

# Large Language Model as a Policy Teacher for Training Reinforcement Learning Agents

Zihao Zhou, Bin Hu, Chenyang Zhao, Pu Zhang and Bin Liu\*

Zhejiang Lab

{zhouzihao,hubin,c.zhao,puz,liubin}@zhejianglab.com

## Abstract

Recent studies have uncovered the potential of Large Language Models (LLMs) in addressing complex sequential decision-making tasks through the provision of high-level instructions. However, LLM-based agents lack specialization in tackling specific target problems, particularly in real-time dynamic environments. Additionally, deploying an LLM-based agent in practical scenarios can be both costly and time-consuming. On the other hand, reinforcement learning (RL) approaches train agents that specialize in the target task but often suffer from low sampling efficiency and high exploration costs. In this paper, we introduce a novel framework that addresses these challenges by training a smaller, specialized student RL agent using instructions from an LLM-based teacher agent. By incorporating the guidance from the teacher agent, the student agent can distill the prior knowledge of the LLM into its own model. Consequently, the student agent can be trained with significantly less data. Moreover, through further training with environment feedback, the student agent surpasses the capabilities of its teacher for completing the target task. We conducted experiments on challenging MiniGrid and Habitat environments, specifically designed for embodied AI research, to evaluate the effectiveness of our framework. The results clearly demonstrate that our approach achieves superior performance compared to strong baseline methods. Our code is available at <https://github.com/ZJLAB-AMMI/LLM4Teach>.

## 1 Introduction

Large Language Models (LLMs) have revolutionized the field of artificial intelligence. These models are trained with an internet-scale text corpus, enabling them to exhibit remarkable capabilities such as natural language generation, question answering, and translation [Brown *et al.*, 2020; Du *et al.*, 2022; Chiang *et al.*, 2023]. Previous work suggests that these models contain vast general knowledge about the world

and are capable of solving complex reasoning problems [Radford *et al.*, 2019; Brown *et al.*, 2020; Wei *et al.*, 2022]. Recently, several works have attempted to use LLMs to generate action plans in an embodied environment [Ahn *et al.*, 2022; Wang *et al.*, 2023a; Driess *et al.*, 2023; Song *et al.*, 2023; Sha *et al.*, 2023; Mao *et al.*, 2023]. However, LLMs face challenges in generating effective end-to-end instructions for specific embodied tasks, especially in real-world dynamic scenarios. This limitation arises from two key factors. Firstly, LLMs do not possess the appropriate task incentives during the training process. Secondly, these models lack the capability to actively interact with the environment and gather real-time data [Carta *et al.*, 2023]. Furthermore, the utilization of LLMs often requires substantial computational resources, e.g., memory and power. These requirements render their deployment impractical and expensive, especially when considering their use on lightweight edge devices. These challenges motivate us to address the following question:

*How do we develop a lightweight, specialized agent that can quickly acquire the capabilities of LLMs for a specific sequential decision-making task?*

A commonly used solution is to train a specialized reinforcement learning (RL) based agent that starts learning from scratch. However, this approach often incurs a significant exploration cost, especially in high-dimensional and complex embodied environments with sparse reward signals, due to the low sampling efficiency of RL methods.

In this paper, we propose a novel approach called *LLM for policy teaching* (LLM4Teach), which utilizes a pre-trained LLM to expedite the training process of a small-scale RL-based student agent specialized for a target task. Specifically, in the early stage of training, the student agent queries the LLM-based teacher agent for action instructions and learns to mimic the behavior of its teacher through minimizing a distillation loss. As the learning process proceeds, the student gradually shifts from learning from its teacher to learning from the environment by upweighting a conventional RL loss. In another word, the objective function used for policy training is defined as a weighted average of the distillation loss and the RL loss. Since it allows the student agent to not only incorporate guidance from its LLM teacher but also learn from online interactions with the environment, LLM4Teach enables the student agent to identify and correct any mistakes made by its teacher, leading to improved performance on the

\*Corresponding author (Email: bins@ieee.org)

target task compared to its teacher. Note that only the student agent is deployed and it shall not interact with the LLM in the test phase. That means the model finally deployed is very lightweight compared to an LLM.

To summarize, our *main contributions* are:

- We propose LLM4Teach, a policy distillation approach to address the limitations of LLM and RL-based agents for embodied sequential decision making.
- We demonstrate the performance of our approach empirically by extensive experiments conducted on challenging embodied environments. In contrast to LLM-based agents, our approach shows improved accuracy and decreased computational workload. In comparison to RL-based agents, it has much greater sample efficiency.
- As a byproduct, we demonstrate that relying solely on LLM can result in various types of incorrect decisions in embodied settings, while LLM4Teach offers an effective approach to mitigate or avoid the influence caused by such incorrect decisions. We also verify that offering uncertainty-aware rather than deterministic guidance through LLM can improve the sample efficiency for the student agent.

## 2 Related Work

In this paper, we consider an algorithmic agent operating in an open dynamic environment. The agent is required to make a series of decisions and take actions based on the current state of the environment, employing a specific policy to successfully complete a designated task. Here we provide a brief overview of relevant research in the literature.

### 2.1 LLM-based Agents

LLMs have exhibited impressive reasoning abilities, motivating researchers to employ them as fundamental components for constructing LLM-based agents in diverse decision-making scenarios [Xi *et al.*, 2023; Yang *et al.*, 2023; Wang *et al.*, 2023b; Biggie *et al.*, 2023; Zhen *et al.*, 2023]. Recent research has demonstrated that LLMs can generate high-level plans in response to natural language descriptions of a given situation [Huang *et al.*, 2022; Shinn *et al.*, 2023; Yao *et al.*, 2022]. However, these plans may propose actions that are not compatible with the acting agent or the environment due to a lack of grounding in the specific problem domain. In addressing this issue, Ahn *et al.* [2022] proposed grounding LLMs through an affordance function of pre-trained skills, which assists LLMs in formulating feasible plans for execution by the agents. Additionally, Carta *et al.* [2023] proposed an approach in which the agent interacts with the environment and subsequently fine-tunes the LLMs using online collected data, thereby enhancing adaptation to the target task. However, frequent interaction with an LLM can be costly. Therefore, Hu *et al.* [2024] suggested an intelligent interaction approach that employs RL to determine when it is necessary to query the LLM, thus avoiding unnecessary interactions. Furthermore, Nottingham *et al.* [2023] optimized the selection of information presented to LLMs, thereby reducing the length of input contexts. While these methods reduce the cost of

utilizing LLMs for decision-making tasks, they all necessitate online access to a pre-trained LLM when the agent is deployed online during the testing phase.

In contrast, our approach involves utilizing the LLM solely during the training phase to distill task-specific knowledge from the LLM into a RL-based agent. Subsequently, during the testing phase, only the lightweight student agent is deployed, which works independently without dependence on the LLM.

### 2.2 LLM Assisted RL

Several studies have investigated the potential of utilizing LLMs to support the standard RL process by tapping into the general knowledge embedded in LLMs. For example, Kwon *et al.* [2023]; Yu *et al.* [2023] and Klissarov *et al.* [2023] employ LLMs to assist in assigning rewards. Kwon *et al.* [2023] use LLMs as proxy reward functions to automatically label trajectory data with rewards, while Yu *et al.* [2023] utilize LLMs to flexibly define reward parameters for optimizing and completing various robot tasks. In a different approach, Klissarov *et al.* [2023] leverage an offline dataset of behaviors and use LLMs’ preferences over pairs of randomly sampled trajectories to construct a reward model. Furthermore, Du *et al.* [2023] and Colas *et al.* [2023] focus on learning diverse behaviors without relying on reward supervision, employing LLMs to generate novel goals during exploration in the environment. In contrast to these previous works, our approach focuses on leveraging prior knowledge about the target task to enhance the initial exploration stage of an RL agent. This allows us to train the policy model with significantly less data, thereby improving the efficiency of the learning process.

### 2.3 Learning from Teacher Agents

Prior research has sought to improve the inefficiencies of tabula rasa RL by utilizing existing teacher agents to guide the learning process of a specialized RL agent for a specific target problem [Da Silva *et al.*, 2020; Agarwal *et al.*, 2022]. These instructions can manifest as demonstrations [Schaal, 1996], scalar feedback [Knox and Stone, 2009], or action advice [Da Silva *et al.*, 2017]. Jump-start RL involves the use of a teacher agent to assist in gathering high-quality data during the initial exploration phase of RL [Uchendu *et al.*, 2023]. Kickstarting RL combines on-policy distillation with RL, prompting the student agent to emulate the teacher’s behavior while optimizing for accumulated returns [Schmitt *et al.*, 2018]. Matthews *et al.* [2022] further extends this approach to hierarchical policies, transferring pre-trained low-level skill policies as teachers and training the student agent alongside a policy-over-teachers from scratch, which weighs the advice from each teacher agent at every time step.

In contrast, our approach does not rely on specialized teacher agents for the target problem. Instead, we harness the extensive general knowledge embedded in LLMs to expedite the learning process of the student agent through on-policy distillation for combining pre-trained fundamental skills to complete the target task.

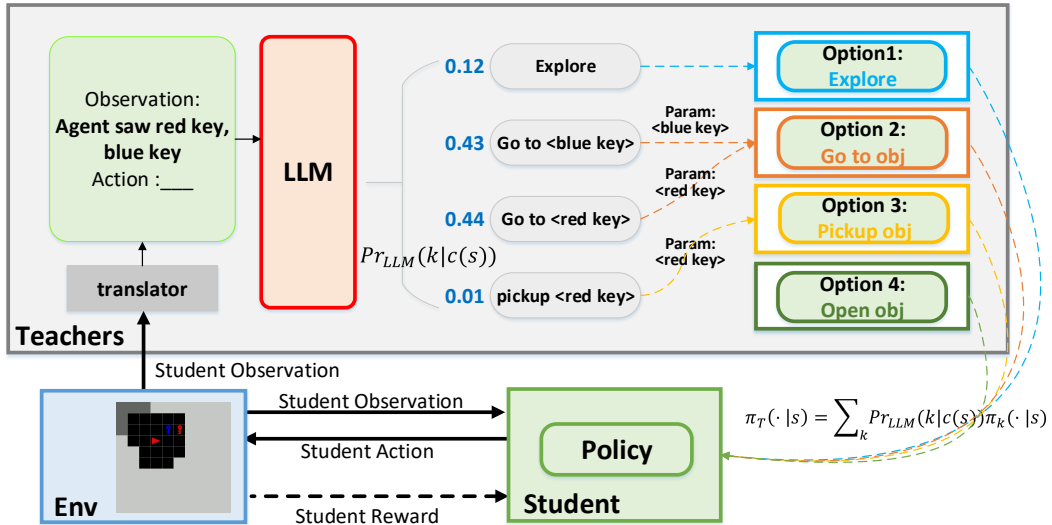


Figure 1: An illustration of our LLM4Teach framework using the MiniGrid environment as an exemplar. The LLM-based teacher agent responds to observations of the state provided by the environment by offering soft instructions. These instructions take the form of a distribution over a set of suggested actions. The student agent is trained to optimize two objectives simultaneously. The first one is to maximize the expected return, the same as in traditional RL algorithms. The other one is to encourage the student agent to follow the guidance provided by the teacher. As the student agent’s expertise increases during the training process, the weight assigned to the second objective gradually decreases over time, reducing its reliance on the teacher.

### 3 LLM4Teach

In this section, we present our methodology LLM4Teach. To begin with, we fix the notations as follows.

We consider a sequential decision-making problem formalized as a Markov Decision Process (MDP), denoted by  $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma \rangle$ , where  $\mathcal{S}$  and  $\mathcal{A}$  denote the state and action spaces, respectively. The transition probability function is denoted as  $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{P}(\mathcal{S})$ , and the reward function is denoted as  $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ . Additionally,  $\gamma$  represents the discount factor. The primary objective is to learn an optimal policy  $\pi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$ , which maximizes the expected cumulative return over time:  $\max_{\pi} \mathbb{E}[\sum_t \gamma^t r_t]$ . The parameter of the policy  $\pi$  is denoted as  $\theta$ . A standard gradient-based RL algorithm minimizes a surrogate loss,  $\mathcal{L}_{\text{RL}}(\theta)$ , using gradient descent with respect to  $\theta$ . This loss is estimated using sampled trajectories, where each trajectory consists of a sequence of tuples of state, action, and reward.

#### 3.1 The LLM4Teach Framework

The core principle of LLM4Teach involves the utilization of a pre-trained LLM as a teacher agent to guide a lightweight student RL agent in swiftly acquiring a policy for real-time decision-making to accomplish a specific embodied task. The student agent is allowed to interact with the environment and receive feedback from these interactions to rectify any errors provided by the teacher agent. Following the training phase, only the lightweight student agent is utilized during the testing phase, yet it owns superior capability in accomplishing the target task compared to its teacher. The conceptual framework of this approach is depicted in Figure 1.

#### 3.2 On the LLM-based Teacher Agent

In accordance with Ahn *et al.* [2022], we first notify the LLM of a set of  $K$  option policies  $\Pi : \{\pi_k : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})\}$  related to the current task using appropriate prompts, where  $k \in \{1, 2, \dots, K\}$  denotes the option index. When presented with a state  $s$ , the student agent requests guidance from the teacher agent for the next step action. The teacher agent initially selects a high-level option  $\pi_k$  from the set  $\Pi$ , prompted by a textual description  $c(s)$  of the state  $s$ . Subsequently, an action suggestion  $a \sim \pi_k(s)$  is generated based on the chosen option, serving as an instruction provided by the teacher.

#### Generating Uncertainty-aware Instructions Using LLM

The process of the student agent learning policies from the teacher agent can be seen as distilling important task-related knowledge from the LLM agent. As demonstrated in Hinton *et al.* [2015], incorporating uncertainty into knowledge distillation can improve sample efficiency and prevent model over-fitting. Consequently, we propose having the LLM offer uncertainty-aware soft instructions to the student agent. When the student agent sends a text description  $c(s)$  to the teacher agent, the teacher agent responds by providing a soft decision  $\pi_T(\cdot|s)$ , i.e., a distribution over available options, in the following way:

$$\pi_T(\cdot|s) = \sum_k Pr_{\text{LLM}}(k|c(s)) \pi_k(\cdot|s), \quad (1)$$

where  $Pr_{\text{LLM}}(k|c(s))$  represents the probability of the LLM teacher selecting the  $k$ th option given the textual description  $c(s)$  of the current state  $s$ , and  $\pi_k(\cdot|s)$  denotes the policy associated with the  $k$ th option. To estimate the uncertainties

---

**Algorithm 1** The student agent’s policy learning algorithm
 

---

**Require:** an LLM agent, pre-trained option policies  $\{\pi_k\}$ , initial policy parameter value  $\theta$ , maximum allowable number of iterations  $T$

```

1: for  $i = 1, 2, \dots, T$  do
2:   Collecting rollouts following the student agent’s initial
   policy and stores the data in a buffer  $\mathcal{D}$ 
3:   for each transition  $(s, a, r) \in \mathcal{D}$  do
4:     Generate a prompt with a textual description  $c(s)$  of
     the state  $s$  for the LLM-based teacher agent
5:     Get the soft decision of the LLM-based teacher
     agent according to Equation (1)
6:   end for
7:   for each gradient descent step do
8:      $\theta \leftarrow \theta - \alpha \nabla_{\theta} (\mathcal{L}_{\text{RL}}(\theta) + \lambda_i \mathbb{E}_s \mathcal{H}(\pi_T(\cdot|s) || \pi_{\theta}(\cdot|s)))$ 
9:   end for
10: end for
    
```

---

$Pr_{\text{LLM}}(k|c(s))$  in our experiments, we query the LLM multiple times with the same prompt to estimate the probability of each decision, similar to Wang *et al.* [2022]. An alternative approach is to access the logits of tokens relevant to option plans and convert them into probabilities [Carta *et al.*, 2023; Ahn *et al.*, 2022]. We conduct an ablation study on these two approaches in subsection 4.1.

### 3.3 On the Learning Process of the Student Agent

The learning process is summarized in Algorithm 1. The policy of the student agent, denoted as  $\pi_{\theta}(\cdot|s)$ , is learned by minimizing the following loss function:

$$\mathcal{L}(\theta) = \mathcal{L}_{\text{RL}}(\theta) + \lambda \mathbb{E}_{s \sim \pi_{\theta}} \mathcal{H}(\pi_T(\cdot|s) || \pi_{\theta}(\cdot|s)), \quad (2)$$

where  $\mathcal{L}_{\text{RL}}(\theta)$  denotes the traditional loss used in RL algorithms to encode the feedback from the environment. This loss is typically designed to maximize the expected return or rewards obtained by the agent. We incorporate the teacher agent’s guidance into the student agent’s learning process by introducing the regularization term  $\mathcal{H}(\pi_T(\cdot|s) || \pi_{\theta}(\cdot|s))$  that describes the difference between the teacher policy  $\pi_T(\cdot|s)$  and the student policy. This term captures the Kullback-Leibler (KL) divergence or Wasserstein distance between those two policies. To control the extent to which the student agent relies on the teacher agent, we introduce an annealing parameter  $\lambda$ . When  $\lambda$  is set to zero, the learning process of the student agent reduces to a standard RL process without any influence from the teacher agent.

We initialize the annealing parameter  $\lambda$  with larger values during the initial stages of training. This setup ensures that the student agent pays more attention to the guidance provided by the LLM-based teacher agent, aiming to align its policy with that of the teacher. As the training progresses, we gradually decay  $\lambda$ , allowing the student agent to shift its focus towards maximizing its expected return. By reducing the influence of the teacher’s guidance, the student agent becomes more independent in its decision-making process and emphasizes its own learned policy. Specially, the annealing

schedule used is designed as follow:

$$\lambda_i = \begin{cases} \lambda_0 - ki & \text{if } i < i_1 \\ \lambda_c & \text{if } i_1 < i < i_2 \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

where  $i$  represents the index of the training iteration,  $k$  represents the decay rate,  $\lambda_0$  is the initial value of  $\lambda$ ,  $\lambda_c$  is a constant value smaller than  $\lambda_0$ , which is maintained from the  $i_1$ th iteration to the  $i_2$ th iteration,  $i_2$  indicates the point at which the connection to the LLM-based teacher agent is closed. For more details on the annealing schedule used in our experiments, see Appendix A.7 of Zhou *et al.* [2023].

This linear reduction of  $\lambda$  enables a smooth transition for the student agent from heavily relying on the teacher’s guidance to prioritizing the RL objective. It provides a balance between learning from the teacher and acquiring autonomy in decision-making. When  $\lambda$  eventually reaches 0, we effectively remove the influence of the teacher’s instructions on the student agent’s policy, then the student agent no longer requires the teacher’s guidance.

## 4 Experiments

We validated the performance of our method, LLM4Teach, through extensive experiments. The aim of the experiments is to demonstrate the specific advantages of LLM4Teach compared to RL baseline methods and approaches that solely rely on LLM for decision-making, and to test its potential in handling real-world sequential decision making problems.

### Simulation Platforms

**MiniGrid** offers a customizable grid world environment with various sizes, object types, and objectives, making it a simple representation of grid-based tasks [Chevalier-Boisvert *et al.*, 2023]. These tasks pose a challenge for RL methods because of their sparse rewards.

**Habitat** is a simulation platform specifically created to support the development of embodied AI systems [Szot *et al.*, 2021]. It provides a comprehensive framework for defining and executing various embodied AI tasks, such as navigation, object rearrangement, and question-answering. Additionally, Habitat enables detailed configuration of embodied agents, including their physical attributes and sensor specifications.

### Baseline Methods

In the experiments, we include three baseline approaches to assess the performance of LLM4Teach.

**LLM soly** operates in two stages. First, the LLM selects an option from a set of available pre-trained options. Then, a low-level action is generated following the selected option policy. In this configuration, only an LLM-based agent is utilized to make real-time decisions, without the involvement of the student agent. This approach allows us to investigate the potential of our proposed LLM4Teach framework in enabling the student agent to outperform its teacher in completing the desired task.

**Hierarchical RL** is an RL baseline approach that involves training the student agent with pre-trained option policies [Matthews *et al.*, 2022]. In light of the hierarchical nature of the tasks, we explore such hierarchical RL approach in our experiments, so that we can assess the benefits of knowledge distillation from a pre-trained LLM that captures world knowledge.

**Baseline RL** is a Tabula rasa RL baseline trained from scratch using the proximal policy optimization (PPO) algorithm [Schulman *et al.*, 2017]. The policy model structure and the training loss function are set the same as our student agent in LLM4Teach.

## 4.1 Experiments on MiniGrid

### Experimental Setting

We created 4 procedurally generated tasks in the MiniGrid environment: {*SimpleDoorKey*, *ColoredDoorKey*, *LavaDoorKey* and *DivergedDoorKey*}. In each task, the agents are situated in rooms with varying layouts and their goal is to unlock the exit door using the correct key. In *SimpleDoorKey*, the agent must explore the room, find a key, and use it to unlock the exit door. In *ColoredDoorKey*, the exit door can only be unlocked with a key that matches its color, adding complexity for the agent to understand task-specific rules. *LavaDoorKey* introduces hazard grids (Lava) to the room, requiring the agent to quickly adapt to new elements. *DivergedDoorKey* presents two exit doors instead of one, allowing the agent to choose either door to escape, emphasizing the importance of using uncertainty-aware instructions to improve overall sample efficiency.

For every task, we incorporate 5 specialized options, which are: {*explore*, *go to*, *pickup*, *drop*, *open*}. All options, with the exception of *explore*, are dependent on specific conditions, such as interacting with an object, for example, *pickup the red key*. These expert policies are compiled under the fundamental task of *SimpleDoorKey*. Each option policy produces a Dirac delta distribution over actions based on the state. Additional information about the environments and options can be found in Appendix A.4 of Zhou *et al.* [2023].

We use ChatGLM-turbo [Du *et al.*, 2022] as the LLM to construct our teacher agent. This powerful model enables our teacher agent to possess complex reasoning capabilities. To leverage these capabilities, we employ Chain-of-thought (CoT) [Wei *et al.*, 2022] style prompts. The CoT prompts consist of multiple stages that guide the LLM’s decision-making process. Firstly, the LLM is prompted to summarize the scene, providing a condensed description of the environment. Secondly, it is instructed to reason about the appropriate course of action based on the given context. Finally, the LLM outputs its decision for the given task. To aid the LLM in understanding the reasoning process and ensuring correct output formatting, an arbitrary example is included in the prompt. This example serves as a reference point and helps the LLM grasp the desired output structure. Figure 2 illustrates an example of the dialogues generated by the LLM using this prompt setup in the *ColoredDoorKey* task.

### Prefix Prompt

**Instruction:**

In a locked 2D grid room, there is an agent whose task is to open the door. The door can only be opened with a key of the same color as the door. The agent can perform the following actions: explore, go to an object, pick up an object, drop an object, or open an object. Agent can only hold one object. Please clarify which object the agent has seen and is holding given the information provided. Then, please determine what the agent should do in its current state. Finally, please select the correct agent action. Do not provide reasoning.

**One-shot Example:**

User: Agent sees box, holds nothing

Assistant:

1. What agent sees and holds: Agent sees a box, holds nothing. Key and door are not observed yet.
2. What should agent do: Key maybe hidden in the box, agent may go to box and open it to check.
3. Action: {go to <box>, open <box>}

### Interaction

**Observation:** Agent sees red key, red door, holds blue key.

**Response:**

1. What agent sees and holds: Agent sees a red key and a red door, and holds a blue key.
2. What should the agent do: Agent should drop the blue key and pick up the red key to open the red door.
3. Actions: {drop <blue key>, go to <red key>, pick up <red key>}

Figure 2: An example of a prefix prompt and an interaction between the student agent and the LLM-based teacher agent for the task *ColoredDoorKey*. The Prefix prompt consists of two blocks: the instruction block briefly introduces the target problem and the CoT reasoning process; and the example block provides one arbitrary example of the expected format of the response from the LLM.

## Results on MiniGrid

The main results in Figure 3 show that the baseline RL struggles to complete tasks, even the simplest one (*SimpleDoorKey*), due to highly sparse rewards. In contrast, hierarchical RL eventually succeeds in the tasks but requires over 10,000 training iterations across all tasks. However, LLM4Teach, guided by the LLM-based teacher, effectively leverages the world knowledge embedded in the LLM, leading to significantly higher sample efficiency compared to prior art RL baselines with sparse rewards.

Results also show that LLM4Teach outperforms *LLM soly* in terms of accumulated returns for all tasks, except for *SimpleDoorKey*. *SimpleDoorKey* is the simplest one, with low reasoning difficulty for LLM. Moreover, all option policies are designed based on this environment, so there is no issue of option policy transfer. Therefore, *LLM soly* can achieve a success rate of nearly 100% for the task.

For the other tasks which are more complex than *SimpleDoorKey*, *LLM soly* performs unsatisfactorily due to the lack of enough task-grounding knowledge. In comparison, LLM4Teach allows the student agent to learn task-grounding knowledge from the environmental feedback, thus performs much better than *LLM soly*. For example, in *Colored-*

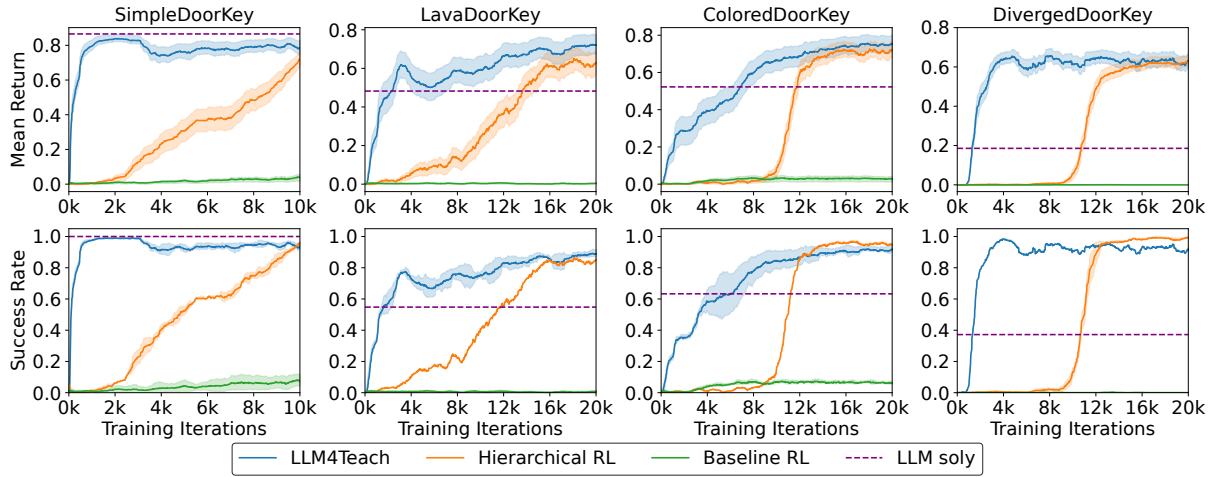


Figure 3: The tested average returns (top row) and task completion success rates (bottom row) vs. the training iteration index of the compared methods across four environments. The dotted vertical line indicates the point at which the teacher’s guidance is diminished, i.e., when  $\lambda_i = 0$ . *LLM soly* does not involve any learning, hence we report its average performance over 500 testing seeds, represented by a dashed horizontal line. For other approaches, we evaluate their policies every 10 iterations with 10 randomly generated testing seeds and report the averaged testing performance here. With our approach, the student agent effectively leverages the knowledge of the LLM-based teacher to bootstrap the early learning stage. Except for the *SimpleDoorKey* task, the student agent in LLM4Teach ultimately outperforms the LLM-based agent by learning from environment feedback through minimizing a traditional RL loss.

*DoorKey*, given the observation “Agent sees a red key, a blue key, a blue door.”, an LLM can suggest “pickup the red key”, while the right option is “pickup the blue key”, since only the the key with the same color of the door can be used to unlock the door. As a result, *LLM soly* only achieves an average return of 0.52. In contrast, utilizing the student agent within LLM4Teach leads to a significantly higher average return of 0.77, as illustrated in Figure 3. This is due to the student agent’s ability to rectify its teacher’s errors and adjust its behavior according to environmental feedback.

We have identified three major categories for the error policies generated by the LLM:

- **Incorrect policies**, which are executable but result in task failure. For example, an incorrect policy could involve moving into the lava, leading to the failure of task completion.
- **Inefficient policies**, which are executable but not necessary for task completion. They can increase the number of steps required to accomplish the task, potentially resulting in time-out errors. For instance, an inefficient policy could involve continuously exploring even after finding the correct key and door, instead of directly proceeding to the door.
- **Inconsistent policies**, which are not executable due to non-compliance with behavioral logic or contextual constraints, e.g., attempting to pick up a new key without first dropping the key that the agent is currently holding.

#### Ablation Study on Uncertainty-aware Instructions

As presented in subsection 3.2, the teacher agent in LLM4Teach offers uncertainty-aware instructions to the student agent, which is a distinguishing feature compared to previous LLM-based agents (e.g., in Ahn *et al.* [2022]), where

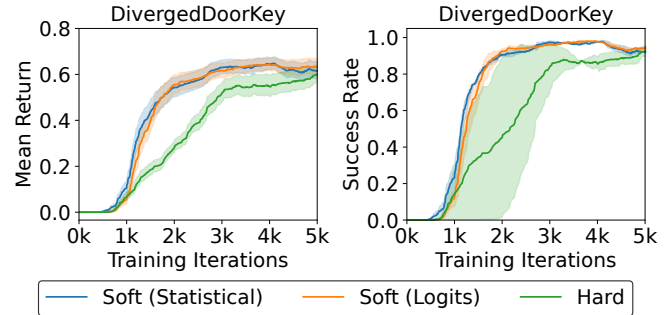


Figure 4: Ablation study on uncertainty-aware instructions. It shows that two types of uncertainty-aware instructions by the teacher both lead to improved sample efficiency for the student agent.

deterministic feedback is provided upon receiving a query.

We conducted ablation studies to investigate the benefits of using uncertainty-aware instructions instead of deterministic ones in the *DivergedDoorKey* task. We considered two approaches for the LLM to provide uncertainty-aware soft instructions. The first one is to query the LLM multiple times with the same prompt to **statistically** estimate the probability of each decision, similar to Wang *et al.* [2022]. The other approach is to access the **logits** of tokens relevant to option plans and convert them into probabilities [Carta *et al.*, 2023; Ahn *et al.*, 2022]. We compare these two approaches with a **hard** instruction baseline, where the LLM’s responses are directly used as deterministic instructions.

The result of the ablation study is shown in Figure 4. It is shown that utilizing uncertainty-aware instructions improves the overall sample efficiency compared to using deterministic ones. Moreover, there is no significant disparity



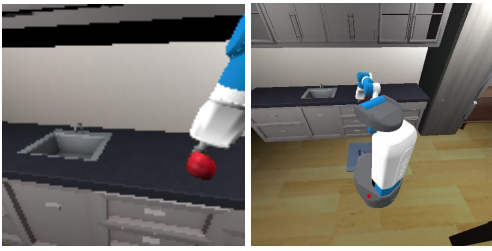


Figure 5: Habitat environment. Left: The visual observation from the onboard camera. Right: A view of the acting robot and its workspace from a third-party camera. Note that the third-party camera mentioned is purely for illustrative purposes and is not utilized during either the training or testing phases.

in performance between the two approaches for generating uncertainty-aware instructions. The first approach is simpler to implement in practical scenarios but consumes more computational resources due to multiple queries to LLMs, particularly when the observation space is large. The second approach necessitates access to logits, making it applicable only to open-source LLMs.

## 4.2 Experiments on Habitat

To evaluate the potential applicability of our method in real-world scenarios, we conducted additional experiments using Habitat [Szot *et al.*, 2021], which involves continuous action spaces and high-dimensional observations.

### Experimental Setting

We focus on a manipulation task called *Nav & Pick*. The objective of the robotic agent is to navigate to the table without any collisions and subsequently perform a precise object pickup. Refer to Figure 5 for a visual representation.

We conduct separate pre-training for two high-level options, namely *Navigate* and *Pick*. These options are utilized by both LLM4Teach and the hierarchical RL baseline [Matthews *et al.*, 2022]. To ensure the effectiveness of option training, we employ ten distinct training environment specifications, each with varying object and target locations. Furthermore, the agent’s initial positions are randomly generated upon environment reset, ensuring diverse training scenarios. For each option, we utilize a ResNet18 backbone in conjunction with a 2-layer MLP architecture to train the corresponding models. For more detailed information about the environments and training parameters, refer to Appendix A.5 of Zhou *et al.* [2023].

We select the Vicuna-7b model [Chiang *et al.*, 2023] as the LLM used in LLM4Teach, following a similar prompt design as in previous experiments on Minigrid. We utilize visual observations captured by the on-board camera as input queries for the LLM. To enable the LLM-based teacher agent to comprehend these visual inputs, we utilize a pre-configured translator in the same way as in Hu *et al.* [2024] to generate natural language descriptions which list identified objects in the visual inputs.

### Results on Habitat

Due to the task being limited to home scenarios, the LLM effectively covers the common-sense reasoning abilities re-

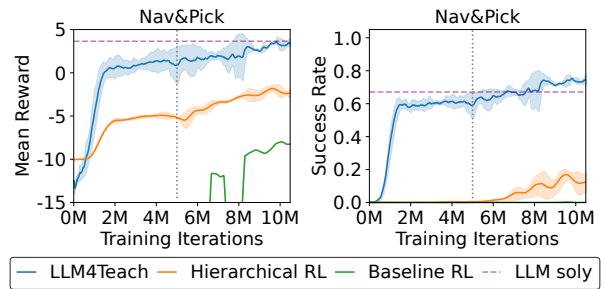


Figure 6: The tested average returns (left) and task completion success rates (right) vs. the training iteration index of the compared methods on the *Nav&Pick* task. For explanations of the lines and curves in the figure, see the caption of Figure 3.

quired to successfully complete the task. This results in few erroneous decision-making during option selection. Consequently, the task completion rate and average returns for *LLM soly*, as depicted in Figure 6, are relatively high. In contrast, the RL baselines struggle to complete the task due to the scarcity of rewards. Our approach, LLM4Teach, consistently outperforms all RL-based baselines in terms of both sample efficiency and asymptotic performance. This highlights the effective utilization of the LLM-based teacher’s knowledge by the student agent in LLM4Teach, facilitating the learning of appropriate policies. Given enough training iterations, our approach exhibits a higher success rate compared to *LLM soly*. The primary advantage of LLM4Teach is that it is an extremely lightweight RL-based student agent specifically designed for utilization in the final online testing phase.

## 5 Concluding Remarks

Both RL and LLMs have limitations in handling complex sequential decision-making problems. RL often lacks sample efficiency and incurs high exploration costs, while LLMs are prone to decision errors and have high deployment costs. Combining LLMs with RL to overcome these limitations is a natural idea, but creating an effective interface between them poses challenges. LLMs utilize texts as input and output, making them suitable for providing high-level instructions, whereas RL operates at a much lower fine-grained level and uses numerical vectors instead of texts as inputs and outputs.

Here we present LLM4Teach, a novel framework that combines LLMs and RL for embodied sequential decision-making tasks. Our approach leverages the reasoning capabilities of LLMs to develop a highly capable RL-based student agent. In particular, we use the LLM to provide high-level suggestions on available options for policy training of the student agent. Extensive experiments demonstrate that our student agent outperforms all RL baselines in sample efficiency. Meanwhile, it achieves superior performance to *LLM soly* in terms of task completion success rate with much fewer computational resources during online testing. For instance, in MiniGrid experiments, the student agent’s model size is 24K compared to LLM’s 130B. Similarly, in Habitat experiments, the student agent’s model size is 10M while an LLM’s is 7B.

## Ethical Statement

There are no ethical issues.

## Acknowledgments

This work was supported by Exploratory Research Project (No.2022RC0AN02) of Zhejiang Lab.

## Contribution Statement

Z. Zhou, B. Hu, C. Zhao and P. Zhang contribute equally. B. Liu is project lead.

## References

- [Agarwal *et al.*, 2022] Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Bellemare. Reincarnating reinforcement learning: Reusing prior computation to accelerate progress. *Advances in Neural Information Processing Systems*, 35:28955–28971, 2022.
- [Ahn *et al.*, 2022] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as I can, not as I say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- [Biggie *et al.*, 2023] Harel Biggie, Ajay Narasimha Mopidevi, Dusty Woods, and Christoffer Heckman. Tell me where to go: A composable framework for context-aware embodied robot navigation. *arXiv preprint arXiv:2306.09523*, 2023.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [Carta *et al.*, 2023] Thomas Carta, Clément Romac, Thomas Wolf, Sylvain Lamprier, Olivier Sigaud, and Pierre-Yves Oudeyer. Grounding large language models in interactive environments with online reinforcement learning. *arXiv preprint arXiv:2302.02662*, 2023.
- [Chevalier-Boisvert *et al.*, 2023] Maxime Chevalier-Boisvert, Bolun Dai, and al et. Minigrad & mineworld: Modular & customizable reinforcement learning environments for goal-oriented tasks. *CoRR*, abs/2306.13831, 2023.
- [Chiang *et al.*, 2023] Wei-Lin Chiang, Zhuohan Li, Zi Lin, and et. al. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality, March 2023.
- [Colas *et al.*, 2023] Cédric Colas, Laetitia Teodorescu, Pierre-Yves Oudeyer, Xingdi Yuan, and Marc-Alexandre Côté. Augmenting autotelic agents with large language models. *arXiv preprint arXiv:2305.12487*, 2023.
- [Da Silva *et al.*, 2017] Felipe Leno Da Silva, Ruben Glatt, and Anna Helena Reali Costa. Simultaneously learning and advising in multiagent reinforcement learning. In *Proceedings of the 16th conference on autonomous agents and multiagent systems*, pages 1100–1108, 2017.
- [Da Silva *et al.*, 2020] Felipe Leno Da Silva, Garrett Warnell, Anna Helena Reali Costa, and Peter Stone. Agents teaching agents: a survey on inter-agent transfer learning. *Autonomous Agents and Multi-Agent Systems*, 34:1–17, 2020.
- [Driess *et al.*, 2023] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023.
- [Du *et al.*, 2022] Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. Glm: General language model pretraining with autoregressive blank infilling. In *Proc. of the 60th Annual Meeting of the Association for Computational Linguistics*, pages 320–335, 2022.
- [Du *et al.*, 2023] Yuqing Du, Olivia Watkins, Zihan Wang, Cédric Colas, Trevor Darrell, Pieter Abbeel, Abhishek Gupta, and Jacob Andreas. Guiding pretraining in reinforcement learning with large language models. *arXiv preprint arXiv:2302.06692*, 2023.
- [Hinton *et al.*, 2015] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [Hu *et al.*, 2023] Bin Hu, Chenyang Zhao, Pu Zhang, Zihao Zhou, Yuanhang Yang, Zenglin Xu, and Bin Liu. Enabling efficient interaction between an agent and an llm: A reinforcement learning approach. *arXiv preprint arXiv:2306.03604*, 2023.
- [Hu *et al.*, 2024] Bin Hu, Chenyang Zhao, Pu Zhang, Zihao Zhou, Yuanhang Yang, Zenglin Xu, and Bin Liu. Enabling intelligent interactions between an agent and an llm: A reinforcement learning approach. *Reinforcement Learning Conference (RLC)*, 2024.
- [Huang *et al.*, 2022] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *International Conference on Machine Learning*, pages 9118–9147. PMLR, 2022.
- [Klissarov *et al.*, 2023] Martin Klissarov, Pierluca D’Oro, Shagun Sodhani, Roberta Raileanu, Pierre-Luc Bacon, Pascal Vincent, Amy Zhang, and Mikael Henaff. Motif: Intrinsic motivation from artificial intelligence feedback. *arXiv preprint arXiv:2310.00166*, 2023.
- [Knox and Stone, 2009] W Bradley Knox and Peter Stone. Interactively shaping agents via human reinforcement: The tamer framework. In *Proceedings of the 11th International conference on Knowledge capture*, pages 9–16, 2009.
- [Kwon *et al.*, 2023] Minae Kwon, Sang Michael Xie, Kalesha Bullard, and Dorsa Sadigh. Reward design with language models. *arXiv preprint arXiv:2303.00001*, 2023.



- [Mao *et al.*, 2023] Jiageng Mao, Yuxi Qian, Hang Zhao, and Yue Wang. Gpt-driver: Learning to drive with gpt. *arXiv preprint arXiv:2310.01415*, 2023.
- [Matthews *et al.*, 2022] Michael Matthews, Mikayel Samvelyan, Jack Parker-Holder, Edward Grefenstette, and Tim Rocktäschel. Hierarchical kickstarting for skill transfer in reinforcement learning. *arXiv preprint arXiv:2207.11584*, 2022.
- [Nottingham *et al.*, 2023] Kolby Nottingham, Yasaman Razeghi, Kyungmin Kim, JB Lanier, Pierre Baldi, Roy Fox, and Sameer Singh. Selective perception: Optimizing state descriptions with reinforcement learning for language model actors. *arXiv preprint arXiv:2307.11922*, 2023.
- [Radford *et al.*, 2019] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, and al et. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [Schaal, 1996] Stefan Schaal. Learning from demonstration. *Advances in neural information processing systems*, 9, 1996.
- [Schmitt *et al.*, 2018] Simon Schmitt, Jonathan J Hudson, Augustin Zidek, Simon Osindero, Carl Doersch, Wojciech M Czarnecki, Joel Z Leibo, Heinrich Kuttler, Andrew Zisserman, Karen Simonyan, et al. Kickstarting deep reinforcement learning. *arXiv preprint arXiv:1803.03835*, 2018.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Sha *et al.*, 2023] Hao Sha, Yao Mu, Yuxuan Jiang, Li Chen, Chenfeng Xu, Ping Luo, Shengbo Eben Li, Masayoshi Tomizuka, Wei Zhan, and Mingyu Ding. Languagempc: Large language models as decision makers for autonomous driving. *arXiv preprint arXiv:2310.03026*, 2023.
- [Shinn *et al.*, 2023] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik R Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [Song *et al.*, 2023] Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2998–3009, 2023.
- [Szot *et al.*, 2021] Andrew Szot, Alexander Clegg, and al et. Habitat 2.0: Training home assistants to rearrange their habitat. *Advances in Neural Information Processing Systems*, 34:251–266, 2021.
- [Uchendu *et al.*, 2023] Ikechukwu Uchendu, Ted Xiao, Yao Lu, Banghua Zhu, Mengyuan Yan, Joséphine Simon, Matthew Bennice, Chuyuan Fu, Cong Ma, Jiantao Jiao, et al. Jump-start reinforcement learning. In *International Conference on Machine Learning*, pages 34556–34583. PMLR, 2023.
- [Wang *et al.*, 2022] Xuezhi Wang, Jason Wei, and al et. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.
- [Wang *et al.*, 2023a] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023.
- [Wang *et al.*, 2023b] Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jikai Tang, Xu Chen, Yankai Lin, et al. A survey on large language model based autonomous agents. *arXiv preprint arXiv:2308.11432*, 2023.
- [Wei *et al.*, 2022] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*, 2022.
- [Xi *et al.*, 2023] Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. The rise and potential of large language model based agents: A survey. *arXiv preprint arXiv:2309.07864*, 2023.
- [Yang *et al.*, 2023] Sherry Yang, Ofir Nachum, Yilun Du, Jason Wei, Pieter Abbeel, and Dale Schuurmans. Foundation models for decision making: Problems, methods, and opportunities. *arXiv preprint arXiv:2303.04129*, 2023.
- [Yao *et al.*, 2022] Shunyu Yao, Jeffrey Zhao, Dian Yu, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *NeurIPS 2022 Foundation Models for Decision Making Workshop*, 2022.
- [Yu *et al.*, 2023] Wenhao Yu, Nimrod Gileadi, Chuyuan Fu, Sean Kirmani, Kuang-Huei Lee, Montse Gonzalez Arenas, Hao-Tien Lewis Chiang, Tom Erez, Leonard Hasenclever, Jan Humplik, et al. Language to rewards for robotic skill synthesis. *arXiv preprint arXiv:2306.08647*, 2023.
- [Zhen *et al.*, 2023] Yue Zhen, Sheng Bi, Lu Xing-tong, Pan Wei-qin, Shi Hai-peng, Chen Zi-rui, and Fang Yi-shu. Robot task planning based on large language model representing knowledge with directed graph structures. *arXiv preprint arXiv:2306.05171*, 2023.
- [Zhou *et al.*, 2023] Zihao Zhou, Bin Hu, Pu Zhang, Chenyang Zhao, and Bin Liu. Large language model as a policy teacher for training reinforcement learning agents. *arXiv preprint arXiv:2311.13373*, 2023.