

REFINING MULTI-TURN REASONING IN LLM AGENTS VIA TURN-LEVEL REWARD DESIGN

Anonymous authors

Paper under double-blind review

ABSTRACT

This paper investigates Reinforcement Learning (RL) approaches to enhance the reasoning capabilities of Large Language Model (LLM) agents in long-horizon, multi-turn scenarios. Such multi-turn agentic tasks can be naturally formalized as turn-level Markov Decision Processes (MDPs). However, most existing methods adopt MDP formulations with trajectory-level rewards, either terminal rewards that provide only a final outcome signal, or delayed rewards that merge intermediate and outcome signals into a single sparse feedback, leading to poor credit assignment. To address this limitation, we reformulate these tasks as MDPs with explicit turn-level rewards and provide theoretical analysis supporting the effectiveness of this design. Building on this formulation, we extend popular RL algorithms, GRPO and PPO, to their respective multi-turn variants, enabling fine-grained credit assignment. We conduct case studies on multi-turn reasoning-augmented search agents, where we carefully design two types of turn-level rewards: verifiable and LLM-as-judge. Our experiments on multi-turn search tasks demonstrate that our proposed formulation, incorporated well-designed turn-level rewards, enables RL algorithms to significantly outperform baseline methods with trajectory-level rewards. Both training and validation reward curves illustrate that our method achieves *greater stability, faster convergence, and higher accuracy*. Numerical results across diverse question-answering datasets further show that our approach consistently delivers highest answer correctness and 100% format correctness.

1 INTRODUCTION

Reinforcement Learning (RL) has recently emerged as a powerful approach for improving the reasoning capabilities of Large Language Models (LLMs), allowing them to explore and refine long Chains of Thought (CoT) (Wei et al., 2022) in complex decision-making tasks. Building on this paradigm, reasoning-based LLMs, such as OpenAI’s o1 (Jaech et al., 2024) and DeepSeek’s R1 (Guo et al., 2025a), demonstrate remarkable performance in textual reasoning tasks by learning analytical thinking and self-reflection. Despite these advancements, LLMs that rely solely on textual reasoning remain limited in tasks that require precise and complex numerical computation, information retrieval from web pages or local databases, or code execution. Equipping LLMs as autonomous agents with access to external tools, such as search engines, scientific calculators, or code interpreters, can significantly extend their capabilities beyond pure text-based reasoning (Gou et al., 2023).

Training LLMs to operate as autonomous agents in interactive environments faces unique challenges. Agent settings often require models to make sequential, multi-turn decisions in complex reasoning tasks. Many existing approaches (Chen et al., 2025b; Jin et al., 2025b; Feng et al., 2025a) formulate these multi-turn interactive tasks as single-turn problems, relying solely on final outcome-level rewards such as answer correctness. Popular RL algorithms, including Group Relative Policy Optimization (GRPO) (Shao et al., 2024) and Proximal Policy Optimization (PPO) (Schulman et al., 2017), are applied in this setting. However, such single-turn formulation is inadequate for long-horizon multi-turn reasoning as it treats the entire trajectory as a single decision step, ignoring the multi-turn structure of the interactive tasks. In particular, it ignores intermediate signals that evaluate each complete agent-environment interaction, such as a tool call and its result, providing feedback at the granularity of a single turn in multi-turn tasks (Lightman et al., 2023; Zhang et al., 2025b; Ma et al., 2023; Choudhury, 2025). Without access to dense turn-level feedback, agents struggle to refine their behavior, making it difficult to interact effectively with dynamic environments over multiple

054 steps. For example, in a search agent (Chen et al., 2025b; Jin et al., 2025a), selecting a good query
 055 early on is crucial for retrieving relevant information; without turn-level retrieval feedback, the agent
 056 may not learn which queries contribute to correct answers.

057 Recent studies (Li et al., 2025a; Qian et al., 2025; Wang et al., 2025a; Labs, 2025; Wang et al., 2025b;
 058 Singh et al., 2025; Zhang et al., 2025a; Jin et al., 2025a) model multi-turn agentic tasks as Markov
 059 Decision Processes (MDPs) and incorporate intermediate rewards like tool execution. However, these
 060 approaches suffer from a credit assignment problem: they merge outcome and intermediate rewards
 061 into a sparse trajectory-level signal for RL training. This aggregation makes advantage estimation
 062 inaccurate and prevents RL algorithms from providing fine-grained supervision across intermediate
 063 rounds of interaction (Guo et al., 2025b; Feng et al., 2025b; Zhang et al., 2025c).

064 Motivated by this, we investigate turn-level reward design for both multi-turn RL algorithms and
 065 agent applications. Our key contributions are as follows:

- 067 • Most existing methods adopt MDP formulations with trajectory-level rewards, either terminal
 068 rewards that provide only a final outcome signal, or delayed rewards that merge intermediate
 069 and outcome signals into a single sparse feedback, leading to poor credit assignment. To
 070 address this limitation, we reformulate these tasks as MDPs with explicit turn-level rewards
 071 and provide theoretical analysis supporting the effectiveness of this design.
- 072 • To train multi-turn LLM agents effectively under our MDP formulation, we propose to
 073 extend GRPO and PPO to their multi-turn variants by incorporating both outcome and
 074 intermediate rewards, enabling fine-grained credit assignment. While multi-turn GRPO
 075 requires exponential rollout samples to compute intermediate advantages, multi-turn PPO
 076 leverages a critic model, offering a more efficient and scalable solution.
- 077 • To highlight the importance of turn-level rewards, we conduct a case study using a reasoning-
 078 augmented search agent that performs multiple rounds of reasoning and search before
 079 producing the final answer. We carefully design turn-level verifiable rewards and turn-level
 080 LLM-as-judge rewards for training the search agent. While verifiable rewards are rigid, the
 081 LLM-as-judge enables a more flexible and nuanced evaluation.
- 082 • Building on this case study, our experiments on multi-turn reasoning-augmented search
 083 tasks show that our proposed MDP formulation integrated turn-level rewards enables RL
 084 algorithms to significantly outperform baseline methods with trajectory-level rewards. Both
 085 training and validation reward curves obtained with the Qwen2.5-7B model demonstrate that
 086 our algorithm with turn-level rewards achieves more stable training, faster convergence, and
 087 higher accuracy for both verifiable and LLM-as-judge rewards. Furthermore, benchmarks
 088 on both in-domain and out-of-domain tasks show that our approach consistently achieves
 089 the highest accuracy and reliably produces outputs with 100% correct format.

090 2 PROBLEM FORMULATION FOR MULTI-TURN AGENT INTERACTION

091 2.1 TURN-LEVEL MDP FORMULATION

092 Let x denote the input prompt sampled from the dataset \mathcal{D} , and $y = [l_1, f_1, \dots, l_K, f_K]$ denote the
 093 complete output response, where l_k is the response generated from an LLM policy π_θ , and f_k is the
 094 corresponding environment feedback at the k -th turn.

095 LLM agents operate in interactive environments where each turn yields stochastic feedback. To
 096 capture these dynamics, we formulate the multi-turn agentic task as a turn-level MDP, which is
 097 formally defined as $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, P, R, \gamma\}$. Here, \mathcal{S} denotes the state space, and \mathcal{A} denotes the
 098 action space; A state $s \in \mathcal{S}$ typically corresponds to an interaction history, while an action $a \in \mathcal{A}$
 099 often corresponds to a sequence of generated tokens; P represents the transition dynamics; R is the
 100 turn-level reward function; γ is the discount factor. At the k -th turn, conditioned on the current state
 101 s_k , the agent makes an action a_k according to the policy π_θ , where $a_k = [l_k, f_k]$ if environment
 102 feedback exists, otherwise $a_k = l_k$. The agent then receives a turn-level reward $R_k = R(s_k, a_k)$, and
 103 transitions to the next state s_{k+1} . A multi-turn rollout trajectory is
 104

$$105 \tau = \{(s_1, a_1, R_1), (s_2, a_2, R_2), \dots, (s_K, a_K, R_K)\}$$

106 Note that the outcome reward is denoted by $R(x, y) = R_K$ for a prompt-response pair (x, y) .

108
109

2.2 REWARD ASSIGNMENT IN TURN-LEVEL MDP

110
111

Based on the granularity of reward assignment, we categorize multi-turn formulations into three types of turn-level MDPs:

112
113
114
115
116
117

1. Turn-level MDP with a terminal reward \mathcal{M}_1 : provides only a final outcome (terminal) reward with no intermediate rewards.
2. Turn-level MDP with a delayed reward \mathcal{M}_2 : provides an accumulated reward that merges both intermediate and outcome rewards into a single delayed signal.
3. Turn-level MDP with explicit turn-level rewards \mathcal{M}_3 : provides explicit rewards at each turn.

118
119
120
121

Here, \mathcal{M}_1 contains only *outcome* rewards, whereas both \mathcal{M}_2 and \mathcal{M}_3 include *intermediate* rewards but differ in how these rewards are distributed across the turns. Moreover, \mathcal{M}_1 and \mathcal{M}_2 provide *trajectory-level* rewards, whereas \mathcal{M}_3 provides explicit *turn-level* rewards. Notably, most existing multi-turn agent studies adopt either \mathcal{M}_1 or \mathcal{M}_2 . In contrast, our paper focuses on \mathcal{M}_3 .

122
123

\mathcal{M}_1 is simple, and many existing studies adopt this formulation for multi-turn agentic tasks, relying on a final outcome reward such as answer correctness:

124
125

$$\max_{\pi_\theta} \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi_\theta(\cdot|x)} [R(x, y)] \quad (1)$$

126
127
128
129
130
131

which can be interpreted as a contextual bandit problem (Bouneffouf & Feraud, 2025; Baheri & Alm, 2023). However, such a single-turn formulation is inadequate because it treats the entire trajectory as a single decision step and ignores intermediate rewards that capture the structure of agent–environment interactions. Without intermediate rewards, the system must simulate entire trajectories before receiving any feedback, leaving it unable to prune or down-weight clearly suboptimal trajectories at early stages (Wu et al., 2023; Singhal et al., 2025) and resulting in poor credit assignment.

132
133

\mathcal{M}_2 and \mathcal{M}_3 are return-equivalent, both maximizing the cumulative discounted return:

134
135
136

$$\max_{\pi_\theta} \mathbb{E}_{s_k, a_k \sim \pi_\theta(\cdot|s_k)} \left[\sum_{k=1}^K \gamma^k R(s_k, a_k) \right] \quad (2)$$

137
138
139
140

in the sense that \mathcal{M}_2 and \mathcal{M}_3 have the same optimal Q -values and therefore share the same optimal policies. However, prior theoretical work (Ng et al., 1999; Arjona-Medina et al., 2019) indicates that when rewards are heavily delayed, \mathcal{M}_2 suffers from severe credit assignment issues, leading to high-variance gradients during policy optimization. See Appendix C for theoretical analysis.

141

3 GRPO WITH TURN-LEVEL REWARDS FOR MULTI-TURN AGENTIC TASKS

142
143

3.1 VANILLA GRPO WITH TRAJECTORY-LEVEL REWARDS

144
145
146
147
148
149
150

GRPO. Recently, the Group Relative Policy Optimization (GRPO) algorithm (Shao et al., 2024) has been widely used to enhance the reasoning capabilities of LLMs, which estimates the advantage in a group-relative manner. Specifically, for each input question x , it samples a group of responses $\{y_1, y_2, \dots, y_G\}$ from the reference policy π_{ref} . GRPO optimizes the policy by maximizing the following objective function:

151
152
153
154
155

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{x \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\text{old}}(\cdot|x)} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|y_i|} \sum_{t=1}^{|y_i|} \min(w_{i,t}(\theta) A_{i,t}, \text{clip}(w_{i,t}(\theta), 1 - \epsilon, 1 + \epsilon) A_{i,t}) - \beta \mathbb{D}_{\text{KL}} [\pi_\theta \parallel \pi_{\text{ref}}] \right], \quad (3)$$

156
157
158
159
160

where $w_{i,t}(\theta) = \frac{\pi_\theta(y_{i,t}|x, y_{i,<t})}{\pi_{\text{old}}(y_{i,t}|x, y_{i,<t})}$ is the token-level importance sampling ratio between the current policy π_θ and the previous policy π_{old} , ϵ is the clipping parameter, and β is the KL divergence coefficient. Given a group of trajectory-level rewards $\{R_i^{\text{traj}}\}_{i=1}^G$, the advantage of the i -th response $A_{i,t}$ is calculated by

161

$$A_{i,t} = A_i^{\text{GRPO}} = \frac{R_i^{\text{traj}} - \text{mean}(\{R_i^{\text{traj}}\}_{i=1}^G)}{\text{std}(\{R_i^{\text{traj}}\}_{i=1}^G)} \quad (4)$$

162 GRPO is naturally suited for MDPs with trajectory-level rewards, namely \mathcal{M}_1 and \mathcal{M}_2 , where
 163 $R_i^{\text{traj}} = R(x, y_i)$ for \mathcal{M}_1 , and $R_i^{\text{traj}} = \sum_{k=1}^K \gamma^k R_{i,k}$ for \mathcal{M}_2 , with $R_{i,k} = R(s_{i,k}, a_{i,k})$ denoting the
 164 intermediate reward given the state $s_{i,k}$ and action $a_{i,k}$ at the k -th turn.
 165

166 **Limitations of GRPO in Multi-Turn Settings.** For GPRO, the advantage function $A_{i,t}$ in Eq. (4)
 167 is computed at the *trajectory level*, which means the same advantage is assigned uniformly across
 168 the entire trajectory, without distinguishing the contributions of individual turns or tokens. For
 169 long-horizon multi-turn tasks, such coarse-grained credit assignment often leads to unstable training
 170 and suboptimal performance (Guo et al., 2025b; Feng et al., 2025b; Zhang et al., 2025c).
 171

3.2 MT-GPRO: TURN-LEVEL CREDIT ASSIGNMENT FOR GRPO

172 **MT-GPRO.** To highlight the importance of fine-grained credit assignment for GRPO, we consider a
 173 simple two-turn agent setting ($K = 2$), where the agent receives a group of intermediate rewards
 174 $\{R_i^I\}_{i=1}^G$ in the first turn and outcome rewards $\{R_i^O\}_{i=1}^G$ in the second turn. Based on these signals,
 175 we present our turn-level credit assignment strategy for GRPO. The resulting turn-level advantages in
 176 the first and second turns are given by:
 177


$$A_{i,1}^{\text{MT-GPRO}} = A_i^I + \alpha A_i^O, \quad A_{i,2}^{\text{MT-GPRO}} = A_i^O, \quad (5)$$

178 where A_i^I and A_i^O denote the intermediate and outcome advantages:
 179

$$A_i^I = \frac{R_i^I - \text{mean}(\{R_i^I\}_{i=1}^G)}{\text{std}(\{R_i^I\}_{i=1}^G)}, \quad A_i^O = \frac{R_i^O - \text{mean}(\{R_i^O\}_{i=1}^G)}{\text{std}(\{R_i^O\}_{i=1}^G)} \quad (6)$$

180 By leveraging intermediate rewards, all tokens within a single turn share a unified advantage signal.
 181 Moreover, the advantage of a turn depends not only on the rewards from that turn but also on the
 182 contributions of subsequent turns. We refer to this algorithm as *multi-turn GRPO (MT-GPRO)*. A
 183 detailed derivation of MT-GPRO for the K -turn setting is provided in Appendix F.
 184

185 **Case Study for MT-GPRO on a Two-Turn Agentic Task.** We conduct experiments to evaluate
 186 the proposed MT-GPRO method in a two-turn agent setting, where the agent first calls the search
 187 tool with reasoning in the initial turn and then produces the final answer in the subsequent turn (see
 188 Appendix G for details). Beyond the outcome-level exact-match reward, we design intermediate
 189 rewards based on tool-execution feedback for MT-GPRO. Figure 1 presents training reward curves
 190 for GRPO and MT-GPRO, which show that MT-GPRO achieves more stable tool usage (left figure)
 191 and higher exact-match accuracy (right figure), highlighting the importance of fine-grained credit
 192 assignment in multi-turn agentic tasks.
 193



207 Figure 1: Curves for different training reward components during training with various algorithms.
 208 GRPO-OR means GRPO with outcome rewards while GPRO-MR means GRPO with merged outcome
 209 and intermediate rewards. GPRO-OR, GPRO-MR, and MT-GPRO correspond to the MDPs \mathcal{M}_1 ,
 210 \mathcal{M}_2 , and \mathcal{M}_3 , respectively. Each plot shows the training reward score over training steps. Dotted
 211 lines represent the average reward across 10 runs, while solid lines show trends smoothed using the
 212 Exponential Moving Average (EMA).
 213

214 *Remark.* MT-GPRO has two limitations: (1) In MT-GPRO, computing the intermediate advantages
 215 requires G rollout samples at each turn. Therefore, over a horizon of K turns, this results in G^{K-1}
 216 rollout trajectories in total (see Appendix F for details). Such *exponential* growth in complexity makes

216 the approach computationally prohibitive for long-horizon multi-turn tasks. (2) This strategy also
 217 assumes that all rollout samples in a group must contain *the same number of turns*, which requires
 218 enforcing this constraint in the system prompt and leads to a fixed-turn setting. Such a restriction
 219 limits the flexibility and applicability of GRPO in more diverse scenarios. For example, in a search
 220 task, one question may be resolved in a single tool call or require multiple calls to retrieve, filter, and
 221 refine results in a sampled group.

223 4 PPO WITH TURN-LEVEL REWARDS FOR MULTI-TURN AGENTIC TASKS

225 In the previous section, we illustrated the importance of fine-grained credit assignment for GPRO,
 226 which improves the performance of LLM agents in multi-turn interactions. However, the exponential
 227 computational cost, together with the fixed-turn constraint, limits the applicability of MT-GRPO to
 228 general agentic tasks. In this section, we present the PPO algorithm with turn-level rewards, aiming
 229 to provide a more flexible, scalable, and efficient solution.

230 **PPO.** Proximal Policy Optimization (PPO) (Schulman et al., 2017) is a popular actor-critic RL
 231 algorithm commonly used for LLM training (Ouyang et al., 2022). PPO updates the policy by
 232 maximizing the following surrogate objective:

$$234 \mathcal{J}_{\text{PPO}}(\theta) = \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi_{\text{old}}(\cdot | x)} \left[\frac{1}{|y|} \sum_{t=1}^{|y|} \min(w_t(\theta) A_t, \text{clip}(w_t(\theta), 1 - \epsilon, 1 + \epsilon) A_t) \right], \quad (7)$$

235 The advantage estimate A_t is computed using Generalized Advantage Estimation (GAE) (Schulman
 236 et al., 2015), based on rewards and a learned value function (critic model). Formally, for a trajectory
 237 of length T , the GAE A_t at time step t is computed as:

$$238 A_t = \sum_{l=0}^{T-t-1} (\gamma \lambda)^l \delta_{t+l}, \quad \delta_t = r_t + \gamma V_{t+1} - V_t \quad (8)$$

239 where γ is the discount factor, $\lambda \in [0, 1]$ is the GAE parameter, δ_t is the temporal-difference error, r_t
 240 is the token-level reward and V_t is the token-level value at step t . Through the mechanism of GAE,
 241 the token-level value function enables token-level advantage estimation.

242 **Turn-Level Reward Assignment for PPO.** With explicit intermediate rewards, GAE provides
 243 fine-grained training signals at each turn. Given both intermediate rewards R^I and the outcome
 244 reward R^O , the token-level reward r_t is assigned as

$$245 r_t = \begin{cases} R^O & \text{if } t \text{ is the last token of the entire trajectory} \\ R^I & \text{if } t \text{ is the last token of the intermediate turn} \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

246 For clarity, we refer to PPO trained with both intermediate and outcome rewards as *multi-turn PPO*
 247 (*MT-PPO*), while PPO trained with only a sparse trajectory-level reward is referred to as *PPO*. To
 248 achieve fine-grained credit assignment with the usage of turn-level rewards, compared to MT-GRPO,
 249 which requires exponential rollout samples to compute intermediate advantages, MT-PPO leverages a
 250 critic model with GAE, offering a more efficient and scalable solution.

251 **Summary.** Table 1 summarizes the granular-
 252 ity of reward assignment and advantage
 253 estimation across different RL algorithms.
 254 As shown, MT-PPO provides fine-grained
 255 turn-level rewards and token-level advan-
 256 tage estimation. This higher granularity en-
 257 ables more precise credit assignment, which
 258 is particularly beneficial for multi-turn LLM
 259 agents where successful outcomes often de-
 260 pend on a sequence of intermediate deci-
 261 sions. In contrast, trajectory-level meth-
 262 ods provide coarser feedback, which often
 263 leads to weaker learning signals and unsta-
 264 ble training. These insights will be empirically validated in the following experiments.

265 Table 1: Comparison of granularity of reward assign-
 266 ment and advantage estimation across different RL
 267 algorithms for multi-turn LLM agents.

RL Algo.	Granularity	
	Reward	Advantage
GRPO	Trajectory-Level	Trajectory-Level
MT-GRPO	Turn-Level	Turn-Level
PPO	Trajectory-Level	Token-Level
MT-PPO	Turn-Level	Token-Level




Figure 2: Overview of the multi-turn reasoning-augmented search agent pipeline. Given a system prompt and a question, each iteration of the LLM-based search agent proceeds as follows: (1) The agent begins with *reasoning*, analyzing the current context to identify missing information. (2) It then formulates a search query to *retrieve* relevant information from an external database, which is integrated into the evolving *context*. (3) This cycle continues until the agent judges that the context is sufficient, at which point it performs a final round of *reasoning* to generate the answer.

5 CASE STUDY: MULTI-TURN REASONING-AUGMENTED SEARCH AGENT

5.1 TASK FORMULATION

We study an LLM agent that performs multi-turn reasoning with search engine interactions. The task can be naturally formulated under the turn-level MDP framework, as discussed in Section 2, where each action a corresponds either to a reasoning-augmented search step or to producing the final answer. The goal is to improve the agent’s performance through effective integration of external search. Specifically, the agent learns to leverage a Wikipedia search engine to retrieve relevant information and generate an accurate answer. Without search calling, the agent must rely solely on its internal knowledge to answer questions, which can limit accuracy, especially for fact-based queries requiring up-to-date or domain-specific information. The overall interaction follows a multi-turn reasoning–search loop, as shown in Figure 2.

5.2 TURN-LEVEL VERIFIABLE REWARD DESIGN

Unlike existing approaches (Chen et al., 2025b; Jin et al., 2025b), which assign binary rewards based solely on final-answer correctness for the entire trajectory, we design turn-level verifiable rewards for both intermediate and final turns to better align with the environment of the LLM-based search agent.

Outcome Verifiable Rewards evaluate the model-generated responses in the last turn, focusing on both the correctness of the answer and the adherence to the required output format.

- *Outcome Exact Match Reward* evaluates whether the extracted answer (from the `<answer>` tag) exactly matches any accepted ground-truth answer after normalization (e.g., lowercasing and whitespace removal):
- *Outcome Format Reward* ensures format correctness by verifying that: (1) only `<think>` and `<answer>` tags appear (no extra tags), (2) each tag appears exactly once, and (3) `<think>` precedes `<answer>`.

324 The outcome reward is defined as
 325

$$326 R^O = \begin{cases} 1 & f_{\text{em}} = \text{True}, f_{\text{format}} = \text{True}, \\ 327 0.2 & f_{\text{em}} = \text{False}, f_{\text{format}} = \text{True}, \\ 328 -1 & f_{\text{format}} = \text{False}, \end{cases}$$

329 where f_{em} and f_{format} denote the indicators of answer (exact-match) correctness and format correctness,
 330 respectively. A smaller positive reward is given when the answer is incorrect but the output follows
 331 the required format, encouraging structural correctness during training. A negative reward (penalty)
 332 is applied when the format is incorrect, ensuring that the agent adheres to the required structure.

333 **Intermediate Verifiable Rewards** guide the agent’s behavior in intermediate turns by evaluating
 334 the presence of ground-truth answers in retrieved content, enforcing proper format usage, and
 335 discouraging excessive search calls.

- 336 • *Intermediate Retrieval Existence Reward* evaluates whether any accepted answer appears in
 337 the one-round search result (from `<information>` tag), using case-insensitive matching.
 338 $R_{\text{retrieval}}^I = 0.3$ if retrieved information contains any ground-truth, otherwise 0.
- 339 • *Intermediate Format Reward* ensures format correctness by verifying that: (1) only
 340 `<think>`, `<search>`, and `<information>` tags appear (no extra tags), (2) each tag
 341 appears exactly once, and (3) `<think>` precedes `<search>` and `<information>`.
 342 $R_{\text{format}}^I = 0.1$ if the format is correct, otherwise -0.2 .
- 343 • *Intermediate Search Count Reward* penalizes excessive search usage.

$$344 R_{\text{search}}^I = -\lambda_s \cdot n_{\text{search}},$$

345 where λ_s is a predefined positive constant controlling the weight of the search count reward,
 346 n_{search} denotes the cumulative number of search invocations from the first turn up to the
 347 current turn.

348 The intermediate reward is defined as $R^I = R_{\text{retrieval}}^I + R_{\text{format}}^I + R_{\text{search}}^I$. Retrieval correctness is
 349 similarly assigned a smaller weight than answer correctness, again to reduce the risk of reward hacking.
 350 In addition, we introduce an intermediate search penalty to discourage excessive or unnecessary
 351 search calls, preventing the agent from either avoiding the avoiding the question answering or failing
 352 due to crashes.

353 5.3 LLM AS JUDGE FOR TURN-LEVEL EVALUATION

354 Verifiable rewards, such as exact match, provide a strict and objective form of evaluation. However,
 355 they can be overly rigid: an agent may produce a correct answer that differs slightly in form from
 356 the ground truth but still receives negative feedback. To complement such verifiable signals, we
 357 adopt the *LLM-as-judge* paradigm, where a strong LLM evaluates agent outputs. The LLM-as-judge
 358 framework consists of two key components: step-by-step reasoning and rubric-based scoring.

359 **Reasoning.** We employ a generative reasoning model (GRM) (Li et al., 2025b) as the judge,
 360 prompting it to generate detailed justifications before assigning a score. The step-by-step reasoning
 361 process encourages the judge to evaluate output quality using rubric-based criteria rather than relying
 362 on shallow correlations.

363 **Rubrics.** Rubric-based scoring provides structured evaluation criteria that improve both consistency
 364 and reliability across assessments. Unlike outcome-level evaluation that only considers the final
 365 answer, our framework assesses each turn’s output. This fine-grained assessment offers richer
 366 feedback and aligns naturally with multi-turn agentic tasks, where intermediate steps critically
 367 influence overall success. The judge model evaluates format correctness, reasoning quality, and
 368 search effectiveness, while also applying a search penalty to discourage excessive or unnecessary
 369 tool calls. Additional implementation details are provided in Appendix D.2.

370 6 EXPERIMENTS

371 In our experiments, we build our codebase upon the open-source project Search-R1 (Jin et al.,
 372 2025b), which trains LLM agents for multi-turn reasoning-augmented search tasks. More details on
 373 experimental settings can be found in Appendix D.1.

378 6.1 EVALUATED METHODS
379380 We compare both training reward dynamics and benchmark performance across different methods.
381382 **Training Dynamics.** We evaluate our MT-PPO against several PPO-based baselines:
383

- 384 • PPO-OR (Jin et al., 2025b): vanilla PPO trained with only outcome rewards, where the
385 trajectory-level reward is a binary signal indicating final-answer correctness, corresponding
386 to the terminal-reward MDP \mathcal{M}_1 .
- 387 • PPO-MR (Jin et al., 2025a): vanilla PPO trained with merged intermediate and outcome
388 rewards, where the trajectory-level reward combines intermediate rewards (retrieval correct-
389 ness) and outcome rewards (answer correctness and format correctness), corresponding to
390 the delayed-reward MDP \mathcal{M}_2 . The detailed reward design is provided in Section 4.1 of (Jin
391 et al., 2025a).
- 392 • MT-PPO (ours): PPO variant trained with both intermediate and outcome rewards, where the
393 turn-level reward design is described in Section 5.2, with $\lambda_s = 0.1$ by default, corresponding
394 to the turn-level-reward MDP \mathcal{M}_3 .

395 We omit GRPO training curves since, as reported in (Jin et al., 2025b), GRPO consistently crashes
396 during training.397 **Benchmark Evaluation.** In addition to the base model and the instruct model, we further compare
398 our method against Search-R1 trained with GRPO and PPO (Jin et al., 2025b;a),¹ OTC trained with
399 GRPO and PPO (Wang et al., 2025a), and StepSearch trained with PPO (Wang et al., 2025c).400 Since PPO baselines often crash, we evaluate them using either the final checkpoint or the last
401 checkpoint prior to collapse.402 **Evaluation Metrics.** We evaluate model performance using three types of rewards: (1) answer
403 correctness (exact match) reward, (2) format correctness reward, and (3) retrieval correctness reward.
404 Each reward is assigned a value of 1.0 if the criterion is satisfied and 0 otherwise.405 6.2 EXPERIMENT SETUP
406407 **Datasets.** These datasets are categorized as follows: (1) General Question Answering: NQ
408 (Karpukhin et al., 2020), TriviaQA (Joshi et al., 2017), and PopQA (Mallen et al., 2022). (2)
409 Multi-Hop Question Answering: HotpotQA (Yang et al., 2018), 2WikiMultiHopQA (Ho et al., 2020),
410 and Musique (Trivedi et al., 2022). These datasets cover a diverse range of search and reasoning
411 challenges, providing a comprehensive basis for evaluation.412 **Training Details.** We use Qwen2.5-7B (Yang et al., 2024) as the base model, E5 (Wang et al., 2022)
413 as the retriever, and 2018 Wikipedia dump (Karpukhin et al., 2020) as the corpus. We set the number
414 of retrieved passages to 3, and the maximum number of turns N_{\max} to 4. The system prompt follows
415 that of Search-R1 (Jin et al., 2025b). We also enable policy loss masking on retrieved tokens.416 6.3 MAIN RESULTS
417418 **Training Dynamics.** Figures 3 and 5 show training and validation reward curves for PPO and
419 MT-PPO. MT-PPO achieves substantially more stable training, converging faster in the early phase
420 (first 100 steps) thanks to intermediate rewards that provide stronger guidance. As training progresses,
421 PPO exhibits high variance and even performance degradation, especially on HotpotQA, while
422 MT-PPO maintains consistent improvement. MT-PPO attains higher average accuracy than PPO,
423 demonstrating greater robustness. Format reward curves show that MT-PPO consistently follows
424 the correct output format, while PPO struggles, especially on HotpotQA, where formatting mistakes
425 prevent correct evaluation. This indicates that turn-level rewards in MT-PPO stabilize training and
426 enforce structural correctness. Retrieval curves further show that MT-PPO achieves more consistent
427 accuracy by leveraging intermediate signals to guide reasoning. Figure 6 presents training curves for
428 MT-PPO and PPO with judge rewards, where MT-PPO again demonstrates stable optimization.429
430 ¹The GRPO baselines (GRPO-OR and GRPO-MR) correspond to the PPO baselines (PPO-OR and PPO-MR)
431 with the same reward design (Jin et al., 2025b;a).




Figure 3: Training reward curves recorded during training for PPO baselines and MT-PPO on the NQ and HotpotQA datasets. The rewards include answer correctness, format correctness, and retrieval correctness. Solid lines show mean reward values, while shaded regions indicate variability across five independent runs.

Table 2: The performance results of different methods on six datasets. Bold numbers indicate the best performance for each dataset and metric. [†]/^{*} denote in-domain/out-of-domain datasets. [‡] indicates results copied from the original paper.

Methods	General QA			Multi-Hop QA			Avg.
	NQ [†]	TriviaQA [*]	PopQA [*]	HotpotQA [†]	2wiki [*]	Musique [*]	
<i>Answer Correctness (Exact Match)</i>							
Qwen2.5-7B-Base	0.177	0.319	0.181	0.160	0.167	0.040	0.174
Qwen2.5-7B-Instruct	0.320	0.563	0.349	0.292	0.277	0.118	0.320
GRPO-OR (Search-R1)	0.391	0.560	0.388	0.331	0.306	0.129	0.351
GRPO-MR (Search-R1) [‡]	0.453	0.628	0.450	0.416	0.375	0.164	0.414
PPO-OR (Search-R1)	0.483	0.639	0.456	0.435	0.382	0.199	0.432
PPO-MR (Search-R1) [‡]	0.472	0.629	0.452	0.436	0.402	0.180	0.429
GRPO (OTC) [‡]	0.444	0.597	0.431	0.366	0.311	0.130	0.380
PPO (OTC) [‡]	0.446	0.623	0.425	0.383	0.363	0.152	0.399
PPO (StepSearch)	0.355	0.570	0.385	0.351	0.396	0.179	0.373
MT-PPO (ours)	0.490	0.647	0.459	0.453	0.424	0.209	0.447
<i>Format Correctness</i>							
Qwen2.5-7B-Base	0.118	0.118	0.105	0.098	0.084	0.082	0.101
Qwen2.5-7B-Instruct	0.183	0.267	0.067	0.109	0.037	0.071	0.122
GRPO-OR (Search-R1)	0.706	0.685	0.597	0.513	0.376	0.328	0.534
PPO-OR (Search-R1)	0.909	0.954	0.952	0.916	0.806	0.834	0.895
PPO (StepSearch)	0.521	0.614	0.668	0.560	0.396	0.571	0.555
MT-PPO (ours)	0.999	0.997	0.999	0.998	0.999	0.999	0.999

Benchmark Performance. Table 2 reports results on six QA datasets, spanning both general and multi-hop reasoning tasks. MT-PPO consistently outperforms PPO and GRPO in answer correctness, with the largest gains on multi-hop tasks such as HotpotQA and 2Wiki. Moreover, MT-PPO nearly perfects format correctness, reaching close to 100% across datasets, underscoring the effectiveness of multi-turn credit assignment in producing both accurate and well-structured outputs.

6.4 ABLATION STUDY




Figure 4: Ablation studies on (1) the search count reward λ_s and (2) the maximum number of turns N_{\max} on the NQ dataset. The left panel reports answer correctness, the middle panel shows the average number of turns, and the right panel illustrates accuracy under different N_{\max} settings.

We conduct two ablation studies to analyze the effects of (1) the search-count reward and (2) the maximum number of turns on training dynamics and final performance. As shown in Figure 4, incorporating a moderate search-count reward (e.g., $\lambda_s = 0.1$) significantly improves training stability and answer correctness. The left panel shows that MT-PPO with $\lambda_s = 0.1$ achieves the highest and most consistent accuracy, while overly strong penalties (e.g., $\lambda_s = 0.3$) degrade performance.

The middle panel illustrates how the search-count reward shapes the agent’s turn usage. With $\lambda_s = 0.1$, the agent learns to reduce unnecessary search calls early in training and eventually stabilizes around an efficient number of turns. In contrast, removing this term ($\lambda_s = 0.0$) leads to unstable behavior, including excessive or erratic turn usage, which ultimately harms convergence.

Finally, the right panel shows the effect of varying the maximum number of allowed turns N_{\max} . The results indicate that MT-PPO is robust across different turn limits: adjusting N_{\max} from 4 to 6 yields nearly identical accuracy curves. This suggests that MT-PPO adapts its strategy effectively without being overly sensitive to the chosen turn budget.

7 CONCLUSION AND FUTURE WORK

In this paper, we highlighted the importance of turn-level rewards for multi-turn agentic tasks. By introducing carefully designed intermediate signals, we extended GRPO and PPO into multi-turn variants, allowing LLM agents to receive more informative feedback at each stage of interaction. Experiments on reasoning-augmented search agents show that incorporating turn-level rewards substantially improves both the stability and accuracy of training across different RL algorithms. We believe that turn-level rewards have broad applicability beyond search, offering a general mechanism for improving the effectiveness of multi-turn agents in diverse interactive environments.

540 REFERENCES
541

542 Jose A Arjona-Medina, Michael Gillhofer, Michael Widrich, Thomas Unterthiner, Johannes Brand-
543 stetter, and Sepp Hochreiter. Rudder: Return decomposition for delayed rewards. *Advances in
544 Neural Information Processing Systems*, 32, 2019.

545 Ali Baheri and Cecilia Alm. Llms-augmented contextual bandit. In *NeurIPS 2023 Foundation Models
546 for Decision Making Workshop*. NeurIPS 2023, 2023.

547 Hao Bai, Yifei Zhou, Jiayi Pan, Mert Cemri, Alane Suhr, Sergey Levine, and Aviral Kumar. Digirl:
548 Training in-the-wild device-control agents with autonomous reinforcement learning. *Advances in
549 Neural Information Processing Systems*, 37:12461–12495, 2024.

550 Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou,
551 and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization,
552 text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023.

553 Djallel Bouneffouf and Raphael Feraud. Multi-armed bandits meet large language models. *arXiv
554 preprint arXiv:2505.13355*, 2025.

555 William Brown. Verifiers: Reinforcement learning with llms in verifiable environments. [https://
556 /github.com/willccbb/verifiers](https://github.com/willccbb/verifiers), 2025.

557 Thomas Carta, Clément Romac, Thomas Wolf, Sylvain Lamprier, Olivier Sigaud, and Pierre-Yves
558 Oudeyer. Grounding large language models in interactive environments with online reinforcement
559 learning. In *International Conference on Machine Learning*, pp. 3676–3713. PMLR, 2023.

560 Kevin Chen, Marco Cusumano-Towner, Brody Huval, Aleksei Petrenko, Jackson Hamburger, Vladlen
561 Koltun, and Philipp Krähenbühl. Reinforcement learning for long-horizon interactive llm agents.
562 *arXiv preprint arXiv:2502.01600*, 2025a.

563 Mingyang Chen, Tianpeng Li, Haoze Sun, Yijie Zhou, Chenzheng Zhu, Fan Yang, Zenan Zhou,
564 Weipeng Chen, Haofen Wang, Jeff Z Pan, et al. Learning to reason with search for llms via
565 reinforcement learning. *arXiv preprint arXiv:2503.19470*, 2025b.

566 Wentse Chen, Jiayu Chen, Hao Zhu, and Jeff Schneider. Context-lite multi-turn reinforcement
567 learning for llm agents. In *ES-FoMo III: 3rd Workshop on Efficient Systems for Foundation Models*,
568 2025c.

569 Jie Cheng, Ruixi Qiao, Lijun Li, Chao Guo, Junle Wang, Gang Xiong, Yisheng Lv, and Fei-Yue Wang.
570 Stop summation: Min-form credit assignment is all process reward model needs for reasoning.
571 *arXiv preprint arXiv:2504.15275*, 2025.

572 Sanjiban Choudhury. Process reward models for llm agents: Practical framework and directions.
573 *arXiv preprint arXiv:2502.10325*, 2025.

574 Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu
575 Yu, Qixin Xu, Weize Chen, et al. Process reinforcement through implicit rewards. *arXiv preprint
576 arXiv:2502.01456*, 2025.

577 Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, Baoquan Zhong, Chengquan Jiang,
578 Jinxin Chi, and Wanjun Zhong. Retool: Reinforcement learning for strategic tool use in llms.
579 *arXiv preprint arXiv:2504.11536*, 2025a.

580 Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. Group-in-group policy optimization for llm
581 agent training. *arXiv preprint arXiv:2505.10978*, 2025b.

582 Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Minlie Huang, Nan Duan, and
583 Weizhu Chen. Tora: A tool-integrated reasoning agent for mathematical problem solving. *arXiv
584 preprint arXiv:2309.17452*, 2023.

585 Daya Guo, Dejian Yang, Huawei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
586 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
587 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025a.

594 Yiran Guo, Lijie Xu, Jie Liu, Dan Ye, and Shuang Qiu. Segment policy optimization: Effective
 595 segment-level credit assignment in rl for large language models. *arXiv preprint arXiv:2505.23564*,
 596 2025b.

597 Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. Constructing a multi-hop
 598 qa dataset for comprehensive evaluation of reasoning steps. *arXiv preprint arXiv:2011.01060*,
 599 2020.

600 Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec
 601 Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint*
 602 *arXiv:2412.16720*, 2024.

603 Bowen Jin, Jinsung Yoon, Priyanka Kargupta, Sercan O Arik, and Jiawei Han. An empirical
 604 study on reinforcement learning for reasoning-search interleaved llm agents. *arXiv preprint*
 605 *arXiv:2505.15117*, 2025a.

606 Bowen Jin, Hansi Zeng, Zhenrui Yue, Dong Wang, Hamed Zamani, and Jiawei Han. Search-r1:
 607 Training llms to reason and leverage search engines with reinforcement learning. *arXiv preprint*
 608 *arXiv:2503.09516*, 2025b.

609 Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. TriviaQA: A large scale
 610 distantly supervised challenge dataset for reading comprehension. In Regina Barzilay and
 611 Min-Yen Kan (eds.), *Proceedings of the 55th Annual Meeting of the Association for Com-
 612 putational Linguistics (Volume 1: Long Papers)*, pp. 1601–1611, Vancouver, Canada, July
 613 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1147. URL <https://aclanthology.org/P17-1147/>.

614 Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick SH Lewis, Ledell Wu, Sergey Edunov, Danqi
 615 Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. In *EMNLP*
 616 (1), pp. 6769–6781, 2020.

617 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph
 618 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model
 619 serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems
 620 Principles*, pp. 611–626, 2023.

621 Bespoke Labs. Improving multi-turn tool use with reinforcement learning. [https://www.bespokelabs.ai/blog/
 622 improving-multi-turn-tool-use-with-reinforcement-learning](https://www.bespokelabs.ai/blog/improving-multi-turn-tool-use-with-reinforcement-learning), 2025.
 623 Accessed: 2025-04-17.

624 Xuefeng Li, Haoyang Zou, and Pengfei Liu. Torl: Scaling tool-integrated rl. *arXiv preprint*
 625 *arXiv:2503.23383*, 2025a.

626 Yi-Chen Li, Tian Xu, Yang Yu, Xuqin Zhang, Xiong-Hui Chen, Zhongxiang Ling, Ningjing Chao,
 627 Lei Yuan, and Zhi-Hua Zhou. Generalist reward models: Found inside large language models.
 628 *arXiv preprint arXiv:2506.23235*, 2025b.

629 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
 630 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth
 631 International Conference on Learning Representations*, 2023.

632 Runze Liu, Fengshuo Bai, Yali Du, and Yaodong Yang. Meta-reward-net: Implicitly differentiable
 633 reward learning for preference-based reinforcement learning. *Advances in Neural Information
 634 Processing Systems*, 35:22270–22284, 2022.

635 Qianli Ma, Haotian Zhou, Tingkai Liu, Jianbo Yuan, Pengfei Liu, Yang You, and Hongxia Yang.
 636 Let's reward step by step: Step-level reward model as the navigators for reasoning. *arXiv preprint*
 637 *arXiv:2310.10080*, 2023.

638 Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi.
 639 When not to trust language models: Investigating effectiveness of parametric and non-parametric
 640 memories. *arXiv preprint arXiv:2212.10511*, 2022.

648 Andrew Y Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations:
 649 Theory and application to reward shaping. In *Icmi*, volume 99, pp. 278–287. Citeseer, 1999.
 650

651 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
 652 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow
 653 instructions with human feedback. *Advances in neural information processing systems*, 35:27730–
 654 27744, 2022.

655 Eduardo Pignatelli, Johan Ferret, Matthieu Geist, Thomas Mesnard, Hado van Hasselt, Olivier
 656 Pietquin, and Laura Toni. A survey of temporal credit assignment in deep reinforcement learning.
 657 *arXiv preprint arXiv:2312.01072*, 2023.

658 Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiusi Chen, Dilek Hakkani-Tür, Gokhan Tur,
 659 and Heng Ji. Toolrl: Reward is all tool learning needs. *arXiv preprint arXiv:2504.13958*, 2025.
 660

661 Max-Philip B. Schrader. gym-sokoban. <https://github.com/mpSchrader/gym-sokoban>, 2018.

662

663 Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon
 664 Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari,
 665 go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.

666

667 John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional
 668 continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*,
 669 2015.

670 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 671 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

672

673 Bytedance Seed. Seed-thinking-v1. 5: Advancing superb reasoning models with re-
 674 inforcement learning. Technical report, Technical report, ByteDance, 2025. URL
 675 <https://github.com/ByteDance-Seed/Seed-Thinking-v1.5/blob/main/seed-thinking-v1.5.pdf>.

676

677 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 678 Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical
 679 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

680 Joykirat Singh, Raghav Magazine, Yash Pandya, and Akshay Nambi. Agentic reasoning and tool
 681 integration for llms via reinforcement learning. *arXiv preprint arXiv:2505.01441*, 2025.

682

683 Raghav Singhal, Zachary Horvitz, Ryan Teehan, Mengye Ren, Zhou Yu, Kathleen McKeown, and
 684 Rajesh Ranganath. A general framework for inference-time scaling and steering of diffusion
 685 models. In *International Conference on Machine Learning (ICML)*, 2025.

686

687 Shengjie Sun, Runze Liu, Jiafei Lyu, Jing-Wen Yang, Liangpeng Zhang, and Xiu Li. A large
 688 language model-driven reward design framework via dynamic feedback for reinforcement learning.
 689 *Knowledge-Based Systems*, 326:114065, 2025.

690

691 Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. Musique: Multihop
 692 questions via single-hop question composition. *Transactions of the Association for Computational
 693 Linguistics*, 10:539–554, 2022.

694

695 Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia
 696 Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process-and
 697 outcome-based feedback. *arXiv preprint arXiv:2211.14275*, 2022.

698

699 Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan
 700 Lambert, Shengyi Huang, Kashif Rasul, and Quentin Gallouédec. Trl: Transformer reinforcement
 701 learning. <https://github.com/huggingface/trl>, 2020.

702

703 Hongru Wang, Cheng Qian, Wanjun Zhong, Xiusi Chen, Jiahao Qiu, Shijue Huang, Bowen Jin,
 704 Mengdi Wang, Kam-Fai Wong, and Heng Ji. Otc: Optimal tool calls via reinforcement learning.
 705 *arXiv preprint arXiv:2504.14870*, 2025a.

702 Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Dixin Jiang, Rangan Majumder,
 703 and Furu Wei. Text embeddings by weakly-supervised contrastive pre-training. *arXiv preprint*
 704 *arXiv:2212.03533*, 2022.

705 Zihan Wang, Kangrui Wang, Qineng Wang, Pingyue Zhang, Linjie Li, Zhengyuan Yang, Kefan Yu,
 706 Minh Nhat Nguyen, Licheng Liu, Eli Gottlieb, et al. Ragen: Understanding self-evolution in llm
 707 agents via multi-turn reinforcement learning. *arXiv preprint arXiv:2504.20073*, 2025b.

708 Ziliang Wang, Xuhui Zheng, Kang An, Cijun Ouyang, Jialu Cai, Yuhang Wang, and Yichao Wu.
 709 Stepsearch: Igniting llms search ability via step-wise proximal policy optimization. *arXiv preprint*
 710 *arXiv:2505.15107*, 2025c.

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

714 Luhuan Wu, Brian L. Trippe, Christian A. Naesseth, David M. Blei, and John P. Cunningham.
 715 Practical and asymptotically exact conditional sampling in diffusion models. In *Advances in*
 716 *Neural Information Processing Systems 36 (NeurIPS 2023)*, 2023.

717 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
 718 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint*
 719 *arXiv:2412.15115*, 2024.

720 Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov,
 721 and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question
 722 answering. *arXiv preprint arXiv:1809.09600*, 2018.

723 Shunyu Yao, Rohan Rao, Matthew Hausknecht, and Karthik Narasimhan. Keep CALM and
 724 explore: Language models for action generation in text-based games. In Bonnie Webber,
 725 Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020 Conference on Empirical*
 726 *Methods in Natural Language Processing (EMNLP)*, pp. 8736–8754, Online, November
 727 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.704. URL
 728 <https://aclanthology.org/2020.emnlp-main.704/>.

729 Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. Webshop: Towards scalable
 730 real-world web interaction with grounded language agents. *Advances in Neural Information*
 731 *Processing Systems*, 35:20744–20757, 2022.

732 Simon Zhai, Hao Bai, Zipeng Lin, Jiayi Pan, Peter Tong, Yifei Zhou, Alane Suhr, Saining Xie, Yann
 733 LeCun, Yi Ma, et al. Fine-tuning large vision-language models as decision-making agents via
 734 reinforcement learning. *Advances in neural information processing systems*, 37:110935–110971,
 735 2024.

736 Shaokun Zhang, Yi Dong, Jieyu Zhang, Jan Kautz, Bryan Catanzaro, Andrew Tao, Qingyun Wu,
 737 Zhiding Yu, and Guilin Liu. Nemotron-research-tool-n1: Tool-using language models with
 738 reinforced reasoning. *arXiv preprint arXiv:2505.00024*, 2025a.

739 Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, Bowen Yu, Dayiheng Liu,
 740 Jingren Zhou, and Junyang Lin. The lessons of developing process reward models in mathematical
 741 reasoning. *arXiv preprint arXiv:2501.07301*, 2025b.

742 Zijing Zhang, Ziyang Chen, Mingxiao Li, Zhaopeng Tu, and Xiaolong Li. Rlvmr: Reinforcement
 743 learning with verifiable meta-reasoning rewards for robust long-horizon agents. *arXiv preprint*
 744 *arXiv:2507.22844*, 2025c.

745

746

747

748

749

750

751

752

753

754

755

756 **A LLM USAGE**
757758 In this work, LLMs were used exclusively for polishing the writing. No part of the technical content,
759 experimental design, or analysis relied on LLMs. The authors retain full responsibility for the
760 correctness and originality of the ideas, methods, and results.
761762 **B RELATED WORK**
763764 **B.1 PROCESS REWARD AND CREDIT ASSIGNMENT IN RL**
765766 Process rewards provide fine-grained credit assignment and enhance both training efficiency and
767 optimization stability in RL. Such dense rewards have proven effective in classical RL domains such
768 as games and robotic control (Schrittwieser et al., 2020; Liu et al., 2022; Sun et al., 2025). Process
769 reward models have also been extensively explored for inference-time scaling in LLMs (Lightman
770 et al., 2023; Uesato et al., 2022). Recent studies have further highlighted the importance of effective
771 credit assignment in RL (Pignatelli et al., 2023; Shao et al., 2024; Cui et al., 2025; Cheng et al., 2025;
772 Feng et al., 2025b; Guo et al., 2025b), particularly for textual reasoning tasks such as mathematical
773 problem solving. In multi-turn agent interaction settings, turn-level rewards evaluate each complete
774 agent-environment interaction, such as a tool invocation and its resulting output, providing feedback
775 at the granularity of a single turn. This setting naturally emphasizes the advantages of process-level
776 rewards and fine-grained credit assignment. However, the design of effective reward functions for
777 multi-turn agents, as well as RL algorithms capable of leveraging such fine-grained credit signals,
778 remains underexplored. A recent study (Chen et al., 2025c) investigates step-level credit assignment
779 for multi-turn LLM agents by introducing a dual-discounting GAE formulation in PPO. While their
780 method relies on critic-derived step-level value estimates for credit assignment, our approach directly
781 incorporates explicit intermediate rewards into the PPO objective, enabling more precise and more
782 stable credit assignment.
783784 **B.2 RL FOR LLM AGENTS**
785786 RL has been applied to train long-horizon multi-turn LLM agents in diverse domains, including
787 search (Chen et al., 2025b; Jin et al., 2025b;a), tool use (Feng et al., 2025a; Li et al., 2025a; Qian
788 et al., 2025; Wang et al., 2025a; Labs, 2025; Zhang et al., 2025a; Singh et al., 2025), text-based
789 games (Yao et al., 2020; Carta et al., 2023; Zhai et al., 2024; Wang et al., 2025b), web shopping (Yao
790 et al., 2022), digital app interaction (Chen et al., 2025a), and mobile device control (Bai et al., 2024).
791 A number of these studies (Jin et al., 2025a; Feng et al., 2025a; Li et al., 2025a; Qian et al., 2025;
792 Wang et al., 2025a; Labs, 2025; Zhang et al., 2025a; Singh et al., 2025) apply RL algorithms such as
793 GRPO and PPO to train tool-using LLM agents, including calculators, code interpreters, and search
794 engines, thus enabling reasoning with external tools. However, these methods generally collapse
795 outcome- and turn-level signals into a single trajectory-level reward, limiting fine-grained credit
796 assignment. One related work is StepSearch (Wang et al., 2025c), which applies PPO with turn-level
797 rewards for multi-turn search. However, it relies heavily on data augmentation and requires prompt
798 modifications during preprocessing. In contrast, our method avoids such prompt engineering and
799 provides a cleaner, more general framework for turn-level reward design. Our approach is orthogonal
800 to existing search-agent methods (Chen et al., 2025b; Jin et al., 2025b;a; Wang et al., 2025c) and
801 applies broadly to multi-turn LLM agents beyond search tasks.
802803 **C THEORETICAL ANALYSIS OF POLICY-GRADIENT VARIANCE**
804805 In this section, we analyze and compare the policy-gradient variance of the two return-equivalent
806 MDPs \mathcal{M}_2 and \mathcal{M}_3 , since the comparison between \mathcal{M}_1 and \mathcal{M}_3 is trivial. Although \mathcal{M}_2 and \mathcal{M}_3
807 share the same optimal Q -values and thus the same optimal policies, they differ in how rewards are
808 assigned across turns. Our theoretical results show that the MDP with explicit turn-level rewards
809 \mathcal{M}_3 yields lower policy-gradient variance than the MDP that relies solely on a delayed accumulated
810 reward \mathcal{M}_2 . These findings help justify our formulation and highlight why turn-level rewards and
811 fine-credit credit assignment leads to more stable and efficient RL training for GRPO and PPO, as
812 demonstrated in Sections 3 and 4.
813

810 In general, the policy gradient of the expected return objective $J(\theta)$ can be written as
 811

$$812 \quad 813 \quad 814 \quad \nabla_{\theta} J(\theta) = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{k=1}^K \nabla_{\theta} \log \pi_{\theta}(a_k \mid s_k) G_k \right],$$

815 where G_k denotes the return associated with timestep k . For the delayed-reward MDP \mathcal{M}_2 , the
 816 timestep return is

$$817 \quad 818 \quad 819 \quad G_k = \sum_{t=1}^K \gamma^t R(s_t, a_t),$$

820 which is the discounted sum of all turn-level rewards over the trajectory. For the turn-level-reward
 821 MDP \mathcal{M}_3 , the timestep return is

$$822 \quad 823 \quad 824 \quad G_k = \sum_{t=k}^K \gamma^{t-k} R(s_t, a_t),$$

825 which corresponds to the discounted future return starting from timestep k .

826 **Lemma 1.** Consider a K -step episodic MDP with random rewards R_1, \dots, R_K . Let

$$827 \quad 828 \quad h_k = \nabla_{\theta} \log \pi_{\theta}(a_k \mid s_k), \quad k = 1, \dots, K.$$

829 Suppose that the discount factor is $\gamma = 1$, and define the returns

$$830 \quad 831 \quad 832 \quad G_k^{\mathcal{M}_2} = \sum_{t=1}^K R_t, \quad G_k^{\mathcal{M}_3} = \sum_{t=k}^K R_t.$$

833 The REINFORCE estimator is

$$834 \quad 835 \quad 836 \quad \hat{g} = \sum_{k=1}^K h_k G_k,$$

837 where G_k can be chosen as either $G_k^{\mathcal{M}_2}$ or $G_k^{\mathcal{M}_3}$. Assume:

- 838 (i) the rewards (R_1, \dots, R_K) are independent of (h_1, \dots, h_K) ;
- 839 (ii) each h_k has finite mean and variance, with $\mathbb{E}[h_k] = m_k$ and $\text{Var}(h_k) < \infty$;
- 840 (iii) $\text{Cov}(h_i G_i, h_j G_j) = 0$ for all $i \neq j$;
- 841 (iv) the rewards are nonnegative and positively correlated;
- 842 (v) each reward has finite mean and variance,

$$843 \quad 844 \quad \mathbb{E}[R_t] = \mu_t, \quad \text{Var}(R_t) = \sigma_t^2.$$

845 Then the following holds:

$$846 \quad 847 \quad \text{Var}(\hat{g}^{\mathcal{M}_2}) \geq \text{Var}(\hat{g}^{\mathcal{M}_3}),$$

848 and the variance gap admits the explicit lower bound

$$849 \quad 850 \quad 851 \quad \text{Var}(\hat{g}^{\mathcal{M}_2}) - \text{Var}(\hat{g}^{\mathcal{M}_3}) \geq \sum_{k=1}^K \left[\text{Var}(h_k) \left(\sum_{t=1}^{k-1} \sigma_t^2 + \left(\sum_{t=1}^{k-1} \mu_t \right)^2 \right) + m_k^2 \sum_{t=1}^{k-1} \sigma_t^2 \right] \geq 0.$$

852 *Proof.* We begin by computing the variance of $h_k G_k$:

$$853 \quad 854 \quad \text{Var}(h_k G_k) = \mathbb{E}[h_k^2 G_k^2] - (\mathbb{E}[h_k G_k])^2.$$

855 Because h_k and G_k are independent by Assumption (i),

$$856 \quad 857 \quad \mathbb{E}[h_k^2 G_k^2] = \mathbb{E}[h_k^2] \mathbb{E}[G_k^2], \quad \mathbb{E}[h_k G_k] = \mathbb{E}[h_k] \mathbb{E}[G_k] = m_k \mathbb{E}[G_k].$$

864 Using $\mathbb{E}[h_k^2] = \text{Var}(h_k) + m_k^2$ gives
 865

$$\begin{aligned} \text{Var}(h_k G_k) &= \mathbb{E}[h_k^2] \mathbb{E}[G_k^2] - m_k^2 \mathbb{E}[G_k]^2 \\ &= (\text{Var}(h_k) + m_k^2) \mathbb{E}[G_k^2] - m_k^2 \mathbb{E}[G_k]^2 \\ &= \text{Var}(h_k) \mathbb{E}[G_k^2] + m_k^2 \mathbb{E}[G_k^2] - m_k^2 \mathbb{E}[G_k]^2 \\ &= \text{Var}(h_k) \mathbb{E}[G_k^2] + m_k^2 (\mathbb{E}[G_k^2] - \mathbb{E}[G_k]^2) \\ &= \text{Var}(h_k) \mathbb{E}[G_k^2] + m_k^2 \text{Var}(G_k), \end{aligned}$$

872 For \mathcal{M}_2 and \mathcal{M}_3 , we have
 873

$$G_k^{\mathcal{M}_2} = S = \sum_{t=1}^K R_t, \quad G_k^{\mathcal{M}_3} = S_k = \sum_{t=k}^K R_t.$$

877 Thus

$$\begin{aligned} \text{Var}(h_k G_k^{\mathcal{M}_2}) &= \text{Var}(h_k) \mathbb{E}[S^2] + m_k^2 \text{Var}(S), \\ \text{Var}(h_k G_k^{\mathcal{M}_3}) &= \text{Var}(h_k) \mathbb{E}[S_k^2] + m_k^2 \text{Var}(S_k). \end{aligned}$$

881 Subtracting yields

$$\text{Var}(h_k G_k^{\mathcal{M}_2}) - \text{Var}(h_k G_k^{\mathcal{M}_3}) = \text{Var}(h_k) (\mathbb{E}[S^2] - \mathbb{E}[S_k^2]) + m_k^2 (\text{Var}(S) - \text{Var}(S_k)). \quad (\star)$$

884 Next we characterize the terms in (\star) . Define the prefix sum

$$P_k = \sum_{t=1}^{k-1} R_t, \quad S = P_k + S_k.$$

888 A direct expansion gives

$$\mathbb{E}[S^2] - \mathbb{E}[S_k^2] = \mathbb{E}[P_k^2] + 2\mathbb{E}[P_k S_k].$$

891 Since rewards are nonnegative and positively correlated by Assumption (iv),

$$\mathbb{E}[P_k S_k] \geq 0.$$

894 Furthermore, Assumption (v) implies

$$\mathbb{E}[P_k^2] = \text{Var}(P_k) + (\mathbb{E}[P_k])^2 \geq \sum_{t=1}^{k-1} \sigma_t^2 + \left(\sum_{t=1}^{k-1} \mu_t \right)^2.$$

899 Similarly,

$$\text{Var}(S) - \text{Var}(S_k) = \text{Var}(P_k) + 2\text{Cov}(P_k, S_k) \geq \text{Var}(P_k) \geq \sum_{t=1}^{k-1} \sigma_t^2.$$

903 Substituting these lower bounds into (\star) yields

$$\text{Var}(h_k G_k^{\mathcal{M}_2}) - \text{Var}(h_k G_k^{\mathcal{M}_3}) \geq \text{Var}(h_k) \left(\sum_{t=1}^{k-1} \sigma_t^2 + \left(\sum_{t=1}^{k-1} \mu_t \right)^2 \right) + m_k^2 \sum_{t=1}^{k-1} \sigma_t^2.$$

908 Because $\text{Cov}(h_i G_i, h_j G_j) = 0$ for $i \neq j$ by Assumption (iii), the variance of each estimator
 909 decomposes into the sum of its per-step variances:
 910

$$\text{Var}(\hat{g}^{\mathcal{M}_2}) = \sum_{k=1}^K \text{Var}(h_k G_k^{\mathcal{M}_2}), \quad \text{Var}(\hat{g}^{\mathcal{M}_3}) = \sum_{k=1}^K \text{Var}(h_k G_k^{\mathcal{M}_3}).$$

914 Since we have shown that each term satisfies

$$\text{Var}(h_k G_k^{\mathcal{M}_2}) \geq \text{Var}(h_k G_k^{\mathcal{M}_3}),$$

915 with an explicit lower bound on the difference, summing over $k = 1, \dots, K$ immediately yields both
 916 the overall variance ordering and the stated lower bound. This concludes the proof. \square
 917

918 The assumptions used in Lemma 1 are mild. Assumptions (ii) and (v) impose finite first and second
 919 moments on h_k and R_k , which are required to ensure that all variance terms are well defined and to
 920 derive the explicit variance gap. Assumption (i) reflects the natural fact that rewards are generated
 921 by the environment and do not directly depend on the stochasticity of the policy gradient estimator.
 922 Assumption (iii) holds when different time steps use independent sampling noise and allows the
 923 variance of the full estimator to decompose into a sum of per-step variances. Finally, Assumption (iv)
 924 states that rewards are nonnegative and positively correlated, a property satisfied in many episodic
 925 tasks where progress or success accumulates over time.

927 D PPO EXPERIMENTS

929 D.1 DETAILS FOR EXPERIMENTAL SETUP (PPO)

931 D.1.1 EVALUATED METHODS

933 We list all evaluated methods.

- 934 • PPO-OR (Jin et al., 2025b): vanilla PPO trained with only outcome rewards, where the
 935 trajectory-level reward is a binary signal indicating final-answer correctness, corresponding
 936 to the terminal-reward MDP \mathcal{M}_1 .
- 937 • PPO-MR (Jin et al., 2025a): vanilla PPO trained with merged intermediate and outcome
 938 rewards, where the trajectory-level reward combines intermediate rewards (retrieval correct-
 939 ness) and outcome rewards (answer correctness and format correctness), corresponding to
 940 the delayed-reward MDP \mathcal{M}_2 . The detailed reward design is provided in Section 4.1 of (Jin
 941 et al., 2025a).
- 942 • MT-PPO (ours): PPO variant trained with both intermediate and outcome rewards, where the
 943 turn-level reward design is described in Section 5.2, with $\lambda_s = 0.1$ by default, corresponding
 944 to the turn-level-reward MDP \mathcal{M}_3 .
- 945 • MT-PPO (ours): PPO variant trained with both intermediate and outcome rewards, where the
 946 turn-level reward design is described in Section 5.2, with $\lambda_s = 0.1$ by default.
- 947 • GRPO-OR (Jin et al., 2025b): vanilla GRPO trained with only outcome rewards, where the
 948 trajectory-level reward is a binary signal indicating final-answer correctness, corresponding
 949 to the terminal-reward MDP \mathcal{M}_1 .
- 950 • GRPO-MR (Jin et al., 2025a): vanilla GRPO trained with merged intermediate and outcome
 951 rewards, where the trajectory-level reward combines intermediate rewards (retrieval correct-
 952 ness) and outcome rewards (answer correctness and format correctness), corresponding to
 953 the delayed-reward MDP \mathcal{M}_2 . The detailed reward design is provided in Section 4.1 of (Jin
 954 et al., 2025a).
- 955 • OTC (Wang et al., 2025a): trains Search-R1 using GRPO and PPO with trajectory-level
 956 rewards jointly consider correctness and tool efficiency.
- 957 • StepSearch (Wang et al., 2025c): trains Search-R1 using PPO with turn-level rewards based
 958 on information gain and redundancy penalty.

959 We evaluate Search-R1 with both GRPO-OR and PPO-OR, and StepSearch using their official public
 960 checkpoints. Since Search-R1 with GRPO-MR and PPO-MR, as well as OTC, have not released their
 961 checkpoints, we directly report the results from their respective papers in Table 2.

963 D.1.2 EVALUATION METRICS

965 For each trajectory, we evaluate the following metrics:

966 **Answer correctness.** The answer correctness reward evaluates whether the extracted answer (from
 967 the `<answer>` tag) exactly matches any accepted ground-truth answer after normalization (e.g.,
 968 lowercasing and whitespace removal).

969 **Format correctness.** The format correctness reward ensures structural validity by verifying that the
 970 outputs in both the final turn and all intermediate turns comply with the specifications described in
 971 Section 5.2.

972 **Retrieval correctness.** The retrieval correctness reward evaluates whether any accepted answer
 973 appears in at least one search result (from the `<information>` tag), using case-insensitive string
 974 matching.

975 Each reward is assigned a value of 1.0 if the criterion is satisfied and 0 otherwise.

977 D.1.3 TRAINING DETAILS

979 We follow most of the experimental settings in Search-R1 (Jin et al., 2025b).

981 **PPO Training.** All experiments are conducted on 8 NVIDIA H100 GPUs. We enable gradient
 982 checkpointing and adopt Fully Sharded Data Parallel with CPU offloading. The learning rates of the
 983 policy and critic models are set to $1e-6$ and $1e-5$, respectively. Training is performed for 500 steps
 984 over 4 epochs, with warm-up ratios of 0.285 and 0.015 for the policy and critic models, respectively.
 985 The total batch size is 512, with a mini-batch size of 256 and a micro-batch size of 64 for policy
 986 updates, and a micro-batch size of 8 for critic updates. We adopt GAE with $\lambda = 1$ and $\gamma = 1$. The
 987 maximum sequence length is set to 4,096 tokens, with a maximum response length of 500 tokens and
 988 a maximum retrieved content length of 500 tokens. The KL-divergence regularization coefficient β
 989 and clipping ratio ϵ are set to 0.001 and 0.2, respectively.

990 **Rollout Generation.** We use vLLM (Kwon et al., 2023) with a tensor parallel size of 4, a GPU
 991 memory utilization ratio of 0.6, a temperature of 1.0, and a top- p value of 1.0.

992 D.2 LLM JUDGE SETUP FOR TURN-LEVEL EVALUATION (PPO)

994 In our experiments, we use gpt-oss-120b² as the judge model. We provide both outcome-level
 995 and turn-level LLM-as-judge prompts, where the outcome-level and turn-level scores are used for
 996 PPO-OR and MT-PPO training.

998 **Outcome-Level LLM-as-Judge Prompt**

1000 You are an expert evaluator for multi-turn search-augmented reasoning systems. Given a user
 1001 prompt, ground truth answer, and multi-turn generated response, determine whether the final
 1002 answer matches the ground truth.

1003 **## EVALUATION TASK**

1004 Evaluate whether the multi-turn response provides a correct final answer that matches the ground
 1005 truth.

1006 **## SCORING CRITERIA**

1007 **Score 1.0 (Correct):**

- The answer within `<answer></answer>` tags matches the ground truth.

1008 **Score 0.0 (Incorrect):**

- No `<answer></answer>` tags found, or
- The answer within `<answer></answer>` tags does not match the ground truth, or
- The answer in `<answer>` tag exceeds 5 tokens.

1015 **## OUTPUT FORMAT**

1016 Provide your evaluation using this format:

- `<reasoning>` Your step-by-step reasoning about whether the answer matches the
 1017 ground truth `</reasoning>`
- `<score>` 1.0 or 0.0 `</score>`

1021 **REQUIREMENTS:**

- First provide reasoning, then the score.
- Score must be exactly 1.0 or 0.0.

1025 ²<https://huggingface.co/openai/gpt-oss-120b>

```

1026
1027 ## EVALUATION DATA
1028 {prompt_text}
1029 {turns_text}
1030 {ground_truth_text}
1031 ## Your Evaluation
1032
1033 Turn-Level LLM-as-Judge Prompt
1034 You are an expert evaluator for multi-turn search-augmented reasoning systems. Given a
1035 user prompt, ground truth answer, and multi-turn generated response, evaluate each turn's
1036 effectiveness and compliance.
1037
1038 ## EVALUATION TASK
1039 Assess each turn's format compliance, content quality, and contribution toward the ground truth
1040 answer.
1041
1042 ## SCORING CRITERIA
1043 FINAL TURN (Last Turn) - Score Range: [-1.0 to 1.0]
1044 Format Compliance:
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

```

Format Compliance:

- Required: <think>...</think><answer>...</answer> (tags only, once each, in order)
- Answer in <answer> tag must not exceed 5 tokens

Answer Correctness:

- Correct and complete answer in <answer> tag that matches the ground truth

Scoring Rules:

- If format is incorrect: Final Turn Score = -1.0
- If format is correct, answer is incorrect: Final Turn Score = 0.2
- If format is correct, answer is correct: Final Turn Score = 1.0

INTERMEDIATE TURNS - Score Range: [-1.0 to 1.0]

Format Compliance:

- Required: <think>...</think><search>...</search><information>...</information> (tags only, once each, in order)
- Correct format: +0.1
- Incorrect format: -0.2

Information Quality:

- Relevant information in <information> tag that helps toward the ground truth answer (e.g., ground truth exists in the retrieved result within <information> tag): +0.3
- Irrelevant or unhelpful information in <information> tag: +0.0

Search Efficiency Penalty:

- Number of searches = Total count of <search> tags across all turns from Turn 1 up to and including the current turn
- Search penalty = Number of searches \times (-0.1)
- Encourages finding answers with fewer searches

Intermediate Turn Score = Format Compliance + Information Quality + Search Penalty

OUTPUT FORMAT

Provide your evaluation using ONLY these XML tags:

<reasoning>

Systematically evaluate each turn: check format compliance, assess content quality, calculate scores with clear explanations

```

1080
1081    </reasoning>
1082
1083    <score>
1084        Turn1: X.X
1085        Turn2: X.X
1086        Turn3: X.X
1087        ...
1088    </score>
1089
1090    REQUIREMENTS:
1091    • Must provide exactly {len(turns)} scores (one per turn)
1092    • Use decimal format (e.g., 0.5, -0.3, 1.0)
1093    • Use only the specified XML tags, no additional text
1094
1095    ## EVALUATION DATA
1096    {prompt_text}
1097    {turns_text}
1098    {ground_truth_text}
1099    1100    TURNS TO EVALUATE: {len(turns)}
1101    1102    ## Your Evaluation

```

D.3 ADDITIONAL EXPERIMENT RESULTS (PPO)




Figure 5: Validate reward curves recorded during training for PPO baselines and MT-PPO on the NQ and HotpotQA datasets. The rewards include answer correctness, format correctness, and retrieval correctness. Solid lines show mean reward values, while shaded regions indicate variability across five independent runs.

D.4 ROLLOUT EXMAPLES (PPO)

```

1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

```




Figure 6: Training reward curves recorded during training for PPO baselines and MT-PPO on the NQ dataset. The rewards include outcome judge reward, answer correctness (exact match), format correctness, and retrieval correctness. Solid lines show mean reward values, while shaded regions indicate variability across five independent runs.

Table 3: A rollout example of generating a response with a correct format.

Question: who will take the throne after the queen dies?

<think> I need to find out who will take the throne after the queen dies. I'll search for it. </think>

<search> who will take the throne after the queen dies </search>

<information> Doc 1>Title: "Succession to the British throne" Catholics are eligible. Queen Elizabeth II is the sovereign, and her .. Doc 2>Title: "Monarchy of the United Kingdom" Monarchy of the United Kingdom The monarchy of the United Kingdom, commonly referred to as ... Doc 3>Title: "Demise of the Crown" the Crown. This erroneous meaning is undermined by the principle in constitutional law of the continuity ... </information>

<think> I found out that after the queen dies, her heir apparent will take the throne. I need to find out who the heir apparent is for Queen Elizabeth II. </think>

<search> heir apparent for Queen Elizabeth II </search>

<information> Doc 1>Title: "Heir apparent" rank behind her brothers regardless of their ages or her age. Thus, normally, even an only ... Doc 2>Title: "Heir apparent" to the 16 thrones of Elizabeth II to absolute primogeniture, except for male heirs born ... Doc 3>Title: "Heir presumptive" or the senior member of a collateral line (if the monarch is childless); the birth ... </information>

<think> I found out that the heir apparent for Queen Elizabeth II is her eldest son, Charles, Prince of Wales. Now I can provide the answer. </think>

<answer> Charles, Prince of Wales </answer>

1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198

1199 Table 4: A rollout example of generating a response with an incorrect format, where there is no
 1200 answer in the last turn

1201

1202

Question: in which sea pearl is found in india?

1203

1204

<think> I need to find out in which sea pearl is found in India. I'll search for it. </think>

1205

1206

<search> in which sea pearl is found in India </search>
 1207 <information> Doc 1>Title: Pearl pearl: saltwater and freshwater pearls come in ... Doc
 1208 2>Title: Pearl is a particularly large one weighing 14 lb (6.4 kg). The largest ... Doc 3>Title:
 1209 "Pearl Fishery Coast" Pearl Fishery Coast The Pearl Fishery Coast refers to a coastal area ...
 1210 </information>

1211

1212

<think> I found out that pearls are found in the Indian Ocean. Now I need to find out in which
 sea pearls are found in India. </think>

1213

1214

<search> in which sea are pearls found in India </search>

1215

1216

<information> Doc 1>Title: Pearl pearl: saltwater and freshwater pearls come in many different
 ... Doc 2>Title: "Geography of India" Great Nicobar Island); Rama's Bridge, and Point Calimere.
 1217 The Arabian ... Doc 3>Title: "Bay of Bengal" the gems of Sri Lanka. Garnet and other precious
 1218 gems are also found in abundance in the ... </information>

1219

1220

<think> I found out that pearls are found in the Indian Ocean. Now I need to find out in which
 sea pearls are found in India. </think>

1221

1222

<search> in which sea pearls are found in India </search>

1223

1224

<information> Doc 1>Title: Pearl pearl: saltwater and freshwater pearls come in many different
 colors. While white, and ... Doc 2>Title: Pearl pearls". The correct definition of a South Sea
 1225 pearl – as described by CIBJO and GIA – is a pearl ... Doc 3>Title: Pearl is a particularly large
 1226 one weighing 14 lb (6.4 kg). The largest known pearl (also from ... </information>

1227

1228

<think> I found out that pearls are found in the Indian Ocean. Now I need to find out in which
 sea pearls are found in India. </think>

1229

1230

<search> in which sea are pearls found in India </search>

1231

1232

1233

1234

1235

1236


1237

1238

1239

1240

1241

1242
1243 E ADDITIONAL CASE STUDY: MULTI-TURN GAME AGENT
12441245
1246 Figure 7: Training curves on the Sokoban environment. Left: success rate over training steps. Right:
1247 effective action rate, measuring the proportion of actions that contribute to valid box movement.
1248

Metric	GRPO-OR	GRPO-MR	PPO-OR	PPO-MR	MT-PPO
Success Rate	0.0781	0.2578	0.3828	0.5078	0.6563
Effective Action Rate	0.2604	0.6771	0.6953	0.8451	0.9870

1249
1250 Table 5: Performance comparison of different RL algorithms on the Sokoban environment at step
1251 300.
12521253
1254 In this section, we present an additional case study on multi-turn game agents, focusing on the
1255 classic grid-based puzzle Sokoban (Schrader, 2018). In this puzzle, the agent must push all boxes
1256 to designated target locations. The environment is represented as a 2D grid, and the action space is
1257 discrete (up, down, left, right). The key challenge is that Sokoban is irreversible: boxes can be pushed
1258 but not pulled, so a single wrong move may lead to an unrecoverable dead-end. As a result, solving
1259 the puzzle requires the agent to reason several steps ahead rather than relying on simple navigation
1260 heuristics.
12611262
1263 In our experiments, we use Qwen2.5-VL-3B (Bai et al., 2023) as the base model. The Sokoban
1264 environment is configured with a (6×6) grid and allows up to 100 steps per episode. Each puzzle
1265 contains one box, and solving it requires at least five actions. During interaction, the agent may
1266 take up to three actions per turn, and it can interact with the environment for a maximum of three
1267 turns. For reward design, the agent receives a success reward of 10 when all boxes are placed on their
1268 target locations in the final state. At each turn, it receives a box-placement reward of 1 for each box
1269 pushed onto a target, along with a format reward of 0.5 to encourage proper visual-state reasoning and
1270 structured output. A failure penalty of -0.1 is applied at each step when the task remains incomplete.
1271 We compare our MT-PPO with GRPO-OR, GRPO-MR, PPO-OR, PPO-MR.
12721273
1274 The experimental results in Figure 7 show the training dynamics of different methods on the Sokoban
1275 environment. MT-PPO consistently achieves a higher success rate throughout training and maintains
1276 a substantially higher effective action rate, demonstrating its ability to cope with the long-horizon
1277 and irreversible structure of Sokoban. These curves highlight how MT-PPO learns more stable and
1278 purposeful action sequences during optimization.
12791280
1281 Table 5 reports the final test performance at step 300. MT-PPO achieves the highest scores on both
1282 success rate and effective action rate, outperforming all PPO and GRPO baselines by a significant
1283 margin. The results illustrate the advantages of explicit turn-level rewards and fine-grained credit
1284 assignment, further validating the generality of our approach beyond language-based tasks.
12851286
1287
1288
1289
1290
1291
1292
1293
1294
1295




Figure 8: Comparison of rollout structures in GRPO variants. GRPO-OR denotes GRPO with outcome-level rewards, while GRPO-MR denotes GRPO with merged outcome and intermediate rewards. The red and green nodes indicate the presence of turn-level rewards at those states.

F DERIVATION OF MT-GPRO FOR THE GENERAL MULTI-TURN SETTING

We now derive the MT-GPRO algorithm for the general K -turn setting. We begin by defining two types of advantages: the outcome advantage and the intermediate advantage.

- The outcome advantage captures global task completion signals. Given a group of the outcome reward $\{R_i^O\}_{i=1}^G$, it is defined as

$$A_i^O = \frac{R_i^O - \text{mean}(\{R_i^O\}_{i=1}^G)}{\text{std}(\{R_i^O\}_{i=1}^G)}, \quad (10)$$

- The intermediate advantage captures local optimization signals by comparing returns across trajectories at the same timestep. At the k -th turn ($k = 1, \dots, K-1$), given a state s_k , the algorithm samples G actions $\{a_{i,(k)}\}_{i=1}^G$, resulting in a group of intermediate rewards $R_{i,(k)}^I = R(s_k, a_{i,(k)})$. The intermediate advantage is defined as

$$A_{i,(k)}^I = \frac{R_{i,(k)}^I - \text{mean}(\{R_{i,(k)}^I\}_{i=1}^G)}{\text{std}(\{R_{i,(k)}^I\}_{i=1}^G)} \quad (11)$$

We combine these into a unified advantage that assigns credit at both global and local scales by aggregating current and future advantages:

$$A_{i,(k)}^{\text{MT-GPRO}} = \sum_{l=k}^{K-1} \alpha^{l-k} A_{i,(l)}^I + \alpha^{K-k} A_i^O \quad (12)$$

where $\alpha \in [0, 1]$ is a discount coefficient controlling the relative weight of current and future terms. This aggregated advantage is uniformly assigned to all tokens generated within the k -th turn, i.e.,

$$A_{i,1} = \dots = A_{i,t} = A_{i,(k)}^{\text{MT-GPRO}}$$

where t indexes tokens within the k -th turn,

In MT-GPRO, computing the intermediate advantages requires G rollout samples at each turn for $k = 1, \dots, K-1$. Note that rollouts are not needed at the final turn; instead, the final advantage is computed after collecting all per-state rollout samples. Therefore, over a horizon of K turns, this results in G^{K-1} rollout trajectories in total. When $k = 2$, only G rollout trajectories are required, which is the same as in the vanilla GRPO setting.

Figure 8 compares the rollout tree structures of GRPO and MT-GPRO. We observe that GRPO-OR and GRPO-MR perform per-trajectory rollouts (chain-based structures), whereas MT-GPRO performs per-state rollouts (tree-based structures). As a result, the computational complexity of GRPO scales linearly with the number of turns, while that of MT-GPRO grows exponentially with respect to the number of turns.

1350 G GRPO EXPERIMENTS

1352 G.1 TASK FORMULATION (GRPO)




Figure 9: Overview of the multi-turn LLM agent pipeline and comparison of different advantage estimation methods. The agent interacts with the tool environment across multiple steps: reasoning, tool use, and answer generation, receiving both intermediate and final rewards. GRPO is used as a representative algorithm to illustrate the different advantage estimation strategies. GRPO-OR and GRPO-MR serve as baselines with trajectory-level advantage estimation, while MT-GRPO is our proposed variant with fine-grained turn-level advantage.

To emphasize the importance of fine-grained credit assignment in multi-turn agent interactions, we formulate the task under the MDP framework, involving multiple steps of reasoning, tool use, and answer summarization for question answering. Specifically, our tool-use environment is modeled on a Wikipedia search setup, where the agent learns to leverage a Wikipedia search engine to retrieve relevant information and generate accurate answers. The goal is to improve the agent’s performance through effective integration of external tool use. Without tool calling, the agent must rely solely on its internal knowledge to answer questions, which can limit accuracy, especially for fact-based queries requiring up-to-date or domain-specific information.

To clearly illustrate the impact of credit assignment, we design a simplified two-turn tool-use environment in which the LLM agent can interact with the search tool environment for a maximum of two turns. In this setup, the agent is allowed to call the Wikipedia search engine at most once before submitting an answer to the question. Figure 9 illustrates the pipeline of the multi-turn, tool-calling LLM agent system. Given a system prompt and a question, the LLM agent first performs a reasoning step and issues a tool call, specifying both the tool name and a query derived from its reasoning. The external tool environment processes the query and returns a search result. Based on the retrieved result, the agent performs a second round of reasoning to summarize the information and generate the final answer. The whole process can be summarized as

$$\text{reasoning} \rightarrow \text{search} \rightarrow \text{result} \rightarrow \text{reasoning} \rightarrow \text{answer}$$

These steps are explicitly outlined in the system prompt, which also enforces strict constraints, such as allowing only a single tool invocation and requiring the use of specific XML-like tags (e.g., `<reasoning>`, `<tool>`, `<result>`, `<answer>`) to delineate each stage of the interaction. The full system prompt is provided in Appendix G.5. Table 7 presents an example rollout in which the agent successfully calls the search tool. If the tool name or argument format is incorrect, the tool environment returns an error message, indicated by the response beginning with “Error:”. If the agent fails to include a tool-calling command in the first reasoning step, the tool environment will not be invoked. If the XML format or tag usage is incorrect—for example, if tags are missing, nested improperly, or misnamed—the environment may fail to parse the agent’s response, resulting in an

1404 error or a skipped tool invocation. Additional rollout examples where the agent fails to call the tool
 1405 correctly are provided in Appendix G.6.

1406 Moreover, following the reformulation strategy proposed in Seed-Thinking-v1.5 (Seed, 2025), which
 1407 converts multiple-choice questions into fill-in-the-blank or short-answer formats to reduce guessing
 1408 and better evaluate reasoning ability, we adopt a similar method. Specifically, we convert our tasks
 1409 into short-answer form and evaluate the model’s responses based on exact match with the ground-truth
 1410 answers.

1411

1412 G.2 REWARD DESIGN (GRPO)

1413

1414 Figure 9 illustrates the pipeline of the multi-turn, tool-calling LLM agent system. To align with the
 1415 environment of the tool-calling LLM agent, we design two types of verifiable reward functions.

1416 **Intermediate Verifiable Rewards:** These depend solely on the first turn performed by the LLM
 1417 agent. To compute intermediate rewards, we incorporate verifiers related to tool execution and search
 1418 results. These verifiers ensure that the search engine is correctly invoked and that the ground-truth
 1419 answer appears in the retrieved results.

- 1420 • *Tool Execution Reward:* Awards 0.2 if the tool is correctly executed, determined by the
 1421 presence of properly formatted tool calls (`<tool>...</tool>`) and successful responses
 1422 (i.e., the environment’s response does not begin with “Error:”).
- 1423 • *Search Result Answer Presence:* Awards 0.5 if any accepted answer appears in the search
 1424 results returned by the tool (extracted from the `<result>...</result>` tag), using a
 1425 case-insensitive comparison.

1426 **Outcome Verifiable Rewards:** These evaluate the final model-generated responses. Specifically,
 1427 they assess both the correctness of the answer and its formatting, ensuring that the output aligns with
 1428 the expected structure and content.

- 1429 • *Final Answer Presence Reward:* Awards 0.5 if any accepted answer is present in the model’s
 1430 final response (extracted from the `<answer>...</answer>` tag).
- 1431 • *Exact Match Reward:* Awards 1.0 if the model’s answer (extracted from
 1432 `<answer>...</answer>`) exactly matches any accepted answer after standard text
 1433 preprocessing (i.e., lowercasing and stripping whitespace).
- 1434 • *XML Format Reward:* Evaluates the structural integrity of the model’s output based
 1435 on the expected schema: `<reasoning>...</reasoning>` followed by either
 1436 `<tool>...</tool>` or `<answer>...</answer>`. See the agent’s pipeline in Figure 9. Checks include: (1) the presence of at least one expected field (`<reasoning>`,
 1437 `<tool>`, `<answer>`), (2) correct spacing (no leading or trailing whitespace within tags),
 1438 (3) message starting with `<reasoning>`, and (4) message ending with `</tool>` or
 1439 `</answer>`. Partial credit is awarded based on these criteria (weighted: 40% field pres-
 1440 ence, 20% spacing, 20% correct starting tag, 20% correct ending tag), and the final score is
 1441 scaled by 0.2.
- 1442 • *XML Tag Usage Reward:* Assesses the correct usage of XML tags for the defined fields. For
 1443 each tag, the reward verifies that exactly one opening and one closing tag are present. The
 1444 reward is the proportion of correctly used tags (normalized by the number of tags checked),
 1445 scaled by 0.2.

1446 Here, both final rewards and intermediate rewards are defined as the summation of their respective
 1447 component rewards. It is easy to observe that intermediate rewards evaluate only the performance
 1448 of the agent’s first turn, whereas outcome rewards assess the quality of the entire trajectory. This
 1449 distinction leads to several characteristic scenarios:

- 1450 • *Tool Invocation with Poor Final Answer:* The agent correctly invokes a tool in the first
 1451 turn, but fails to produce a correct or well-formatted final answer, resulting in intermediate
 1452 rewards but little or no outcome reward.
- 1453 • *Incorrect or Absent Tool Use with Valid Final Answer:* The agent either skips tool usage
 1454 or invokes a tool incorrectly (e.g., due to malformed syntax or an error response), yet still

1458 generates a correct and well-structured final answer. In this case, the agent receives partial
 1459 or full outcome rewards despite earning no intermediate rewards.

1460 • *Failure Across Both Levels*: The agent neither invokes a tool correctly nor produces a valid
 1461 final answer, resulting in zero rewards and a strong negative learning signal.

1463 **G.3 EXPERIMENT SETUP (GRPO)**

1465 In our experiments, we build our codebase upon the open-source project verifiers (Brown, 2025),
 1466 which trains LLM agents for multi-turn tool-use tasks, including math calculators, code interpreters,
 1467 and search engines.

1468 **Task & Dataset.** We focus on the multi-turn reasoning and search-based tool-use task. We use the
 1469 TriviaQA dataset (Joshi et al., 2017) to train the LLM agent for answering questions by interacting
 1470 with a Wikipedia search engine. TriviaQA offers a diverse set of challenging questions, making it a
 1471 suitable benchmark for evaluating multi-turn reasoning capabilities.

1472 **Evaluated Methods** We compare our proposed MT-GPRO with vanilla GRPO.

1474 • **GRPO**: vanilla GRPO with trajectory-level advantage estimation
 1475 – **GRPO-OR**: GRPO using only outcome rewards
 1476 – **GRPO-MR**: GRPO using merged outcome and intermediate rewards
 1477 • **MT-GPRO** (ours): GPRO variant with turn-level advantage estimation using both outcome
 1478 and intermediate rewards

1480 **Training Details.** We use Qwen2.5-7B (Yang et al., 2024) as the base model. Experiments are
 1481 conducted on a node equipped with 8 NVIDIA H100 GPUs: one GPU is dedicated to rollout
 1482 generation, while the remaining seven GPUs are used for model training. Rollout generation is
 1483 handled by vLLM (Kwon et al., 2023). Model training is performed using the Huggingface TRL
 1484 implementation of GRPO (von Werra et al., 2020).

1485 **Hyperparameters.** For all methods, the number of rollout generations is set to 21. The maximum
 1486 completion length during generation is set to 1024 tokens. The KL divergence penalty is disabled by
 1487 setting $\beta = 0$. The learning rate is fixed at 1×10^{-6} . We use a per-device batch size of 12 and set
 1488 gradient accumulation steps to 4. Each batch undergoes two training iterations. The total number of
 1489 training steps is set to 300.

1491 **G.4 MAIN RESULTS (GRPO)**

1493 Figure 10 shows reward component curves during training across various algorithms. From the
 1494 answer presence and exact match reward curves, it is evident that MT-GPRO outperform GRPO-OR
 1495 and GRPO-MR, demonstrating that fine-grained credit assignment enhances the performance of
 1496 multi-turn LLM agents.

1497 The intermediate rewards, including tool execution and search result answer presence rewards, reveal
 1498 that MT-GPRO achieves 100% success in tool execution while GRPO-OR gradually stops calling
 1499 search tools in question answering tasks and achieves worse final performance. This is because
 1500 GRPO-OR does not incorporate turn-level rewards effectively in its advantage estimation, which
 1501 indicates the importance of turn-level feedback in multi-turn interaction tasks.


1502 Figures 11, 12, and 13 illustrate reward component curves during training with different algorithms,
 1503 where shaded regions represent the range between the maximum and minimum values across 10
 1504 runs, showcasing the variability in learning performance. Notably, the proposed MT-GPRO method
 1505 demonstrates lower variance during training, while GRPO-OR and GRPO-MR exhibit greater insta-
 1506 bility. An interesting observation is that the tool execution curve of MT-GPRO temporarily drops
 1507 sharply around step 230–250 but subsequently recovers and stabilizes. This demonstrates that even if
 1508 the agent forgets to call search tools in the middle of the training, it eventually learns to incorporate
 1509 them in the final stages. This finding further emphasizes the significance of credit assignment in our
 1510 proposed algorithms, contributing to more stable training.

1511 Table 6 presents the validation reward scores across different models. MT-GPRO achieves the highest
 1512 performance in all reward metrics. Compared to GRPO-MR, which reaches 0.3724 in final search

1512 Table 6: Performance comparison across different methods on reward scores evaluated on the
 1513 validation set. Values in parentheses indicate the reward range for each metric. Bold numbers indicate
 1514 the best performance for each reward type.

1516 1517 1518 1519 Model	1520 1521 1522 1523 1524 1525 1526 1527 1528 1529 1530 1531 1532 1533 1534 1535 1536 1537 1538 1539 1540 1541 1542 1543 1544 1545 1546 1547 1548 1549 1550 1551 1552 1553 1554 1555 1556 1557 1558 1559 1560 1561 1562 1563 1564 1565		1520 1521 1522 1523 1524 1525 1526 1527 1528 1529 1530 1531 1532 1533 1534 1535 1536 1537 1538 1539 1540 1541 1542 1543 1544 1545 1546 1547 1548 1549 1550 1551 1552 1553 1554 1555 1556 1557 1558 1559 1560 1561 1562 1563 1564 1565	
	Tool Execution (0-0.2)	Search Answer (0-0.5)	XML Format (0-0.2)	Exact Match (0-1)
Qwen2.5-7B-Base	0.0559	0.0934	0.1562	0.0469
Qwen2.5-7B-Instruct	0.1626	0.2814	0.1982	0.1559
Qwen2.5-7B-Base + GRPO-OR	0	0	0.04	0
Qwen2.5-7B-Base + GRPO-MR	0.2	0.3724	0.1994	0.3346
Qwen2.5-7B-Base + MT-GRPO	0.2	0.3926	0.1996	0.5010

1528 answer and 0.3346 in exact match, MT-GRPO demonstrates clear improvements, especially in exact
 1529 match with a margin of +0.1664. In contrast, GRPO-OR performs poorly across all metrics, scoring
 1530 0 in intermediate rewards and only 0.04 in XML format. These results confirm that fine-grained
 1531 credit assignment in MT-GRPO leads to better turn-level decision-making and more accurate final
 1532 outcomes in multi-turn tasks.



1620 G.5 SYSTEM PROMPT (GRPO)
16211622 System Prompt
16231624 You have access to the following tools to help solve problems:
1625 {tool_descriptions}
16261627 Follow these steps exactly once:
1628

1. Think through your reasoning inside <reasoning> tags
2. Use a tool by writing a JSON command inside <tool> tags with:
 - "name": the tool to use
 - "args": the arguments for the tool
3. You will see the tool's output inside <result> tags
4. Think through the tool's output inside <reasoning> tags
5. Based on your reasoning, provide your final answer inside <answer> tags

1636 Important:
1637

- Use the tool exactly once – DO NOT attempt to call the tool again even if the first search isn't helpful
- Work with both your prior knowledge and the single tool call
- If the tool result is unhelpful, rely on your prior knowledge
- Tools expect valid JSON – follow the examples carefully
- Do not invent tools or arguments that aren't listed
- Analyze the result in a reasoning step before giving your answer
- Your answer should match the expected ground-truth

1648 Wiki Search Tool Descriptions
16491650 wiki_search: Searches Wikipedia and returns the top matching article
1651 content.
16521653 Arguments:
1654 – query1655 G.6 ROLLOUT EXAMPLES (GRPO)
16561657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

1674

1675

1676

1677

1678

1679

1680

1681

1682

1683

1684

1685

1686

1687

1688

1689

1690

1691

1692

1693

1694

1695

1696

1697

1698

1699

1700

1701

1702

1703

1704

1705

1706

1707

1708

1709

1710

1711

1712

1713

1714

1715

1716

1717

1718

1719

1720

1721

1722

1723

1724

1725

1726

1727

1728

Table 7: A rollout example of successfully calling the search tool during training.

Question: What serial killer, who buried the remains of 26 of his victims in his Chicago crawl space, was known as the “Killer Clown”?

<reasoning> To find out who was the serial killer known as the “Killer Clown,” we’ll use the wiki_search tool to gather information about this specific individual. </reasoning>

<tool>

```

  {
    "name": "wiki_search",
    "args": {
      "query": "Killer Clown serial killer buried remains in Chicago crawl space"
    }
  }
</tool>
```

<result> John Wayne Gacy. John Wayne Gacy (March 17, 1942 – May 10, 1994) was an American serial killer who raped, tortured, and murdered at least 33 teenage boys and young men between 1972 and 1978 in Cook County, Illinois (a part of metropolitan Chicago). All of Gacy’s known murders were committed inside his Norwood Park ranch house... </result>

<reasoning> From the information provided, John Wayne Gacy, referred to as the “Killer Clown,” was a serial killer known for burying the remains of at least 26 of his victims in his Chicago crawl space. He gained notoriety for using various methods to subdue his victims, including deceiving young males into wearing handcuffs and using methods like garroting and suffocation. His final confirmed victim was identified through DNA testing in 2011. Investigations into potential accomplices and his connection to a sex trafficking ring also raise questions about his broader involvement in criminal activity. </reasoning>

<answer> John Wayne Gacy </answer>

Table 8: A rollout example of failing to call the search tool and producing a wrong answer.

Question: In the horror film Them, what type of creatures were Them?

<reasoning> In the horror film “Them!”, the main creatures were giant ants. I know this because I remember from my general knowledge that “Them!” is a 1954 science fiction horror film about giant ants and the damage they cause in a coastal town. So, to answer the question, the creatures in “Them!” were giant ants. </reasoning>

<answer> Giant </answer>

Table 9: A rollout example of generating a response with an incorrect format.

Question: Who was the Canadian jazz pianist (1925-2007), winner of eight Grammy Awards who released over 200 recordings?

<reasoning> I need to identify the Canadian jazz pianist, born in 1925 and died in 2007, who won eight Grammy Awards and released over 200 recordings. This person’s name should be easy to find with a Wikipedia search since the question specifies Grammy Awards and a long recording history. Once I find the Wikipedia article, I can read about the pianist’s career to confirm the details mentioned in the question. </reasoning>

Error: Tool command not found or invalid XML format. Please ensure correct formatting.

<answer><tool>< reasoning> From the Wikipedia search I found that Oscar Peterson, born in 1925 and deceased in 2007, was a Canadian jazz pianist who won eight Grammy Awards and released over 200 recordings. </reasoning><answer>Oscar Peterson<answer><result><answer>Oscar Peterson</answer>


1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
G.7 ADDITIONAL EXPERIMENTAL RESULTS (GRPO)

Figure 11: Curves for different training reward components during training using GRPO-OR, where shaded regions represent the range between the maximum and minimum values across 10 runs.




Figure 12: Curves for different training reward components during training using GRPO-MR, where shaded regions represent the range between the maximum and minimum values across 10 runs.




Figure 13: Curves for different training reward components during training using MT-GRPO, where shaded regions represent the range between the maximum and minimum values across 10 runs.