

Speculate While You Reason: Teaching Agents to Predict Their Next Tool Call via Joint Agent–Speculator RL

Jiabao Ji^{1*} Yujian Liu^{1*} Li An^{1*}
 Rohit Jain² Gungor Polatkan² Siyu Zhu² Shiyu Chang¹

¹University of California, Santa Barbara ²LinkedIn Inc
 {jiabaoji,yujianliu,li_an,chang87}@ucsb.edu

Abstract

Large language model agents often spend substantial wall-clock time waiting for tool call results. Tool-call speculation can hide this latency by predicting and pre-executing an agent’s next tool call if the prediction matches the agent’s eventual tool call, but existing speculators are typically separate draft models or cached traces that are poorly aligned with the deployed agent’s own behavior. We identify this *speculator–agent gap* and show that the target agent itself is a strong next-call speculator. This points to a simpler design: unifying the agent and speculator within the same model. In this paper, we introduce the *self-speculating agent*, a single model that both solves tasks in agent mode and predicts its next tool call from partial trajectories in speculator mode, fully reusing the prefix KV cache. To enable this dual-mode agent without degrading performance, we propose a joint agent-speculator reinforcement learning method, which derives speculation targets from the agent’s own rollouts and alternates agent and speculator updates. Across agentic search QA and conversational tool-use tasks, our method improves average next tool-call Hit@1 from 44.1 to 61.2 for Qwen3-4B and from 48.9 to 66.3 for Qwen3.5-4B, while preserving agent task success.

1 Introduction

Large language model (LLM) agents solve complex tasks by interleaving natural-language reasoning with calls to external tools, including web search engines, databases, code interpreters, and task-specific APIs (Schick et al., 2023; Patil et al., 2023; Chen et al., 2023; Liu et al., 2024). Recent work further trains and evaluates such agents in search, web, function-calling, and long-horizon interaction settings (Jin et al., 2025; Song et al.,




Figure 1: External versus self tool-call speculation. After the previous tool observation, both branches share h_t and run in parallel: an external draft model with separate parameters and KV cache (top), or self-speculation with a shared prefix KV cache (bottom).

2025; Tan et al., 2025; Qi et al., 2025; Qian et al., 2025; Zhang et al., 2025; Liu et al., 2026; Luo et al., 2025). These tools extend agents beyond their parametric knowledge by allowing them to retrieve up-to-date information, execute computations, inspect structured records, and interact with external environments. This flexibility, however, comes at the cost of latency. Unlike token generation, tool calls often involve remote services, network I/O, or even invoking another LLM as a sub-agent. As a result, agents may spend a substantial fraction of their inference time waiting for tool results rather than performing token generation (Nichols et al., 2025; Huang et al., 2025; Hooper et al., 2026), so reducing this wait is important for responsive agents.

A promising approach is *tool-call speculation*: issue likely future tool calls before the agent explicitly produces them. After the previous tool observation (or, on the first turn, the user query), a speculator predicts the agent’s next structured tool call and executes that call in parallel while the agent continues reasoning. If the prediction matches the agent’s eventual tool call exactly, the pre-executed result can be reused and the tool la-

*Equal contribution. Work done when Jiabao, Yujian, and Li were interning at LinkedIn Inc.

tency is hidden; otherwise it is discarded. The speculator must therefore match the call this agent will actually issue, not any plausible call under the same context.

This requirement makes the design of the speculator especially important. Prior work typically builds the speculator outside the deployed agent, either by using a smaller LLM as an external draft model (Ye et al., 2026; Nichols et al., 2025; Huang et al., 2025) or by consulting cached tool call traces derived from previous agent trajectories (Zhong et al., 2026; Sui et al., 2026). These methods can reduce latency, but they share a fundamental limitation: the speculator is not the agent. An off-the-shelf draft model approximates a generic tool-using assistant rather than the particular agent being deployed. A cache-based method replays calls from previous trajectories, which may correspond to different user queries, different intermediate reasoning states, or outdated environments. In both cases, the speculation methods remain separated from the actual deployed agent.

We refer to this separation as the *speculator-agent gap*: conditioned on the same intermediate agent trajectory, the speculator and the deployed agent may choose different next tool calls. Figure 1 illustrates this gap, where an external draft model maintains a separate set of model parameters and KV-cache and predicts a different call than the target agent eventually issues. In Section 2, we find that the gap appears even within the same model family: smaller Qwen draft models often fail to match the 4B target agent’s next tool call, while the 4B agent itself is already a stronger off-the-shelf speculator when prompted to predict its own next call. External speculators also introduce practical systems overhead, including additional model weights, a separate KV cache, or historical traces.

These observations suggest a simpler speculator design: rather than aligning a separate speculator to the agent, the agent itself can speculate. This leads to the central research question: *Can we train an agent to better anticipate its own next tool call while preserving task solving performance?*

We therefore present a *self-speculating agent*: a single model trained to operate in two modes. In *agent mode*, the model follows the standard agentic trajectory, producing reasoning, tool calls, observations, and final answers. In *speculator mode*, the same model receives an intermediate trajectory together with a short speculation suffix and

directly predicts the next structured tool call. Because both modes share the trajectory prefix, inference can reuse the agent’s prefix KV cache and branch only for the speculation suffix and candidate call. Unlike external speculation, this approach does not require serving an auxiliary draft model or maintaining historical tool-call traces.

However, training such an agent is nontrivial. Self-speculation couples two objectives that are naturally imbalanced: the agent must remain effective at solving tasks, while also becoming accurate at predicting its own next tool call from partial trajectories. The two modes differ substantially in output distribution, and naively training the agent to improve speculation can hurt agent performance. To address this problem, we introduce *joint agent-speculator RL*. During the RL process, we sample agent rollouts from the current policy, construct speculator queries from those rollouts, and optimize each mode in an alternating procedure. This keeps the speculator mode on-policy: it is trained on the same contexts that arise from the current agent, rather than on stale rollouts or data from a different model. We further stabilize this joint training with several training-time techniques, such as optimizer resets at mode switches.

Experiments across agentic search QA and conversational tool-use tasks show that joint agent-speculator RL substantially improves speculation performance while preserving task performance. For Qwen3-4B and Qwen3.5-4B, average next tool-call Hit@1 score improves from 44.1 to 61.2 and from 48.9 to 66.3, respectively, while average downstream task success remains stable or slightly improves.

Our contributions are as follows:

- We identify the speculator-agent gap in agentic tool-use tasks where off-the-shelf draft models often fail to match the agent’s next structured call and add serving overhead.
- We introduce a self-speculating agent, where one LLM switches between task-solving and next-call prediction modes.
- We propose joint agent-speculator RL, which trains speculation from the current agent’s own rollouts while preserving task success.

2 Evaluating Off-the-Shelf Models for Tool-Call Speculation

Off-the-shelf draft models are a natural starting point for tool-call speculation. We evaluate their speculator-agent gap in realistic agentic settings, where the speculator observes the shared prefix h_t after the previous tool observation, or the user query on the first turn, and predicts the next tool call, including both the tool name and arguments. This setup measures whether an off-the-shelf model can match the deployed agent’s tool decisions under the same inference-time context.

2.1 Evaluation Protocol

For each target agent, we run speculation online during the agent’s rollout. In each domain, we sample 50 evaluation queries and execute the deployed agent in the environment. After the previous tool observation, or after the user query on the first turn, the serving controller switches from the target agent model to an off-the-shelf speculator model. The speculator receives the current trajectory prefix with a fixed speculation suffix appended:

```
<think> Okay, let's see. The user
provided what I need. I'll look it up.
The next step is to make the tool call.
</think>
```

The speculator then decodes one candidate tool call. The controller switches back to the target agent, lets the agent produce the actual call, and executes the agent’s call in the environment. Thus speculation is evaluated under the same online state that the agent sees, while the environment state is still determined only by the deployed agent. Figure 2 illustrates this online comparison in a τ -bench trajectory: the speculator is prompted from the current prefix, the agent continues independently, and the pre-executed result is reusable only when the speculated tool name and arguments exactly match the agent’s eventual call.

We instantiate this protocol in two same-family comparisons. For the Qwen3-4B agent, we evaluate Qwen3-0.6B and Qwen3-1.7B as speculators; for the Qwen3.5-4B agent, we evaluate Qwen3.5-0.8B and Qwen3.5-2B. In both families, we also prompt the target 4B model to predict its own next call from the same online prefix. We run the protocol on 50 queries from MuSiQue in the agentic SearchQA environment (Jin et al., 2025) and the airline domain of τ -bench (Yao et al., 2024; Barres et al., 2025).




Figure 2: Online tool-call speculation example. From the shared trajectory prefix after the previous tool observation, the left stream continues to the agent’s reasoning and actual tool call, while the right stream predicts the tool call in parallel.

Table 1 reports both prediction quality and serving cost. We use Hit@1 exact match as the primary speculation metric because a pre-executed tool call is reusable only when both the tool name and argument dictionary match the agent’s eventual call. We also report speculation wall-clock time and peak GPU memory to capture the overhead of alternating between the target agent and a separate draft model during the same rollout. The memory estimate includes both resident model weights and peak KV-cache state: external speculators require the target agent weights, the draft model weights, and two active KV caches, whereas self-speculation reuses the target model weights and the target prefix KV cache, adding only the small speculation branch. All measurements in this off-the-shelf setting use a single H100 GPU. We implement the comparison in the SGLang inference engine with an explicit parameter switch between the external drafter model and the target agent model, so each query serially alternates agent generation, speculator generation, and agent continuation while recording time and memory usage.

2.2 Findings

External speculators add serving overhead. Even small draft models require additional weights and a second KV cache, so their cost is not just the extra decoding step. In Table 1, Qwen3-4B on MuSiQue increases from 8.70GB and 8.3s with self-speculation to 10.71GB/27.4s with Qwen3-0.6B and 12.76GB/33.1s with Qwen3-1.7B. The same pattern appears in τ -bench and in the Qwen3.5 series. External speculation therefore

Speculator	MuSiQue			τ -bench		
	H@1 $\uparrow$	Time $\downarrow$	Mem. $\downarrow$	H@1 $\uparrow$	Time $\downarrow$	Mem. $\downarrow$
Target agent: Qwen3-4B						
Qwen3-0.6B	4.3	27.4	10.71	16.8	45.2	14.59
Qwen3-1.7B	14.7	33.1	12.76	25.4	52.1	16.64
Qwen3-4B self	25.3	8.3	8.70	37.2	24.1	10.90
Target agent: Qwen3.5-4B						
Qwen3.5-0.8B	8.7	22.3	9.31	21.3	39.2	9.98
Qwen3.5-2B	16.2	30.6	11.55	32.8	54.1	12.22
Qwen3.5-4B self	27.4	6.7	7.73	44.1	19.1	8.22

Table 1: Online off-the-shelf speculation results. For each target 4B agent, we compare smaller same-family external draft speculators with the target model prompted to speculate its own next tool call. GPU memory includes model weights and peak KV-cache.

needs large Hit@1 gains to justify the extra serving cost, especially when memory is limited or agentic contexts are long.

The agent is the strongest off-the-shelf speculator for itself. The target agent outperforms smaller same-family draft models, indicating that family similarity alone does not recover the target agent’s next-call distribution. For Qwen3-4B, self-speculation reaches 25.3 Hit@1 on MuSiQue and 37.2 on τ -bench, compared with 14.7 and 25.4 for the strongest external speculator. For Qwen3.5-4B, self-speculation improves over Qwen3.5-2B from 16.2 to 27.4 on MuSiQue and from 32.8 to 44.1 on τ -bench. This supports the speculator-agent gap: the best off-the-shelf predictor of an agent’s next call is often the agent itself. These findings motivate training the agent itself as the speculator instead of deploying a separate external LLM.

3 Method

Motivated by the findings in Section 2.2, we train a single tool-calling LLM agent π_θ to act both as the task-solving agent and as its own next-call speculator. We first define the agent task reward and the speculation reward used to score predicted tool calls. Second, we describe the joint agent-speculator RL method, which turns current agent trajectories into on-policy speculation targets. Finally, we describe the training recipe used to stabilize the shared policy for both modes.

3.1 Task and Reward Formulation

Given a user query x , the agent interacts with an environment through a sequence of

reasoning steps, tool calls, and tool observations. We denote an agent trajectory by $\tau = (x, r_1, a_1, o_1, \dots, r_T, a_T, o_T, y)$, where r_t is the reasoning before the t -th tool call, a_t is the tool call, o_t is the returned observation, and y is the final answer. Each tool call is a structured action $a_t = (n_t, \sigma_t)$, with tool name n_t and argument dictionary σ_t . The agent reward $R_{\text{ag}}(\tau)$ is evaluated on the final answer y against the ground-truth answer y^* defined in the environment.

The speculation task is derived from the same trajectory and is implemented by the same policy π_θ . Agent and speculator modes therefore share model parameters, with the speculator mode triggered simply by appending a fixed speculation suffix to the input context. For each tool-call turn t , the shared state is the prefix immediately after the previous tool observation,

$$h_t = (x, r_1, a_1, o_1, \dots, r_{t-1}, a_{t-1}, o_{t-1}), \quad (1)$$

with $h_1 = (x)$ on the first turn. From this state, the agent generates the next reasoning r_t and actual call a_t . In parallel, the speculator appends s_{sp} and predicts $\hat{a}_t \sim \pi_\theta(\cdot \mid h_t \oplus s_{\text{sp}})$. The target is the agent’s own next call a_t from the sampled trajectory. Since a_t comes from the current policy, the targets change as the agent changes, keeping speculation training on-policy with the deployed agent.

We score a predicted call $\hat{a}_t = (\hat{n}_t, \hat{\sigma}_t)$ against the target call $a_t = (n_t, \sigma_t)$ using a name-conditioned argument reward. The tool-name score is $R_{\text{name}} = \mathbf{1}[\hat{n}_t = n_t]$. If the tool name matches, we compare argument dictionaries with macro token-F1 over target argument keys:

$$R_{\text{args}} = \frac{1}{|\text{keys}(\sigma_t)|} \sum_{k \in \text{keys}(\sigma_t)} \text{F1}(\hat{\sigma}_t[k], \sigma_t[k]).$$

Missing predicted arguments receive score zero, and extra predicted arguments do not receive credit. Argument values are normalized to strings before token-F1. For tools with no arguments, we set $R_{\text{args}} = 1$ after a correct tool-name match. The final speculation reward is

$$R_{\text{sp}} = R_{\text{name}} \cdot R_{\text{args}}. \quad (2)$$

This construction matches the exact-match reuse requirement for speculation: a wrong tool name cannot reuse the pre-executed result, while a correct tool name can still receive partial credit for

close argument values during training. Macro averaging gives each target argument key equal weight, avoiding pooled token-F1 behavior where long argument values dominate shorter but equally important keys. At serving time we do not use a confidence threshold: the pre-executed result is reused if and only if $\hat{a}_t = a_t$. Any mismatch discards it and executes the actual call, which adds one extra request and fallback delay.

3.2 Joint Agent–Speculator RL

Agent and speculator samples differ substantially. The agent mode generates long, reasoning-heavy trajectories, whereas the speculator mode generates a single structured tool call from a partial trajectory. Their different output distributions and rewards make direct joint training difficult. We therefore keep the data collection coupled but the policy updates separated. Coupled data collection means each iteration forms (h_t, a_t) pairs from fresh agent rollouts, so the speculator tracks the current policy rather than a fixed offline set. Separated updates mean $m(k)$ applies only one reward per optimizer step, with an optimizer reset at each mode switch so that momentum from one objective does not carry into the other.

At RL iteration k , we sample a query batch $Q \subset \mathcal{D}$ and generate G_{ag} agent trajectories for each $x \in Q$. These trajectories serve as standard agent rollouts: their final rewards R_{ag} are normalized within each query group to produce DAPO-style relative advantages for the agent objective (Yu et al., 2025). At the same time, they provide the speculation supervision induced by the current policy. For each realized tool-call turn, we take the prefix h_t and the actual next call a_t from the sampled trajectory, form the speculator input $h_t \oplus s_{\text{sp}}$, and treat a_t as the target call.

A fixed schedule $m(k) \in \{\text{ag}, \text{sp}\}$ selects which objective is optimized at iteration k . If $m(k) = \text{ag}$, we update π_θ with the agent rollouts and advantages from R_{ag} . If $m(k) = \text{sp}$, we update the same π_θ with the speculator rollouts and advantages from R_{sp} . Both modes use the same DAPO policy optimization (Yu et al., 2025); they differ in the input context, rollout structure, and reward function. Unless otherwise noted, $m(k)$ follows a repeated 4:8 schedule, selected on SearchQA in Section 4.4. Algorithm 1 summarizes the full procedure.

Algorithm 1 Joint Agent-Speculator RL

Require: Policy π_θ , query set $\mathcal{D}$, speculation suffix s_{sp} , alternation schedule $m(k)$, group sizes $G_{\text{ag}}, G_{\text{sp}}$

- 1: $m_{\text{prev}} \leftarrow \emptyset$
- 2: **for** RL iteration $k = 1, 2, \dots$ **do**
- 3: Sample query batch $Q \sim \mathcal{D}$
- 4: Generate agent groups $\mathcal{T} = \{\tau_i(x)\}_{x \in Q, i=1}^{G_{\text{ag}}}$
- 5: Build speculation pairs $\mathcal{S} = \{(h_t, a_t) : t \in \tau, \tau \in \mathcal{T}\}$
- 6: Generate speculator groups $\hat{\mathcal{A}} = \{\hat{a}_{t,j} \sim \pi_\theta(\cdot \mid h_t \oplus s_{\text{sp}})\}_{(h_t, a_t) \in \mathcal{S}, j=1}^{G_{\text{sp}}}$
- 7: **if** $m(k) \neq m_{\text{prev}}$ **then**
- 8: Reset optimizer state
- 9: **end if**
- 10: **if** $m(k) = \text{ag}$ **then**
- 11: DAPO update π_θ on grouped agent rewards $R_{\text{ag}}(\mathcal{T})$
- 12: **else**
- 13: DAPO update π_θ on grouped speculation rewards $\{R_{\text{sp}}(\hat{a}_{t,j}, a_t)\}_{(h_t, a_t) \in \mathcal{S}, j=1}^{G_{\text{sp}}}$
- 14: **end if**
- 15: $m_{\text{prev}} \leftarrow m(k)$
- 16: **end for**

3.3 Stabilizing Dual-Mode RL

Training a shared policy for both task solving and self-speculation is non-trivial. The two modes produce different output distributions and optimize different rewards, so direct joint training can suffer from mode interference or collapse into unproductive tool-calling behavior. We therefore use a training recipe with three components. ❶ Before RL, we run a short SFT warmup on a mixture of successful agent trajectories and their corresponding intermediate-trajectory speculation examples. The agent examples supervise the full successful trajectory, while each speculation example takes a partial trajectory, appends s_{sp} , and supervises the model to produce the ground-truth next tool call. This gives the model well-formed behavior in both modes before RL begins. ❷ We reset optimizer state whenever the schedule switches between agent and speculator updates. Since the two modes optimize different rewards, momentum and adaptive statistics differ between the two modes, which can strongly affect subsequent optimization steps (Sutskever et al., 2013; Kingma and Ba, 2015); the reset reduces carryover from one mode’s objective to the next. ❸ The agent reward includes a penalty for repeated meaningless tool calls, discouraging a common failure mode in which the agent continues calling tools without making task progress.

Model	HotpotQA		MuSiQue		BCP		τ -bench Airline		τ -bench Retail		Avg.	
	H@1 $\uparrow$	Succ. $\uparrow$	H@1 $\uparrow$	Succ. $\uparrow$	H@1 $\uparrow$	Succ. $\uparrow$	H@1 $\uparrow$	Succ. $\uparrow$	H@1 $\uparrow$	Succ. $\uparrow$	H@1 $\uparrow$	Succ. $\uparrow$
Qwen3 series												
Qwen3-0.6B	3.6	52.1	4.2	15.1	2.8	7.0	16.5	21.6	18.5	36.4	9.1	26.4
Qwen3-1.7B	13.7	52.1	15.8	15.1	11.9	7.0	24.1	21.6	27.3	36.4	18.6	26.4
Qwen3-4B	23.2	52.1	26.4	15.1	20.7	7.0	35.6	21.6	39.7	36.4	29.1	26.4
Qwen3-4B-SFT	37.8	51.3	40.9	16.4	35.6	7.4	51.2	22.0	55.1	35.8	44.1	26.6
Qwen3-4B-RL	55.2	53.0	58.4	17.1	52.1	8.1	68.0	23.1	72.3	37.0	61.2	27.7
Qwen3.5 series												
Qwen3.5-0.8B	8.4	63.1	9.5	25.1	7.8	29.6	20.2	56.4	22.4	69.7	13.7	48.8
Qwen3.5-2B	14.9	63.1	16.8	25.1	13.7	30.2	31.6	56.4	34.7	69.7	22.3	48.9
Qwen3.5-4B	25.4	63.1	28.3	25.1	23.6	30.7	42.1	56.4	46.8	69.7	33.2	49.0
Qwen3.5-4B-SFT	40.7	62.2	43.9	26.0	38.2	31.2	58.6	56.4	62.9	70.2	48.9	49.2
Qwen3.5-4B-RL	58.6	64.3	61.8	27.0	55.7	32.4	76.1	58.1	79.4	71.2	66.3	50.6

Table 2: Main self-speculation results. We compare smaller same-family speculators, the base 4B agent, the SFT-initialized 4B agent, and our dual-mode RL agent across search and τ -bench benchmarks. For external speculator rows, task success is the success of the fixed target 4B agent whose next calls are being predicted, not the smaller speculator run as an independent agent.

4 Experiments

Our experiments evaluate whether a single model can improve its own next-tool-call speculation while preserving downstream task performance. We first describe the evaluation setup, then report main results across search and conversational tool-use benchmarks, analyze cross-domain generalization, and ablate the stabilization techniques used in joint agent-speculator RL.

4.1 Experimental Setup

Tasks and environments. We evaluate on two families of structured tool-use tasks. Agentic SearchQA (Jin et al., 2025; Song et al., 2025) requires the model to answer multi-hop questions by issuing search calls and reasoning over returned evidence. τ -bench conversational tool-use tasks (Yao et al., 2024; Barres et al., 2025) require the agent to interact with a simulated user and complete airline or retail tasks through schema-constrained APIs. These settings stress complementary aspects of speculation: SearchQA emphasizes query formulation under partial information, whereas τ -bench emphasizes exact API selection and argument filling in long-horizon interactions.

Training and evaluation data. For Agentic SearchQA settings, we train on FlashQA and evaluate on HotpotQA, MuSiQue (Trivedi et al., 2022), and BrowseComp-Plus (BCP) (Chen et al., 2025b). For BCP, we follow Sun et al. (2025)

and use the 150-query evaluation split from their 680/150 train-test partition. We also follow the official BCP setup to use Qwen3-Embed-0.6B as the dense retriever. For conversational tool-use tasks, we train on ToolScale (Su et al., 2025) and evaluate on τ -bench Airline and τ -bench Retail (Barres et al., 2025).

Models and baselines. We instantiate our method with Qwen3-4B (Qwen Team, 2025) and Qwen3.5-4B (Qwen Team, 2026). We compare against smaller same-family LLMs as the external speculators, the base 4B agent prompted to predict its own next call, and the SFT-warmup checkpoint before RL. Specifically, the external speculators are Qwen3-0.6B and Qwen3-1.7B for Qwen3-4B, and Qwen3.5-0.8B and Qwen3.5-2B for Qwen3.5-4B. We employ the same speculation suffix for all models.

Training protocol. We follow the joint agent-speculator RL recipe in Section 3.3. Each run begins with a short SFT warmup on a mix of successful trajectories and speculation examples so that the model can emit well-formed tool calls under the speculation suffix before RL begins. We then instantiate the alternating RL between agent updates and speculator updates, resetting optimizer state when the schedule switches modes. Unless otherwise noted, we use DAPO-style grouped policy optimization (Yu et al., 2025) with $G_{\text{ag}} = 8$ agent rollouts, $G_{\text{sp}} = 8$ speculator rollouts, learn-

ing rate 3×10^{-6} , and a fixed 200-step budget per RL training. Appendix A.1 lists the remaining implementation details, such as the training hardware.

Evaluation protocol and metrics. For speculation evaluation, we first run the model in agent mode and collect complete trajectories. At each tool-call turn, we take the shared prefix h_t in Equation 1, after the previous observation and before the next reasoning step, append the speculation suffix, and ask the speculator to predict the next call. We report Hit@1 exact match, where a hit requires both the tool name and the full argument dictionary to match the agent’s eventual call. Serving uses the same exact-match rule, with no confidence threshold (Section 3.1). Wall-clock time under synthetic tool latency and a breakdown of misses are in Appendix D.1 and Appendix D.2. For task quality, we report task success on HotpotQA, MuSiQue, BCP, τ -bench Airline, and τ -bench Retail. For external speculators, task success refers to the fixed target 4B agent whose calls are being predicted, not to the smaller draft model as an independent agent. Together, these metrics capture the central trade-off: speculation is useful only if it improves reusable call prediction without degrading task behavior.

4.2 Main Results

Table 2 reports the main results across both model families and evaluation domains.

Joint agent-speculator RL improves speculation performance. Table 2 extends the finding from Section 2: the deployed agent is already a stronger speculator for itself than smaller same-family draft models. In the Qwen3 series, the strongest external speculator averages 18.6 Hit@1, while the base 4B agent reaches 29.1 when prompted to predict its own next call. SFT warmup raises this self-speculation Hit@1 to 44.1 for Qwen3-4B and 48.9 for Qwen3.5-4B. Joint agent-speculator RL further improves average Hit@1 to 61.2 and 66.3, respectively, with gains in the same direction on every benchmark. Since Hit@1 requires both the tool name and full argument dictionary to match, these gains indicate more reusable next-call predictions. We also evaluate the same SearchQA recipe at 8B (Appendix D.3) and measure next-call Hit@1 on WideSearch (Appendix D.4).

Agent performance is largely preserved.

Since the same parameter set is used in both agent and speculator modes for our self-speculating agent, improving next-call prediction can risk degrading task solving performance. In Table 2, we observe that task success remains stable after RL: average success changes from 26.6 to 27.7 for Qwen3-4B and from 49.2 to 50.6 for Qwen3.5-4B, with no systematic degradation across search QA or conversational tool-use benchmarks. The speculation gains above are therefore not obtained by sacrificing the deployed agent’s end-task performance.

4.3 Generalization Across Domains

An interesting question to ask is whether the ability to self-speculate generalizes across domains. We investigate this by evaluating checkpoints trained on one agent task in the other: the SearchQA-trained Qwen3-4B checkpoint is tested on τ -bench, and the ToolScale-trained Qwen3.5-4B checkpoint is tested on SearchQA. Table 3 shows that next-call prediction transfers across domains, while end-task success is less stable. For example, the SearchQA-trained Qwen3-4B checkpoint improves Airline/Retail Hit@1 from 35.6/39.7 to 45.6/50.1 after RL, while task success drops from 21.6/36.4 to 17.6/29.8. In the reverse direction, the ToolScale-trained Qwen3.5-4B checkpoint improves HotpotQA/MuSiQue Hit@1 from 25.4/28.3 to 35.4/37.9 after RL, while task success also decreases. Next-call prediction can transfer as call-level behavior, for example matching common argument patterns, while task solving remains more domain-specific. This is consistent with specialization and catastrophic forgetting during continued training (Li et al., 2024; Lin et al., 2024). Direct in-domain training remains strongest for both metrics.

4.4 Ablation Study

We ablate the stabilization techniques in the joint training recipe on SearchQA with Qwen3-4B. Figure 3 shows training-time reward traces when removing SFT warmup, optimizer reset, or alternating schedule, while Table 4 compares different agent:speculator update schedules under the same iteration budget.

Effect of SFT warmup. The w/o SFT warmup variant starts RL directly from the base model, testing whether the speculation suffix alone is enough

Model	HotpotQA		MuSiQue		Tau-Air		Tau-Ret	
	H@1	Succ.	H@1	Succ.	H@1	Succ.	H@1	Succ.
Qwen3-4B								
Base	23.2	52.1	26.4	15.1	35.6	21.6	39.7	36.4
SFT	37.8	49.4	40.9	13.8	41.2	18.9	45.8	32.0
RL	55.2	48.6	58.4	12.9	45.6	17.6	50.1	29.8
Qwen3.5-4B								
Base	25.4	63.1	28.3	25.1	42.1	56.4	46.8	69.7
SFT	31.8	60.0	34.2	22.8	58.6	53.1	62.9	66.0
RL	35.4	57.9	37.9	21.6	76.1	51.4	79.4	63.7

Table 3: Generalization across domains. We evaluate whether a self-speculator trained in one interaction family transfers to the other.




Figure 3: Reward trend over the first 100 steps for each mode: (left) agent reward and (right) speculator reward. The full method preserves agent reward while improving speculator reward. Removing SFT warmup, optimizer resets, or alternating updates leads to weaker or less stable reward trends.

to induce well-formed call-prediction behavior. In Figure 3, this variant lags the full method in speculation reward and is less stable early in training. This suggests that warmup provides a useful initialization before RL optimizes next-call prediction accuracy, because early updates can otherwise spend capacity on discovering the output format rather than matching the agent’s eventual call.

Effect of optimizer reset. The w/o optimizer reset variant keeps the alternating schedule but preserves optimizer state across mode switches. In Figure 3, it produces weaker and less stable reward trends than the full method. The two modes optimize different targets: agent updates improve task-solving behavior, while speculator updates improve next-call prediction. If the same optimizer state is carried from one mode into the other, momentum and adaptive statistics from the previous target can push the new updates in a less useful direction (Sutskever et al., 2013; Kingma and Ba, 2015). Resetting the optimizer at each mode switch gives each block a fresh optimizer state and makes training more stable by limiting cross-mode interference through optimizer history.

Effect of alternating updates. The w/o alternating updates trace in Figure 3 shows that mixing agent and speculator updates leads to early collapse in training. Table 4 compares block schedules under the same iteration budget. Switching every step is worst: the 1:1 schedule achieves only 31.8 average Hit@1 and 10.3 average task success. Performance improves as each mode is trained for longer consecutive updates, reaching 55.2 average Hit@1 and 26.1 average task success with the 4:8 schedule. This suggests that the speculator needs enough consecutive signal to track the agent’s current call distribution, while the agent still needs periodic updates to preserve task performance. We treat 4:8 as a SearchQA default rather than a universal rule; Appendix D.5 compares 4:4 and 4:8 on ToolScale.

$m(k)$	Hotpot		MuSiQ		BCP		Avg.	
schedule	H@1	Succ.↑	H@1	Succ.↑	H@1	Succ.↑	H@1	Succ.↑
1:1	31.8	18.4	34.6	6.8	29.1	5.6	31.8	10.3
2:2	41.9	39.6	44.7	13.9	38.8	15.2	41.8	22.9
4:4	50.1	49.8	53.2	16.3	47.5	7.2	50.3	24.4
4:8	55.2	53.0	58.4	17.1	52.1	8.1	55.2	26.1

Table 4: Ablation over agent:speculator update schedules on SearchQA. We report Hit@1 and task success on the three search datasets used in the main table.

5 Related Work

Reinforcement learning for LLM agents. Reinforcement learning is widely used to post-train LLMs, from RLHF and AI-feedback alignment (Ouyang et al., 2022; Bai et al., 2022) to recent reasoning-oriented RL methods and systems (Shao et al., 2024; DeepSeek-AI et al., 2025; Kimi Team et al., 2025; Yu et al., 2025), including efficient reasoning, verifier-based code RL, and external tool use (Hou et al., 2025; Liu et al., 2025; He et al., 2026; Schick et al., 2023; Patil et al., 2023; Chen et al., 2023; Liu et al., 2024, 2026). More recent work uses RL to optimize multi-turn tool behavior, including search and web agents (Jin et al., 2025; Song et al., 2025; Tan et al., 2025; Qi et al., 2025; Zheng et al., 2025), general tool-integrated reasoning and function calling (Qian et al., 2025; Zhang et al., 2025; Li et al., 2025; Feng et al., 2025; Singh et al., 2025), embodied multi-robot control (Ji et al., 2025), and long-horizon interactive agent training (Chen et al., 2025a; Wei et al., 2025; Luo et al., 2025). Our work also uses RL for tool-using agents, but optimizes a different objective: the same policy must

both solve the task and predict its own future tool calls from partial trajectories.

Tool call speculation. Speculative decoding accelerates language-model inference by verifying drafts from a faster model (Leviathan et al., 2023). Tool-call speculation applies the same idea at the action level: likely future tool calls are issued early so tool latency can be hidden behind model generation. Recent work explores this direction with external speculators, tool caches, action-specific strategies, workflow-pattern reuse, and system-level scheduling (Ye et al., 2026; Nichols et al., 2025; Zhong et al., 2026; Sui et al., 2026; Huang et al., 2025; Hooper et al., 2026). These methods show that action-level speculation can reduce latency, but they rely on a separate draft model, historical traces, or action-specific heuristics. In contrast, our self-speculating agent uses the deployed model itself as the speculator, avoiding a second model or trace cache while keeping speculation targets on-policy.

6 Conclusion

In this paper, we introduced self-speculating agents, where the deployed model also predicts its own next tool call through a speculation suffix. This design addresses the speculator-agent gap of external draft models while avoiding separate model and KV-cache overhead. We train the dual-mode agent with online joint agent-speculator RL using the agent’s own rollouts as speculation targets. Across agentic SearchQA and conversational tool-use benchmarks, self-speculation improves next-call prediction while preserving downstream task success. The present evidence covers structured tool calls; extending the method to open-ended actions remains future work.

Limitations

While our results show that a self-speculating agent can improve next-tool-call prediction while preserving downstream task performance, the present study is an initial step rather than a complete treatment of speculative tool use.

Our method assumes that speculative tool calls can be issued without changing the external environment. This is appropriate for read-only tools such as search, retrieval, database lookup, and many information-gathering API calls, where an incorrect speculation can simply be discarded. It is less suitable for tools that mutate external state,

such as placing an order, updating a database record, sending a message, or triggering an irreversible workflow. In such settings, speculative execution would require additional safeguards, such as dry-run modes, transactional rollback, or restricting speculation to the read-only prefix of a task.

The current reward and reuse rules are also limited to structured tool calls. A call is a tool name plus an argument dictionary, which makes exact matching a well-defined serving rule and lets argument F1 provide a dense training signal. This does not establish transfer to open-ended actions such as code, browser interaction, or free-form plans, where many outputs may be equivalent, token overlap is a poor proxy for correctness, and unsafe reuse is harder to detect. Such settings would need different rewards, equivalence checks, and safety constraints, which we leave to future work.

Our empirical evaluation is also limited in scope. We evaluate self-speculation on two agentic task families: multi-hop search question answering and conversational tool use. These settings cover different tool-call patterns, but they do not exhaust the diversity of agent environments, especially code execution, web navigation, long-running workflows, and multi-agent tool use. Moreover, our main training experiments use 4B-scale backbone models. Appendix D.3 reports a Qwen3-8B SearchQA run with the same recipe; we do not extrapolate to much larger models such as 70B. Larger models may exhibit different optimization dynamics and a different speculator-agent gap. Extending the evaluation to more domains and model scales is therefore an important direction for future work.

Acknowledgments

The UCSB team acknowledges support from the National Science Foundation (NSF) under Grant Nos. 2338252, 2302730, and 2619240.

References

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson El-

- hage, Nicholas Schiefer, Noemi Mercado, Nova Das-Sarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022. [Constitutional AI: Harmlessness from AI feedback](#). *arXiv preprint arXiv:2212.08073*.
- Victor Barres, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan. 2025. [\$\tau^2\$ -Bench: Evaluating conversational agents in a dual-control environment](#). *arXiv preprint arXiv:2506.07982*.
- Baian Chen, Chang Shu, Ehsan Shareghi, Nigel Collier, Karthik Narasimhan, and Shunyu Yao. 2023. [FireAct: Toward language agent fine-tuning](#). *arXiv preprint arXiv:2310.05915*.
- Kevin Chen, Marco Cusumano-Towner, Brody Huval, Aleksei Petrenko, Jackson Hamburger, Vladlen Koltun, and Philipp Krähenbühl. 2025a. [Reinforcement learning for long-horizon interactive LLM agents](#). *arXiv preprint arXiv:2502.01600*.
- Zijian Chen, Xueguang Ma, Shengyao Zhuang, Ping Nie, Kai Zou, Andrew Liu, Joshua Green, Kshama Patel, Ruoxi Meng, Mingyi Su, Sahel Sharifymoghaddam, Yanxi Li, Haoran Hong, Xinyu Shi, Xuye Liu, Nandan Thakur, Crystina Zhang, Luyu Gao, Wenhui Chen, and Jimmy Lin. 2025b. [BrowseComp-Plus: A more fair and transparent evaluation benchmark of deep-research agent](#). *arXiv preprint arXiv:2508.06600*.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Hanwei Xu, Honghui Ding, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jingchang Chen, Jingyang Yuan, Jinhao Tu, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaichao You, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingxu Zhou, Meng Li, Miaojuan Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shutong Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yao-hui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yudian Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. [DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning](#). *Nature*, 645:633–638.
- Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, Baoquan Zhong, Chengquan Jiang, Jinxin Chi, and Wanjuan Zhong. 2025. [ReTool: Reinforcement learning for strategic tool use in LLMs](#). *arXiv preprint arXiv:2504.11536*.
- Zhongmou He, Yee Man Choi, Kexun Zhang, Ivan Berovich, Jiabao Ji, Juntong Zhou, Dejian Xu, Aidan Zhang, Yixiao Zeng, and Lei Li. 2026. [HARDTEST-GEN: A high-quality RL verifier generation pipeline for LLM algorithmic coding](#). In *International Conference on Learning Representations*.
- Coleman Hooper, Minwoo Kang, Suhong Moon, Nicholas Lee, Eric Wen, John Wawrzyniec, Michael W. Mahoney, Yakun Sophia Shao, Amir Gholami, and Kurt Keutzer. 2026. [Speculative interaction agents: Building real-time agents with asynchronous I/O and speculative tool calling](#). *arXiv preprint arXiv:2605.13360*.
- Bairu Hou, Yang Zhang, Jiabao Ji, Yujian Liu, Kaizhi Qian, Jacob Andreas, and Shiyu Chang. 2025. [ThinkPrune: Pruning long chain-of-thought of LLMs via reinforcement learning](#). *arXiv preprint arXiv:2504.01296*.
- Zixiao Huang, Wen Zeng, Tianyu Fu, Tengxuan Liu, Yizhou Sun, Ke Hong, Xinhao Yang, Chengchun Liu, Yan Li, Quanlu Zhang, Guohao Dai, Zhenhua Zhu, and Yu Wang. 2025. [Reducing latency of LLM search agent via speculation-based algorithm-system co-design](#). *arXiv preprint arXiv:2511.20048*.
- Jiabao Ji, Yongchao Chen, Yang Zhang, Ramana Rao Kompella, Chuchu Fan, Gaowen Liu, and Shiyu Chang. 2025. [Collision- and reachability-aware](#)

- multi-robot control with grounded LLM planners. *arXiv preprint arXiv:2505.20573*.
- Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan O. Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. [Search-R1: Training LLMs to reason and leverage search engines with reinforcement learning](#). In *Proceedings of the 2nd Conference on Language Modeling*.
- Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, Chuning Tang, Congcong Wang, Dehao Zhang, Enming Yuan, Enzhe Lu, Fengxiang Tang, Flood Sung, Guangda Wei, Guokun Lai, Haiqing Guo, Han Zhu, Hao Ding, Hao Hu, Hao Yang, Hao Zhang, Haotian Yao, Haotian Zhao, Haoyu Lu, Haoze Li, Haozhen Yu, Hongcheng Gao, Huabin Zheng, Huan Yuan, Jia Chen, Jianhang Guo, Jianlin Su, Jianzhou Wang, Jie Zhao, Jin Zhang, Jingyuan Liu, Junjie Yan, Junyan Wu, Lidong Shi, Ling Ye, Longhui Yu, Mengnan Dong, Neo Zhang, Ningchen Ma, Qiwei Pan, Qucheng Gong, Shaowei Liu, Shengling Ma, Shupeng Wei, Sihan Cao, Siying Huang, Tao Jiang, Weihao Gao, Weimin Xiong, Weiran He, Weixiao Huang, Weixin Xu, Wenhao Wu, Wenyang He, Xianguhui Wei, Xianqing Jia, Xingzhe Wu, Xinran Xu, Xinxing Zu, Xinyu Zhou, Xuehai Pan, Y. Charles, Yang Li, Yangyang Hu, Yangyang Liu, Yanru Chen, Yejie Wang, Yibo Liu, Yidao Qin, Yifeng Liu, Ying Yang, Yiping Bao, Yulun Du, Yuxin Wu, Yuzhi Wang, Zaida Zhou, Zhaoji Wang, Zhaowei Li, Zhen Zhu, Zheng Zhang, Zhexu Wang, Zhilin Yang, Zhiqi Huang, Zihao Huang, Ziyao Xu, Zonghan Yang, and Zongyu Lin. 2025. [Kimi k1.5: Scaling reinforcement learning with LLMs](#). *arXiv preprint arXiv:2501.12599*.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *International Conference on Learning Representations*.
- Yaniv Leviathan, Matan Kalman, and Yossi Matias. 2023. [Fast inference from transformers via speculative decoding](#). In *Proceedings of the 40th International Conference on Machine Learning*, pages 19274–19286.
- Hongyu Li, Liang Ding, Meng Fang, and Dacheng Tao. 2024. [Revisiting catastrophic forgetting in large language model tuning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4297–4308.
- Xuefeng Li, Haoyang Zou, and Pengfei Liu. 2025. [ToRL: Scaling tool-integrated RL](#). *arXiv preprint arXiv:2503.23383*.
- Yong Lin, Hangyu Lin, Wei Xiong, Shizhe Diao, Jianmeng Liu, Jipeng Zhang, Rui Pan, Haoxiang Wang, Wenbin Hu, Hanning Zhang, Hanze Dong, Renjie Pi, Han Zhao, Nan Jiang, Heng Ji, Yuan Yao, and Tong Zhang. 2024. [Mitigating the alignment tax of RLHF](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 580–606.
- Yujian Liu, Jiabao Ji, Li An, Tommi Jaakkola, Yang Zhang, and Shiyu Chang. 2026. [How well do agentic skills work in the wild: Benchmarking LLM skill usage in realistic settings](#). *arXiv preprint arXiv:2604.04323*.
- Yujian Liu, Jiabao Ji, Yang Zhang, Wenbo Guo, Tommi Jaakkola, and Shiyu Chang. 2025. [HarnessLLM: Automatic testing harness generation via reinforcement learning](#). *arXiv preprint arXiv:2511.01104*.
- Zuxin Liu, Thai Hoang, Jianguo Zhang, Ming Zhu, Tian Lan, Shirley Kokane, Juntao Tan, Weiran Yao, Zhiwei Liu, Yihao Feng, Rithesh Murthy, Liangwei Yang, Silvio Savarese, Juan Carlos Niebles, Huan Wang, Shelby Heinecke, and Caiming Xiong. 2024. [APIGen: Automated pipeline for generating verifiable and diverse function-calling datasets](#). In *Advances in Neural Information Processing Systems*, volume 37.
- Xufang Luo, Yuge Zhang, Zhiyuan He, Zilong Wang, Siyun Zhao, Dongsheng Li, Luna K. Qiu, and Yuqing Yang. 2025. [Agent lightning: Train ANY AI agents with reinforcement learning](#). *arXiv preprint arXiv:2508.03680*.
- Daniel Nichols, Prajwal Singhanian, Charles Jekel, Abhinav Bhatele, and Harshitha Menon. 2025. [Optimizing agentic language model inference via speculative tool calls](#). *arXiv preprint arXiv:2512.15834*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744.
- Shishir G. Patil, Tianjun Zhang, Xin Wang, and Joseph E. Gonzalez. 2023. [Gorilla: Large language model connected with massive APIs](#). *arXiv preprint arXiv:2305.15334*.
- Zehan Qi, Xiao Liu, Iat Long Iong, Hanyu Lai, Xueqiao Sun, Jiadai Sun, Xinyue Yang, Yu Yang, Shuntian Yao, Wei Xu, Jie Tang, and Yuxiao Dong. 2025. [WebRL: Training LLM web agents via self-evolving online curriculum reinforcement learning](#). In *International Conference on Learning Representations*.
- Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiushi Chen, Dilek Hakkani-Tür, Gokhan Tur, and Heng Ji. 2025. [ToolRL: Reward is all tool learning needs](#). *arXiv preprint arXiv:2504.13958*.

- Qwen Team. 2025. [Qwen3 technical report](#). *arXiv preprint arXiv:2505.09388*.
- Qwen Team. 2026. Qwen3.5-4B model card. <https://huggingface.co/Qwen/Qwen3.5-4B>. Accessed 2026-05-25.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language models can teach themselves to use tools](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 68539–68551.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [DeepSeekMath: Pushing the limits of mathematical reasoning in open language models](#). *arXiv preprint arXiv:2402.03300*.
- Joykirat Singh, Raghav Magazine, Yash Pandya, and Akshay Nambi. 2025. [Agentic reasoning and tool integration for LLMs via reinforcement learning](#). *arXiv preprint arXiv:2505.01441*.
- Huatong Song, Jinhao Jiang, Yingqian Min, Jie Chen, Zhipeng Chen, Wayne Xin Zhao, Lei Fang, and Ji-Rong Wen. 2025. [R1-Searcher: Incentivizing the search capability in LLMs via reinforcement learning](#). *arXiv preprint arXiv:2503.05592*.
- Hongjin Su, Shizhe Diao, Ximing Lu, Mingjie Liu, Jiacheng Xu, Xin Dong, Yonggan Fu, Peter Belcak, Hanrong Ye, Hongxu Yin, Yi Dong, Evelina Bakhurina, Tao Yu, Yejin Choi, Jan Kautz, and Pavlo Molchanov. 2025. [ToolOrchestra: Elevating intelligence via efficient model and tool orchestration](#). *arXiv preprint arXiv:2511.21689*.
- Yifan Sui, Han Zhao, Rui Ma, Zhiyuan He, Hao Wang, Jianxun Li, Kaiqiang Xu, Kai Chen, and Yuqing Yang. 2026. [Parallelizing tool execution and LLM generation for low-latency agent serving](#). *arXiv preprint arXiv:2603.18897*.
- Weiwei Sun, Miao Lu, Zhan Ling, Kang Liu, Xuesong Yao, Yiming Yang, and Jiecao Chen. 2025. [Scaling long-horizon LLM agent via context-folding](#). *arXiv preprint arXiv:2510.11967*.
- Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. 2013. [On the importance of initialization and momentum in deep learning](#). In *Proceedings of the 30th International Conference on Machine Learning*, pages 1139–1147.
- Zhiwen Tan, Jiaming Huang, Qintong Wu, Hongxuan Zhang, Chenyi Zhuang, and Jinjie Gu. 2025. [RAG-R1: Incentivizing the search and reasoning capabilities of LLMs through multi-query parallelism](#). *arXiv preprint arXiv:2507.02962*.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. [MuSiQue: Multi-hop questions via single-hop question composition](#). *Transactions of the Association for Computational Linguistics*, 10:539–554.
- Quan Wei, Siliang Zeng, Chenliang Li, Zhongruo Wang, William Brown, Oana Frunza, Wei Deng, Anderson Schneider, Yuriy Nevmyvaka, Yang Katie Zhao, Alfredo Garcia, and Mingyi Hong. 2025. [Reinforcing multi-turn reasoning in LLM agents via fine-grained reward structure and credit assignment](#). *arXiv preprint arXiv:2505.11821*.
- Ryan Wong, Jiawei Wang, Junjie Zhao, Li Chen, Yan Gao, Long Zhang, Xuan Zhou, Zuo Wang, Kai Xiang, Ge Zhang, Wenhao Huang, Yang Wang, and Ke Wang. 2025. [WideSearch: Benchmarking agentic broad info-seeking](#). *arXiv preprint arXiv:2508.07999*.
- Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. 2024. [\$\tau\$ -bench: A benchmark for tool-agent-user interaction in real-world domains](#). *arXiv preprint arXiv:2406.12045*.
- Naimeng Ye, Arnav Ahuja, Georgios Liargkovas, Yunan Lu, Kostis Kaffes, and Tianyi Peng. 2026. [Speculative actions: A lossless framework for faster agentic systems](#). In *International Conference on Learning Representations*.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Juncai Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Ru Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Yonghui Wu, and Mingxuan Wang. 2025. [DAPO: An open-source LLM reinforcement learning system at scale](#). In *Advances in Neural Information Processing Systems*, volume 38.
- Yabo Zhang, Yihan Zeng, Qingyun Li, Zhen Hu, Kevin Han, and Wangmeng Zuo. 2025. [Tool-R1: Sample-efficient reinforcement learning for agentic tool use](#). *arXiv preprint arXiv:2509.12867*.
- Yuxiang Zheng, Dayuan Fu, Xiangkun Hu, Xiaojie Cai, Lyumanshan Ye, Pengrui Lu, and Pengfei Liu. 2025. [DeepResearcher: Scaling deep research via reinforcement learning in real-world environments](#). *arXiv preprint arXiv:2504.03160*.
- Shuzhang Zhong, Baotong Lu, Qi Chen, Chuanjie Liu, Fan Yang, and Meng Li. 2026. [DualSpec: Accelerating deep research agents via dual-process action speculation](#). *arXiv preprint arXiv:2603.07416*.

A Implementation Details

During evaluation, speculation is run from the shared prefix h_t in Equation 1, immediately after the previous tool observation and before the agent’s next reasoning step. The speculator predicts a candidate call, but the Hit@1 protocol executes only the agent’s actual call. This keeps the evaluated trajectory identical to the ordinary agent trajectory and makes each speculation label the current agent’s own next action. During RL, we use the same prefix–next-call construction to derive speculation examples from freshly sampled agent rollouts, as described in Section 3.2.

A.1 Training Hyperparameters

Table 5 lists the default training settings used for joint agent-speculator RL. Unless stated otherwise, ablations use the same rollout budget, learning rate, schedule, and evaluation protocol as the corresponding full run.

Setting	Value
Backbone	Qwen3-4B, Qwen3.5-4B
RL method	DAPO-style GPO
$G_{\text{ag}} / G_{\text{sp}}$	8 / 8
Schedule	4 agent / 8 speculator
Reset	Each mode switch
Learning rate	3×10^{-6}
Length limit	16K tokens (SearchQA); 32K (τ -bench)
RL budget	200 steps per setting
Hardware	8 H100 (SearchQA); 8 H200 (ToolScale)
SFT data	Successful trajectories and speculation examples
SFT target	Agent trajectories; next calls under the suffix

Table 5: Default training hyperparameters.

A.2 Prompt and Tool Schema Details

The fixed speculation suffix used in all experiments is:

```
<think> Okay, let's see. The user
provided what I need. I'll look it up.
The next step is to make the tool call.
</think>
```

For each speculation query, the suffix is appended after the shared prefix h_t and before the model decodes the candidate tool call. We use the same suffix for off-the-shelf speculator evaluation, SFT warmup example construction, joint agent-speculator RL, and final speculation evaluation. The suffix does not name a benchmark, tool, argument key, or target answer.

We represent every tool call as a structured action $a = (n, \sigma)$, where n is the tool name and σ is a dictionary of named arguments. SearchQA actions contain a search tool name and query string. τ -bench-style actions use the API name and schema-constrained argument dictionary supplied by the environment. Calls that cannot be parsed into the expected structured format are treated as invalid predictions and receive zero speculation reward.

For Hit@1 evaluation, a prediction is correct only if the parsed tool name and the complete parsed argument dictionary exactly match the agent’s eventual call. For the shaped speculation reward used during RL, we first require the tool name to match and then compute macro token-F1 over target argument keys, as defined in Equation 2. Missing predicted argument keys receive zero score, and extra predicted keys do not compensate for missing target keys. This keeps the training reward dense enough to provide partial credit for close arguments, while the deployment reuse metric remains exact structured-call equality.

B Artifact Licenses

Table 6 summarizes the licenses reported by the upstream publishers for the external artifacts used in our experiments. We use these artifacts for research evaluation and model training under their published terms, and we retain the corresponding citations and license notices when redistributing metadata or scripts.

Artifact	Reported license
Qwen3-0.6B, Qwen3-1.7B, Qwen3-4B, Qwen3-8B	Apache-2.0
Qwen3.5-0.8B, Qwen3.5-2B, Qwen3.5-4B	Apache-2.0
Qwen3-Embedding-0.6B	Apache-2.0
HotpotQA	CC BY-SA 4.0
MuSiQue	CC BY 4.0
BrowseComp-Plus	MIT
τ -bench / τ^2 -bench	MIT
ToolScale	Upstream dataset terms

Table 6: Licenses or access terms reported by the upstream model, dataset, or benchmark publishers for artifacts used in the experiments.

C Data Statistics

Table 7 reports the train, evaluation, and SFT warmup counts used in our experiments. SFT

counts are separated into agent-trajectory rows and speculation rows; slash-separated counts denote Qwen3-4B / Qwen3.5-4B when both warmup files are used. Following the subset-based evaluation practice in R1-Searcher (Song et al., 2025), SearchQA evaluation uses 500 examples per dataset sampled from the FlashRAG development splits, rather than the full public development or test sets. We do not assume that these are the identical examples released by R1-Searcher unless their instance IDs match explicitly in metadata.

Resource	Train	Eval	SFT agent	SFT spec
FlashQA / SearchQA	283,338	2,500	1,658 / 1,351	1,658 / 1,351
HotpotQA	–	500	–	–
MuSiQue	–	500	–	–
BrowseComp-Plus	680	150	–	–
ToolScale	4,063	–	0	5,000
τ -bench Airline	–	50	–	–
τ -bench Retail	–	114	–	–

Table 7: Dataset, benchmark, and SFT warmup counts. The ToolScale warmup source contains speculation prompt-label rows only.

D Additional Experimental Results

The main paper reports next-call Hit@1 and task success for 4B self-speculating agents on SearchQA and τ -bench. Hit@1 remains exact match on the tool name and argument dictionary, and task success follows Section 4.1 unless noted.

D.1 End-to-End Serving under Synthetic Tool Latency

Hit@1 in Tables 1 and 2 measures whether a predicted call would be reusable. End-to-end time from the user query to the final answer further depends on tool latency. On our local SearchQA and τ -bench setups, tools return in microseconds, far below the seconds-scale remote APIs that motivate speculation, so native timing cannot reflect that serving regime.

To model remote tools, we attach an independent delay to every tool call, $T_{\text{tool}} \sim \mathcal{U}(0.5\text{s}, 1.5\text{s})$, with mean 1.0s. The delay is sampled per call and frozen, so a matched no-speculation baseline sees the same latency draws as the speculative run. We then run the serving loop from Section 3.1: from h_t , the speculator predicts $\hat{a}_t$ and pre-executes that call while the agent generates r_t and a_t . An exact match reuses the returned observation; a mismatch drops the speculative result and executes a_t , paying both the extra request and the sampled delay

of the true call. There is no confidence threshold. We evaluate Qwen3-4B self, SFT, and RL on 50 MuSiQue queries and 50 τ -bench Airline queries (the same splits as Section 2). Reported times include model generation, speculative requests, misses, and fallback. Speedup is the ratio of the matched no-speculation time to the speculative time for the same checkpoint and queries.

	MuSiQue			Airline		
	H@1 $\uparrow$	E2E	Spd. $\uparrow$	H@1 $\uparrow$	E2E	Spd. $\uparrow$
Self	25.3	3.65	1.04 $\times$	37.2	8.74	1.03 $\times$
SFT	40.9	3.52	1.08 $\times$	51.2	8.41	1.07 $\times$
RL	58.4	3.36	1.13$\times$	68.0	8.04	1.12$\times$

Table 8: Qwen3-4B end-to-end serving under synthetic tool latency $T_{\text{tool}} \sim \mathcal{U}(0.5\text{s}, 1.5\text{s})$. E2E is seconds; speedup is versus a matched no-speculation run.

Table 8 links Hit@1 to wall-clock time. From self to SFT to RL, Hit@1 rises and end-to-end time falls on both benchmarks. On MuSiQue, Qwen3-4B-RL reaches 58.4 Hit@1 and 3.36s (1.13 $\times$ over no speculation), compared with 25.3 Hit@1 and 3.65s for the prompted base model. On Airline, the same checkpoint reaches 68.0 Hit@1 and 8.04s (1.12 $\times$), compared with 37.2 Hit@1 and 8.74s. The relative speedups are similar even though Airline trajectories are longer in absolute time, which is consistent with hiding a shared per-call delay behind reasoning rather than shortening generation. The wall-clock gains remain after counting misses and fallback, but they are modest: a miss still pays the sampled tool delay and an extra request, and speculation does not remove generation time. These measurements use synthetic remote-tool latency.

D.2 Failure Analysis of Near-Miss Predictions

The speculation reward in Equation 2 is name-conditioned argument F1, so a nearly correct call can receive partial credit during training. Serving reuse is stricter: $\hat{a}_t$ is reusable only when it matches a_t exactly, because a single wrong field can change the tool result. Hit@1 therefore collapses every non-exact prediction into a miss, and the main tables do not show whether those misses are malformed calls, wrong tools, or close argument errors.

We take every next-call prediction from the Qwen3-4B-RL checkpoints on SearchQA and

ToolScale and assign it to one of four disjoint outcomes: exact match, invalid or non-executable output, a valid call to the wrong tool, or the correct tool with at least one incorrect argument.

Prediction outcome	SearchQA	ToolScale
Exact match / reusable	55.2	70.2
Invalid / non-executable	3.5	2.6
Wrong tool / action	0.0	9.4
Right tool, wrong arguments	41.3	17.8

Table 9: Qwen3-4B-RL prediction breakdown (%). Each column sums to 100%. Exact-match shares recover the Hit@1 averages in Table 2.

The exact-match shares recover the main-table averages: 55.2 Hit@1 on SearchQA and 70.2 on ToolScale. Invalid calls are rare (3.5% and 2.6%), so the suffix-conditioned format is largely in place after joint RL. The dominant miss is the right tool with wrong arguments: 41.3% on SearchQA and 17.8% on ToolScale. These predictions are often close, such as a rewritten search query or one incorrect API value, which is why argument F1 can still give a useful training signal, but they remain unsafe to reuse. Wrong-tool errors are 0.0% on SearchQA, which exposes a single search tool, and 9.4% on ToolScale, where the agent must choose among many APIs. SearchQA is therefore almost entirely an argument-matching problem, while ToolScale still has a non-trivial API selection error. In both cases, most misses are nearly correct rather than malformed.

D.3 Scale to Qwen3-8B

The main experiments use two 4B families, which test transfer across model lines but not scale. We therefore repeat the SearchQA recipe at 8B. We train Qwen3-8B on FlashQA with the same mixed SFT warmup, 200 DAPO steps, $G_{\text{ag}} = G_{\text{sp}} = 8$, learning rate 3×10^{-6} , and 4:8 schedule as the Qwen3-4B SearchQA runs, and evaluate on HotpotQA and MuSiQue with the same Hit@1 and task-success protocol. We do not retrain 8B on ToolScale or BCP in this study.

	HotpotQA		MuSiQue		Avg.	
	H@1↑	Succ.↑	H@1↑	Succ.↑	H@1↑	Succ.↑
Base	28.7	52.8	31.9	16.9	30.3	34.9
SFT	44.3	52.1	47.8	17.6	46.1	34.9
RL	62.4	54.2	65.8	19.3	64.1	36.8

Table 10: Qwen3-8B SearchQA results with the same FlashQA recipe as the 4B runs.

Table 10 follows the 4B pattern. The prompted 8B agent already reaches 30.3 average Hit@1. SFT warmup raises this to 46.1 without changing average task success (34.9). Joint RL further raises average Hit@1 to 64.1 and average success to 36.8, with gains on both HotpotQA (62.4 Hit@1, 54.2 success) and MuSiQue (65.8 Hit@1, 19.3 success). Relative to the 4B SearchQA columns in Table 2, the 8B RL checkpoint is stronger on these two datasets, with the same qualitative pattern: large Hit@1 gains and stable or slightly higher task success.

D.4 Transfer to WideSearch

HotpotQA, MuSiQue, and BCP cover multi-hop and deep-research search, but not broad information seeking in which the agent must collect many atomic facts and assemble a structured output. We therefore evaluate the FlashQA-trained 4B checkpoints on WideSearch (Wong et al., 2025) without further training, using the same prefix h_t , speculation suffix, and exact-match Hit@1 rule as in Section 4.1. We report Hit@1 only. WideSearch also defines an Item F1 metric for the final structured output; we do not report it here.

Model	Hit@1↑
Qwen3-4B	18.7
Qwen3-4B-RL	34.8
Qwen3.5-4B	21.9
Qwen3.5-4B-RL	39.2

Table 11: Next-call Hit@1 on WideSearch (FlashQA-trained 4B checkpoints).

Table 11 shows that next-call prediction transfers, though the absolute Hit@1 is lower than on HotpotQA and MuSiQue, as expected for longer and broader search trajectories. RL raises Hit@1 from 18.7 to 34.8 for Qwen3-4B and from 21.9 to 39.2 for Qwen3.5-4B. The prompted base models already sit below their in-domain SearchQA Hit@1, so the benchmark is not a replay of the training distribution. The RL gains show that the

learned speculator mode is not confined to the original evaluation sets.

D.5 ToolScale Update Schedule

Table 4 selects 4:8 as the best agent:speculator block schedule on Qwen3-4B SearchQA under a fixed 200-step budget. ToolScale uses more APIs and longer conversations, so we also compare 4:4 and 4:8 there while holding the remaining training settings fixed, and evaluate on τ -bench Airline and Retail.

	Airline		Retail		Avg.	
	H@1 $\uparrow$	Succ. $\uparrow$	H@1 $\uparrow$	Succ. $\uparrow$	H@1 $\uparrow$	Succ. $\uparrow$
4:4	66.7	22.8	71.0	36.5	68.9	29.7
4:8	68.0	23.1	72.3	37.0	70.2	30.1

Table 12: ToolScale 4:4 versus 4:8 on Qwen3-4B. The 4:8 row matches Table 2.

Table 12 shows that the two schedules are close, and that 4:8 is slightly better on every column: Airline 68.0/23.1 versus 66.7/22.8, Retail 72.3/37.0 versus 71.0/36.5, and 70.2/30.1 versus 68.9/29.7 on average. The 4:8 row matches the Qwen3-4B-RL ToolScale numbers in Table 2. The gap is smaller than on SearchQA, where 4:8 led 4:4 by 4.9 Hit@1 and 1.7 success. The two experiments support 4:8 as a default rather than a constant. We recommend choosing the ratio on held-out data using both Hit@1 and task success.

E LLM Usage Statement

We used large language models to assist with paper grammar editing and with code implementation and debugging.