

TRACE: A Unified Rollout Budget Allocation Framework for Efficient Agentic Reinforcement Learning

Heming Zou^{1,2*}, Qi Wang^{1†}, Yun Qu^{1,2*}, Yuhang Jiang¹, Lizhou Cai¹, Yixiu Mao¹, Ru Peng^{2*}, Xin Xu², Weijie Liu², Kai Yang², Saiyong Yang^{2†}, Xiangyang Ji^{1†}

¹Tsinghua University, ²LLM Department, Tencent

Correspondence: cheemswang@mail.tsinghua.edu.cn, stevesyang@tencent.com, xyji@tsinghua.edu.cn

Abstract

Reinforcement learning with verifiable rewards (RLVR) is a promising approach for enhancing reasoning and agentic behavior in large language models. However, rollout-intensive policy optimization is often limited by insufficient reward contrast, arising when overly simple or complex prompts generate low-variance feedback and when outcome-only rewards assign the same terminal assessment to every decision in a multi-turn rollout. Past efforts have focused on allocating available rollout resources to promising prompts, yet they only leverage sample informativeness at the prompt level and neglect variation in prefix-level informativeness across turns within the same rollout. This work targets multi-turn agentic RL by modeling each ReAct-style thought-action-observation turn as a semantically distinct node, allowing budget allocation to extend from prompt roots to turn-level prefixes with further continuations, which naturally forms tree-structured rollouts. We introduce Tree Rollout Allocation for Contrastive Exploration (TRACE), a unified rollout allocation framework that enhances reward contrast within a fixed sampling budget. Technically, TRACE allocates rollout budget to both prompt roots and intermediate prefixes that are most likely to yield mixed terminal rewards. A shared generalizable predictor estimates conditional success probability at these anchors from prefix histories to guide this allocation. The resulting adaptive tree structure enriches outcome-only feedback and amplifies the policy-update signal. Empirically, TRACE achieves competitive performance and efficiency gains on typical agentic benchmarks, e.g., improving Qwen3-14B Multi-Hop QA average accuracy by 2.8 points over competitive baselines at equal sampling cost. Code is available at <https://github.com/gfyddha/TRACE>.

*Work completed during an internship at Tencent.

†Corresponding authors.




Figure 1: TRACE redirects a fixed rollout budget toward contrast-rich roots and prefixes, converting scalable outcome rewards into denser mixed-reward contrast and stepwise preference pairs than uniform allocation.

1 Introduction

Reinforcement learning with verifiable rewards (RLVR) has become a key mechanism for incentivizing reasoning and agentic behavior in large language models (LLMs), using verifiable feedback to shape exploration over solution paths and environment interactions (Jaech et al., 2024; Guo et al., 2025a). Despite its success, RLVR is computationally expensive because rollouts require long Chain-of-Thought (CoT) (Wei et al., 2022) generation for reasoning and interleaved environment interactions for agentic tasks. Deciding where to spend limited rollout budget is therefore a central design problem (Zheng et al., 2025; Lin et al., 2025).

In practice, most scalable RLVR pipelines rely on outcome rewards rather than process rewards, avoiding hand-crafted intermediate reward rules that are vulnerable to reward hacking (Shao et al., 2024). The learning signal provided by such rollouts, however, is highly uneven in informativeness and sparse for credit assignment. At the prompt level, overly easy or hard instances tend to produce low-variance outcome rewards (Yu et al., 2025). At the trajectory level, assigning a single terminal

reward to a long rollout leaves little local contrast for credit assignment across generated tokens and agent turns (Lightman et al., 2024). Prior work alleviates prompt-level inefficiency through prompt selection (Zheng et al., 2025; Gao et al., 2025; Bae et al., 2026), which decides which prompts to sample, and rollout allocation (Li et al., 2025c), which decides how many rollouts each prompt receives. Together, they can be unified as scoring a candidate set and assigning rollout budget toward prompts likely to produce contrastive outcome variation: a prompt assigned zero rollouts is skipped, whereas a selected prompt assigned a larger positive budget receives more independent root rollouts. Yet this allocation strategy acts only at the prompt root. Once a prompt receives budget, each rollout is still generated as an atomic trajectory.

Multi-turn agentic RL exposes a finer allocation unit, the turn. While single-turn reasoning can be segmented at token or sentence boundaries, ReAct-style (Yao et al., 2022) interaction naturally packages each decision as a thought-action-observation node. A complete rollout therefore exposes semantically meaningful prefixes that can be revisited for branching. These visited prefixes become candidate branching points for extra exploration. Some prefixes are unlikely to produce new outcome variation, whereas others are likely to yield counterfactual continuations with different final rewards. Technically, allocating budget at this level is not a matter of sampling more independent trajectories; it is a choice of which prefix occurrences should receive additional continuation branches.

This viewpoint naturally turns flat rollout collection into tree-structured rollouts. Prompt roots are depth-zero anchors and internal prefixes are non-root turn-level anchors, so root-level allocation becomes a special case of allocating budget over tree nodes. The shared principle is *mixed-reward contrast construction*: allocate rollout budget to anchors whose descendant sets are most likely to contain both successful and failed outcomes. At the root, this principle unifies prompt filtering and rollout-count allocation, with zero budget rejecting a prompt and positive group budget setting its root rollout count. Inside an active prompt, the same principle allocates extra branches to prefixes likely to produce opposite-outcome siblings.

Figure 1 summarizes this view: under a fixed rollout budget, allocating branches to contrast-rich roots and prefixes yields more mixed terminal outcomes than uniform allocation over the rollout tree.

We instantiate this idea as **Tree Rollout Allocation for Contrastive Exploration (TRACE)**, a unified rollout budget allocation framework for efficient RLVR in multi-turn agentic tasks. TRACE first performs global root allocation over a large candidate pool, then local tree expansion over the prefix occurrences exposed by the resulting rollouts. Both stages are guided by a shared generalizable predictor that estimates conditional success probability from prefix histories. The resulting rollout trees turn scalable outcome rewards into implicit stepwise preference pairs for tree-aware policy optimization. Across three agentic settings: Mathematical Reasoning, Multi-Hop QA, and Function Calling, TRACE improves accuracy over uniform sampling under the same rollout budget by directing samples toward anchors with reward contrast.

In summary, our contributions are three-fold:

1. We formulate **mixed-reward contrast allocation** as a unified view of sample-efficient RLVR, incorporating prompt filtering, rollout-count allocation, and prefix-level branching as strategic budget decisions over rollout-tree anchors.
2. We derive **root and prefix allocation utilities** based on a shared descendant-set objective, distributing budget according to the likelihood that an anchor’s descendants will yield both successful and failed outcomes.
3. We instantiate this view in **TRACE**, a predictor-guided, budget-constrained tree rollout framework integrating global root allocation with local prefix expansion, enhancing implicit credit signals for sample-efficient policy optimization.

2 Preliminaries

2.1 Multi-Turn RLVR as Prefix Histories

This work investigates multi-turn agentic reinforcement learning within the ReAct framework (Yao et al., 2022). We define the input prompt as the root node x of the rollout tree. At each turn $t = 1, \dots, T$, the model produces a thought τ_t and an action a_t , to which the environment responds with an output o_t . We then encapsulate the complete turn as $n_t := \langle \tau_t, a_t, o_t \rangle$. The history available after turn t is described as

$$\mathcal{H}_t := (x, n_1, \dots, n_t), \quad \mathcal{H}_0 := x. \quad (1)$$

With policy model π_θ and environment transition P_{env} , the rollout dynamics can be expressed as the

following hierarchical generative process:

$$\begin{aligned} (\tau_t, a_t) &\sim \pi_\theta(\cdot \mid \mathcal{H}_{t-1}), \\ o_t &\sim P_{\text{env}}(\cdot \mid \mathcal{H}_{t-1}, \tau_t, a_t). \end{aligned} \quad (2)$$

A complete rollout is identified by its terminal history $\mathcal{H}_T$, and the RLVR framework assigns a binary terminal reward $r(\mathcal{H}_T) \in \{0, 1\}$. Thus, the underlying training objective is formulated as

$$\max_{\theta} \mathbb{E}_{x \sim \mathcal{D}, \mathcal{H}_T \sim \mathbb{P}_{\pi_\theta}(\cdot \mid x)} [r(\mathcal{H}_T)]. \quad (3)$$

For any prefix $\mathcal{H}_t$, we define the conditional success probability of solving the prompt as

$$V_t^\pi := \mathbb{E}_\pi[r(\mathcal{H}_T) \mid \mathcal{H}_t]. \quad (4)$$

A prompt is the shortest history $\mathcal{H}_0$, while later prefixes add the thoughts, actions, and observations already produced during the rollout.

2.2 Initialize-then-Expand Rollout Trees

To remain suited for parallelized LLM inference engines, tree rollout methods commonly follow an initialize-then-expand construction (Ji et al., 2025; Hou et al., 2025). After an allocation rule assigns each prompt x_i a root rollout count m_i , Stage 1 samples complete bare rollouts. For rollout j , let $T_{i,j}$ be its terminal turn index and let $\mathcal{H}_{i,j,t}$ denote its prefix after turn t :

$$\mathcal{H}_{i,j,T_{i,j}} \sim \pi_\theta(\cdot \mid x_i), \quad j = 1, \dots, m_i, \quad (5)$$

and records terminal rewards $r_{i,j}$. Setting $m_i = 0$ skips prompt x_i in Stage 1, so Eq. (5) is not executed and no bare rollout index j is defined. Each visited prefix $\mathcal{H}_{i,j,t}$ can then serve as a branching anchor: if prefix expansion assigns $K_{i,j,t}$ continuations to it, Stage 2 samples terminal suffixes

$$S_{i,j,t}^{(b)} \sim \pi_\theta(\cdot \mid \mathcal{H}_{i,j,t}), \quad b = 1, \dots, K_{i,j,t}. \quad (6)$$

The original suffix is the factual branch under the prefix, while newly sampled suffixes $S_{i,j,t}^{(b)}$ are counterfactual branches. This abstraction makes prompt roots and internal prefixes comparable budget-allocation anchors: both are conditioning contexts with terminal descendants below them. For an active prompt x_i , the corresponding nonterminal prefix-anchor index set is

$$\mathcal{A}_i := \{(j, t) : 1 \leq j \leq m_i, 1 \leq t < T_{i,j}\}, \quad (7)$$

which excludes terminal leaves and records prefix occurrences continuation branching. For a visited prefix occurrence $(j, t) \in \mathcal{A}_i$, setting $K_{i,j,t} = 0$ skips Stage 2 branching at that anchor, so Eq. (6) is not executed and no continuation index b is defined.

3 Rollout Allocation through the Lens of Contrast Construction

In outcome-reward RLVR (Guo et al., 2025a), more rollouts are not automatically more informative. A broad line of process-supervision, preference-learning, and tree-based RL methods reflects the same lesson: learning becomes denser when the model can compare alternatives rather than consume isolated terminal labels (Lightman et al., 2024; Lai et al., 2024; Zhang et al., 2024). If all descendants of an anchor receive the same terminal reward, group-relative or tree-aware updates have little local preference signal (Shao et al., 2024; Yang et al., 2025b): sampled continuations do not tell the optimizer which decisions should be preferred under the same conditioning context. When rollouts sharing the same prompt root or prefix disagree in terminal outcome while only the continuation differs, they yield implicit pairwise preferences over sibling continuations rather than a single sparse terminal reward shared across the entire rollout. We therefore treat mixed-reward contrast construction as the allocation objective: a fixed rollout budget is useful to the extent that it induces such pairwise comparison at prompt roots and visited prefixes.

3.1 Prefix Information for Better Difficulty Prediction

Prior sample-efficient RLVR methods use prompt-level difficulty prediction to filter prompts and allocate root rollouts (Zheng et al., 2025; Gao et al., 2025; Bae et al., 2026; Li et al., 2025c). This motivates the pre-sampling prediction problem at prefix anchors in our tree setting: just as prompt roots are scored before Stage 1 rollout allocation, visited prefixes should be scored before Stage 2 continuation allocation. Each prefix history $\mathcal{H}_t$ refines the prompt root with observed interaction, so it conditions the same conditional success probability V_t^π on richer state information and supports sharper forecasts of group outcome variation and downstream reward contrast.

Proposition 1 (Prefix information improves group difficulty prediction). *For $m \geq 1$, let $\bar{R}_{t,m}$ be the average terminal reward of m independent continuations sampled from prefix $\mathcal{H}_t$. For any squared-loss predictor f that maps a history to a scalar estimate, define $\mathcal{E}_{t,m}^* := \inf_f \mathbb{E}[(\bar{R}_{t,m} - f(\mathcal{H}_t))^2]$. Then $\mathbb{E}[\bar{R}_{t,m} \mid \mathcal{H}_t] = V_t^\pi$, and for any $t < T$,*

$$\mathcal{E}_{t+1,m}^* \leq \mathcal{E}_{t,m}^*. \quad (8)$$




Figure 2: **Contrastive-allocation diagnostics.** Using rollout results collected over several steps with Qwen3-8B under the Multi-Hop QA (HotpotQA) tree-sampling setting, panel (a) shows that many prompt-root and prefix anchors have empirical success rate $\hat{p}$ near 0 or 1, where outcome contrast is scarce. Panel (b) measures each anchor’s pair contrast as $\hat{p}_h(1 - \hat{p}_h)$; the x -axis is the fraction of rollout budget assigned to the top-ranked anchors, and the y -axis is the fraction of total pair contrast $\sum_h \hat{p}_h(1 - \hat{p}_h)$ covered by those anchors, compared with uniform allocation. Panel (c) evaluates the same prefix-conditioned predictor at different observed prefix depths, showing lower group-success MSE at deeper depths and more reliable prefix-conditioned allocation.

Consequently, $\mathcal{E}_{t,m}^* \leq \mathcal{E}_{0,m}^*$ for all t . $t = 0$ and $m = 1$ recover prompt-level difficulty prediction and single-rollout reward prediction, respectively.

Intuitively, Proposition 1 asks how well one can forecast the average terminal reward of m fresh continuations sampled after a history $\mathcal{H}_t$. The optimal mean-squared error $\mathcal{E}_{t,m}^*$ measures this predictability using all information in $\mathcal{H}_t$, with target $\mathbb{E}[\bar{R}_{t,m} | \mathcal{H}_t] = V_t^\pi$. Because $\mathcal{H}_{t+1}$ extends $\mathcal{H}_t$ with one additional turn, prediction cannot worsen as more interaction is observed ($\mathcal{E}_{t+1,m}^* \leq \mathcal{E}_{t,m}^*$), so prefix-level scoring is at least as informative as prompt-only scoring and strictly better once intermediate turns are available. This justifies scoring visited prefixes before Stage 2 continuation allocation. We attach the proof in Appendix D.1.

3.2 A Shared Mixed-Outcome Principle for Roots and Prefixes

Given such root- and prefix-level predictions, the allocation target is the same at both scales: at a prompt root, the allocator asks whether fresh rollouts will form a diverse outcome group; at a visited prefix, it asks whether additional continuations from the same history can disagree with the factual suffix. This gives a single contrast-seeking rule: anchors near deterministic success or failure provide little contrast, while anchors with intermediate conditional success probability are more likely to expose both successful and failed descendants.

This prefix-level view explains how outcome-only rewards can still induce local credit. Shared-prefix branches keep the past fixed and vary only the future continuation, so opposite terminal outcomes become a local comparison for tree-aware

optimizers. From a martingale view of conditional success, this remaining contrast has a simple form: each additional turn revises the forecast of terminal success, and the squared forecast revisions accumulate the remaining uncertainty along the suffix rather than only at the terminal reward.

Proposition 2 (Prefix Uncertainty as Remaining Contrast Potential). *Let $\mathcal{F}_t := \sigma(\mathcal{H}_t)$, $Z_t := V_t^\pi = \mathbb{E}_\pi[r(\mathcal{H}_T) | \mathcal{F}_t]$, and let $[Z]_{t:T} := \sum_{s=t}^{T-1} (Z_{s+1} - Z_s)^2$ denote the quadratic variation of the conditional success forecast along the future rollout suffix. For any prefix $\mathcal{H}_t$,*

$$\mathbb{E}_\pi [[Z]_{t:T} | \mathcal{F}_t] = V_t^\pi (1 - V_t^\pi). \quad (9)$$

Proposition 2 identifies the conditional expected quadratic variation with Bernoulli variance: $\mathbb{E}_\pi [[Z]_{t:T} | \mathcal{F}_t] = V_t^\pi (1 - V_t^\pi)$. Thus $V_t^\pi (1 - V_t^\pi)$ is not merely a static uncertainty score, but the expected accumulated movement in conditional success probability below the prefix, and therefore measures downstream contrast potential. Fig. 2 confirms the empirical picture: prefix information improves group-success prediction, and a small subset of ranked anchors captures much of the pair contrast. We prove Proposition 2 in Appendix D.2.

3.3 Contrast Values as Update Activation

The allocation rules are chosen before the optimizer sees gradients, but they target the activation frontier where gradients become informative. Under binary rewards, a root or prefix anchor activates local gradient contrast only when its descendants expose both success and failure.

This observation connects contrast allocation to optimizer-visible gradients. Let $\mathcal{A}_{\text{root}} = \{x_i\}_{i=1}^B$

and let $\mathcal{A}_{\text{pref}}$ denote the local prefix-anchor set $\mathcal{A}_i$ for an active prompt; aggregation over prompts sums these local terms. For stage $q \in \{\text{root}, \text{pref}\}$ and allocation b , we formulate the aggregate local policy-gradient contribution in general form as

$$\begin{aligned} G_q(b) &= \sum_{h \in \mathcal{A}_q} \mathbb{I}\{\text{Act}(h, b_h)\} \tilde{G}_h(b_h), \\ \tilde{G}_h(b_h) &= \sum_{\mathcal{H} \in \mathcal{D}_h(b_h)} \omega_{\mathcal{H}} \nabla_{\theta} \log \pi_{\theta}(\mathcal{H} \mid h), \end{aligned} \quad (10)$$

Here $\mathcal{D}_h(b_h)$ is the sampled terminal-descendant set below anchor h , $\omega_{\mathcal{H}}$ is the optimizer-dependent contrast weight, $\mathbb{I}\{\cdot\}$ is the indicator function, and $\text{Act}(\cdot, \cdot)$ is the activation event: $\text{Act}(h, b_h)$ holds when the sampled descendants below anchor h under budget b_h contain both outcomes.

Proposition 3 (Activation allocation dominates uniform). *Let b_q^* be the TRACE allocation for stage $q \in \{\text{root}, \text{pref}\}$, and let b_q^u be any feasible uniform allocation at the same budget. Write $G_q^* := G_q(b_q^*)$ and $G_q^u := G_q(b_q^u)$. Assume normalized conditional gradient energy, meaning $\mathbb{E}[\|\tilde{G}_h(b_h)\|^2 \mid \text{Act}(h, b_h), h] = c_q > 0$ for all anchors in stage q . If the combined update also satisfies $\mathbb{E}\langle G_{\text{root}}^*, G_{\text{pref}}^* \rangle \geq \mathbb{E}\langle G_{\text{root}}^u, G_{\text{pref}}^u \rangle$, then*

$$\begin{aligned} \mathbb{E}\|G_q^*\|^2 &\geq \mathbb{E}\|G_q^u\|^2, \quad q \in \{\text{root}, \text{pref}\}, \\ \mathbb{E}\|G_{\text{root}}^* + G_{\text{pref}}^*\|^2 &\geq \mathbb{E}\|G_{\text{root}}^u + G_{\text{pref}}^u\|^2. \end{aligned} \quad (11)$$

Inequalities are strict when uniform is not optimal for the corresponding stage.

Proposition 3 suggests the main allocation consequence: the first line of Eq. (11) shows that each stage strengthens the expected squared local gradient signal over uniform allocation, and the second line shows that their combination strengthens the full tree update when the two gradient sources do not cancel. The proof in Appendix D.3 shows the factorization behind this statement: under standard outcome-reward updates, the local gradient contribution is zero unless opposite outcomes are observed below the same anchor, so the expected squared gradient norm equals an activation probability times a conditional gradient scale.

4 TRACE

The previous section identifies the useful sampling targets: roots and prefixes whose descendants are likely to disagree in terminal reward. TRACE instantiates this principle as a global-to-local allocation procedure, summarized in Fig. 3 and Algorithm 1. It solves root allocation over a candidate

prompt pool and local prefix expansion over visited prefix occurrences. The predictor $\tilde{V}_{\psi}(\mathcal{H})$ estimates the same conditional success probability written as V_t^{π} for a turn-indexed history $\mathcal{H}_t$ in Section 2. The resulting trees supervise this predictor through recursive targets and provide root-/prefix-level comparisons for tree-aware policy optimization.

4.1 Global Root Allocation

At step s , TRACE allocates a root rollout budget over a candidate prompt set $\mathcal{B}_s = \{x_1, \dots, x_B\}$, where $B = |\mathcal{B}_s|$ is the batch size and each prompt receives a count in $\mathcal{M} := \{0\} \cup \{2, \dots, M\}$. Let $v_i := \tilde{V}_{\psi}(x_i)$ be the predicted root success probability. For $m \geq 2$, TRACE assigns

$$V_{\text{root}}(x_i, m) = 1 - v_i^m - (1 - v_i)^m, \quad (12)$$

and sets $V_{\text{root}}(x_i, 0) = V_{\text{root}}(x_i, 1) = 0$. This is the predicted probability that m root rollouts for x_i contain success and failure. Formally, root allocation solves the budget-constrained problem:

$$\begin{aligned} \max_{m_1, \dots, m_B} \quad & \sum_{i=1}^B V_{\text{root}}(x_i, m_i) \\ \text{s.t.,} \quad & \sum_{i=1}^B m_i = M, \quad m_i \in \mathcal{M}. \end{aligned} \quad (13)$$

Setting $m_i = 0$ skips the prompt. We exclude $m_i = 1$ from $\mathcal{M}$ because typical tree-aware policy optimization is group-based and requires at least two root rollouts per prompt. Any $m_i \geq 2$ activates the prompt and sets the number of bare rollouts generated below it.

4.2 Local Prefix Expansion

After active prompt x_i finishes its m_i bare rollouts, let $r_{i,j} \in \{0, 1\}$ be the terminal reward of bare rollout j . Its prefix after turn t is $\mathcal{H}_{i,j,t}$, with $\mathcal{H}_{i,j,0} = x_i$. Recall the prefix-anchor index set $\mathcal{A}_i$ defined in Eq. (7), each index $(j, t) \in \mathcal{A}_i$ denotes a visited nonterminal prefix together with the factual terminal outcome already observed below it. We allocate a fixed local budget

$$\sum_{(j,t) \in \mathcal{A}_i} K_{i,j,t} = m_i N, \quad (14)$$

so Stage 2 does not wait for or compete with other prompts once prompt x_i is ready.




Figure 3: **Framework overview of TRACE.** A prefix value predictor scores prompt roots and visited prefixes, TRACE solves budgeted root allocation and prompt-local prefix expansion, and the resulting rollout trees provide recursive value targets and root-/prefix-level comparisons for tree-aware policy optimization.

Using the same predictor for the prefix success probability, assigning k continuations to prefix occurrence (j, t) has value

$$V_{\text{pref}}(i, j, t, k) := 1 - \left[r_{i,j} \tilde{V}_\psi(\mathcal{H}_{i,j,t}) + (1 - r_{i,j})(1 - \tilde{V}_\psi(\mathcal{H}_{i,j,t})) \right]^k, \quad (15)$$

namely the probability that at least one new continuation flips the observed reward. For each active prompt x_i , local prefix expansion solves

$$\begin{aligned} \max_{\{K_{i,j,t}\}_{(j,t) \in \mathcal{A}_i}} \sum_{(j,t) \in \mathcal{A}_i} V_{\text{pref}}(i, j, t, K_{i,j,t}) \\ \text{s.t.} \quad \sum_{(j,t) \in \mathcal{A}_i} K_{i,j,t} = m_i N. \end{aligned} \quad (16)$$

Both allocation problems are solved by efficient dynamic programming tools, whose cost is negligible relative to rollout generation. Expanding (j, t) fixes the prefix $\mathcal{H}_{i,j,t}$ and resamples a new continuation after that prefix. By design, terminal leaves are meaningless to expand.

This factorization is system-aware: a fully global prefix allocator over $\bigcup_i \mathcal{A}_i$ would wait for all active prompts to finish bare rollouts before launching any continuation. TRACE only requires prompt-level completion. Since the m_i bare rollouts for prompt x_i are co-located on the same worker, their generation does not suffer from intra-prompt long-tail

waiting. Consequently, the expansion for prompt x_i can be enqueued immediately upon their return, bypassing the inter-prompt waiting caused by other prompts distributed across different workers.

4.3 Online Prefix Value Estimation

The allocation objectives require the conditional success probability at both prompt roots and visited prefixes. We learn a shared generalizable predictor that estimates it from prefix histories at both levels.

$$\tilde{V}_\psi : \mathcal{H}_t \mapsto \tilde{V}_\psi(\mathcal{H}_t) \in [0, 1], \quad (17)$$

which provides $\tilde{V}_\psi(x_i)$ for Eq. (12) and $\tilde{V}_\psi(\mathcal{H}_{i,j,t})$ for Eq. (15). It is used only by the allocator, not by the downstream policy optimizer.

Recursive tree-backed targets. After a rollout tree is collected, its terminal leaves provide binary rewards. For any node y , let $\mathcal{C}(y)$ be its children and let n_y be the number of executed terminal descendants below it:

$$n_y = \begin{cases} 1, & y \text{ is a terminal leaf,} \\ \sum_{c \in \mathcal{C}(y)} n_c, & \text{otherwise,} \end{cases}$$

We set $\hat{V}(y) = r_y$ at terminal leaves and compute internal targets by a bottom-up empirical average:

$$\hat{V}(y) = \frac{1}{n_y} \sum_{c \in \mathcal{C}(y)} n_c \hat{V}(c). \quad (18)$$

Algorithm 1: TRACE rollout allocation

Input: Candidate pool sequence $\{\mathcal{B}_s\}_{s=1}^S$
with $|\mathcal{B}_s| = B$; initial policy π_θ ;
predictor $\tilde{V}_\psi$; root budget M ;
expansion factor N .

Output: Finetuned LLM policy π_θ .

for $s = 1$ **to** S **do**

- Initialize collected tree set $\mathcal{G}_s \leftarrow \emptyset$;
- // **Global root allocation**
- Estimate $\tilde{V}_\psi(x_i)$ for all $x_i \in \mathcal{B}_s$;
- Solve Eq. (13) for root counts $\{m_i\}_{i=1}^B$;
- foreach** *prompt* x_i **with** $m_i > 0$ **do**
- Generate m_i bare rollouts from x_i
 and collect terminal rewards;
- Build the prefix-anchor index set $\mathcal{A}_i$;
- Estimate $\tilde{V}_\psi(\mathcal{H}_{i,j,t})$ for each
 $(j, t) \in \mathcal{A}_i$;
- // **Local prefix expansion**
- Solve Eq. (16) for continuation
 counts $\{K_{i,j,t}\}$;
- foreach** (j, t) **with** $K_{i,j,t} > 0$ **do**
- Generate $K_{i,j,t}$ continuations
 from prefix $\mathcal{H}_{i,j,t}$;
- Add the resulting prompt-local
 rollout trees to $\mathcal{G}_s$;
- // **Predictor and policy updates**
- Compute recursive targets $\hat{V}$ on $\mathcal{G}_s$ and
update $\tilde{V}_\psi$ by Eq. (19);
- Update π_θ with any tree-aware
optimizer using $\mathcal{G}_s$;

Thus $\hat{V}(y)$ is the empirical success rate of terminal descendants below y . We train $\tilde{V}_\psi$ by mean-squared regression over a supervision set $\mathcal{S}$ of roots and informative internal nodes:

$$\mathcal{L}_{\text{value}} := \frac{1}{|\mathcal{S}|} \sum_{y \in \mathcal{S}} \left(\tilde{V}_\psi(y) - \hat{V}(y) \right)^2. \quad (19)$$

The supervision set and lightweight predictor implementation details are reported in Appendix E.6.

4.4 Tree-Aware Policy Optimization

TRACE passes the completed rollout trees to a tree-aware policy optimizer as in Algorithm 1. This optimizer can use any prefix-level credit rule that propagates leaf rewards through the tree, including process-reward backups or group-relative backups. In our experiments, we use the concrete instance described in Appendix E.4.

5 Experiments

5.1 Experimental Setup

Datasets and models. We evaluate TRACE on three multi-turn settings. (1) **Mathematical Reasoning**: trains on the DeepScaler corpus (Luo et al., 2025) with a Python interpreter and evaluates on AIME24, AMC23, MATH500 (Lightman et al., 2024), MinervaMath (Minerva) (Lewkowycz et al., 2022), OlympiadBench (Olympiad) (He et al., 2024a), and three out-of-distribution benchmarks: MMLU-Pro (Wang et al., 2024), ARC-c (Clark et al., 2018), and GPQA-diamond (GPQA) (Rein et al., 2023). (2) **Multi-Hop QA**: trains on HotpotQA (Hotpot)’s training split (Yang et al., 2018) and evaluates on its validation split, 2WikiMultiHopQA (2Wiki) (Ho et al., 2020), MuSiQue (Musiq) (Trivedi et al., 2022), and Bamboogle (Bamb) (Press et al., 2023) with a local E5 retrieval server (Wang et al., 2022) built over a Wikipedia dump (Karpukhin et al., 2020). (3) **Function Calling**: trains on 80% of the BFCL v4 (Patil et al., 2025) multi-turn split and evaluates on the 20% test portion across the base (Base), long-context (Long), missing-function (Miss-Func), and missing-parameter (Miss-Param) subsets. Main experiments use Qwen3-8B and Qwen3-14B policy backbones (Yang et al., 2025a). Appendix C reports results using Llama-3.2-3B-Instruct (Grattafiori et al., 2024) as a backbone.

Baselines. We compare TRACE against four baselines. (1) **ReAct** (Yao et al., 2022): evaluates the base model in the same multi-turn agent and tool scaffold. (2) **GRPO** (Shao et al., 2024): samples prompts randomly at the prompt level. (3) **PCL** (Gao et al., 2025): a pure CoT prompt-selection baseline that actively filters prompts by difficulty. (4) **TreePO**: adds tree-structured rollouts and tree-aware updates on top of GRPO, but still selects turn-level continuations randomly, as instantiated in Appendix E.4. All methods, including TRACE, use the same rollout-budget accounting described in Appendix E.5, so gains reflect allocation rather than additional samples.

Metrics. We report final-answer accuracy for **Mathematical Reasoning**, HotpotQA-style exact match and token-level F1 with partial credit at $F1 \geq 0.3$ for **Multi-Hop QA**, and official BFCL success rate for **Function Calling**, averaging each benchmark over a task-specific number of trajectories and reporting the average across benchmarks




Figure 4: **Test accuracy during training.** The six panels cover Mathematical Reasoning (DeepScaler), Multi-Hop QA (HotpotQA), and Function Calling (BFCL v4) with Qwen3-8B (top) and Qwen3-14B (bottom). We compare GRPO, PCL, random TreePO allocation, and TRACE under the same rollout-budget setting. Higher curves indicate stronger final policies under identical sampling budgets.

within each setting. Appendix C reports allocation and predictor diagnostics, while Appendix E gives evaluation protocols, $\text{avg}@k$ multiplicities, dataset splits, model settings, and rollout details.

5.2 Main Results

5.2.1 Improved Accuracy under the Same Budget

Figure 4 shows the training curves across all three settings and two model sizes. Overall, TRACE improves average performance across Mathematical Reasoning, Multi-Hop QA, and Function Calling for both Qwen3-8B and Qwen3-14B, consistently yielding higher evaluation curves than GRPO, PCL, and TreePO under the same rollout budget. On Mathematical Reasoning (DeepScaler), it improves the in-distribution average over GRPO from 70.0 to 71.1 for Qwen3-8B and from 73.5 to 74.9 for Qwen3-14B; similar gains appear on Multi-Hop QA and Function Calling. The TreePO comparison is especially revealing because both methods share tree rollouts and tree-aware updates, yet TRACE allocates branches to roots and prefixes predicted to be more informative rather than sampling them randomly. Full per-benchmark breakdowns are reported in Appendix C.1.

5.2.2 More Informative Rollouts under the Same Budget

Figure 5 reports the effective ratio. At the prompt level, this metric is the within-batch fraction of prompts whose rollout trees contain terminal leaves with both successful and failed outcomes. Under a

matched rollout budget, TRACE achieves a higher training-time average effective ratio than GRPO, PCL, and TreePO across all three settings and both model scales. On Mathematical Reasoning (DeepScaler), it raises the average over GRPO from 26.8% to 60.6% for Qwen3-8B and from 34.7% to 59.7% for Qwen3-14B; similar gains appear on Multi-Hop QA and Function Calling. By directing rollout budget toward contrast-rich roots and prefixes rather than saturated or low-variance anchors, TRACE increases the frequency of non-degenerate terminal groups and amplifies the group-relative policy-update signal per rollout unit. The full metric definition appears in Appendix C.1.

5.2.3 Reliable Difficulty Prediction

Figures 6 and 7 show that the predictor learns difficulty rankings at both prompt and node levels. The prompt-level signal is expected, since root outcomes provide the densest supervision. More importantly, the signal transfers to internal prefixes: even when training is dominated by prompt-level supervision, the same predictor ranks node difficulty reliably enough to guide turn-level branching. This suggests that prefix histories carry a learnable notion of local uncertainty, not just noisy fragments of full trajectories. Appendix C reports the Spearman definition and per-domain trends.

5.3 Additional Analysis

5.3.1 Ablation Studies

We isolate the two allocation stages on Qwen3-8B Multi-Hop QA (HotpotQA) by replacing each




Figure 5: **Effective ratio during training.** The six panels cover Mathematical Reasoning (DeepScaler), Multi-Hop QA (HotpotQA), and Function Calling (BFCL v4) with Qwen3-8B (top) and Qwen3-14B (bottom). The panels compare the fraction of contrastive samples selected by each allocation strategy. Higher values indicate more non-degenerate reward groups per update.

Table 1: **Ablation of active allocation stages on Qwen3-8B Multi-Hop QA (HotpotQA).** Stage 1 allocates prompt-root budget; Stage 2 allocates continuation budget over visited prefixes.

Stage 1	Stage 2	Avg. Acc.	Avg. Eff. Ratio
Uniform	Uniform	49.5	42.8
Active	Uniform	49.8	49.1
Uniform	Active	50.0	47.3
Active	Active	50.6	52.3

active allocator with uniform allocation while keeping the tree budget and optimizer fixed. Table 1 shows that both stages help and their gains stack: root allocation selects prompts likely to become non-degenerate, while prefix allocation spends continuation budget on prefixes where additional descendants can still reveal contrast.

5.3.2 Compatibility with Different Budgets

We also vary the tree budget on Qwen3-8B Multi-Hop QA (HotpotQA). Since the expected rollout budget is $M(1 + N/2)$, (512, 2) uses budget 1024, while (512, 6) and (1024, 2) both use budget 2048. Table 2 shows that TRACE improves TreePO at every budget, lifting the effective ratio by about 9.5–10.2 points and gaining roughly one accuracy point. It also shows that budget shape matters, not only total size: at budget 2048, broader root coverage with (1024, 2) is stronger than deeper continuation sampling from 512 roots. The bottleneck is therefore not just rollout count, but whether the budget reaches states where rewards can form contrast.

Table 2: **Compatibility with different budgets on Qwen3-8B Multi-Hop QA (HotpotQA).** We vary the root budget M and continuation budget N and compare random TreePO allocation with TRACE.

M	N	Method	Avg. Acc.	Avg. Eff. Ratio
512	2	TreePO	48.8	32.2
512	2	TRACE	49.7	42.4
512	6	TreePO	49.4	37.7
512	6	TRACE	50.3	47.8
1024	2	TreePO	49.5	42.8
1024	2	TRACE	50.6	52.3

6 Conclusion

This work formulates mixed-reward contrast allocation as a unified view of sample-efficient RLVR for multi-turn agents. We show that prompt filtering, rollout-count allocation, and prefix branching are all budget decisions over rollout-tree anchors, and present TRACE to allocate budget guided by a shared generalizable predictor of conditional success probability. TRACE directs budget toward anchors whose descendants are likely to expose mixed terminal outcomes, strengthening the group-relative update signal and implicit stepwise preference pairs under outcome-only rewards. Experiments on Mathematical Reasoning, Multi-Hop QA, and Function Calling confirm improved accuracy and rollout informativeness under a matched budget. Moreover, this contrast-allocation view reframes agentic RLVR sampling as a question of where to branch within the rollout tree, not only how many independent rollouts to draw.

Limitations

TRACE is mainly developed for outcome-only RLVR, and tasks without clear terminal verification may therefore require revisiting the mixed-outcome contrast objective. For allocation, higher prediction accuracy can lead to better budget placement, while predictor instantiations remain orthogonal to the TRACE framework. In this work, we instantiate the predictor in a vanilla setting, and stronger alternatives for different agentic tasks remain future work. Empirically, our evaluation is mainly conducted on Mathematical Reasoning, Multi-Hop QA, and Function Calling with Qwen3-8B and Qwen3-14B, leaving more complex and non-stationary agentic scenarios for future exploration.

Acknowledgments

This work was supported by the National Science and Technology Major Project 2025ZD1606303, the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China JYB2025XDXM503, and the National Natural Science Foundation of China (NSFC) under Grants 62495092 and 62306326.

References

- Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. 2017. Hindsight experience replay. *Advances in neural information processing systems*, 30.
- Jose A Arjona-Medina, Michael Gillhofer, Michael Widrich, Thomas Unterthiner, Johannes Brandstetter, and Sepp Hochreiter. 2019. Rudder: Return decomposition for delayed rewards. *Advances in Neural Information Processing Systems*, 32.
- Sanghwan Bae, Jiwoo Hong, Min Young Lee, Hanbyul Kim, JeongYeon Nam, and Donghyun Kwak. 2026. Online difficulty filtering for reasoning oriented reinforcement learning. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 700–719.
- Xiaoyin Chen, Jiarui Lu, Minsu Kim, Dinghuai Zhang, Jian Tang, Alexandre Piché, Nicolas Gontier, Yoshua Bengio, and Ehsan Kamalloo. 2025. Self-evolving curriculum for llm reasoning. *arXiv preprint arXiv:2505.14970*.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu Yu, Qixin Xu, Weize Chen, et al. 2025. Process reinforcement through implicit rewards. *arXiv preprint arXiv:2502.01456*.
- Guanting Dong, Hangyu Mao, Kai Ma, Licheng Bao, Yifei Chen, Zhongyuan Wang, Zhongxia Chen, Jiazhen Du, Huiyang Wang, Fuzheng Zhang, et al. 2025. Agentic reinforced policy optimization. *arXiv preprint arXiv:2507.19849*.
- Mehdi Fatemi, Banafsheh Rafiee, Mingjie Tang, and Kartik Talamadupula. 2025. Concise reasoning via reinforcement learning. *arXiv preprint arXiv:2504.05185*.
- Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. 2025. Group-in-group policy optimization for llm agent training. *Advances in Neural Information Processing Systems*, 38:46375–46408.
- Xidong Feng, Ziyu Wan, Muning Wen, Stephen Marcus McAleer, Ying Wen, Weinan Zhang, and Jun Wang. 2023. Alphazero-like tree-search can guide large language model decoding and training. *arXiv preprint arXiv:2309.17179*.
- Zhaolin Gao, Joongwon Kim, Wen Sun, Thorsten Joachims, Sid Wang, Richard Yuanzhe Pang, and Liang Tan. 2025. Prompt curriculum learning for efficient llm post-training. *arXiv preprint arXiv:2510.01135*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. 2025a. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638.
- Yiran Guo, Lijie Xu, Jie Liu, Dan Ye, and Shuang Qiu. 2025b. Segment policy optimization: Effective segment-level credit assignment in rl for large language models. *Advances in Neural Information Processing Systems*, 38:114399–114431.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, et al. 2024a. Olympiad-bench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3828–3850.

- Mingqian He, Yongliang Shen, Wenqi Zhang, Zeqi Tan, and Weiming Lu. 2024b. Advancing process verification for large language models via tree-based preference learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2086–2099.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. Constructing a multi-hop qa dataset for comprehensive evaluation of reasoning steps. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6609–6625.
- Zhenyu Hou, Ziniu Hu, Yujiang Li, Rui Lu, Jie Tang, and Yuxiao Dong. 2025. Treerl: Llm reinforcement learning with on-policy tree search. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12355–12369.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Jingcheng Hu, Yinmin Zhang, Qi Han, Daxin Jiang, Xiangyu Zhang, and Heung-Yeung Shum. 2025a. Open-reasoner-zero: An open source approach to scaling up reinforcement learning on the base model. *arXiv preprint arXiv:2503.24290*.
- Zengjie Hu, Jiantao Qiu, Tianyi Bai, Haojin Yang, Binghan Yuan, Qi Jing, Conghui He, and Wentao Zhang. 2025b. Vade: Variance-aware dynamic sampling via online sample-level difficulty estimation for multimodal rl. *arXiv preprint arXiv:2511.18902*.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Yuxiang Ji, Ziyu Ma, Yong Wang, Guanhua Chen, Xiangxiang Chu, and Liaoni Wu. 2025. Tree search for llm agent reinforcement learning. *arXiv preprint arXiv:2509.21240*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pages 6769–6781.
- Amirhossein Kazemnejad, Milad Aghajohari, Eva Portelance, Alessandro Sordani, Siva Reddy, Aaron Courville, and Nicolas Le Roux. 2024. Vineppo: Unlocking rl potential for llm reasoning through refined credit assignment. *arXiv preprint arXiv:2410.01679*.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526.
- Jing Yu Koh, Stephen McAleer, Daniel Fried, and Ruslan Salakhutdinov. 2024. Tree search for language model agents. *arXiv preprint arXiv:2407.01476*.
- Xin Lai, Zhuotao Tian, Yukang Chen, Senqiao Yang, Xiangru Peng, and Jiaya Jia. 2024. Step-dpo: Step-wise preference optimization for long-chain reasoning of llms. *arXiv preprint arXiv:2406.18629*.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. 2022. Solving quantitative reasoning problems with language models. *Advances in neural information processing systems*, 35:3843–3857.
- Pengxiang Li, Zhi Gao, Bofei Zhang, Yapeng Mi, Xiaojian Shawn Ma, Chenrui Shi, Tao Yuan, Yuwei Wu, Yunde Jia, Song-Chun Zhu, et al. 2025a. Iterative tool usage exploration for multimodal agents via step-wise preference tuning. *Advances in Neural Information Processing Systems*, 38:59496–59528.
- Ruilin Li, Heming Zou, Xiufeng Yan, Zheming Liang, Jie Yang, Chenliang Li, and Xue Yang. 2026. Enhancing pretrained model-based continual representation learning via guided random projection. *arXiv preprint arXiv:2603.19145*.
- Xuefeng Li, Haoyang Zou, and Pengfei Liu. 2025b. Limr: Less is more for rl scaling. *arXiv preprint arXiv:2502.11886*.
- Zhizhong Li and Derek Hoiem. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947.
- Ziniu Li, Congliang Chen, Tianyun Yang, Tian Ding, Ruoyu Sun, Ge Zhang, Wenhao Huang, and Zhi-Quan Luo. 2025c. Knapsack rl: Unlocking exploration of llms via optimizing budget allocation. *arXiv preprint arXiv:2509.25849*.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2024. Let’s verify step by step. In *International Conference on Learning Representations*, volume 2024, pages 39578–39601.
- Zhihang Lin, Mingbao Lin, Yuan Xie, and Rongrong Ji. 2025. Cppo: Accelerating the training of group relative policy optimization-based reasoning models. *Advances in Neural Information Processing Systems*, 38:61043–61068.
- Mingjie Liu, Shizhe Diao, Ximing Lu, Jian Hu, Xin Dong, Yejin Choi, Jan Kautz, and Yi Dong. 2025.

- Prorl: Prolonged reinforcement learning expands reasoning boundaries in large language models. *arXiv preprint arXiv:2505.24864*.
- Jieyi Long. 2023. Large language model guided tree-of-thought. *arXiv preprint arXiv:2305.08291*.
- Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y Tang, Manan Roongta, Colin Cai, Jeffrey Luo, Tianjun Zhang, Li Erran Li, et al. 2025. Deepscaler: Surpassing o1-preview with a 1.5 b model by scaling rl. *Notion Blog*, 3(5).
- Yixiu Mao, Yun Qu, Qi Wang, Heming Zou, and Xiangyang Ji. 2026a. Dynamics-predictive sampling for active rl finetuning of large reasoning models. *arXiv preprint arXiv:2603.10887*.
- Yixiu Mao, Yun Qu, Qi Wang, Heming Zou, and Xiangyang Ji. 2026b. Rlvr without ineffective samples: Group prioritized off-policy optimization for llm reasoning. *arXiv preprint arXiv:2606.01281*.
- Fanqing Meng, Lingxiao Du, Zongkai Liu, Zhixiang Zhou, Quanfeng Lu, Daocheng Fu, Botian Shi, Wenhai Wang, Junjun He, Kaipeng Zhang, et al. 2025. Mm-eureka: Exploring visual aha moment with rule-based large-scale reinforcement learning. *CoRR*.
- Andrew Y Ng, Daishi Harada, and Stuart Russell. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, volume 99, pages 278–287. Citeseer.
- Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. 2017. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*, pages 2778–2787. PMLR.
- Shishir G Patil, Huanzhi Mao, Fanjia Yan, Charlie Cheng-Jie Ji, Vishnu Suresh, Ion Stoica, and Joseph E Gonzalez. 2025. The berkeley function calling leaderboard (bfcl): From tool use to agentic evaluation of large language models. In *Forty-second International Conference on Machine Learning*.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah A Smith, and Mike Lewis. 2023. Measuring and narrowing the compositionality gap in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5687–5711.
- Yun Qu, Boyuan Wang, Yuhang Jiang, Jianzhun Shao, Yixiu Mao, Heming Zou, Chang Liu, Qi Wang, Meiqin Liu, and Xiangyang Ji. 2026a. Stop wandering, find the keys: LLMs discriminate key states for efficient multi-agent exploration. *SCIENCE CHINA Information Sciences*.
- Yun Qu, Qi Wang, Yixiu Mao, Vincent Tao Hu, Björn Ommer, and Xiangyang Ji. 2025. Can prompt difficulty be online predicted for accelerating rl finetuning of reasoning models? *Preprint*, arXiv:2507.04632.
- Yun Qu, Qi Wang, Yixiu Mao, Heming Zou, Yuhang Jiang, Yingyue Li, Wutong Xu, Lizhou Cai, Weijie Liu, Clive Bai, et al. 2026b. Listwise policy optimization: Group-based rlvr as target-projection on the llm response simplex. *arXiv preprint arXiv:2605.06139*.
- Yun Qu, Qi Wang, Yixiu Mao, Heming Zou, Yuhang Jiang, Weijie Liu, Clive Bai, Kai Yang, Yangkun Chen, Saiyong Yang, et al. 2026c. Small generalizable prompt predictive models can steer efficient rl post-training of large reasoning models. *arXiv preprint arXiv:2602.01970*.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2023. Gpqa: A graduate-level google-proof q&a benchmark. *arXiv preprint arXiv:2311.12022*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Philip Sedgewick. 2014. Spearman’s rank correlation coefficient. *Bmj*, 349.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Qianli Shen, Daoyuan Chen, Yilun Huang, Zhenqing Ling, Yaliang Li, Bolin Ding, and Jingren Zhou. 2025. Bots: A unified framework for bayesian online task selection in llm reinforcement finetuning. *arXiv preprint arXiv:2510.26374*.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2024. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint arXiv:2409.19256*.
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*.
- Sijun Tan, Michael Luo, Colin Cai, Tarun Venkat, Kyle Montgomery, Aaron Hao, Tianhao Wu, Arnab Balyan, Manan Roongta, Chenguang Wang, Li Erran Li, Raluca Ada Popa, and Ion Stoica. 2025. rllm: A framework for post-training language agents. *Notion Blog*.
- Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. 2025. Kimi k1.5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. Musique: Multi-hop questions via single-hop question composition.

- Transactions of the Association for Computational Linguistics*, 10:539–554.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022. Text embeddings by weakly-supervised contrastive pre-training. *arXiv preprint arXiv:2212.03533*.
- Yiping Wang, Qing Yang, Zhiyuan Zeng, Liliang Ren, Liyuan Liu, Baolin Peng, Hao Cheng, Xuehai He, Kuan Wang, Jianfeng Gao, et al. 2025. Reinforcement learning for reasoning in large language models with one training example. *arXiv preprint arXiv:2504.20571*.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, et al. 2024. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37:95266–95290.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Liang Wen, Yunke Cai, Fenrui Xiao, Xin He, Qi An, Zhenyu Duan, Yimin Du, Junchen Liu, Lifu Tang, Xiaowei Lv, et al. 2025. Light-rl: Curriculum sft, dpo and rl for long cot from scratch and beyond. *arXiv preprint arXiv:2503.10460*.
- Yuxi Xie, Anirudh Goyal, Wenyue Zheng, Min-Yen Kan, Timothy P Lillicrap, Kenji Kawaguchi, and Michael Shieh. 2024. Monte carlo tree search boosts reasoning via iterative preference learning. *arXiv preprint arXiv:2405.00451*.
- Huajian Xin, ZZ Ren, Junxiao Song, Zhihong Shao, Wanbiao Zhao, Haocheng Wang, Bo Liu, Liyue Zhang, Xuan Lu, Qiusi Du, et al. 2025. Deepseek-prover-v1.5: Harnessing proof assistant feedback for reinforcement learning and monte-carlo tree search. In *International Conference on Learning Representations*, volume 2025, pages 72274–72303.
- Haobo Xu, Sirui Chen, Ruizhong Qiu, Yuchen Yan, Chen Luo, Monica Cheng, Jingrui He, and Hanghang Tong. 2026. Prune as you generate: Online rollout pruning for faster and better rlvr. *arXiv preprint arXiv:2603.24840*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, et al. 2024. Qwen2.5-math technical report: Toward mathematical expert model via self-improvement. *arXiv preprint arXiv:2409.12122*.
- Zhicheng Yang, Zhijiang Guo, Yinya Huang, Xiaodan Liang, Yiwei Wang, and Jing Tang. 2025b. Treerpo: Tree relative policy optimization. *arXiv preprint arXiv:2506.05183*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 2369–2380.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.
- Zhiyuan Yao, Yi-Kai Zhang, Yuxin Chen, Yueqing Sun, Zishan Xu, Yu Yang, Tianhao Hu, Qi Gu, Hui Su, and Xunliang Cai. 2026. Cobra-rl: Capability-oriented budget allocation for reinforcement learning in llms. *arXiv preprint arXiv:2602.03048*.
- Yixin Ye, Zhen Huang, Yang Xiao, Ethan Chern, Shijie Xia, and Pengfei Liu. 2025. Limo: Less is more for reasoning. *arXiv preprint arXiv:2502.03387*.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. 2025. Dapo: An open-source llm reinforcement learning system at scale. *Advances in Neural Information Processing Systems*, 38:113222–113244.
- Yongcheng Zeng, Zexu Sun, Bokai Ji, Erxue Min, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, Haifeng Zhang, Xu Chen, and Jun Wang. 2025. Cures: From gradient analysis to efficient curriculum learning for reasoning llms. *Preprint*, arXiv:2510.01037.
- Dan Zhang, Sining Zhou, Ziniu Hu, Yisong Yue, Yuxiao Dong, and Jie Tang. 2024. Rest-mcts*: Llm self-training via process reward guided tree search. *Advances in Neural Information Processing Systems*, 37:64735–64772.
- Wenlin Zhang, Xiangyang Li, Kuicai Dong, Yichao Wang, Pengyue Jia, Xiaopeng Li, Yingyi Zhang, Derong Xu, Zhaocheng Du, Huifeng Guo, et al. 2025a. Process vs. outcome reward: Which is better for agentic rag reinforcement learning. *Advances in Neural Information Processing Systems*, 38:58701–58729.
- Xiaojiang Zhang, Jinghui Wang, Zifei Cheng, Wenhao Zhuang, Zheng Lin, Minglei Zhang, Shaojie Wang, Yinghan Cui, Chao Wang, Junyi Peng, et al. 2025b. Srpo: A cross-domain implementation of large-scale reinforcement learning on llm. *arXiv preprint arXiv:2504.14286*.

- Yi-Kai Zhang, Yueqing Sun, Hongyan Hao, Qi Gu, Xunliang Cai, De-Chuan Zhan, and Han-Jia Ye. 2026a. $v_{\{0.5\}}$: Generalist value model as a prior for sparse rl rollouts. *arXiv preprint arXiv:2603.10848*.
- Zhi Zhang, Zhen Han, Costas Mavromatis, Qi Zhu, Yunyi Zhang, Sheng Guan, Dingmin Wang, Xiong Zhou, Shuai Wang, Soji Adeshina, et al. 2026b. Train less, learn more: Adaptive efficient rollout optimization for group-based reinforcement learning. *arXiv preprint arXiv:2602.14338*.
- Yubao Zhao, Weiquan Huang, Sudong Wang, Ruochen Zhao, Chen Chen, Yao Shu, and Chengwei Qin. 2026. Training multi-turn search agent via contrastive dynamic branch sampling. *arXiv preprint arXiv:2602.03719*.
- Haizhong Zheng, Yang Zhou, Brian Bartoldson, Bhavya Kaikhura, Fan Lai, Jiawei Zhao, and Beidi Chen. 2025. Act only when it pays: Efficient reinforcement learning for llm reasoning via selective rollouts. *Advances in Neural Information Processing Systems*, 38:124321–124346.
- Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. 2023. Language agent tree search unifies reasoning acting and planning in language models. *arXiv preprint arXiv:2310.04406*.
- Heming Zou, Yixiu Mao, Yun Qu, Qi Wang, and Xiangyang Ji. 2025a. Utility-diversity aware online batch selection for llm supervised fine-tuning. *arXiv preprint arXiv:2510.16882*.
- Heming Zou, Yunliang Zang, and Xiangyang Ji. 2025b. Structural features of the fly olfactory circuit mitigate the stability-plasticity dilemma in continual learning. *arXiv preprint arXiv:2502.01427*.
- Heming Zou, Yunliang Zang, Wutong Xu, and Xiangyang Ji. 2026. Fly-CL: A fly-inspired framework for enhancing efficient decorrelation and reduced training time in pre-trained model-based continual representation learning. In *The Fourteenth International Conference on Learning Representations*.
- Heming Zou, Yunliang Zang, Wutong Xu, Yao Zhu, and Xiangyang Ji. 2025c. FlyloRA: Boosting task decoupling and parameter efficiency via implicit rank-wise mixture-of-experts. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

Appendix Overview

This appendix provides supplementary material for TRACE, including related work, discussion, extended experimental results, theoretical proof, experimental details, and representative data examples. The appendix is organized as follows:

- **Appendix A (Related Work):** reviews prior work on RLVR and agentic RLVR, sample-efficient prompt selection and allocation, tree search for LLMs, and credit assignment under sparse rewards.
- **Appendix B (Discussion):** discusses allocation as an orthogonal layer, the axes of training-time sampling in Table 3, and root-to-prefix predictor generalization.
- **Appendix C (Extended Experimental Results):** reports the full benchmark tables and diagnostic analyses, including effective-ratio behavior, prefix-predictor rank correlation, transfer to Llama-3.2-3B-Instruct, computational overhead, and TRACE allocation behavior.
- **Appendix D (Theoretical Proof):** gives detailed proofs for prefix predictability, the energy identity connecting prefix uncertainty to reward variance, and the TRACE dominance result.
- **Appendix E (Experimental Details):** describes the tasks, models, training settings, policy optimizers and TreePO instances, rollout construction and budget accounting, and prefix-predictor supervision and implementation.
- **Appendix F (Data Examples):** presents representative real data examples from the three training domains, including the system prompts and tool schemas seen by the agents.

A Related Work

RLVR and Agentic RLVR. Reinforcement learning with verifiable rewards (RLVR) has become a dominant post-training paradigm for LLMs on tasks with automatically checkable outcomes (Jaech et al., 2024; Guo et al., 2025a; Team et al., 2025). Proximal Policy Optimization (PPO) (Schulman et al., 2017) remains a general policy-gradient backbone, while Group Relative Policy Optimization (GRPO) (Shao et al., 2024; Qu et al., 2026b) is especially attractive for RLVR because it removes the expensive value model and

estimates advantages from groups of sampled rollouts. Much of RLVR still lives in a flat regime where a prompt is sampled, complete responses are generated, and a terminal reward drives the update. Agentic environments fracture this abstraction. A rollout is a chain of thought-action-observation turns (Yao et al., 2022), and intermediate prefixes such as search queries, tool calls, and partial solution states can alter the entire downstream trajectory (Qu et al., 2026a). TRACE uses this structure to turn rollout budget from a scalar sampling count into a structural allocation decision across roots and prefixes.

Sample Efficient RLVR. Sample-efficient RLVR improves policy updates by deciding which examples deserve rollout budget. Offline curation filters prompts by difficulty, diversity, length, or estimated solution quality (Ye et al., 2025; Li et al., 2025b; Wen et al., 2025; Hu et al., 2025a; Yang et al., 2024; Fatemi et al., 2025; Wang et al., 2025), reducing training cost but adding preprocessing and remaining static as the policy evolves (Qu et al., 2025; Gao et al., 2025; Mao et al., 2026b). Online prompt selection addresses this limitation by adapting to the current model, either through rollout-based filtering of uninformative prompts (Yu et al., 2025; Liu et al., 2025; Cui et al., 2025; Meng et al., 2025; Bae et al., 2026; Xu et al., 2026; Zhang et al., 2026b) or through predictors that estimate prompt difficulty (Zhang et al., 2025b; Zheng et al., 2025; Gao et al., 2025; Qu et al., 2025; Zeng et al., 2025; Chen et al., 2025; Mao et al., 2026a; Qu et al., 2026c; Zou et al., 2025a; Zhang et al., 2026a; Hu et al., 2025b; Shen et al., 2025). A complementary prompt-level allocation line asks how many samples each selected prompt should receive (Li et al., 2025c; Yao et al., 2026). This line is effective for dynamic rollout accounting, but it typically relies on pre-calculated per-prompt accuracy and has limited ability to generalize to unseen prompts without explicit rollouts. TRACE treats filtering, rollout counts, and prefix branching as boundary cases of one allocation operation, using a shared generalizable predictor to score candidate anchors from prompt roots to internal prefixes and an allocator to assign budget according to predicted conditional success probability.

Tree Search for LLMs. Tree search lets LLMs branch over interaction prefixes and compare with alternative continuations, rather than sampling a

single complete trajectory. Test-time methods such as Tree-of-Thought (Yao et al., 2023; Long, 2023; Snell et al., 2024; Koh et al., 2024; Zhou et al., 2023) and MCTS-style decoding (Xin et al., 2025) spend additional computation at test time to search for a stronger answer for each prompt. Offline variants sample or score reasoning trees and distill the induced comparisons into SFT or preference-learning data (Feng et al., 2023; He et al., 2024b; Xie et al., 2024; Zhang et al., 2025a; Lai et al., 2024; Li et al., 2025a). Online tree-based RL brings branching into training, using tree rollouts for value estimation or finer credit assignment (Zhang et al., 2024; Hou et al., 2025; Ji et al., 2025; Yang et al., 2025b; Kazemnejad et al., 2024; Guo et al., 2025b; Dong et al., 2025; Zhao et al., 2026). TRACE is closest to this online lineage, but repurposes the tree. It branches over agent interaction prefixes rather than tokens or generic reasoning steps, and uses the tree not as an inference solver, but as an allocation substrate for discovering reward contrast under a fixed rollout budget.

Credit Assignment under Sparse Rewards. Sparse outcome rewards are coarse signals for multi-turn agents because one terminal success or failure can be assigned back to many earlier decisions while intermediate decisions can carry sharply different responsibility for the final result (Feng et al., 2025). Classical RL attacks this with exploration bonuses, replay, goal relabeling, reward shaping, or value-based credit estimates (Pathak et al., 2017; Andrychowicz et al., 2017; Ng et al., 1999; Arjona-Medina et al., 2019). In LLM agents, process reward models (PRMs) can generate step- or token-level feedback, but dense process rewards are costly to annotate and brittle across open-ended tool-use settings. Instead, tree rollouts already provide a form of implicit credit assignment by comparing continuations that share a prefix (Ji et al., 2025; Zhao et al., 2026). TRACE strengthens this mechanism without explicit process supervision by allocating more branches to prefixes whose alternative continuations are likely to diverge in outcome, turning sparse final rewards into sharper local preference signals.

B Discussion

B.1 Allocation as an Orthogonal Layer

TRACE separates rollout acquisition from policy optimization. Tree-structured objectives such as TreeRPO (Yang et al., 2025b) and Tree-GRPO (Ji

et al., 2025) define how an existing tree becomes gradients through branch comparisons, value targets, or group-relative advantages. TRACE operates upstream by deciding which anchors deserve additional descendants before the optimizer sees the tree. Thus the same policy objective can receive a contrast-poor or contrast-rich tree under the same rollout budget.

The predictor and TreePO instance used in our experiments are practical instantiations rather than primary technical commitments. TRACE only requires a predictor that ranks anchors by expected reward contrast and a tree-structured optimizer that consumes the resulting rollout tree. Better predictors should improve budget placement, but the framework is not tied to a specific predictor architecture or tree-policy objective.

B.2 Axes of Training-Time Sampling

Table 3 organizes training-time sampling methods by the allocation axes they make controllable. GRPO exposes only a flat sampling axis, PCL turns prompt admission into a learned decision, and TreePO exposes a tree substrate while leaving the tree acquisition policy random. The missing axis is not another optimizer choice, but a structural sampling decision inside the rollout itself. Once a trajectory has interacted with a search engine, a Python interpreter, or an API environment, its visited prefixes become decision states with heterogeneous downstream uncertainty. Some have already collapsed into near-certain success or failure, while others remain poised at points where an alternative query, tool call, or reasoning continuation can alter the final reward. TRACE treats prompt roots and such prefixes as anchors in a single allocation geometry: a lightweight generalizable predictor scores where contrast is still latent, and the allocator decides whether to skip a root, widen it, or branch from an unresolved prefix.

B.3 Root-to-Prefix Predictor Generalization

Using one shared predictor for roots and prefixes is not merely an implementation shortcut. A prefix history is an information-refined prompt: it contains the current plan, the executed actions, and the environment observations that shape what remains possible. Root-level supervision is denser because each root aggregates more terminal descendants and yields more reliable tree-backed targets, anchoring the predictor’s global difficulty scale, while prefix-level supervision is sparser because

Table 3: **Qualitative comparison of training-time allocation methods.** GRPO and PCL operate over flat rollouts, TreePO introduces a tree substrate with random branching, and TRACE uses learned allocation at both prompt roots and visited prefixes.

Method	Prompt Selection	Root Budget	Prefix Expansion	Tree-Structured Update	Learned Allocator
GRPO	Random	Uniform	×	×	×
PCL	Predictive	Fixed	×	×	Root only
TreePO	Random	Uniform	Random	✓	×
TRACE	Predictive	Adaptive	Predictive	✓	Root + prefix

Table 4: **Performance comparison on Mathematical Reasoning.** We train on the DeepScaler corpus and evaluate Qwen3-8B and Qwen3-14B on four in-distribution mathematical reasoning benchmarks and three out-of-distribution benchmarks. TRACE consistently outperforms competitive baselines on average. The best results are indicated in bold.

Model	Method	In-distribution						Out-of-distribution			
		AIME24	AMC23	MATH500	Minerva	Olympiad	Avg	MMLU-Pro	ARC-c	GPQA	Avg
Qwen3-8B	ReAct	52.3	87.0	89.8	37.8	58.6	65.1	68.5	95.3	52.4	72.1
	GRPO	63.6	91.0	92.3	40.2	62.8	70.0	72.7	95.3	55.9	74.6
	PCL	64.0	91.5	92.8	40.7	63.0	70.4	72.9	95.4	56.2	74.8
	TreePO	64.3	91.3	92.6	40.9	63.1	70.4	73.0	95.2	56.4	74.9
	TRACE	63.9	91.2	93.4	41.2	65.8	71.1	73.4	95.7	56.9	75.3
Qwen3-14B	ReAct	61.5	89.4	91.8	40.0	62.3	69.0	75.1	96.2	59.2	76.8
	GRPO	65.6	94.3	94.2	45.2	68.4	73.5	75.9	96.1	59.4	77.1
	PCL	66.2	95.1	95.0	45.1	67.6	73.9	76.1	96.2	59.7	77.3
	TreePO	66.4	95.0	94.8	45.4	67.8	74.0	76.0	96.4	59.8	77.4
	TRACE	66.1	94.7	94.9	47.3	71.5	74.9	76.5	96.5	60.4	77.8

fewer descendants sit below each prefix and the targets are noisier, but closer to the decisions that create local contrast. For training efficiency, predictor updates use all prompt-root examples together with only a small stream of prefix examples; nevertheless, the positive prefix-level correlations in Appendix C indicate that prefix difficulty is still ranked accurately enough for branching and that the predictor is not memorizing prompt identities. It learns a transferable history-conditioned notion of uncertainty, which is precisely the signal needed for targeted branching.

C Extended Experimental Results

C.1 Detailed Benchmark Results

Figure 4 shows training curves with respect to optimization steps, while Tables 4 and 5 summarize the final endpoint performance. Overall, TRACE enables more effective RLVR training under the same rollout-budget setting. Compared with GRPO, TRACE improves the average scores by 0.7–2.8 points across Mathematical Reasoning, Multi-Hop QA, and Function Calling scenarios. Compared

with random TreePO allocation, TRACE also consistently improves the reported averages, indicating that the gains come from learned allocation rather than from introducing a tree substrate alone. These improvements come with negligible predictor overhead, analyzed in Appendix C.5.

C.2 Effective Ratio

For the root or prefix stage $q \in \{\text{root}, \text{pref}\}$, we use the anchor set $\mathcal{A}_q$, descendant set $\mathcal{D}_h(b_h)$, and activation event $\text{Act}(h, b_h)$ from Eq. (10). The effective ratio is the fraction of sampled terminal descendants attached to activated anchors:

$$\text{Eff}_q(b) = \frac{\sum_{h \in \mathcal{A}_q} |\mathcal{D}_h(b_h)| \mathbb{I}\{\text{Act}(h, b_h)\}}{\sum_{h \in \mathcal{A}_q} |\mathcal{D}_h(b_h)|}. \quad (20)$$

Higher values indicate that more sampled trajectories can produce non-degenerate group-relative or local contrast. In our diagnostic figures and ablation tables, we report the prompt-level instantiation because it is shared by flat and tree rollouts. For a training batch $\mathcal{B}_s$, let $\mathcal{D}_i := \mathcal{D}_{x_i}(b_{x_i})$ be the sampled terminal-descendant set below prompt root x_i (Eq. (10)), and let $R_i = \{r(\mathcal{H}) : \mathcal{H} \in \mathcal{D}_i\}$ denote

Table 5: **Performance comparison on Multi-Hop QA and Function Calling.** We train on HotpotQA and BFCL v4, and evaluate Qwen3-8B and Qwen3-14B on four multi-hop reasoning benchmarks and four tool-calling scenarios, respectively. TRACE consistently outperforms competitive baselines on average. The best results are indicated in **bold**.

Model	Method	Multi-Hop QA					Function Calling				
		Hotpot	2Wiki	Musiq	Bamb	Avg	Base	Miss-Func	Miss-Param	Long	Avg
Qwen3-8B	ReAct	41.9	37.6	18.4	44.7	35.7	36.7	21.5	31.1	22.3	28.0
	GRPO	53.2	48.7	27.4	64.7	48.5	58.5	31.0	45.8	38.7	43.5
	PCL	53.6	49.2	27.9	65.0	48.9	59.0	32.1	46.8	39.3	44.3
	TreePO	54.0	49.7	28.4	65.9	49.5	58.9	32.0	46.5	39.4	44.2
	TRACE	55.7	50.9	29.5	66.3	50.6	61.2	34.4	48.8	40.4	46.2
Qwen3-14B	ReAct	46.4	44.2	22.8	55.3	42.2	51.1	23.6	35.6	22.8	33.6
	GRPO	55.3	52.7	30.1	66.8	51.2	70.6	32.6	40.9	36.4	46.1
	PCL	55.6	52.9	30.2	67.2	51.5	70.0	31.8	41.4	40.4	45.9
	TreePO	55.1	55.3	31.0	70.6	53.0	70.8	33.0	43.2	39.4	46.6
	TRACE	57.8	53.1	33.7	71.3	54.0	72.4	34.8	44.6	40.2	48.0




Figure 6: **Prompt-level prediction quality during training.** The six panels cover Mathematical Reasoning (DeepScaler), Multi-Hop QA (HotpotQA), and Function Calling (BFCL v4) with Qwen3-8B (top) and Qwen3-14B (bottom). The panels show Spearman’s rank correlation between predicted prompt difficulty and empirical success rate, with the associated significance shown on the right axis.

their terminal rewards. We compute

$$\text{Eff}_{\text{root}}(\mathcal{B}_s) = \frac{1}{|\mathcal{B}_s|} \sum_{x_i \in \mathcal{B}_s} \mathbb{I}[\text{Var}(R_i) > 0]. \quad (21)$$

Here, $\mathbb{I}[\cdot]$ is the indicator function. This metric measures the fraction of prompt-root groups that contain both successful and failed outcomes. It is an end-to-end batch informativeness measure: root allocation and useful prefix expansion can both increase the chance that a prompt produces mixed outcomes.

Figure 5 shows that TRACE consistently produces more informative prompt-root groups than GRPO and random TreePO allocation. This metric isolates whether the prompts entering an up-

date yield non-degenerate outcome variation, so higher values indicate that learned allocation supplies denser group-relative signal under the same rollout budget.

C.3 Spearman Correlation

To evaluate the quality of difficulty prediction, we report Spearman’s rank correlation coefficient (Sedgwick, 2014). It measures the strength of the monotonic relationship between predicted and empirical difficulty rankings and is invariant to monotonic transformations, making it the right metric for allocation, where relative ordering matters more than calibration. For a root-level or prefix-level evaluation set $\mathcal{S}_q$, where $q \in \{\text{root}, \text{pref}\}$,




Figure 7: **Prefix-level prediction quality during training.** The six panels cover Mathematical Reasoning (DeepScaler), Multi-Hop QA (HotpotQA), and Function Calling (BFCL v4) with Qwen3-8B (top) and Qwen3-14B (bottom). The panels show Spearman’s rank correlation between predicted prefix difficulty and empirical continuation success, with the associated significance shown on the right axis.

Table 6: **Llama-3.2-3B Multi-Hop QA accuracy.** To test transfer beyond Qwen backbones, we report answer-match reward scores on four multi-hop QA benchmarks and their average.

Method	Hotpot	2Wiki	Musiq	Bamb	Avg
ReAct	6.7	7.3	2.9	6.6	5.9
GRPO	38.2	30.5	13.0	32.3	28.5
PCL	39.3	32.5	13.4	30.5	28.9
TreePO	39.7	33.6	13.6	29.9	29.2
TRACE	41.0	35.3	15.2	37.7	32.3




Figure 8: **Llama-3.2-3B Multi-Hop QA average accuracy.** We report the four-benchmark average over HotpotQA, 2WikiMultiHopQA, MuSiQue, and Bamboogle.

let $R_\psi(y)$ and $R_{\hat{V}}(y)$ be the ranks of the predictor score $\tilde{V}_\psi(y)$ from Eq. (17) and the tree-backed empirical target $\hat{V}(y)$ from Eq. (18) within $\mathcal{S}_q$. We compute

$$\rho_q = \frac{\text{cov}_{y \in \mathcal{S}_q}(R_\psi(y), R_{\hat{V}}(y))}{\text{std}_{y \in \mathcal{S}_q}(R_\psi(y)) \text{std}_{y \in \mathcal{S}_q}(R_{\hat{V}}(y))}. \quad (22)$$

We also report the corresponding p -value under the null hypothesis that $\tilde{V}_\psi(y)$ and $\hat{V}(y)$ are independent over $y \in \mathcal{S}_q$, which assesses whether the observed rank agreement is statistically significant.

Figures 6 and 7 show that the predictor learns stable difficulty rankings at both levels. Prompt-level correlation rises first because roots provide the densest supervision. Prefix-level correlation remains positive as well, even though predictor training is dominated by prompt-level examples and uses only limited prefix-level supervision to expose the history format. This transfer is consis-

tent with the root-to-prefix generalization view in Appendix B.3, suggesting that the predictor is not memorizing prompt IDs but naturally extending a root-trained difficulty model to serialized interaction histories, which is exactly the signal needed for targeted branching.

C.4 Applicability to Other Model Families

We further test TRACE on Llama-3.2-3B-Instruct to examine whether the allocation interface in Section 4 transfers beyond the Qwen3 family. Consistent with the Multi-Hop QA trend in Section 5.2, Table 6 shows that TRACE achieves the best average accuracy, improving over GRPO by 3.8 points and over random TreePO allocation by 3.1 points. The gains appear across HotpotQA, 2WikiMultiHopQA, MuSiQue, and Bamboogle, suggesting that the benefit is not tied to a single benchmark instance.




Figure 9: **Llama-3.2-3B allocation and predictor diagnostics on Multi-Hop QA (HotpotQA)**. The panels follow Figures 5–7, showing effective ratio, prompt-level correlation, and prefix-level correlation. Correlation panels show the predictor diagnostics.




Figure 10: **Training time breakdown on Multi-Hop QA (HotpotQA)**. We report average wall-clock time for TRACE with Qwen3-8B and Qwen3-14B. The current implementation uses an unoptimized predictor scoring path, while predictor parameter updates remain a small fraction of total runtime.

Figure 8 shows that the advantage emerges during training rather than from a stronger starting point. Figure 9 provides the corresponding allocation diagnostics, showing that TRACE maintains a higher effective ratio and that the predictor preserves useful prompt- and prefix-level rankings. Together, these results support the central design claim that TRACE is a portable root-prefix allocation layer rather than a policy-backbone-specific heuristic.

C.5 Computational Overhead

As a representative case, we profile Multi-Hop QA training runs on HotpotQA to quantify where wall-clock time is spent. Figure 10 decomposes the runtime into LLM rollout, policy update, predictor scoring, predictor update, and residual overhead. The decomposition uses nonoverlapping critical path timers in which predictor scoring is subtracted from trajectory collection and predictor update is separated from trajectory transformation. This avoids double counting nested diagnostic timers.

The dominant cost is still policy optimization, which takes 66.5% of the Qwen3-8B runtime and 75.4% of the Qwen3-14B runtime. The learnable predictor update is small, accounting for 1.3% and

0.9% of total runtime. Predictor scoring during root and prefix allocation accounts for 1.9% and 1.4% of total runtime. Because TRACE uses the same Qwen3-0.6B predictor for both policy scales, total predictor overhead decreases from 3.2% on Qwen3-8B to 2.3% on Qwen3-14B, while the policy update becomes increasingly dominant. These results indicate that TRACE trades a negligible allocation cost for more informative rollouts, and that further engineering improvements can primarily target batched prefix scoring.

C.6 Allocation Behavior

We inspect Function Calling (BFCL v4) as a representative allocation trace because its longer function-calling episodes make turn-wise prefix allocation smoother than in shorter QA rollouts. Figure 11 shows two complementary views of Stage 2. The relative-position view normalizes each anchor by the length of its bare rollout, while the absolute-depth view reports the share of continuation budget assigned to each turn index. In Figure 11(a), the budget-weighted prefix position remains near the middle of the bare rollout, matching the half-trajectory continuation approximation used in the budget accounting of Appendix E.5. The absolute-




Figure 11: **Stage 2 allocation behavior on Function Calling (BFCL v4).** We average over the common late training window. The left panel bins anchors by relative position within their bare rollout and weights each anchor by the assigned continuation count $K_{i,j,t}$. The right panel reports the share of the total Stage 2 continuation budget assigned to each absolute anchor depth.




Figure 12: **Stage 1 root allocation summary on Function Calling (BFCL v4).** Panel (a) shows the fraction of candidate prompts receiving $m_i = 0$ or $m_i \geq 2$; by construction TRACE uses $m_i \in \{0\} \cup \{2, 3, \dots\}$, so $m_i = 1$ has zero mass. Panel (b) shows active-prompt m_i mean, median, and two-standard-deviation range over the same late training window. Exact per-integer counts within the active set require logging the full integer histogram of m_i .

depth view further shows that nontrivial budget remains across later prefixes up to depth 15, indicating that Stage 2 is not merely adding extra root-like samples.

Figure 12 summarizes Stage 1 root allocation over the candidate prompt pool. Most candidates receive $m_i = 0$ and are skipped, while active prompts receive multiple bare rollouts, typically around five to seven in Function Calling (BFCL v4). The current logs record exact inactive/active ratios and active-set quantiles; per-integer counts within the active set require logging the full integer histogram of m_i .

D Theoretical Proof

D.1 Proof of Proposition 1

Proof. Conditioned on $\mathcal{H}_t$, the m terminal rewards in $\bar{R}_{t,m}$ are independent Bernoulli variables with mean V_t^π . Hence

$$\mathbb{E}[\bar{R}_{t,m} \mid \mathcal{H}_t] = V_t^\pi, \quad (23)$$

so the Bayes-optimal predictor of group correctness is V_t^π . Its optimal risk is

$$\mathcal{E}_{t,m}^* = \frac{1}{m} \mathbb{E}[V_t^\pi (1 - V_t^\pi)]. \quad (24)$$

Because $V_t^\pi = \mathbb{E}[V_{t+1}^\pi \mid \mathcal{H}_t]$, the law of total variance gives

$$\mathbb{E}[V_{t+1}^\pi (1 - V_{t+1}^\pi)] \leq \mathbb{E}[V_t^\pi (1 - V_t^\pi)]. \quad (25)$$

Therefore,

$$\mathcal{E}_{t,m}^* - \mathcal{E}_{t+1,m}^* = \frac{1}{m} \mathbb{E}[\text{Var}(V_{t+1}^\pi | \mathcal{H}_t)] \geq 0. \quad (26)$$

The inequality is strict exactly when the next prefix changes the conditional success probability with positive probability under $\mathcal{H}_t$. Summing the one-step inequalities gives $\mathcal{E}_{t,m}^* \leq \mathcal{E}_{0,m}^*$ for every t . $\square$

D.2 Proof of Proposition 2

Proof. Let $Z_t := V_t^\pi = \mathbb{E}_\pi[r(\mathcal{H}_T) | \mathcal{F}_t]$, where $\mathcal{F}_t := \sigma(\mathcal{H}_t)$. Since (Z_t) is a Doob martingale, we have $\mathbb{E}_\pi[Z_{s+1} | \mathcal{F}_s] = Z_s$ and therefore $\mathbb{E}_\pi[A_s^\pi | \mathcal{F}_s] = 0$ for $A_s^\pi := Z_{s+1} - Z_s$. Using

$$\begin{aligned} Z_{s+1}^2 - Z_s^2 &= (Z_{s+1} - Z_s)^2 \\ &\quad + 2Z_s(Z_{s+1} - Z_s), \end{aligned} \quad (27)$$

and summing from $s = t$ to $T - 1$, we obtain

$$Z_T^2 - Z_t^2 = \sum_{s=t}^{T-1} (A_s^\pi)^2 + 2 \sum_{s=t}^{T-1} Z_s A_s^\pi. \quad (28)$$

Taking conditional expectation with respect to $\mathcal{F}_t$, the second term vanishes by the tower property:

$$\begin{aligned} \mathbb{E}_\pi[Z_s A_s^\pi | \mathcal{F}_t] &= \mathbb{E}_\pi[\mathbb{E}_\pi[Z_s A_s^\pi | \mathcal{F}_s] | \mathcal{F}_t] \\ &= \mathbb{E}_\pi[Z_s \mathbb{E}_\pi[A_s^\pi | \mathcal{F}_s] | \mathcal{F}_t] = 0. \end{aligned} \quad (29)$$

Hence

$$\mathbb{E}_\pi[[Z]_{t:T} | \mathcal{F}_t] = \mathbb{E}_\pi[Z_T^2 | \mathcal{F}_t] - Z_t^2. \quad (30)$$

Because the terminal reward is binary, $Z_T = r(\mathcal{H}_T) \in \{0, 1\}$ and thus $Z_T^2 = Z_T$. Therefore,

$$\mathbb{E}_\pi[Z_T^2 | \mathcal{F}_t] = \mathbb{E}_\pi[Z_T | \mathcal{F}_t] = Z_t, \quad (31)$$

which yields

$$\mathbb{E}_\pi[[Z]_{t:T} | \mathcal{F}_t] = Z_t - Z_t^2 = V_t^\pi(1 - V_t^\pi). \quad (32)$$

This proves Eq. (9). $\square$

D.3 Proof of Proposition 3

Proof. Fix an anchor a and a budget decision b . Let $\text{Act}(a, b)$ denote the event that the sampled descendants below a contain at least one successful and one failed terminal history.

Step 1: outcome contrast is the activation event. For a root group of size m , activation means that the group is non-degenerate:

$$\Pr(\text{Act}(a, m) | a) = 1 - p_a^m - (1 - p_a)^m. \quad (33)$$

This is Eq. (12) after replacing p_a by the predictor. For a visited prefix occurrence with factual reward r , a new continuation matches the factual outcome with probability

$$q_a(r) = rp_a + (1 - r)(1 - p_a). \quad (34)$$

With k new continuations, activation means at least one opposite-outcome sibling appears:

$$\Pr(\text{Act}(a, k) | a, r) = 1 - q_a(r)^k, \quad (35)$$

which is Eq. (15).

Step 2: common optimizers have zero local gradient without activation. For pairwise Bradley–Terry or DPO-style losses, an informative pair consists of two terminal histories $\mathcal{H}^+, \mathcal{H}^-$ below a with different rewards. Let

$$\begin{aligned} \ell_\theta(\mathcal{H} | a) &= \log \frac{\pi_\theta(\mathcal{H} | a)}{\pi_{\text{ref}}(\mathcal{H} | a)}, \\ \Delta_\theta &= \ell_\theta(\mathcal{H}^+ | a) - \ell_\theta(\mathcal{H}^- | a). \end{aligned}$$

Then the Bradley–Terry gradient is

$$\begin{aligned} \nabla_\theta \log \sigma(\beta \Delta_\theta) &= \beta \sigma(-\beta \Delta_\theta) \nabla_\theta \Delta_\theta \\ &= \beta \sigma(-\beta \Delta_\theta) [\nabla_\theta \log \pi_\theta(\mathcal{H}^+ | a) \\ &\quad - \nabla_\theta \log \pi_\theta(\mathcal{H}^- | a)] \end{aligned} \quad (36)$$

If all sampled rewards below a are identical, no such success–failure pair exists and the local pairwise contribution is zero.

For group-relative updates, let $\mathcal{H}_1, \dots, \mathcal{H}_m$ be sampled descendants with rewards $R_s \in \{0, 1\}$ and mean $\bar{R}$. Ignoring clipping and constants, the local group-relative gradient contribution is

$$G_{\text{grp}}(a) = \sum_{s=1}^m (R_s - \bar{R}) \nabla_\theta \log \pi_\theta(\mathcal{H}_s | a). \quad (37)$$

For any vectors $g_1, \dots, g_m$, the algebraic identity

$$\begin{aligned} \sum_{s=1}^m (R_s - \bar{R}) g_s \\ = \frac{1}{m} \sum_{u < v} (R_u - R_v)(g_u - g_v), \end{aligned} \quad (38)$$

holds. Taking $g_s = \nabla_\theta \log \pi_\theta(\mathcal{H}_s | a)$ shows that $G_{\text{grp}}(a)$ is a sum of success–failure score differences; if the group is all-success or all-failure, every coefficient vanishes.

Now let $\tilde{G}_a(b)$ denote the local gradient shape that the chosen optimizer would produce from the available success–failure contrasts below anchor

a , and let $G_a(b)$ be the realized local gradient contribution. The pairwise and group-relative arguments above show that no such contribution is available when no opposite-outcome descendants are observed. Tree-aware optimizers only redistribute such leaf-level contrast through the tree, so they also have zero local contrast contribution below a when $\text{Act}(a, b)$ fails. Hence, for all optimizer classes considered here,

$$G_a(b) = \mathbb{I}\{\text{Act}(a, b)\} \tilde{G}_a(b). \quad (39)$$

Step 3: activation factors the expected squared gradient norm. Taking squared norms and conditional expectation in Eq. (39) gives

$$\begin{aligned} \mathbb{E}[\|G_a(b)\|^2 \mid a] &= \Pr(\text{Act}(a, b) \mid a) \cdot \\ \mathbb{E}[\|\tilde{G}_a(b)\|^2 \mid \text{Act}(a, b), a]. \end{aligned} \quad (40)$$

By definition, $V_a(b)$ denotes the appropriate activation probability: $\Pr(\text{Act}(a, b) \mid a)$ for roots and $\Pr(\text{Act}(a, b) \mid a, r)$ for prefix anchors with factual reward r . If the conditional squared gradient norm lies in $[c_-, c_+]$ across candidate anchors, then the expected squared local gradient norm is bounded between $c_- V_a(b)$ and $c_+ V_a(b)$.

Step 4: dominance of local gradient energy. For a stage $q \in \{\text{root}, \text{pref}\}$, let h index the candidate anchors in that stage and let b_h be the budget assigned to anchor h . Step 3 gives the total expected squared local gradient norm in the form

$$\mathcal{E}_q(b) = \sum_h V_h(b_h) C_h(b_h),$$

where $C_h(b_h)$ is the conditional squared gradient scale after activation. This factor is optimizer- and anchor-dependent, while $V_h(b_h)$ is exactly the activation probability controlled by rollout allocation. Under the normalized-gradient condition in Proposition 3, $C_h(b_h) = c_q$ for every anchor in stage q . TRACE therefore optimizes the gradient-energy component

$$\mathcal{J}_q(b) := \sum_h V_h(b_h), \quad (41)$$

where V_h is Eq. (12) for $q = \text{root}$ and Eq. (15) for $q = \text{pref}$. TRACE solves the resulting combinatorial optimization over the feasible set $\mathcal{F}_q$. If $b_q^* \in \arg \max_{b \in \mathcal{F}_q} \mathcal{J}_q(b)$ and $b_q^u \in \mathcal{F}_q$ is a uniform baseline, then

$$\mathcal{E}_q(b_q^*) = c_q \mathcal{J}_q(b_q^*) \geq c_q \mathcal{J}_q(b_q^u) = \mathcal{E}_q(b_q^u). \quad (42)$$

This proves the first line of Eq. (11) for both $q = \text{root}$ and $q = \text{pref}$.

For the combined tree update, write $G_r := G_{\text{root}}(b_{\text{root}})$ and $G_p := G_{\text{pref}}(b_{\text{pref}})$. Expanding the squared norm gives

$$\begin{aligned} \mathbb{E}\|G_r + G_p\|^2 &= \mathcal{E}_{\text{root}}(b_{\text{root}}) + \mathcal{E}_{\text{pref}}(b_{\text{pref}}) \\ &\quad + 2\mathbb{E}\langle G_r, G_p \rangle. \end{aligned}$$

The two energy terms are no smaller under TRACE by the stage-wise result above. The cross-term condition in Proposition 3 states that TRACE does not reduce $\mathbb{E}\langle G_r, G_p \rangle$ relative to uniform, so the full expected squared gradient norm is also no smaller. Strictness holds for any stage whose uniform allocation is not an optimizer of $\mathcal{J}_q$. This proves Proposition 3. $\square$

E Experimental Details

E.1 Tasks

E.1.1 Mathematical Reasoning

Training dataset. We train on the DeepScaler (Luo et al., 2025) corpus, a collection of 40,315 competition-level mathematics problems. We use the public DeepScaleR preview dataset hosted at <https://huggingface.co/datasets/agentica-org/DeepScaleR-Preview-Dataset>. Each rollout is executed by a tool-augmented agent with access to a Python interpreter for calculation and verification.

Evaluation benchmarks. We evaluate in-distribution mathematical reasoning on AIME24, AMC23, MATH500 (Lightman et al., 2024), MinervaMath (Lewkowycz et al., 2022), and OlympiadBench (He et al., 2024a), using the math evaluation sets hosted at <https://huggingface.co/datasets/math-ai>. Per-dataset scores use the same sampling multiplicities as our evaluation scripts: AIME24 is reported as avg@32, AMC23 as avg@16, and MATH500, MinervaMath, and OlympiadBench as avg@4. To test out-of-distribution generalization, we additionally evaluate on MMLU-Pro (Wang et al., 2024), ARC-Challenge (Clark et al., 2018), and GPQA-diamond (Rein et al., 2023), reported as avg@1, avg@4, and avg@8 respectively. The ‘‘Avg’’ columns in the tables are arithmetic averages across the corresponding benchmark columns.

Reward function. We use a binary verifier reward: a trajectory receives reward 1 if the final

extracted answer matches the ground truth and 0 otherwise. The same reward interface is used for base-model evaluation and RL checkpoints.

E.1.2 Multi-Hop QA

Training dataset. We train on the training split of HotpotQA (Yang et al., 2018), using the HuggingFace version hosted at https://huggingface.co/datasets/hotpotqa/hotpot_qa. The agent follows a ReAct-style search-and-answer loop, issuing search queries to a local E5 retrieval server (Wang et al., 2022) built over a Wikipedia dump (Karpukhin et al., 2020). This keeps retrieval deterministic and avoids external search-engine drift during RL training.

Evaluation benchmarks. We evaluate on the HotpotQA validation split, 2WikiMultiHopQA (Ho et al., 2020), MuSiQue (Trivedi et al., 2022), and Bamboogle (Press et al., 2023). The additional held-out sets are prepared from the FlashRAG benchmark repository at https://huggingface.co/datasets/RUC-NLPIR/FlashRAG_datasets, with canonical sources at <https://github.com/Alab-NII/2wikimultihop>, <https://github.com/StonyBrookNLP/musique>, and <https://huggingface.co/datasets/chiaeyewken/bamboogle>. These datasets stress different forms of multi-hop evidence aggregation, from entity chaining to compositional question decomposition. Multi-Hop QA evaluation uses one sampled trajectory per question, so each per-dataset score is avg@1. The reported “Avg” is the arithmetic average over the four QA benchmarks when per-checkpoint evaluation is available.

Reward function. We use the implemented HotpotQA-style search reward as both the training reward and evaluation score. The evaluator first extracts a concise final answer and normalizes it following SQuAD/HotpotQA conventions. An exact match receives reward 1; otherwise, the evaluator computes token-level F1 against all reference answers, treats predictions with F1 at least 0.3 as partially correct, and assigns reward equal to the F1 score. Predictions below this threshold receive reward 0.

E.1.3 Function Calling

Dataset. We use the multi-turn split of BFCL v4 (Patil et al., 2025), rather than the full BFCL v4 suite, which also contains many single-turn,

live, memory, and web-search categories. We use the four agentic multi-turn subsets from <https://huggingface.co/datasets/gorilla-llm/Berkeley-Function-Calling-Leaderboard>: base, long-context, missing-function, and missing-parameter. We split these multi-turn examples into an 80% training portion and a 20% held-out test portion. The agent observes API specifications, calls functions over multiple turns, and receives environment feedback after each call.

Evaluation protocol. We follow the official BFCL evaluation protocol and report success rates over the base, missing-function, missing-parameter, and long-context subsets. Validation uses four sampled trajectories per example, so subset scores are reported as avg@4. The reported “Avg” is the arithmetic average over the four multi-turn BFCL subsets. A trajectory is correct only when the executed API sequence and final task outcome satisfy the evaluator.

E.2 Models

All main experiments use Qwen3-8B and Qwen3-14B (Yang et al., 2025a) policy backbones, loaded from the official Qwen model family. We additionally run Multi-Hop QA experiments with Llama-3.2-3B-Instruct (Grattafiori et al., 2024) as an extra backbone in Appendix C. The ReAct baseline evaluates the corresponding base or instruction-tuned model in the same multi-turn agent scaffold without RL fine-tuning. TRACE uses a lightweight Qwen3-0.6B (Yang et al., 2025a) model as the prefix predictor. The relevant model repositories are:

- Qwen3-8B: <https://huggingface.co/Qwen/Qwen3-8B>;
- Qwen3-14B: <https://huggingface.co/Qwen/Qwen3-14B>;
- Llama-3.2-3B-Instruct: <https://huggingface.co/meta-llama/Llama-3.2-3B-Instruct>;
- Qwen3-0.6B predictor backbone: <https://huggingface.co/Qwen/Qwen3-0.6B>.

E.3 Training Details

All RL experiments are implemented in the rLLM/verl (Sheng et al., 2024; Tan et al., 2025) stack with asynchronous vLLM rollout workers. Unless otherwise stated, actor optimization uses Adam with learning rate $1e-6$, no KL penalty,

sequence-mean token aggregation, gradient clipping at 1.0, and Clip-Higher ranges $\epsilon_{\text{low}} = 0.2$ and $\epsilon_{\text{high}} = 0.28$. We use FSDP with parameter and optimizer offloading for the actor and reference model. Main training runs use two nodes with eight H20 GPUs per node.

The environment horizon and sequence lengths are task-specific. For Multi-Hop QA, we use maximum prompt length 2048, maximum response length 6144, and at most 5 agent turns. For Mathematical Reasoning (DeepScaler), we use maximum prompt length 2048, maximum response length 14336, and at most 5 agent turns. For Function Calling (BFCL v4), we use maximum prompt length 12800, maximum response length 8192, and at most 40 agent turns to accommodate long multi-turn tool-use traces.

Rollout sampling follows each benchmark’s baseline setting. Multi-Hop QA uses temperature 1.0 and top- $p = 1.0$ during training. Mathematical Reasoning (DeepScaler) uses temperature 0.6 and top- $p = 0.95$, matching the math evaluation style. Function Calling (BFCL v4) uses temperature 0.9 and top- $p = 1.0$.

E.4 Policy Optimizers and TreePO Instances

TRACE is optimizer-agnostic and outputs tree-structured rollout data that can be consumed by different tree-aware RL objectives. In the flat baseline, GRPO normalizes terminal rewards within a response group and removes the value network used in PPO. Let $\mathcal{B}_G$ denote $x \sim \mathcal{D}$ and $\{y_i\}_{i=1}^k \sim \pi_{\theta_{\text{old}}}(\cdot | x)$, and write $\bar{\rho} = \text{clip}(\rho, 1 - \epsilon_{\text{low}}, 1 + \epsilon_{\text{high}})$. GRPO optimizes

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{\mathcal{B}_G} \left[\frac{1}{k} \sum_{i=1}^k \frac{1}{|y_i|} \sum_{t=1}^{|y_i|} \left(\min(\rho_{i,t} \hat{A}_i, \bar{\rho}_{i,t} \hat{A}_i) - \beta D_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}) \right) \right], \quad (43)$$

where $\rho_{i,t} = \frac{\pi_{\theta}(y_{i,t}|x, y_{i,<t})}{\pi_{\theta_{\text{old}}}(y_{i,t}|x, y_{i,<t})}$ and

$$\hat{A}_i = \frac{r_i - \text{mean}(\{r_j\}_{j=1}^k)}{\text{std}(\{r_j\}_{j=1}^k)}. \quad (44)$$

For tree-structured training, we use two concrete optimizers on the same rollout-tree batch $\mathcal{B}_T$, which samples $x \sim \mathcal{D}$ and a rollout tree $\mathcal{G}(x)$ from $\pi_{\theta_{\text{old}}}$ under the initialize-then-expand construction in §2.2. For stable training, Mathematical Reasoning (DeepScaler) and Multi-Hop QA (HotpotQA)

use TreeRPO (Yang et al., 2025b), while Function Calling (BFCL v4) uses Tree-GRPO (Ji et al., 2025). For any tree node y , let $\mathcal{C}(y)$ denote its child set as in Eq. (18) and let $R(y)$ be the scalar reward backed up by the chosen optimizer; at a terminal leaf with history $\mathcal{H}_T$, we set $R(y) = r(\mathcal{H}_T)$ from Eq. (3). In TreeRPO, each internal node y backs up the average child reward

$$R(y) = \frac{1}{|\mathcal{C}(y)|} \sum_{c \in \mathcal{C}(y)} R(c). \quad (45)$$

For every parent p , its child set $\mathcal{C}(p)$ forms a local group, with child advantage

$$\begin{aligned} \hat{A}_c^{\text{TR}} &= \frac{R(c) - \mu_p}{\sigma_p}, \\ \mu_p &= \frac{1}{|\mathcal{C}(p)|} \sum_{c' \in \mathcal{C}(p)} R(c'). \end{aligned} \quad (46)$$

Let $\mathcal{E}_T$ denote the set of parent-child edges (p, c) in $\mathcal{G}(x)$ used for prefix-level sibling comparisons, and let o_c be the token sequence generated on child branch c . TreeRPO applies the same clipped GRPO-style update as Eq. (43):

$$\begin{aligned} \mathcal{J}_{\text{TR}}(\theta) &= \mathbb{E}_{\mathcal{B}_T} \left[\frac{1}{|\mathcal{E}_T|} \sum_{(p,c) \in \mathcal{E}_T} \frac{1}{|o_c|} \sum_{t=1}^{|o_c|} \left(\min(\rho_{c,t} \hat{A}_c^{\text{TR}}, \bar{\rho}_{c,t} \hat{A}_c^{\text{TR}}) \right. \right. \\ &\quad \left. \left. - \beta D_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}) \right) \right], \end{aligned} \quad (47)$$

where $\rho_{c,t} = \frac{\pi_{\theta}(o_{c,t}|p, o_{c,<t})}{\pi_{\theta_{\text{old}}}(o_{c,t}|p, o_{c,<t})}$.

Tree-GRPO normalizes terminal outcome rewards within an intra-tree group (sibling branches sharing the same parent prefix) and an inter-tree group (all complete trajectories for the same prompt). Let $\{y_i\}_{i=1}^G$ denote the G complete trajectories produced by tree search for prompt x , with terminal rewards $R(y_i) = r(\mathcal{H}_T^{(i)})$. Their combined advantage is

$$\begin{aligned} \hat{A}_i^{\text{intra/inter}} &= \frac{R(y_i) - \mu_{\text{intra/inter}}}{\sigma_{\text{intra/inter}}}, \\ \hat{A}_i^{\text{TG}} &= \hat{A}_i^{\text{intra}} + \hat{A}_i^{\text{inter}}. \end{aligned} \quad (48)$$

Tree-GRPO then applies the same clipped token-level update as GRPO:

$$\mathcal{J}_{\text{TG}}(\theta) = \mathbb{E}_{\mathcal{B}_T} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|y_i|} \sum_{t=1}^{|y_i|} \left(\min \left(\rho_{i,t} \hat{A}_i^{\text{TG}}, \bar{\rho}_{i,t} \hat{A}_i^{\text{TG}} \right) - \beta D_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right) \right], \quad (49)$$

where $\rho_{i,t} = \frac{\pi_\theta(y_{i,t}|x, y_{i,<t})}{\pi_{\theta_{\text{old}}}(y_{i,t}|x, y_{i,<t})}$. In all cases, TreePO denotes the same tree-aware optimizer family with random root and prefix allocation, while TRACE changes only the allocation policy that decides where tree branches are collected.

E.5 Rollout Construction and Budget Accounting

All methods draw from the same task-dependent candidate pool, but they consume it differently. For the fixed-size prompt baselines, Multi-Hop QA selects 256 prompts from 512 candidates, while Mathematical Reasoning (DeepScaler) and Function Calling (BFCL v4) select 128 prompts from 256 candidates. GRPO selects this subset uniformly, and PCL selects it with its difficulty predictor. Each selected prompt then receives 8 flat rollouts. TreePO uses the same selected-prompt count and total budget, but collects a random tree by sampling 4 bare rollouts per selected prompt and randomly choosing two branch points in each bare rollout.

TRACE does not first commit to a fixed number of selected prompts. It solves Stage 1 over the full candidate pool and assigns integer root counts m_i , where $m_i = 0$ skips a candidate and $m_i > 0$ determines how many bare rollouts are generated. Figure 12 shows that the resulting active-prompt fraction does not exceed the fixed-size prompt baselines, so TRACE does not gain an advantage from exposure to more distinct prompts. In Multi-Hop QA we use $M = 1024$ root rollout budget units and $N = 2$ continuation budget units per active root; in Mathematical Reasoning (DeepScaler) and Function Calling (BFCL v4) we use $M = 512$ and $N = 2$. After the m_i bare rollouts in Eq. (5) finish, each trajectory induces visited prefix histories $\mathcal{H}_{i,j,t}$ for $1 \leq t < T_{i,j}$. Stage 2 allocates $m_i N$ continuation slots over these prefix occurrences and samples the additional suffixes in Eq. (6). We use one expansion round in the main experiments, so newly generated suffixes are evaluated as terminal branches and are not recursively expanded.

For budget accounting, we count a root rollout

as one trajectory unit and approximate a continuation from an internal prefix as one half trajectory in expectation (Ji et al., 2025; Hou et al., 2025). Thus, ignoring rare fallback cases, TRACE uses approximately $M(1 + N/2)$ trajectory units and is matched to a flat baseline with P selected prompts and G rollouts per prompt by choosing (M, N) such that $M(1 + N/2) \approx PG$. If no valid prefix candidate is available for a prompt, we convert the local continuation budget into additional root rollouts. In implementation, Stage 2 can be triggered per prompt as soon as that prompt’s bare rollouts finish, which avoids a batch-wide straggler barrier while preserving the logical two-stage semantics.

E.6 Prefix Predictor Supervision and Implementation

The prefix predictor $\tilde{V}_\psi$ is implemented as a Qwen3-0.6B critic model, extending the lightweight value-model setup used by PCL (Gao et al., 2025; Qu et al., 2026c) from prompt roots to serialized prefix histories. For each queried anchor, the input is the serialized prompt and the interaction history up to that prompt root or prefix. The predictor outputs a scalar success score in $[0, 1]$, which is used by both Stage 1 root allocation and Stage 2 prefix allocation. For multiple prefixes under the same prompt, prefix-cache reuse amortizes the shared prompt and early-history computation, keeping the additional scoring overhead small.

Following Eqs. (18)–(19), the predictor is updated online after each rollout step by fitting scalar tree-backed success targets with an MSE value-model objective. Each update uses the current step’s prompt-root examples together with a small stream of prefix examples, with prefix samples accounting for 6% of the predictor batch. Predictor learning rates are $3\text{e}-5$ for Multi-Hop QA and Function Calling (BFCL v4), and $1.5\text{e}-5$ for Mathematical Reasoning (DeepScaler). We train for two predictor epochs per rollout on Multi-Hop QA and Function Calling (BFCL v4), and four epochs per rollout on Mathematical Reasoning (DeepScaler). Because tree-backed targets are refreshed as the policy improves, the predictor learns under a non-stationary target distribution and thus faces the stability–plasticity dilemma between adapting to new targets and mitigating catastrophic forgetting. As future work, continual-learning techniques could better manage this dilemma (Kirkpatrick et al., 2017; Li and Hoiem, 2017; Zou et al., 2025b, 2026; Li et al., 2026), and parameter-efficient fine-

tuning (Hu et al., 2022; Zou et al., 2025c) could reduce update overhead without sacrificing ranking quality.

F Data Examples

This section shows representative examples from the actual rLLM data and agent interfaces used by our training scripts. We include the system message seen by the agent, including the tool schemas injected by ToolAgent.

Mathematical Reasoning (DeepScaler) Example

System prompt:

You are a math assistant that can write python to solve math problems.
Tools
You may call one or more functions to assist with the user query.
You are provided with function signatures within <tools></tools> XML tags:
<tools>
[
{
 "type": "function",
 "function": {
 "name": "python",
 "description": "Execute Python code in a sandboxed environment. Returns results and standard output/error.",
 "parameters": {
 "type": "object",
 "properties": {
 "code": {
 "type": "string",
 "description": "Execute Python code in a sandboxed environment. Returns results and standard output/error."
 },
 },
 "timeout": {
 "type": "integer",
 "description": "Maximum execution time in seconds before timing out",
 "default": 12
 }
 },
 "required": [
 "code"
]
 }
}]
</tools>
For each function call, return a json object with function name and arguments within <tool_call></tool_call> XML tags:
<tool_call>
{
 "name": <function-name>,
 "arguments": <args-json-object>
}
</tool_call>

Prompt:

The operation $\otimes$ is defined for all nonzero numbers by $a \otimes b = \frac{a^2}{b}$.
Determine $[(1 \otimes 2) \otimes 3] - [1 \otimes (2 \otimes 3)]$.

Ground-Truth Answer:

$-\frac{2}{3}$

Multi-Hop QA (HotpotQA) Example

System prompt:

You are a helpful AI assistant that can search for information to answer questions accurately.
When answering questions:
1. Use the available search tools to find relevant and reliable information
2. Synthesize information from multiple sources when needed
3. Provide accurate and comprehensive answers based on your search results
4. Always put your final answer in \boxed{} format
For example:
- If the answer is "American", write: \boxed{American}
- If the answer is "yes", write: \boxed{yes}
- If the answer is a year like "1985", write: \boxed{1985}
Remember to search thoroughly and provide your final answer clearly within the \boxed{} format.
Tools
You may call one or more functions to assist with the user query.
You are provided with function signatures within <tools></tools> XML tags:
<tools>
[
{
 "type": "function",
 "function": {
 "name": "local_search",
 "description": "Search for information using a dense retrieval server with Wikipedia corpus",
 "parameters": {
 "type": "object",
 "properties": {
 "query": {
 "type": "string",
 "description": "Search query to retrieve relevant documents"
 },
 },
 "top_k": {
 "type": "integer",
 "description": "Number of results to return (default: 10)",
 "minimum": 1,
 "maximum": 50
 }
 },
 "required": [
 "query"
]
 }
}]
</tools>

```

}
]
</tools>
For each function call, return a json object
with function name and arguments within
<tool_call></tool_call> XML tags:
<tool_call>
{"name": <function-name>, "arguments":
<args-json-object>}
</tool_call>

```

Prompt:

Were Scott Derrickson and Ed Wood of the same nationality?

Ground-Truth Answer:

yes

Function Calling (BFCL v4) Example

System prompt:

```

# Tools
You may call one or more functions to assist
with the user query.
You are provided with function signatures
within <tools></tools> XML tags:
<tools>
Important: Always use only the latest tool
list provided, ignoring any functions
mentioned in previous messages.
{"name": "add_to_watchlist", "description":
"This tool belongs to the trading system,
which allows
users to trade stocks, manage their account,
and view stock information. Tool description:
Add a
stock to the watchlist.", "parameters":
{"type": "dict", "properties": {"stock":
{"type": "string",
"description": "the stock symbol to add to
the watchlist. "}}, "required": ["stock"]},
"response":
{"type": "dict", "properties": {"watchlist":
{"type": "array", "description": "the
watchlist.",
"items": {"type": "string"}}}}}
{"name": "cancel_order", "description": "This
tool belongs to the trading system, which
allows users
to trade stocks, manage their account, and
view stock information. Tool description:
Cancel an
order.", "parameters": {"type": "dict",
"properties": {"order_id": {"type":
"integer",
"description": "ID of the order to cancel.
"}}, "required": ["order_id"]}, "response":
{"type":
"dict", "properties": {"order_id": {"type":
"integer", "description": "ID of the
cancelled order."},
"status": {"type": "string", "description":
"New status of the order after cancellation
attempt."}}}
{"name": "filter_stocks_by_price",
"description": "This tool belongs to the
trading system, which

```

```

allows users to trade stocks, manage their
account, and view stock information. Tool
description:
Filter stocks based on a price range.",
"parameters": {"type": "dict", "properties":
{"stocks":
{"type": "array", "items": {"type": "string"}},
"description": "List of stock symbols to
filter."},
"min_price": {"type": "float", "description":
"Minimum stock price."}, "max_price":
{"type":
"float", "description": "Maximum stock price.
"}}, "required": ["stocks", "min_price",
"max_price"]}, "response": {"type": "dict",
"properties": {"filtered_stocks": {"type":
"array",
"description": "Filtered list of stock
symbols within the price range.", "items":
{"type":
"string"}}}}}
{"name": "fund_account", "description": "This
tool belongs to the trading system, which
allows users
to trade stocks, manage their account, and
view stock information. Tool description:
Fund the
account with the specified amount.",
"parameters": {"type": "dict", "properties":
{"amount":
{"type": "float", "description": "Amount to
fund the account with. "}}, "required":
["amount"]},
"response": {"type": "dict", "properties":
{"status": {"type": "string", "description":
"Status of
the funding operation."}, "new_balance":
{"type": "float", "description": "Updated
account balance
after funding."}}}
{"name": "get_account_info", "description":
"This tool belongs to the trading system,
which allows
users to trade stocks, manage their account,
and view stock information. Tool description:
Get
account information.", "parameters": {"type":
"dict", "properties": {}, "required": []},
"response":
{"type": "dict", "properties": {"account_id":
{"type": "integer", "description": "ID of the
account."}, "balance": {"type": "float",
"description": "Current balance of the
account."},
"binding_card": {"type": "integer",
"description": "Card number associated with
the account."}}}
{"name": "get_available_stocks",
"description": "This tool belongs to the
trading system, which
allows users to trade stocks, manage their
account, and view stock information. Tool
description:
Get a list of stock symbols in the given
sector.", "parameters": {"type": "dict",
"properties":
{"sector": {"type": "string", "description":
"The sector to retrieve stocks from (e.g.,

```

```
'Technology'). "}}, "required": ["sector"]],
"response": {"type": "dict", "properties":
{"stock_list": {"type": "array",
"description": "List of stock symbols in the
specified sector.",
"items": {"type": "string"}}}}}
{"name": "get_current_time", "description":
"This tool belongs to the trading system,
which allows
users to trade stocks, manage their account,
and view stock information. Tool description:
Get the
current time.", "parameters": {"type":
"dict", "properties": {}, "required": []},
"response":
{"type": "dict", "properties":
{"current_time": {"type": "string",
"description": "Current time in
HH:MM AM/PM format."}}}
{"name": "get_order_details", "description":
"This tool belongs to the trading system,
which allows
users to trade stocks, manage their account,
and view stock information. Tool description:
Get the
details of an order.", "parameters": {"type":
"dict", "properties": {"order_id": {"type":
"integer",
"description": "ID of the order. "}},
"required": ["order_id"]}, "response":
{"type": "dict",
"properties": {"id": {"type": "integer",
"description": "ID of the order."},
"order_type": {"type":
"string", "description": "Type of the
order."}, "symbol": {"type": "string",
"description": "Symbol
of the stock in the order."}, "price":
{"type": "float", "description": "Price at
which the order
was placed."}, "amount": {"type": "integer",
"description": "Number of shares in the
order."},
"status": {"type": "string", "description":
"Current status of the order. [Enum]:
[\"Open\",
\"Pending\", \"Completed\",
\"Cancelled\"]}}}}
{"name": "get_order_history", "description":
"This tool belongs to the trading system,
which allows
users to trade stocks, manage their account,
and view stock information. Tool description:
Get the
stock order ID history.", "parameters":
{"type": "dict", "properties": {}, "required":
[]},
"response": {"type": "dict", "properties":
{"order_history": {"type": "array",
"description": "List
of orders ID in the order history.", "items":
{"type": "integer"}}}}}
{"name": "get_stock_info", "description":
"This tool belongs to the trading system,
which allows
users to trade stocks, manage their account,
and view stock information. Tool description:
Get the
```

```
details of a stock.", "parameters": {"type":
"dict", "properties": {"symbol": {"type":
"string",
"description": "Symbol that uniquely
identifies the stock. "}}, "required":
["symbol"]}, "response":
{"type": "dict", "properties": {"price":
{"type": "float", "description": "Current
price of the
stock."}, "percent_change": {"type": "float",
"description": "Percentage change in stock
price."},
"volume": {"type": "float", "description":
"Trading volume of the stock."}, "MA(5)":
{"type":
"float", "description": "5-day Moving Average
of the stock."}, "MA(20)": {"type": "float",
"description": "20-day Moving Average of the
stock."}}}
{"name": "get_symbol_by_name", "description":
"This tool belongs to the trading system,
which allows
users to trade stocks, manage their account,
and view stock information. Tool description:
Get the
symbol of a stock by company name.",
"parameters": {"type": "dict", "properties":
{"name": {"type":
"string", "description": "Name of the
company. "}}, "required": ["name"]},
"response": {"type":
"dict", "properties": {"symbol": {"type":
"string", "description": "Symbol of the stock
or \"Stock
not found\" if not available."}}}
{"name": "get_transaction_history",
"description": "This tool belongs to the
trading system, which
allows users to trade stocks, manage their
account, and view stock information. Tool
description:
Get the transaction history within a
specified date range.", "parameters":
{"type": "dict",
"properties": {"start_date": {"type":
"string", "description": "Start date for the
history (format:
'YYYY-MM-DD').", "default": "None"},
"end_date": {"type": "string", "description":
"End date for the
history (format: 'YYYY-MM-DD').", "default":
"None"}}, "required": [], "response":
{"type":
"dict", "properties": {"transaction_history":
{"type": "array", "description": "List of
transactions
within the specified date range.", "items":
{"type": "dict", "properties": {"type":
{"type":
"string", "description": "Type of
transaction. [Enum]: [\"deposit\",
\"withdrawal\"]"}, "amount":
{"type": "float", "description": "Amount
involved in the transaction."}, "timestamp":
{"type":
"string", "description": "Timestamp of the
transaction, formatted as 'YYYY-MM-DD
HH:MM:SS'."}}}}}}}
```

```

{"name": "get_watchlist", "description":
"This tool belongs to the trading system,
which allows
users to trade stocks, manage their account,
and view stock information. Tool description:
Get the
watchlist.", "parameters": {"type": "dict",
"properties": {}, "required": []},
"response": {"type":
"dict", "properties": {"watchlist": {"type":
"array", "description": "List of stock
symbols in the
watchlist.", "items": {"type": "string"}}}}}
{"name": "notify_price_change",
"description": "This tool belongs to the
trading system, which
allows users to trade stocks, manage their
account, and view stock information. Tool
description:
Notify if there is a significant price change
in the stocks.", "parameters": {"type":
"dict",
"properties": {"stocks": {"type": "array",
"items": {"type": "string"}, "description":
"List of
stock symbols to check."}, "threshold":
{"type": "float", "description": "Percentage
change
threshold to trigger a notification. "}},
"required": ["stocks", "threshold"]},
"response": {"type":
"dict", "properties": {"notification":
{"type": "string", "description":
"Notification message
about the price changes."}}}
{"name": "place_order", "description": "This
tool belongs to the trading system, which
allows users
to trade stocks, manage their account, and
view stock information. Tool description:
Place an
order.", "parameters": {"type": "dict",
"properties": {"order_type": {"type":
"string",
"description": "Type of the order
(Buy/Sell)."}, "symbol": {"type": "string",
"description": "Symbol
of the stock to trade."}, "price": {"type":
"float", "description": "Price at which to
place the
order."}, "amount": {"type": "integer",
"description": "Number of shares to trade.
"}}, "required":
["order_type", "symbol", "price", "amount"]},
"response": {"type": "dict", "properties":
{"order_id": {"type": "integer",
"description": "ID of the newly placed
order."}, "order_type":
{"type": "string", "description": "Type of
the order (Buy/Sell)."}, "status": {"type":
"string",
"description": "Initial status of the
order."}, "price": {"type": "float",
"description": "Price at
which the order was placed."}, "amount":
{"type": "integer", "description": "Number of
shares in the
order."}}}

```

```

{"name": "remove_stock_from_watchlist",
"description": "This tool belongs to the
trading system,
which allows users to trade stocks, manage
their account, and view stock information.
Tool
description: Remove a stock from the
watchlist.", "parameters": {"type": "dict",
"properties":
{"symbol": {"type": "string", "description":
"Symbol of the stock to remove. "}},
"required":
["symbol"]}, "response": {"type": "dict",
"properties": {"status": {"type": "string",
"description":
"Status of the removal operation."}}}
{"name": "trading_get_login_status",
"description": "This tool belongs to the
trading system, which
allows users to trade stocks, manage their
account, and view stock information. Tool
description:
Get the login status.", "parameters":
{"type": "dict", "properties": {}, "required":
[]},
"response": {"type": "dict", "properties":
{"status": {"type": "boolean", "description":
"Login
status."}}}
{"name": "trading_login", "description":
"This tool belongs to the trading system,
which allows
users to trade stocks, manage their account,
and view stock information. Tool description:
Handle
user login.", "parameters": {"type": "dict",
"properties": {"username": {"type": "string",
"description": "Username for
authentication."}, "password": {"type":
"string", "description":
"Password for authentication. "}},
"required": ["username", "password"]},
"response": {"type":
"dict", "properties": {"status": {"type":
"string", "description": "Login status
message."}}}
{"name": "trading_logout", "description":
"This tool belongs to the trading system,
which allows
users to trade stocks, manage their account,
and view stock information. Tool description:
Handle
user logout for trading system.",
"parameters": {"type": "dict", "properties":
{}, "required": []},
"response": {"type": "dict", "properties":
{"status": {"type": "string", "description":
"Logout
status message."}}}
{"name": "withdraw_funds", "description":
"This tool belongs to the trading system,
which allows
users to trade stocks, manage their account,
and view stock information. Tool description:
Withdraw
funds from the account balance.",
"parameters": {"type": "dict", "properties":
{"amount": {"type":

```

```

"float", "description": "Amount to withdraw
from the account. "}}, "required":
["amount"]},
"response": {"type": "dict", "properties":
{"status": {"type": "string", "description":
"Status of
the transaction."}, "new_balance": {"type":
"float", "description": "Updated account
balance after
the transaction."}}}}
</tools>
For each function call, return a json object
with function name and arguments within
<tool_call></tool_call> XML tags:
<tool_call>
{"name": <function-name>, "arguments":
<args-json-object>}
</tool_call>

```

Prompt:

Turn 1: Can you provide the latest trading details for Quasar Ltd.?

Turn 2: Furthermore, I'd like to delve into the technology sector. Would you be able to compile a comprehensive list of stock symbols pertinent to this industry and add all of them to my watchlist

Ground-Truth API sequence:

```

[
[
"get_stock_info(symbol='QUAS')"
],
[
"get_available_stocks('Technology')",
"add_to_watchlist(stock='AAPL')",
"add_to_watchlist(stock='GOOG')",
"add_to_watchlist(stock='MSFT')"
]
]

```