

STeCa: Step-level Trajectory Calibration for LLM Agent Learning

Hanlin Wang, Jian Wang[†], Chak Tou Leong, Wenjie Li
Department of Computing, The Hong Kong Polytechnic University
{hanlin-henry.wang, chak-tou.leong}@connect.polyu.hk
jian51.wang@polyu.edu.hk cswjli@comp.polyu.edu.hk

Abstract

Large language model (LLM)-based agents have shown promise in tackling complex tasks by interacting dynamically with the environment. Existing work primarily focuses on behavior cloning from expert demonstrations or preference learning through exploratory trajectory sampling. However, these methods often struggle to address long-horizon tasks, where suboptimal actions accumulate step by step, causing agents to deviate from correct task trajectories. To address this, we highlight the importance of *timely calibration* and the need to automatically construct calibration trajectories for training agents. We propose **Step-Level Trajectory Calibration (STeCa)**, a novel framework for LLM agent learning. Specifically, STeCa identifies suboptimal actions through a step-level reward comparison during exploration. It constructs calibrated trajectories using LLM-driven reflection, enabling agents to learn from improved decision-making processes. We finally leverage these calibrated trajectories with successful trajectories for reinforced training. Extensive experiments demonstrate that STeCa significantly outperforms existing methods. Further analysis highlights that timely calibration enables agents to complete tasks with greater robustness. Our code and data are available at <https://github.com/WangHanLinHenry/STeCa>.

1 Introduction

Large language models (LLMs) have recently demonstrated impressive capabilities in reasoning and planning, leading to the development of various LLM-based agents that tackle real-world tasks such as household assistance (Puig et al., 2018; Shridhar et al., 2020b), web browsing (Yao et al., 2022; Deng et al., 2023), and complex scientific reasoning (Wang et al., 2022a). These agents typically engage in long-horizon interactions with en-




Figure 1: Step-level calibration enables LLM agents to construct calibrated trajectories and learn to mitigate the accumulation of suboptimal actions.

vironments, making sequential decisions to accomplish specific goals. Despite their promise, LLM agents still face significant challenges in generating high-quality task plans for complex scenarios (Xie et al., 2024; Wang et al., 2024a). This highlights the need for more effective approaches to enhance their decision-making capabilities over time.

Previous work has investigated agent learning by leveraging augmented exploratory data (Chen et al., 2023; Yin et al., 2023; Zeng et al., 2023; Xiang et al., 2024). These methods primarily rely on behavior cloning from expert demonstrations, training agents exclusively on successful trajectories. However, this approach prevents agents from proactively self-correcting mistakes, leading to the accumulation of errors and ultimately suboptimal task performance (Xie et al., 2024). To address this limitation, another line of work focuses on preference learning (Song et al., 2024; Xiong et al., 2024) and reinforcement learning (Carta et al., 2023; Tan et al., 2024), integrating failure trajectories additionally to refine decision-making. These approaches train LLM agents using explicit error signals or reward functions. However, many long-horizon agentic tasks involve multi-turn interactions, where errors often only become evident at the terminal state (Yuan et al., 2025). As a result, these methods fail to address early-stage deviations, which may not be immediately apparent but accumulate incrementally over time, ultimately leading to significant errors.

[†]Corresponding author.

To address the above limitations, we highlight the importance of **timely calibration**, which allows agents to *adjust suboptimal actions as they occur, rather than deferring correction until after the entire exploration*. As illustrated in Figure 1, when an early suboptimal action occurs, the subsequent actions are prone to deviate from the optimal trajectory, significantly increasing the risk of task failure. If an agent can engage in self-reflection and calibrate its behavior in real time, it stands a much better chance of successfully completing the task. However, implementing step-level calibrations presents significant challenges. 1) Unlike in mathematical reasoning tasks (Kumar et al., 2024; Xi et al., 2025), where well-defined rules simplify error detection, identifying deviations at each step in long-horizon agentic tasks is considerably more complex. This complexity stems from the dynamic and diverse nature of task execution in interactive environments. 2) As far as we know, the lack of step-level calibration trajectory data poses a major obstacle to training agents to effectively recognize and mitigate deviations.

In this work, we propose **Step-level Trajectory Calibration (STeCa)**, a novel agent learning framework that enables LLM agents to perform real-time calibration. STeCa operates by interacting with the environment to perform explorations and utilizes Monte Carlo (MC) sampling (Kakade and Langford, 2002) to estimate step reward for each action. By comparing the rewards of adjacent actions, STeCa effectively identifies deviated actions that lead to suboptimal performance. Then, we utilize off-the-shelf LLMs for reflection, which revises a deviated action into its ground-truth counterpart while generating a reflective thought. The resulting action and its thought, along with subsequent expert trajectory, form a *calibrated trajectory*. All calibrated trajectories, combined with successful trajectories during exploration, are then used to reinforce the agent’s training, optimizing its learning process. We evaluate STeCa on two widely-used agent benchmarks (Puig et al., 2018; Shridhar et al., 2020b). Extensive experimental results demonstrate that STeCa significantly outperforms existing methods, achieving higher success rates across a variety of tasks.

In summary, our contributions are as follows:

- We highlight the importance of timely calibration in interactive agentic tasks, a crucial aspect largely overlooked by previous methods. Unlike existing approaches that rely on terminal-state

error signals or reward functions, we emphasize the need for real-time adjustments to prevent the accumulation of deviations, which can lead to significant errors in long-horizon tasks.

- We introduce STeCa, a novel learning framework that enhances LLM agents by integrating an automated deviation detection mechanism and calibrated trajectory construction. It equips agents with essential calibration capabilities for improvement during task execution.
- Extensive experiments demonstrate that STeCa significantly outperforms existing methods. By detecting deviations in real-time, STeCa enables agents to effectively mitigate the accumulation of suboptimal actions and handle long-horizon tasks more robustly.

2 Preliminaries

Task Formulation. This work investigates how LLM-based agents tackle long-horizon tasks within specific environments through interactions. Following previous studies (Song et al., 2024; Xiong et al., 2024), we formalize these agentic tasks as a partially observable Markov decision process (POMDP), which contains the key elements $(\mathcal{U}, \mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{T}, \mathcal{R})$. Here, \mathcal{U} denotes the instruction space, \mathcal{S} the state space, \mathcal{A} the action space, \mathcal{O} the observation space, \mathcal{T} the transition function ($\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$), and \mathcal{R} the reward function ($\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$). Since the task planning capability of LLM agents is our main focus, $\mathcal{U}, \mathcal{A}, \mathcal{O}$ are subsets of natural language space.

Given a task instruction $u \in \mathcal{U}$, the LLM agent π_θ at time step t takes an action $a_t \sim \pi_\theta(\cdot | u, e_{t-1})$ and receives the environmental feedback as the observation $o_t \in \mathcal{O}$. e_{t-1} denotes the historical interaction trajectory $(a_1, o_1, \dots, a_{t-1}, o_{t-1})$. Each action a_t incurs the environment state to $s_t \in \mathcal{S}$. The interaction loop terminates when either the agent completes the task or the maximum step is reached. The final trajectory is $e_m = (u, a_1, o_1, \dots, a_m, o_m)$, where m denotes the trajectory length. The outcome reward $r_o(u, e_m) \in [0, 1]$ indicates the success or failure of the task.

Step-level Reward Acquisition. It is crucial to acquire step-level rewards as feedback to improve decision-making for LLM agents. Following prior work (Kakade and Langford, 2002; Salimans and Chen, 2018; Xiong et al., 2024), we leverage expert trajectories as demonstrations and ask an LLM

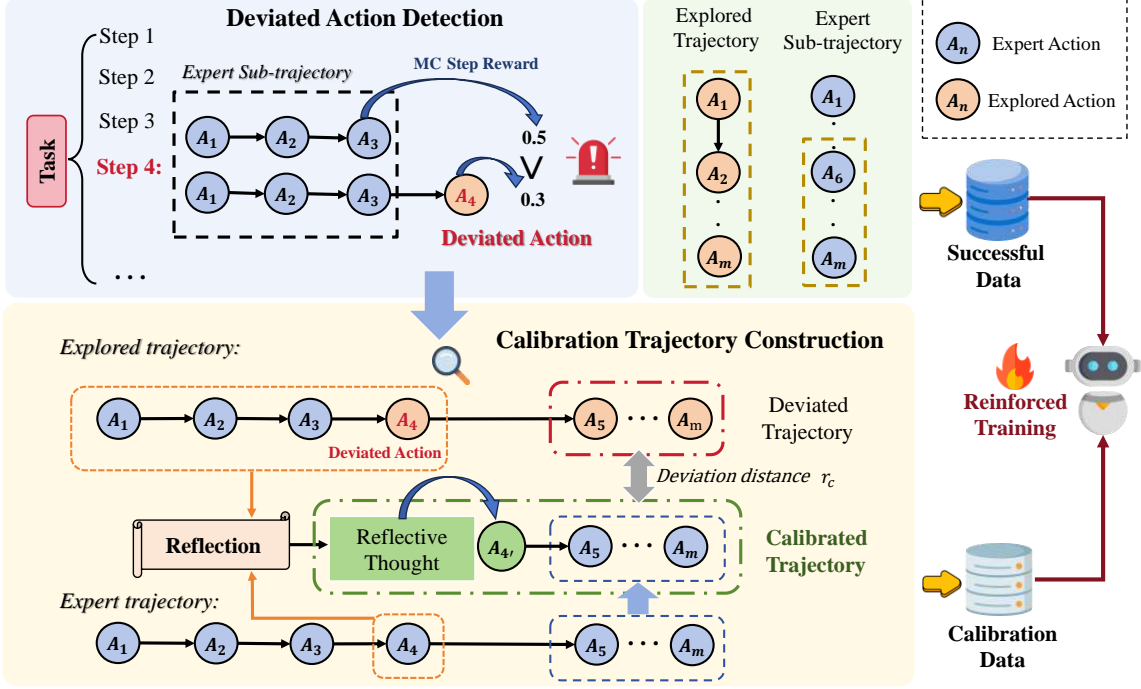


Figure 2: Overview of the Step-level Trajectory Calibration (STeCa) framework for LLM agent learning.

agent to begin exploration from the specific state $s_{t-1} \in \mathcal{S}$ toward the target state for a given demonstration. At each t -step, the agent’s policy π_θ generates an action a_t , and we define a step-level reward $r_{step}(s_{t-1}, a_t)$ to quantify the contribution of a_t to future success. Specifically, at t -th step, the agent generates N new subsequent trajectories $\{e_{t+1:m}^{(i)}\}_{i=1}^N$ using the widely-used Monte Carlo sampling, conditioned on the historical trajectory e_t . Each trajectory receives an outcome reward $r_o(u, e_m)$ from the environment. The step-level reward $r_{step}(s_{t-1}, a_t)$ is computed as the expected value of these outcome rewards:

$$r_{step}(s_{t-1}, a_t) = \mathbb{E}_{e_m \sim \pi_\theta(e_{t+1:m}|e_t)}[r_o(u, e_m)]. \quad (1)$$

Normalized Dynamic Time Warping. The normalized Dynamic Time Warping (nDTW) algorithm (Müller, 2007), implemented via dynamic programming (DP), effectively measures the distance between two trajectories containing multiple time steps. Formally, given a pair of trajectories (x, y) , this computation process is computed as:

$$D(i, j) = d(x_i, y_j) + \min \begin{cases} D(i-1, j) \\ D(i, j-1) \\ D(i-1, j-1) \end{cases} \quad (2)$$

where $d(x_i, y_j)$ denotes a distance function such as L_2 or cosine distance, $D(0, 0) = d(x_0, y_0)$. x_i denotes the action at the i -th step in the trajectory x , while y_j denotes the action at the j -th step in the trajectory y . With a normalization operation, the nDTW distance d_{nDTW} is given by:

$$d_{nDTW}(x, y) = \frac{D(x-1, y-1)}{\sqrt{n_x^2 + n_y^2}}, \quad (3)$$

where $d_{nDTW} \in [0, 1]$, n_x and n_y denote the number of steps in the trajectory x and y , respectively.

3 Method

In this section, we present **Step-level Trajectory Calibration (STeCa)**, a novel learning framework for LLM agents. First, we warm up agent training with supervised fine-tuning (§3.1), equipping LLM agents with necessary task planning capabilities. Then, we focus on calibration trajectory construction (§3.2), which detects deviated actions for an explored trajectory through step-level reward comparison and calibrates them by reflection. Finally, we utilize these calibrated trajectories as a crucial part of data for reinforced training (§3.3). Figure 2 illustrates the overview of STeCa.

3.1 Warm-up via Supervised Fine-tuning

Supervised fine-tuning (SFT) on the expert trajectory data has demonstrated promising results, serv-

ing as an effective initial step for developing strong agents. We employ ReAct-style (Yao et al., 2023) trajectory to conduct SFT, which additionally generates a Chain-of-Thought (CoT) (Wei et al., 2022) rationale before each action. Considering that the CoT and the corresponding action are generated together, we represent both as a single unit, denoted as a_t , for simplicity. Given an expert trajectory dataset $\mathcal{D} = \left\{ (u, e)^{(i)} \right\}_{i=1}^{|\mathcal{D}|}$, where each trajectory $e = (u, a_1, o_1, \dots, a_m, o_m)$, u represents the initial task instruction, a_t denotes the action (including its rationale) at step t , o_t is the corresponding observation, and $|\mathcal{D}|$ is the number of trajectories, the SFT loss function is formulated as:

$$\mathcal{L}_{\text{SFT}}(\theta) = -\mathbb{E}_{e \sim \mathcal{D}} \left[\sum_{t=1}^n \log \pi_{\theta}(a_t | e_{t-1}) \right]. \quad (4)$$

This warm-up process equips the LLM agent with the necessary task-planning capabilities, enabling it to generate both rationales and actions, resulting in a base agent π_{base} .

3.2 Calibration Trajectory Construction

To construct the calibration trajectories, we utilize the base agent π_{base} to explore the environment through interaction. During this exploration, sub-optimal actions often lead to a cascade of further suboptimal decisions, causing the trajectory to deviate from successful task completion. We define these actions, which are likely to cause deviations from the optimal trajectory and increase the risk of task failure, as *deviated actions*. Below, we introduce the details of detecting deviated actions and constructing calibrated trajectories accordingly.

Deviated Action Detection via Step-level Reward Comparison. Since long-horizon tasks can be modeled as a partially observable Markov decision process (POMDP), where the future action in a task execution process depends on the current action, we must consider this Markov property when detecting deviated actions. To illustrate this, we define the probability of an agent successfully completing a task based on a ‘‘good’’ historical trajectory (e.g., an expert trajectory) at time step t as $p(a_{>t} | a_{\leq t})$, where we omit the environmental states for simplicity. After executing a subsequent action a_{t+1} , the probability of task completion becomes $p(a_{>t+1} | a_{\leq t+1})$. If a_{t+1} is a ‘‘good’’ action (e.g., an expert action), $p(a_{>t+1} | a_{\leq t+1})$ will generally be greater than $p(a_{>t} | a_{\leq t})$. This is because

agentic tasks typically consist of sequential actions, where each action contributes to task completion as the sequence progresses. Thus, by comparing the task completion probabilities before and after executing an action, we can determine whether the action is deviated.

Specifically, we employ step-level rewards, calculated via Monte Carlo (MC) sampling introduced in §2, as an approximate estimation of the task completion probabilities. The base agent would conduct an explored action \hat{a}_{t+1} based on the expert sub-trajectory $e_{1:t}$. The explored action \hat{a}_{t+1} is classified as a deviated action if its step reward is significantly lower than that of the previous expert action a_t by a predefined threshold δ ; otherwise, it is considered a non-deviated action. The formal detection criterion is defined as follows:

$$\begin{cases} \text{Deviated Action:} \\ r_{\text{step}}(s_t, \hat{a}_{t+1}) - r_{\text{step}}(s_{t-1}, a_t) < \delta, \\ \text{Non-deviated Action:} \\ r_{\text{step}}(s_t, \hat{a}_{t+1}) - r_{\text{step}}(s_{t-1}, a_t) \geq \delta, \end{cases} \quad (5)$$

where $r_{\text{step}}(s_{t-1}, a_t)$ represents the step reward for the expert action a_t at the t -th step, $r_{\text{step}}(s_t, \hat{a}_{t+1})$ denotes the step reward for the explored action \hat{a}_{t+1} , and $\delta \geq 0$ is a threshold parameter. Step rewards r_{step} constructed by MC sampling are only utilized to detect the deviated actions.

Calibrated Trajectory Collection with Reflective Thoughts. As shown in Figure 2, after identifying a deviated action in an explored trajectory, our goal is to enable the LLM agent to ‘‘know’’ that the action is deviated and learn how to realign with the task objective. Achieving this goal requires calibrated trajectories for training the agent. Inspired by many previous studies on LLM reflections (Shinn et al., 2023), we employ off-the-shelf LLMs to generate reflective thoughts for calibration. Formally, we concatenate the preceding expert sub-trajectory $e_{1:t-1}$, the deviated action \hat{a}_t , and the corresponding ground-truth action a_t in the expert trajectory, and prompt a state-of-the-art LLM (e.g., GPT-4o (OpenAI, 2024)) for reflection, transforming the deviated action \hat{a}_t into the ground-truth action along with its reflective thought, which is denoted as a'_t . This formulates the subsequent *calibrated trajectory* $e_{c(t:m)} = (a'_t, e_{t+1:m})$, where $e_{t+1:m}$ represents the expert sub-trajectory from the step $t + 1$ to the end step m . The detailed prompt for this reflection is provided in Appendix E.2.

Our calibration dataset \mathcal{D}_c is constructed as:

$$\mathcal{D}_c = \{e_{c(t:m)}^{(j)}\} \cup \{e_{d(1:m)}^{(j)}\}, \quad (6)$$

where $e_{d(1:m)} = (e_{1:t-1}, \hat{a}_t, \hat{e}_{t+1:m})$ denotes a deviated trajectory, which will be used in subsequent reinforced training. Note that we perform trajectory calibration immediately when detecting the first deviated action, rather than waiting until the trajectory concludes. This approach ensures timely calibration and reduces unnecessary exploration.

3.3 Reinforced Training

While training on calibration trajectories enhances an agent’s calibration capability, relying exclusively on these trajectories may hinder their ability to recognize correctness. To mitigate this, we introduce two types of successful data during exploration. First, we construct the *explored successful trajectory* dataset, \mathcal{D}_e , by collecting successful trajectories that the base agent independently explores from the beginning, along with their corresponding expert trajectories. Note that these explored successful trajectories are not completely the same as the expert trajectories, because they include trial-and-error actions during explorations. Second, we build the *expert sub-trajectory* dataset, \mathcal{D}_s . Specifically, for a failed trajectory, where the first erroneous action occurs at step t , we extract the corresponding expert action and the subsequent trajectory, following Xiong et al. (2024). These sub-trajectories guide the agent in learning from challenging cases more effectively.

Using the collected data, we perform reinforced training to enhance LLM agents. Our goal is to guide the agent toward generating optimal trajectories that maximize task performance while minimizing suboptimal outcomes. We introduce *trajectory deviation distance* (TDD), a measure that quantifies how much a suboptimal trajectory deviates from an optimal one at the trajectory level. Drawing inspiration from Xu et al. (2024), we utilize the nDTW distance d_{nDTW} (as detailed in §2), to quantify the deviation distance between a suboptimal trajectory e_s and its corresponding optimal trajectory e_o . A smaller $d_{\text{nDTW}}(e_s, e_o)$ indicates a lower deviation. This deviation distance will be utilized as a trajectory-level reward signal in reinforced training.

To ensure balanced training across the datasets, we refine the reward mechanism by incorporating the trajectory deviation distance. The reward func-

tions for each type of data are defined as follows:

$$r_c = 1 + \eta \cdot d_{\text{nDTW}}(e_{c(t:m)}, e_{d(t:m)}), \quad (7)$$

$$r_s = 1 + \eta \cdot d_{\text{nDTW}}(e_{t:m}, \hat{e}_{t:m}), \quad (8)$$

$$r_e = 1 - \eta \cdot d_{\text{nDTW}}(\tilde{e}_{1:m}, e_{1:m}), \quad (9)$$

where for the calibration trajectory $e_{c(t:m)}$ and the expert sub-trajectory $e_{t:m}$, we increase the reward as the deviation distance grows, encouraging the agent to calibrate larger deviations. For the explored successful trajectory $\tilde{e}_{1:m}$, we reduce the reward for unnecessary explorations when the deviation distance increases, discouraging deviations from optimal behavior. η is a temperature coefficient that controls the impact of deviation distance on the reward. Finally, we integrate these rewards into reinforcement training using the policy gradient (Peters and Schaal, 2007) algorithm. The overall training objective is given by:

$$\begin{aligned} \mathcal{L}(\theta) = & \mathbb{E}_{(e_{c(t:m)}, e_{1:t-1}) \sim \mathcal{D}_c} \left[r_c \cdot \log \pi_{\theta}(e_{c(t:m)} \mid e_{1:t-1}) \right] \\ & + \mathbb{E}_{(e_{t:m}, e_{1:t-1}) \sim \mathcal{D}_s} \left[r_s \cdot \log \pi_{\theta}(e_{t:m} \mid e_{1:t-1}) \right] \\ & + \mathbb{E}_{(\tilde{e}_{1:m}, u) \sim \mathcal{D}_e} \left[r_e \cdot \log \pi_{\theta}(\tilde{e}_{1:m} \mid u) \right]. \end{aligned} \quad (10)$$

4 Experiments

4.1 Experimental Settings

Datasets. We conduct experiments on two representative agentic task datasets: **VirtualHome** (Puig et al., 2018) and **ALFWorld** (Shridhar et al., 2020b). For ALFWorld, we utilize datasets constructed by Song et al., 2024. For the VirtualHome dataset, we leverage the predefined tasks from the ActivityPrograms knowledge base (Puig et al., 2018) and construct a corresponding dataset in a manner closely aligned with the ALFWorld dataset. Please refer to Appendix A for further details regarding the dataset construction process and associated statistical information.

Baseline Methods. We evaluate STeCa against the following two categories of baselines: (1) **prompting-based** approaches, including GPT-3.5-turbo (Ouyang et al., 2022) and GPT-4 (Achiam et al., 2023). (2) **tuning-based** methods, which include supervised fine-tuning (SFT) methods, such as pure SFT (Chen et al., 2023), RFT (Yuan et al., 2023), and E²CL (Wang et al., 2024a), reinforcement learning-based methods such as PPO (Schulman et al., 2017) and Step-PPO (Wang et al.,

Paradigm	Method	VirtualHome		ALFWorld		Average
		Seen	Unseen	Seen	Unseen	
Prompting-based	GPT-3.5-Turbo (Ouyang et al., 2022)	6.3	2.6	7.9	10.5	6.8
	GPT-4 (Achiam et al., 2023)	34.2	9.4	42.9	38.1	31.2
Tuning-based	Llama-2-7B-Chat + PPO (Schulman et al., 2017)	23.9	25.0	22.1	29.1	25.0
	Llama-2-7B-Chat + SFT (Chen et al., 2023)	64.9	57.7	60.0	67.2	63.3
	Llama-2-7B-Chat + RFT (Yuan et al., 2023)	65.1	58.3	62.9	66.4	63.2
	Llama-2-7B-Chat + Step-PPO (Wang et al., 2024b)	65.7	59.6	65.7	69.4	65.1
	Llama-2-7B-Chat + ETO (Song et al., 2024)	66.6	60.1	68.6	72.4	66.9
	Llama-2-7B-Chat + E ² CL (Wang et al., 2024a)	67.1	61.8	70.1	73.9	68.2
	Llama-2-7B-Chat + IPR (Xiong et al., 2024)	67.6	61.9	70.3	74.7	68.6
	Llama-2-7B-Chat + STeCa (Ours)	69.6	63.6	74.3	76.1	70.9
	Llama-2-7B-Chat + STeCa w/ SFT+DPO	66.8	63.5	74.1	75.5	70.0
Llama-2-7B-Chat + STeCa w/o RT	68.8	62.4	72.1	74.9	69.6	

Table 1: Performance of different methods on two agent datasets. “Seen” refers to the held-out test set containing tasks present during training, while “Unseen” refers to the test set with unseen task variations.

2024b), as well as preference learning methods like ETO (Song et al., 2024) and IPR (Xiong et al., 2024). Additional details about the baselines are provided in Appendix B.

Implementation Details. We utilize Llama-2-7B-Chat (Touvron et al., 2023) as the base model for training LLM agents. We set $\delta = 0$ as the threshold to detect deviated actions in two environments. We use $\eta = 1$ for VirtualHome and $\eta = 0.01$ for ALFWorld to weight the contribution of trajectory deviation to the reward. To obtain step-level rewards with MC sampling, we set the temperature to 1 and the number of samples N to 5. More details are presented in Appendix C.

Evaluation Metrics. Following existing studies (Song et al., 2024; Xiong et al., 2024), we adopt the **Average Final Reward** as our evaluation metric. This metric measures the success rate of test tasks. In ALFWorld, the environment provides binary final rewards, where a reward of 1 indicates task completion and 0 indicates failure. Similarly, in VirtualHome, a trajectory is deemed successful if the final environment state aligns with a predefined target state and yields a reward of 1; otherwise, the reward is 0. This metric ensures a consistent measure of task performance across both environments.

4.2 Main Results

Table 1 summarizes the performance of various methods on long-horizon tasks in the VirtualHome and ALFWorld environments. The proposed method, STeCa, achieves the highest overall performance, with an average reward of 70.9, signifi-

cantly outperforming the baseline methods. Compared to prompting-based methods, which exhibit relatively poor performance, STeCa demonstrates a significant improvement. These results highlight the inherent limitations of closed-source LLMs relying solely on prompt engineering. As a tuning-based method, STeCa demonstrates consistent superiority over prior approaches. Specifically, it achieves an average final reward of 70.9, surpassing IPR, the previous state-of-the-art method with an average reward of 68.6, by 3.4%. This improvement highlights the effectiveness of trajectory calibration in enhancing generalization and overall performance. Moreover, STeCa outperforms E²CL, a method that incorporates self-reflection mechanisms, by 4.0%. Notably, STeCa achieves this without requiring additional iterative training, underscoring its superior training efficiency.

To further validate our method, we conducted an ablation study by designing two variants of the training process. In the first variant, we applied supervised fine-tuning (SFT) followed by Direct Preference Optimization (DPO) (Rafailov et al., 2024) on both optimal and suboptimal trajectories in the collected dataset (denoted as w/ SFT + DPO). In the second variant, we performed SFT on the collected dataset but restricted the training to only the optimal trajectories, omitting reward tuning (denoted as w/o RT). The experimental results showed that employing both SFT and DPO led to a slight decrease in the average reward to 70.0, while SFT without reward tuning (w/o RT) resulted in a further drop to 69.6. Although both variants demonstrated competitive performance, neither surpassed the re-

Base Model	Method	Seen	Unseen
Mistral-7B	SFT	67.1	58.2
	IPR	71.4	73.9
	STeCa	73.3	75.3
Llama-3-8B-Instruct	SFT	68.6	62.7
	IPR	72.3	75.8
	STeCa	74.9	77.0

Table 2: Performance of SFT, IPR, and our STeCa with different base models on the ALFWorld dataset.

sults achieved by STeCa, thereby underscoring the effectiveness of the mechanisms proposed in our method.

5 Discussions and Analyses

In this section, we provide detailed analyses of the STeCa framework from the following aspects.

5.1 Effectiveness with Different Base Models

To validate the broad effectiveness of our method, we evaluate STeCa with different base models, including Mistral-7B and Llama-3-8B-Instruct, on the ALFWorld environment. We compare its performance with SFT and IPR across both seen and unseen tasks. As shown in Table 2, STeCa consistently outperforms both SFT and IPR across multiple base models. Notably, in unseen tasks, STeCa achieves a 17.1% improvement over SFT on Mistral-7B, highlighting its generalization ability for developing LLM-based agents. Furthermore, with Llama-3-8B-Instruct, a stronger backbone model, STeCa further achieves better performance, underscoring its potential for building advanced agents in the future.

5.2 Comparisons between Variants of STeCa

Variants of Step-Level Reward Acquisition. To investigate the impact of different methods for reward acquisition, we conducted experiments using GPT-4o and a trained reward model to annotate reward for each step action while keeping all other processes unchanged. The detailed process of step action reward acquisition is described in Appendix D.1. As shown in Table 3, our method utilizing MC sampling for step reward acquisition achieves superior performance compared to alternative variants, highlighting the effectiveness of MC sampling for reward acquisition. Notably, employing GPT-4o to directly annotate rewards for step actions demonstrates performance comparable

Variants	VirtualHome		ALFWorld	
	Seen	Unseen	Seen	Unseen
Step-level Reward Acquisition				
MC Sampling	69.6	63.6	74.3	76.1
GPT-4o Annotation	69.1	62.5	74.1	74.9
RM Prediction	68.2	61.8	74.0	73.3
Reflective Thought Generation				
GPT-4o Generation	69.6	63.6	74.3	76.1
Self-generation	66.0	61.1	71.4	73.3

Table 3: Comparisons between variants of STeCa.




Figure 3: Variations in Monte Carlo (MC) step rewards with respect to the number of remaining steps until task completion for expert trajectories.

to our method, suggesting that step rewards can be effectively obtained through more computationally efficient approaches. This finding provides a promising direction for future research into optimizing the efficiency of reward acquisition.

Variants of Reflective Thought Generation. In STeCa, we employ GPT-4o to generate reflection thoughts for constructing calibration trajectories. To evaluate the impact of reflection quality on performance, we attempted to prompt the base agent π_{base} to generate reflections while keeping all other processes unchanged. The results, summarized in Table 3, reveal a significant performance degradation, underscoring the critical importance of high-quality reflection generation for achieving optimal results. This finding also suggests that the base agent, trained solely on expert trajectories, lacks effective reflection capabilities.

5.3 Analyses of Deviated Action Detection

To validate the empirical Markov property introduced in Section 3.2, i.e., non-deviated “good” actions increase the likelihood of task completion, we conducted a statistical analysis using expert trajectories. Specifically, we compared task completion probabilities at varying distances from task completion, employing MC step rewards as a proxy for

these probabilities. As illustrated in Figure 3, MC step rewards monotonically increase as the agent progresses toward task completion in both environments. This trend demonstrates that the accumulation of optimal actions significantly contributes to task completion. Conversely, deviated actions consistently reduce the task completion probability, further supporting our approach of using step reward comparisons between adjacent steps as a reliable criterion for detecting deviated actions.

5.4 Analyses of Calibration

In this section, we compare the calibration capabilities of STeCa and baseline methods. We evaluate calibration performance using the average final reward achieved upon successful task completion in the presence of deviated actions. To enable this analysis, we constructed datasets containing historical trajectories with deviated actions, categorized into seen and unseen scenarios across both environments. Additional details are provided in Appendix D.2. As shown in Figure 4, STeCa outperforms baseline methods by a significant margin. For instance, it achieves a 14.8% relative improvement over IPR on unseen tasks in the VirtualHome environment, highlighting its superior calibration performance. To comprehensively illustrate the performance of different agent learning methods, we provide some concrete examples in Appendix F.

To examine the impact of deviated actions on LLM agent performance, we evaluate these agents trained with different methods under two settings: one with deviated actions included in historical trajectories and another without. The construction process of the testing datasets is detailed in Appendix D.2. As shown in Figure 5, STeCa exhibits minimal performance variation between the two settings, unlike other methods, which show significant differences. This suggests that the presence of deviated actions has little effect on STeCa’s performance, as it effectively calibrates subsequent trajectories. These findings highlight STeCa’s robustness and its potential for reliable performance in diverse and challenging environments.

6 Related Work

LLM Agent Learning. LLM agents are widely used for tackling complex real-world tasks (Wang et al., 2023a; Hu et al., 2024; Wang et al., 2024a), relying on iterative interactions with their environment guided by task objectives and constraints.




Figure 4: Calibration performance of different methods on the VirtualHome and ALFWorld datasets.




Figure 5: Correlation between the deviation distance and success rate (measured by average final reward).

However, in long-horizon planning, excessive interactions make them prone to suboptimal actions, increasing the risk of failure. While closed-source LLMs demonstrate strong intelligence, open-source counterparts still lag behind (Liu et al., 2023; Wang et al., 2023b). To address this gap, some studies focus on improving task success rates by increasing the likelihood of generating optimal actions (Chen et al., 2023; Yuan et al., 2023). Alternatively, another line of research seeks to mitigate suboptimal actions by collecting them and applying preference learning methods to reduce their occurrence (Song et al., 2024; Xiong et al., 2024). Recently, researchers have explored the capacity of LLM agents to self-correct errors, enhancing their ability to ensure successful task completion (Wang et al., 2024a; Qu et al., 2024). However, these methods primarily focus on self-correction after errors have already occurred, lacking the ability to detect suboptimal actions in advance and calibrate subsequent planning accordingly.

Process Supervision. Process supervision provides fine-grained guidance, making it a promising approach for addressing long-horizon problems (Uesato et al., 2022). Early studies have explored obtaining step-level rewards and using them

to optimize intermediate processes through reinforcement learning (Lightman et al., 2023; Deng et al., 2024; Wang et al., 2024b). Others have focused on constructing step-level positive and negative data pairs and applying preference learning techniques to achieve more precise optimization (Xiong et al., 2024; Jiao et al., 2024). However, existing studies have yet to address the construction of step-level reflection data. Such data could empower LLM agents to detect suboptimal actions, analyze the reasons for their suboptimality, and determine how to calibrate them to ensure successful task completion.

7 Conclusion

In this paper, we introduce STeCa, a novel agent learning framework designed to enhance the performance of LLM agents in long-horizon tasks. STeCa identifies deviated actions through step-level reward comparisons and constructs calibration trajectories via reflection. These trajectories serve as critical data for reinforced training. Extensive experiments demonstrate that STeCa significantly outperforms baseline methods, with additional analyses underscoring its robust calibration capabilities.

Limitations

While our approach demonstrates superior performance compared to baseline methods, it is important to acknowledge the limitations of our current work as follows:

(1) Computational Inefficiency: Although the Monte Carlo (MC) sampling approach in STeCa achieves superior performance in constructing step rewards compared to alternative methods, it requires a substantial number of sampling iterations, resulting in significant computational overhead. This inefficiency represents a notable limitation of our current implementation. Future work should focus on developing more efficient methods for constructing step rewards while preserving the performance advantages of our approach.

(2) Limited Utilization of Step Rewards: While our approach leverages step rewards to identify and evaluate deviated actions effectively, it does not fully exploit the potential of step rewards for broader decision-making or optimization tasks. This constrained utilization may limit the overall performance improvements that could be achieved by incorporating step rewards into other aspects of the framework. Future research should explore

strategies to better harness the rich information embedded in step rewards to enhance the overall effectiveness and adaptability of the system.

(3) Handling Multiple Deviated Actions: Our method primarily focuses on “timely” calibration, wherein the LLM agent identifies a deviated action and immediately adjusts its subsequent trajectory to prevent further deviations. This approach effectively mitigates the accumulation of errors over time. Nonetheless, our current framework does not explicitly address multi-step calibration for multiple deviated actions. Systematically handling such cases presents an opportunity for future work, which could further enhance the robustness of our method in more complex scenarios.

Ethics Statement

This work aims to develop LLM agents within simulated environments. The VirtualHome and ALFWorld environment setup and related data strictly follow the specifications of VirtualHome (Puig et al., 2018) and ALFWorld (Shridhar et al., 2020b). We utilize VirtualHome v2.3.0¹ (MIT license²) and ALFWorld³ (MIT license⁴) to conduct our experiments. All the LLMs we use for fine-tuning are open-source, and we strictly follow the protocols for the academic use of these models. Additionally, we acknowledge the use of AI assistants, including GitHub Copilot and ChatGPT, in supporting our coding and writing processes.

Acknowledgements

This work was supported by the Research Grants Council of Hong Kong (15209724). The authors would like to thank the anonymous reviewers for their valuable feedback and constructive suggestions.

References

Marwa Abdulhai, Isadora White, Charlie Snell, Charles Sun, Joey Hong, Yuexiang Zhai, Kelvin Xu, and Sergey Levine. 2023. Lmrl gym: Benchmarks for multi-turn reinforcement learning with language models. *arXiv preprint arXiv:2311.18232*.

¹<https://github.com/xavierpuigf/virtualhome/tree/master>

²<https://github.com/xavierpuigf/virtualhome/blob/master/LICENSE>

³<https://github.com/alfworld/alfworld>

⁴<https://github.com/alfworld/alfworld/blob/master/LICENSE>

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Thomas Carta, Clément Romac, Thomas Wolf, Sylvain Lamprier, Olivier Sigaud, and Pierre-Yves Oudeyer. 2023. Grounding large language models in interactive environments with online reinforcement learning. In *International Conference on Machine Learning*, pages 3676–3713. PMLR.
- Baian Chen, Chang Shu, Ehsan Shareghi, Nigel Collier, Karthik Narasimhan, and Shunyu Yao. 2023. Fireact: Toward language agent fine-tuning. *arXiv preprint arXiv:2310.05915*.
- Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su. 2023. Mind2web: Towards a generalist agent for the web. *Advances in Neural Information Processing Systems*, 36:28091–28114.
- Zhirui Deng, Zhicheng Dou, Yutao Zhu, Ji-Rong Wen, Ruibin Xiong, Mang Wang, and Weipeng Chen. 2024. From novice to expert: Llm agent policy optimization via step-wise reinforcement learning. *arXiv preprint arXiv:2411.03817*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Siyuan Hu, Mingyu Ouyang, Difei Gao, and Mike Zheng Shou. 2024. The dawn of gui agent: A preliminary case study with claude 3.5 computer use. *arXiv preprint arXiv:2411.10323*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Fangkai Jiao, Chengwei Qin, Zhengyuan Liu, Nancy F Chen, and Shafiq Joty. 2024. Learning planning-based reasoning by trajectories collection and process reward synthesizing. *arXiv preprint arXiv:2402.00658*.
- Sham Kakade and John Langford. 2002. Approximately optimal approximate reinforcement learning. In *Proceedings of the Nineteenth International Conference on Machine Learning*, pages 267–274.
- Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, et al. 2024. Training language models to self-correct via reinforcement learning. *arXiv preprint arXiv:2409.12917*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. *arXiv preprint arXiv:2305.20050*.
- Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, et al. 2023. Agentbench: Evaluating llms as agents. *arXiv preprint arXiv:2308.03688*.
- Chang Ma, Junlei Zhang, Zhihao Zhu, Cheng Yang, Yujie Yang, Yaohui Jin, Zhenzhong Lan, Lingpeng Kong, and Junxian He. 2024. Agentboard: An analytical evaluation board of multi-turn llm agents. *arXiv preprint arXiv:2401.13178*.
- Meinard Müller. 2007. Dynamic time warping. *Information retrieval for music and motion*, pages 69–84.
- OpenAI. 2024. Hello GPT-4o. <https://openai.com/index/hello-gpt-4o/>.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Jan Peters and Stefan Schaal. 2007. Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the 24th international conference on Machine learning*, pages 745–750.
- Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio Torralba. 2018. Virtualhome: Simulating household activities via programs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8494–8502.
- Yuxiao Qu, Tianjun Zhang, Naman Garg, and Aviral Kumar. 2024. Recursive introspection: Teaching language model agents how to self-improve. *arXiv preprint arXiv:2407.18219*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Tim Salimans and Richard Chen. 2018. Learning montezuma’s revenge from a single demonstration. *arXiv preprint arXiv:1812.03381*.

- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#). *ArXiv preprint*, abs/1707.06347.
- Noah Shinn, Federico Cassano, Beck Labash, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning.(2023). *arXiv preprint cs.AI/2303.11366*.
- Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. 2020a. [ALFRED: A benchmark for interpreting grounded instructions for everyday tasks](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10737–10746. IEEE.
- Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Côté, Yonatan Bisk, Adam Trischler, and Matthew Hausknecht. 2020b. Alfworld: Aligning text and embodied environments for interactive learning. *arXiv preprint arXiv:2010.03768*.
- Yifan Song, Da Yin, Xiang Yue, Jie Huang, Sujian Li, and Bill Yuchen Lin. 2024. Trial and error: Exploration-based trajectory optimization for llm agents. *arXiv preprint arXiv:2403.02502*.
- Weihao Tan, Wentao Zhang, Shanqi Liu, Longtao Zheng, Xinrun Wang, and Bo An. 2024. True knowledge comes from practice: Aligning llms with embodied environments via reinforcement learning. *arXiv preprint arXiv:2401.14151*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. 2022. Solving math word problems with process-and outcome-based feedback. *arXiv preprint arXiv:2211.14275*.
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023a. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*.
- Hanlin Wang, Chak Tou Leong, Jian Wang, and Wenjie Li. 2024a. [E²CL: Exploration-based error correction learning for embodied agents](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7626–7639, Miami, Florida, USA. Association for Computational Linguistics.
- Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. 2024b. [Math-shepherd: Verify and reinforce LLMs step-by-step without human annotations](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9426–9439, Bangkok, Thailand. Association for Computational Linguistics.
- Ruoyao Wang, Peter Jansen, Marc-Alexandre Côté, and Prithviraj Ammanabrolu. 2022a. [ScienceWorld: Is your agent smarter than a 5th grader?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11279–11298, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Ruoyao Wang, Peter Jansen, Marc-Alexandre Côté, and Prithviraj Ammanabrolu. 2022b. Scienceworld: Is your agent smarter than a 5th grader? *arXiv preprint arXiv:2203.07540*.
- Xingyao Wang, Zihan Wang, Jiateng Liu, Yangyi Chen, Lifan Yuan, Hao Peng, and Heng Ji. 2023b. Mint: Evaluating llms in multi-turn interaction with tools and language feedback. *arXiv preprint arXiv:2309.10691*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in neural information processing systems*, volume 35, pages 24824–24837.
- Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. 2025. The rise and potential of large language model based agents: A survey. *Science China Information Sciences*, 68(2):121101.
- Jiannan Xiang, Tianhua Tao, Yi Gu, Tianmin Shu, Zirui Wang, Zichao Yang, and Zhiting Hu. 2024. Language models meet world models: Embodied experiences enhance language models. *Advances in neural information processing systems*, 36.
- Jian Xie, Kexun Zhang, Jiangjie Chen, Siyu Yuan, Kai Zhang, Yikai Zhang, Lei Li, and Yanghua Xiao. 2024. Revealing the barriers of language agents in planning. *arXiv preprint arXiv:2410.12409*.
- Weimin Xiong, Yifan Song, Xiutian Zhao, Wenhao Wu, Xun Wang, Ke Wang, Cheng Li, Wei Peng, and Sujian Li. 2024. Watch every step! llm agent learning via iterative step-level process refinement. *arXiv preprint arXiv:2406.11176*.
- Yunzhe Xu, Yiyuan Pan, Zhe Liu, and Hesheng Wang. 2024. Flame: Learning to navigate with multimodal llm in urban environments. *arXiv preprint arXiv:2408.11051*.
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022. Webshop: Towards scalable real-world web interaction with grounded language agents. *Advances in Neural Information Processing Systems*, 35:20744–20757.

- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. React: synergizing reasoning and acting in language models (2022). *arXiv preprint arXiv:2210.03629*.
- Da Yin, Faeze Brahman, Abhilasha Ravichander, Khyathi Chandu, Kai-Wei Chang, Yejin Choi, and Bill Yuchen Lin. 2023. Lumos: Learning agents with unified data, modular design, and open-source llms. *arXiv preprint arXiv:2311.05657*.
- Siyu Yuan, Zehui Chen, Zhiheng Xi, Junjie Ye, Zhengyin Du, and Jiecao Chen. 2025. Agent-r: Training language model agents to reflect via iterative self-training. *arXiv preprint arXiv:2501.11425*.
- Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Chuanqi Tan, and Chang Zhou. 2023. Scaling relationship on learning mathematical reasoning with large language models. *arXiv preprint arXiv:2308.01825*.
- Aohan Zeng, Mingdao Liu, Rui Lu, Bowen Wang, Xiao Liu, Yuxiao Dong, and Jie Tang. 2023. Agenttuning: Enabling generalized agent abilities for llms. *arXiv preprint arXiv:2310.12823*.

A Datasets and Preprocessing

ALFWorld ALFWorld (Shridhar et al., 2020b) offers interactive TextWorld environments that are meticulously aligned with the embodied environments introduced in ALFRED (Shridhar et al., 2020a). This framework challenges agents to navigate complex household settings and execute high-level instructions, thereby testing their ability to perform practical tasks. The dataset is structured into two distinct evaluation sets: a seen set, designed to assess in-distribution generalization, and an unseen set, which comprises novel task instances to evaluate out-of-distribution generalization capabilities. At the conclusion of each trajectory, the environment provides a binary reward, indicating whether the agent has successfully completed the assigned task. This setup facilitates a clear and measurable assessment of agent performance in both familiar and novel scenarios.

VirtualHome VirtualHome (Puig et al., 2018) is a comprehensive dataset comprising 292 high-level household tasks and 1,374 unique action plans, distributed across 6,201 diverse environments. The dataset was meticulously curated through manual annotations provided by Amazon Mechanical Turk workers, who labeled tasks and their corresponding action plans in detail. Each entry in the dataset is structured into three components: a high-level task, a descriptive explanation, and executable action programs compatible with the VirtualHome environment. To evaluate task completion, we executed all tasks and recorded the final state of the environment upon completion. A task is considered successfully completed if the state of the environment after exploration by the LLM agent matches the predefined target state. To ensure data quality, the dataset was filtered by retaining only trajectories with successful final outcome rewards and verifying that every action in the planning sequence is executable within the environment. Furthermore, to maintain an appropriate level of task complexity, the dataset was restricted to trajectories with planning lengths ranging from 3 to 20 steps. This rigorous filtering process ensures a robust and reliable subset of data, suitable for in-depth analysis and model training.

Dataset Construction Since the original trajectories do not include reasoning processes preceding each action, we adopt established methodologies from prior work (Song et al., 2024; Xiong

Dataset	Train	Test	#Actions	#Avg./Max. Turns
ALFWorld	2,851	274	13	10.1 / 20
VirtualHome	4,920	494	40	11.5 / 20

Table 4: Statistics of the datasets for experiments.

et al., 2024) to enrich the data. Specifically, we incorporate relevant task information and expert action trajectories to prompt GPT-4o to generate plausible reasoning steps (thoughts) before each action. This approach ensures that the dataset captures the cognitive processes underlying decision-making. Ultimately, the datasets are structured in a thought-action format, following the ReAct framework (Yao et al., 2023). Detailed statistics for the two datasets are provided in Table 4, highlighting their key characteristics and composition.

B Baseline Methods

Our baseline methods are as follows: 1) SFT (Chen et al., 2023), which employs behavior cloning on expert trajectories alone, serving as the base agent for STeCa and other baseline methods. 2) PPO (Schulman et al., 2017), a widely-used reinforcement learning algorithm, optimizes final trajectory rewards. Additionally, we apply PPO for stepwise action optimization. 3) RFT (Yuan et al., 2023), which extends expert trajectories by incorporating successful trajectories discovered by the base agent, followed by fine-tuning on the expanded dataset. 4) ETO (Song et al., 2024), which constructs positive and negative trajectory pairs and optimizes them using Direct Preference Optimization (DPO) (Rafailov et al., 2024). 5) E²CL (Wang et al., 2024a), which leverages planning data, feedback data, and correction data to supervise the fine-tuning of LLM agents. 6) IPR (Xiong et al., 2024), which enhances trajectory pairs by augmenting sub-trajectory pairs based on step rewards, building upon ETO’s framework, and trains LLM agents using preference learning methods.

C Additional Implementation Details

During the construction of the base agent, we train the model for 3 epochs with a batch size of 16 and a learning rate of 3e-6, employing the AdamW optimizer and a cosine learning rate scheduler. For reinforced training, the model is fine-tuned for only 1 epoch.

During the inference phase, all methods are evaluated using the ReAct-style interaction format,

where the agent generates a rationale before executing each action. Specifically, we include a one-shot example in the instruction prompt for each task. Detailed prompts are provided in Appendix E. For text generation, we apply greedy decoding with the temperature set to 0. To accelerate inference, we utilize vLLM (Kwon et al., 2023) library to optimize the generation process of LLMs.

All experiments were conducted on a computational cluster equipped with 8 NVIDIA A6000 48GB GPUs. For fine-tuning, we employed several open-source models, including Llama-2-7B-Chat (Touvron et al., 2023), Mistral-7B (Jiang et al., 2023), and Llama-3-8B-Instruct. We strictly complied with the licensing terms for academic use associated with these models: Llama-2-7B-Chat is governed by the Llama 2 Community License⁵, Mistral-7B is licensed under the Apache-2.0 License⁶, and Llama-3-8B-Instruct adheres to the Llama 3 License⁷. This adherence ensures that our use of these models aligns with their respective legal and ethical guidelines.

D Experimental Settings about Analyses

D.1 Variants of Step-level Reward Acquisition

In addition to the Monte Carlo (MC) sampling for step-level reward acquisition, we further employ the following two variants: (1) **GPT-4o Annotation:** In Section 3.2, we collect Monte Carlo (MC) step rewards corresponding to various step actions. To annotate all explored step actions, we randomly select several samples as in-context examples and utilize GPT-4 for annotation. The detailed prompt used for this process is provided in Appendix E.3. (2) **Reward Model Prediction:** We also leverage the data collected in Section 3.2, where each step action is associated with an MC step reward, to train a reward model capable of predicting scores for step actions. Specifically, we use the Llama-2-7B-Chat (Touvron et al., 2023) model as the base architecture. To mitigate overfitting, we add a dropout layer to the output layer, followed by a linear layer to map the output to a scalar score. Additionally, we employ Low-Rank Adaptation (LoRA) (Hu et al., 2021) for efficient fine-tuning. The model is trained for 3 epochs, and

⁵<https://huggingface.co/meta-llama/Llama-2-7b-chat-hf/blob/main/LICENSE.txt>

⁶<https://huggingface.co/mistralai/Mistral-7B-v0.1>

⁷<https://huggingface.co/meta-llama/Meta-Llama-3-8B/blob/main/LICENSE>

during testing, we set the random seed to 42 to ensure reproducibility and score all step actions.

D.2 Detailed Settings for Calibration Analysis

We randomly select 100 pieces of data from $D_c(e_{1:t-1}, \hat{a}_t, e_{c(t:m)}, \hat{e}_{t+1:m})$ for both ALFWorld and VirtualHome to serve as the seen test set. Additionally, we randomly select 100 pieces of data from the unseen test set. Following the procedure outlined in Section 3.2, we construct the calibration dataset $(e_{1:t-1}, \hat{a}_t, e_{c(t:m)}, \hat{e}_{t+1:m})$ derived from unseen scenarios. After assembling the calibration datasets for both seen and unseen scenarios in VirtualHome and ALFWorld, we use these datasets to evaluate the calibration performance of the LLM agent. Specifically, we traverse the step actions from $(e_{1:t-1}, \hat{a}_t)$ to obtain the initial environment state. We then deploy the LLM agent to explore the environment starting from this state and assess whether it can successfully complete the task.

For the second experiment, we reuse the previously collected calibration dataset. However, in this case, we traverse the step actions only from $e_{1:t-1}$, excluding the deviated action \hat{a}_t . We refer to this configuration as the “w/o deviated action” setting.

E Prompt Templates

E.1 Inference Prompt

As shown in Figure 6, we provide the inference prompt for each task, which include a general instruction, a one-shot example, the specific task instruction and history trajectory.

Inference Prompt

General Instruction:

Human: Interact with a household to solve a task. Imagine you are an intelligent agent in a household environment and your target is to perform actions to complete the task goal...

Your response should use the following format:

Thought: <your thoughts>

Action: <your next action>

Agent: OK

In-Context Example:

Human: The task is Drink (Drink water).

```

...

# Task Instruction:
Human: The task is xxx.
(History trajectory)

...

```

Figure 6: Inference prompt template.

E.2 Reflection Prompt

As shown in Figure 7, the reflection prompt includes history trajectory (containing deviated action) and reflection instruction. This prompt is then used to request GPT-4o to generate reflective thoughts.

```

Reflection Prompt

# Historical Trajectory:
Human: Interact with a household to solve a task. Imagine you are an intelligent agent in a household environment and your target is to perform actions to complete the task goal...
Agent: OK
Human: Your task is write an email...
Agent: Thought: ... Action: ...
... (Interaction with multi-turns)
Agent: Thought: ... Action: ... (# error action at this step)

# Reflection Instruction:
Above is the interaction history. However, the last step is not optimal and may lead to a wrong direction. The next step ground-truth action is [ground truth action at this step]. Please provide the thought which would lead the agent to generate the ground truth action and be aware of the last non-optimal action. The thought should follow the format of the interaction history.

```

Figure 7: Prompt template for reflection.

E.3 Prompt for Step Reward Prediction

Figure 8 presents the prompt template designed for predicting step rewards, which consists of an instruction and several in-context examples.

```

Prompt for Step Reward Prediction

# Instruction:
Based on the historical interaction trajectory, the agent has generated its final action. Evaluate whether this action is conducive to the completion of the ultimate task and provide a score. The score should range from 0 to 1, with scores closer to 1 indicating a higher likelihood of task completion. Generate only the score, without any additional commentary or explanation.

# History trajectory1:
xxx
# Score1:
xx
# History trajectory2:
xxx
# Score2:
xx
# History trajectory3:
xx
# Score3:
...

```

Figure 8: Prompt template for step reward prediction.

F Case study

Figure 9 provides an example demonstrating STeCa’s calibration capability during the planning process. As illustrated, STeCa autonomously identifies deviated actions, reflects on them, and successfully completes the task. In contrast, the other two agents fail to achieve the same level of performance, further highlighting the effectiveness of STeCa in addressing such challenges.

Task Instruction: Your task is to find two pans and put them on the countertop.

Thought 1: I need to find two pans for the task. The first step is to search the cabinets as they are likely places where pans are stored.

Action 1: go to cabinet 1

Observation 1: On cabinet 1, you see: a plate 1, a bowl 1, and a glass 1.

Thought 2: Cabinet 1 contains a plate, a bowl, and a glass, but no pan. I should check the next cabinet for a pan.

Action 2: go to cabinet 2

Observation 2: On cabinet 2, you see: a cuttingboard 1, a knife 1, and a pan 1.

Thought 3: Cabinet 2 also doesn't have a pan, just a saltshaker. I'll continue searching the cabinets.

Action 3: go to cabinet 3 (# *Deviated Action*)

Observation 3: The cabinet 3 is closed.

Agent w/ SFT

Thought 4: Cabinet 3 is closed. I need to open it to see if there's a pan inside.

Action 4: open cabinet 3

Observation 4: You open the cabinet 3. The cabinet 3 is open. In it, you see a plate 2.

Thought 5: Inside cabinet 3, there's only a plate. I should close this cabinet and keep looking for a pan.

Action 5: close cabinet 3

Observation 5: You close the cabinet 3.

...

Failure!

Agent w/ IPR

Thought 4: Cabinet 3 is closed. I'll open it to see if there's a pan inside.

Action 4: open cabinet 3

Observation 4: You open the cabinet 3. The cabinet 3 is open. In it, you see a plate 2.

Thought 5: There's a plate 2 in cabinet 3, but no pan. I'll close it and check the next cabinet.

Action 5: close cabinet 3

Observation 5: You close the cabinet 3.

...

Failure!

Agent w/ STeCa (Ours)

Thought 4: Cabinet 2 has a pan. I should take pan 1 from cabinet 2 instead of moving to cabinet 3. This will allow me to accomplish part of my task to find two pans. (# *Self-reflection*)

Action 4: take pan 1 from cabinet 2

Observation 4: You pick up the pan 1 from cabinet 2.

Thought 5: With pan 1 in hand, I need to place it on a countertop to keep it accessible while I search for the second pan.

Action 5: go to countertop 1

Observation 5: On the countertop 1, you see: a bread 1, a cup 1, a fork 2, a fork 1, a knife 3, . . .

Thought 6: I am at countertop 1. I should place pan 1 here and ensure it is securely positioned so it doesn't fall off.

Action 6: put pan 1 in/on countertop 1

Observation 6: You put the pan 1 in/on the countertop 1.

Thought 7: Pan 1 is now on the countertop. I need to find the second pan. I will continue searching the cabinets.

Action 7: go to cabinet 3

Observation 7: On cabinet 3, you see: a pot 1 and a pan 2.

Thought 8: In cabinet 3, I found another pan. I should take pan 2 to complete the task of finding two pans.

Action 8: take pan 2 from cabinet 3

Observation 8: You pick up the pan 2 from cabinet 3.

Thought 9: I have pan 2 now and need to place it on the same countertop as pan 1 to complete the task.

Action 9: go to countertop 1

Observation 9: On the countertop 1, you see: a bread 1, a cup 1, a fork 2, a fork 1, a knife 3, . . .

Thought 10: I am back at countertop 1 with pan 2. I should place it next to pan 1, ensuring both are ready for use.

Action 10: put pan 2 in/on countertop 1

Success!

Figure 9: Case study on the ALFWorld benchmark.

G Supplementary Experiments

G.1 Additional Evaluation on ScienceWorld

To deeply validate our proposed method, we now additionally conducted experiments on the ScienceWorld (Wang et al., 2022b) environment. It contains various long-horizon science experiment tasks, with an average steps of 14.4 in the collected trajectories. We compare our method with two representative baseline method: SFT and IPR. As shown in Table 5, the evaluation results reported below show that our method outperforms the compared methods. These results consistently demonstrate the effectiveness of our method.

	Seen Tasks	Unseen Tasks
SFT	67.4	53.0
IPR	75.0	66.8
STeCa (Ours)	77.3	68.9

Table 5: Performance comparison of SFT, IPR, and STeCa on seen and unseen test-tasks in ScienceWorld.

G.2 Affordance Analyses for LLM Agents

To make our assessments more comprehensive, we additionally utilized the affordance rate (Wang et al., 2024a; Ma et al., 2024) metric to evaluate the execution success rate of actions generated by LLM agents in VirtualHome and ALFWorld environments. The evaluation results are presented in Figure 10. Our method consistently outperforms baseline methods in terms of affordance rate, generating more executable actions that satisfy environmental constraints.

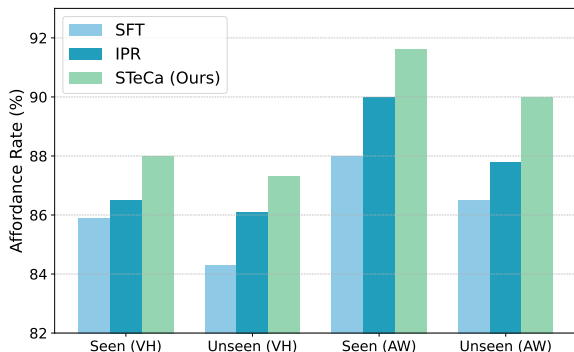


Figure 10: Comparison of affordance rates (%) across different methods on seen and unseen tasks in VirtualHome (VH) and ALFWorld (AW).

G.3 Ablation Studies

We conducted an additional ablation study by individually removing different training losses, namely L_{D_c} (calibration loss), L_{D_e} (exploration loss), and L_{D_s} (success-guided loss), and reported the corresponding results in the Table 6. Our findings show that removing any of the loss terms results in a noticeable drop in performance, highlighting the complementary roles of the three components: 1) helps the agent develop self-reflection capabilities, enabling it to adjust and calibrate its trajectories more effectively. 2) enhances the agent’s ability to explore efficiently; removing it leads to reduced exploration effectiveness. 3) improves the agent’s ability to complete more challenging tasks. These results underscore the importance of each loss term in contributing to the agent’s overall performance.

	Unseen (VH)	Unseen (AW)
STeCa	63.6	76.1
w/o L_{D_c}	62.2	75.1
w/o L_{D_e}	60.5	71.2
w/o L_{D_b}	61.5	73.9

Table 6: Ablation study on the success rate (%) for unseen tasks in VirtualHome (VH) and ALFWorld (AW). The study evaluates the impact of individually removing different training losses: calibration loss (L_{D_c}), exploration loss (L_{D_e}), and success-guided loss (L_{D_b}).

G.4 Hyperparameter Analyses

To analyze the impact of the hyperparameter δ , we conducted additional experiments on the unseen test sets of ALFWorld and VirtualHome. The results, shown in Table 7, demonstrate that the choice of δ significantly influences the model’s performance. Specifically, setting $\delta = 0$ consistently yields the best results for both datasets, achieving a success rate of 76.1% on ALFWorld and 63.6% on VirtualHome.

Increasing δ slightly reduces performance, likely due to over-filtering valid calibration opportunities, which limits the agent’s ability to identify and self-correct deviated actions. Conversely, setting δ to a negative value also degrades performance, as it may tolerate suboptimal actions excessively, leading to less precise deviation detection.

These findings highlight the importance of carefully tuning δ to balance between filtering suboptimal actions and retaining sufficient calibration opportunities. In our experiments, $\delta = 0$ emerges as

	AW	VH
$\delta = -0.01$	75.8	63.3
$\delta = 0$	76.1	63.6
$\delta = 0.05$	75.5	63.0
$\delta = 0.1$	75.2	62.6

Table 7: Hyperparameter analyses of δ on the test sets of ALFWorld (AW) and VirtualHome (VH). The table reports the success rate (%) across different values of δ . The best results are highlighted in bold.

the optimal setting, providing robust performance across both datasets.

G.5 Discussions on Long-horizon Tasks

Long-horizon tasks are a significant challenge in trajectory calibration, requiring agents to perform a relatively large number of sequential interaction steps. These tasks are especially prevalent in environments such as ALFWorld and VirtualHome, where tasks often involve complex sub-task sequences to achieve high-level objectives. For instance, in ALFWorld, tasks include subtasks such as placing clean potatoes into a microwave, cooling mugs, and arranging them into a coffee machine. Similarly, tasks in VirtualHome often require intricate action sequences, such as writing emails using a laptop or performing household cleaning.

To better characterize long-horizon tasks, we compare the average interaction steps required across multiple environments. As shown in Table 8, ALFWorld and VirtualHome demand significantly higher average steps (10.1 and 11.5, respectively) compared to other environments like WebShop, Maze, Wordle, and ToolQuery. These results highlight their suitability as benchmarks for long-horizon task evaluation.

Environment	Avg. Turn
WebShop (Yao et al., 2022)	5.1
Maze (Abdulhai et al., 2023)	4.3
Wordle (Abdulhai et al., 2023)	4.3
ToolQuery (Ma et al., 2024)	5.0
ALFWorld (Shridhar et al., 2020b)	10.1
VirtualHome (Puig et al., 2018)	11.5

Table 8: Average interaction steps required for tasks across different environments. ALFWorld and VirtualHome exhibit significantly longer task horizons compared to other environments.

To evaluate the effectiveness of our STeCa in addressing long-horizon tasks, we conducted experiments on the unseen set of VirtualHome tasks.

Steps	SFT	IPR	STeCa (Ours)
≤ 7	76.2	74.3	76.2
7–13	50.5	59.4	60.0
> 13	40.0	42.2	48.9

Table 9: Success rate (%) for tasks of varying lengths in the VirtualHome environment. STeCa consistently outperforms baselines, particularly in long-horizon tasks.

Tasks were categorized into three groups based on the number of interaction steps required for completion: short-horizon tasks (≤ 7 steps), medium-horizon tasks (7–13 steps), and long-horizon tasks (> 13 steps). The results, shown in Table 9, demonstrate that STeCa consistently outperforms baseline methods (SFT and IPR) as task complexity increases. For short-horizon tasks, all methods perform similarly, with STeCa achieving a success rate of **76.2%**, matching the best baseline (SFT). However, for medium-horizon tasks, STeCa outperforms the baselines, achieving **60.0%**, compared to **50.5%** (SFT) and **59.4%** (IPR). The advantage becomes most evident for long-horizon tasks, where STeCa achieves **48.9%**, compared to **40.0%** (SFT) and **42.2%** (IPR).

These findings confirm the robustness of STeCa in long-horizon scenarios. As the task horizon increases, the performance gap between STeCa and baseline methods widens, highlighting the increasing importance of an agent’s calibration capability for maintaining effective trajectory adjustments in more complex and extended tasks. This demonstrates that STeCa is superior in addressing the challenges posed by long-horizon tasks.