal, D. Xu, J. Tremblay, D. Fox, J. Thomason, and A. Garg, “Progprompt: Generating situated robot task plans using large language models,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 11 523–11 530.
- [17] Y. Ding, X. Zhang, C. Paxton, and S. Zhang, “Task and motion planning with large language models for object rearrangement,” 2023.
- [18] K. Lin, C. Agia, T. Migimatsu, M. Pavone, and J. Bohg, “Text2motion: From natural language instructions to feasible plans,” 2023.
- [19] Y. Chen, J. Arkin, Y. Zhang, N. Roy, and C. Fan, “Autotamp: Autoregressive task and motion planning with llms as translators and checkers,” 2023.
- [20] P. Jansen, “Visually-grounded planning without vision: Language models infer detailed plans from high-level instructions,” in *Findings of the Association for Computational Linguistics: EMNLP 2020*. Online: Association for Computational Linguistics, 2020, pp. 4412–4417. [Online]. Available: <https://aclanthology.org/2020.findings-emnlp.395>
- [21] W. Huang, C. Wang, R. Zhang, Y. Li, J. Wu, and L. Fei-Fei, “Voxposer: Composable 3d value maps for robotic manipulation with language models,” 2023.
- [22] Z. Mandi, S. Jain, and S. Song, “Roco: Dialectic multi-robot collaboration with large language models,” 2023.
- [23] Y. Cui, S. Karamcheti, R. Palleti, N. Shivakumar, P. Liang, and D. Sadigh, “No, to the right: Online language corrections for robotic manipulation via shared autonomy,” in *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI ’23. New York, NY, USA: Association for Computing Machinery, 2023, p. 93–101. [Online]. Available: <https://doi.org/10.1145/3568162.3578623>
- [24] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafraan, K. Narasimhan, and Y. Cao, “React: Synergizing reasoning and acting in language models,” *ArXiv preprint*, vol. abs/2210.03629, 2022. [Online]. Available: <https://arxiv.org/abs/2210.03629>
- [25] P. Sharma, B. Sundaralingam, V. Blukis, C. Paxton, T. Hermans, A. Torralba, J. Andreas, and D. Fox, “Correcting robot plans with natural language feedback,” 2022.
- [26] Z. Liu, A. Bahety, and S. Song, “Reflect: Summarizing robot experiences for failure explanation and correction,” *arXiv preprint arXiv:2306.15724*, 2023.
- [27] A. Z. Ren, A. Dixit, A. Bodrova, S. Singh, S. Tu, N. Brown, P. Xu, L. Takayama, F. Xia, J. Varley, Z. Xu, D. Sadigh, A. Zeng, and A. Majumdar, “Robots that ask for help: Uncertainty alignment for large language model planners,” 2023.
- [28] X. Chen, S. Zhang, P. Zhang, L. Zhao, and J. Chen, “Asking before action: Gather information in embodied decision making with language models,” 2023.
- [29] 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.

- [30] R. Taori, I. Gulrajani, T. Zhang, Y. Dubois, X. Li, C. Guestrin, P. Liang, and T. B. Hashimoto, "Stanford alpaca: An instruction-following llama model," https://github.com/tatsu-lab/stanford_alpaca, 2023.
- [31] S. G. Patil, T. Zhang, X. Wang, and J. E. Gonzalez, "Gorilla: Large language model connected with massive apis," 2023.
- [32] Y. Shen, K. Song, X. Tan, D. Li, W. Lu, and Y. Zhuang, "Hugginggpt: Solving ai tasks with chatgpt and its friends in hugging face," 2023.
- [33] Y. Qin, S. Liang, Y. Ye, K. Zhu, L. Yan, Y. Lu, Y. Lin, X. Cong, X. Tang, B. Qian, S. Zhao, R. Tian, R. Xie, J. Zhou, M. Gerstein, D. Li, Z. Liu, and M. Sun, "Toollm: Facilitating large language models to master 16000+ real-world apis," 2023.
- [34] Q. Tang, Z. Deng, H. Lin, X. Han, Q. Liang, B. Cao, and L. Sun, "Toolalpaca: Generalized tool learning for language models with 3000 simulated cases," 2023.
- [35] H. Yang, S. Yue, and Y. He, "Auto-gpt for online decision making: Benchmarks and additional opinions," 2023.
- [36] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa, "Large language models are zero-shot reasoners," 2023.
- [37] K. Shridhar, A. Stolfo, and M. Sachan, "Distilling reasoning capabilities into smaller language models," 2023.
- [38] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "Lora: Low-rank adaptation of large language models," in *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. [Online]. Available: <https://openreview.net/forum?id=nZeVKeeFYf9>
- [39] P. Gao, J. Han, R. Zhang, Z. Lin, S. Geng, A. Zhou, W. Zhang, P. Lu, C. He, X. Yue, H. Li, and Y. Qiao, "Llama-adapter v2: Parameter-efficient visual instruction model," 2023.
- [40] B. Rozière, J. Gehring, F. Gloeckle, S. Sootla, I. Gat, X. E. Tan, Y. Adi, J. Liu, T. Remez, J. Rapin, A. Kozhevnikov, I. Evtimov, J. Bitton, M. Bhatt, C. C. Ferrer, A. Grattafiori, W. Xiong, A. Défossez, J. Copet, F. Azhar, H. Touvron, L. Martin, N. Usunier, T. Scialom, and G. Synnaeve, "Code llama: Open foundation models for code," 2023.
- [41] H. Walke, K. Black, A. Lee, M. J. Kim, M. Du, C. Zheng, T. Zhao, P. Hansen-Estruch, Q. Vuong, A. He, V. Myers, K. Fang, C. Finn, and S. Levine, "Bridgedata v2: A dataset for robot learning at scale," 2023.
- [42] H. Bharadhwaj, J. Vakil, M. Sharma, A. Gupta, S. Tulsiani, and V. Kumar, "Roboagent: Towards sample efficient robot manipulation with semantic augmentations and action chunking," *arxiv*, 2023.
- [43] H.-S. Fang, H. Fang, Z. Tang, J. Liu, J. Wang, H. Zhu, and C. Lu, "Rh20t: A robotic dataset for learning diverse skills in one-shot," 2023.
- [44] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. Chi, Q. Le, and D. Zhou, "Chain-of-thought prompting elicits reasoning in large language models," 2023.
- [45] N. Ho, L. Schmid, and S.-Y. Yun, "Large language models are reasoning teachers," 2023.
- [46] V. Makovychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, and G. State, "Isaac gym: High performance gpu-based physics simulation for robot learning," 2021. [Online]. Available: <https://arxiv.org/abs/2108.10470>
- [47] E. Olson, "AprilTag: A robust and flexible visual fiducial system," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2011, pp. 3400–3407.
- [48] D. Zhu, J. Chen, X. Shen, X. Li, and M. Elhoseiny, "Minigt-4: Enhancing vision-language understanding with advanced large language models," *ArXiv preprint*, vol. abs/2304.10592, 2023. [Online]. Available: <https://arxiv.org/abs/2304.10592>
- [49] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," *ArXiv preprint*, vol. abs/2304.08485, 2023. [Online]. Available: <https://arxiv.org/abs/2304.08485>
- [50] S. Yenamandra, A. Ramachandran, K. Yadav, A. Wang, M. Khanna, T. Gervet, T.-Y. Yang, V. Jain, A. W. Clegg, J. Turner, Z. Kira, M. Savva, A. Chang, D. S. Chaplot, D. Batra, R. Mottaghi, Y. Bisk, and C. Paxton, "Homerobot: Open vocab mobile manipulation," 2023.