How well can LLMs provide planning feedback in grounded environments?

Abstract

Learning to plan in grounded environments typically requires carefully designed reward functions or high-quality annotated demonstrations. Recent works show that pretrained foundation models, such as large language models (LLMs) and vision language models (VLMs), capture background knowledge helpful for planning, which reduces the amount of reward design and demonstrations needed for policy learning. We evaluate how well LLMs and VLMs provide feedback across symbolic, language, and continuous control environments. We consider prominent types of feedback for planning, including binary feedback, preference feedback, action advising, goal advising, and delta action feedback. We also consider inference methods that impact feedback performance, including in-context learning, chain-of-thought, and access to environment dynamics. We find that foundation models can provide diverse and useful feedback across many domains, yet their reliability is highly dependent on model scale, reasoning capabilities, and domain representation. Larger models and reasoning models benefit more from enhanced inference methods. Finally, feedback quality degrades for environments with complex dynamics or continuous state spaces and action spaces.

1 Introduction

Training instruction following agents in grounded environments often requires a large number of trials (Mnih et al., 2013; Schulman et al., 2017) or a large set of high-quality annotated demonstrations (Stepputtis et al., 2020; Mandlekar et al., 2021; Ahn et al., 2022). Recent progress in large-scale pretraining shows that large language models (LLMs) capture a significant amount of world-knowledge that is helpful for reasoning and planning (Brown et al., 2020; Ahn et al., 2022). While these models struggle as direct policies in grounded environments (Mandlekar et al., 2021; Xie et al., 2024), they can provide helpful feedback on agent behaviour. For instance, LLMs can generate and refine executable reward functions with human feedback (Kwon et al., 2023; Xie et al., 2023), as well as identify productive behaviour for imitation (Colas et al., 2023; Zhong et al., 2024).

We conduct a comprehensive study on how well LLMs and VLMs can provide feedback in grounded environments. Given a LLM/VLM-based Feedback Model (FM), a language instruction, observation from the environment, and agent behaviour, we compare FM feedback on the agent behaviour to ground truth feedback derived from an optimal policy. We evaluate state-of-the-art FMs across model sizes, reasoning capabilities, multiple domains, observation modality, and inference methods. First, we consider feedback types commonly used in planning: **binary feedback** on whether the action is optimal; **preference** on which action out of a pair results in faster task completion; **action advising** to choose the correct action; **goal advising** to identify the next subgoal; and **delta action** to provide action correction. Second, we consider a diverse selection of domains from symbolic games, verbalised text environments, and continuous control of a robotic arm. Third, we evaluate how the quality of feedback changes under inference methods that affect model performance, such as observation modality, hints about domain dynamics, in-context learning, and chain-of-thought.

We find Feedback Models can achieve high feedback accuracy across many domains. First, feedback quality improves with larger and reasoning models. The latter exhibit less class and positional bias, and benefit the most from test-time techniques such as hints about domain dynamics, in-context learning, and chain-of-thought. Second, out of the feedback types we consider, FMs tend to provide high-quality binary and preference feedback more easily — this is especially true for non-reasoning

Figure 1: Our evaluation pipeline for different types of feedback. We sample trajectories according to a policy π and construct snapshots x from the trajectory that corresponds to a certain type of feedback. For example, in action advising, we only need to sample states. Then, we compute the ground truth feedback y^*. expert policy π^* For each (x, y^*), we prompt the Feedback Model and compare its response y to the ground truth feedback y^* to report the accuracy of feedback and F1 score (if possible).

models. Third, feedback quality degrades in domains with complex, hierarchical dynamics, or domains with continuous state and action spaces. On predominantly visual environments, we find that VLMs consistently outperform LLMs on preference and binary feedback in cases such as continuous control, where image observations are much easier to reason about than textual descriptions.

To the best of our knowledge, this is the first comprehensive analysis on how well LLMs and VLMs provide across diverse feedback types, domains, and inference methods. We show that language model feedback is a promising means to extract learning signals for planning in a diverse selection of environments, without hand-engineered reward functions or human demonstrations. Future research should investigate how to effectively elicit robust feedback, as well as how to learn from such feedback efficiently. Given the success of high-quality language model feedback across diverse domains, we advocate for further research into how this feedback can be used to improve policy learning efficiency and effectiveness. Additionally, we advocate for further research in how to extract high-quality feedback in domains with complex dynamics and continuous state and action spaces — features that make it particularly challenging to design rewards and label demonstrations.

2 METHODOLOGY

We evaluate how well LLMs and VLMs (which we will refer to as LMs) provide feedback in grounded environments. Specifically, an LM serves as a **Feedback Model** (FM) that provides judgment on how well an agent is performing a task given the task instruction and snapshots of the agent's behaviour. Figure 1 presents an overview of how we evaluate the Feedback Model.

We formalize feedback collection using terminology from sequential decision-making. Let S denote the state space, A the action space, $T : S \times A \times S \to [0, 1]$ the transition function, $R : S \times A \times S \to \mathbb{R}$ the reward function, and γ the discount factor. We assume a Markov Decision Process $\mathcal{M} = \langle S, A, T, R, \gamma \rangle$, where an agent acts according to policy $\pi_\theta : S \to A$. In planning, we find an optimal policy π^* that maximizes expected cumulative discounted future reward $G = \mathbb{E}_{(s,a) \sim \pi}[\sum_t \gamma^t R(s_t, a_t, s_{t+1})]$. To characterize how well a state maximizes G, we construct the value function $V_\pi(s) = \mathbb{E}_{(s,a) \sim \pi}[G | s_0 = s]$ and the state-action value function $Q_\pi(s, a) = \mathbb{E}_{(s,a) \sim \pi}[G | s_0 = s, a_0 = a]$. In order to evaluate the Feedback Model, we assume access to an optimal policy π_θ^*. The Feedback Model is given x, a **snapshot** consisting of policy behaviour and environment observations. The FM then provides $F(x) = y$, its **feedback** on

the behaviour x. We evaluate the correctness of the judgment by comparing it with **ground truth feedback** y^* computed using the optimal policy.

2.1 TYPES OF FEEDBACK

We consider five types of feedback commonly used in planning. For each, we describe the composition of snapshot x and how the ground truth feedback y^* is computed.

Binary feedback. The FM observes a snapshot containing the state and an action $x = (s, a)$ and provides a binary judgment $F(x) = y_{\text{bin}} \in \{-1, 1\}$ representing whether the action a is optimal in state s. The ground truth binary feedback is 1 if $\pi^*(s) = a$ and -1 otherwise. For cases where there exist several optimal actions, we consider y_{bin} to be correct if a is any of the optimal actions. Appendix G details how we compute y^* for different domains. Binary feedback estimates binarized advantage value $A_\pi(s, a) = Q_\pi(s, a) - V_\pi(s)$. Pérez-Dattari et al. (2020) and Warnell et al. (2018) use binary feedback to directly learn policies. Zhong et al. (2024) use binary feedback to identify productive behaviour for imitation learning.

Action advising. The FM directly provides the optimal action given the state presented in the snapshot $x = s$. The ground truth feedback is the output of the optimal policy $y_{\text{act}}^* = \pi^*(s)$. For cases where there exist several optimal actions, we consider y_{act} to be correct if it is any of the optimal actions. LLM has been used to predict optimal actions for long-horizon games (Wang et al., 2023), robotic control (Ahn et al., 2022), and computer use (Xie et al., 2023).

Preference feedback. The FM observes a snapshot consisting of the state and a pair of potential actions $x = (s, a_1, a_2)$, and chooses the action more conducive for solving the task (e.g. higher value $Q(s, a)$). The ground truth feedback uses Q^* under the optimal policy in order to compute the ground truth preference:

$$y_{\text{pref}}^* = \begin{cases} +1, & \text{if } Q_{\pi^*}(s, a_1) > Q_{\pi^*}(s, a_2) \\ -1, & \text{if } Q_{\pi^*}(s, a_1) < Q_{\pi^*}(s, a_2) \\ 0 & \text{otherwise} \end{cases}$$

Lee et al. (2021) use human preference to learn a reward model for reinforcement learning. Yang et al. (2023) also use preferences as constraints for policy updates.

Goal advising. Goal advising is similar to action advising in that the FM is given a snapshot consisting of the state $x = s$. However, instead of predicting the next action, the FM predicts the ideal next state s' to best solve the task (e.g. higher value $V(s')$). The ground truth feedback is defined as: $y_{\text{goal}} = \text{argmax}_{s' \in \text{Reachable}(s)} V_{\pi^*}(s')$, where V_{π^*} is the value function of the optimal policy. Here, we use the transition function T to define $\text{Reachable} : s \to \{s' \in S \mid \exists a \in A, T(s, s', a) > 0\}$ as a mapping from a state s to a set of states with non-zero transition probability from s. If there are multiple optimal next states, we consider y_{goal} to be correct if it is any of the optimal next states. Du et al. (2023) use LLMs to generate subgoals and give intermediate rewards to the agent when it achieves the generated subgoals. Zhao et al. (2024) similarly decompose goals using LLMs, however the low-level goals are then further given to policy LLMs to predict the appropriate action.

Delta action feedback. The FM is given a snapshot consisting of the state and an action $x = (s, a)$, and provides a corrective delta action Δ such that $a + \Delta$ is more conducive for solving the task (e.g.higher value $Q(s, a + \Delta) > Q(s, a)$). The ground truth feedback uses the optimal policy in order to compute the ground truth delta action $y_\Delta^* = \pi^*(s) - a$. We consider y_Δ to be correct if it is within some threshold ϵ of the ground truth y_Δ^*. Liu et al. (2023) map verbal feedback provided by human observers to a set of delta actions, which are used to correct actions produced by a policy. Colas et al. (2023) convert agent trajectories to language descriptions, and use LLMs to retroactively provide corrections to actions in the trajectory.

2.2 TASK DOMAINS

In order to measure how well Feedback Models provide feedback across diverse settings, we collect policy behaviours from several distinct domains that vary in observation space, action space, degree of observability, and task settings. Figure 2 illustrates these domains, which include symbolic games

(a) Cliff Walking   (b) ALFWorld   (c) HierarchyCraft

(d) MiniGrid Door Key   (e) MiniGrid Four Rooms   (f) Robomimic Lift Object

Figure 2: Six evaluation domains for language model feedback.

CliffWalking (Sutton, 2018) and MiniGrid (Chevalier-Boisvert et al., 2023), navigation and high-level manipulation text simulators ALFWorld (Shridhar et al., 2021) and HierchyCraft (Federico, 2023), and continuous control Robomimic (Mandlekar et al., 2021). We evaluate LLM feedback for all domains, and VLM feedback for predominantly visual domains (MiniGrid Door Key and Robomimic). Appendix F shows examples of FM prompts for each environment.

Cliff Walking   The agent navigates to a fixed goal location in a grid of size 4x12 without falling off the cliff in the last row. The agent has full observability of the grid and its current position (e.g. `(3, 0)`), and chooses among the actions `right`, `left`, `up`, and `down` to navigate to the goal.

MiniGrid (Chevalier-Boisvert et al., 2023)   We consider two MiniGrid tasks, where the agent navigates between rooms to reach a goal location (green rectangle). In Four Rooms (Fig 2e), the agent needs to navigate four rooms to reach the goal. It chooses among the actions `turn left`, `turn right`, `move forward`. In Door Key (Fig 2d), the agent must first pick up the key, unlock the door, then reach the goal. The agent chooses among actions `turn left`, `turn right`, `move forward`, `pick up`, and `unlock door`. For both tasks, the agent has full observability of the room layout and content. We evaluate VLM feedback on Door Key, as it involves more visual understanding components like locked door, unlocked door, wall and key. We use a screenshot of the current grid as the observation.

HierarchyCraft (Federico, 2023)   We consider the task of finding a diamond. The agent must navigate to appropriate locations to collect materials (e.g. iron, wood) and craft items (e.g. iron pick axe) helpful for procuring the diamond. The action space consists of 201 discrete actions (e.g. searching for wood in the forest). The agent observes its inventory and the items in the environment via text description. This environment is partially observable, as objects in a location are hidden by default.

ALFWorld (Shridhar et al., 2021)   The agent navigates and performs high-level interactions with objects in a house in order to accomplish high-level tasks (e.g. wash mugs and put them in the cabinet). The agent observes the goal and text descriptions of the scene, and chooses among a list of valid actions described in language. This environment is partially observable, as objects in a location are hidden by default (e.g. cups in a cabinet).

Robomimic (Mandlekar et al., 2021)   We study the lift object task, where a 7-DoF robot arm lifts an object. The agent has full observability of the end effector position, the object position. It chooses from a list of control vectors that move the end effector and adjust the gripper state. Similar to Mandlekar et al. (2021), we assume an external policy that can accurately control the gripper

Domain	Average Legal Action Size	Percentage of States with Multiple Optimal Action	Optimal Action Percentage
Door Key	5	0.03	0.21
Four Rooms	3	0.00	0.33
ALFWorld	26.6	0.51	0.25
HierarchyCraft	11.1	0.41	0.13

Table 1: Statistics about states where multiple expert actions frequently exist in the expert advising dataset in domains of Door Key, Four Rooms, ALFWorld and HierarchyCraft. The percentage of optimal actions is the number of optimal actions divided by the total legal actions in the dataset.

rotation. We represent the observations via text description and provide a list of delta actions for the agent to choose from. For VLM feedback, we instead use a single image as the current observation, which is captured from a camera in front of the desk and robot arm (shown in Figure 2f).

2.3 DATASET COLLECTION

For each domain, we collect a dataset of policy behaviour $\mathcal{D} = \{x_1, x_2 \ldots\}$ by sampling from a policy π. For each snapshot x_i, we first compute the set of optimal actions using the optimal policy π^*, then compute the ground truth feedback y^*. For Cliff Walking, we enumerate all states/ actions for a dataset of 37 snapshots for action/goal advising, 148 for binary feedback, and 444 for preference feedback. For other domains, we use 3 choices of π: expert $\pi = \pi^*$, random policy that uniformly samples actions, and half-expert which alternates between the two with 50% chance. For each state, we enumerate all actions to construct $\mathcal{D}$. For example, for Four rooms, we sample N states $s_1 \ldots s_N$ from π and create a state-action pair for each of the 3 actions to build a binary feedback dataset of $3N$ examples. We cannot compute the optimal actions for Robomimic Lift Object, as it is impossible to enumerate all possible continuous actions. Following Mandlekar et al. (2021), we instead discretize the action space using a threshold based on empirical evaluation of a behavioural cloning expert, and use this threshold to enumerate the set of delta actions. More details about dataset collection is shown in the Appendix Section G .

2.4 IMPACT OF ENVIRONMENT DYNAMICS

In some domains, such as MiniGrid DoorKey, multiple optimal actions can exist at a given state (e.g., turning left or right is equivalent when the agent simply needs to turn around), which can, in principle, make the task easier. The statistics for domains with multiple optimal actions are summarized in Table 1. For Four Rooms, in our full expert action advising dataset of 10,000 states, only 17 states had multiple optimal actions (e.g., moving forward and then turning right versus turning right and then moving forward). In the downsampled 500 points, 100% of states had exactly 1 optimal action. For ALFWorld, although 51% of states have multiple optimal actions (due to the partial ordering of sub-goals), the problem remains challenging: the average branching factor (number of legal actions) is 26.6, and the proportion of optimal actions (optimal actions / total legal actions) is only 25%. Thus, a random policy would fail 75For HierarchyCraft, the average legal action space is 11.14, and the proportion of strictly optimal actions is only 13%. Overall, in our domains, a random policy cannot easily succeed, and under our requirement that feedback be optimal (shortest-path to the goal), planning remains challenging in both HierarchyCraft and ALFWorld.

2.5 FEEDBACK CONDITIONS

We evaluate FM inference methods that can affects LLM performance in grounded environments.

In-Context Learning (ICL) ICL improves LLM performance by providing input and output examples (Brown et al., 2020), from which LLMs deduce patterns that help generalize to new tasks (Dong et al., 2024; Chen et al., 2025). In this setting, we combine instructions, observations, and ground truth feedback as ICL examples for the FM.

Domain	Preference		Action Adv		Bin Feedback		Goal Adv	
Cliff Walking	0.33	0.77	0.68	0.73	0.61	0.73	0.37	0.77
	0.91	1.00	0.99	1.00	1.00	0.98	1.00	1.00
Door Key	0.66	0.75	0.44	0.44	0.78	0.75	N/A	
	0.89	0.98	0.81	0.99	0.92	1.00		
Four Rooms	0.59	0.66	0.32	0.42	0.59	0.69	N/A	
	0.65	0.75	0.61	0.78	0.66	0.81		
Robomimic	0.55	0.67	0.01	0.00	0.47	0.50	N/A	
	0.55	0.56	0.00	0.00	0.51	0.46		
ALFWorld	0.49	0.67	0.60	0.83	0.76	0.74	N/A	
	0.78	0.83	0.85	0.83	0.77	0.76		
HierachyCraft	0.51	0.51	0.58	0.61	0.81	0.85	N/A	
	0.81	0.89	0.65	0.78	0.88	0.91		

Table 2: Feedback accuracy for Llama 3.1 8B , Llama 3.1 70B , QwQ 32B , and R1 . Delta action is only applicable for continuous state spaces (Robomimic), its feedback column is not shown due to space constraint — Robomimic Delta Action numbers are 0.07 for Llama 3.1 8B, 0.12 for Llama 3.1 70B, 0.03 for QwQ, and 0.07 for R1. For goal advising, we only consider cases ground truth feedback is computable — there exists an intermediate state that the FM can propose which the optimal policy must encounter (Cliff Walking).

Chain-of-Thought and Thinking Guides (CoT)   CoT techniques encourage LLMs to generate intermediate steps towards the final solution (Wei et al., 2022), and consistently improve performance for reasoning problems such as program synthesis (Chen et al., 2023) and planning (Yao et al., 2023). We use CoT in all experiments. Additionally, we design planning-specific questions, called **Thinking Guides**, for the FM to answer to encourage the completion of helpful high-level intermediate steps (e.g. reasoning about its inventory). Instead of trying to deduce the reasoning procedure from ICL examples, Thinking Guides directly gives FMs high-level hints for how to reason.

Access to High-level Domain Dynamics   Kaiser et al. (2020) and Hafner et al. (2023) show that knowledge of environment dynamics improve model-based reinforcement learning. In this setting, we enumerate the transition dynamics of the environment and either provide a full description of the dynamics or a limited partial version explaining only legal actions. We additionally explore a second setting, where instead of directly providing environment dynamics, we provide **history context** to the FM model. This evaluates the FM's ability to infer relevant dynamics from the history. In order to provide history context for Cliff Walking, we use an expert policy to sample trajectories.

3   EXPERIMENTS AND ANALYSES

We design experiments to quantify how well LMs provide feedback across different feedback types, environment domains, observation modality, and inference methods. We conduct our study with LLama 3.1 7B, Llama 3.1 70B (Grattafiori et al., 2024), QwQ 32B (Qwen Team, 2024), and DeepSeek R1 (DeepSeek-AI et al., 2025). The Llama models are non-reasoning models, while QwQ and R1 are reasoning models with performance similar to the best closed-source models. We evaluate VLM feedback with Qwen 2.5 VL 7B, Qwen 2.5 VL 32B, Qwen 2.5 VL 72B (Bai et al., 2025) and GPT o4 mini. We show Feedback Model prompts in Appendix F for each feedback type and for each feedback condition. We consider malformatted feedback and feedback with illegal values to be incorrect. In the evaluation, we down-sample the large datasets for ALFWorld and Robomimic to 1000. For preference feedback, we only evaluate feedback on pairs of optimal and non-optimal actions, as opposed to pairs of equally optimal actions.

The accuracy of Feedback Models across domains and feedback types under the expert policy is shown in Table 2. We show results on snapshots from expert policy, except for the smaller Cliff Walking domain where we enumerate all possible states. Additional results using random policies and mixed policies are shown in Appendix C.

Domain	Preference	Action Adv	Bin Feedback	Goal Adv
Cliff Walking	0.59 ↑.26 0.78 ↑.01
0.95 ↑.04 0.97 ↓.03	0.57 ↓.15 0.62 ↓.11
1.00 ↑.01 0.97 ↓.03	0.69 ↑.08 0.80 ↑.07
0.95 ↓.05 0.98 =	0.58 ↑.21 0.97 ↑.20
0.97 ↓.03 0.98 ↓.02
Door Key	0.50 ↓.11 0.71 ↓.04
0.94 ↑.05 0.98 =	0.35 ↓.09 0.55 ↑.11
0.87 ↑.06 0.98 ↑.01	0.68 ↓.10 0.72 ↓.03
0.92 = 0.98 ↓.02	N/A
Four Rooms	0.58 ↓.01 0.61 ↓.05
0.65 = 0.79 ↑.04	0.53 ↑.20 0.54 ↑.12
0.60 ↓.01 0.83 ↑.05	0.42 ↓.17 0.50 ↓.19
0.54 ↓.11 0.86 ↑.05	N/A
ALFWorld	0.50 ↑.01 0.71 ↑.09
0.86 ↑.08 0.85 ↑.02	0.69 ↑.09 0.85 ↑.01
0.88 ↑.03 0.93 ↑.10	0.74 ↓.02 0.77 ↑.03
0.76 ↓.01 0.77 ↑.01	N/A
HierachyCraft	0.56 ↑.05 0.93 ↑.42
0.89 ↑.08 0.93 ↑.04	0.63 ↑.05 0.68 ↑.07
0.52 ↓.13 0.84 ↑.06	0.83 ↑.02 0.87 ↑.02
0.87 ↑.01 0.89 ↓.02	N/A

Table 3: Feedback accuracy with in-context learning, the arrows show accuracy change with this condition relative to baseline.

Domain	Preference	Action Adv	Bin Feedback	Delta Action
Door Key	0.39 0.52
0.60 0.61	0.42 0.27
0.32 0.43	0.48 0.65
0.64 0.77	N/A
In Context Learning	0.40 ↑.01 0.53 ↑.01
0.58 ↓.02 0.58 ↓.03	0.47 ↑.05 0.18 ↓.09
0.24 ↓.08 0.37 ↓.06	0.74 ↑.22 0.75 ↑.10
0.67 ↑.03 0.70 ↓.07	N/A
Thinking Guides	0.39 = 0.49 ↓.03
0.45 ↓.15 0.60 ↓.01	0.51 ↑.09 0.18 ↓.09
0.23 ↓.09 0.28 ↓.15	0.40 ↓.08 0.31 ↓.34
0.42 ↓.22 0.48 ↓.29	N/A
Unk	0.50 ↑.11 0.44 ↓.08
0.54 ↓.06 0.45 ↓.16	0.27 ↓.15 0.15 ↓.12
0.26 ↓.06 0.15 ↓.28	0.54 ↑.06 0.51 ↓.14
0.58 ↓.06 0.57 ↓.20	N/A
Unk, History	0.50 ↑.11 0.71 ↑.19
0.66 ↑.06 0.55 ↓.06	0.40 ↓.02 0.50 ↑.23
0.40 ↑.08 0.42 ↓.01	0.74 ↑.26 0.67 ↑.02
0.67 ↑.03 0.67 ↓.10	N/A
History	0.22 ↓.17 0.72 ↑.20
0.72 ↑.12 0.61 =	0.49 ↑.07 0.59 ↑.22
0.47 ↑.15 0.49 ↑.06	0.72 ↑.24 0.75 ↑.10
0.74 ↑.10 0.79 ↑.02	N/A
Robomimic	0.82 0.77
0.82 0.70	0.00 0.00
0.00 0.00	0.43 0.18
0.04 0.09	0.01 0.05
0.04 0.01
History	0.54 ↓.28 0.75 ↓.02
0.80 ↓.02 0.76 ↑.06	0.00 = 0.00 =
0.00 = 0.00 =	0.56 ↑.13 0.62 ↑.44
0.62 ↑.58 0.45 ↑	0.02 ↑.01 0.02 ↓.03
0.03 ↑.01 0.01 ↑.03

Table 4: VLM Feedback accuracy for Qwen 2.5 VL 7B , Qwen 2.5 VL 32B , Qwen 2.5 VL 72B , and o4-mini . The arrows show accuracy change with this condition relative to baseline.

LLMs can provide high-quality feedback Table 2 shows that LLMs can provide accurate feedback across domains. Moreover, larger models and reasoning models provide more accurate feedback regardless of feedback type, across different domains. For instance, the smaller QwQ consistently outperforms the larger non-reasoning Llama 3.1 70B. One source of small FM errors is that small FMs are prone to bias. First, Appendix D shows they are biased towards predicting actions as non-optimal, resulting in much worse recall than precision for binary feedback. Second, Appendix E shows they are biased towards preferring the first action, resulting in much worse accuracy when the ground truth action is located second for preference feedback. In contrast, larger FMs exhibit far less positive/negative bias in binary feedback and positional bias in preference feedback. Notably, reasoning models achieve strong performance across all of our discrete domains. With the rapid recent progress of reasoning-focused models such as Kimi K2 Thinking, Gemini 3, and GPT-5.1, there is an increasing risk of benchmark saturation. Indeed, our evaluation of GPT-5.1 (non-thinking) in Appendix Section I already shows performance comparable to that of reasoning models. However, our results also indicate that models still struggle in continuous-action domains such as Robomimic Lift, where our best preference provider (Qwen 2.5 VL 72B) still has an 18% error rate. Moreover, in safety-critical settings such as robotic planning (ALFWorld) and robotic manipulation (Robomimic), error rates remain non-negligible (e.g., 17% for action advising in ALFWorld), and thus cannot be ignored.

Preference and Binary Feedback are easier to provide   Preference and Binary Feedback are generally the easiest forms of feedback to accurately provide across domains, especially for non-reasoning models. For instance, Llama 3.1 8B and 70B provide significantly more accurate binary feedback and preference feedback than action advising. In contrast, larger models such as QwQ and R1 are able to provide action advising with high quality similar to binary and preference feedback.

Image observation allows better feedback in continuous domains   VLM feedback results are shown in Table 4. For predominantly visual domains, VLMs do not necessarily outperform LLMs if one can verbalize the environment succinctly. For example, in the Minigrid Door Key, we observed that even one of our best models, Qwen 2.5VL 72B shows worse performance than the model with similar sizes (Llama 3 70B) in text feedback. However, for Robomimic, in which it is difficult to provide succinct descriptions of states and actions, VLMs give preference up to 82% accuracy and binary feedback up to 62% accuracy, outperforming even bigger LLMs (see Table 2). Here, text observation requires a deep understanding of mechanical dynamics, while image observation is more intuitive to understand and hence gives higher accuracy.

Feedback quality degrades significantly in continuous state/action spaces   Table 2 shows that in the Robomimic continuous control domain, no model provides accurate feedback. [1] We observe similar patterns for VLMs, with the exception of preference feedback. We also experiment with egocentric verbalization, which has been shown to improve LLM planning in embodied environments (Shridhar et al., 2021; Schumann et al., 2024). Our results in Appendix A show that this does not significantly improve performance of any FMs. Our results suggest current foundation models struggle with interpreting continuous state and action spaces without human prior knowledge.

ICL and Thinking Guides respectively help bigger and reasoning models   Tables 3 show performance gains from ICL. Results for Thinking guides can be found in Appendix B. Smaller FMs generally do not benefit from either technique and are occasionally confused by the additional complexities in the prompt. For instance, Llama 3.1 8B degrades by ∼0.1 across feedback types in Door Key with ICL, and by as much as 0.21 in HierarchyCraft with Thinking Guides. In contrast, for larger FMs, ICL gains tend to be consistently positive, with negative changes being very limited. For instance, in HeirarchyCraft, ICL improves action advising of R1 by 0.1 and preference of Llama 3.1 70B by 0.42. Thinking Guides consistently improve the performance of reasoning FMs such as QwQ and R1, though by less than ICL. On non-reasoning models, its gains are inconsistent across feedback types. For VLMs, neither ICL nor Thinking Guides consistently improves performance, apart from ICL binary feedback for small VLMs in Qwen 2.5 VL 7B and 32B.

Access to environment dynamics hints can improve feedback accuracy.   Table 5 shows FM performance given different hints about high-level domain transition dynamics. Full descriptions of dynamics generally improve feedback quality, especially for small domains such as CliffWalking and Door Key. For complex domains such as HierarchyCraft, long prompts that describe complex dynamics degrade the performance of small FMs, but consistently improve the performance of large FMs. For instance, the 8B Llama feedback worsens across feedback types while the 70B model consistently improves, by as much as 0.24 for preference feedback. Similar trends hold when FMs induce dynamics from history context: large models benefit from context while small models do not. A more significant improvement is observed in VLM feedback. Allowing access to history context allows even up to 22% increase in action advising accuracy in Minigrid Doorkey and 44% increase on binary feedback accuracy in Robomimic. In Door Key, feedback accuracy with history context and unknown dynamics shows similar performance with fully known dynamics, suggesting that the VLM is able to infer dynamics from the history context.

Analysis of wrong responses   We identify four main types of feedback errors as **format**, **insufficient knowledge**, **incorrect reasoning**, and **incorrect visual understanding** and show detailed examples in Appendix H. **Format** errors overwhelmingly occur with small FMs, which are prone to output invalid JSON and invent answers not in the given options (e.g `BOTH` during preference feedback). This also happens to larger reasoning models like DeepSeek R1. Surprisingly, we observe that reasoning models occasionally overthink, generating excessively long chains-of-thought that

---

[1]OLAF (Liu et al., 2023) similarly extract language feedback for Robomimic, however unlike in OLAF, humans do not facilitate querying feedback in our experiments.

Domain	Preference		Action Adv		Bin Feedback		Goal Adv	
Cliff Walking Unk	0.49	0.53	0.19	0.21	0.19	0.39	0.11	0.22
	0.71	0.85	0.38	0.58	0.41	0.74	0.26	0.52
Full	0.51 ↑.02	0.77 ↑.24	0.20 ↑.01	0.44 ↑.23	0.47 ↑.28	0.39 =	0.15 ↑.04	0.69 ↑.45
	0.95 ↑.24	1.00 ↑.15	1.00 ↑.62	1.00 ↑.42	1.00 ↑.59	0.98 ↑.24	1.00 ↑.74	1.00 ↑.48
History	0.48 ↓.01	0.53 =	0.19 =	0.20 ↓.01	0.25 ↑.06	0.41 ↑.02	0.25 ↑.14	0.41 ↑.19
	0.70 ↓.01	0.85 =	0.47 ↑.06	0.68 ↓.06	0.42 ↑.01	0.80 ↑.06	0.42 ↑.16	0.84 ↑.32
Door Key Unk	0.53	0.76	0.41	0.51	0.79	0.75	N/A	
	0.91	0.92	0.82	0.87	0.92	0.94		
Full	0.66 ↑.13	0.75 ↓.01	0.44 ↑.03	0.44 ↓.07	0.78 ↓.01	0.79 ↑.04	N/A	
	0.89 ↓.02	0.98 ↑.06	0.81 ↓.01	0.99 ↑.12	0.92 =	1.00 ↑.06		
History	0.54 ↑.01	0.75 ↓.01	0.40 ↓.04	0.50 ↓.01	0.79 =	0.75 =	N/A	
	0.89 ↓.02	0.91 ↓.01	0.78 ↓.02	0.86 ↓.01	0.90 ↓.02	0.93 ↓.01		
HierachyCraft Unk	0.51	0.51	0.58	0.61	0.81	0.85	N/A	
	0.81	0.89	0.65	0.78	0.88	0.91		
Full	0.49 ↓.02	0.83 ↑.32	0.30 ↓.28	0.65 ↑.04	0.79 ↓.02	0.82 ↓.03	N/A	
	0.90 ↑.09	0.87 ↓.02	0.78 ↑.13	0.80 ↑.02	0.92 ↑.04	0.92 ↑.01		
Legal Action	0.51 =	0.80 ↑.29	0.46 ↓.12	0.60 ↓.01	0.81 =	0.83 ↓.02	N/A	
	0.87 ↑.06	0.92 ↑.03	0.65 =	0.80 ↑.02	0.88 =	0.91 =		
History	0.52 ↑.01	0.66 ↑.15	0.45 ↓.12	0.64 ↑.03	0.78 ↓.03	0.84 ↓.01	N/A	
	0.90 ↑.09	0.91 ↑.02	0.84 ↑.19	0.88 ↑.10	0.91 ↑.03	0.92 ↑.01		
Robomimic	0.55	0.67	0.01	0.00	0.47	0.50	0.07	0.12
	0.55	0.56	0.00	0.00	0.51	0.46	0.03	0.07
History	0.42 ↓.13	0.53 ↓.14	0.00 ↓.01	0.00 =	0.46 ↓.01	0.60 ↑.10	0.08 ↓.01	0.08 ↓.04
	0.51 ↓.04	0.67 ↑.11	0.00 =	0.00 =	0.48 ↓.03	0.50 ↑.04	0.00 ↓.03	0.05 ↓.02

Table 5: Feedback accuracy with the FM not given domain dynamics (Unk), given fully domain dynamics via text description (Full), given domain dynamics only for legal actions in the current state (Legal Action), and inferring dynamics from history context (History). For Robomimic, the last column is Delta action. For CliffWalking, the last column is Goal Advising.

violate the output format. **Insufficient knowledge** errors occur when FMs lack the world knowledge necessary to deduce the correct judgment. This type of error occurs most in complex domains such as HierarchyCraft, where FMs express incorrect assumptions about the domain, and reason (correctly) towards the wrong conclusion. A less frequent source of such errors is from powerful FMs that make extraneous assumptions about the domain. For example, an action is preferable in Cliff Walking if it reduces the number of steps to the goal, yet R1 sometimes misclassifies optimal actions because R1 deems it important to maintain a long distance from the cliff. Giving the FMs hints about domain dynamics reduces the frequency of these errors. **Incorrect reasoning** errors occur when the FM demonstrate a correct understanding of domain dynamics, but incorrectly reasons towards a wrong conclusion. This is common in domains with complex state spaces such as continuous sensors in Robomic, where FMs tend to correctly state domain dynamics but incorrectly calculate state transitions. For instance, the FM may calculate the incorrect end effector location in Robomic after taking a delta action despite knowing the formula, or picking up unnecessary items in HierarchyCraft despite knowing the correct recipes. **Incorrect observation understanding** happens when the vision language model cannot correctly understand the image observation. For example, in Door Key, weaker VLMs frequently misinterpret the orientation of the agent and the location of walls. This can be improved by adding reference images in the prompt as examples shown in the Appendix Section F. Furthermore, we observed that for MiniGrid domains, where the observation is represented with an ASCII art of the grid, the model often struggles to understand the spatial layout and misunderstands the position.

4 RELATED WORK

Learning from feedback Learning from feedback pre-dates LLMs. Ross et al. (2011b) directly uses expert demonstrations to learn policies. Ho & Ermon (2016) induce policies from offline expert

behaviour, without interacting with experts during training. Torabi et al. (2018) induce policies from observations in expert trajectories, without annotated actions. Instead of direct policy induction, expert behaviour and preferences can also be used to infer reward functions (Ng & Russell, 2000; Ouyang et al., 2022; Lee et al., 2021; Metcalf et al., 2023). Vision language models have also been used to provide preference for RL (Wang et al., 2024b; Tu et al., 2025). Compared to prior work, we evaluate LM feedback without expert policies and without annotated demonstrations. Extracting high quality automatic feedback significantly reduces the required human labour for training.

Learning from AI feedback   Instead of extracting feedback from domain experts, recent work extracts feedback from domain-agnostic pretrained foundation models. Bai et al. (2022) outline a framework for alignment using AI feedback in pretraining and fine-tuning of LLMs. Du et al. (2023) use LLMs to generate intermediate goals for RL. Xie et al. (2023) and Wang et al. (2024a) generate reward functions respectively using LLMs and vision-language models. Huang et al. (2024) uses image observations to query VLMs for robotic manipulation constraints. Colas et al. (2023) use LLMs to retroactive relabel actions. Zhong et al. (2024) use LLMs to identify desirable behaviours for imitation (Zhong et al., 2024). Several works use LLMs to generate preference data for exploration (Klissarov et al., 2023) and policy search (Lee et al., 2023; Liu et al., 2024) or trajectory judgment (Guan et al., 2024). Sharma et al. (2024) critically evaluate this last family of preference approaches, noting that gains largely result from combining superior teachers and weaker students. Our work is a unified, comprehensive study of AI feedback approaches, specifically evaluating the limitations of current state-of-the-art LLMs for providing planning feedback in grounded environments. We consider the primary types AI feedback in prior works, and evaluate these techniques on a shared, diverse set of domains across model classes and inference methods.

LLMs for planning   LLMs capture world knowledge helpful for planning (Brown et al., 2020; Liu et al., 2023; Stamper et al., 2024; Wu et al., 2023; DeepSeek-AI et al., 2025). Wang et al. (2023) use LLMs as policy for learning Minecraft. Ahn et al. (2022) and Vemprala et al. (2024) use LLMs to compose executable code using predefined APIs for robotic control. These are but a small representative sample recent works that use LLMs as a policy. LLMs can also be used to provide action priors (Zhang et al., 2024), subgoals (Du et al., 2023), and reward functions (Xie et al., 2023) for reinforcement learning. These methods for using LLMs in planning are two specific types of feedback considered in our study. The action advising feedback is identical to using LLM as policy. The goal advising feedback is similar to using LLM for next-goal generation or value estimation. Beyond action and goal advising, we evaluates how well LLMs can provide other types of feedback under diverse inference methods and environment settings.

5   CONCLUSION AND FUTURE WORK

We conduct a comprehensive study on evaluating language model feedback across feedback types, environment domains, and inference methods. We find that feedback quality improves with larger models and reasoning models, with large reasoning models providing diverse feedback with 80+% accuracy for most domains and showing reduced class and positional bias. Inference-time techniques such as in-context learning, chain-of-thought, and dynamics hints improve the performance of large reasoning models, however the complex prompt injections they require degrade performance of smaller or non-reasoning models. For domains such as continuous control, where it is challenging to provide succinct textual descriptions of state and actions, VLMs can provide higher quality feedback than LLMs. Given the evidence of high-accuracy LLM feedback in several domains, we encourage future research in how to leverage such feedback to systematically improve planning sample efficiency and quality, as well as improve reliability in safety-critical domains. In addition, future study is required in how to improve LLM feedback in domains with complex dynamics and with continuous state and action spaces — both common characteristics of challenging environments for which it is even more difficult to engineer rewards and label demonstrations. Lastly, as language model capabilities continue to improve rapidly (Li et al., 2025), we plan to extend this benchmark to include longer-horizon tasks such as BabyAI BossLevel (Hui et al., 2020) and more complex robotic manipulation tasks such as Robomimic Square (Mandlekar et al., 2021).

ETHICS STATEMENT

Though currently we only evaluate LLM feedback quality in grounded domains, the released feedback dataset can be used to train decision-making policies. This feedback dataset can be misused (e.g. training to reinforce incorrect feedback) to mislead policy training or create misleading feedback models, both of which may result in harmful model behaviour if deployed. We consider this work as a first step towards understanding LLM feedback, such that we can quantify limitations and mitigate risks when using such feedback to train models with less human annotation and fewer labelled demonstrations.

REFERENCES

Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario Jauregui Ruano, Kyle Jeffrey, Sally Jesmonth, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee, Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander Toshev, Vincent Vanhoucke, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Mengyuan Yan, and Andy Zeng. Do as i can and not as i say: Grounding language in robotic affordances. In *arXiv preprint arXiv:2204.01691*, 2022.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. Constitutional AI: Harmlessness from AI Feedback. *arXiv [cs.CL]*, 2022.

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.

Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W Cohen. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *Transactions in Machine Learning*, 2023.

Wentong Chen, Yankai Lin, ZhenHao Zhou, HongYun Huang, Yantao Jia, Zhao Cao, and Ji-Rong Wen. ICLeval: evaluating in-context learning ability of large language models. In *COLING*, 2025.

Maxime Chevalier-Boisvert, Bolun Dai, Mark Towers, Rodrigo de Lazcano, Lucas Willems, Salem Lahlou, Suman Pal, Pablo Samuel Castro, and Jordan Terry. Minigrid & miniworld: Modular & customizable reinforcement learning environments for goal-oriented tasks. In *NeurIPS*, 2023.

Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.

Cédric Colas, Laetitia Teodorescu, Pierre-Yves Oudeyer, Xingdi Yuan, and Marc-Alexandre Côté. Augmenting autotelic agents with large language models. In *CoLLAs*, 2023.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025.

Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Tianyu Liu, et al. A survey on in-context learning. In *ACL*, 2024.

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. In *International Conference on Machine Learning*, 2023.

Mathis Federico. HierarchyCraft, 2023.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi,

Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vítor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ra-

maswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The llama 3 herd of models, 2024.

Lin Guan, Yifan Zhou, Denis Liu, Yantian Zha, Heni Ben Amor, and Subbarao Kambhampati. Task success is not enough: Investigating the use of video-language models as behavior critics for catching undesirable agent behaviors. *arXiv preprint arXiv:2402.04210*, 2024.

Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models, 2023. *CoRR*, abs/2301.04104, 2023.

Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in neural information processing systems*, 2016.

Wenlong Huang, Chen Wang, Yunzhu Li, Ruohan Zhang, and Li Fei-Fei. Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation. *arXiv preprint arXiv:2409.01652*, 2024.

David Yu-Tung Hui, Maxime Chevalier-Boisvert, Dzmitry Bahdanau, and Yoshua Bengio. Babyai 1.1, 2020.

Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H Campbell, Konrad Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, et al. Model-based reinforcement learning for atari. In *ICLR*, 2020.

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 [cs.AI]*, 2023.

Minae Kwon, Sang Michael Xie, Kalesha Bullard, and Dorsa Sadigh. Reward design with language models. In *International Conference on Learning Representations*, 2023.

Aqeel Labash, Jaan Aru, Tambet Matiisen, Ardi Tampuu, and Raul Vicente. Perspective taking in deep reinforcement learning agents. *Frontiers in Computational Neuroscience*, 14:69, 2020.

Harrison Lee, Samrat Phatale, Hassan Mansoor, Thomas Mesnard, Johan Ferret, Kellie Lu, Colton Bishop, Ethan Hall, Victor Carbune, Abhinav Rastogi, and Sushant Prakash. RLAIF vs. RLHF: Scaling reinforcement learning from human feedback with AI feedback. *arXiv [cs.CL]*, 2023.

Kimin Lee, Laura Smith, and Pieter Abbeel. Pebble: Feedback-efficient interactive reinforcement learning via relabeling experience and unsupervised pre-training. In *International Conference on Machine Learning*, 2021.

Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, et al. From system 1 to system 2: A survey of reasoning large language models. *arXiv preprint arXiv:2502.17419*, 2025.

Aiwei Liu, Haoping Bai, Zhiyun Lu, Xiang Kong, Simon Wang, Jiulong Shan, Meng Cao, and Lijie Wen. Direct large language model alignment through self-rewarding contrastive prompt distillation. In *Annual Meeting of the Association for Computational Linguistics*, 2024.

Huihan Liu, Alice Chen, Yuke Zhu, Adith Swaminathan, Andrey Kolobov, and Ching-An Cheng. Interactive Robot Learning from Verbal Correction, 2023.

Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. In *CoRL*, 2021.

Katherine Metcalf, Miguel Sarabia, Natalie Mackraz, and Barry-John Theobald. Sample-Efficient Preference-based Reinforcement Learning with Dynamics Aware Rewards. In *CoRL*, 2023.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin A. Riedmiller. Playing atari with deep reinforcement learning. *CoRR*, abs/1312.5602, 2013.

Andrew Y Ng and Stuart J Russell. Algorithms for inverse reinforcement learning. In *International Conference on Machine Learning*, 2000.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. *CoRR*, abs/2203.02155, 2022.

Rodrigo Pérez-Dattari, Carlos Celemin, Javier Ruiz-del Solar, and Jens Kober. Interactive learning with corrective feedback for policies based on deep neural networks. In *Proceedings of the 2018 International Symposium on Experimental Robotics*, pp. 353–363. Springer, 2020.

Qwen Team. Qwq: Reflect deeply on the boundaries of the unknown, 2024.

Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 627–635. JMLR Workshop and Conference Proceedings, 2011a.

Stephane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. In *International Conference on Artificial Intelligence and Statistics*, 2011b.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017.

Raphael Schumann, Wanrong Zhu, Weixi Feng, Tsu-Jui Fu, Stefan Riezler, and William Yang Wang. VELMA: Verbalization embodiment of LLM agents for Vision and Language Navigation in Street View. In *AAAI*, 2024.

Archit Sharma, Sedrick Keh, Eric Mitchell, Chelsea Finn, Kushal Arora, and Thomas Kollar. A critical evaluation of AI feedback for aligning large language models. In *Advances in Neural Information Processing Systems*, 2024.

Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Côté, Yonatan Bisk, Adam Trischler, and Matthew Hausknecht. ALFWorld: Aligning Text and Embodied Environments for Interactive Learning. In *Proceedings of the International Conference on Learning Representations*, 2021.

John Stamper, Ruiwei Xiao, and Xinying Hou. Enhancing LLM-Based Feedback: Insights from Intelligent Tutoring Systems and the Learning Sciences. In *International Conference on Artificial Intelligence in Education*, 2024.

Simon Stepputtis, Joseph Campbell, Mariano Phielipp, Stefan Lee, Chitta Baral, and Heni Ben Amor. Language-conditioned imitation learning for robot manipulation tasks. In *Advances in Neural Information Processing Systems*, 2020.

Richard S Sutton. *Reinforcement learning: An introduction*. Bradford, 2018.

Andrew Szot, Max Schwarzer, Harsh Agrawal, Bogdan Mazoure, Rin Metcalf, Walter Talbott, Natalie Mackraz, R Devon Hjelm, and Alexander T Toshev. Large language models as generalizable policies for embodied tasks. In *The Twelfth International Conference on Learning Representations*, 2023.

Faraz Torabi, Garrett Warnell, and Peter Stone. Generative adversarial imitation from observation. *arXiv preprint arXiv:1807.06158*, 2018.

Songjun Tu, Jingbo Sun, Qichao Zhang, Xiangyuan Lan, and Dongbin Zhao. Online preference-based reinforcement learning with self-augmented feedback from large language model. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '25, pp. 2069–2077, Richland, SC, 2025. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400714269.

Sai H. Vemprala, Rogerio Bonatti, Arthur Bucker, and Ashish Kapoor. ChatGPT for Robotics: Design Principles and Model Abilities. *IEEE Access*, 12:55682–55696, 2024. ISSN 2169-3536.

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, 2023.

Yufei Wang, Zhanyi Sun, Jesse Zhang, Zhou Xian, Erdem Biyik, David Held, and Zackory Erickson. RL-VLM-F: Reinforcement Learning from Vision Language Foundation Model Feedback. In *International Conference on Machine Learning*, 2024a.

Yufei Wang, Zhanyi Sun, Jesse Zhang, Zhou Xian, Erdem Biyik, David Held, and Zackory Erickson. Rl-vlm-f: Reinforcement learning from vision language foundation model feedback. In *Proceedings of the 41th International Conference on Machine Learning*, 2024b.

Garrett Warnell, Nicholas Waytowich, Vernon Lawhern, and Peter Stone. Deep tamer: Interactive agent shaping in high-dimensional state spaces. In *Proceedings of the AAAI conference on artificial intelligence*, 2018.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.

Xiaoqian Wu, Yong-Lu Li, Jianhua Sun, and Cewu Lu. Symbol-LLM: Leverage Language Models for Symbolic System in Visual Human Activity Reasoning. In *NeurIPS*, 2023.

Tianbao Xie, Siheng Zhao, Chen Henry Wu, Yitao Liu, Qian Luo, Victor Zhong, Yanchao Yang, and Tao Yu. Text2reward: Automated dense reward function generation for reinforcement learning. In *ICLR*, 2023.

Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, et al. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments. In *NeurIPS*, 2024.

Qisen Yang, Shenzhi Wang, Matthieu Gaetan Lin, Shiji Song, and Gao Huang. Boosting offline reinforcement learning with action preference query. In *International Conference on Machine Learning*, pp. 39509–39523. PMLR, 2023.

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822, 2023.

Shenao Zhang, Sirui Zheng, Shuqi Ke, Zhihan Liu, Wanxin Jin, Jianbo Yuan, Yingxiang Yang, Hongxia Yang, and Zhaoran Wang. How Can LLM Guide RL? A Value-Based Approach. *CoRR*, abs/2402.16181, 2024.

Qi Zhao, Haotian Fu, Chen Sun, and George Konidaris. EPO: Hierarchical LLM agents with environment preference optimization. In *EMNLP*, 2024.

Victor Zhong, Dipendra Misra, Xingdi Yuan, and Marc-Alexandre Côté. Policy Improvement using Language Feedback Models. In *Advances in Neural Information Processing Systems*, 2024.

A   EGOCENTRIC OBSERVATION

Labash et al. (2020) show that egocentric observation ability as well as perspective-taking ability is important for policy learning (Labash et al., 2020). For LLMs specifically, Szot et al. (2023) and Zhong et al. (2024) show that egocentric verbalization improves the performance of LLM agents in embodied planning. Table 6 shows that unlike prior results for planning, egocentric verbalization generally degrades feedback performance across feedback types and across most domains.

Domain	Preference		Action Adv		Bin Feedback		Delta Act	
Cliff Walking	0.50 ↑0.05	0.55 ↓0.08	0.44 ↓0.14	0.49 ↓0.14	0.50 ↓0.04	0.50 ↓0.09	N/A	
	0.95 ↑.04	1.00 =	1.00 =	1.00 =	0.88 ↓.12	0.89 ↓.11		
Door Key	0.53 ↓.08	0.72 ↓.03	0.52 ↑.08	0.30 ↓.13	0.80 ↑.02	0.72 ↓.03	N/A	
	0.67 ↓.12	0.70 ↓.28	0.95 ↑.14	1.00 ↑.01	0.73 ↓.19	0.76 ↓.24		
Four Rooms	0.37 ↓.22	0.63 ↓.03	0.34 ↑.02	0.33 ↓.11	0.60 ↑.01	0.68 ↓.01	N/A	
	0.72 ↑.07	0.69 ↓.06	0.83 ↑.22	0.83 ↑.05	0.70 ↑.04	0.72 ↓0.09		
Robomimic	0.52 ↓.03	0.53 ↓.14	0.00 ↓.01	0.00 =	0.52 ↑.05	0.52 ↑.02	0.15 ↑.08	0.06 ↓.06
	0.47 ↓.08	0.45 ↓.11	0.00 =	0.00 =	0.49 ↓.03	0.52 ↑.06	0.06 ↑.03	0.07 =

Table 6: Feedback accuracy with egocentric observations.

Domain	Preference	Action Adv	Bin Feedback	Goal Adv
Cliff Walking	0.49 ↑.16 0.47 ↓.30
0.99 ↑.08 1.00 =	0.68 = 0.92 ↑.19
0.99 = 1.00 =	0.52 ↓0.09 0.50 ↓.23
1.00 = 0.93 ↓.05	0.57 ↑.20 0.92 ↑.15
1.00 = 0.97 ↓.03
Door Key	0.53 ↓.13 0.72 ↓.03
0.94 ↑.05 0.98 =	0.47 ↑.03 0.41 ↓.03
0.89 ↑.07 0.99 =	0.78 = 0.76 ↑.01
0.94 ↑.02 1.00 =	N/A
Four Rooms	0.53 ↓.03 0.58 ↓.08
0.71 ↑.06 0.81 ↑.06	0.47 ↑.15 0.43 ↑.01
0.59 ↓.02 0.75 ↓.03	0.63 ↑.04 0.62 ↓.07
0.71 ↑.05 0.81 =	N/A
ALFWorld	0.51 ↑.02 0.68 ↑.01
0.82 ↑.03 0.88 ↑.05	0.59 ↓.01 0.83 =
0.83 ↓.02 0.83 =	0.72 ↓.04 0.76 ↓.02
0.76 ↓.01 0.76 =	N/A
HierachyCraft	0.51 = 0.78 ↑.27
0.84 ↑.03 0.89 =	0.37 ↓.21 0.63 ↑.02
0.73 ↑.08 0.78 =	0.79 ↑.02 0.85 =
0.88 = 0.90 ↓.01	N/A

Table 7: Feedback accuracy with Thinking Guides, the arrows show accuracy change with this condition relative to baseline.

B EXPERIMENT RESULTS IN THINKING GUIDES

The results for the thinking guides are shown in the Table 7. It shows that most thinking guides performance gains are in thinking models and especially QwQ 32B, while in some domains, Deepseek R1, the accuracy is already saturated and is very close to one.

Domain	Preference		Action Adv		Bin Feedback		Delta Action	
Door Key	0.56	0.73	0.24	0.29	0.72	0.66	N/A	
	0.96	0.99	0.80	0.99	0.89	0.95		
Four Rooms	0.51	0.59	0.32	0.33	0.57	0.61	N/A	
	0.64	0.67	0.40	0.54	0.66	0.75		
Robomimic	0.53	0.59	0.00	0.00	0.54	0.59	0.07	0.09
	0.31	0.12	0.00	0.00	0.31	0.24	0.00	0.04
ALFWorld	0.50	0.62	0.52	0.52	0.50	0.62	N/A	
	0.71	0.77	0.52	0.52	0.71	0.77		
HierachyCraft	0.48	0.76	0.53	0.64	0.49	0.59	N/A	
	0.84	0.86	0.66	0.71	0.58	0.60		

Table 8: Feedback accuracy in random dataset for Llama 3.1 8B , Llama 3.1 70B , QwQ 32B , and R1 .

Domain	Preference		Action Adv		Bin Feedback		Delta Action	
Door Key	0.57	0.71	0.37	0.43	0.78	0.71	N/A	
	0.91	0.98	0.82	0.99	0.91	0.99		
Four Rooms	0.47	0.59	0.32	0.38	0.59	0.64	N/A	
	0.73	0.75	0.56	0.74	0.71	0.77		
Robomimic	0.60	0.63	0.00	0.00	0.48	0.51	0.06	0.13
	0.51	0.23	0.00	0.00	0.46	0.27	0.04	0.03
ALFWorld	0.52	0.71	0.49	0.64	0.77	0.80	N/A	
	0.75	0.85	0.78	0.77	0.83	0.83		
HierachyCraft	0.51	0.81	0.39	0.60	0.76	0.82	N/A	
	0.85	0.89	0.64	0.77	0.84	0.88		

Table 9: Feedback accuracy in half-expert dataset for Llama 3.1 8B , Llama 3.1 70B , QwQ 32B , and R1 .

C RESULTS UNDER DIFFERENT DISTRIBUTIONS

In this subsection, we present similar results acquired with datasets collected from other sampling policies (random policy and half-expert policy). The results under the half-expert dataset and the random dataset are shown in Table 9 and Table 8, and show similar trends to results under the expert dataset. We also observe random dataset results seem to have a larger variance since the dataset's state distribution often only covers a small part of the state space. For example, in Cliff Walking, a random policy has 75% of the chance to take any of the actions RIGHT, LEFT or DOWN which results in hitting the wall or falling off the cliff. Consequently, under the random policy, the agent almost never reaches the goal and mostly wanders around its start point. The image observation results on VLMs under the half-expert dataset and the random dataset are shown in Table 10 and Table 11. The results, similar to LLM feedback, are consistent with results in the expert dataset, while the random dataset poses more challenges, as it contains unusual states. For example, the action advising accuracy for Qwen 2.5 VL 72B dropped to 8%.

Domain	Preference		Action Adv		Bin Feedback		Delta Action	
Door Key	0.27	0.56	0.22	0.48	0.55	0.72	N/A	
	0.52	0.46	0.08	0.28	0.62	0.64		
Robomimic	0.92	0.76	0.00	0.00	0.38	0.17	0.02	0.08
	0.85	0.57	0.00	0.00	0.06	0.07	0.07	0.00

Table 10: Feedback accuracy in random dataset for Qwen 2.5 VL 7B , Qwen 2.5 VL 32B , Qwen 2.5 VL 72B , and o4-mini .

Domain	Preference		Action Adv		Bin Feedback		Delta Action	
Door Key	0.37	0.53	0.39	0.24	0.50	0.65	N/A	
	0.52	0.40	0.28	0.31	0.63	0.69		
Robomimic	0.85	0.81	0.00	0.00	0.42	0.24	0.01	0.09
	0.82	0.72	0.00	0.00	0.01	0.03	0.06	0.02

Table 11: Feedback accuracy in half-expert dataset for Qwen 2.5 VL 7B , Qwen 2.5 VL 32B , Qwen 2.5 VL 72B , and o4-mini .

Domain	F1			Precision			Recall		
Cliff Walking	0.44	0.88		1.00	1.00		0.28	0.33	
	0.66	0.93		1.00	1.00		0.42	0.79	
Door Key	0.12	0.51		0.47	0.46		0.07	0.57	
	0.83	1.00		0.89	0.98		0.90	1.00	
Four Rooms	0.28	0.56		0.37	0.55		0.22	0.56	
	0.53	0.71		0.49	0.66		0.57	0.76	
Robomimic	0.05	0.13		0.52	0.68		0.03	0.07	
	0.00	0.00		0.00	0.00		0.00	0.00	
ALFWorld	0.18	0.37		0.66	0.49		0.11	0.29	
	0.29	0.25		0.63	0.66		0.19	0.15	
HierachyCraft	0.31	0.47		0.34	0.37		0.28	0.65	
	0.55	0.63		0.51	0.59		0.60	0.67	

Table 12: F1 scores in traverse dataset (Cliff Walking) or expert dataset (all other domains) for Llama 3.1 8B , Llama 3.1 70B , QwQ 32B , and R1 .

D POSITIVE VS NEGATIVE BIAS IN BINARY FEEDBACK

In this subsection, we evaluate Feedback Model binary feedback performance using F1 score in order to evaluate model performance in the event of class imbalance. This imbalance in all datasets, since each domain has an action space with more than two actions. Consequently, for binary feedback, a larger portion of examples contain non-optimal actions. Table 12 shows binary feedback F1 scores across domains. We see that non-reasoning models, particularly Llama 3.1 8b, often obtain significantly worse recall than precision across domains. In other words, non-reasoning models are strongly biased towards identifying actions as non-optimal.

Location of ground truth option	Llama 3.1 8B accuracy	Llama 3.1 70B accuracy
First	0.76	0.23
Second	0.69	0.64

Table 13: Positional bias of Llama 3.1 models on ALFWorld preference feedback using sample snapshots from expert policy

E POSITIONAL BIAS IN PAIRWISE PREFERENCES

In our dataset collection process for preference, we enumerate all possible combinations of legal actions at each state. Because of the sheer size of the dataset, we randomly downsample 1000 samples for experiments. We conduct an extra study to investigate biases in positional preferences by Feedback Model in providing preference feedback. Specifically, for two options A and B, where A is the ground truth preferable option, we ensure that both (A, B) and (B, A) are in the evaluation dataset. The results are shown in Table 13. We find that smaller models like Llama 3.1 8B significantly favour whichever option is presented first, consequently resulting in high accuracy when the correct option is first but low accuracy when the correction option is second. This bias is significantly reduced for bigger models, which are largely agnostic to the position of the correct option. In other words, smaller models exhibit stronger positional bias than larger models.

F   VERBALISATION AND PROMPT EXAMPLES

In this section, we provide prompt examples for selected domains and types of feedback using chain-of-thought prompting. We also provide prompt examples for different experimental conditions.

F.1   STANDARD PROMPTS

In this section, we provide prompt examples corresponding to Table 2, with no specific inference enhancement applied. One relevant experiment we ran for binary feedback specifically considers choices for how to describe optimal actions in the prompt. Our early results showed that asking "is this action good or bad" lacks precision and resulted identification of desirable behaviour that are not optimal as "good". Our final experiments instead asks the Feedback Model "is action ACTION the best action you can take?", which resulted in significantly fewer false positives.

Binary Feedback in MiniGrid Door Key

You need to solve Minigrid Doorkey.
Minigrid Doorkey involves crossing a gridworld, where you need to pick up the key, unlock the door and reach the goal. Do not hit the wall.
The state is described with an ASCII art. In the ASCII art, "#" means wall. "K" means key. "D" mean door. "G" means goal. "." means empty tile. "^" mean the agent who is facing upwards. "V" mean the agent who is facing downwards. "<" means the agent who is facing leftwards. ">" means the agent who is facing rightwards.
The possible actions are:
1. **TURN LEFT**, changing the direction to left, for example, the agent "^" will become "<".
2. **TURN RIGHT**, changing the direction to right, for example, the agent "^" will become ">".
3. **MOVE FORWARD**, moving forwards for one tile
4. **PICK UP THE KEY**, picking up the key, this is only possible when you are facing the tile with key on it
5. **UNLOCK THE DOOR**, this action is only possible when you have the key.
    You will be asked to judge an action whether it is the optimal action, given a state.
You are in the following state represented in ASCII art:

```
#####
#.#.#
#>D.#
#.#G#
#####
```

The door is locked. You have the key.
Is action **UNLOCK THE DOOR** the best action you can take? Please think step by step.
Only give the answer in a new line in JSON format:
`{"reasoning":  <REASONING>, "feedback":  <FEEDBACK>}`
Where <FEEDBACK> is one of "YES" or "NO", <REASONING> is a string of your thinking steps.

Preference in ALF World

You received instruction: Put a chilled potato in the microwave... Your observation is: You open the fridge 1. The fridge 1 is open. In it, you see a egg 1, a lettuce 2, and a pan 1.
Below is your history:
You received instruction: Put a chilled potato in the microwave. Your observation is: You are in the middle of a room. Looking quickly around you, you see a cabinet 6, a cabinet 5, a cabinet 4, a cabinet 3, a cabinet 2, a cabinet 1, a coffeemachine 1, a countertop 3, a countertop 2, a countertop 1, a drawer 3, a drawer 2, a drawer 1, a fridge 1, a garbagecan 1, a microwave 1, a shelf 3, a shelf 2, a shelf 1, a sinkbasin 1, a stoveburner 4, a stoveburner 3, a stoveburner 2, a stoveburner 1, and a toaster 1. . You took action: look.
You received instruction: Put a chilled potato in the microwave. Your observation is: You are in the middle of a room. Looking quickly around you, you see nothing.. You took action: go to sinkbasin 1.
You received instruction: Put a chilled potato in the microwave. Your observation is: You arrive at loc 29. On the sinkbasin 1, you see a butterknife 3, a cup 1, a egg 2, a knife 1, a spatula 2, and a tomato 1.. You took action: go to microwave 1.
You received instruction: Put a chilled potato in the microwave. Your observation is: You arrive at loc 17. The microwave 1 is closed.. You took action: open microwave 1.
You received instruction: Put a chilled potato in the microwave. Your observation is: You open the microwave 1. The microwave 1 is open. In it, you see a apple 1, and a mug 1.. You took action: close microwave 1.
You received instruction: Put a chilled potato in the microwave. Your observation is: You close the microwave 1.. You took action: go to garbagecan 1.
You received instruction: Put a chilled potato in the microwave. Your observation is: You arrive at loc 31. On the garbagecan 1, you see a apple 2, and a glassbottle 2.. You took action: go to fridge 1.
You received instruction: Put a chilled potato in the microwave. Your observation is: You arrive at loc 30. The fridge 1 is closed. You took action: open fridge 1.
You can take the following actions:
close fridge 1
examine fridge 1
go to cabinet 1
go to cabinet 2
go to cabinet 3
go to cabinet 4
go to cabinet 5
go to cabinet 6
go to coffeemachine 1
go to countertop 1
go to countertop 2
go to countertop 3
go to drawer 1
go to drawer 2
go to drawer 3
go to garbagecan 1
go to microwave 1
go to shelf 1
go to shelf 2
go to shelf 3
go to sinkbasin 1
<Some possible action omitted here for ease of read>
Given go to countertop 2 or close fridge 1, which action is better? Please think step by step. Only give the answer in a new line in JSON format: {"reasoning": <REASONING>, "preference": <PREFERENCE>} Where <PREFERENCE> is one of "FIRST" or "SECOND", <REASONING> is a string of your thinking steps.

Action Advising in Hierachy Craft

You are playing MineCraft and you want to craft diamond.
You are in forest.
Items in your current zone are:
crafting table: 1
furnace: 1
You have the following items in your inventory:
wood plank: 3
wood pickaxe: 7
stone pickaxe: 12
iron ore: 3
You can take these actions:
pickup crafting table
pickup furnace
smelt iron ore with wood plank
craft stick
search for dirt at forest
search for wood at forest
search for cobblestone with stone pickaxe at forest
search for cobblestone with wood pickaxe at forest
move to swamp from forest
move to meadow from forest
move to underground from forest with stone pickaxe
Which action do you choose? Please think step by step.
Only give the answer in a new line in JSON format: {"reasoning": <REASONING>, "action": <ACTION>} Where <ACTION> is one of possible actions, <REASONING> is a string of your thinking steps.

Delta Action in Robomimic

You are a robot engineer. You need to control a robot arm to lift an object. You need to observe the robot's current state and choose which action to do next in order to lift the object. You can move the robot in the X, Y, or Z direction, as well as open/close its gripper. The robot's state is given in the following format:

```
Gripper position: [x y z ]
Object position: [x y z]
Object relative distance to gripper: [dx dy dz]
Gripper finger state: open or closed
```

You can choose actions that move the gripper in the X, Y, Z direction, with additional control of the gripper. The format of the action is:

```
[resulting x, resulting y, resulting z, ignore, ignore, ignore,
resulting gripper state]
```

Note on the gripper: Set grip to 1 to open the gripper and -1 to close the gripper. The robot's gripper should be closed if it is beginning to grasp the object, or when it is holding the object. When it is approaching the object, the gripper is open. If the robot gripper needs to be closed, you should continue to close the gripper, even if it is closed. Similarly, if the robot gripper needs to be open, you should continue to open the gripper, even if it is already open.

You current observation is:
Gripper position: [-0.01261643 0.02076922 0.82920151]
Object position: [-0.02050141 0.02279622 0.82019789]
Object relative position to gripper: [0.00788497 -0.002027 0.00900362]
Gripper finger position: [0.03368153 -0.03366622]
Gripper status: +1
Given the following actions, which action is better? Please think step by step.
Index 0: [0.25368214 0.02193805 -0.23200776 0.00478249 0.25400147 0.01591277 1.]
Index 1: [-0.30631787 0.02193805 -0.23200776 0.00478249 0.25400147 0.01591277 1.]
Index 2: [-0.02631785 0.30193806 -0.23200776 0.00478249 0.25400147 0.01591277 1.]
Index 3: [-0.02631785 -0.25806195 -0.23200776 0.00478249 0.25400147 0.01591277 1.]
Index 4: [-0.02631785 0.02193805 0.04799224 0.00478249 0.25400147 0.01591277 1.]
Index 5: [-0.02631785 0.02193805 -0.5120078 0.00478249 0.25400147 0.01591277 1.]
Index 12: [-0.02631785 0.02193805 -0.23200776 0.00478249 0.25400147 0.01591277 1.]
Index 13: [-0.02631785 0.02193805 -0.23200776 0.00478249 0.25400147 0.01591277 -1.]
You should consider the position of the robot gripper and object, and how they are related to each other. For example, if the robot gripper is on the left of the object, you should consider moving the robot gripper to the right.
Only give the answer in a new line in JSON format:
{"reasoning": <REASONING>, "index": <INDEX>}
Where <INDEX> is one of action choice index, <REASONING> is a string of your thinking steps.

Goal Advising in Cliff Walking

You are an agent trying to move to position (3, 11).
You can take 4 possible actions:

- going UP
- going DOWN
- going LEFT
- going RIGHT

The grid world is 4x12 with zero-indexed positions. There are holes on the grid that you do not know. Do not exceed the boundaries of the world, and do not fall into the holes. You will be asked to predict the best next state given a state.

You are in position (3, 0). What coordinates should you move to in the next step? Please think step by step.
Only give the answer in a new line in JSON format:
```
{"reasoning":  <REASONING>, "row":  <ROW>, "column":
<COLUMN>}
```
Where <ROW> is the number of the row you should move, and <COLUMN> is the number of column you should move, <REASONING> is a string of your thinking steps.

F.2 IN CONTEXT LEARNING PROMPT EXAMPLES

In this section, we provide prompt examples for our ICL experiments. We insert these examples after the description of the task, right before providing the current observation and querying for feedback.

Cliff Walking Goal Advising ICL Example

Below are examples about reasoning over the best next state.

Question: You are in position (3, 0). What coordinates should you move to in the next step?
Answer: {"reasoning": "There are 4 neighbouring coordinates:
(2, 0), (4, 0), (3, -1), (3, 1). If I move to (2, 0), it
is a safe place. If I move to (4, 0), I will hit the wall
as I am in a zero-indexed 4x12 grid and (4, 0) is outside
of boundary. If I move to (3, -1), I will similarly hit the
wall as (3, -1) is outside of boundary. If I move to (3, 1)
and since there are holes in (3, 1) ... (3, 10), I will fall
into a hole. Therefore moving to (2, 0) is the only safe
option.","row": "2", "column": "0"}

Question: You are in position (2, 11). What coordinates should you move to in the next step?
Answer: {"reasoning": "There are 4 neighbouring coordinates:
(1, 11), (3, 11), (2, 10) and (2, 12). If I move to (1, 11)
and then I will need 2 steps as (1, 11) -> (2, 11) -> (3, 11)
to reach the goal (3, 11). If I move to (3, 11) I will reach
the goal. If I move to (2, 10) and then I need 3 steps as
(2, 10) -> (3, 10) -> (3, 11) to reach the goal. If I move
to (2, 12) I will hit the wall as it is outside of boundary.
Therefore (3, 11) is the best option.","row": "3", "column":
"11"}

Question: You are in position (1, 5). What coordinates should you move to in the next step?
Answer: {"reasoning": "There are 4 possible action: (0, 5),
(2, 5), (1, 6), (1, 4). If I move to (0, 5) and I will need
9 steps to reach goal: (0, 5) -> (0, 6) -> (0, 7) -> (0, 8)
-> (0, 9) -> (0, 10) -> (0, 11) -> (1, 11) -> (2, 11) -> (3,
11). If I move to (2, 5) and I will need 7 steps to reach
goal: (2, 5) -> (2, 6) -> (2, 7) -> (2, 8) -> (2, 9) -> (2,
10) -> (2, 11) -> (3, 11). If I move to (1, 6) and I will
need 7 steps to reach goal: (1, 6) -> (1, 7) -> (1, 8) ->
(1, 9) -> (1, 10) -> (1, 11) -> (2, 11) -> (3, 11). If I
move to (1, 4) and I will need 9 steps to reach goal: (1, 4)
-> (1, 5) -> (1, 6) -> (1, 7) -> (1, 8) -> (1, 9) -> (1, 10)
-> (1, 11) -> (2, 11) -> (3, 11). Therefore, (1, 6) and (2,
5) are the best next coordinates as they give shortest path
to goal. I just need one good coordinates so I will pick one
from (1, 6) and (2, 5)","row": "2", "column": "5"}

Door Key Binary Feedback ICL

Answer: `{"reasoning": "The agent is facing a wall, so moving forward does nothing and fails to progress toward picking up the key. A turn to face the key is a better action.", "feedback": "NO"}`

Question: You are in the following state represented in ascii art:

```
#####
#.D.#
#^#.#
#.#G#
#####
```

The door is locked. You have the key.
Is action MOVE FORWARD the best action you can take? Please think step by step.
Only give the answer in a new line in JSON format:
`{"reasoning": <REASONING>, "feedback": <FEEDBACK>}`
Where <FEEDBACK> is one of "YES" or "NO", <REASONING> is a string of your thinking steps.

Answer: `{"reasoning": "You already have the key, so the immediate goal is to move adjacent to the locked door and unlock it. The tile directly ahead is free, and moving forward places you in position to turn and unlock the door. Therefore, moving forward is the best action.", "feedback": "YES"}`

Question: You are in the following state represented in ascii art:

```
#####
#.#.#
#.>.#
#.#G#
#####
```

The door is unlocked. You have the key.
Is action TURN LEFT the best action you can take? Please think step by step.
Only give the answer in a new line in JSON format:
`{"reasoning": <REASONING>, "feedback": <FEEDBACK>}`
Where <FEEDBACK> is one of "YES" or "NO", <REASONING> is a string of your thinking steps.

Answer: `{"reasoning": "The agent is facing right and needs to move toward the goal on its right and then turn down. Turning left would make the agent face upwards into a wall. Moving forward is the better move.", "feedback": "NO"}`

Action Advising ICL Four Rooms

Below are examples about reasoning over the best action.

Question:
You are in the following state represented in ascii art:

```
###################
#........#........#
#........#........#
#........#........#
#........#........#
#........#........#
#........#........#
#........#........#
#.................#
######.##########.#
#........#....G.<.#
#........#........#
#........#........#
#.................#
#........#........#
#........#........#
#........#........#
#........#........#
###################
```

Which action do you choose? Please think step by step.
Only give the answer in a new line in JSON format:
`{"reasoning": <REASONING>, "action": <ACTION>}`
Where <ACTION> is one of "TURN LEFT", "TURN RIGHT" or "MOVE FORWARD",
<REASONING> is a string of your thinking steps.

Answer: `{"reasoning": "The agent is already facing left and the goal is two tiles to its left, so moving forward brings it closer to the goal.", "action": "MOVE FORWARD"}`

HierachyCraft ICL in Preference

Question:
You are playing MineCraft and you want to craft diamond. You will be asked to compare two actions given a state and provide preference.

You are in forest.
Items in your current zone are:

- crafting table: 1

You have the following items in your inventory:

- wood plank: 4
- stick: 8

You can take these actions:

- pickup crafting table
- craft stick
- craft crafting table
- craft wood pickaxe
- craft wood axe
- craft wood shovel
- craft wood sword
- search for dirt at forest
- search for wood at forest
- move to swamp from forest
- move to meadow from forest

Given `craft crafting table` or `craft wood pickaxe`, which action is better?
Please think step by step.
Only give the answer in a new line in JSON format:
`{"reasoning":  <REASONING>, "preference":  <PREFERENCE>}`
Where <PREFERENCE> is one of "FIRST" or "SECOND", <REASONING> is a string of your thinking steps.

Answer: `{"reasoning":  "We already have access to a crafting table in the forest, so crafting another is not immediately useful.  We need a wood pickaxe to gather stone and progress toward diamonds.", "preference":  "SECOND"}`

ALF World ICL Binary Feedback Example

Question:
You received instruction: Place a chilled tomato in a microwave.
Your observation is: You arrive at loc 3. On the shelf 2, you see a creditcard 2.
Your inventory: tomato 1. Below is your history:

```
You received instruction: Place a chilled tomato in a microwave.
Your observation is: You are in the middle of a room. Looking quickly
around  you, you see a cabinet 6, a cabinet 5, a cabinet 4, a cabinet
3, a cabinet 2, a cabinet 1, a coffeemachine 1, a countertop 3, a
countertop 2, a countertop 1, a drawer 3, a drawer 2, a drawer 1, a
fridge 1, a garbagecan 1, a microwave 1, a shelf 3, a shelf 2, a
shelf 1, a sinkbasin 1, a stoveburner 4, a stoveburner 3, a
stoveburner 2, a stoveburner 1, and a toaster 1.
Your inventory: empty.
You took action: go to cabinet 1.

You received instruction: Place a chilled tomato in a microwave.
Your observation is: You arrive at loc 13. On the cabinet 1, you see
a bowl 1.
Your inventory: empty.
You took action: go to countertop 1.

<More History Steps Ommited for Ease of Reading>
```

You can take the following actions:
```
examine shelf 2
examine tomato 1
go to cabinet 1
go to cabinet 2
go to cabinet 3
go to cabinet 4
go to cabinet 5
go to cabinet 6
go to coffeemachine 1
go to countertop 1
go to countertop 2
go to countertop 3
go to drawer 1
go to drawer 2
go to drawer 3
go to fridge 1
go to garbagecan 1
go to microwave 1
go to shelf 1
go to shelf 3
go to sinkbasin 1
go to stoveburner 1
go to stoveburner 2
go to stoveburner 3
go to stoveburner 4
go to toaster 1
inventory
look
```

Is action `go to cabinet 3` the best action you can take? Please think step by step.
Only give the answer in a new line in JSON format:
```
{"reasoning":  <REASONING>, "feedback":  <FEEDBACK>}
```
Where <FEEDBACK> is one of "YES" or "NO", <REASONING> is a string of your thinking steps.

Answer: `{"reasoning":  "We already have the chilled tomato in our inventory and our goal is to place it in the microwave.  The next logical step is to go back to the microwave and open it, then put the tomato inside.  Going to cabinet 3 does not help achieve this goal.", "feedback":  "NO"}`

F.3 THINKING GUIDES PROMPT EXAMPLES

In this subsection, we show the example prompts of thinking guides. We insert these thinking guides inside standard prompts, right after we describe the domain and before we provide the current observation and query for feedback.

Thinking Guides for Door Key and Four Rooms

You should think about these questions.
What is your current goal?
Based on the rules you know, what subgoals do you need to achieve?
Based on what you see and what you have, what subgoals have you achieved?
If you have not achieved all of the subgoals, based on the rules and doable actions, can you achieve any subgoals now?
If yes, which subgoal do you want to achieve next?
If no, what do you need or where should you go next?

Thinking Guides for HierachyCraft

You should think about these questions.
You need to get a diamond.
Based on the rules you know, what subgoals do you need to achieve?
Based on what you see and what you have, what subgoals have you achieved?
If you have not achieved all of the subgoals, based on the rules and doable actions, can you achieve any subgoals now?
If yes, which subgoal do you want to achieve next?
If no, what do you need or where should you go next?

Thinking Guides for Cliff Walking

You should think about these questions.
What is your current goal?
Based on the rules you know, which coordinates do you need to reach the goal?
Based on what you see and what you known, can you reach any of the coordinates now?
If no, comparing the possible actions you can take, which one leads you closer to the goal?

Thinking Guides for ALF World

You should think about these questions.
What is your current goal?
Based on the rules you know, what subgoals do you need to achieve?
Based on what you see and what you have, what subgoals have you achieved?
If you have not achieved all of the subgoals, based on the rules and doable actions, can you achieve any subgoals now?
If yes, which subgoal do you want to achieve next?
If no, what do you need or where should you go next?

F.4 DYNAMICS HINT PROMPT EXAMPLES

In this subsection, we show prompt examples of experiments with/without dynamics hints. For unknown dynamics, we alter the task description at the beginning of the prompt and remove anything related to dynamics. For fully known dynamics, we append the rules of the game right after we describe the task before we provide states and query for feedback. For history context, we append the history steps at the end of the prompt, right after we provide the current observation and query for feedback.

Cliff Walking Unknown Dynamics

You are an agent trying to move to position (3, 11).
You can take 4 possible actions:

```
- going UP
- going DOWN
- going LEFT
- going RIGHT
```

The grid world is 4x12 with zero-indexed positions. There are holes on the grid that you do not know. Do not exceed the boundaries of the world, and do not fall into the holes.

Cliff Walking History Context Example

Here is the history of your actions:
Step 1: You were at position (3, 0), and you took action UP, then you reached (2, 0).
Step 2: You were at position (2, 0), and you took action DOWN, then you reached (3, 0).
Step 3: You were at position (3, 0), and you took action RIGHT, then you reached (3, 0).
You fell into the holes.

Door Key Unknown Dynamics Example

You need to solve Minigrid Doorkey. Minigrid Doorkey involves crossing a gridworld, where you need to pick up the key, unlock the door and reach the goal. Do not hit the wall. The state is described with an ASCII art. In the ASCII art, "#" means wall. "K" means key. "D" means door. "G" means goal. "." means empty tile. "^" means the agent who is facing upwards. "V" means the agent who is facing downwards. "<" means the agent who is facing leftwards. ">" means the agent who is facing rightwards. The possible actions are:

```
    1. TURN LEFT
    2. TURN RIGHT
    3. MOVE FORWARD
    4. PICK UP THE KEY
    5. UNLOCK THE DOOR
```

Door Key History Context Example

Here is the history of your actions:
Step 1: You see:

```
#####
#K#.#
#.D.#
#<#G#
#####
```

You don't have the key and the door is locked. You took action `TURN RIGHT`.
Step 2: You see:

```
#####
#K#.#
#.D.#
#^#G#
#####
```

You don't have the key and the door is locked. You took action `MOVE FORWARD`.
Step 3: You see:

```
#####
#K#.#
#^D.#
#.#G#
#####
```

You don't have the key. The door is locked. You took action `PICK UP THE KEY`.
Step 4: You see:

```
#####
#.#.#
#^D.#
#.#G#
#####
```

You have the key. The door is locked. You took action `TURN RIGHT`.

HierachyCraft Full Dynamics Example

Below are the rules for the minecraft.

```
You need a crafting table to craft something.
You need a furnace to smelt something.
You can pick up a crafting table or a furnace and take them
around.
You need to place one in your current area if you want to use
them.
You can get 3 iron ingot after smelting with 3 iron ore, 2
wood plank.
You can get 3 gold ingot after smelting with 3 gold ore, 2
wood plank.
You can get 8 iron ingot after smelting with 8 iron ore, 1
coal.
You can get 8 gold ingot after smelting with 8 gold ore, 1
coal.
You can craft 4 wood plank with 1 wood.
You can craft 4 stick with 2 wood plank.
You can craft 1 furnace with 8 cobblestone.
You can craft 3 paper with 3 reeds.
You can craft 1 book with 3 paper, 1 leather.
You can craft 1 enchanting table with 1 book, 4 obsidian, 2
diamond.
You can craft 1 clock with 4 gold ingot, 1 redstone.
You can craft 1 flint with 10 gravel.
You can craft 4 flint and steel with 1 iron ingot, 1 flint.
You can craft 2 blaze powder with 1 blaze rod.
You can craft 1 ender eye with 1 blaze powder, 1 ender pearl.
You can craft 10 wood pickaxe with 3 wood plank, 2 stick.
You can craft 10 wood axe with 3 wood plank, 2 stick.
You can craft 10 wood shovel with 1 wood plank, 2 stick.
You can craft 10 wood sword with 2 wood plank, 1 stick.
You can craft 20 stone pickaxe with 3 cobblestone, 2 stick.
You can craft 20 stone axe with 3 cobblestone, 2 stick.
You can craft 20 stone shovel with 1 cobblestone, 2 stick.
You can craft 20 stone sword with 2 cobblestone, 1 stick.
You can craft 40 iron pickaxe with 3 iron ingot, 2 stick.
You can craft 40 iron axe with 3 iron ingot, 2 stick.
You can craft 40 iron shovel with 1 iron ingot, 2 stick.
You can craft 40 iron sword with 2 iron ingot, 1 stick.
You can craft 10 gold pickaxe with 3 gold ingot, 2 stick.
You can craft 10 gold axe with 3 gold ingot, 2 stick.
You can craft 10 gold shovel with 1 gold ingot, 2 stick.
You can craft 10 gold sword with 2 gold ingot, 1 stick.
You can craft 100 diamond pickaxe with 3 diamond, 2 stick.
You can craft 100 diamond axe with 3 diamond, 2 stick.
You can craft 100 diamond shovel with 1 diamond, 2 stick.
You can craft 100 diamond sword with 2 diamond, 1 stick.
You can collect 2 dirt at forest.
You can collect 2 dirt at swamp.
You can collect 2 dirt at meadow.
You need to have at least 1 diamond shovel to collect 16 dirt
at forest.
1 diamond shovel will be consumed.
<The rest are omitted>
```

F.5 EGOCENTRIC PROMPT EXAMPLES

In MiniGrids Egocentric Observation experiments, the ASCII arts are rotated so that the agent is always facing forward. In Robomimic experiments, we verbalise the observation with relative position to the object.

Door Key with Egocentric Observation

You need to solve Minigrid Doorkey.
Minigrid Doorkey involves crossing a gridworld, where you need to pick up the key, unlock the door and reach the goal. Do not hit the wall.
The state is described with an ASCII art in first-person perspective. In the ASCII art, "#" means wall. "K" means key. "D" means door. "G" means goal. "." means empty tile.
"^" is where you are. You are always facing forward.
The possible actions are:

1. TURN LEFT.

2. TURN RIGHT.

3. MOVE FORWARD, moving forwards for one tile

4. PICK UP THE KEY, picking up the key, this is only possible when you are facing the tile with key on it.

5. UNLOCK THE DOOR, this action is only possible when you have the key.

You will be asked to predict the best action given a state.
You are in the following state represented in ASCII art:

```
#####
#..G#
##D##
#.^.#
#####
```

The door is locked. You have the key.
Which action do you choose? Please think step by step.
Only give the answer in a new line in JSON format:
`{"reasoning": <REASONING>, "action": <ACTION>}`
Where `<ACTION>` is one of "TURN LEFT", "TURN RIGHT", "MOVE FORWARD", "PICK UP THE KEY" or "UNLOCK THE DOOR", `<REASONING>` is a string of your thinking steps.

Door Key with Egocentric Observation

You are a robot arm trying to lift an object. You can move your gripper in the X, Y, or Z direction, as well as open/close your gripper. Each step, you see:

- Relative distance from your gripper to object in x direction
- Relative distance from your gripper to object in y direction
- Relative distance from your gripper in z direction
- Gripper finger state: open or closed

You can choose actions that move the gripper in the X, Y, Z direction, with additional control of the gripper.

Note on the gripper: You should close it when beginning to grasp the object, or when you are holding the object. When you are approaching the object, the gripper should be open. If the gripper needs to be closed, you should continue to close the gripper, even if it is closed. Similarly, if the gripper needs to be open, you should continue to open the gripper, even if it is already open.

You current observation is:

- The object is 0.002046926351863232 to the right of your gripper.
- The object is 0.024008545831402408 in front of your gripper.
- The object is 0.8217841891593334 below your gripper.

Given the following action: [-0.2087612 -0.02484818 -0.16213629 0.01804361 -0.04799016 0.00671699 0.9998945]
Here is a list of corrective actions that you can take. Each action will modify the current gripper state by the action amount. Which action do you choose? Please think step by step.

- Action 0: Move gripper forward by 0.28
- Action 1: Move gripper backward by 0.28
- Action 2: Move gripper right by 0.28
- Action 3: Move gripper left by 0.28
- Action 4: Move gripper vertically upward by 0.28
- Action 5: Move gripper vertically downward by 0.28
- Action 6: Open gripper
- Action 7: Close gripper

You should consider the position of the robot gripper and object, and how they are related to each other. For example, if the robot gripper is on the left of the object, you should consider moving the robot gripper to the right.

Only give the answer in a new line in JSON format:
{"reasoning": <REASONING>, "index": <INDEX>}
Where <INDEX> is one of action choice index, <REASONING> is a string of your thinking steps.

F.6 VLM FEEDBACK PROMPT EXAMPLES

Door Key with Image Observation

You need to solve Minigrid Doorkey. Minigrid Doorkey involves crossing a gridworld, where you need to pick up the key, unlock the door and reach the goal. Do not hit the wall. The state is depicted in an image.

This is you facing upwards. ▲

This is you facing right. ▶

This is you facing down. ▼

This is you facing left. ◀

This is an empty tile. ■

This is a key. 🔑

This is a locked door. ▤

This is an unlocked door. ▐

This is your goal. ■

This is a wall. ■ The possible actions are: 1. TURN LEFT, changing the direction to left 2. TURN RIGHT, changing the direction to right 3. MOVE FORWARD, moving forwards for one tile 4. PICK UP THE KEY, picking up the key, this is only possible when you are facing the tile with key on it 5. UNLOCK THE DOOR, this action is only possible when you have the key. Is action TURN LEFT the best action you can take? Please think step by step. Only give the answer in a new line in JSON format: `{"reasoning":` `<REASONING>,"feedback":` `<FEEDBACK>}`
where ¡FEEDBACK¿ is one of "YES" or "NO", ¡REASONING¿ is a string of your thinking steps.
This is your image observation.

Robomimic with Image Observation

You are a robot engineer. You need to control a robot arm to lift an object. You need to observe the robot's current state and choose which action to do next in order to lift the object. You can move the robot in the X, Y, or Z direction, as well as open/close its gripper. You will also recieve a image as your observation. You can choose actions that move the gripper in the X, Y, Z direction, with additional control of the gripper. The format of the action is:

n [resulting x, resulting y, resulting z, ignore, ignore, ignore, resulting gripper state]

nIn the image, you will see you are facing the robot and a red cube (the object) is on the table. When dx ¡ 0, the gripper moves away from you. When dx ¿ 0, the gripper moves towards you. When dy ¡ 0, the gripper moves towards the left of the picture. When dy ¿ 0, the gripper moves towards the right of the picture. When dz ¡ 0, the gripper moves towards the table. When dz ¿ 0, the gripper moves away from the table.

n Note on the gripper: Set grip to 1 to open the gripper and -1 to close the gripper. The robot's gripper should be closed if it is beginning to grasp the object, or when it is holding the object. When it is approaching the object, the gripper is open. If the robot gripper needs to be closed, you should continue to close the gripper, even if it is closed. Similarly, if the robot gripper needs to be open, you should continue to open the gripper, even if it is already open.

Your image observation is:

Which action do you choose? Please think step by step. You should consider the position of the robot gripper and object, and how they are related to each other. For example, if the robot gripper is on the left of the object, you should consider moving the robot gripper to the right. Only give the answer in a new line in JSON format:

`{"reasoning":` <REASONING>, `"action":` <ACTION>} where ¡ACTION¿ is the optimal action following the format of (dx, dy, dz), ¡REASONING¿ is a string of your thinking steps.

Domain	Pref	Act	Binary
Cliff Walking	444	37	148
Door Key	2000	100	500
Four Rooms	3000	500	1500
Robomimic (Lift)	2015	2015	2015
ALFWorld	2778k	3942	105k
HierarchyCraft	139k	1000	11k

Table 14: Dataset size for each domain and feedback type. The listed types are preference feedback, action advising, and binary feedback. Delta action is supported in Robomimic with 2015 examples. Goal advising is supported in Cliff Walking with 37 examples.

G DOMAINS AND DATASET COLLECTION DETAILS

Multiple Optimal Actions and Expert Policies During the collection of our dataset, we need an expert policy. However, in many domains, there might exist multiple optimal actions at the same time. In Cliff Walking, the expert policy is hand-coded and we enumerated all optimal actions for all states. In MiniGrid Door key and MiniGrid Four Rooms, we use exhaustive search to enumerate all possible optimal actions. In ALFWorld and HeirachyCraft, a built-in planner can serve as an expert policy. Although these planners give an optimal action for each state, they do not yield the full set of all optimal actions. Consequently, we cannot use these planners to compute the ground truth feedback for cases where there are multiple optimal actions. For these two domains, we also use exhaustive search to enumerate all possible actions, and consider action optimal if, after taking this action, there exists some sequence of remaining actions that reaches the goal in the same number of steps as the planner.

The final dataset size of each domain and feedback type can be seen in Table 14.

Preference and Binary Feedback We use the sampling policies only to sample states. When constructing our preference and preference dataset for discrete domains, given a state, we enumerate all legal actions (for binary feedback) and all combinations of legal actions (for preference).

Robomimic In Robomimic, the continuous action space of the domain means we cannot enumerate all optimal actions. In Robomimic experiments, we use an imitation learning policy as an expert policy, which gives 100% success rate in the task. We also conducted an empirical study to decide the threshold distance for optimal actions. We injected different scales of random noise into the actions and inspect the task success rate and found 0.28 is the threshold of noise where bigger noise starts to negatively impact the success rate. Therefore, we use this threshold as the biggest possible distance from the imitation learning policy action for an action to be optimal. We also use this threshold for delta actions.

H FEEDBACK MODEL ERROR EXAMPLES

In this section, we present some wrong responses from LLM.

Wrong response example in Cliff Walking Binary Feedback

You are in position (0, 4). Is action DOWN the best action you can take?

```
{ "reasoning": "The agent's goal is to reach (3, 11) from (0,
4).
Action DOWN moves to (1, 4).
While this starts descending, reaching row 3 safely requires
avoiding holes in columns 1-10. Moving right from (0, 4)
toward column 11 first ensures a direct path to the goal
without risking holes. Moving down now doesn't shorten the
path and delays aligning with the safe column (11). Thus,
moving right is better than down.",
"feedback": "NO"
}
```

In the example above, the LM assumes that they also need to keep away from the cliff, as well as give a wrong expectation of steps to the goal.

Wrong response example in HiearchyCraft Action Advising

```
<Prompt Omitted>
{"reasoning": "I need to find a way to get diamond, and I need
to craft a furnace to smelt the ore. I have wood plank and
stick, so I can craft a furnace. I also need to find iron ore
or gold ore to smelt. I can search for wood at the swamp to
get more wood and craft more items.",
"action": "search for wood at swamp"}'
```

In the example above, the LM does not have knowledge of system dynamics. Even if it correctly reasons that it needs to collect more wood, it still picks the non-optimal action. This action is not optimal because, moving to the forest and subsequently searching for wood leads to a shorter path to the goal.

Wrong response example in Door Key Preference

You need to solve Minigrid Doorkey.

Minigrid Doorkey involves crossing a gridworld, where you need to pick up the key, unlock the door and reach the goal. Do not hit the wall.

The state is described with an ASCII art. In the ASCII art, "#" means wall. "K" means key. "D" means door. "G" means goal. "." means empty tile. "^" means the agent who is facing upwards. "V" means the agent who is facing downwards. "<" means the agent who is facing leftwards. ">" means the agent who is facing rightwards.

The possible actions are:

1. TURN LEFT, changing the direction to left, for example, the agent "^" will become "<".

2. TURN RIGHT, changing the direction to right, for example, the agent "^" will become ">".

3. MOVE FORWARD, moving forwards for one tile

4. PICK UP THE KEY, picking up the key, this is only possible when you are facing the tile with key on it

5. UNLOCK THE DOOR, this action is only possible when you have the key.

You will be asked to compare two actions given a state and provide preference.

You are in the following state represented in ASCII art:

```
#####
#.#.#
#>D.#
#.#G#
#####
```

The door is unlocked. You have the key.

Given `MOVE FORWARD` or `TURN LEFT`, which action is better? Please think step by step. Only give the answer in a new line in JSON format:

`{"reasoning": <REASONING>, "preference": <PREFERENCE>}`

Where <PREFERENCE> is one of "FIRST" or "SECOND", <REASONING> is a string of your thinking steps.

```
{
    "reasoning": "In the current state, the agent is facing
    downwards and the door is unlocked. Moving forward is not
    a good option as it will lead to hitting the wall.
    Turning left is a better option as it allows the agent to
    change direction without moving and potentially see new
    possibilities to reach the goal.",
    "preference": "SECOND"
}
```

In the wrong response example above, the LM misinterprets the observation and deems itself facing downwards, which leads to a wrong response in preferring the non-optimal action `TURN LEFT`.

Wrong response example in Door Key Image Observation Action Advising

```
<Prompt Omitted>
```

This is your image observation.

```
{"reasoning": "I am currently facing right and need to move
towards the goal which is directly ahead. Since there's no
obstacle in my path, I can move forward.",
"action": "MOVE FORWARD"}
```

In this wrong response example, the VLM misrecognises the position of the goal and disregards the wall ahead and decides to move forward, which will lead the agent to hit the wall.

I STUDY ON CLOSED-WEIGHT MODEL

We also conducted the experiments on one of the state-of-the-art closed-weight model GPT5.1. The performance can be seen in Table 15. GPT-5.1 demonstrates superior performance across most domains, achieving 92.4% preference accuracy in ALFWorld. Notably, it rivals reasoning-specialised models in some settings, confirming that frontier APIs remain a strong source of feedback.

Domain	Preference	Action Adv	Bin Feedback	Delta Action
Cliff Walking	1.00	1.00	0.99	N/A
Door Key	0.69	0.50	0.83	N/A
Four Rooms	0.53	0.61	0.73	N/A
Robomimic	0.63	0.00	0.02	0.05
ALFWorld	0.92	0.90	0.76	N/A
HierachyCraft	0.88	0.72	0.90	N/A

Table 15: Feedback accuracy in GPT 5.1 in expert dataset. The goal advising accuracy for Cliff Walking is 0.98.

J STOCHASTICITY AMONG DIFFERENT SEEDS

We additionally conducted a study on generation stochasticity in Language Model Feedback in QwQ 32B. We conduct this study with 6 different seeds, and the standard variances are reported in Table 16. We see a standard deviation less than 0.058, indicating that the feedback signal is statistically reliable and not an artifact of lucky sampling.

Domain	Action Advising	Binary Feedback	Preference
CliffWalking	0.998 ± 0.004	0.995 ± 0.006	0.960 ± 0.028
DoorKey	0.844 ± 0.058	0.934 ± 0.011	0.936 ± 0.011
HCraft	0.596 ± 0.024	0.890 ± 0.009	0.837 ± 0.008

Table 16: Average feedback accuracy ($\pm$ standard deviation) across domains and feedback types in QwQ 32B over 6 seeds.

K ANALYSIS ON CoT LENGTH

We analysed the average Chain-of-Thought token length over different domains, models and types of feedback. The results are shown in the Table 17. We find a nuanced relationship. Correct feedback often correlates with concise CoT generation (fewer tokens), suggesting that when the model "understands" the physics, it reasons efficiently. Extremely long CoT chains often indicate hallucination loops or confusion. We also find that binary feedback often costs more tokens than other types of feedback.

L LEARNING FROM LLM FEEDBACK

In this study, we established the "upper bound" of signal quality available to learning from model feedback. We conducted an additional preliminary study to do online learning from LM Feedback. We integrated feedback from GPT-OSS-120B during online policy training (e.g. there is no offline dataset). GPT-OSS-120B shows high accuracy during on-policy feedback (77% preference, 88% binary feedback, 42% action advising on HierarchyCraft "get diamond"). For preference, we use PPO with a reward function learning from the preferences following (Christiano et al., 2017). For binary feedback, we directly use binary feedback as rewards in addition to terminal rewards. For action advising, we use an online version of DaggerRoss et al. (2011a) and conduct behaviour cloning on the

Domain	Type	Model	Correct Response CoT Length	Wrong Response CoT Length	Correct–Wrong CoT Length Difference
HierachyCraft	Action Advising	GPT5.1	80.79	86.42	-5.63
HierachyCraft	Binary Feedback	GPT5.1	112.58	133.36	-20.78
HierachyCraft	Preference	GPT5.1	118.56	139.43	-20.87
Four Rooms	Action Advising	GPT5.1	75.21	102.33	-27.11
Four Rooms	Binary Feedback	GPT5.1	126.39	141.19	-14.79
Four Rooms	Preference	GPT5.1	102.92	116.40	-13.48
ALFWorld	Action Advising	GPT5.1	63.66	62.12	1.54
ALFWorld	Binary Feedback	GPT5.1	93.01	96.29	-3.28
ALFWorld	Preference	GPT5.1	108.28	117.76	-9.48
HierachyCraft	Action Advising	QwQ-32B	70.69	66.76	3.93
HierachyCraft	Binary Feedback	QwQ-32B	86.20	87.84	-1.65
HierachyCraft	Preference	QwQ-32B	77.87	79.91	-2.04
Four Rooms	Action Advising	QwQ-32B	69.78	71.36	-1.58
Four Rooms	Binary Feedback	QwQ-32B	72.99	76.75	-3.76
Four Rooms	Preference	QwQ-32B	73.71	75.58	-1.87
ALFWorld	Action Advising	QwQ-32B	68.51	68.26	0.25
ALFWorld	Binary Feedback	QwQ-32B	90.26	90.85	-0.58
ALFWorld	Preference	QwQ-32B	88.27	92.13	-3.86
HierachyCraft	Action Advising	DeepSeek-R1	106.32	107.46	-1.14
HierachyCraft	Binary Feedback	DeepSeek-R1	121.77	126.58	-4.80
HierachyCraft	Preference	DeepSeek-R1	105.56	115.00	-9.44
Four Rooms	Action Advising	DeepSeek-R1	94.63	104.60	-9.96
Four Rooms	Binary Feedback	DeepSeek-R1	102.62	118.48	-15.87
Four Rooms	Preference	DeepSeek-R1	100.17	102.83	-2.66
ALFWorld	Action Advising	DeepSeek-R1	80.13	77.28	2.84
ALFWorld	Binary Feedback	DeepSeek-R1	105.87	108.42	-2.55
ALFWorld	Preference	DeepSeek-R1	96.97	104.01	-7.05
HierachyCraft	Action Advising	LLama-3.1-70B	102.66	95.24	7.42
HierachyCraft	Binary Feedback	LLama-3.1-70B	114.21	125.97	-11.76
HierachyCraft	Preference	LLama-3.1-70B	95.77	107.29	-11.52
Four Rooms	Action Advising	LLama-3.1-70B	89.14	95.45	-6.31
Four Rooms	Binary Feedback	LLama-3.1-70B	95.25	102.38	-7.13
Four Rooms	Preference	LLama-3.1-70B	80.63	80.15	0.48
ALFWorld	Action Advising	LLama-3.1-70B	79.87	80.15	-0.28
ALFWorld	Binary Feedback	LLama-3.1-70B	102.14	102.47	-0.33
ALFWorld	Preference	LLama-3.1-70B	90.80	98.98	-8.18
HierachyCraft	Action Advising	LLama-3.1-8B	99.01	98.00	1.02
HierachyCraft	Binary Feedback	LLama-3.1-8B	99.05	102.30	-3.25
HierachyCraft	Preference	LLama-3.1-8B	73.54	70.79	2.75
Four Rooms	Action Advising	LLama-3.1-8B	59.25	55.63	3.62
Four Rooms	Binary Feedback	LLama-3.1-8B	81.65	77.58	4.06
Four Rooms	Preference	LLama-3.1-8B	66.17	62.80	3.37
ALFWorld	Action Advising	LLama-3.1-8B	87.60	91.62	-4.01
ALFWorld	Binary Feedback	LLama-3.1-8B	110.37	112.68	-2.31
ALFWorld	Preference	LLama-3.1-8B	82.60	87.93	-5.34

Table 17: Average CoT token length of correct responses and wrong responses across domains and types of feedback.

states' policy visited with action advising from LLM. Each experiment consists of 5 seeds, and we compare the results with reinforcement learning with terminal reward only and with oracle feedback. For action advising, we compare the results with oracle action advising as a baseline. The results are shown in Figure 3. As we can see, feedback types where the model achieved high accuracy (Binary Feedback and Preference) significantly improved sample efficiency compared to learning from sparse terminal rewards alone. On the contrary, as predicted by our benchmarking, Action Advising

proved difficult for the model (low accuracy). Consequently, the downstream policy performance degraded when relying on this noisy signal.

(a) Reinforcement Learning from LLM Preference

(b) Reinforcement Learning from LLM Binary Feedback

(c) Behaviour Cloning from LLM Action Advising

Figure 3: Learning from online Language Model Feedback in HierachyCraft Get Diamond

M   REPRODUCIBILITY STATEMENT

The experiment details are converted throughout the paper as well as the appendix. For reproducibility, we included details in the dataset collection and evaluation pipelines in Section 2. We provide the prompt examples in Appendix F and more implementation details in Appendix G. We use models that are open-source and publicly available and we also provide a code and dataset appendix. We will also open-source our codebase.