PROCESS REWARD INFORMED TREE ROLLOUT FOR EFFECTIVE MULTI-TURN RL

Anonymous authors
Paper under double-blind review

ABSTRACT

Reinforcement learning (RL) has become a key approach for training LLM agents, yet popular methods such as GRPO/RLOO rely on multiple independently sampled complete trajectories for advantage estimation. In long-horizon agentic tasks, such a uniform rollout strategy can waste budget on uninformative dead-end attempts, while promising intermediate states do not receive sufficient exploration. We propose Process-Scorer Guided Adaptive Tree Rollout (PATR), a quality-aware rollout framework for multi-turn agent RL. PATR uses task-appropriate process feedback to score partial trajectories, selectively branches from promising states, reuses shared prefixes, and conservatively stops degenerate paths to reduce wasted sampling. The resulting rollout groups remain compatible with standard policy optimization while providing more efficient exploration under the same training budget. We evaluate PATR on FrozenLake and the challenging SWE-Bench, which is largely unexplored by prior tree-rollout agent RL methods. Experiments show that PATR improves performance by up to $+5.0$ points on SWE-Bench and $+9.3$ points on FrozenLake, highlighting process-guided tree rollouts as an effective strategy for scalable multi-turn RL.

1 INTRODUCTION

Large language models (LLMs) are increasingly trained as agents that solve tasks through multi-turn interactions with tools, feedback, and external observations (Yao et al., 2022; Schick et al., 2023; Shinn et al., 2023). Reinforcement learning (RL) such as GRPO (Shao et al., 2024) has become a standard approach for improving such agentic capabilities. However, most existing GRPO-style agent training still relies on independently sampled complete trajectories (Li et al., 2026a). This rollout strategy allocates computation uniformly across trajectories, regardless of their intermediate quality. For long-horizon agentic tasks, such uniform allocation can waste budget on uninformative trajectories, such as repeated tool-use loops, while promising intermediate states receive insufficient exploration. As a result, the rollout group may contain redundant or low-value trajectories, leading to inefficient exploration and noisy advantage estimation.

Recent work has begun to improve rollout generation through tree-structured sampling, dynamic branching, and step-level feedback. In single-turn reasoning, tree-based methods reuse shared prefixes and explore multiple continuations from intermediate states, showing that rollout construction itself can be an important component of RL training (Surana et al., 2026; Xing et al., 2025). This idea is even more critical for multi-turn agents, where each action changes the future context and final rewards are often sparse or delayed (Feng et al., 2026; Djuhera et al., 2026).

Nevertheless, existing multi-turn tree rollout methods exhibit key limitations in how they guide exploration. One line of work branches around high-entropy tool-use steps (Dong et al., 2025b;a) or allocate rollout budget using value-based uncertainty estimates (Cao et al., 2026). While these signals encourage exploration, they do not directly measure whether a partial trajectory is making meaningful progress toward the task objective. Another line of work applies per-turn search to select the highest-scoring action at every step (Djuhera et al., 2026), which can bias the rollout distribution away from the current policy and discard informative negative examples. These limitations motivate a rollout mechanism that allocates budget based on process-level trajectory quality while preserving diverse outcomes, including failures, for policy learning.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

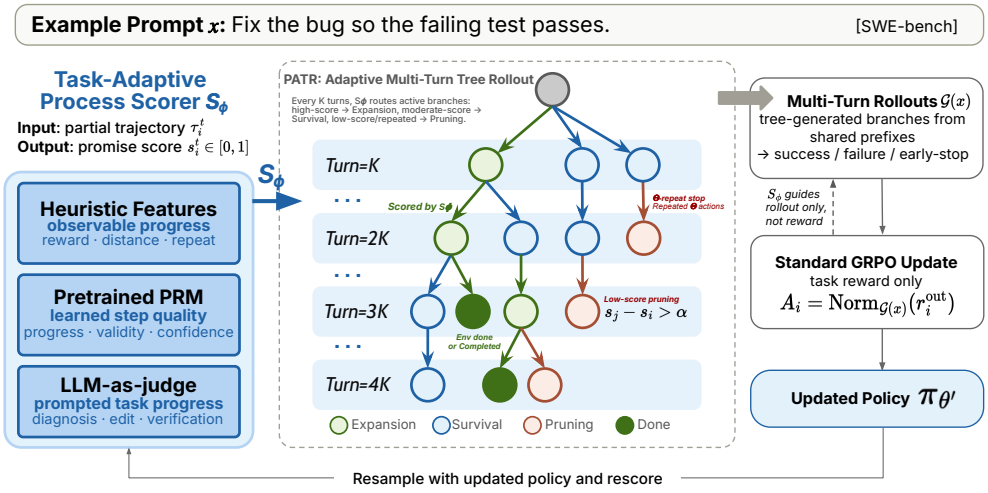


Figure 1: Overview of PATR. A task-adaptive process scorer guides expansion, survival, and pruning of partial trajectories; the rollout group is then optimized with standard GRPO on task rewards.

We propose Process-Scorer Guided Adaptive Tree Rollout (PATR), a cost-effective rollout generation framework for multi-turn agent RL that allocates more rollout budget to promising partial trajectories while terminating unpromising ones early. For each training instance, PATR constructs a tree of partial trajectories instead of sampling independent full trajectories from scratch. We periodically evaluate active branches by a task-adaptive process scorer and route into one of three paths: high-scoring branches are *expanded* by sampling multiple child continuations from the same intermediate state, moderate-scoring branches *survive* into the next iteration, and low-scoring or degenerate branches are *pruned* early. Pruned branches are retained rather than discarded, preserving rollout diversity and serving as negative examples for policy optimization. Since child branches inherit their parent prefixes, PATR naturally reuses shared computation and produces more informative rollout trajectories under a comparable sampling budget. We then employ standard GRPO-style policy optimization on the resulting rollout group with the same outcome rewards, avoiding direct optimization against the auxiliary process scorer.

We evaluate PATR on both FrozenLake and the more challenging SWE-Bench. FrozenLake provides a controlled testbed where process quality can be estimated with simple progress heuristics, while SWE-Bench evaluates whether process-guided tree rollouts can support realistic long-horizon coding-agent tasks. Experiments show that PATR improves performance by up to $+5.0$ points on SWE-Bench and $+9.3$ points on FrozenLake, highlighting the value of process-guided rollout construction for effective multi-turn RL.

Our contributions are summarized as follows:

- We propose PATR, a process-guided adaptive tree rollout framework that expands promising partial trajectories, preserves surviving branches, and prunes degenerate paths.
- We keep PATR compatible with GRPO by using process feedback only for rollout allocation while optimizing the policy with standard outcome rewards.
- We validate PATR on FrozenLake and SWE-Bench, showing consistent gains over GRPO and other rollout-generation baselines.

2 Related Work

Adaptive Rollout Generation for Agent RL. GRPO (Shao et al., 2024) and subsequent RL-based reasoning and agent methods, such as DeepSeek-R1 (Guo et al., 2025) and ARPO (Dong et al., 2025b), commonly sample multiple trajectories per task and compute group-relative advantages from outcome rewards. Recent work improves this rollout process through structured or adaptive sampling. ARPO (Dong et al., 2025b) and AEPO (Dong et al., 2025a) branch around high-entropy tool-use steps, ATPO (Cao et al., 2026) allocates rollout budget using uncertainty estimates in med-

ical dialogue, and TSR (Djuhera et al., 2026) applies tree-style search during training-time rollout generation. Tree-based reasoning methods such as Tree of Thoughts (Yao et al., 2023), RAP (Hao et al., 2023), and tree-search agents (Koh et al., 2024) also explore multiple continuations, but mainly for inference-time solution selection. The closest works are TSR and Tree-GRPO. TSR (Djuhera et al., 2026) searches over candidate actions during rollout generation and keeps only the highest-scoring candidates for training, which shifts the training distribution toward successful trajectories. Tree-GRPO (Ji et al., 2026) builds rollout trees without a process scorer and modifies GRPO with intra- and inter-tree advantage estimates. In contrast, PATR uses process scores only to decide *where* to branch, retains every sampled trajectory, including pruned ones, and leaves the GRPO advantage and objective unchanged; retaining pruned branches proves important in practice (Section 6.3).

Process Feedback and Step-Level Supervision. Process reward models, introduced in step-by-step verification work such as PRM (Lightman et al., 2024), Math-Shepherd (Wang et al., 2024), and Skywork PRM (He et al., 2025), provide intermediate quality estimates for multi-step reasoning. Several methods incorporate such signals into training objectives: PRIME (Cui et al., 2025) derives implicit process rewards, iStar (Liu et al., 2025) learns step-level rewards during agent training, WS-GRPO (Mundada et al., 2026) constructs weakly supervised prefix rewards, and Step-GRPO (Li et al., 2026b) modifies GRPO with structured process supervision. LLM-as-judge methods (Zheng et al., 2023) offer a flexible alternative when task-specific PRMs are unavailable, especially for open-ended agentic tasks where intermediate progress is difficult to specify manually. In contrast, our method uses process feedback only to allocate rollout budget, leaving the policy objective and outcome reward unchanged.

3 PRELIMINARIES

3.1 MULTI-TURN AGENTIC REINFORCEMENT LEARNING

We consider reinforcement learning for language agents that solve tasks through multi-turn interaction with external systems. Given a task prompt x, after t turns the agent observes a history $h_t = (x, a_1, o_1, \ldots, a_t, o_t)$, where a_t is the agent action and o_t is the corresponding observation. The policy samples the next action as $a_{t+1} \sim \pi_\theta(\cdot \mid h_t)$. A completed rollout $\tau = (x, a_1, o_1, \ldots, a_T, o_T)$ receives an outcome reward $r^{\text{out}}(\tau)$ upon termination. In agentic tasks, this reward is often sparse or delayed, making rollout quality critical (Wang et al., 2025).

3.2 GROUP-RELATIVE POLICY OPTIMIZATION

GRPO (Shao et al., 2024) samples a group of G rollouts for each task and computes group-normalized rewards, avoiding the need for a learned value function. For trajectory τ_i in the rollout group of task x, the outcome advantage is,

$$A_i^{\text{out}} = \frac{r_i^{\text{out}} - \mu_{\text{out}}(x)}{\sigma_{\text{out}}(x) + \varepsilon}, \tag{1}$$

where $r_i^{\text{out}} = r^{\text{out}}(\tau_i)$, and $\mu_{\text{out}}(x)$ and $\sigma_{\text{out}}(x)$ are the mean and standard deviation of outcome rewards within the group.

The policy is optimized with the clipped importance-weighted objective

$$\mathcal{L}(\theta) = \mathbb{E}_x \left[\frac{1}{G} \sum_{i=1}^{G} \frac{1}{|\tau_i|} \sum_{t=1}^{|\tau_i|} \min\Big(\rho_{i,t}(\theta) A_i^{\text{out}},\ \text{clip}\big(\rho_{i,t}(\theta), 1-\epsilon, 1+\epsilon\big) A_i^{\text{out}} \Big) \right],$$

where $\rho_{i,t}(\theta) = \frac{\pi_\theta(a_{i,t}|h_{i,t-1})}{\pi_{\theta_{\text{old}}}(a_{i,t}|h_{i,t-1})}$ is the importance ratio. Since advantages are computed relative to the sampled group, policy learning depends directly on rollout-group quality. However, standard GRPO samples complete trajectories independently from the initial task state, without considering intermediate trajectory quality.

4 PATR: PROCESS-SCORER GUIDED ADAPTIVE TREE ROLLOUT

We propose **PATR**(Process-scorer guided Adaptive Tree Rollout), a cost-effective rollout generation framework for multi-turn agent RL, that allocates more rollout budget to promising partial trajectories and terminate unpromising ones early, producing groups with diverse outcomes from shared prefixes. Given a task prompt x, PATR constructs a tree of partial trajectories and adaptively allocates rollout budget using a process scorer. Algorithm 1 summarizes the procedure.

4.1 ADAPTIVE TREE CONSTRUCTION

PATR initializes each task prompt x with B_0 active branches sampled from the current policy. Let $\mathcal{A}$ and $\mathcal{C}$ denote the sets of active and completed branches, respectively. At each iteration, every active branch is advanced for up to K interaction steps. Branches that terminate during this interval, such as by successful completion or reaching the maximum step budget, are moved from $\mathcal{A}$ to $\mathcal{C}$. For each remaining active branch τ_i^t at step t ($t \in \{K, 2K, 3K, \ldots\}$), a process scorer S_ϕ produces a scalar score

$$s_i^t = S_\phi(x, \tau_i^t) \in [0, 1], \tag{2}$$

where higher scores indicate more promising partial trajectories for further exploration.

After scoring, each active branch is routed into one of three paths based on its process score: **expansion**, **survival**, or **pruning**.

Expansion. High-scoring branches are selected into an expansion set $\mathcal{E} \subseteq \mathcal{A}$. For each selected branch τ_i^t, PATR samples M child continuations from the same intermediate state:

$$\tau_i^{t,m} \sim \pi_\theta(\cdot \mid x, \tau_i^t), \qquad m = 1, \ldots, M. \tag{3}$$

These children share the parent prefix and explore different continuations, allowing PATR to allocate additional samples to promising intermediate states. The parent branch is then removed from $\mathcal{A}$ and replaced by its children.

Survival. Branches with moderate scores remain active. They are advanced for another K steps in the next iteration and re-evaluated with updated process scores.

Pruning. PATR terminates a branch τ_i^t when its process score is substantially lower than the current active set:

$$s_i^t < \text{median}_{j \in \mathcal{A}}\, s_j^t - \alpha, \tag{4}$$

where α is a margin threshold. PATR also stops deterministically degenerate branches, such as those repeating the same action for Θ consecutive turns. Early-stopped branches are moved to $\mathcal{C}$ rather than discarded, so they remain available as negative trajectories for policy optimization.

The rollout process continues until no active branch remains. The final rollout group is constructed from all completed trajectories.

$$\mathcal{G}(x) = \mathcal{C}. \tag{5}$$

Importantly, process scores are used only to decide where to allocate additional rollout budget; they do not replace the task reward. The resulting group contains both successful and failed trajectories from shared intermediate states, providing diverse samples for relative advantage estimation.

4.2 TASK-ADAPTIVE PROCESS SCORING

PATR abstracts task-specific intermediate feedback into a scalar process score as defined in Eq. 2. This interface supports different scorer instantiations, including domain-specific heuristics, pre-trained PRMs, and LLM judges, without changing the rollout algorithm or the GRPO objective. Because PATR only uses these scores to rank branches within the same rollout group and to compare them against the group median (Eq. 4), the scorer need not be calibrated across tasks or scorer types; it only needs to order the branches of a group by how promising they are.

Heuristic scorer. When intermediate progress is directly observable, we instantiate S_ϕ as a lightweight heuristic over the current state and accumulated step rewards. For example, in navigation-style tasks, the scorer can combine the reward accumulated so far with a normalized progress measure such as distance to the goal. This instantiation requires no learned model and provides a controlled testbed when reliable progress signals are available.

Algorithm 1 PATR Rollout Generation

Require: Task prompt x; policy π_θ; process scorer S_ϕ; budgets (B_0, K, M); thresholds (α, Θ)
Initialize $\mathcal{A}$ with B_0 branches from x under π_θ; $\mathcal{C} \leftarrow \emptyset$
while $\mathcal{A} \neq \emptyset$ **do**
  Roll out each branch in $\mathcal{A}$ for up to K steps under π_θ
  Move terminated branches from $\mathcal{A}$ to $\mathcal{C}$
  if $\mathcal{A} \neq \emptyset$ **then**
    Compute $s_i^t = S_\phi(x, \tau_i^t)$ for each $\tau_i^t \in \mathcal{A}$ Eq. 2
    Select high-scoring expansion set $\mathcal{E} \subseteq \mathcal{A}$
    Let $\mathcal{L} \subseteq \mathcal{A} \setminus \mathcal{E}$ be branches satisfying the low-score or Θ-repeat criterion Eq. 4
    Sample child set $\mathcal{B} = \{\tau_i^{t,m} : \tau_i^t \in \mathcal{E},\ m = 1, \ldots, M\}$ Eq. 3
    $\mathcal{A} \leftarrow (\mathcal{A} \setminus (\mathcal{E} \cup \mathcal{L})) \cup \mathcal{B}$; $\mathcal{C} \leftarrow \mathcal{C} \cup \mathcal{L}$
  end if
end while
return $\mathcal{G}(x) = \mathcal{C}$

Pretrained PRM scorer. For tasks where a pretrained process reward model is available, we convert each partial trajectory into a problem–response format. The problem field contains the task prompt, while the response field serializes recent interaction steps, including agent actions, observations, and available step rewards. The PRM produces step-level scores $\{c_\ell\}_{\ell=1}^{L_i}$ at designated step boundaries, which are aggregated into a branch score:

$$s_i^t = f_{\text{agg}}(c_1, \ldots, c_{L_i}), \tag{6}$$

where f_{agg} can be the last-step score, the mean score, the maximum score, or a weighted combination of the last and mean scores. For long trajectories, we truncate the serialized input while preserving the task prompt and recent interaction steps, which are most relevant for deciding whether to expand the current branch.

LLM-as-judge scorer. For complex agentic tasks where no reliable hand-crafted signal or pretrained PRM is sufficient, we instantiate S_ϕ with an LLM judge. The judge receives the task prompt and a compact representation of the partial trajectory, including recent actions, observations, and automatically extracted trajectory signals. Specifically, to make scoring sensitive to different stages of SWE-style problem solving, we use phase-specific judge prompts for diagnosis, editing, and verification. The judge is instructed to output a normalized JSON score indicating whether continuing from the current state is likely to lead to a successful solution. We further blend the judge output with a lightweight heuristic based on trajectory features such as repeated actions, error or test signals in observations, evidence before editing, and targeted verification:

$$s_i^t = \lambda s_{\text{judge}}(\tau_i^t) + (1 - \lambda) s_{\text{heur}}(\tau_i^t). \tag{7}$$

This design provides robust process scores for long-horizon tool-use tasks while keeping the interface to PATR identical to the heuristic and pretrained-PRM cases. Appendix C provides the full prompts and describes the heuristic scoring signals in detail.

4.3 GRPO Training with Tree-Generated Rollout Groups

After tree rollout construction, PATR uses $\mathcal{G}(x)$ from Eq. 5 as the rollout group for GRPO. For each trajectory $\tau_i \in \mathcal{G}(x)$, we compute its outcome reward $r_i^{\text{out}} = r^{\text{out}}(\tau_i)$ and normalize rewards within the tree-generated group,

$$A_i^{\text{tree}} = \frac{r_i^{\text{out}} - \mu_{\mathcal{G}}(x)}{\sigma_{\mathcal{G}}(x) + \varepsilon}, \tag{8}$$

where $\mu_{\mathcal{G}}(x)$ and $\sigma_{\mathcal{G}}(x)$ are computed over all valid trajectories in $\mathcal{G}(x)$. The policy is then updated with the same clipped GRPO objective in Eq. 2, replacing the independently sampled rollout group with the tree-generated group.

Remark. PATR departs from the independent sampling of vanilla GRPO in two ways. First, it allocates more continuations to selected partial histories, introducing a controlled and localized bias in the rollout group; this mirrors the efficiency–bias trade-off in RL, where non-uniform allocation of samples focuses effort on more informative states Kocsis & Szepesvári (2006); Schaul et al. (2016);

Espeholt et al. (2018). Second, sibling branches share a prefix, so their outcomes are correlated and the group statistics $\mu_{\mathcal{G}}(x)$ and $\sigma_{\mathcal{G}}(x)$ in Eq. 8 are weighted toward expanded subtrees; the effective number of independent samples is therefore smaller than $|\mathcal{G}(x)|$. Both effects are bounded by design: every action is still sampled from the current policy conditioned on a realized history, the scorer only changes how many continuations each history receives, and all completed, failed, and pruned branches are retained with their outcome rewards. Moreover, siblings that diverge after a shared prefix and reach different outcomes provide exactly the within-group contrast that relative advantages rely on, while retained pruned branches keep low-reward outcomes represented in the group (Table 4). Reweighting trajectories by their branching probability to correct for this allocation is a natural extension that we leave for future work.

5 Experiments

5.1 Experimental Setup

Tasks, datasets, and models. We evaluate PATR on two settings with different levels of interaction complexity. **FrozenLake** (Brockman et al., 2016) is a grid-world navigation task in which an agent must reach a goal while avoiding holes, providing a controlled testbed for multi-turn exploration. We use the rLLM FrozenLake environment (Tan et al., 2025) with procedurally generated task prompts. **SWE-agent training** targets long-horizon software-engineering tasks, where agents must inspect repositories, edit code, and verify fixes through tool interaction. For training, we use a filtered subset of R2E-Gym-Lite-with-Difficulty, which is derived from R2E-Gym (Jain et al., 2025) by removing the most difficult instances. We evaluate on SWE-Bench Verified (Jimenez et al., 2024), a human-filtered set of 500 instances designed for reliable evaluation of coding agents and language models. For FrozenLake, we train Qwen2.5-0.5B-Instruct and Qwen2.5-3B-Instruct (Yang et al., 2024); for SWE-agent training, we use Qwen3-4B-Instruct-2507 (Yang et al., 2025).

Baselines. We compare against representative GRPO-based and tree-style rollout strategies under the same policy optimization objective. **GRPO** (Shao et al., 2024) samples independent complete trajectories for each task. **DAPO** (Yu et al., 2026) strengthens GRPO with token-level policy-gradient refinements, rejection sampling, and asymmetric clipping. **ARPO** (Dong et al., 2025b) performs adaptive branching for agentic rollouts based on uncertainty signals rather than process-level trajectory quality. **Tree-Random** uses the same tree rollout structure as PATR but replaces process-guided selection with random branch allocation.

Evaluation. For FrozenLake, we follow the standard success-rate evaluation. Each test prompt is sampled four times, and we report the average fraction of successful task instances. For SWE-Bench evaluation, we use the SWE-agent scaffold with Docker-based execution environments. Each trained agent generates one trajectory per SWE-Bench Verified instance, and we report the resolved rate, following the standard SWE-Bench protocol.

Implementation Details. All experiments are implemented with the rLLM framework (Tan et al., 2025). Across both tasks, PATR starts from four initial branches and samples two child continuations when expanding a branch. Branching is performed at fixed intervals, with $K = 5$ for FrozenLake and $K = 13$ for SWE-agent training. We set the maximum interaction length to 10 turns for FrozenLake and 50 turns for SWE-agent training, with maximum response lengths of 10K and 32K tokens, respectively. Early stopping uses the score-margin rule in Eq. 4 with $\alpha = 0.5$, together with an action-loop criterion that terminates branches repeating the same action six times.

The process scorer is instantiated according to the task. For FrozenLake, we use a lightweight progress heuristic based on accumulated rewards and shortest-path distance to the goal, treating holes as blocked cells. For SWE-agent training, we evaluate two scorer variants: **PATR-PRM**, which uses Skywork-o1-Open-PRM-Qwen-2.5-7B (He et al., 2025), and **PATR-Judge**, which uses Qwen2.5-Coder-7B-Instruct (Hui et al., 2024) as an LLM judge with phase-specific prompts. Both scorers operate on recent trajectory context; additional scoring details are provided in Appendix A. To ensure a fair comparison, GRPO and DAPO use a rollout group size of 8, while ARPO, Tree-Random, and PATR use the same number of initial branches. Detailed training hyperparameters and hardware settings are reported in Appendix A.

Method	Succ. Rate ↑	Pass@4 ↑	RLen ↓	Turns ↓
Qwen2.5-0.5B-Instruct				
GRPO	66.5	67.0	513	3.4
DAPO	71.0	71.0	562	4.2
ARPO	69.7	72.0	377	**3.3**
Tree-Random	56.3	63.0	400	3.9
PATR (ours)	**75.8**	**78.0**	**344**	3.6
Qwen2.5-3B-Instruct				
GRPO	66.8	71.0	367	**2.6**
DAPO	72.7	73.0	533	2.7
ARPO	68.7	72.0	497	3.5
Tree-Random	72.7	75.0	419	3.6
PATR (ours)	**74.3**	**79.0**	**327**	2.8

Table 1: FrozenLake results across two model scales. We report success rate, Pass@4, and average response length (RLen) for each rollout method.

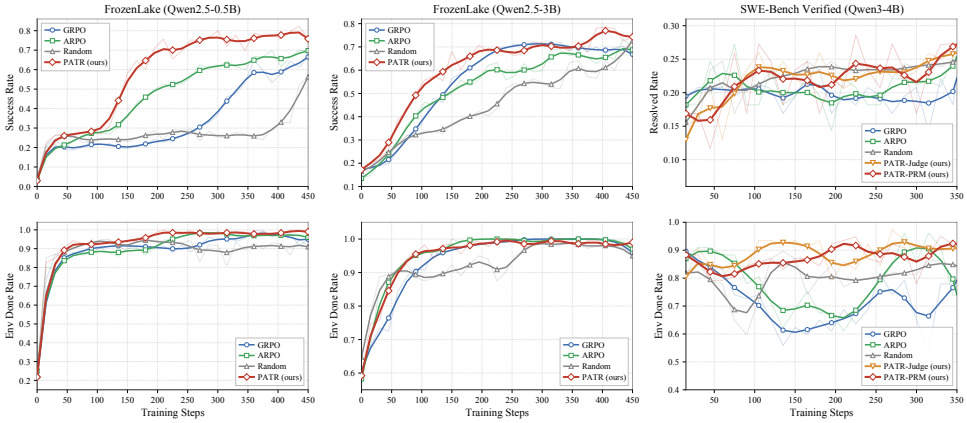


Figure 2: Training curves on FrozenLake and SWE-Bench. PATR improves performance earlier in training and maintains stronger trajectory completion throughout. "Random" denotes Tree-Random.

5.2 Main Results

FrozenLake. We report the final FrozenLake performance across two model scales in Table 1. PATR achieves the highest success rate and Pass@4 for both Qwen2.5-0.5B-Instruct and Qwen2.5-3B-Instruct. The gain is especially pronounced for the smaller model, where PATR improves the success rate from 66.5% with GRPO to 75.8%, suggesting that process-guided rollout allocation is particularly useful when the base policy is less reliable. Beyond accuracy, PATR also produces the shortest average response length on both model scales, indicating that the gains do not come from simply exploring longer trajectories. The comparison with Tree-Random further isolates the effect of process guidance. Although Tree-Random benefits from shared-prefix rollout construction and is competitive in some settings, it remains below PATR in both success rate and Pass@4. This shows that PATR gains from combining tree-structured exploration with quality-aware expansion, rather than from tree construction alone. DAPO also improves over GRPO but often produces longer trajectories, highlighting the difference between objective-level improvements and rollout-level allocation. PATR also remains robust with a deliberately weaker, hole-blind scorer, still outperforming all baselines (Appendix B).

SWE-Bench. We next evaluate the more challenging SWE-Bench Verified setting in Table 2. Both PATR variants outperform all baselines in the resolved rate. PATR-PRM achieves the best result, improving over GRPO by $+5.0$ points and over the strongest baseline ARPO by $+1.8$ points. PATR-Judge also improves over all baselines, showing that the framework can benefit from either pretrained PRM scoring or LLM-as-judge feedback. Both variants complete substantially more task instances than the baselines with fewer turns than DAPO, ARPO, and Tree-Random. This sug-

Method	Resolved Rate ↑	Env Done ↑	Turns ↓	RLen ↓
GRPO	22.2 ±1.2	79.2	25.8	18,990
DAPO	24.6 ±1.0	71.4	31.6	19,068
ARPO	25.4 ±0.9	74.0	27.3	17,848
Tree-Random	24.8 ±1.1	84.4	30.1	19,593
PATR-Judge (ours)	26.0 ±0.8	90.8	**24.1**	18,459
PATR-PRM (ours)	**27.2** ±0.9	**92.2**	24.4	**17,504**

Table 2: SWE-Bench results with Qwen3-4B-Instruct-2507. We report resolved rate, environment completion rate (Env Done), average turns, and average training response length (RLen). Resolved rates show the variation over repeated inference runs with different sampling seeds.

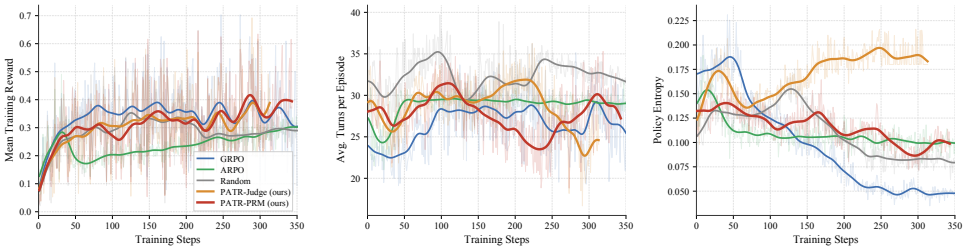


Figure 3: Training dynamics on SWE-Bench. PATR improves the balance between reward acquisition, interaction length, and policy entropy compared with flat and unguided tree-rollout baselines. "Random" denotes Tree-Random.

gests that process-guided expansion helps the policy learn more directed interactions. Tree-Random achieves a high environment-done rate but lower resolved rate, indicating that tree rollouts alone can improve completion behavior, while process-guided branch selection is needed to better align exploration with successful issue resolution. The stronger performance of PATR-PRM suggests that step-level PRM scores provide more stable branch rankings for code editing, while the judge variant remains a flexible alternative when task-specific PRMs are unavailable.

5.3 Training Dynamics

As shown in Figure 2, PATR reaches higher success rates earlier than the baselines on FrozenLake, with the largest gap for the smaller Qwen2.5-0.5B model, suggesting that process-guided branching improves sample efficiency and can compensate for weaker initial policies. On SWE-Bench, both PATR variants maintain higher resolved and environment-done rates across training, whereas flat rollout methods improve more slowly and unguided tree baselines show less consistent gains.

6 Analysis

6.1 Effect of Tree Hyperparameters

We examine how tree construction choices affect PATR-PRM on SWE-Bench Verified. As shown in Table 3, the default setting, with branch factor $M{=}2$, scoring interval $K{=}13$, and top-$k{=}2$, achieves the best resolved rate. Expanding more branches at each checkpoint slightly reduces performance, suggesting that overly broad expansion weakens the process-score filter. Increasing the branch factor to $M{=}3$ also hurts performance, especially when expansion is concentrated on a single selected branch, indicating reduced rollout-group diversity. Longer scoring intervals further degrade performance, showing that delayed rescoring limits the ability to redirect exploration during long-horizon interaction. Overall, PATR benefits from a moderate branching strategy that scores frequently while keeping expansion selective.

6.2 Effect of Pruning Margin

We further vary the pruning margin α in Eq. 4, which controls how far below the median score a branch must fall before it is stopped early. As shown in Table 5, a small margin ($\alpha{=}0.2$) prunes

Setting	M	K	Top-k	Resolved ↑
Default	**2**	**13**	**2**	**27.2**
Top-k	2	13	4	26.6
Branch factor	3	13	2	25.8
	3	13	1	22.0
Interval	2	15	2	23.4
	2	20	2	24.6

Table 3: Ablation of tree rollout hyperparameters for PATR-PRM on SWE-Bench Verified. We vary the branch factor M, scoring interval K, and number of expanded branches k per checkpoint.

Method	Group	Resolved ↑	Turns ↓
PATR-PRM	Best_n	25.8	26.4
	Keep All	**27.2**	**24.4**
PATR-Judge	Best_n	23.4	31.2
	Keep All	**26.0**	**24.1**

Table 4: Effect of rollout group construction on SWE-Bench Verified. *Keep All* retains pruned trajectories in the GRPO group, while *Best_n* excludes them.

Pruning Margin (α)	Succ. Rate ↑
0.2	61.7
0.5 (default)	**75.8**
0.8	73.5

Table 5: Effect of the pruning margin α in Eq. 4 on FrozenLake with Qwen2.5-0.5B-Instruct.

too aggressively and removes branches that could still recover, reducing the success rate to 61.7%. A large margin (α=0.8) rarely triggers pruning and moves PATR toward unguided tree rollout, yet remains competitive at 73.5%. The default α=0.5 gives the best balance, consistent with the conservative stopping design in Section 4.1. Since α only affects rollout allocation and never enters the reward, advantage, or objective, its effect is confined to which branches continue to be sampled.

6.3 Effect of Rollout Group Construction

We study whether pruned trajectories should be retained in the rollout group for GRPO training. Table 4 compares *Best_n*, which excludes pruned trajectories, with *Keep All*, which retains them together with completed branches. Keeping all trajectories consistently improves both PATR-PRM and PATR-Judge, with a larger gain for PATR-Judge from 23.4% to 26.0% resolved rate. This supports our design choice of preserving the full tree output for GRPO. Because GRPO computes advantages relative to trajectories from the same task, failed and early-pruned branches can provide useful reward contrast rather than being discarded. The lower average turns under *Keep All* suggest that pruned trajectories serve as informative negatives without full interaction budgets.

6.4 Training Signal and Policy Dynamics

Figure 3 compares training reward, interaction length, and policy entropy on SWE-Bench. Both PATR variants obtain competitive training rewards while avoiding the excessive interaction length observed in some baselines. PATR-PRM maintains strong reward trends with relatively short episodes, consistent with its higher resolved rate in Table 2. The entropy curves show different exploration patterns: flat GRPO rapidly loses entropy, while PATR-Judge maintains higher entropy for longer and PATR-PRM follows a more conservative profile. These dynamics suggest that process-guided tree rollout improves the balance between exploration, trajectory length, and reward acquisition during training.

7 Conclusion

We introduced PATR, a process-guided adaptive tree rollout framework for multi-turn agent RL that uses task-adaptive process feedback to decide where to expand, preserve, or prune partial trajectories. Note that PATR only uses process feedback to guide rollout construction without touching the provided rewards, which makes the resulting rollout group compatible with standard GRPO. Experiments on FrozenLake and SWE-Bench show consistent gains over independent-rollout and other rollout-generation baselines, highlighting process-guided tree rollouts as a promising direction for effective and scalable multi-turn agent learning.

AI USE STATEMENT

In this work, we used generative AI tools for "Implement methods". We have not used generative AI tools for "Help develop theoretical models or conceptual frameworks", "Formulate mathematical claims", "Propose or refine hypotheses", "Design or provide feedback on research methodology or experiments", "Support qualitative and thematic data analysis", "Interpret results" and "Generate synthetic data sets", "Provide critical ingredients for proving mathematical claims", "Assist in the writing of proofs", "Assist with translation", "Clean and reformat dataset" are not applicable to this work. Additionally, we used generative AI tools for "Edit a research paper to improve readability". We have reviewed all AI-assisted work. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We have made efforts to ensure that our results can be reproduced. Section 4 and Algorithm 1 specify the full PATR rollout procedure, and Section 5.1 describes the environments, datasets, models, baselines, and main hyperparameters. Appendix A reports the remaining training hyperparameters, hardware, and process-scorer implementation details, and Appendix C provides all LLM-judge prompts together with the heuristic scoring signals. All datasets and environments (the rLLM FrozenLake environment, R2E-Gym-Lite, and SWE-Bench Verified) and all models (Qwen2.5-Instruct, Qwen3-4B-Instruct-2507, Skywork-o1-Open-PRM-Qwen-2.5-7B, and Qwen2.5-Coder-7B-Instruct) are publicly available, and our implementation builds on the open-source rLLM framework. We will release our code upon publication.

REFERENCES

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

Ruike Cao, Shaojie Bai, Fugen Yao, Liang Dong, Jian Xu, and Li Xiao. Atpo: Adaptive tree policy optimization for multi-turn medical dialogue. *arXiv preprint arXiv:2603.02216*, 2026.

Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Yuchen Zhang, Jiacheng Chen, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu Yu, et al. Process reinforcement through implicit rewards. *arXiv preprint arXiv:2502.01456*, 2025.

Aladin Djuhera, Swanand Ravindra Kadhe, Farhan Ahmed, Heiko Ludwig, and Holger Boche. Tsr: Trajectory-search rollouts for multi-turn rl of llm agents. *arXiv preprint arXiv:2602.11767*, 2026.

Guanting Dong, Licheng Bao, Zhongyuan Wang, Kangzhi Zhao, Xiaoxi Li, Jiajie Jin, Jinghan Yang, Hangyu Mao, Fuzheng Zhang, Kun Gai, et al. Agentic entropy-balanced policy optimization. *arXiv preprint arXiv:2510.14545*, 2025a.

Guanting Dong, Hangyu Mao, Kai Ma, Licheng Bao, Yifei Chen, Zhongyuan Wang, Zhongxia Chen, Jiazhen Du, Huiyang Wang, Fuzheng Zhang, et al. Agentic reinforced policy optimization. *arXiv preprint arXiv:2507.19849*, 2025b.

Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymir Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures, 2018. URL https://arxiv.org/abs/1802.01561.

Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. Group-in-group policy optimization for llm agent training. *Advances in Neural Information Processing Systems*, 38:46375–46408, 2026.

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

Shibo Hao, Yi Gu, Haodi Ma, Joshua Hong, Zhen Wang, Daisy Wang, and Zhiting Hu. Reasoning with language model is planning with world model. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 8154–8173, 2023.

Jujie He, Jiacai Liu, Chris Yuhao Liu, Rui Yan, Chaojie Wang, Peng Cheng, Xiaoyu Zhang, Fuxiang Zhang, Jiacheng Xu, Wei Shen, et al. Skywork open reasoner 1 technical report. *arXiv preprint arXiv:2505.22312*, 2025.

Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, et al. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*, 2024.

Naman Jain, Jaskirat Singh, Manish Shetty, Liang Zheng, Koushik Sen, and Ion Stoica. R2e-gym: Procedural environments and hybrid verifiers for scaling open-weights swe agents. *arXiv preprint arXiv:2504.07164*, 2025.

Yuxiang Ji, Ziyu Ma, Yong Wang, Guanhua Chen, Xiangxiang Chu, and Liaoni Wu. Tree search for llm agent reinforcement learning, 2026. URL https://arxiv.org/abs/2509.21240.

Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? In *International Conference on Learning Representations*, volume 2024, pp. 54107–54157, 2024.

Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *European conference on machine learning*, pp. 282–293. Springer, 2006.

Jing Yu Koh, Stephen McAleer, Daniel Fried, and Ruslan Salakhutdinov. Tree search for language model agents. *arXiv preprint arXiv:2407.01476*, 2024.

Jiazheng Li, Yawei Wang, Qiaojing Yan, Yijun Tian, Zhichao Xu, Huan Song, Panpan Xu, and Lin Lee Cheong. Salt: Step-level advantage assignment for long-horizon agents via trajectory graph. In *Findings of the Association for Computational Linguistics: EACL 2026*, pp. 4709–4725, 2026a.

Weijie Li, Jin Wang, Liang-Chih Yu, and Xuejie Zhang. Step-grpo: Enhancing reasoning quality and efficiency via structured prm-based reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 31734–31742, 2026b.

Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *International Conference on Learning Representations*, volume 2024, pp. 39578–39601, 2024.

Xiaoqian Liu, Ke Wang, Yuchuan Wu, Fei Huang, Yongbin Li, Junge Zhang, and Jianbin Jiao. Agentic reinforcement learning with implicit step rewards. *arXiv preprint arXiv:2509.19199*, 2025.

Gagan Mundada, Zihan Huang, Rohan Surana, Sheldon Yu, Jennifer Yuntong Zhang, Xintong Li, Tong Yu, Lina Yao, Jingbo Shang, Julian McAuley, et al. Ws-grpo: Weakly-supervised group-relative policy optimization for rollout-efficient reasoning. *arXiv preprint arXiv:2602.17025*, 2026.

Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay, 2016. URL https://arxiv.org/abs/1511.05952.

Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *Advances in neural information processing systems*, 36:68539–68551, 2023.

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

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. *Advances in neural information processing systems*, 36:8634–8652, 2023.

Rohan Surana, Gagan Mundada, Xunyi Jiang, Chuhan Wang, Zhenwei Tang, Difan Jiao, Zihan Huang, Yuxin Xiong, Junda Wu, Sheldon Yu, Xintong Li, Raghav Jain, Nikki Kuang, Sizhe Zhou, Bowen Jin, Zhendong Chu, Tong Yu, Ryan Rossi, Kuan-Hao Huang, Jingbo Shang, Jiawei Han, and Julian McAuley. Generate, filter, control, replay: A comprehensive survey of rollout strategies for llm reinforcement learning, 2026. URL https://arxiv.org/abs/2605.02913.

Sijun Tan, Michael Luo, Colin Cai, Tarun Venkat, Kyle Montgomery, Aaron Hao, Tianhao Wu, Arnav Balyan, Manan Roongta, Chenguang Wang, et al. rllm: A framework for post-training language agents, 2025.

Hanlin Wang, Chak Tou Leong, Jiashuo Wang, Jian Wang, and Wenjie Li. Spa-rl: Reinforcing llm agents via stepwise progress attribution. *arXiv preprint arXiv:2505.20732*, 2025.

Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 9426–9439, 2024.

Shangyu Xing, Siyuan Wang, Chenyuan Yang, Xinyu Dai, and Xiang Ren. Lookahead tree-based rollouts for enhanced trajectory-level exploration in reinforcement learning with verifiable rewards. *arXiv preprint arXiv:2510.24302*, 2025.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. Qwen2 technical report, 2024. URL https://arxiv.org/abs/2407.10671.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022.

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

Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system at scale. *Advances in Neural Information Processing Systems*, 38:113222–113244, 2026.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623, 2023.

A Additional Implementation Details

Training details. We use the same optimization hyperparameters across methods unless otherwise specified. The learning rate is 1×10^{-6}, the clipping ratio is 0.28, the KL coefficient is 0.001, and the rollout sampling temperature is 1.0. FrozenLake is trained with batch size 64 for 400 steps, while SWE-agent training uses batch size 8 for 300 steps. FrozenLake experiments are run on one A100 node, and SWE-agent experiments are run on one H200 node.

Process scoring details. For FrozenLake, the process scorer combines accumulated step rewards with a normalized BFS shortest-path distance to the goal, where holes are treated as blocked cells. The grid, including the start, goal, and holes, is fully observable and given to the agent in the task prompt, so the scorer uses no hidden information; its score is used only for branch allocation and is never included in the policy input or the reward. For SWE-agent training, both PATR-PRM and PATR-Judge score each branch using recent trajectory context. PATR-PRM serializes the most recent 10 interaction steps and feeds them to Skywork-o1-Open-PRM-Qwen-2.5-7B, which returns step-level process scores; we use the score of the most recent step as the branch score. PATR-Judge provides Qwen2.5-Coder-7B-Instruct with the task context and a compact representation of recent actions, observations, tool types, and extracted trajectory signals. The judge outputs a trajectory-level score, which is blended with a lightweight heuristic that rewards useful signals such as targeted testing and evidence-based editing while penalizing repeated actions and premature termination. We describe the heuristic and provide all judge prompts in Appendix C.

Scoring infrastructure. For SWE-bench experiments, the PRM and LLM-judge models are hosted on separate vLLM inference servers and queried during expansion checkpoints. This separates policy rollout generation from process scoring and avoids GPU contention during training.

B Robustness to Scorer Quality

To test whether the FrozenLake gains rely on an exact progress signal, we replace the default BFS shortest-path scorer with a deliberately weaker variant that uses the Manhattan distance to the goal and ignores holes. This scorer overestimates progress for states blocked by holes and can therefore favor branches heading toward dead ends. As shown in Table 6, PATR with the Manhattan scorer still outperforms all baselines at both model scales and closely matches the default BFS scorer, while Tree-Random, which uses the same tree structure without process guidance, remains substantially lower. This indicates that the gains stem from quality-aware rollout allocation rather than access to an exact progress signal.

Method	Qwen2.5-0.5B	Qwen2.5-3B
GRPO	66.5	66.8
DAPO	71.0	72.7
ARPO	69.7	68.7
Tree-Random	56.3	72.7
PATR (Manhattan scorer)	74.9	**74.5**
PATR (BFS scorer, default)	**75.8**	74.3

Table 6: FrozenLake success rate (%) with different process scorers. The Manhattan scorer ignores holes and is strictly less informative than the default BFS scorer.

C LLM-as-Judge Scoring Details

This section provides additional details about the LLM-as-judge scorer used in PATR-Judge for SWE-Bench. At each scoring checkpoint, the judge receives the task description and a compact representation of the current partial trajectory, including recent agent actions, environment observations, inferred tool types, and automatically extracted trajectory signals. The judge outputs a scalar score in $[0, 1]$ indicating how promising the current branch is for further expansion. To make the evaluation more stage-aware, we use different prompts for diagnosis, editing, and verification phases.

Phase selection. The prompt is selected automatically from recent interaction patterns. If the recent trajectory contains test-related actions or observations, such as running `pytest`, `unittest`, or observing pass/fail messages, we use the verification prompt. If the recent trajectory contains file modification actions, such as patching, insertion, or replacement, we use the editing prompt. If the recent trajectory mainly contains repository exploration, such as search, file inspection, traceback

```
You evaluate partial SWE-agent trajectories for tree-search expansion.
Output ONLY valid JSON: {"score": <0.0 to 1.0>}. Nothing else.
```

Code 1: System prompt for the LLM-as-judge scorer.

```
Score this partial SWE-agent trajectory for tree-search branching.

Consider:
- Is the agent making measurable progress toward solving the issue?
- Are its actions grounded in evidence from the code or observations?
- Is it avoiding repetitive or aimless commands?
- If we continue from this state, how likely will this branch lead to a
    correct fix?

0.0 = stuck, looping, or wrong direction
0.5 = some relevant work but unclear progress
1.0 = clearly on track, strong evidence of progress

Task:
{task}

Trajectory:
{trajectory}

Output ONLY: {"score": <0.0 to 1.0>}
```

Code 2: Default judge prompt.

analysis, or command-line reading, we use the diagnosis prompt. When no specific phase is detected, we use the default prompt. This design allows the judge to evaluate the type of progress that is most relevant to the current stage of problem solving.

Heuristic score. In addition to the LLM score, we compute a lightweight heuristic score from recent trajectory signals. The heuristic starts from a neutral value and adjusts the score according to observable behaviors. It penalizes repeated identical actions, premature finish actions, command failures, and edits made without prior evidence from search, file inspection, or testing. It rewards signals associated with useful progress, including targeted test execution, passing tests, informative tracebacks, code or file information discovered during exploration, and edits made after relevant evidence has been collected. The final branch score is a weighted combination of the LLM judgment and the heuristic score, with the LLM judgment receiving the larger weight. The heuristic acts as a stabilizing signal, especially when the LLM judge assigns overly optimistic scores to repetitive or weakly grounded trajectories.

Judge prompts. We list the system prompt and all phase-specific scoring prompts below 1, 2, 3, 4, and 5. The placeholders `{task}` and `{trajectory}` are filled with the task description and the compact trajectory representation at each scoring checkpoint.

```
Score this partial trajectory during EXPLORATION / DIAGNOSIS.

Consider:
- Is the agent searching in the right files and directories?
- Is it narrowing down the bug location with each step?
- Is it gathering evidence (reading code, reproducing the bug) rather
    than guessing?
- Is it avoiding repeated searches that yield no new information?

0.0 = aimless browsing, wrong files, repetitive commands
0.5 = relevant exploration but slow progress
1.0 = efficiently locating the root cause

Task:
{task}

Trajectory:
{trajectory}

Output ONLY: {"score": <0.0 to 1.0>}
```

Code 3: Judge prompt for the diagnosis phase.

```
Score this partial trajectory during EDITING.

Consider:
- Is the agent editing the file that evidence pointed to?
- Is the change minimal and tied to the diagnosed root cause?
- Did the agent collect enough evidence before editing, or is it
    guessing?
- Could this edit introduce new bugs or break other functionality?

0.0 = editing blind, wrong file, speculative changes
0.5 = reasonable edit but unclear if it addresses root cause
1.0 = precise, evidence-based fix in the right location

Task:
{task}

Trajectory:
{trajectory}

Output ONLY: {"score": <0.0 to 1.0>}
```

Code 4: Judge prompt for the editing phase.

```
Score this partial trajectory during TESTING / VERIFICATION.

Consider:
- Is the agent running tests relevant to the issue (not random test
    suites)?
- If tests fail, does the agent appear to understand the failure?
- If tests pass, do they actually cover the bug that was fixed?
- Is the agent verifying the fix rather than declaring premature success?

0.0 = irrelevant tests, ignoring failures, false success
0.5 = running some tests but coverage of the fix is unclear
1.0 = targeted verification that confirms the fix works

Task:
{task}

Trajectory:
{trajectory}

Output ONLY: {"score": <0.0 to 1.0>}
```

Code 5: Judge prompt for the verification phase.