

Failure makes the agent stronger: Enhancing Accuracy through Structured Reflection for Reliable Tool Interactions

Junhao Su^{1*} Yuanliang Wan^{1*} Junwei Yang^{1*}

Hengyu Shi² Tianyang Han² Yurui Qiu^{1†} Junfeng Luo^{1,2}

¹Vision Agent Team, Meituan ²MeiGen AI Team, Meituan
{qiuyurui, luojunfeng}@meituan.com

*Equal Contrib ✉Corresponding Authors †Project Leader

Abstract

Tool-augmented large language models (LLMs) are typically trained via supervised imitation learning or coarse-grained reinforcement learning, approaches that primarily optimize one-shot tool calls. Existing practices of self-reflection largely rely on heuristic prompting or unidirectional reasoning traces: the model is encouraged to “think more,” rather than to treat error diagnosis and correction as a learnable capability. This makes them fragile in multi-turn interaction settings—once a call fails, the model tends to repeat the same mistake instead of recovering. To address this issue, we propose structured reflection, which transforms the “from error to repair” process into a first-class, controllable, and trainable action. The agent produces a concise yet precise reflection process: specifically, the model diagnoses the error based on evidence from the previous step and then proposes a correct and executable follow-up call. During training, we combine DAPO and GSPO’s objective functions and design a more principled reward mechanism tailored to tool calling, optimizing the stepwise strategy Reflect → Call → Final. To evaluate this capability, we introduce Tool-Reflection-Bench, a lightweight benchmark dataset that programmatically verifies structural validity, executability, parameter correctness, and result consistency. Tasks in the benchmark are constructed as miniature trajectories of Erroneous Call → Reflection → Corrected Call and are split into disjoint training and testing sets. Experiments on BFCL v3 and Tool-Reflection-Bench show that our method achieves significant improvements in multi-turn tool-call success rates and error recovery, while also reducing redundant calls. These results demonstrate that making reflection explicit and treating it as an optimization objective can substantially enhance the reliability of tool interaction, providing a reproducible pathway for agents to grow stronger by learning from failure. Our Code

and Dataset is available at: <https://github.com/MeiGen-AI/Tool-Reflection-Bench>

1 Introduction

The integration of external tools with large language models through tool calling represents a significant breakthrough in the development of agents. It transforms large language models from mere text generators into highly practical tools for interacting with humans (WANG et al., 2025; Qu et al., 2024a), significantly enhancing the ability of AI agents to solve complex real-world tasks (Huang et al., 2024; Qin et al., 2023; Qu et al., 2024b). Tool calling bridges the gap between the vast internal knowledge of LLMs and external resources, enabling LLMs to access up-to-date information, perform delicate computations, and more, thereby unlocking their broad potential for applications across multiple domains (Zhong et al., 2023; Theuma and Shareghi, 2024; Hao et al., 2024).

Currently, the training of tool-call capabilities in large language models typically relies on supervised fine-tuning and reinforcement learning (Chen et al., 2025b; Qian et al., 2025), where these methods optimize the ability for single-turn tool calls through carefully designed reward mechanisms. However, these approaches face several challenges in the context of tool calling. First, the issue of rewards in tool calling is particularly prominent—small errors in parameter selection or formatting often render the entire function call invalid, thus limiting the effective learning signal (Lattimer et al., 2024). Second, existing methods generally rely on unidirectional reasoning, which, while sufficient for simpler scenarios, has clear limitations: when LLMs make mistakes during tool calls, they often struggle to locate the root cause of the error (Li et al., 2025). While generating correct function calls is crucial, it is even more important for LLMs to learn how to identify and correct their




Figure 1: In the experiments on BFCL v3 and Tool-Reflection-Bench, our method significantly improves the multi-turn tool-calling accuracy of several open-source LLMs on BFCL v3. At the same time, it substantially enhances the error-repair rate for tool calls on the Tool-Reflection-Bench test set, achieving performance that even surpasses that of closed-source LLMs with comparable parameter sizes.

own mistakes (Ye et al., 2024).

To address the above-mentioned issues, we propose an innovative reflection process aimed at error localization and correction through explicit reflection steps, which differs from existing forward reasoning methods. Specifically, we design a process in which the LLM intentionally makes mistakes during tool calls, carefully crafts reflection content based on the errors, and then generates the correct call. Through this approach, we transform the self-correction ability of large models from a heuristic process (Yang et al., 2024) into a clear, trainable capability. Our training approach is primarily reinforcement learning-based. During the reinforcement learning process, we specifically design a customized reward mechanism tailored for tool-calling scenarios, with a particular emphasis on multi-turn interactions. Concretely, the reward design encompasses multiple dimensions, including format reward, tool-name reward, parameter reward, and semantic reward of reflection, which together provide the model with multi-dimensional feedback and effectively guide its learning, and we further combine DAPO’s decoupled clipping range and dynamic sampling—expanding exploration while skipping near-zero-advantage roll-outs—with GSPO’s sequence-level importance sampling and same-granularity clipping, which avoids token/sequence mismatch and stabilizes optimization. With this training methodology, our approach equips LLMs with genuine self-reflection and error-correction capabilities. On the BFCL v3 benchmark, our method yields significant im-

provements in LLM accuracy for multi-turn tool calling, thereby demonstrating its effectiveness in real-world applications.

We construct a Tool-Reflection-Bench based on the BUTTON dataset (Chen et al., 2024) style. First, we collected tool-call failure cases from real-world scenarios and various benchmarks, analyzing and summarizing several common failure patterns. Next, We selected several existing tool-call datasets (Qin et al., 2023; Liu et al., 2024b) and randomly combined them according to the call style of the BUTTON dataset and introduced these failure patterns into the data, disrupting the originally correct call processes to generate failure cases. Finally, we meticulously designed a reflection process to repair these failures, resulting in successful tool calls. The training set includes the complete process described above to train LLMs to achieve true self-correction capabilities, while the test set only contains the first two steps, used to evaluate the self-correction abilities of the LLMs. By constructing the Tool-Reflection-Bench in this manner, combined with our custom reward mechanism for tool calling, we have made breakthroughs in LLMs’ self-correction abilities during training. Particularly in multi-turn tool-calling scenarios, we observed significant improvements in accuracy. Through the reasoning process from failure to correction, LLMs can more effectively identify and learn from potential mistakes, thus enhancing the model’s stability and robustness in interactions. This makes the agent’s behavior more robust and powerful.

In summary, our contributions are as follows:

- We introduce an explicit, trainable reflection process that diagnoses the cause of a failed tool call using prior evidence and proposes a corrected, executable call. This transforms the "from failure to repair" process from a heuristic method into a learnable action strategy, enabling LLMs to genuinely possess self-reflection and error-correction capabilities, thereby enhancing the agent’s multi-turn interactions with users.
- We design a more effective reward mechanism for tool call, tailored for RL training, using a GRPO-style objective function. This approach employs multi-dimensional rewards—format executability, tool name accuracy, parameter correctness, and semantic consistency—to mitigate sparse rewards and propagate signals across multi-turn trajectories.
- We propose Tool-Reflection-Bench, which collects failure patterns from real interaction scenarios and benchmark datasets, injects perturbations into correct calls, and attaches a reflection process to repair the calls. This allows for training LLMs in their Self-Correction ability in tool-calling scenarios.
- Our method significantly improves the accuracy of multi-turn tool calls and the ability to recover from tool call errors, while maintaining competitive single-turn tool call performance. We validate this by experiments on BFCL v3 (Patil et al.) and Tool-Reflection-Bench.

2 Method

2.1 Tool-Reflection-Bench

The construction of Tool-Reflection-Bench consists of the following steps: perturbation-based disruptions, positive samples transformations, and the reflection repair process. The original positive samples are derived from BUTTON (Chen et al., 2024) transformations and self-constructed based on few-shot prompts. The entire benchmark is divided into a training set and a test set, with approximately 5,000 samples in the training set, in addition to the reflection-augmented data constructed as described above, the training set also contains a very

small portion of original data drawn from BUTTON (Chen et al., 2024) and XLAM (Zhang et al., 2024). And around 1,000 samples in the test set, the test set is exclusively composed of perturbation-derived items and does not include raw, unperturbed positives from BUTTON or XLAM.

2.1.1 Perturbation-based Disruptions

Let the initial correct message sequence be

$$D^+ = \left(m_0^{\text{sys}}, m_1^{\text{usr}}, m_2^{\text{ast}}, m_3^{\text{tool}}, m_4^{\text{ast}}, m_5^{\text{tool}}, \dots, m_{2k}^{\text{ast}}, m_{2k+1}^{\text{tool}}, \dots, m_n^{\text{final}} \right). \quad (1)$$

where m_0^{sys} is the system prompt, m_1^{usr} the user query, m_{2i}^{ast} the assistant’s i -th tool call in structured form (e.g., `<call>[{ . . . }, { . . . }, . . .]</call>`), m_{2i+1}^{tool} the tool return (JSON), and m_n^{final} the final answer.

We define a set of disruption operators

$$\mathcal{P} = \{P_1, P_2, P_3, P_4\}, \quad (2)$$

each operating on an assistant call m_{2k}^{ast} and instantiating a common failure mode:

1. **P_1 call-order swap**: replace the current tool call with the next-round tool call dialogue and force an error.
2. **P_2 redundant call**: repeat the same tool at the step (unchanged/irrelevant arguments) and force an error.
3. **P_3 missing call**: replace the intended tool by another tool and force an error.
4. **P_4 argument error**: randomly corrupt the arguments of a call (missing/typed/alias/boundary) and force an error.

These operators specify how a correct tool call can be broken.

2.1.2 Positive Samples Transformations

Given a clean trajectory D^+ and a chosen operator $P_j \in \mathcal{P}$ acting on step $2k$, we produce the negative (erroneous) context; no repair is performed in this step. We construct the erroneous call

$$\tilde{m}_{2k}^{\text{ast}} = \text{ApplyPerturbation}(m_{2k}^{\text{ast}}, P_j), \quad (3)$$

and simulate the tool’s error feedback with a LLM \mathcal{L} :

$$\tilde{m}_{2k+1}^{\text{tool}} = \mathcal{L}(\tilde{m}_{2k}^{\text{ast}}, \mathcal{L}). \quad (4)$$

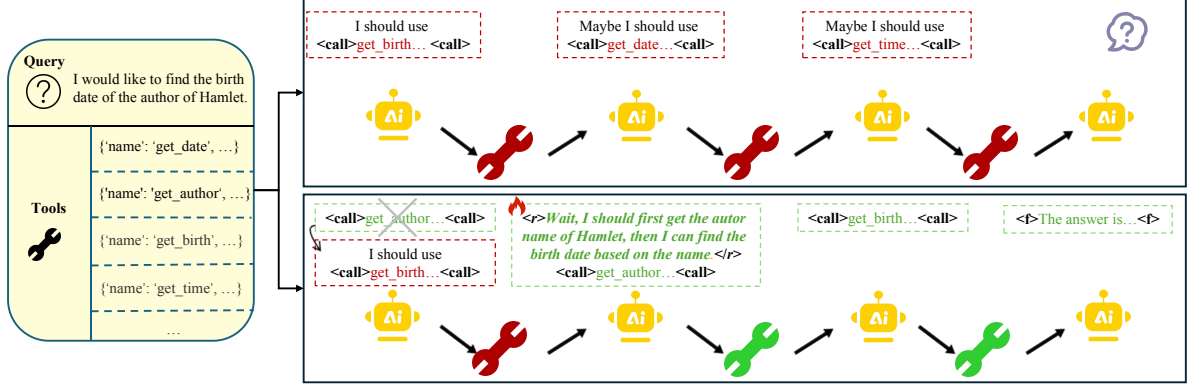


Figure 2: We illustrate the effectiveness of our method with an example. As shown in the figure, the left side presents the tool panel, while the upper-right part depicts industry-standard self-correction approaches, where models attempt to fix errors through heuristic trial-and-error reasoning or by relying on external feedback. In contrast, the lower-right part shows our approach: we introduce an explicit forced reflection process $\langle r \rangle$, enabling the model to truly master the ability to repair errors based on its own failures.

This yields the **negative** trajectory prefix

$$\begin{aligned} D^- &= \text{Perturb}(D^+, P_j) \\ &= (m_0^{\text{sys}}, m_1^{\text{usr}}, \dots, \tilde{m}_{2k}^{\text{ast}}, \\ &\quad \tilde{m}_{2k+1}^{\text{tool}}), \end{aligned} \quad (5)$$

which will later serve as evidence of failure. At this stage, the item consists only of the broken call and its error signal.

2.1.3 Reflection Repair Process

Given a clean trajectory D^+ and its perturbed prefix D^- , we present the LLM with a paired view of the step- $2k$ evidence:

$$\text{clean: } (m_{2k}^{\text{ast}}, m_{2k+1}^{\text{tool}}) \quad \text{vs.} \quad \text{broken: } (\tilde{m}_{2k}^{\text{ast}}, \tilde{m}_{2k+1}^{\text{tool}}). \quad (6)$$

The model outputs a response:

$$\langle \text{reflect} \rangle \text{ref} \langle / \text{reflect} \rangle, \quad (7)$$

where ref briefly diagnoses the discrepancy, and c proposes the fixed tool call. We then apply human supervision to obtain (ref^*, c^*) , where c^* is set to the original correct call from the clean trajectory:

$$(\text{ref}, c) \xrightarrow[\text{human supervision}]{\text{post-editing}} (\text{ref}^*, c^*). \quad (8)$$

$$\mathcal{L}_\Sigma(c^*) = \text{Success Call}. \quad (9)$$

Human supervision cost. Our supervision is performed at the *trajectory* level. Specifically, we construct and retain approximately 5k multi-turn trajectories, and for each trajectory we only require

post-editing the reflection text ref (while c^* is directly copied from the clean call). Two annotators completed the post-editing process over 18 days, making the overall supervision cost controllable.

The finalized item is packaged as

$$x = (D^-, \text{ref}^*, c^*, D_{>2k+1}^+), \quad (10)$$


where $D_{>2k+1}^+$ is the untouched suffix of D^+ (including m_n^{final}). We retain x only if: (i) tags/JSON are well-formed; (ii) c^* is executable; (iii) ref^* correctly cites the clean–broken contrast.

2.2 Reward Design

Preliminary. Given a model completion C and a ground truth G , we decompose both into three (possibly empty) parts:

$$\begin{aligned} C &\mapsto (c_{\text{ref}}, C_{\text{calls}} = \{c_i\}_{i=1}^m, c_{\text{final}}), \\ G &\mapsto (g_{\text{ref}}, G_{\text{calls}} = \{g_j\}_{j=1}^n, g_{\text{final}}). \end{aligned} \quad (11)$$

Here c_{ref} is the diagnosis text wrapped in $\langle \text{reflect} \rangle \langle / \text{reflect} \rangle$, C_{calls} is the *multiset* of tool calls wrapped in one or more $\langle \text{call} \rangle \langle / \text{call} \rangle$ blocks, and c_{final} is the message wrapped in $\langle \text{final} \rangle \langle / \text{final} \rangle$. We treat C_{calls} as a multiset because the same tool may be invoked multiple times within one completion, and we do not assume any canonical order when matching calls to the reference. In our data format, each training example follows *wrong call* \rightarrow *reflection* \rightarrow *corrected call*; the ground truth always contains a reflection and at least one valid tool call (i.e., $n \geq 1$), while the final message may be empty. The ground truth can also be decomposed into the same three parts.



(a) The reward curve of llama-3.1-8b-Instruct during RL training (b) The reward curve of qwen2.5-7b-Instruct during RL training

Figure 3: The reward curves of llama-3.1-8B and Qwen2.5-7B during training, showing an overall upward trend.

Component scores. We compute three component scores:

$$\begin{aligned} s_{\text{ref}} &= \text{Sim}(c_{\text{ref}}, g_{\text{ref}}), \\ s_{\text{call}} &= \mathbb{I}[\text{EqualCalls}(C_{\text{calls}}, G_{\text{calls}})], \\ s_{\text{final}} &= \text{Sim}(c_{\text{final}}, g_{\text{final}}), \end{aligned} \quad (12)$$

where $\text{Sim} \in [0, 1]$ is a semantic similarity function, and $\mathbb{I}[\cdot]$ is the indicator:

$$\mathbb{I}[P] = \begin{cases} 1, & \text{if } P \text{ is true,} \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

We say $\text{EqualCalls}(C_{\text{calls}}, G_{\text{calls}})$ holds iff the two multisets of calls can be put in a one-to-one correspondence such that for every matched pair the **tool name** is identical and the **argument** is identical. We perform multiset matching: each produced call is matched to at most one reference call, and order is ignored.

Normalization with presence masks. Our goal is to keep the aggregated score in $[0, 1]$ even when an instance specifies only a subset of targets (e.g., only `<call>` without `<reflect>` or `<final>`). To this end we use normalization to renormalize over the parts that actually appear in the ground truth, so the maximum remains 1 regardless of how many parts are present.

We define

$$\begin{aligned} I_r &= \mathbb{I}[g_{\text{ref}} \neq \emptyset], \\ I_c &= \mathbb{I}[|G_{\text{calls}}| > 0], \\ I_f &= \mathbb{I}[g_{\text{final}} \neq \emptyset]. \end{aligned} \quad (14)$$

Let $(w_r, w_c, w_f) \geq 0$ be *normalized* base weights (e.g., $w_r + w_c + w_f = 1$). We renormalize over the active parts via

$$W_{\text{act}} = w_r I_r + w_c I_c + w_f I_f. \quad (15)$$

The aggregated structure/semantics score is then

$$S = \frac{w_r I_r s_{\text{ref}} + w_c I_c s_{\text{call}} + w_f I_f s_{\text{final}}}{W_{\text{act}}}. \quad (16)$$

This normalization yields a **consistent** scoring standard across fully and partially supervised instances, avoiding artificial deflation of scores when some targets are absent.

Format/penalty factor. We designed structural penalties tailored for tool-call data formats. Specifically, P_{miss} accounts for cases where the tool is not invoked at all, while P_{extra} and P_{count} penalize redundant calls and mismatches in the total number of calls, respectively. Let

$$n = |G_{\text{calls}}|, \quad m = |C_{\text{calls}}|, \quad (17)$$

Here n and m denote the number of tools invoked in the ground truth and completion calls. Define the three components:

$$\begin{aligned} P_{\text{miss}} &= w_{\text{ref}} \mathbb{I}[g_{\text{ref}} \neq \emptyset \wedge c_{\text{ref}} = \emptyset] \\ &\quad + w_{\text{final}} \mathbb{I}[g_{\text{final}} \neq \emptyset \wedge c_{\text{final}} = \emptyset] \\ &\quad + w_{\text{calls}} \mathbb{I}[n > 0 \wedge m = 0]. \end{aligned} \quad (18)$$

$$\begin{aligned} P_{\text{extra}} &= w_{\text{ref}} \mathbb{I}[c_{\text{ref}} \neq \emptyset \wedge g_{\text{ref}} = \emptyset] \\ &\quad + w_{\text{final}} \mathbb{I}[c_{\text{final}} \neq \emptyset \wedge g_{\text{final}} = \emptyset] \\ &\quad + w_{\text{calls}} \mathbb{I}[m > 0 \wedge n = 0]. \end{aligned} \quad (19)$$

$$\begin{aligned} P_{\text{count}} &= w_{\text{calls}} \mathbb{I}[n > 0 \wedge m > 0 \wedge n \neq m] \\ &\quad \cdot \frac{|n - m|}{\max\{n, m\}}. \end{aligned} \quad (20)$$

Let EqualCalls be the schema-strict equality on bags of calls. And let $E =$

Table 1: Comparison across dimensions (Base, Miss_Func, Miss_Param, Long_Context, Multi-turn Overall) on BFCL v3, all results in the table are reported as the average over three runs with different random seeds.

Models	Method	Base	Miss_Func	Miss_Param	Long_Context	Multi-turn Overall
Llama-3.1-8B-Instruct-FC	Origin	5.0	6.5	4.5	4.5	5.12
	Ours	9.5 (↑95%)	7.0 (↑8%)	5.0 (↑11%)	7.0 (↑56%)	7.12 (↑39%)
Qwen2.5-7B-Instruct-FC	Origin	16.5	11.0	9.0	7.5	11.00
	Ours	22.0 (↑33%)	13.0 (↑18%)	13.5 (↑50%)	11.0 (↑47%)	14.88 (↑35%)
Qwen3-4B-Instruct	Origin	18.0	19.0	13.5	14.5	16.25
	Ours	25.0 (↑39%)	19.5 (↑3%)	17.0 (↑26%)	21.5 (↑48%)	20.75 (↑28%)

EqualCalls($C_{\text{calls}}, G_{\text{calls}}$) We use a reduction factor

$$r = \begin{cases} r_{\text{reduce}}, & \text{if } E, \\ 1, & \text{else,} \end{cases}, \quad r_{\text{reduce}} \in (0, 1]. \quad (21)$$

Aggregate the penalty as

$$P_{\text{total}} = P_{\text{miss}} + \beta_{\text{extra}} P_{\text{extra}} + \gamma_{\text{count}} P_{\text{count}}, \quad (22)$$

and define the instance-wise format factor. Let $P_{\text{fmt}} := P_{\text{miss}} + P_{\text{extra}} + P_{\text{count}}$, $[x]_0^1 := \min\{1, \max\{0, x\}\}$, and $F := \text{FormatFactor}(C, G)$. Then

$$F = \begin{cases} 1, & \text{if } P_{\text{fmt}} = 0, \\ [1 - \lambda_m P_{\text{total}} r]_0^1, & \text{otherwise.} \end{cases} \quad (23)$$

Here $\beta_{\text{extra}}, \gamma_{\text{count}}, \lambda_m \geq 0$ control the strength of the extra-part penalty, the count-mismatch penalty, and the overall scaling, respectively; ($w_{\text{ref}}, w_{\text{calls}}, w_{\text{final}} \geq 0$ are part weights).

Core reward and backoff. The core reward is

$$R_{\text{core}} = S \cdot F. \quad (24)$$

Early in training, S contains a binary component ($s_{\text{call}} \in \{0, 1\}$) and F applies hard penalties; small formatting or argument errors can drive R_{core} close to zero. This yields sparse or unstable gradients and large variance across samples. To stabilize learning and provide a dense shaping signal when the exact-match objective is not yet achieved, we introduce a similarity backoff. Let $[x]_0^1 := \max(0, \min(1, x))$ and $S_b := \text{Sim}(\text{concat}(C), \text{concat}(G))$:

$$R_{\text{total}} = \begin{cases} [R_{\text{core}}]_0^1, & R_{\text{core}} \geq \varepsilon, \\ [w_b S_b]_0^1, & \text{otherwise.} \end{cases} \quad (25)$$

Here $w_b \in (0, 1]$ and $\text{concat}(\cdot)$ linearizes the messages.

2.3 RL for Tool-Reflection-Bench

We adopt a reinforcement-learning objective for tool calling that combines two complementary ideas: **(i) DAPO-style decoupled clipping** (Yu et al., 2025): we use a decoupled clipping range with different lower/upper bounds ($\varepsilon_{\text{low}}, \varepsilon_{\text{high}}$) and a clip-higher policy (a looser upper bound when $r > 1$ for positive advantages), and we skip uninformative prompt groups whose rollouts carry negligible learning signal; **(ii) GSPO-style sequence-level importance sampling** (Zheng et al., 2025): we compute the importance ratio at the sequence level and apply clipping at the same granularity as the sequence-level reward, which avoids the mismatch between token-wise importance sampling and sequence-level rewards and stabilizes optimization.

Objective. Let (q, a) denote the dialog context and the ground-truth targets, and let $\{o_i\}_{i=1}^G$ be G candidates sampled from the behavior policy $\pi_{\theta_{\text{old}}}(\cdot | q)$. Each completion o_i is scored by the reward in Sec. §2.2, yielding $R_i \in [0, 1]$. We maximize a *sequence-level, asymmetrically clipped* objective and minimize its negative as the loss:

$$\mathcal{J}_{\text{RL}}(\theta) = \mathbb{E} \left[\frac{1}{G} \sum_{i=1}^G \min(r_i(\theta), \bar{r}_i(\theta)) \hat{A}_i \right], \quad (26)$$

$$\bar{r}_i(\theta) := \text{clip}(r_i(\theta), 1 - \varepsilon_{\text{low}}, 1 + \varepsilon_{\text{high}}).$$

Here the expectation is over $(q, a) \sim \mathcal{D}$ and $\{o_i\} \sim \pi_{\theta_{\text{old}}}(\cdot | q)$. We define $\text{clip}(x, a, b) = \min\{b, \max\{a, x\}\}$ and typically take $\varepsilon_{\text{high}} > \varepsilon_{\text{low}}$ (“clip-higher”).

Prompt-group dynamic filtering. DAPO skips prompt groups whose candidates provide almost no learning signal (e.g., all-correct or all-wrong). Concretely, let $\mu_R := \text{mean}(\{R_j\}_{j=1}^G)$ and $\sigma_R := \text{std}(\{R_j\}_{j=1}^G)$. We define batch-normalized advan-

tages and a group-level acceptance set:

$$\hat{A}_i = \frac{R_i - \mu_R}{\sigma_R},$$

$$\mathcal{S}(q, a) = \left\{ i \in \{1, \dots, G\} : |\hat{A}_i| > \tau_{\text{adv}} \right\}. \quad (27)$$

We further require sufficient reward dispersion within the group:

$$\begin{aligned} \text{Var}(\{R_i\}_{i=1}^G) &> \tau_{\text{var}}, \\ 0 &< |\mathcal{S}(q, a)| < G. \end{aligned} \quad (28)$$

If (28) fails, we drop the zero-information rollouts and (optionally) draw up to K additional candidates from $\pi_{\theta_{\text{old}}}$, then re-apply the filter. Only indices in $\mathcal{S}(q, a)$ contribute to the expectation in (26).

Sequence-level importance ratio. For a completion $o_i = (o_{i,1}, \dots, o_{i,|o_i|})$, we use the geometric-mean, length-normalized importance ratio:

$$r_i(\theta) = \left(\prod_{t=1}^{|o_i|} \frac{\pi_{\theta}(o_{i,t} \mid q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t} \mid q, o_{i,<t})} \right)^{1/|o_i|}, \quad (29)$$

and perform clipping at the same sequence granularity as the reward (see (26)), thereby avoiding token/sequence granularity mismatch.

3 Experiments

3.1 Experiment Settings

In this part, we will detail the experimental setup, including datasets, hyperparameters, base models, and evaluation metrics.

Datasets. We conduct training on our self-constructed Tool-Reflection-Bench. After human supervision and post-editing, we retained approximately 5k samples in JSONL format to ensure compatibility with RL training under the Swift (Zhao et al., 2025b) framework.

Implementation Details. We train models for 1 epoch (a total of 1,000 steps) on 5,000 training samples, using the reward function defined in Sec.2.2. For each training instance, 4 completions were sampled to form a group. The training parameters were set as follows: temperature = 0.85, repetition penalty = 1.1, epsilon = 0.2, epsilon-high = 0.28, with a dynamic sampling strategy adopted.

Table 2: Experimental Results of Open-Source and Closed-Source Models on the Tool-Reflection-Bench Test Set.

Models	Repair@1 (%)	Repair@3 (%)	Repair@5 (%)
Close-Sourced Models			
LongCat-Lite-8K-Chat	2.3	3.4	4.9
GPT-4o-mini	6.1	8.7	9.0
GPT-4.1-mini	3.1	4.3	5.1
Open-Sourced Models			
Llama-3.1-8B-Instruct	0.7	5.1	6.8
Qwen2.5-7B-Instruct	2.4	6.1	8.0
Qwen3-4B-Instruct	9.6	10.6	10.6
Open-Sourced Models Trained on Our Method			
Llama-3.1-8B-Instruct	4.7	20.5	26.4
Qwen2.5-7B-Instruct	9.3	10.3	11.4
Qwen3-4B-Instruct	14.9	18.5	19.5

BFCL v3 Evaluation Metrics. We follow the official BFCL v3 multi-turn evaluation. Each subset contains 200 multi-turn conversations, and we report **conversation-level accuracy (pass@1)**: a conversation is correct only if it passes BFCL’s end-of-turn checks for all turns. BFCL executes the predicted tool calls and verifies correctness via both state-based (backend state) and response-based checks. We use the default termination rules (stop when no valid tool call is produced; force-terminate at 20 tool steps and mark incorrect). Multi-turn Overall is the unweighted average over the four subsets.

Tool-Reflection-Bench Evaluation Metrics. To assess the model’s repair capability when tool calls fail, we used Tool-Reflection-Bench, with the evaluation metric being repair rate, Repair@n denotes that for the same data instance, if at least one out of n trials succeeds, the metric is recorded as 1; otherwise, it is 0.

Base Models. To verify the generalizability of Tool-Reflection-Bench and our training methodology, we conducted experiments using Llama3.1-8B (Dubey et al., 2024), Qwen2.5-7B-Instruct (Hui et al., 2024), and Qwen3-4B (Yang et al., 2025) as base models.

3.2 Experiment Results

3.2.1 Result on BFCL v3

Comparison with base models. Table 1 reports BFCL v3 multi-turn accuracy for the base models and our post-training. Overall, our method consistently improves multi-turn tool-calling across all three backbones, with the largest gains on error patterns that require argument repair and retrieving missing information from long histories.

Table 3: Ablation study on BFCL v3 multi-turn (conversation-level Pass@1 accuracy, %). We compare different RL variants on **Qwen2.5-7B-Instruct-FC**.

Base Model	RL Method	Base	Miss_Func	Miss_Param	Long_Context	Overall
Qwen2.5-7B-Instruct-FC	/	16.50	11.00	9.00	7.50	11.00
	DAPO	19.50	12.50	12.25	10.75	13.75
	GSPO	20.25	11.50	11.75	9.50	13.25
	Ours	22.00	13.00	13.50	11.00	14.88

Table 4: Ablation study on Tool-Reflection-Bench (Repair@n, %) for **Llama-3.1-8B-Instruct**. Prompt-only uses the inference prompt provided in the supplementary material (Appendix A.4).

Base Model	Method	Repair@1	Repair@3	Repair@5
Llama-3.1-8B-Instruct	/	0.70	5.10	6.80
	Prompt Only	1.50	9.30	12.60
	SFT	3.20	12.40	16.90
	RL	4.70	20.50	26.40

On **Llama-3.1-8B**, Base increases from 5.0 to 9.5 (+95%) and Long_Context from 4.5 to 7.0 (+56%). On **Qwen2.5-7B**, we observe the largest improvement on Miss_Param (9.0 \rightarrow 13.5, +50%), indicating stronger parameter correction under tool feedback. On **Qwen3-4B**, Multi-turn Overall increases from 16.25 to 20.75 (+28%), accompanied by a substantial Long_Context gain (+48%). In contrast, improvements on Miss_Func are smaller (e.g., 19.0 \rightarrow 19.5 on Qwen3-4B), suggesting that selecting the correct tool/function remains harder than repairing arguments given a plausible tool choice, which matches the focus of our reflection-driven repair training.

3.2.2 Result on Tool-Reflection-Bench

Tool-Reflection-Bench evaluates repair under tool-call failures; the test split consists solely of perturbation-derived failure cases (rather than clean trajectories), thus measuring recovery rather than memorization. As shown in Table 2, open-source baselines are weak at one try (Repair@1 \leq 9.6%) and improve only slightly with more attempts. Our post-training consistently boosts repair across all backbones: **Llama-3.1-8B-Instruct** improves from 0.7/5.1/6.8 to **4.7/20.5/26.4**, **Qwen2.5-7B-Instruct** from 2.4/6.1/8.0 to **9.3/10.3/11.4**, and **Qwen3-4B-Instruct** from 9.6/10.6/10.6 to **14.9/18.5/19.5** (Repair@1/3/5). The larger gains at Repair@3/5 suggest more robust reflection-to-repair behavior under repeated attempts. Our fine-tuned models also outperform closed-source baselines (e.g., **LongCat-Lite-8K-Chat** (Team et al.,

2025), **GPT-4o-mini** (OpenAI, 2024b,a), **GPT-4.1-mini** (OpenAI, 2025)) across $n \in \{1, 3, 5\}$.

3.2.3 Ablation Studies

Analysis. Table 3 shows that both DAPO and GSPO improve over the base model on BFCL v3 multi-turn, while our full method achieves the best overall performance across all four subsets. In particular, our gains are most pronounced on **Miss_Param** and **Long_Context**, consistent with our training signal that emphasizes reflection-driven argument repair and recovering missing information from long interaction histories. Table 4 further validates the necessity of learning-based post-training for repair: prompt-only self-correction provides a limited improvement, SFT yields larger gains, and RL delivers the strongest repair capability, especially under multiple trials (Repair@3/5), indicating that the reward-driven optimization better aligns the model with robust correction behavior beyond imitation alone.

4 Conclusion

This paper proposes a structured reflection method for handling tool call failures, transforming the “from error to repair” process into an explicit, controllable, and trainable action. Our approach overcomes the limitations of previous heuristic, feedback-based self-correction methods in terms of controllability and stability. We further construct Tool-Reflection-Bench for both training and evaluation, and design a task-specific reward function tailored to the tool-calling scenario. In the reinforcement learning stage, we combine the strengths of DAPO and GSPO to enhance training effectiveness. Experimental results show that the proposed method significantly improves multi-turn tool call accuracy on BFCL v3 as well as error repair performance on Tool-Reflection-Bench. Overall, our method and dataset effectively enhance the reliability of tool interactions and offer a new perspective on enabling agents to acquire new capabilities by learning from failure.

5 Limitations

Generalization beyond designed failures. Our training data is constructed via a set of perturbation-based disruptions that cover several common categories of tool-calling failures (e.g., missing/incorrect arguments and long-context dependencies). While our results on BFCL v3 multi-turn and Tool-Reflection-Bench demonstrate consistent improvements, these evaluations still represent a bounded set of failure modes. In real-world deployments, agents may encounter additional error types, such as ambiguous user intents, non-stationary tool APIs, partial tool outages, noisy tool outputs, or multi-tool plans with stronger temporal/causal dependencies. Extending perturbations and supervision to better approximate real failure distributions, as well as incorporating logs of naturally occurring failures, are promising directions.

Scaling to larger models. We validate our method on 4B–8B class models and observe substantial gains. However, the behavior of much larger models (e.g., 70B+) may differ: stronger base capabilities can reduce headroom, and the optimal balance between supervised signals and RL shaping may change. Evaluating and adapting the training recipe for larger backbones—including whether lighter post-training (e.g., SFT-only or fewer RL steps) suffices—remains future work.

References

- Chen Chen, Xinlong Hao, Weiwen Liu, Xu Huang, Xingshan Zeng, Shuai Yu, Dexun Li, Shuai Wang, Weinan Gan, Yuefeng Huang, and 1 others. 2025a. Acebench: Who wins the match point in tool learning? *arXiv e-prints*, pages arXiv–2501.
- Hardy Chen, Haoqin Tu, Fali Wang, Hui Liu, Xianfeng Tang, Xinya Du, Yuyin Zhou, and Cihang Xie. 2025b. Sft or rl? an early investigation into training rl-like reasoning large vision-language models. *arXiv preprint arXiv:2504.11468*.
- Mingyang Chen, Haoze Sun, Tianpeng Li, Fan Yang, Hao Liang, Keer Lu, Bin Cui, Wentao Zhang, Zenan Zhou, and Weipeng Chen. 2024. Facilitating multi-turn function calling for llms via compositional instruction tuning. *arXiv preprint arXiv:2410.12952*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407.
- Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, Baoquan Zhong, Chengquan Jiang, Jinxin Chi, and Wanjun Zhong. 2025. Retool: Reinforcement learning for strategic tool use in llms. *arXiv preprint arXiv:2504.11536*.
- Bingguang Hao, Maolin Wang, Zengzhuang Xu, Cunyin Peng, Yicheng Chen, Xiangyu Zhao, Jinjie Gu, and Chenyi Zhuang. 2025. Funreason: Enhancing large language models’ function calling via self-refinement multiscale loss and automated data refinement. *arXiv preprint arXiv:2505.20192*.
- Yilun Hao, Yongchao Chen, Yang Zhang, and Chuchu Fan. 2024. Large language models can plan your travels rigorously with formal verification tools. *CoRR*.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2023. Large language models cannot self-correct reasoning yet. *arXiv preprint arXiv:2310.01798*.
- Shijue Huang, Wanjun Zhong, Jianqiao Lu, Qi Zhu, Jiahui Gao, Weiwen Liu, Yutai Hou, Xingshan Zeng, Yasheng Wang, Lifeng Shang, and 1 others. 2024. Planning, creation, usage: Benchmarking llms for comprehensive tool utilization in real-world complex scenarios. *arXiv preprint arXiv:2401.17167*.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, and 1 others. 2024. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*.
- Yuhua Jiang, Yuwen Xiong, Yufeng Yuan, Chao Xin, Wenyan Xu, Yu Yue, Qianchuan Zhao, and Lin Yan. 2025. Pag: Multi-turn reinforced llm self-correction with policy as generative verifier. *arXiv preprint arXiv:2506.10406*.
- Barrett Martin Lattimer, Varun Gangal, Ryan McDonald, and Yi Yang. 2024. Sparse rewards can self-train dialogue agents. *arXiv preprint arXiv:2409.04617*.
- Minghao Li, Yingxiu Zhao, Bowen Yu, Feifan Song, Hangyu Li, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023. Api-bank: A comprehensive benchmark for tool-augmented llms. *arXiv preprint arXiv:2304.08244*.
- Xuefeng Li, Haoyang Zou, and Pengfei Liu. 2025. Torl: Scaling tool-integrated rl. *arXiv preprint arXiv:2503.23383*.
- Fengyuan Liu, Nouar AlDahoul, Gregory Eady, Yasir Zaki, and Talal Rahwan. 2024a. Self-reflection makes large language models safer, less biased, and ideologically neutral. *arXiv preprint arXiv:2406.10400*.
- Weiwen Liu, Xu Huang, Xingshan Zeng, Xinlong Hao, Shuai Yu, Dexun Li, Shuai Wang, Weinan Gan, Zhengying Liu, Yuanqing Yu, and 1 others. 2024b. Toolace: Winning the points of llm function calling. *arXiv preprint arXiv:2409.00920*.

- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, and 1 others. 2023. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36:46534–46594.
- OpenAI. 2024a. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276. Accessed: 2025-09-25.
- OpenAI. 2024b. Hello gpt-4o. <https://openai.com/index/hello-gpt-4o/>. Accessed: 2025-09-25.
- OpenAI. 2025. Introducing gpt-4.1 in the api. <https://openai.com/index/gpt-4-1/>. Accessed: 2025-09-25.
- Shishir G Patil, Huanzhi Mao, Fanjia Yan, Charlie Cheng-Jie Ji, Vishnu Suresh, Ion Stoica, and Joseph E Gonzalez. The berkeley function calling leaderboard (bfcl): From tool use to agentic evaluation of large language models. In *Forty-second International Conference on Machine Learning*.
- Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiusi Chen, Dilek Hakkani-Tür, Gokhan Tur, and Heng Ji. 2025. Toolrl: Reward is all tool learning needs. *arXiv preprint arXiv:2504.13958*.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, and 1 others. 2023. Toolllm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*.
- Changle Qu, Sunhao Dai, Xiaochi Wei, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, Jun Xu, and J Wen. 2024a. Tool learning with large language models: A survey. corr abs/2405.17935(2024). *arXiv preprint arXiv:2405.17935*.
- Changle Qu, Sunhao Dai, Xiaochi Wei, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, Jun Xu, and J Wen. 2024b. Tool learning with large language models: A survey. corr abs/2405.17935(2024). *arXiv preprint arXiv:2405.17935*.
- Changle Qu, Sunhao Dai, Xiaochi Wei, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, Jun Xu, and Ji-Rong Wen. 2025. From exploration to mastery: Enabling llms to master tools via self-driven interactions. In *International Conference on Learning Representations (ICLR)*.
- Matthew Renze and Erhan Guven. 2024. Self-reflection in llm agents: Effects on problem-solving performance. *arXiv preprint arXiv:2405.06682*.
- RI Saveliev and MV Dendiuk. 2024. Self-reflective retrieval-augmented generation (self-rag) in analytical systems. In *Forestry Education and Science: Current Challenges and Development Prospects. International Science-Practical Conference, October 23-25, 2024, Lviv, Ukraine*.
- Meituan LongCat Team, Bei Li, Bingye Lei, Bo Wang, Bolin Rong, Chao Wang, Chao Zhang, Chen Gao, Chen Zhang, Cheng Sun, and 1 others. 2025. Longcat-flash technical report. *arXiv preprint arXiv:2509.01322*.
- Adrian Theuma and Ehsan Shareghi. 2024. Equipping language models with tool use capability for tabular data analysis in finance. *arXiv preprint arXiv:2401.15328*.
- Juraj Vladika, Ihsan Soydemir, and Florian Matthes. 2025. Correcting hallucinations in news summaries: Exploration of self-correcting llm methods with external knowledge. *arXiv preprint arXiv:2506.19607*.
- MAOLIN WANG, YINGYI ZHANG, CUNYIN PENG, YICHENG CHEN, WEI ZHOU, JINJIE GU, CHENYI ZHUANG, RUOCHENG GUO, BOWEN YU, WANYU WANG, and 1 others. 2025. Function calling in large language models: Industrial practices, challenges, and future directions.
- Renxi Wang, Xudong Han, Lei Ji, Shu Wang, Timothy Baldwin, and Haonan Li. 2024. Toolgen: Unified tool retrieval and calling via generation. *arXiv preprint arXiv:2410.03439*.
- Zhenyu Wu, Qingkai Zeng, Zhihan Zhang, Zhaoxuan Tan, Chao Shen, and Meng Jiang. 2024. Large language models can self-correct with key condition verification. *arXiv preprint arXiv:2405.14092*.
- Qiancheng Xu, Yongqi Li, Heming Xia, and Wenjie Li. 2024. Enhancing tool retrieval with iterative feedback from large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 9609–9619.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Ling Yang, Zhaochen Yu, Tianjun Zhang, Minkai Xu, Joseph E Gonzalez, Bin Cui, and Shuicheng Yan. 2024. Supercorrect: Supervising and correcting language models with error-driven insights. *arXiv preprint arXiv:2410.09008*, 9.
- Junjie Ye, Yilong Wu, Sixian Li, Yuming Yang, Tao Gui, Qi Zhang, Xuanjing Huang, Peng Wang, Zhongchao Shi, Jianping Fan, and 1 others. 2024. Tl-training: A task-feature-based framework for training large language models in tool use. *arXiv preprint arXiv:2412.15495*.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, and 1 others. 2025. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*.

Jianguo Zhang, Tian Lan, Ming Zhu, Zuxin Liu, Thai Hoang, Shirley Kokane, Weiran Yao, Juntao Tan, Akshara Prabhakar, Haolin Chen, and 1 others. 2024. xlam: A family of large action models to empower ai agent systems. *arXiv preprint arXiv:2409.03215*.

Xutong Zhao, Tengyu Xu, Xuewei Wang, Zhengxing Chen, Di Jin, Liang Tan, Zishun Yu, Zhuokai Zhao, Yun He, Sinong Wang, and 1 others. 2025a. Boosting llm reasoning via spontaneous self-correction. *arXiv preprint arXiv:2506.06923*.

Yuze Zhao, Juntao Huang, Jinghan Hu, Xingjun Wang, Yunlin Mao, Daoze Zhang, Zeyinzi Jiang, Zhikai Wu, Baole Ai, Ang Wang, and 1 others. 2025b. Swift: a scalable lightweight infrastructure for fine-tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 29733–29735.

Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, and 1 others. 2025. Group sequence policy optimization. *arXiv preprint arXiv:2507.18071*.

Ruizhe Zhong, Xingbo Du, Shixiong Kai, Zhentao Tang, Siyuan Xu, Hui-Ling Zhen, Jianye Hao, Qiang Xu, Mingxuan Yuan, and Junchi Yan. 2023. Llm4eda: Emerging progress in large language models for electronic design automation. *arXiv preprint arXiv:2401.12224*.

A Appendix

A.1 Use of LLMs

This work leveraged LLMs to verify the mathematical soundness and symbolic accuracy of a few formulas in Sec.A.8.

A.2 Related Works

A.2.1 Tool-augmented Large Language Models

Integrating external tools into large language models has become a key approach to enhancing their functionality, surpassing the simple task of text generation. Traditional LLMs are limited by static knowledge, constrained to the data they were trained on. However, tool-augmented models extend the capabilities of LLMs by enabling them to interact with external resources (Zhang et al., 2024; Hao et al., 2025) (such as APIs (Li et al., 2023), databases, and computational engines) through tool calls. This extension allows LLMs to access real-time data, perform external computations, and even interface with external hardware, making them more practical for solving complex real-world tasks that require dynamic information or specific external operations (Chen et al., 2025a). ToolBench

(Qin et al., 2023) demonstrates the feasibility of integrating external tool calls into LLMs. Through such systems, LLMs can handle more specialized tasks. However, one major challenge of tool augmentation is how to effectively train LLMs to use these tools. Existing training methods, such as supervised fine-tuning and reinforcement learning, typically focus on optimizing single tool calls. This type of interaction often does not involve multi-turn tool calls or responses, which makes the limitations of current methods particularly apparent when errors occur during tool usage. In such cases, the model’s ability to recover from errors becomes crucial.

A.2.2 Self-Correction in LLMs

Self-correction in large language models refers to the model’s ability to diagnose its own errors and correct them based on previous actions (Huang et al., 2023; Liu et al., 2024a). However, this area has not been fully explored. Existing self-correction techniques mostly rely on heuristic methods or unidirectional reasoning processes (Renze and Guven, 2024).

Self-Refine framework (Madaan et al., 2023), which involves having LLMs provide an initial response, followed by a reflection process where the model identifies flaws and makes improvements. Specifically, the same LLM acts as both the responder and the evaluator: the model first generates an initial response, then self-reflects and iteratively revises the output. This approach has been shown to enhance the performance of LLMs in certain domains. However, subsequent studies (Wu et al., 2024; Vladika et al., 2025) have found that relying solely on the model itself often fails to detect subtle errors. Some research (Jiang et al., 2025; Zhao et al., 2025a) has introduced auxiliary verifiers (such as additional models or mechanisms (Saveliev and Dendiuk, 2024; Feng et al., 2025)) to help check the correctness of the initial response. This external self-checking assistance avoids unnecessary repeated revisions, improving efficiency and enhancing the model’s reasoning and verification capabilities. However, this approach remains highly sensitive to the specific phrasing of the prompts, with different prompt wordings leading to varying results (Liu et al., 2024a).

Beyond prompt-only self-revision, recent work in tool-augmented agents improves reliability via interaction/execution feedback loops, e.g., unifying tool retrieval and calling via generation (Wang

et al., 2024), refining tool retrieval with iterative LLM feedback (Xu et al., 2024), and trial-and-error frameworks that leverage tool interaction signals to iteratively improve tool understanding/documentation (Qu et al., 2025).

However, even though these methods improve tool-use robustness, they primarily focus on leveraging online interaction signals (e.g., tool returns, success/failure) to drive behavior improvements. In contrast, our work targets self-correction itself as an explicit, trainable, and controllable capability: we formulate reflection-based error localization, diagnosis, and repair as a learning objective, and optimize it via supervised signals and reward shaping so that the model can reliably trigger and perform corrections during tool calling.

A.3 Cost-Benefit Analysis on Tool-Reflection-Bench

Analysis. Table 5 quantifies the trade-off between RL complexity and performance under a fixed training budget. Across DAPO, GSPO, and our method, the wall-clock cost is nearly identical (all ≈ 8 hours on $4 \times A100$ for 1000 steps), indicating that our pipeline does not introduce meaningful additional training overhead in this setting. Despite comparable cost, our method achieves the best repair capability, improving Repair@5 from 20.8 (DAPO) / 20.5 (GSPO) to 26.4, and yielding a notably larger gain at higher n (Repair@3/5), which suggests more reliable multi-try recovery when tool calls fail.

A.4 Prompt for Prompt-only Self-correction

In this section, we provide a simplified prompt used for the *Prompt Only* baseline in our ablation studies. The prompt enforces a fixed output format (`<reflect>` + `<call>`) to encourage explicit diagnosis before proposing a corrected tool call.

How to run prompt-only self-correction

System

Role. You are a tool-calling assistant.

Goal. A previous tool call failed. Given the user goal, recent context, the failed tool call, and the tool error message, produce (1) a concise reflection diagnosing the failure, and (2) exactly one corrected tool call that is executable under the provided tool schema.

Inputs. You will be given:

- [User Goal]: the task the assistant is trying to accomplish.
- [Tool API / Schema]: function names and argument schemas.

- [Context]: recent messages and tool results.
- [Failed Tool Call]: the tool call that triggered the error.
- [Tool Error / Feedback]: the tool response/error message.

Output format (strict). Output *must* follow this exact format with no extra text:

```
<reflect>
- Error type: (wrong tool / missing tool / wrong arguments
  ↳ / missing arguments / wrong values / redundant
  ↳ call / long-context mismatch).
- Evidence: cite the tool error and/or context that
  ↳ reveals the issue.
- Fix: state the minimal change needed to repair the call.
</reflect>
<call>{"name": "...", "arguments": {...}}</call>
```

Constraints.

1. Produce exactly one `<reflect>` block and exactly one `<call>` block.
2. The `<call>` must be valid JSON and executable under the provided schema.
3. Do *not* output `<final>...</final>` or any additional natural-language response.
4. Prefer minimal edits: only change what is necessary to fix the error.

User

```
[User Goal]
{USER_GOAL}

[Tool API / Schema]
{TOOL_SCHEMA_OR_DOCS}

[Context (recent messages and tool results)]
{CONTEXT}

[Failed Tool Call]
<call>{FAILED_CALL_JSON}</call>

[Tool Error / Feedback]
{ERROR_MESSAGE}
```

A.5 Prompt for Perturbation-based Disruptions

In this section, we provide simplified prompts for generating the four types of tool call perturbations, enabling the community to reproduce our setting. The full prompts and implementation code will be released upon the paper’s acceptance.

A.5.1 Prompt for Call-Order Swap

How to construct an error tool call example

System

Goal. Prepend a controlled erroneous `<call>` and a consistent tool-error message before the first assistant message, so the model must diagnose and repair.

Procedure.

Table 5: **Cost–benefit comparison on Tool-Reflection-Bench for Llama-3.1-8B-Instruct.** We report Repair@ n (% , higher is better) and training cost measured in wall-clock time on $4\times\text{A100}$ GPUs. All RL variants are trained for **1000 steps** with identical data, sampling, and optimization settings; “/” denotes the untrained base model.

Base Model	RL Method	Repair@1	Repair@3	Repair@5	GPU Time ($4\times\text{A100}$)
Llama-3.1-8B-Instruct	/	0.70	5.10	6.80	/
	DAPO	3.90	14.90	20.80	7h48m
	GSPO	3.20	13.70	20.50	7h55m
	Ours	4.70	20.50	26.40	7h56m

1. *Extract calls:* Traverse messages and collect all assistant `<call>...</call>` blocks (regex).
2. *Choose function name:* Parse the last call’s JSON to get “name”; fall back to a regex if needed.
3. *Synthesize wrong call (empty args):*

```
<call>[{"name": "<FUNC_FROM_LAST_CALL>", "arguments": "{}"}]</call>
```
4. *Fabricate tool error (pretty JSON string):*

```
{ "tool": "<FUNC_FROM_LAST_CALL>", "status": "warning", "message": "The called function executed but\n    ↳ returned partial/mismatched data because\n    ↳ the arguments did not match the expected\n    ↳ schema for this call.", "result": null }
```
5. *Insert pair:* Place the wrong assistant call and the tool error *before* the original first assistant message.
6. *Elicit reflection:* Query the LLM with the System/User prompts above to obtain the reflection text, then prepend `<reflect>...</reflect>` to the original assistant message (the original correct call remains).

Notes. Using the last call’s function ensures schema plausibility; empty arguments induce a controlled failure; the synthetic tool message supplies concrete evidence for the subsequent reflection and repair.

How to generate a reflection

System

You are an AI assistant that analyzes failed tool calls and provides reflective summaries. Given an original tool call and a fabricated error response, generate a brief reflection explaining why the call likely failed and how to correct it. Be concrete and concise.

User

Fill the placeholders `{{...}}` exactly.

Original tool call:

```
{{ORIGINAL_CALL}}
```

Error response:

```
{{FAKE_RESPONSE}}
```

Please provide a short reflection on the failure cause and the corrective action.

An Example

User

Original tool call:

```
<call>[{"name": "searchArtistsByArtStyle", "arguments": "{}"}]</call>
```

Error response:

```
{ "tool": "searchArtistsByArtStyle", "status": "warning", "message": "The called function executed but returned\n    ↳ partial/mismatched data because the arguments did\n    ↳ not match the expected schema for this call.", "result": null }
```

Please provide a brief reflection on why this tool call failed and what could be improved. Keep it concise and helpful.

A.5.2 Prompt for Redundant Call

How to construct a redundant tool call example

System

Goal. Inject a *redundant* tool call inside an existing `<call>` list and a matching redundant tool response, so the agent must identify and remove the duplication.

Procedure.

1. *Extract calls:* Traverse the dialogue and collect all assistant-side `<call>...</call>` blocks (regex).
2. *Pick a target (not the first):* Uniformly sample an assistant call position from $\{2, \dots, |C|\}$.
3. *Duplicate within the list:* Parse the target call’s JSON. If it is a list, append a deep-copied first element; if it is a single dict, make a two-element list by duplicating it.
4. *Fabricate a redundant tool response:* Parse the following tool message. Duplicate its first item (or the dict itself) and mark it as redundant, e.g.

```
{ "status": "redundant", "message": "This item\n    ↳ duplicates a previous result." }
```
5. *Keep the ground-truth call:* The *correct* call is the original (non-duplicated) first element of the target call list.
6. *Place the repair evidence:* After the redundant tool message, insert an assistant message with `<reflect>` diagnosing the redundancy and a correct `<call>` (the non-duplicated one), followed by a *clean* tool response (the original, without the redundant copy).

Notes. This perturbation preserves schema but injects duplication at both call and response sides, creating a realistic “over-call” pattern for reflection-and-repair.

How to generate a reflection

System

You are an AI assistant that analyzes *redundant* tool calls and provides reflective summaries. Given a tool-call list and its redundant tool response, write a brief reflection that (i) identifies the duplication, and (ii) states the correct next action (use only the necessary call with proper arguments). Keep the reflection concise and actionable.

User

Fill the placeholders `{{...}}` exactly.

Tool call list (after duplication):

`{{TOOL_CALL_LIST}}`

Redundant tool response:

`{{REDUNDANT_RESPONSE}}`

Please provide a short reflection that points out the redundancy and explains how to proceed correctly.

An Example

User

Tool call list (after duplication):

```
<call>[
  {"name": "searchArtistsByArtStyle", "arguments": {"style": "
    ↳ impressionism"}},
  {"name": "searchArtistsByArtStyle", "arguments": {"style": "
    ↳ impressionism"}}
]</call>
```

Redundant tool response:

```
[
  {"tool": "searchArtistsByArtStyle", "status": "ok", "items
    ↳ ": [...]},
  {"tool": "searchArtistsByArtStyle", "status": "redundant",
    "message": "This item duplicates a previous result.", "
    ↳ items": [...]}
]
```

Please provide a brief reflection on why this redundant call occurred and how to proceed. Keep it concise and helpful.

A.5.3 Prompt for Missing Call

How to construct a missing-call perturbation example

System

Goal. Remove a necessary assistant `<call>` and make the subsequent call fail due to missing context, so the agent must *recover the omitted call* and then proceed correctly.

Procedure.

1. *Extract calls:* Parse all assistant-side `<call>...</call>` blocks (regex).
2. *Select a removable call (not the last):* Uniformly sample an index $i \in \{1, \dots, |C| - 1\}$.
3. *Find paired tool messages:* Locate the tool reply immediately after call i (the one to remove), and the tool reply after call $i+1$ (the “next” call).
4. *Delete call i and its tool reply.*

5. *Degrade the next call:* For the assistant `<call>` at (original) $i+1$, keep the function but set “arguments”: `{}` (empty).

6. *Return an error for the next tool:* Replace that tool reply with an error JSON indicating “missing required arguments”.

7. *Reflection and repair insertion:* After the error tool reply, insert:

- (a) an assistant message containing `<reflect>` that explains the omission and a *reinstated* correct `<call>` (the removed call i);
- (b) the original tool reply for the removed call i ;
- (c) the corrected next assistant call (its original, non-empty arguments);
- (d) the corrected next tool reply (its original content).

Notes. This perturbation creates a realistic “missing prerequisite call” failure: the subsequent step cannot execute without information from the omitted call. The reflection must (i) identify the omission and (ii) restore the correct call before proceeding.

How to generate a reflection

System

You are an AI assistant that analyzes *missing* tool calls and provides reflective summaries. Given the omitted call (that should have been executed) and the resulting error response from the next step, write a concise reflection that (i) identifies what was missing, and (ii) states how to proceed: first reinstate the omitted call with correct arguments, then continue.

User

Fill the placeholders `{{...}}` exactly.

Missing tool call (the one that should have been made):

`{{MISSING_CALL}}`

Error response (from the next step):

`{{ERROR_RESPONSE}}`

Please provide a short reflection that explains the omission and the corrective sequence of actions.

An Example

User

Missing tool call:

```
<call>[{"name": "fetchUserProfile", "arguments": {"user_id": "
  ↳ u_1293"}}]</call>
```

Error response (from the next step):

```
[
  {"status": "error",
    "message": "Missing required arguments. The function
      ↳ call failed because necessary parameters were not
      ↳ provided.",
    "result": null}
]
```

Please provide a brief reflection on what was missing and how to proceed. Keep it concise and helpful.

A.5.4 Prompt for Argument Error

How to construct an argument-error perturbation example

System

Goal. Corrupt the arguments of an existing assistant `<call>` so that the paired tool reply returns a parameter-validation error, forcing the agent to *diagnose mismatched/invalid arguments* and repair with the correct call.

Procedure.

1. *Extract calls:* Parse all assistant-side `<call>...</call>` blocks via regex.
2. *Select a call:* Uniformly sample one index $i \in \{1, \dots, |C|\}$ and locate its immediate tool reply.
3. *Corrupt arguments:* Keep "name" unchanged; replace "arguments" with perturbed values (e.g., wrong types, out-of-range numbers, empty strings, unknown keys). The JSON stays well-formed:

```
<call>[{"name": "<FUNC_NAME>", "arguments": {<WRONG_ARGS>}}]</call>
```
4. *Synthesize error reply:* Replace the paired tool message with a structured error indicating invalid parameters (e.g., "error_code": "INVALID_PARAMETERS" and an informative message).
5. *Reflection and repair insertion:* Immediately after the error, insert:
 - (a) an assistant message with `<reflect>` that contrasts the wrong vs. correct arguments and states the fix;
 - (b) the *original* (correct) call and its original (successful) tool reply.

Notes. Do not alter the function name; only arguments are corrupted. Keep JSON/tags valid to isolate the failure mode to argument errors.

How to generate a reflection

System

You are an AI assistant that analyzes incorrect tool-call *parameters* and provides a reflective summary. Given the correct call, the wrong call (with corrupted arguments), and the error response, write a brief reflection that (i) pinpoints which arguments are incorrect and why, and (ii) states the corrected call. Be concrete and concise.

User

Fill the placeholders `{{...}}` exactly.

Correct tool call (ground truth):

`{{CORRECT_CALL}}`

Wrong tool call made (arguments corrupted):

`{{WRONG_CALL}}`

Error response:

`{{ERROR_RESPONSE}}`

Please provide a short reflection that identifies the parameter issues and the corrective action.

An Example

User

Correct tool call:

```
<call>[{"name": "bookFlight",  
  "arguments": {"from": "SFO", "to": "JFK", "date":  
    ↳ "2025-11-02", "passengers": 1}}]</call>
```

Wrong tool call made:

```
<call>[{"name": "bookFlight",  
  "arguments": {"from": 999999, "to": "", "date": null, "  
    ↳ passengers": "many"}}]</call>
```

Error response:

```
[{"status": "error",  
  "message": "Parameter validation failed for bookFlight.  
    ↳ One or more arguments are invalid.",  
  "result": null,  
  "error_code": "INVALID_PARAMETERS"}]
```

Please provide a brief reflection on which parameters are incorrect and how to fix them. Keep it concise and helpful.

A.6 Training Data Case Study

In this section, we extract one complete sample from each of the four perturbation modes for analysis. Due to space limitations, these four samples are provided in the supplementary material for reference, while here we only present a brief analysis of the data.

A.6.1 Case Study of Call-Order Swap

Setup. The user requests end-to-end logistics for a 10-person business trip (NYC → MIA): search and *book* round-trip flights, search and *book* hotel rooms, and arrange airport-hotel ground transportation. The toolset exposes `search_flights`, `book_flight`, `search_hotels`, `book_hotel`, and `arrange_transportation`, where the last function *depends* on confirmed itinerary details (arrival time and hotel).

Failure trace (order dependency violated). The agent correctly issues two discovery calls, then prematurely attempts to arrange transportation *before* booking:

```
<tool_call>{"name": "search_flights", ...}</tool_call>  
<tool_call>{"name": "search_hotels", ...}</tool_call>  
  
<tool_call>{"name": "arrange_transportation",  
  "arguments": {"arrival_date": "2023-03-21",  
    "passenger_count": 10,  
    "pickup_location": "MIA Airport",  
    "dropoff_location": "(to be decided)"  
    ↳ "}}</tool_call>
```

The tool returns a dependency error because `dropoff_location` cannot be finalized prior to `book_hotel` and pickup timing should be aligned to `book_flight`:

```
Tool: {"status": "error",
      "message": "Transportation requires confirmed flight and
      ↳ hotel details.",
      "result": null}
```

Model reflection and repair. The agent emits a concise `<reflect>` that identifies the *order-dependency* (transport must follow booking), then repairs by booking first and only then arranging transport:

```
<reflect>
Failure stems from an order dependency: transportation must
↳ follow confirmed
flight and hotel, as pickup/dropoff depend on the booked
↳ itinerary.
Correct plan: (1) book flight; (2) book hotel; (3) arrange
↳ transportation.
</reflect>

<tool_call>{"name": "book_flight", ...}</tool_call>
<tool_call>{"name": "book_hotel", ...}</tool_call>
<tool_call>{"name": "arrange_transportation",
  "arguments": {"arrival_date": "2023-03-21",
               "passenger_count": 10,
               "pickup_location": "MIA",
               "dropoff_location": "Oceanview Hotel
  ↳ "}}</tool_call>
```

Why the reward favors the fix. Under our reward in Sec. §2.2: (i) the reflection receives a high semantic score s_{ref} for explicitly diagnosing the *order* constraint; (ii) the repaired calls exactly match schema and parameters, yielding $s_{\text{call}} = 1$ (order-invariant matching across the multiset of calls), while the *format/penalty factor* F upweights the repaired trajectory because executability and dependency checks now pass; (iii) the premature `arrange_transportation` in the failed rollout triggers penalties in F (dependency/executability), so R_{core} is suppressed and the dynamic sampling filter (Sec. §2.3) down-weights such uninformative rollouts. Overall, the corrected plan yields a strictly larger R_{total} than the erroneous one.

Summary. This example shows that explicit reflection converts a latent, long-range constraint (“book before transport”) into a short textual diagnosis that the policy can optimize against. Our order-invariant call equality avoids over-penalizing benign output permutations, while dependency-aware format shaping in F enforces *inter-call* preconditions. Empirically, such cases account for a substantial share of multi-turn gains, whereas single-turn metrics remain stable.

A.6.2 Case Study of Redundant Call

Setup and failure mode. The user requests the list of services for all NYC banks operating for > 10 years. The agent first queries the bank roster with `get_operational_banks`, then correctly

queries services for B001, but erroneously *repeats* the same service lookup for B001 instead of proceeding to unseen banks (B002–B004). This is a canonical *redundant call* error: wasted tool budget without adding new information.

```
...
<tool_call>
{"name": "get_operational_banks",
 "arguments": {"city": "New York City", "min_years_operational": 10}}
</tool_call>

User (tool): [{"name": "get_operational_banks", "results": [{"name": "get_operational_banks", "results": [{"bank_id": "B001"}, {"bank_id": "B002"}, {"bank_id": "B003"}, {"bank_id": "B004"}]}]}]

<tool_call>
{"name": "get_bank_services", "arguments": {"bank_id": "B001"}}
</tool_call>

User (tool): [{"name": "get_bank_services", "arguments": {"bank_id": "B001"}, "results": {"services": [...]}]}]

# Redundant (erroneous) repeat:
<tool_call>
{"name": "get_bank_services", "arguments": {"bank_id": "B001"}}
</tool_call>

User (tool): [{"name": "get_bank_services", "arguments": {"bank_id": "B001"}, "results": {"services": [...], "note": "Duplicate request"}]}]
...
```

Reflection and repair (ground truth). The reflection diagnoses duplication and prescribes deduplication; the corrected plan proceeds with the remaining, unseen IDs:

```
<reflect>
The last call repeated B001, yielding redundant information.
Avoid duplicates; continue with B002, B003, B004 exactly once
↳ each.
</reflect>

<tool_call>{"name": "get_bank_services", "arguments": {"bank_id": "B002"}}</tool_call>
<tool_call>{"name": "get_bank_services", "arguments": {"bank_id": "B003"}}</tool_call>
<tool_call>{"name": "get_bank_services", "arguments": {"bank_id": "B004"}}</tool_call>
```

Why the model failed. The failure arises from (i) insufficient state tracking over the set of already-seen entities (here, bank IDs), and (ii) weak inductive bias against issuing calls whose *marginal information gain* is near zero. In multi-turn settings, local myopic policies often re-issue the last successful pattern without cross-step deduplication.

How the reward steers recovery. Our scoring treats call sets as order-invariant but schema-strict; redundant calls trigger the count-mismatch component in the format factor F (penalizing $|C_{\text{calls}}| \neq |G_{\text{calls}}|$) while `EqualCalls` fails due to multiset mismatch. The reflection text receives a positive

semantic score if it explicitly identifies the duplication and prescribes the missing IDs, encouraging concise, actionable self-correction. Together, the structure score S and format factor F downweight redundant completions and upweight the repaired sequence.

Summary. This case shows that explicit reflection converts a silent efficiency bug into a supervised correction step: the agent (1) cites the duplicated identifier, (2) enumerates the remaining targets, and (3) completes them exactly once. Empirically, such reflection-shaped supervision reduces redundant tool usage and improves multi-turn success without harming single-turn accuracy.

A.6.3 Case Study of Missing Call

Setup. The user asks to register *four* tax documents: (i) W-2 (ABC Corp), (ii) 1099-INT (First National Bank), (iii) property tax statement (county assessor), and (iv) Form 1098 (mortgage lender). The tool schema exposes a single function `add_tax_documents(name, value, category, priority)` with `name`, `value` required.

Baseline failure (*missing calls*). The baseline assistant emits only two `<tool_call>`s (W-2, 1099-INT) and then stops, yielding a 50% recall on required calls. Formally, let G_{calls} contain the four intended calls and C_{calls} the two produced calls. Then $|G_{\text{calls}}| = 4$, $|C_{\text{calls}}| = 2$, and the call-set equality test fails: $\text{EqualCalls}(C_{\text{calls}}, G_{\text{calls}}) = 0$. This is a typical *missing-call* error in multi-item requests: the model recognizes the pattern “one item \rightarrow one call” but truncates the sequence, leaving later items unprocessed.

Structured reflection (*diagnosis*). Our method takes the partially executed trajectory as *negative evidence* and the original request as *positive intent* and generates an explicit reflection:

```
<reflect> “I missed 2 tool call(s). The user
listed multiple items, and each item requires a
separate call. I should enumerate all items and
complete the remaining calls.” </reflect>
```

The reflection correctly localizes the failure (undercounting of required calls), quantifies the deficit (missed= 2), and states the repair rule (enumerate all items \Rightarrow one call per item).

Repairs (*corrective calls*). Conditioned on the reflection, the agent appends the missing tool calls for the remaining items:

- name: Property tax statement; value: county assessor record; category: personal;
- name: Form 1098; value: mortgage interest statement; category: personal.

The assignments work \rightarrow W-2, 1099-INT and personal \rightarrow property tax, 1098 are semantically consistent: the former are employment/bank income records; the latter are household liabilities/taxes. (Any schema-compatible categorization would pass executability; ours also preserves natural semantics.)

Why this matters. This case highlights a frequent multi-turn brittleness: once the agent produces a plausible prefix of calls, it prematurely concludes and fails to cover all requested items. By making *missingness* an explicit, trainable concept, structured reflection converts a sparse binary signal (success/failure) into actionable supervision:

1. **Detection:** Compare item cardinalities and arguments; compute $\mathbb{I}[|C_{\text{calls}}| < |G_{\text{calls}}|]$ and list uncovered entities.
2. **Diagnosis:** Attribute the error to *enumeration/coverage* rather than formatting or parameters.
3. **Repair:** Synthesize the exact missing calls with schema-valid arguments; preserve already-correct calls.

Summary. Empirically, such instances improve the model’s *coverage discipline*: after training, we observe higher multi-item completion rates with negligible increase in redundant calls, indicating that the model learned “one-mentioned-item \Rightarrow one-call” as a robust policy rather than overcalling.

Setup. The tool schema exposes multiple functions with *schema-strict* parameters:

- `|check_plant_water_level(plant_location:string)|`
- `|start_watering(plant_location:string, duration:number)|`
- `|start_trimming(hedge_location:string)|`
- ...

The user requests two primary actions in the backyard: (i) trim hedges and (ii) water all potted plants for about 10 minutes; afterwards ensure plants have enough water and dispose clippings

Baseline failure (*argument error*). The assistant issues

```
<call>[{"name": "check_plant_water_level", "arguments": {}}]</
  ↳ call>
```

omitting the *required* key `plant_location`. The tool returns a schema warning that the arguments “did not match expected schema.” Under our reward, the call-level indicator s_{call} is 0 because the produced call fails schema equality (tool name matches, but the argument map does not).

Structured reflection (*diagnosis*). The reflection generated by our process states that the call “failed because it did not include the required arguments needed by the function’s schema,” and prescribes: “ensure all necessary parameters are provided according to the function’s documentation.” This localizes the error to **parameter mis-specification** (not tool selection or ordering), and points to the concrete fix—satisfy the schema.

Repairs (*efficient plan consistent with the request*). Given the user’s 10-minute target and the backyard scope, the corrected action set executes the two core operations with schema-valid arguments:

- `|start_watering(plant_location="backyard", duration=10)|`
- `|start_trimming(hedge_location="backyard")|`

These can be dispatched in parallel (independent resources), achieving the requested time budget while ensuring plants receive sufficient water and hedges are trimmed. This replaces the invalid pre-check with a direct, time-bounded watering call that already satisfies the user’s constraint.

Why this matters. Argument errors are common in tool use and typically yield *sparse* feedback (“schema mismatch”). By forcing the model to (i) recognize the missing required field and (ii) restate the schema-conformant fix, the reflection step converts a low-information error into actionable supervision. In our benchmark, such instances consistently improve:

1. **Schema adherence:** higher exact-match rate on name/arguments.
2. **Planning under constraints:** selection of parameterized calls (`duration=10`) aligned with user constraints instead of brittle pre-checks with empty arguments.
3. **Stability:** fewer retries and warnings downstream because calls are executable on the first attempt.

Summary. This case illustrates how reflection-guided repair turns a malformed `<call>` into a compact, correct, and time-efficient action plan.

A.7 Test Data Case Study

In this section, we present two representative test cases and their corresponding evaluation results as a case study, providing an intuitive demonstration of the effectiveness of our method and the model’s self-reflection capability for tool-call repair. Since the original cases are relatively long, we include their full content in the supplementary material for reference and provide only the analysis here.

A.7.1 Case I

Setting. The tool set exposes three functions: `getRecipes(max_time, meal_type)`, `getSmoothieIngredients(max_time)`, and `findComplementaryRecipes(recipes, ingredients)`. The user asks for *breakfast* recipes under 15 minutes and smoothie pairings under 5 minutes.

Failure mode (pre-training). The baseline model immediately issues

```
[{"name": "findComplementaryRecipes", "parameters": {}}]
```

which violates the function schema (both `recipes` and `ingredients` are required). The tool returns a schema-warning. Under our reward, this yields $s_{\text{call}} = 0$ and triggers format penalties $F < 1$ due to missing required parameters.

Reflection-driven repair (post-training). After RL on Tool-Reflection-Bench, the model first *reflects* that the failure arises from absent inputs, then correctly decomposes the task into *produce inputs* \rightarrow *compose*:

```
[{"name": "getRecipes", "parameters": {"max_time": 15, "meal_type": "breakfast"}}]
[{"name": "getSmoothieIngredients", "parameters": {"max_time": 5}}]
[{"name": "findComplementaryRecipes", "parameters": {"recipes": "<from getRecipes>", "ingredients": "<from getSmoothieIngredients>"}}]
```

This satisfies the schema strictly (tool names and parameter maps match), making the call set correct and executable.

Why our method helps. (i) **Reward shaping:** The instance accrues a hard penalty when required fields are absent; after repair, s_{call} flips to 1 and $F \rightarrow 1$, raising $R_{\text{core}} = S \cdot F$ substantially. (ii) **Sequence-level RL:** The GSPO-style sequence-ratio with dual clipping aligns the optimization granularity with the sequence reward, while DAPO-style dynamic filtering removes near-zero-advantage rollouts (all-wrong/all-correct), sharpening learning signals for this failure mode.

Takeaway. Compared to the baseline that *jumps* to composition with empty inputs, the trained policy learns to (a) diagnose the schema error, (b) *stage* prerequisite calls to produce the missing inputs, and (c) complete the composition with a schema-valid call set. This precisely matches our benchmark’s objective: enable robust, multi-turn tool use via reflection and repair.

A.7.2 Case II

Setting. Available tools include `get_current_season()`, `get_seeds_by_season(season)`, `filter_seeds_by_availability(seeds)`, `purchase_seeds(seeds, quantity)`, and `calculate_total_cost(purchased_seeds)`. The user asks to *buy 10 packets of seasonal vegetable seeds* and *report the total cost*.

Failure mode (pre-training). The baseline calls the aggregator first, with no inputs:

```
[{"name": "calculate_total_cost", "parameters": {}}]
```

This violates the required schema (`purchased_seeds` missing), producing a warning and yielding $s_{\text{call}} = 0$ and a strong format penalty $F < 1$ in our reward.

Reflection-driven repair (post-training). After RL on Tool-Reflection-Bench, the model first *reflects* that costing requires purchased items, then executes a staged pipeline to materialize prerequisites before aggregation:

```
[{"name": "get_current_season", "parameters": {}}]
```

```
[{"name": "get_seeds_by_season", "parameters": {"season": "<CUR_SEASON>"}}]
```

```
[{"name": "filter_seeds_by_availability", "parameters": {"seeds": "<SEASONAL_SEEDS>"}}]
```

```
[{"name": "purchase_seeds", "parameters": {"seeds": "<AVAILABLE_SEEDS>", "quantity": 10}}]
```

```
[{"name": "calculate_total_cost", "parameters": {"purchased_seeds": "<PURCHASED>"}}]
```

Each call now matches tool name and parameter map exactly (schema-strict), so $s_{\text{call}} = 1$ and $F \rightarrow 1$.

Why it works. *Reward design* penalizes missing required fields and redundant structure, while granting full credit only when the `<call>` set exactly matches the ground truth (schema-strict, order-invariant). The *sequence-level RL objective* (GSPO-style ratio, dual clipping) aligns optimization with

sequence rewards, and *DAPO-style dynamic filtering* removes near-zero-advantage groups, concentrating updates on informative failures. Together these guide the policy to diagnose schema errors, stage prerequisite calls, and complete the costing correctly.

Takeaway. The trained policy no longer “guesses” totals from empty inputs. Instead, it *plans* \rightarrow *acquires data* \rightarrow *purchases* \rightarrow *aggregates*, a behavior precisely targeted by our reflection-and-repair rewards.

A.8 Theoretical Analysis

We analyze the main design choices of our reward in Sec. §2.2 and the RL objective in Sec. §2.3. Throughout, $\text{Sim} \in [0, 1]$, all weights are non-negative, presence masks are indicators, and $\text{clip}(x, a, b) = \min\{b, \max\{a, x\}\}$. To avoid symbol overloading, we denote by r_{fmt} the format-penalty attenuation scalar used in Sec. §2.2 (called r there), and by r_{seq} the sequence-level importance ratio in Sec. §2.3.

A.8.1 Consistency of Presence-Mask Normalization

Recall

$$W_{\text{act}} := w_r I_r + w_c I_c + w_f I_f, \\ S := \frac{w_r I_r s_{\text{ref}} + w_c I_c s_{\text{call}} + w_f I_f s_{\text{final}}}{W_{\text{act}}}. \quad (30)$$

where $w_{\bullet} \geq 0$, $I_{\bullet} \in \{0, 1\}$, at least one $I_{\bullet} = 1$, $s_{\text{ref}}, s_{\text{final}} \in [0, 1]$, and $s_{\text{call}} \in \{0, 1\}$.

Lemma 1 (Convex-combination form). Let $\mathcal{A} = \{k \in \{r, c, f\} : I_k = 1\}$ and define

$$\alpha_k = \frac{w_k}{\sum_{j \in \mathcal{A}} w_j} \quad \text{for } k \in \mathcal{A}. \quad (31)$$

Then $\alpha_k \geq 0$, $\sum_{k \in \mathcal{A}} \alpha_k = 1$, and

$$S = \sum_{k \in \mathcal{A}} \alpha_k s_k, \quad (32)$$

where $s_r = s_{\text{ref}}$, $s_c = s_{\text{call}}$, $s_f = s_{\text{final}}$.

Proof. Since $I_k = 1$ iff $k \in \mathcal{A}$, the numerator equals $\sum_{k \in \mathcal{A}} w_k s_k$ and $W_{\text{act}} = \sum_{k \in \mathcal{A}} w_k > 0$. Divide both to obtain the stated form.

Proposition 1 (Boundedness, stability, and scale invariance). With $W_{\text{act}} > 0$:

(a) $S \in [0, 1]$ and, more sharply, $S \in [\min_{k \in \mathcal{A}} s_k, \max_{k \in \mathcal{A}} s_k]$.

- (b) If one only toggles *absent* parts (keeps \mathcal{A} and $\{w_k\}_{k \in \mathcal{A}}$ unchanged), then S is unchanged.
- (c) For any $\lambda > 0$, replacing each active weight by λw_k leaves S unchanged.

Proof. (a) By Lemma 1, S is a convex combination of $\{s_k\}_{k \in \mathcal{A}}$; the interval bound follows from $s_k \in [0, 1]$. (b) Absent-part toggles do not change \mathcal{A} nor the active w_k . (c) Common scaling cancels in numerator/denominator.

Corollary 1 (Continuity and Lipschitzness). Fix \mathcal{A} and w_k for $k \in \mathcal{A}$. Then S is an affine (hence continuous) map of $(s_k)_{k \in \mathcal{A}}$ with

$$|S - S'| \leq \sum_{k \in \mathcal{A}} \alpha_k |s_k - s'_k| \leq \max_{k \in \mathcal{A}} |s_k - s'_k|, \quad (33)$$

so S is 1-Lipschitz w.r.t. the ℓ_∞ -norm on the active scores.

Remark. The definition via $\text{clip}_{[0,1]}(\cdot)$ in (38) is not needed for S since the convex-combination form already implies $S \in [0, 1]$.

A.8.2 Format Factor: Boundedness, Monotonicity, and EqualCalls Attenuation

Let

$$P_{\text{total}} = P_{\text{miss}} + \beta_{\text{extra}} P_{\text{extra}} + \gamma_{\text{count}} P_{\text{count}}, \quad (34)$$

where $\beta_{\text{extra}}, \gamma_{\text{count}} \geq 0$ and $P_\bullet \geq 0$.

and define the attenuation scalar

Let $E := \text{EqualCalls}(C_{\text{calls}}, G_{\text{calls}})$.

$$r_{\text{fmt}} = \begin{cases} r_{\text{reduce}}, & E, \\ 1, & \text{otherwise,} \end{cases} \quad r_{\text{reduce}} \in (0, 1]. \quad (35)$$

Consider

$$F = \text{clip}_{[0,1]}(1 - \lambda_m P_{\text{total}} r_{\text{fmt}}), \quad \lambda_m \geq 0. \quad (36)$$

This is equivalent to the piecewise definition in (23) since $P_{\text{miss}}=P_{\text{extra}}=P_{\text{count}}=0$ implies the inner value equals 1.

Proposition 2 (Core properties of F).

- (a) *Boundedness and regularity.* $F \in [0, 1]$ for all inputs; F is continuous, piecewise affine in $(P_{\text{miss}}, P_{\text{extra}}, P_{\text{count}})$ and 1-Lipschitz w.r.t. its scalar argument before clipping.

- (b) *Monotonicity.* For fixed $(\lambda_m, r_{\text{fmt}})$, F is non-increasing in $P_{\text{miss}}, P_{\text{extra}}, P_{\text{count}}$ and non-increasing in λ_m and in r_{fmt} .

- (c) *EqualCalls attenuation improves F .* If EqualCalls holds so that r_{fmt} is replaced by $r_{\text{reduce}} \leq 1$, then F weakly increases.

- (d) *Plateau characterization.* $F = 1$ iff $\lambda_m P_{\text{total}} r_{\text{fmt}} = 0$ (e.g., $P_{\text{total}} = 0$ or $\lambda_m = 0$). If $\lambda_m > 0$ and $r_{\text{fmt}} > 0$, then $F = 0$ iff $P_{\text{total}} \geq 1/(\lambda_m r_{\text{fmt}})$.

Corollary 2 (Sensitivity bound). Off the plateaus ($1 - \lambda_m P_{\text{total}} r_{\text{fmt}} \in (0, 1)$),

$$|\Delta F| \leq \lambda_m r_{\text{fmt}} \left(|\Delta P_{\text{miss}}| + \beta_{\text{extra}} |\Delta P_{\text{extra}}| + \gamma_{\text{count}} |\Delta P_{\text{count}}| \right). \quad (37)$$

A.8.3 Core Reward with Similarity Backoff: Signal and Variance Control

Let $R_{\text{core}} = S \cdot F$ as in (24). To stabilize learning when R_{core} is very small, we introduce a similarity backoff. Define $[x]_0^1 := \min\{1, \max\{0, x\}\}$ and $S_b := \text{Sim}(\text{concat}(C), \text{concat}(G))$:

$$R_{\text{total}} = \begin{cases} [R_{\text{core}}]_0^1, & R_{\text{core}} \geq \varepsilon, \\ [w_b S_b]_0^1, & \text{otherwise.} \end{cases} \quad (38)$$

Here $w_b \in (0, 1]$ and $\varepsilon > 0$. Note that $R_{\text{core}} \in [0, 1]$ already, so clipping is redundant but harmless, and it keeps the two branches notationally symmetric.

We analyze its effect under a standard policy-gradient estimator $\nabla_\theta \mathbb{E}[R_{\text{total}}] = \mathbb{E}[R_{\text{total}} \nabla_\theta \log \pi_\theta(\cdot)]$.

Lemma 3 (Uniform bounded variance of the reward). Since $R_{\text{total}} \in [0, 1]$, we have $\text{Var}(R_{\text{total}}) \leq \frac{1}{4}$ for any data distribution.

Lemma 4 (Non-degenerate gradient second moment on the backoff branch). Let $\mathcal{B} := \{R_{\text{core}} < \varepsilon\}$ with $\mathbb{P}(\mathcal{B}) = p > 0$. Let $g_\theta := \nabla_\theta \log \pi_\theta(\cdot)$. Assume $S_b \geq \sigma$ a.s. on \mathcal{B} for some $\sigma > 0$ and $\mathbb{E}[\|g_\theta\|^2 \mathbf{1}_{\mathcal{B}}] > 0$. Then

$$\mathbb{E}[\|R_{\text{total}} g_\theta\|^2] \geq (w_b \sigma)^2 \mathbb{E}[\|g_\theta\|^2 \mathbf{1}_{\mathcal{B}}] > 0. \quad (39)$$

Implication. When R_{core} frequently approaches 0 (early in training), the backoff branch prevents the gradient second moment from degenerating; together with the variance upper bound in Lemma 3, this stabilizes the optimization updates.

A.8.4 Sequence-Level Importance Sampling and Clipping

Let the sampled completion be $o = (o_1, \dots, o_T)$, and define the sequence-level (geometric-mean, length-normalized) ratio:

$$\begin{aligned} r_{\text{seq}}(\theta) &= \left(\prod_{t=1}^T \frac{\pi_{\theta}(o_t | q, o_{<t})}{\pi_{\theta_{\text{old}}}(o_t | q, o_{<t})} \right)^{1/T} \\ &= \exp \left(\frac{1}{T} \sum_{t=1}^T \log \rho_t \right), \\ \rho_t &:= \frac{\pi_{\theta}(o_t | q, o_{<t})}{\pi_{\theta_{\text{old}}}(o_t | q, o_{<t})}. \end{aligned} \quad (40)$$

Proposition 3 (Length-independent ratio range under bounded log-ratios). If $\log \rho_t \in [-L, L]$ a.s. for some $L > 0$, then

$$e^{-L} \leq r_{\text{seq}}(\theta) \leq e^L \quad \text{for all } T \geq 1, \quad (41)$$

whereas the unnormalized product ratio ranges in $[e^{-LT}, e^{LT}]$.

Implication. The geometric mean aligns the ratio granularity with the sequence-level reward in (26), prevents exponential blow-up with T , and—together with dual clipping—reduces variance at the sequence level.

A.8.5 Dynamic Filtering of Prompt Groups (DAPO-style)

Let a prompt group produce G rollouts $\{o_i\}_{i=1}^G$ with rewards $R_i \in [0, 1]$ and batch z -scored advantages

$$\begin{aligned} \hat{A}_i &:= \frac{R_i - \bar{R}}{s_R}, \\ \bar{R} &:= \frac{1}{G} \sum_{j=1}^G R_j, \\ s_R &:= \sqrt{\frac{1}{G} \sum_{j=1}^G (R_j - \bar{R})^2} > 0. \end{aligned} \quad (42)$$

Define the *accepted* set

$$\begin{aligned} \mathcal{S} &:= \{i : |\hat{A}_i| > \tau_{\text{adv}}\}, \\ 0 &< |\mathcal{S}| < G, \\ \text{Var}(\{R_i\}_{i=1}^G) &> \tau_{\text{var}} > 0. \end{aligned} \quad (43)$$

Write the per-sample (sequence-level, dual-clipped) PPO-like term as

$$\begin{aligned} \ell_i(\theta) &:= \min(r_{\text{seq},i}(\theta), \bar{r}_{\text{seq},i}(\theta)) \hat{A}_i, \\ \bar{r}_{\text{seq},i}(\theta) &:= \text{clip}(r_{\text{seq},i}(\theta), 1 - \varepsilon_{\text{low}}, 1 + \varepsilon_{\text{high}}). \end{aligned} \quad (44)$$

and denote its gradient by $g_i(\theta) = \nabla_{\theta} \ell_i(\theta)$. Assume the usual score-function bound and clipped ratio range:

$$\begin{aligned} \|\nabla_{\theta} \log \pi_{\theta}(o_{i,t} | q, o_{i,<t})\| &\leq B_{\pi}, \\ r_{\text{seq},i}(\theta) &\in [1 - \varepsilon_{\text{low}}, 1 + \varepsilon_{\text{high}}]. \end{aligned} \quad (\star)$$

A uniform bound on per-rollout gradients.

Since $r_{\text{seq},i}(\theta)$ is the geometric mean of token ratios, let $T_i := |o_i|$ and $h_{i,t} := (q, o_{i,<t})$. Then

$$\begin{aligned} \nabla_{\theta} r_{\text{seq},i}(\theta) &= r_{\text{seq},i}(\theta) \frac{1}{T_i} \sum_{t=1}^{T_i} \nabla_{\theta} \log \pi_{\theta}(o_{i,t} | h_{i,t}) \\ &= \frac{r_{\text{seq},i}(\theta)}{T_i} \sum_{t=1}^{T_i} \nabla_{\theta} \log \pi_{\theta}(o_{i,t} | h_{i,t}). \end{aligned} \quad (45)$$

Using (\star) and that the clipped branch is constant on plateaus, there exists a finite $C_{\psi} = (1 + \varepsilon_{\text{high}}) B_{\pi}$ such that

$$\|g_i(\theta)\| \leq C_{\psi} |\hat{A}_i| \quad \text{for all } i, \theta. \quad (46)$$

Lemma 5 (Zero or near-zero advantages).

- (a) If $\hat{A}_i = 0$, removing o_i leaves the group-wise expected gradient unchanged.
- (b) If $|\hat{A}_i| \leq \tau_{\text{adv}}$, then for any θ ,

$$\begin{aligned} \|\mathbb{E}[g_i(\theta)]\| &\leq C_{\psi} \tau_{\text{adv}}, \\ \mathbb{E}[\|g_i(\theta)\|^2] &\leq C_{\psi}^2 \tau_{\text{adv}}^2. \end{aligned} \quad (47)$$

Proof. (a) The contribution is proportional to \hat{A}_i . (b) Apply (46) and take expectations.

Bias and variance effects with $\frac{1}{G}$ normalization.

Let the *filtered* group gradient be

$$\begin{aligned} \tilde{g}(\theta) &:= \frac{1}{G} \sum_{i \in \mathcal{S}} g_i(\theta), \\ g(\theta) &:= \frac{1}{G} \sum_{i=1}^G g_i(\theta) \quad (\text{unfiltered}). \end{aligned} \quad (48)$$

Define the discarded set $\mathcal{S}^c = \{1, \dots, G\} \setminus \mathcal{S}$. Then

$$\mathbb{E}[\tilde{g}(\theta)] - \mathbb{E}[g(\theta)] = -\frac{1}{G} \sum_{i \in \mathcal{S}^c} \mathbb{E}[g_i(\theta)], \quad (49)$$

$$\|\mathbb{E}[\tilde{g}(\theta)] - \mathbb{E}[g(\theta)]\| \leq \frac{|\mathcal{S}^c|}{G} C_{\psi} \tau_{\text{adv}} \leq C_{\psi} \tau_{\text{adv}},$$

using Lemma 5(b). Moreover,

$$\mathbb{E} \left[\left\| \frac{1}{G} \sum_{i \in \mathcal{S}^c} g_i(\theta) \right\|^2 \right] \leq \frac{|\mathcal{S}^c|}{G^2} C_\psi^2 \tau_{\text{adv}}^2, \quad (50)$$

thus, discarding near-zero advantageous terms induces at most an $O(\tau_{\text{adv}}^2)$ -level change in the second moment; with respect to the $\frac{1}{G}$ normalization, it does not introduce any additional scaling bias.

Acceptance constraints avoid degeneracy. The constraints $0 < |\mathcal{S}| < G$ and $\text{Var}(\{R_i\}) > \tau_{\text{var}}$ ensure: (i) the batch standardization s_R is well-defined; (ii) both positive and negative (or at least non-identical) signals are present, preventing the trivial zero-gradient case where all \hat{A}_i are identical. Consequently, $\tilde{g}(\theta)$ is a non-degenerate direction whenever useful learning signal exists.

Asymptotic unbiasedness with vanishing threshold. If the threshold decays $\tau_{\text{adv}}^{(t)} \downarrow 0$ and the law of \hat{A}_i has a continuous density at 0, then the discard probability $\mathbb{P}(|\hat{A}_i| \leq \tau_{\text{adv}}^{(t)}) \rightarrow 0$, and

$$\lim_{t \rightarrow \infty} \left\| \mathbb{E}[\tilde{g}_t(\theta)] - \mathbb{E}[g(\theta)] \right\| = 0, \quad (51)$$

i.e., the dynamic filtering becomes asymptotically unbiased while retaining finite-time variance-reduction benefits.

Summary. Dynamic filtering deletes rollouts whose contributions are provably negligible (zero or $O(\tau_{\text{adv}})$), thereby reducing variance and compute without altering the expected update in the limit $\tau_{\text{adv}} \rightarrow 0$; using the same $1/G$ normalization as (26) avoids spurious scaling bias.

A.8.6 Convergence Considerations for the Clipped Sequence-Level Objective

Consider the surrogate objective $\mathcal{J}_{\text{RL}}(\theta)$ in (26), where rewards are bounded in $[0, 1]$ and the sequence-level importance ratios are dual-clipped to $[1 - \varepsilon_{\text{low}}, 1 + \varepsilon_{\text{high}}]$.

Assumptions.

- (A1) **Bounded scores.** There exists $B_\pi < \infty$ such that for all histories $(q, o_{<t})$ and tokens o_t , $\|\nabla_\theta \log \pi_\theta(o_t | q, o_{<t})\| \leq B_\pi$.
- (A2) **Bounded rewards & finite clipping.** For each rollout o_i , $R_i \in [0, 1]$ and $r_{\text{seq},i}(\theta) \in [1 - \varepsilon_{\text{low}}, 1 + \varepsilon_{\text{high}}]$ with $0 < \varepsilon_{\text{low}}, \varepsilon_{\text{high}} < \infty$.

(A3) **Non-degenerate batch dispersion.** On accepted groups, $\text{Var}(\{R_i\}_{i=1}^G) \geq \tau_{\text{var}} > 0$, so $\hat{A}_i = (R_i - \bar{R})/\text{std}(R)$ are well-defined.

(A4) **Vanishing filtering.** $\tau_{\text{adv}}^{(t)} \downarrow 0$ and the law of \hat{A}_i has a continuous density at 0, so $\mathbb{P}(|\hat{A}_i| \leq \tau_{\text{adv}}^{(t)}) \rightarrow 0$.

(A5) **Sizesizes.** Robbins–Monro conditions: $\sum_t \eta_t = \infty$ and $\sum_t \eta_t^2 < \infty$.

Lemma 6 (Bounds on per-sample gradients and second moments). Let $o = (o_1, \dots, o_T)$ with $T := |o|$, and let $h_t := (q, o_{<t})$. Let $r_{\text{seq}}(\theta)$ denote the (clipped) sequence ratio. Then

$$\begin{aligned} \nabla_\theta r_{\text{seq}}(\theta) &= \frac{r_{\text{seq}}(\theta)}{T} \sum_{t=1}^T \nabla_\theta \log \pi_\theta(o_t | h_t), \\ \|\nabla_\theta r_{\text{seq}}(\theta)\| &\leq (1 + \varepsilon_{\text{high}}) B_\pi. \end{aligned} \quad (52)$$

Moreover, the PPO-style term is piecewise smooth and its gradient magnitude is bounded by $C_1 := (1 + \varepsilon_{\text{high}}) B_\pi |\hat{A}|$; together with (A3), $|\hat{A}| \leq \frac{1}{\sqrt{\tau_{\text{var}}}}$ yields a uniform second-moment bound $\mathbb{E}[\|\nabla_\theta \ell_i(\theta)\|^2] \leq C_2 < \infty$.

Lemma 7 (Asymptotic unbiasedness under vanishing filtering). Let $g(\theta)$ denote the full (unfiltered) stochastic gradient and $\tilde{g}_\tau(\theta) = \frac{1}{G} \sum_{i: |\hat{A}_i| > \tau} g_i(\theta)$ the filtered version with $\frac{1}{G}$ normalization. Under (A4) and the bounded second moments above,

$$\lim_{\tau \downarrow 0} \left\| \mathbb{E}[\tilde{g}_\tau(\theta)] - \mathbb{E}[g(\theta)] \right\| = 0 \quad \text{for all } \theta. \quad (53)$$

Theorem 1 (Convergence to a stationary point of the surrogate). Suppose (A1)–(A5) hold. Then the iterates of stochastic gradient ascent on $\mathcal{J}_{\text{RL}}(\theta)$ with the dynamic filtering scheme converge almost surely to the set of stationary points of the surrogate objective.

Proof sketch. By Lemma 6 and the reward boundedness (Lemma 3), the stochastic gradients have uniformly bounded second moments; the objective is bounded and piecewise smooth (kinks of measure zero). Lemma 7 guarantees that the bias due to filtering vanishes as $\tau_{\text{adv}}^{(t)} \rightarrow 0$. Therefore the noisy gradient process forms a Robbins–Monro stochastic approximation with asymptotically unbiased gradients and square-summable noise, yielding a.s. convergence to stationary points of \mathcal{J}_{RL} (e.g., Kushner–Yin/Bottou).

Remarks. (i) The min-with-clipping introduces bias w.r.t. the *true* off-policy objective, but ensures variance control and stability; the theorem concerns the surrogate we optimize. (ii) Sequence-level ratios and sequence-level clipping align the gradient scale with the sequence reward, avoiding token/sequence granularity mismatch and contributing to the boundedness needed above. (iii) In practice, we keep τ_{var} and the clip window fixed and decay τ_{adv} , which satisfies the lemmas' conditions and matches our training protocol.