

Beyond ReAct: A Planner-Centric Framework for Complex Tool-Augmented LLM Reasoning

Xiaolong Wei^{1*}, Yuehu Dong^{2*}, Xingliang Wang³, Xingyu Zhang⁴, Zhejun Zhao^{2†}, Dongdong Shen², Long Xia², Dawei Yin²

¹Beihang University

²Baidu Inc.

³Beijing University of Posts and Telecommunications

⁴Beijing Jiaotong University

xiaolongwei@buaa.edu.cn, wangxingliang@bupt.edu.cn, zhangxingyu@bjtu.edu.cn,

{dongyuehu, zhaozhejun, shendongdong}@baidu.com, long.phil.xia@gmail.com, yindawei@acm.org

Abstract

Existing tool-augmented large language models (LLMs) encounter significant challenges when processing complex queries. Current frameworks such as ReAct are prone to local optimization traps due to their reliance on incremental decision-making processes. To address these limitations, we propose a novel Planner-centric Plan-Execute paradigm that fundamentally resolves local optimization bottlenecks through architectural innovation. Central to our approach is a novel Planner model that performs global Directed Acyclic Graph (DAG) planning for complex queries, enabling optimized execution beyond conventional tool coordination. We also introduce ComplexTool-Plan, a large-scale benchmark dataset featuring complex queries that demand sophisticated multi-tool composition and coordination capabilities. Additionally, we develop a two-stage training methodology that integrates Supervised Fine-Tuning (SFT) with Group Relative Policy Optimization (GRPO), systematically enhancing the Planner’s tool selection accuracy and global planning awareness through structured DAG-based planning. When integrated with a capable executor, our framework achieves state-of-the-art performance on the StableToolBench benchmark for complex user queries, demonstrating superior end-to-end execution capabilities and robust handling of intricate multi-tool workflows.

Introduction

Large Language Models (LLMs) have demonstrated remarkable prowess in language processing, yet their inherent knowledge is static and they lack direct interaction with the external world. A transformative solution is tool augmentation, which empowers LLMs to act as autonomous agents by calling external APIs and tools (Schick et al. 2023; Patil et al. 2024). This paradigm shifts LLMs from being mere text generators to active problem-solvers capable of tackling complex, real-world tasks, a trend exemplified by systems like HuggingGPT (Shen et al. 2023) and the rise of multi-agent frameworks (Hong et al. 2023). The evaluation

*Equal contribution.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

of these agentic capabilities has itself become a major research thrust, with new benchmarks like AgentBench (Liu et al. 2023) setting the stage for more rigorous assessment.

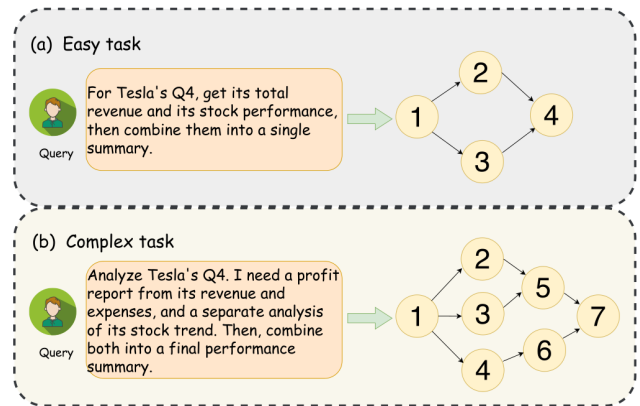


Figure 1: An example of a simple versus a complex task. A simple query results in a basic, parallel DAG, while a complex query involving nested logic is translated into a more elaborate, multi-level DAG.

The dominant approach for orchestrating tool usage relies on reactive, step-by-step reasoning frameworks like ReAct (Yao et al. 2023b), often augmented by self-reflection techniques (Shinn et al. 2023). However, this paradigm suffers from inherent local optimization traps due to its incremental decision-making process. While potentially effective for simple queries, its reactive nature often falters on complex tasks. As illustrated in Fig. 1, complex tasks involve more intricate dependencies than simple ones. To mitigate this, advanced methods employ tree search (Yao et al. 2023a) or explicit planning prompts (Wang et al. 2023) to explore multiple paths. While these methods improve upon purely linear reasoning, they fail to architecturally resolve the fundamental bottleneck of local optimization. Critically, they remain focused on finding an optimal sequential path, overlooking opportunities for parallelism inherent in complex workflows. Moreover, the computational overhead of constructing and traversing search trees, requiring numerous LLM

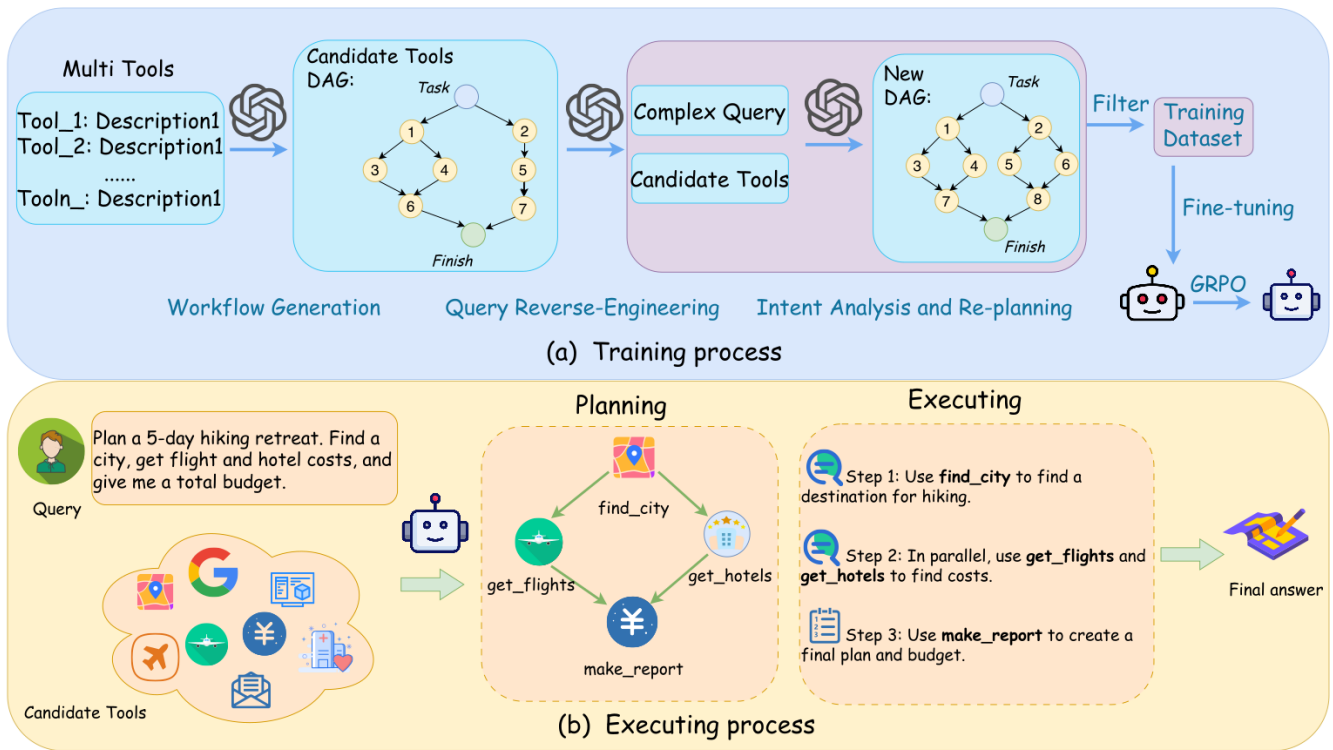


Figure 2: The figure illustrates our proposed framework. (a) The Training Process shows our automated pipeline for creating a training dataset and then training the Planner model via fine-tuning and reinforcement learning (GRPO). (b) The Executing Process demonstrates how the trained Planner takes a user query, generates a parallelizable execution plan as a Directed Acyclic Graph (DAG), and orchestrates the tools to produce the final answer.

calls, can be prohibitive.

To address this fundamental challenge, we adopt a framework that decouples planning from execution (Gao et al. 2023; Chen et al. 2022). Our approach employs a dedicated Planner model, trained to translate complex user queries into an execution plan structured as a DAG. In this DAG, nodes represent selected tools, and directed edges capture dependency relationships between them. By optimizing the Planner to improve the accuracy of both node selection (tool choices) and edge prediction (dependency relationships) in the DAG, we achieve more effective plans for complex queries, improving workflow execution compared to sequential models.

Realizing such a global planner presents two significant hurdles. First, there is a lack of large-scale, structured data required to train a model to generate complex plans, a challenge addressed by recent large-scale tool-use benchmarks like ToolBench (Qin et al. 2023) and API-Bank (Li et al. 2023b). Second, evaluating the quality of a plan is non-trivial. To overcome these, we introduce **ComplexTool-Plan**, a new, large-scale benchmark with graded difficulty levels. It serves as both the training ground and the evaluation standard for sophisticated planning agents. We then devise a two-stage hybrid training strategy for our Planner, combining Supervised Fine-Tuning (SFT) with subsequent

refinement via Reinforcement Learning (RL) (Ouyang et al. 2022), a strategy that aligns with recent findings on using self-play and RL to elicit stronger reasoning (Chen et al. 2024; Wei et al. 2025).

Our extensive experiments validate the efficacy of our decoupled approach. On ComplexTool-Plan, our Planner substantially outperforms leading baselines. Crucially, when our Planner is paired with a powerful executor model like GPT-4o, the integrated system establishes a new state-of-the-art (SOTA) on the widely recognized end-to-end benchmark, StableToolBench (Guo et al. 2024), underscoring the practical superiority of our framework.

In summary, our key contributions are:

- We enhance task success in tool-augmented LLMs by proposing a planner-centric plan-execute paradigm framework and developing a DAG-optimized Planner, which significantly improves the accuracy of tool selection (nodes) and dependency modeling (edges).
- We construct ComplexTool-Plan, a large-scale, difficulty-graded benchmark designed to foster and evaluate complex agentic planning capabilities.
- Through extensive experiments, we demonstrate that our framework achieves state-of-the-art performance, offering a more scalable, efficient, and robust solution for complex multi-tool orchestration.

Related Works

We survey the rapidly evolving landscape of tool-augmented LLMs (Li et al. 2025c,b) along three primary axes: the evolution of tool-calling paradigms, RL-based policy optimization, and the development of evaluation benchmarks.

Tool Calling Paradigms

The paradigms for integrating LLMs with external tools have evolved from the interleaved "Reasoning-Acting" steps of ReAct (Yao et al. 2023b) to more reliable, natively supported structured API calls. A significant branch of research leverages code as a universal tool interface, with works like Program-of-Thought (PoT) (Chen et al. 2022) and PAL (Gao et al. 2023) offloading complex computations to code interpreters. This trend has produced highly specialized agents like SWE-agent (Yang et al. 2024), which can autonomously resolve software engineering tasks. For instance, advanced models are increasingly required for challenging real-world applications like multi-modal fake news detection, which involves reasoning over diverse data from various sources and domains (Tong et al. 2025; Lu, Tong, and Ye 2025).

Recent work is also moving beyond general-purpose models towards those specifically fine-tuned for tool use (Zhao et al. 2025). For instance, Toolformer (Schick et al. 2023) demonstrates that a language model can teach itself to use tools by learning to insert beneficial API calls into text, thus significantly improving its ability to handle complex, tool-dependent instructions in a zero-shot manner. Concurrently, the interaction paradigm is expanding to visual environments, with powerful GUI agents like CogAgent (Hong et al. 2024) demonstrating the ability to understand and interact with graphical user interfaces. Beyond reactive or single-agent frameworks, Li et al. (Li et al. 2025a) introduce an *AI Search Paradigm* that formalizes reasoning as coordinated planning among multiple LLM-based agents, while Chen et al. (Chen et al. 2025) propose a proactive multi-agent orchestration framework for complex information-seeking tasks. Our work contributes to this evolution by focusing on the upfront, holistic planning of tool interactions, a critical component for efficient orchestration.

RL for Tool Policy Optimization

While Supervised Fine-Tuning (SFT) can teach models to mimic tool-use trajectories, Reinforcement Learning (RL) is superior for strategic decision-making in unseen scenarios. The latest SOTA models, such as Llama 3 (Dubey et al. 2024), heavily rely on advanced RLHF techniques to refine their policy. Building on foundational work demonstrating that RL can enhance reasoning (Guo et al. 2025), recent research has shown that RL teaches not just the "syntax" of tool calling, but the "strategy" (Feng et al. 2025; Jin et al. 2025).

The field is advancing towards more sophisticated, outcome-driven optimization loops. For instance, AlphaCodium (Ridnik, Kredo, and Friedman 2024) employs a test-based, multi-stage iterative process. Similarly, influential frameworks like Tree of Thoughts (ToT) (Yao et al. 2023a), which explores multiple reasoning paths, and Reflexion (Shinn et al. 2023), which leverages self-reflection,

highlight the trend towards more deliberate and verifiable reasoning. Our hierarchical reward function aligns with this trend by providing a structured, outcome-driven signal to teach the model a robust planning strategy.

Benchmarks for Tool Calling

The evaluation of LLM tool-calling capabilities relies on increasingly sophisticated benchmarks, which have evolved from early interactive environments like ALFWorld (Shridhar et al. 2020) and WebShop (Yao et al. 2022). In recent years, benchmarks have grown in scale and realism, with examples like ToolBench (Qin et al. 2023) and its more stable successor, StableToolBench (Guo et al. 2024). The focus has also expanded towards greater realism, as seen in the web environment of WebArena (Zhou et al. 2023), and more challenging tasks, as in the GAIA benchmark (Mialon et al. 2023).

To support this research, new platforms are emerging that focus on specific aspects like evaluating the correctness of API calls in API-Bank (Li et al. 2023b) and facilitating multi-agent collaboration in open-source frameworks like MetaGPT (Hong et al. 2023). Our work contributes a new benchmark, ComplexTool-Plan, which fills a specific gap in this landscape by targeting the evaluation of an agent's ability to generate complex, non-linear plans.

Methodology

Our method trains a language model to plan complex multi-tool tasks, as shown in Fig. 2. The framework includes: a formal **Problem Formulation**, our **ComplexTool-Plan** data generation pipeline, and a **Model Training** process using both Supervised Fine-Tuning (SFT) and Reinforcement Learning (RL).

Problem Formulation

We formalize multi-tool task planning as a structured prediction task of learning a policy $\pi : Q \times T \rightarrow \mathcal{G}$. Given a natural language query Q and a toolset T , the policy generates an execution plan as a Directed Acyclic Graph (DAG), $G = (V, E)$, where vertices $V \subseteq T$ are tools and edges E represent data dependencies. The goal is to train a model M_θ to find the optimal policy π^* that maximizes the expected plan utility $U(G)$:

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{(Q,T) \sim p(Q,T)} [U(M_\theta(Q,T))] \quad (1)$$

The utility $U(\cdot)$ is realized through a reward function.

ComplexTool-Plan

To overcome data scarcity, we developed ComplexTool-Plan, a three-stage automated pipeline to generate our training dataset, D_{train} .

- **Workflow Generation:** We initiate the process by leveraging a powerful LLM-DeepSeek-V3, to author a large set of workflows. For each workflow, the model is provided with a subset of tools from our predefined library \mathcal{T} and prompted to generate a structurally complex and logically consistent execution plan in the form

of a DAG, G_{orig} . This LLM-driven approach produces the diverse, high-quality ground-truth solutions for our planning tasks.

- **Query Reverse-Engineering:** Next, we employ a powerful teacher LLM ($M_{teacher}$)-DeepSeek-V3 to reverse-engineer a natural language query Q for each generated workflow G_{orig} . The model’s objective is to capture the user’s intent behind the workflow, effectively transforming the difficult problem of plan generation into a more manageable text-to-text task.
- **Intent Analysis and Re-planning:** The final stage acts as a crucial quality filter. The reverse-engineered query Q might be ambiguous or fail to fully capture the intent of the original plan G_{orig} . Therefore, we use the same teacher model $M_{teacher}$, now acting as an expert planner, to re-solve the problem based solely on Q . This ensures that the final DAG plan, G_{final} , is a faithful and optimal solution derivable from the query, thus guaranteeing the high fidelity of our resulting (Q, G_{final}) training pairs in D_{train} .

Model Training

Our training process commences with Supervised Fine-Tuning (SFT) on a suite of Qwen3 models—specifically, the 0.6B, 1.7B, 4B, and 8B variants. This initial SFT phase serves to provide an effective initialization (i.e., a cold start), which is subsequently followed by a GRPO training phase for further policy refinement.

Base Models and Supervised Fine-Tuning (SFT) We first perform Supervised Fine-Tuning (SFT) on our dataset D_{train} to initialize the model. The objective is to learn to generate the ground-truth plan G_{gt} by minimizing the Negative Log-Likelihood (NLL) loss:

$$\mathcal{L}_{SFT}(\theta) = -\mathbb{E}_{(Q, G_{gt}) \sim D_{train}} [\log P(G_{gt} | Q, T; \theta)] \quad (2)$$

RL Training Set Curation To ensure a stable and efficient RL phase, we curate the training data to focus on problems at the frontier of the model’s current capabilities (Chen et al. 2024). We use the SFT-trained model as a filter: tasks that the model already consistently solves (offering no learning signal) or consistently fails (intractably hard) are excluded from the RL training set. This curation of high-variance instances, where the outcome is uncertain, concentrates the training on the most informative samples, preventing policy degradation and fostering robust optimization.

Reinforcement Learning with Hierarchical Rewards

To move beyond SFT’s limitations, we use Reinforcement Learning (RL). This, however, requires a nuanced reward signal to properly evaluate complex DAGs, distinguishing structural from strategic errors. We address this by introducing a Hierarchical Reward Function $R(y)$. To optimize this complex, multi-dimensional reward, we utilize the Group Relative Policy Optimization (GRPO) algorithm (Shao et al. 2024), which is specifically designed for such scenarios and ensures robust policy improvement.

Our reward function, $R(y)$, where y represents the generated plan, evaluates plan quality through a fail-fast hierar-

chical process. It prioritizes structural correctness, with critical errors incurring large negative penalties and terminating evaluation.

- **Level 1 and 2: Structural and Semantic Penalties.** We first apply penalties for critical errors that render a plan invalid. This includes ill-formed syntax (e.g., non-JSON) or the presence of a **cycle**, both resulting in a large penalty (-10.0). A less severe penalty (-2.0) is given for a lack of **connectivity** (e.g., isolated nodes), which indicates a semantic flaw.
- **Level 3: Planning Fidelity Rewards.** If a plan passes all structural checks, it receives a positive reward based on its quality. This reward consists of two parts: (i) an **Edge F1 Score Reward** (R_{F1}), calculated as 5 times the edge-level F1 score against the ground truth to reward partial structural correctness, and (ii) a large **Perfect Match Bonus** (R_{bonus}) of +5.0 if the plan is identical to the ground truth, incentivizing perfect solutions.

The final reward, $R(y)$, is computed based on these components. If any penalty is applied, the reward is the value of that penalty, checked in a specific order of precedence. Otherwise, it is the sum of the fidelity rewards:

$$R(y) = \begin{cases} P_{syntax}(y) & \text{if invalid syntax} \\ P_{cycle}(y) & \text{else if has cycle} \\ P_{connectivity}(y) & \text{else if disconnected} \\ R_{F1}(y) + R_{bonus}(y) & \text{otherwise} \end{cases} \quad (3)$$

The domain of this reward function, $[-10.0, 10.0]$, provides a rich, multi-faceted learning signal that guides the model not only on *whether* its plan is correct, but also on the *nature* and *severity* of its errors.

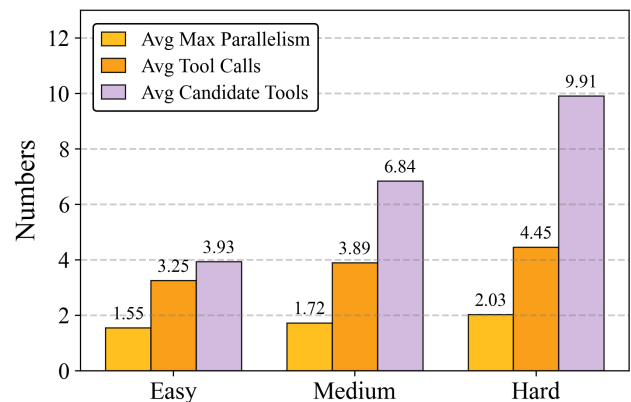


Figure 3: This chart shows our three task difficulties: Easy, Medium, and Hard. Harder tasks have more available tools to choose from (blue) and also require more tools to be used (orange).

Experimental Setup

Dataset

We conduct a comprehensive evaluation of our framework from two primary perspectives: (1) the quality of the gen-

erated plans, and (2) the end-to-end task execution success rate. For the former, we use our newly constructed ComplexTool-Plan benchmark. For the latter, we evaluate our system on the established StableToolBench benchmark.

ComplexTool-Plan Benchmark We evaluate our Planner on ComplexTool-Plan, our benchmark constructed from a comprehensive library of 4,535 tool APIs sourced from ModelScope (Li et al. 2023a). The benchmark features three difficulty levels (Easy, Medium, and Hard) and provides a curated SFT dataset of 3,000 instances. For RL training, the dataset is filtered to comprise 787 instances, which are subsequently divided into training and test sets using an 8:2 ratio. Further details are available in Fig. 3.

End-to-End Execution Benchmark To evaluate the practical effectiveness of our framework, we use StableToolBench (Guo et al. 2024) for end-to-end evaluation. StableToolBench is an enhanced version of the test set from ToolBench (Qin et al. 2023), specifically designed to mitigate execution instability via a caching system and API simulators. It categorizes test cases across two axes: tool generalization and scenario complexity.

- **Tool Generalization: Inst.** (unseen instructions for seen tools), **Tool** (unseen tools in seen categories), and **Cat.** (unseen tool categories).
- **Scenario Complexity: I1** (single-tool), **I2** (multi-tool in-category), and **I3** (multi-tool in-collection).

Task difficulty escalates from I1-Inst. to I3-Cat., providing a comprehensive testbed for agentic capabilities.

In our end-to-end evaluation on StableToolBench, we pair our Planner with GPT-4o as the executor.

Baselines

To ensure a comprehensive evaluation, we benchmark our framework against two categories of state-of-the-art baselines:

- **Proprietary Planners:** For direct planning evaluation on our ComplexTool-Plan benchmark, we compare against powerful proprietary models prompted to generate DAG plans directly. These include GPT-4o, Claude-3.7, DeepSeek-V3, and Ernie-X1, which serve as a practical upper bound for the task.
- **Open-Source Frameworks:** For end-to-end evaluation on StableToolBench (Guo et al. 2024), we compare against established methods applied to the ToolLLaMA model (Qin et al. 2023). This includes the reactive framework ReAct, the planning-based method DFSDT, and a highly relevant concurrent work, LLMCompiler (Kim et al. 2024).

Evaluation Metrics

Planning Quality Metrics To comprehensively evaluate planning quality, we assess three key aspects: the correctness of tool selection (node-level), the accuracy of inferred dependencies (edge-level), and the holistic accuracy of the entire plan. Our evaluation is conducted on a test set comprising 500 samples from the Easy split and 500 from the Hard

split. For each generated Directed Acyclic Graph (DAG), we compute the following metrics against the ground-truth plan, with all final scores macro-averaged.

Let S_{pred} be the set of predicted items (nodes or edges) and S_{true} be the set of ground-truth items.

- **Node-level Performance (P/R/F1):** Measures the accuracy of tool selection. We compute Precision, Recall, and F1-score between the set of predicted nodes (N_{pred}) and ground-truth nodes (N_{true}).

- **Precision (P):** Measures fidelity by calculating the fraction of predicted tools that are correct.

$$P = \frac{|N_{pred} \cap N_{true}|}{|N_{pred}|}$$

- **Recall (R):** Measures completeness by calculating the fraction of required tools that were identified.

$$R = \frac{|N_{pred} \cap N_{true}|}{|N_{true}|}$$

- **F1-Score:** The harmonic mean of Precision and Recall.

$$F1 = 2 \cdot \frac{P \cdot R}{P + R}$$

- **Edge-level Performance (P/R/F1):** Measures the structural correctness of the plan. We apply the same P/R/F1 metrics to compare the set of predicted edges (E_{pred}) against the ground-truth edges (E_{true}). This assesses the model’s ability to infer the correct execution flow.

- **DAG Exact Match (EM):** Our most stringent metric, measuring the holistic correctness of the plan. A plan receives a score of 1 if and only if both its node set and edge set perfectly match the ground truth ($N_{pred} = N_{true}$ and $E_{pred} = E_{true}$), and 0 otherwise.

End-to-End Evaluation Metrics To evaluate the end-to-end task-solving capabilities of our framework, we adopt the two primary metrics established by the StableToolBench benchmark. The first, Solvable Pass Rate (SoPR), serves as the primary measure of absolute capability and is defined as the percentage of tasks that an agent successfully completes according to the benchmark’s ground-truth criteria. The second, Solvable Win Rate (SoWR), is a comparative metric that evaluates performance relative to a strong baseline. Specifically, it measures the win rate of our agent against the benchmark’s provided GPT-3.5 (ReAct) baseline, offering a clear signal of relative improvement. To ensure statistical robustness and mitigate performance fluctuations, all reported scores are the average of three independent runs for each experimental setting.

Results and Discussion

RQ1: How does Reinforcement Learning enhance performance over Supervised Fine-Tuning?

Our results demonstrate that the RL stage systematically improves holistic plan accuracy over SFT alone, as measured by our most direct metric: DAG Exact Match Accuracy. As shown in Tab. 1, on the Easy set, our Qwen3-8B model’s

Method	Node-level Performance			Edge-level Performance			DAG Exact Match Acc.
	Precision	Recall	F1-Score	Precision	Recall	F1-Score	
GPT-4o	0.916	0.943	0.929	0.765	0.793	0.779	0.635
Deepseek-V3	0.777	0.763	0.770	0.641	0.645	0.643	0.511
Claude-3.7	0.939	0.959	0.949	0.801	0.829	0.815	0.644
Ernie-X1	0.925	0.973	0.948	0.716	0.744	0.730	0.562
Qwen3-0.6B (SFT)	0.951	0.986	0.968	0.826	0.872	0.848	0.671
Qwen3-1.7B (SFT)	0.960	0.987	0.973	0.856	0.890	0.873	0.723
Qwen3-1.7B (SFT+RL)	0.974	0.985	0.979	0.874	0.885	0.879	0.756
Qwen3-4B (SFT)	0.972	0.986	0.979	0.882	0.895	0.888	0.768
Qwen3-4B (SFT+RL)	0.978	0.988	0.983	0.887	0.898	0.892	0.778
Qwen3-8B (SFT)	0.974	0.992	0.983	0.889	0.909	0.899	0.781
Qwen3-8B (SFT+RL)	0.982	0.987	0.984	0.905	0.907	0.906	0.803

Table 1: Results of our Planner models and proprietary baselines on the Easy evaluation set. Metrics include node-level, edge-level, and DAG Exact Match accuracy. Our fine-tuned models show superior performance compared to the baselines, with performance scaling positively with model size and the addition of RL. The Qwen3-0.6B(SFT+RL) variant is excluded due to training instability.

Method	Node-level Performance			Edge-level Performance			DAG Exact Match Acc.
	Precision	Recall	F1-Score	Precision	Recall	F1-Score	
GPT-4o	0.850	0.863	0.856	0.447	0.482	0.464	0.098
Deepseek-V3	0.842	0.787	0.814	0.414	0.415	0.414	0.082
Claude-3.7	0.884	0.910	0.897	0.476	0.508	0.491	0.106
Ernie-X1	0.862	0.844	0.853	0.400	0.411	0.405	0.052
Qwen3-0.6B (SFT)	0.891	0.866	0.878	0.561	0.541	0.551	0.160
Qwen3-1.7B (SFT)	0.899	0.880	0.889	0.599	0.591	0.595	0.213
Qwen3-1.7B (SFT+RL)	0.914	0.861	0.887	0.616	0.579	0.597	0.218
Qwen3-4B (SFT)	0.918	0.894	0.906	0.632	0.623	0.627	0.241
Qwen3-4B (SFT+RL)	0.918	0.876	0.896	0.631	0.611	0.621	0.259
Qwen3-8B (SFT)	0.925	0.896	0.910	0.661	0.654	0.657	0.295
Qwen3-8B (SFT+RL)	0.937	0.873	0.904	0.678	0.641	0.659	0.319

Table 2: Results of our Planner models and proprietary baselines on the Hard evaluation set. Our fine-tuned models show superior performance compared to the baselines, with performance scaling positively with model size and the addition of RL. The Qwen3-0.6B(SFT+RL) variant is excluded due to training instability.

accuracy increases from **0.781** (SFT) to **0.803** (SFT+RL). This trend is consistent across all model sizes and becomes even more critical as task complexity increases. Notably, the training instability of our smallest model (Qwen3-0.6B) suggests this trend has a lower bound: sufficient model capacity is crucial to prevent reward hacking, where a model learns a degenerate, low-effort policy simply to avoid penalties rather than solving the task.

On the challenging **Hard set** (Tab. 2), the impact of RL is more pronounced. The Qwen3-8B model’s DAG Exact Match accuracy jumps from **0.295** (SFT) to **0.319** (SFT+RL), an **8.1%** relative improvement. This shows RL is not merely polishing plans but is essential for correcting subtle structural errors that emerge in complex scenarios, a conclusion supported by consistent gains in Edge-level F1-Scores. Essentially, while SFT effectively teaches the model to select the right tools (nodes), RL is crucial for orchestrating them correctly (edges). By optimizing for structural

integrity, RL moves the model beyond plausible mimicry to generating functionally correct plans.

RQ2: How does model scaling affect performance robustness as task complexity increases?

Our experiments reveal a clear, positive relationship between model size and planning capability, with larger models demonstrating significantly greater robustness against increasing task complexity.

First, our framework shows strong scalability. This is most evident on the challenging Hard set, where even top-tier models like GPT-4o falter. On this set, the DAG Exact Match Accuracy of our SFT+RL models rises monotonically with scale: from **0.218** (1.7B) to **0.259** (4B) and finally **0.319** (8B). This stark performance gap over general-purpose models validates our specialized training approach.

Second, and more importantly, model scaling directly enhances robustness. While all models degrade on harder

Method	I1-Inst.		I1-Tool		I1-Cat.		I2-Inst.		I2-Cat.		I3-Inst.		Average	
	SoPR	SoWR	SoPR	SoWR	SoPR	SoWR	SoPR	SoWR	SoPR	SoWR	SoPR	SoWR	SoPR	SoWR
GPT-series														
GPT-3.5 (ReAct)	53.0	–	53.0	–	51.2	–	37.6	–	43.9	–	48.6	–	47.9	–
GPT-3.5 (DFSDT)	63.8	58.9	73.9	65.8	65.8	60.1	57.1	72.6	69.8	68.5	69.9	67.2	66.7	65.5
GPT-4 (ReAct)	54.4	53.4	44.1	60.1	48.8	52.9	50.6	69.8	48.9	62.1	42.6	54.1	48.2	58.7
GPT-4 (DFSDT)	69.0	57.1	69.6	66.5	68.1	61.4	70.8	73.6	68.0	62.9	76.0	63.9	70.3	64.2
Open-source														
ToolLLaMA (ReAct)	42.7	36.2	35.4	36.1	38.6	34.6	39.9	49.1	40.9	38.7	29.8	41.0	37.9	39.3
ToolLLaMA† (ReAct)	26.7	22.1	25.0	27.2	31.7	29.4	23.1	32.1	24.5	28.2	20.5	24.6	25.3	27.3
ToolLLaMA (DFSDT)	56.6	39.9	55.5	46.8	56.5	41.8	49.7	53.8	53.4	49.2	53.6	50.8	54.2	47.1
ToolLLaMA† (DFSDT)	41.8	35.6	39.9	37.3	44.9	39.9	36.0	47.2	39.1	39.5	33.3	26.2	39.2	37.6
LLMCompiler	39.2	35.6	35.1	36.0	39.8	35.3	37.5	45.6	38.4	38.1	27.0	36.5	36.2	37.9
Qwen3-1.7B (RL)	61.4	52.1	51.7	42.5	58.1	49.3	48.7	45.8	46.8	48.2	44.7	40.5	51.9	46.4
Qwen3-4B (RL)	62.1	54.6	55.4	47.1	61.4	53.5	59.2	55.7	48.7	51.3	50.8	48.6	56.3	51.8
Qwen3-8B (RL)	67.0	57.9	56.6	49.8	59.3	51.7	60.1	58.2	54.7	56.4	61.0	55.9	59.8	55.0

Table 3: End-to-end task-solving results on the StableToolBench benchmark. The table compares our Qwen3 models with GPT-series and other open-source methods across various task complexities. Our Qwen3 series, particularly the 8B model, demonstrates the best performance among all open-source baselines in both absolute success (SoPR) and win rate against GPT-3.5(ReAct) (SoWR).

tasks, larger models exhibit a more graceful decline. For instance, when moving from the Easy to the Hard set, the accuracy of our Qwen3-1.7B model plummets by **71.2%** (from 0.756 to 0.218). In contrast, our largest Qwen3-8B model sees a more contained **60.3%** drop (from 0.803 to 0.319). This widening performance gap is crucial, demonstrating that scaling up our specialized Planner not only boosts absolute performance but also enhances its **resilience** against complex, real-world challenges.

RQ3: What is the end-to-end effectiveness and efficiency of our framework?

Our framework demonstrates highly competitive end-to-end performance. As shown in Tab. 3, our Qwen3-8B (RL) model achieves an average SoPR of **59.8%**, substantially outperforming the widely-used reactive approach of GPT-4 (ReAct) at 48.2%. This highlights that a superior planning strategy can be more impactful than raw model capability, unlocking the potential of smaller models.

While iterative paradigms like DTA-Llama (Zhu et al. 2025) report a higher SoPR, their success is largely driven by a multi-turn framework that allows for in-process error correction and adaptation. In contrast, our non-iterative, plan-then-execute paradigm has only one chance to generate a correct global plan upfront. This architectural choice prioritizes predictability and planning quality over the adaptive resilience of iterative feedback loops.

Crucially, this design choice leads to state-of-the-art efficiency. As detailed in Tab. 4, our method requires an average of just **2.29 inference steps** to complete a task, significantly fewer than all other methods, including iterative parallel frameworks like DTA-Llama (2.48 steps). This demonstrates that our global planning approach is not only effective but also architecturally more efficient, solving complex tasks with fewer high-level decision rounds.

Method	I1-I.	I1-T	I1-C.	I2-I.	I2-C.	I3-I.
GPT-series						
GPT-3.5 (ReAct)	4.28	4.75	4.48	5.16	5.05	5.31
GPT-3.5 (DFSDT)	11.60	13.36	11.77	16.60	14.06	12.54
GPT-3.5 (Parallel)	25.33	28.06	26.12	31.79	31.04	38.10
GPT-4 (ReAct)	3.27	3.64	3.87	4.04	4.19	4.23
GPT-4 (DFSDT)	5.90	8.09	6.67	9.97	18.13	14.05
GPT-4 (Parallel)	4.66	9.18	12.90	3.63	5.98	10.38
Open-source						
ToolLLaMA (ReAct)	3.42	3.47	3.50	3.67	3.63	3.64
ToolLLaMA (DFSDT)	8.09	8.51	8.10	10.20	9.93	9.23
LLMCompiler	5.48	5.56	6.07	5.36	5.68	5.62
Qwen2.5 (Parallel)	9.07	9.47	12.01	14.58	14.56	12.38
DTA-Llama	2.41	2.41	2.51	2.32	2.34	2.48
Qwen3-8B (RL)	2.26	2.33	2.29	2.16	2.29	2.41

Table 4: Inference steps for different methods.

Conclusion

In this work, we introduced a Planner-centric framework that decouples planning from execution to address the local optimization traps of reactive tool-use agents. Our core contribution is a Planner model trained via a two-stage SFT-GRPO strategy on our new **ComplexTool-Plan** benchmark. It generates a global, parallelizable DAG plan in a single forward pass.

Our experiments demonstrate that this paradigm is highly effective. The Planner itself produces higher-quality plans than strong proprietary models. When integrated with an executor, our framework achieves state-of-the-art results for open-source models on the challenging **StableToolBench** benchmark, while requiring the fewest inference steps.

Crucially, our findings show that a sophisticated planning strategy can be more impactful than raw model capability, offering a more scalable and predictable path toward capable autonomous agents.

References

- Chen, W.; Ma, X.; Wang, X.; and Cohen, W. W. 2022. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *arXiv preprint arXiv:2211.12588*.
- Chen, X.; Li, Y.; Cai, H.; Ma, Z.; Chen, X.; Xiong, H.; Wang, S.; He, B.; Sun, L.; and Yin, D. 2025. Multi-Agent Proactive Information Seeking with Adaptive LLM Orchestration for Non-Factoid Question Answering. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, 4341–4352.
- Chen, Z.; Deng, Y.; Yuan, H.; Ji, K.; and Gu, Q. 2024. Self-Play Fine-Tuning Converts Weak Language Models to Strong Language Models. In *International Conference on Machine Learning*, 6621–6642. PMLR.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv e-prints*, arXiv–2407.
- Feng, J.; Huang, S.; Qu, X.; Zhang, G.; Qin, Y.; Zhong, B.; Jiang, C.; Chi, J.; and Zhong, W. 2025. Retool: Reinforcement learning for strategic tool use in llms. *arXiv preprint arXiv:2504.11536*.
- Gao, L.; Madaan, A.; Zhou, S.; Alon, U.; Liu, P.; Yang, Y.; Callan, J.; and Neubig, G. 2023. Pal: Program-aided language models. In *International Conference on Machine Learning*, 10764–10799. PMLR.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Guo, Z.; Cheng, S.; Wang, H.; Liang, S.; Qin, Y.; Li, P.; Liu, Z.; Sun, M.; and Liu, Y. 2024. StableToolBench: Towards Stable Large-Scale Benchmarking on Tool Learning of Large Language Models. In *Findings of the Association for Computational Linguistics ACL 2024*, 11143–11156.
- Hong, S.; Zhuge, M.; Chen, J.; Zheng, X.; Cheng, Y.; Wang, J.; Zhang, C.; Wang, Z.; Yau, S. K. S.; Lin, Z.; et al. 2023. MetaGPT: Meta programming for a multi-agent collaborative framework. In *The Twelfth International Conference on Learning Representations*.
- Hong, W.; Wang, W.; Lv, Q.; Xu, J.; Yu, W.; Ji, J.; Wang, Y.; Wang, Z.; Dong, Y.; Ding, M.; et al. 2024. Cogagent: A visual language model for gui agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14281–14290.
- Jin, B.; Zeng, H.; Yue, Z.; Yoon, J.; Arik, S.; Wang, D.; Zamani, H.; and Han, J. 2025. Search-r1: Training llms to reason and leverage search engines with reinforcement learning. *arXiv preprint arXiv:2503.09516*.
- Kim, S.; Moon, S.; Tabrizi, R.; Lee, N.; Mahoney, M. W.; Keutzer, K.; and Gholami, A. 2024. An llm compiler for parallel function calling. In *Forty-first International Conference on Machine Learning*.
- Li, C.; Chen, H.; Yan, M.; Shen, W.; Xu, H.; Wu, Z.; Zhang, Z.; Zhou, W.; Chen, Y.; Cheng, C.; et al. 2023a. Modelscope-agent: Building your customizable agent system with open-source large language models. *arXiv preprint arXiv:2309.00986*.
- Li, M.; Zhao, Y.; Yu, B.; Song, F.; Li, H.; Yu, H.; Li, Z.; Huang, F.; and Li, Y. 2023b. Api-bank: A comprehensive benchmark for tool-augmented llms. *arXiv preprint arXiv:2304.08244*.
- Li, Y.; Cai, H.; Kong, R.; Chen, X.; Chen, J.; Yang, J.; Zhang, H.; Li, J.; Wu, J.; Chen, Y.; et al. 2025a. Towards AI Search Paradigm. *arXiv preprint arXiv:2506.17188*.
- Li, Y.; Xiong, H.; Zhang, Y.; Bian, J.; Peng, T.; Li, X.; Wang, S.; Kong, L.; and Yin, D. 2025b. Rankelectra: Semi-supervised pre-training of learning-to-rank electra for web-scale search. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, 2415–2425.
- Li, Y.; Zhang, H.; Zhang, Y.; Cai, H.; Cai, M.; Wang, S.; Xiong, H.; Kong, L.; Yin, D.; and Chen, L. 2025c. Rankexpert: A mixture of textual-and-behavioral experts for multi-objective learning-to-rank in web search. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, 4578–4589.
- Liu, X.; Yu, H.; Zhang, H.; Xu, Y.; Lei, X.; Lai, H.; Gu, Y.; Ding, H.; Men, K.; Yang, K.; et al. 2023. Agentbench: Evaluating llms as agents. *arXiv preprint arXiv:2308.03688*.
- Lu, W.; Tong, Y.; and Ye, Z. 2025. DAMMFND: Domain-Aware Multimodal Multi-view Fake News Detection. In *Proceedings of the AAI Conference on Artificial Intelligence*, volume 39, 559–567.
- Mialon, G.; Fourrier, C.; Wolf, T.; LeCun, Y.; and Scialom, T. 2023. Gaia: a benchmark for general ai assistants. In *The Twelfth International Conference on Learning Representations*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Patil, S. G.; Zhang, T.; Wang, X.; and Gonzalez, J. E. 2024. Gorilla: Large language model connected with massive apis. *Advances in Neural Information Processing Systems*, 37: 126544–126565.
- Qin, Y.; Liang, S.; Ye, Y.; Zhu, K.; Yan, L.; Lu, Y.; Lin, Y.; Cong, X.; Tang, X.; Qian, B.; et al. 2023. Toolllm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*.
- Ridnik, T.; Kredo, D.; and Friedman, I. 2024. Code generation with alphacodium: From prompt engineering to flow engineering. *arXiv preprint arXiv:2401.08500*.
- Schick, T.; Dwivedi-Yu, J.; Dessì, R.; Raileanu, R.; Lomeli, M.; Hambro, E.; Zettlemoyer, L.; Cancedda, N.; and Scialom, T. 2023. Toolformer: Language models can teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36: 68539–68551.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath:

Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.

Shen, Y.; Song, K.; Tan, X.; Li, D.; Lu, W.; and Zhuang, Y. 2023. Hugginggpt: Solving ai tasks with chatgpt and its friends in hugging face. *Advances in Neural Information Processing Systems*, 36: 38154–38180.

Shinn, N.; Cassano, F.; Gopinath, A.; Narasimhan, K.; and Yao, S. 2023. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36: 8634–8652.

Shridhar, M.; Yuan, X.; Côté, M.-A.; Bisk, Y.; Trischler, A.; and Hausknecht, M. 2020. Alfworld: Aligning text and embodied environments for interactive learning. *arXiv preprint arXiv:2010.03768*.

Tong, Y.; Lu, W.; Cui, X.; Mao, Y.; and Zhao, Z. 2025. DAPT: Domain-Aware Prompt-Tuning for Multimodal Fake News Detection. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 7902–7911.

Wang, L.; Xu, W.; Lan, Y.; Hu, Z.; Lan, Y.; Lee, R. K.-W.; and Lim, E.-P. 2023. Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language models. *arXiv preprint arXiv:2305.04091*.

Wei, X.; Lu, B.; Zhang, X.; Zhao, Z.; Shen, D.; Xia, L.; and Yin, D. 2025. Igniting Creative Writing in Small Language Models: LLM-as-a-Judge versus Multi-Agent Refined Rewards. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 17171–17197.

Yang, J.; Jimenez, C. E.; Wettig, A.; Lieret, K.; Yao, S.; Narasimhan, K.; and Press, O. 2024. Swe-agent: Agent-computer interfaces enable automated software engineering. *Advances in Neural Information Processing Systems*, 37: 50528–50652.

Yao, S.; Chen, H.; Yang, J.; and Narasimhan, K. 2022. Webshop: Towards scalable real-world web interaction with grounded language agents. *Advances in Neural Information Processing Systems*, 35: 20744–20757.

Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T.; Cao, Y.; and Narasimhan, K. 2023a. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36: 11809–11822.

Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K.; and Cao, Y. 2023b. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*.

Zhao, Z.; Dong, Y.; Liu, A.; Zheng, L.; Liu, P.; Shen, D.; Xia, L.; Zhao, J.; and Yin, D. 2025. TURA: Tool-Augmented Unified Retrieval Agent for AI Search. *arXiv:2508.04604*.

Zhou, S.; Xu, F. F.; Zhu, H.; Zhou, X.; Lo, R.; Sridhar, A.; Cheng, X.; Ou, T.; Bisk, Y.; Fried, D.; et al. 2023. Webarena: A realistic web environment for building autonomous agents. *arXiv preprint arXiv:2307.13854*.

Zhu, D.; Shi, W.; Shi, Z.; Ren, Z.; Wang, S.; Yan, L.; and Yin, D. 2025. Divide-Then-Aggregate: An Efficient Tool Learning Method via Parallel Tool Invocation. *arXiv preprint arXiv:2501.12432*.