
FlowSteer: Interactive Agentic Workflow Orchestration via End-to-End Reinforcement Learning

Anonymous Authors¹

Abstract

In recent years, a variety of powerful agentic workflows have been applied to solve a wide range of human problems. However, existing workflow orchestration still faces key challenges, including high manual cost, reliance on specific operators/large language models (LLMs), and sparse reward signals. To address these challenges, we propose FlowSteer, an end-to-end reinforcement learning framework that takes a lightweight policy model as the agent and an executable canvas environment, automating workflow orchestration through multi-turn interaction. In this process, the policy model analyzes execution states and selects editing actions, while the canvas executes operators and returns feedback for iterative refinement. Moreover, FlowSteer provides a plug-and-play framework that supports diverse operator libraries and interchangeable LLM backends. To effectively train this interaction paradigm, we propose Canvas Workflow Relative Policy Optimization (CWRPO), which introduces diversity-constrained rewards with conditional release to stabilize learning and suppress shortcut behaviors. Experimental results on twelve datasets show that FlowSteer significantly outperforms baselines across various tasks. Our code is available at <https://anonymous.4open.science/r/FlowSteer-9B2E>.

1. Introduction

In recent years, a variety of powerful agentic systems have been applied to solve a wide range of human problems (Ruan et al., 2023; Shen et al., 2023; Hong et al., 2024), gradually moving beyond single-turn question answering (QA) toward executable end-to-end task completion. In this

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

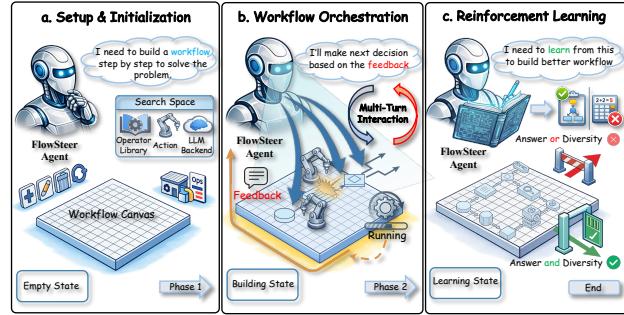


Figure 1. Overview of the FlowSteer framework pipeline. The agent first initializes with the task and explores the search space. Then, through multi-turn interaction with the canvas, it analyzes workflow states, selects editing actions, and receives execution feedback to iteratively build and refine the workflow. Finally, the agent learns from diversity-constrained rewards to continuously improve its workflow orchestration strategies across diverse tasks.

process, workflow orchestration has become a key bridge from task goals to reproducible execution: by organizing operators into an executable workflow graph, systems can complete complex tasks with improved controllability, debuggability, and reusability (Zeng et al., 2023; Qian et al., 2024), as shown in Figure 1. However, in practice, workflow construction still heavily relies on manual drag-and-drop and rule-based configuration (Li et al., 2025a), making it costly to transfer across new tasks, new operator libraries, new model backends, or different application domains.

To address these issues, three main paradigms of workflow orchestration have emerged, as shown in Figure 2. First, static workflow selection retrieves pre-defined workflows from a library based on task similarity (Wang et al., 2024). Second, offline workflow generation trains policy models via supervised fine-tuning (SFT) or group relative policy optimization (GRPO) (Shao et al., 2024a) to generate workflows. Third, automated workflow optimization—such as AFlow (Zhang et al., 2024a), GPTSwarm (Zhuge et al., 2024), and LATS (Zhou et al., 2023)—combines search and execution feedback to iteratively improve workflow structures.

However, these methods still face several challenges: **(i) High manual/heuristic dependence**—rules and templates require continual maintenance and are often tightly coupled to specific scenarios, limiting reuse and generaliza-

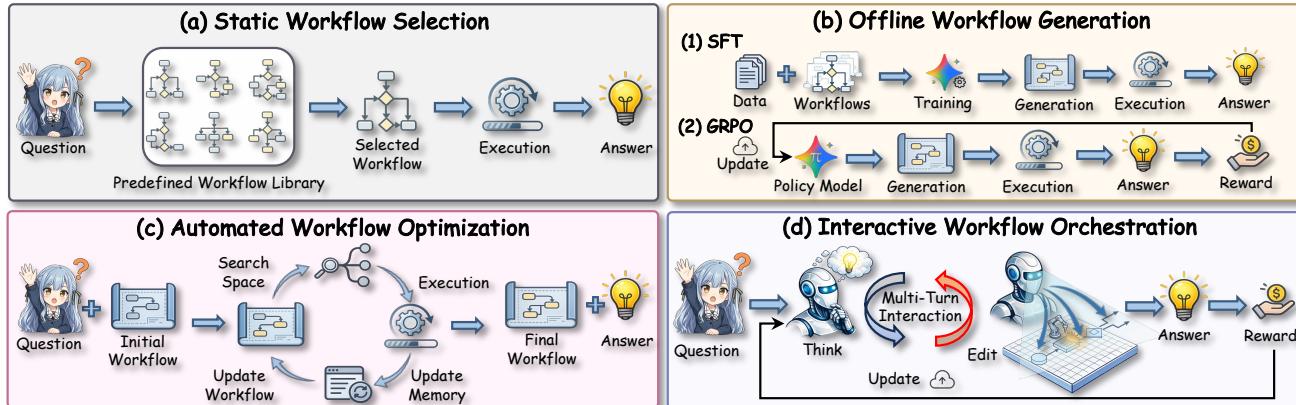


Figure 2. Comparison of different workflow orchestration paradigms: static workflow selection, offline workflow generation, automated workflow optimization, and our interactive workflow orchestration framework FlowSteer.

tion (Schick et al., 2023; Wang et al., 2024); **(ii) Operator/backend lock-in (path lock-in)**—existing approaches tend to rely on fixed operators or a single strong large language model (LLM) backend, making it difficult for a policy model to compose operators in a plug-and-play manner, so performance and robustness drop sharply when the operator library or backend environment changes (Yao et al., 2023; Zhou et al., 2024); **(iii) Sparse and unstable learning signals**—training with only terminal correctness rewards can lead to shortcut behaviors (e.g., premature termination and oversimplified graphs) and reward hacking, and may further suffer from intra-group advantage collapse, making long-horizon credit assignment unstable (Yu et al., 2025; DeepSeek-AI, 2025).

To address these challenges, we propose FlowSteer, an end-to-end reinforcement learning (RL)-enhanced framework for workflow orchestration that supports plug-and-play deployment across different operator libraries and interchangeable LLM backends. In our framework, a lightweight policy model acts as an agent that selects and revises compositions of operators within an executable Workflow Canvas environment. The canvas executes the selected nodes and returns execution traces and feedback, enabling the policy to learn transferable and diverse orchestration strategies from real execution loops. We further design a diversity-constrained reward and a conditionally released answer-based reward to jointly improve workflow structure quality and task correctness. Overall, FlowSteer provides a low-cost, transferable, training-stable, and scalable paradigm for automated orchestration of agentic workflows across diverse tasks.

We evaluate FlowSteer on three task categories: QA, mathematical reasoning, and code generation. Experimental results demonstrate that FlowSteer enhances both generation quality and reasoning accuracy through its RL-driven multi-turn workflow interaction framework, surpassing baselines and existing methods, as shown in Figure 2. Beyond

strengthening the reasoning capabilities of large-scale LLM backends, FlowSteer also improves the performance of smaller-scale LLMs, while reducing token consumption and interaction turns through more efficient orchestration strategies. Furthermore, FlowSteer can adapt across diverse task types without task-specific fine-tuning, demonstrating broad adaptability and strong practical potential.

2. Related Work

Agent Workflows. Before the rise of LLM agents, workflow automation was largely rule- or template-driven, with operators and control flow specified by hand. In the era of LLM agents, agent workflows improve long-horizon reliability via a plan–act–feedback loop (Erdogan et al., 2025; Shang et al., 2025; Hong et al., 2024). Existing studies can be grouped into three lines: single-agent decision making, which models tool use as sequential decisions (Erdogan et al., 2025) or interleaved reasoning and acting (Yang et al., 2025); orchestration, which uses LLM controllers for tool/model routing (Shang et al., 2025) or constrained application programming interface (API) planning to ground intent (Wang et al., 2024); and multi-agent collaboration, which relies on standard operating procedures (SOPs)/roles (Hong et al., 2024) or cross-team orchestration (Du et al., 2025). For sustained capability, agentic tool reasoning (Wu et al., 2025) and reusable workflow memory further improve efficiency (Wang et al., 2024), supporting scalable automation in practice (Qian et al., 2024). In this paper, we propose FlowSteer, an end-to-end RL framework that learns workflow orchestration from executable canvas feedback.

Reinforcement Learning for Agents. With recent progress in RL for LLM-based agents, agent RL models interaction as a long-horizon Markov decision process (MDP) (Zhou et al., 2024). In hierarchical multi-turn RL, coupling value learning with token-level policy learning eases delayed credit

110 *Table 1.* Operator library \mathcal{O} and action space \mathcal{A} in FlowSteer. Each row shows a category of operators, their outputs, action types, and the
 111 corresponding graph update operations.

Category	Operator $o \in \mathcal{O}$	Output	Type	Action	Graph Update
Planning	Plan, Decompose	strategy, sub-problems	Editing	add	$V_t \leftarrow V_{t-1} \cup \{v\}$
Solving	Programmer, Custom, AnswerGen	code, response, answer	Editing	delete	$V_t \leftarrow V_{t-1} \setminus \{v\}$
Verification	Test, Review, Verify	correct, feedback	Editing	modify	$op(v) \leftarrow o'$
Revision	Revise	revised solution	Config	set_prompt	$prompt(v) \leftarrow p$
Ensemble	ScEnsemble, Aggregate	voted, combined	Terminal	finish	$y_q = \text{Execute}(\mathcal{G}_T, q)$
Formatting	Format	final answer y_q	Control	parallel, cond, loop	branch/merge in E_t

120 assignment (Zhou et al., 2024). For tool-chain control, step-
 121 grained shaping (Yu et al., 2025) and outcome feedback im-
 122 prove tool selection and self-correction (Feng et al., 2025). When retrieval is treated as an action, search rollouts learn
 123 think-then-search behaviors (Jin et al., 2025). Moreover,
 124 large-scale RL motivates GRPO-style objectives based on
 125 verifiable or group-relative signals (DeepSeek-AI, 2025;
 126 Shao et al., 2024a).

3. Preliminaries

131 **Definition 1: Workflow Graph.** A workflow graph is
 132 a directed acyclic graph $\mathcal{G} = (V, E, \text{attr})$, where $V =$
 133 $\{v_1, \dots, v_n\}$ is a set of n operator nodes, $E \subseteq V \times V$ en-
 134 codes data dependencies and execution order, and $\text{attr}(v) =$
 135 $(op(v), \text{param}(v), \text{prompt}(v))$ specifies the operator type
 136 from library \mathcal{O} , parameter configuration, and execution
 137 prompt for each node $v \in V$ in the workflow.

138 **Definition 2: Orchestration Trajectory.** An orchestration
 139 trajectory is a complete sequence of T interaction steps that
 140 uniquely determines a workflow graph \mathcal{G}_τ :

$$\tau = \{(a_t^{\text{think}}, a_t, o_t^{\text{exec}})\}_{t=1}^T \Rightarrow \mathcal{G}_\tau, \quad (1)$$

141 where a_t^{think} denotes the reasoning reflection at step t , $a_t =$
 142 $(\alpha_t, a_t^{\text{out}})$ is the editing action with type $\alpha_t \in \mathcal{A}_{\text{type}}$ and
 143 content a_t^{out} , and o_t^{exec} is the execution feedback from the
 144 canvas environment. The complete definitions of operators
 145 and actions are summarized in Table 1.

146 **Problem Statement.** Given a task $q \in \mathcal{D}_Q$, an operator
 147 library \mathcal{O} , and a pluggable backend LLM $\mathcal{M}_{\text{backend}}$, the
 148 workflow orchestration problem aims to learn a policy π_θ
 149 that generates trajectory τ . The corresponding workflow \mathcal{G}_τ
 150 is then executed to produce the answer:

$$y_q = \text{Execute}(\mathcal{G}_\tau, q, \mathcal{M}_{\text{backend}}). \quad (2)$$

151 The overall learning objective is to maximize expected re-
 152 ward over the task distribution:

$$\max_{\theta} \mathcal{J}(\theta) = \mathbb{E}_{q \sim \mathcal{D}_Q} \mathbb{E}_{\tau \sim P_\theta(\cdot | q)} [R(\tau)], \quad (3)$$

153 where $P_\theta(\tau | q)$ is the trajectory distribution induced by
 154 policy π_θ , and $R(\tau)$ is the trajectory-level reward measuring
 155 both structural quality and answer correctness.

4. Methodology: FlowSteer

156 As illustrated in Figure 3, this section introduces the Flow-
 157 Steer framework, including Workflow Canvas and agent
 158 initialization (Section 4.1), workflow orchestration via multi-
 159 turn interaction (Section 4.2), and outcome-directed end-to-
 160 end reinforcement learning (Section 4.3).

4.1. Workflow Canvas and Agent Initialization

161 FlowSteer follows a ReAct-based agent paradigm (Yao et al.,
 162 2023), where a lightweight policy model (Flow-Director)
 163 interacts with an executable canvas environment (Workflow
 164 Canvas) to construct the workflow graph \mathcal{G} (Definition 1). We define the canvas environment, operator library, action
 165 space, state space, and orchestration target.

166 **Workflow Canvas \mathcal{C} .** The Workflow Canvas is the envi-
 167 ronment that maintains the workflow graph state \mathcal{G}_t and
 168 provides executable feedback at each orchestration step:

$$\mathcal{C} = (\mathcal{G}_t, \mathcal{O}, \mathcal{M}_{\text{backend}}, d^{\text{lib}}), \quad (4)$$

169 where $\mathcal{G}_t = (V_t, E_t, \text{attr})$ is the workflow graph at step t ,
 170 $\mathcal{O} = \{o_1, \dots, o_K\}$ is the operator library with K operators,
 171 $\mathcal{M}_{\text{backend}}$ is the pluggable LLM backend, and d^{lib} is the
 172 operator library description that enables the policy to learn
 173 available operators.

174 **Operator Library \mathcal{O} and Action Space \mathcal{A} .** As summarized
 175 in Table 1, we employ 12 functional operators organized
 176 into six categories (planning, solving, verification, revision,
 177 ensemble, and formatting), along with control structures
 178 (parallel, conditional, loop). At each step t , the
 179 agent generates a reflection a_t^{think} that analyzes the current
 180 state, followed by an editing action $a_t = (\alpha_t, a_t^{\text{out}})$, where
 181 the action type $\alpha_t \in \mathcal{A}_{\text{type}}$ covers node insertion, removal,
 182 modification, prompt configuration, and termination. Or-
 183 chestration terminates when $\alpha_t = \text{finish}$.

184 **State Space \mathcal{S} .** The agent state $H_t \in \mathcal{S}$ is defined by
 185 its complete interaction history. Given task q , operator
 186 description d^{lib} , and template a^{tmpl} , the initial state is $H_0 =$
 187 $[q \oplus d^{\text{lib}} \oplus a^{\text{tmpl}}]$, where \oplus denotes sequence concatenation.
 188 The state updates as:

$$H_t = H_{t-1} \oplus (a_t^{\text{think}}, a_t, o_t^{\text{exec}}). \quad (5)$$

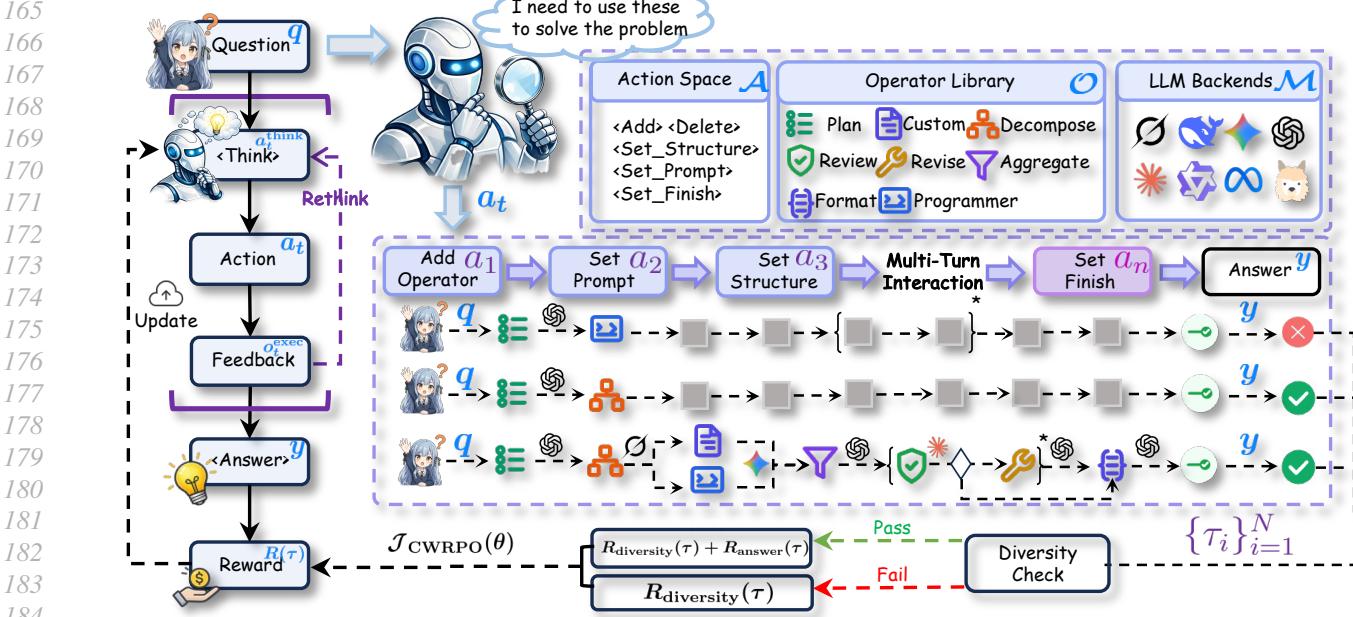


Figure 3. An overview of the FlowSteer framework. The policy model (Flow-Director) interacts with Workflow Canvas through multi-turn interactions and learns from diversity-constrained rewards via CWRPO.

Orchestration Target. The agent interacts with canvas \mathcal{C} until reaching `finish` or maximum turns T_{\max} . Following Definition 2, the complete trajectory τ determines workflow \mathcal{G}_τ , and the answer is obtained via execution (Eq. 2).

Proposition 1. *The operator-action space $(\mathcal{O}, \mathcal{A})$ covers diverse workflow patterns across task types through canvas-grounded orchestration.*

Proof. We provide experimental results in Section 5.2 and Section 5.3, and theoretical proofs in Appendix B.1. \square

4.2. Workflow Orchestration via Multi-Turn Interaction

Building on the canvas \mathcal{C} and state H_0 defined above, we model workflow orchestration as a multi-turn interaction between policy π_θ (Flow-Director) and canvas \mathcal{C} (Workflow Canvas) to iteratively construct the workflow. The interaction follows the prompt template shown in Table 2.

Modeling the Step-wise Orchestration Policy. Each turn submits one atomic editing action, decomposing long-horizon planning into checkable and repairable local decisions. At each step t , Flow-Director outputs reflection a_t^{think} and editing action a_t , modeled as a hierarchical policy conditioned on history H_{t-1} :

$$\begin{aligned} \pi_\theta(a_t^{\text{think}}, a_t | H_{t-1}) &= \pi_\theta(a_t^{\text{think}} | H_{t-1}) \\ &\cdot \pi_\theta(a_t | a_t^{\text{think}}, H_{t-1}) \cdot \pi_\theta(a_t^{\text{out}} | a_t, a_t^{\text{think}}, H_{t-1}). \end{aligned} \quad (6)$$

Canvas-Feedback-Driven Workflow Interaction. Given action a_t from Flow-Director, the Workflow Canvas ex-

ecutes the action and returns feedback o_t^{exec} , forming an iterative closed-loop of “diagnose-edit-verify”. This interaction comprises three stages: (i) *Action Execution and Feedback Generation*. The canvas performs syntax parsing and constraint checking, then generates feedback o_t^{exec} including action success status, failure reasons, and repair suggestions:

$$o_t^{\text{exec}} \sim \mathcal{C}_{\text{exec}}(\cdot | \mathcal{G}_{t-1}, a_t), \quad (7)$$

where $\mathcal{C}_{\text{exec}}$ denotes the canvas feedback distribution. (ii) *State Update and History Accumulation*. The canvas updates the workflow graph (see Table 1 for update operations) and the state evolves via Eq. 5:

$$\mathcal{G}_t = \text{Update}(\mathcal{G}_{t-1}, a_t, o_t^{\text{exec}}). \quad (8)$$

(iii) *Iterative Correction until Termination*. Flow-Director continues decision-making based on H_t until $\alpha_t = \text{finish}$ or $t = T_{\max}$, enabling error discovery and repair without manual rules. **Trajectory Distribution and Optimization.** Following Definition 2, a complete trajectory τ records the interaction from initialization to termination. The trajectory distribution is jointly determined by the policy and environment:

$$P_\theta(\tau) = \prod_{t=1}^T \left[\pi_\theta(a_t^{\text{think}}, a_t | H_{t-1}) \cdot \mathcal{C}_{\text{exec}}(o_t^{\text{exec}} | \mathcal{G}_{t-1}, a_t) \right], \quad (9)$$

where only π_θ is optimized. The training objective follows Eq. 3, with $R(\tau)$ detailed in Section 4.3.

220 You are building a workflow step by step to solve the problem. In each turn, output **EXACTLY ONE** XML action
 221 (add/delete/modify/set_prompt/finish or a structure action). The goal is to build a reliable workflow, not a single-shot
 222 answer. Keep your thinking brief and focus on choosing the next action, not solving the whole problem yourself. Let the operators
 223 do the computation. Proceed step by step with the following rules: **<think>** (brief reasoning for the current state and next action)
 224 **</think> <action>** (exactly one editing action with operator/target) **</action>** After the first step, in each interaction with the canvas,
 225 write: **<think>** (your reasoning based on the **<feedback>...</feedback>** from canvas) **</think> <action>** (new action to extend or refine
 226 the workflow) **<action>** Each **<action>** must build on what came before. Let the content of the workflow evolve naturally (for example:
 227 plan → programmer → verify → format). Continue producing think within **<think>...</think>** and action within **<action>...</action>**
 228 until the workflow is complete. Once the workflow is ready, write: **<think>** (final check that Format operator is added) **</think>**
<action>finish</action> Task: **{task}**.

Table 2. The system prompt template utilized by Flow-Director to interact with the Workflow Canvas.

232 **Proposition 2.** *The canvas-feedback-driven multi-turn interaction replaces manual configuration, improving orchestration efficiency and reliability.*

236 *Proof.* We provide experimental results in Section 5.5 and
 237 theoretical proofs in Appendix B.2. \square

4.3. Outcome-Directed End-to-End Reinforcement Learning

243 To optimize policy π_θ toward generating well-structured
 244 workflows, we propose Canvas Workflow Relative Policy
 245 Optimization (CWRPO), which optimizes through multi-
 246 turn agent-canvas interactions.

247 **CWRPO Objective** $\mathcal{J}_{\text{CWRPO}}(\theta)$. Given task $q \in \mathcal{D}_Q$,
 248 the agent interacts with canvas \mathcal{C} to generate a
 249 group of N trajectories $\{\tau_i\}_{i=1}^N$, where each $\tau_i =$
 250 $\{(a_t^{\text{think},(i)}, a_t^{(i)}, o_t^{\text{exec},(i)})\}_{t=1}^T$. The CWRPO objective is:

$$\begin{aligned}
 \mathcal{J}_{\text{CWRPO}}(\theta) = \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \frac{1}{|\tau_i|_{\text{mask}}} \sum_{t=1}^{|\tau_i|} \text{mask}_t^{(i)} \right. \\
 \left. \cdot \min \left(\rho_\theta^{(i,t)} \hat{A}_i, \text{clip}(\rho_\theta^{(i,t)}, 1 \pm \epsilon) \hat{A}_i \right) - \beta D_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right], \quad (10)
 \end{aligned}$$

where the importance ratio is defined as:

$$\rho_\theta^{(i,t)} = \frac{\pi_\theta(a_t^{\text{think},(i)}, a_t^{(i)} | H_{t-1}^{(i)})}{\pi_{\theta_{\text{old}}}(a_t^{\text{think},(i)}, a_t^{(i)} | H_{t-1}^{(i)}), \quad (11)$$

264 which measures the ratio between current policy π_θ and
 265 behavior policy $\pi_{\theta_{\text{old}}}$ used for trajectory sampling, and the
 266 normalized advantage $\hat{A}(\tau_i) = (R(\tau_i) - \mu_{\text{src}})/(\sigma_{\text{src}} + \epsilon)$
 267 uses within-group statistics (Shao et al., 2024a) μ_{src} , σ_{src}
 268 partitioned by data source. The mask $\text{mask}_t^{(i)} \in \{0, 1\}$ is
 269 set to 1 for policy-generated tokens and 0 for environment
 270 feedback, ensuring gradients only affect policy decisions;
 271 $|\tau_i|_{\text{mask}}$ is the masked token count. The $\text{clip}(\cdot)$ operator
 272 stabilizes updates (Schulman et al., 2017), and the KL term
 273 regularizes toward reference π_{ref} with strength β .

Outcome-directed Reward Function $R(\tau)$. We decompose the reward into diversity constraint $R_{\text{diversity}}(\tau)$ and answer reward $R_{\text{answer}}(\tau)$, encouraging both structural reliability and answer correctness. (i) *Diversity Constraint Reward.* This ensures workflows possess necessary structural “skeleton”:

$$\begin{aligned}
 R_{\text{diversity}}(\tau) = \min (1.0, R_{\text{checker}} + R_{\text{format}} \\
 + R_{\text{operator}} + R_{\text{control}}), \quad (12)
 \end{aligned}$$

where R_{checker} encourages verification operators, R_{format} encourages formatting operators, R_{operator} requires minimum operator count, and R_{control} encourages control structures. These components jointly ensure structural completeness for reliable workflow execution. Component weights are detailed in Appendix C.1. (ii) *Answer Reward.* This measures semantic correctness between execution output y_q (obtained via Eq. 2) and ground-truth y_q^* :

$$R_{\text{answer}}(\tau) = \text{Evaluate}(y_q, y_q^*, \text{type}(q)), \quad (13)$$

where $\text{type}(q)$ indicates the task type for metric selection, such as exact match for QA, pass rate for code generation, and accuracy for math reasoning. (iii) *Overall Outcome Reward.* The total reward employs conditional release, using diversity as prerequisite for answer reward:

$$\begin{aligned}
 R(\tau) = -1.0 + R_{\text{diversity}}(\tau) \\
 + \mathbb{I}\{R_{\text{diversity}}(\tau) = 1.0\} \cdot R_{\text{answer}}(\tau), \quad (14)
 \end{aligned}$$

where $\mathbb{I}\{\cdot\}$ is the indicator function that implements a curriculum-like gating mechanism. When $R_{\text{diversity}} < 1.0$, the answer reward is not released, prompting the policy to first construct qualified skeletons; when $R_{\text{diversity}} = 1.0$, the answer reward is released, shifting focus to correctness and suppressing shortcut behaviors.

Proposition 3. *The diversity constraint and conditional release mechanism stabilize learning signals and suppress shortcut behaviors.*

Proof. We provide experimental results in Section 5.5 and Section 5.6, and theoretical proofs in Appendix B.3. \square

Dataset	Metric	Baseline		SFT		GRPO		AFlow		Agent+RL (4o-mini)			Ours (4o-mini)
		Qwen3-8B	4o-mini	Qwen3-8B	Qwen3-8B	4o-mini	Agentflow	Router-R1	Orchestrator	FlowSteer ($\Delta \uparrow$)			
GSM8K	Acc.	91.41 \pm 0.4	92.97 \pm 0.6	92.19 \pm 0.3	92.97 \pm 0.8	94.53 \pm 0.5	93.75 \pm 0.7	94.01 \pm 0.4	93.94 \pm 0.9	96.09 (+3.12)			
MATH	Acc.	66.41 \pm 0.7	60.94 \pm 0.5	61.72 \pm 0.9	68.75 \pm 0.4	70.31 \pm 0.8	71.87 \pm 0.6	76.56 \pm 0.3	72.26 \pm 0.5	81.25 (+20.31)			
HotPotQA	EM	67.19 \pm 0.8	63.28 \pm 0.4	70.31 \pm 0.6	59.38 \pm 0.7	68.75 \pm 0.3	67.19 \pm 0.9	72.00 \pm 0.5	67.97 \pm 0.4	78.12 (+14.84)			
	F1	74.05 \pm 0.3	73.03 \pm 0.8	75.25 \pm 0.5	64.95 \pm 0.6	77.90 \pm 0.7	77.88 \pm 0.4	79.84 \pm 0.6	75.61 \pm 0.8	84.98 (+11.95)			
SQuAD v2	EM	54.69 \pm 0.6	47.66 \pm 0.7	73.44 \pm 0.4	66.41 \pm 0.5	73.44 \pm 0.9	64.06 \pm 0.3	59.84 \pm 0.8	70.34 \pm 0.7	78.12 (+30.46)			
	F1	61.54 \pm 0.5	59.42 \pm 0.6	77.31 \pm 0.8	72.00 \pm 0.4	82.41 \pm 0.3	72.45 \pm 0.9	65.29 \pm 0.5	75.24 \pm 0.6	83.67 (+24.25)			
MBPP	Pass@1	63.28 \pm 0.4	64.84 \pm 0.9	57.03 \pm 0.6	77.34 \pm 0.8	83.20 \pm 0.4	79.69 \pm 0.5	73.43 \pm 0.7	74.22 \pm 0.3	84.38 (+19.54)			
HumanEval	Pass@1	81.25 \pm 0.8	82.81 \pm 0.3	61.72 \pm 0.7	86.72 \pm 0.6	90.62 \pm 0.5	87.50 \pm 0.4	85.15 \pm 0.9	89.06 \pm 0.5	92.96 (+10.15)			
Avg.	EM	60.94 \pm 0.5	55.47 \pm 0.6	71.88 \pm 0.4	62.90 \pm 0.7	71.10 \pm 0.8	65.63 \pm 0.5	65.92 \pm 0.4	69.16 \pm 0.6	78.12 (+22.65)			
	F1	67.80 \pm 0.4	66.23 \pm 0.5	76.28 \pm 0.7	68.48 \pm 0.8	80.16 \pm 0.6	75.17 \pm 0.4	72.57 \pm 0.5	75.43 \pm 0.7	84.33 (+18.10)			
	Acc./Pass	75.59 \pm 0.6	75.39 \pm 0.4	68.17 \pm 0.8	81.45 \pm 0.5	84.67 \pm 0.7	83.20 \pm 0.6	82.29 \pm 0.3	82.37 \pm 0.4	88.67 (+13.28)			

Table 3. Main results on reasoning and code generation benchmarks. “Baseline” uses the original model without fine-tuning. “SFT” and “GRPO” apply supervised fine-tuning and group relative policy optimization on Qwen3-8B respectively. “AFlow” uses search-based workflow generation. “Agent+RL” denotes reinforcement learning-based agent methods. All workflow methods use GPT-4o-mini as backend.

5. Experiments

We evaluate FlowSteer through the following research questions (RQs): **RQ1**: Can FlowSteer outperform existing workflow orchestration methods? **RQ2**: How does FlowSteer generalize to out-of-distribution benchmarks? **RQ3**: How transferable is FlowSteer across different LLM backends? **RQ4**: What are the contributions of core components such as the multi-turn interaction paradigm and RL modules? **RQ5**: How does CWRPO compare against other RL algorithms (GRPO, DAPO)?

5.1. Experimental Setup

Datasets. We evaluate FlowSteer on six in-distribution (IID) benchmarks covering three task categories: (i) mathematical reasoning (GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021b)), (ii) question answering (HotPotQA (Yang et al., 2018), SQuAD v2 (Rajpurkar et al., 2018)), and (iii) code generation (MBPP (Austin et al., 2021), HumanEval (Chen et al., 2021)). To assess generalization, we further test on six out-of-distribution (OOD) benchmarks: TriviaQA (Joshi et al., 2017), NaturalQuestions (Kwiatkowski et al., 2019), MathQA (Amini et al., 2019), AIME 2025, APPS (Hendrycks et al., 2021a), and DS-1000 (Lai et al., 2023). More details of the datasets are provided in Appendix D.

Baselines. We compare FlowSteer against the following baselines: (i) direct LLM baselines (Qwen3-8B (Qwen Team, 2025), GPT-4o-mini (OpenAI, 2024)), (ii) standard fine-tuning methods (SFT, GRPO (Shao et al., 2024a)), (iii) workflow-based methods (AFlow (Zhang et al., 2024a)), and (iv) agent with RL methods (AgentFlow (Li et al., 2025b), Router-R1 (Zhang et al., 2025), Orchestrator (Dang et al., 2025)). More details are provided in Appendix E.

Evaluation Metrics. Following standard practice, we use

Exact Match (EM) and F1 score for question answering tasks, Accuracy (Acc.) for mathematical reasoning tasks, and Pass@1 for code generation benchmarks. More details of the evaluation metrics are shown in Appendix F.

Implementation Details. We use Qwen3-8B (Qwen Team, 2025) as the policy model (Flow-Director), with GPT-4o-mini (OpenAI, 2024) as the default workflow execution backend. For transferability experiments, we also test GPT-OSS-120B as an alternative backend. More implementation details are illustrated in Appendix G.

5.2. Main Results (RQ1)

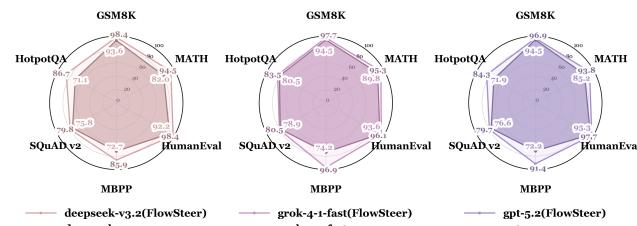
As illustrated in Table 3, FlowSteer achieves consistent improvements over the backend model (GPT-4o-mini) across six IID benchmarks and outperforms multiple baseline categories. The largest improvements are observed on mathematical reasoning datasets, while more stable performance is achieved on QA tasks compared to other learning or search-based methods. Strong Pass@1 results are also maintained on code generation benchmarks. These results confirm the broad applicability of FlowSteer and its ability to amplify backend model capabilities in complex reasoning scenarios.

5.3. Generalization Results (RQ2)

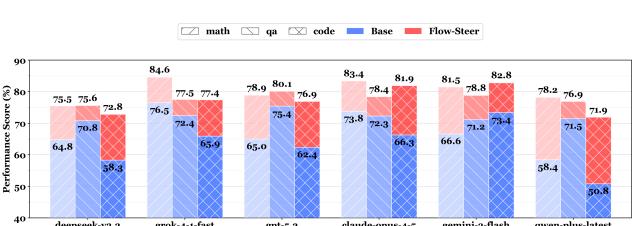
As shown in Table 4, FlowSteer maintains overall superiority over baseline methods across six OOD benchmarks, with more stable improvement trends in question answering and mathematical reasoning tasks. In comparison, directly using large-scale LLM backends typically exhibits stronger zero-shot capabilities but lacks transferable orchestration strategies, while search-based workflow methods are prone to efficiency limitations in OOD scenarios. These results demonstrate that FlowSteer achieves robust cross-task and cross-distribution improvements without task-specific fine-

Dataset	Metric	Baseline		SFT		GRPO		AFlow		Agent+RL (4o-mini)			Ours (4o-mini)
		Qwen3-8B	4o-mini	Qwen3-8B	Qwen3-8B	4o-mini	Agentflow	Router-R1	Orchestrator	FlowSteer ($\Delta \uparrow$)			
TriviaQA	EM	60.16 \pm 0.5	71.09 \pm 0.8	60.94 \pm 0.4	59.38 \pm 0.6	73.44 \pm 0.7	75.00 \pm 0.3	75.78 \pm 0.9	77.36 \pm 0.5	79.69 (+8.60)	84.11 (+2.71)	84.11 (+2.71)	84.11 (+2.71)
	F1	69.17 \pm 0.9	81.40 \pm 0.4	69.88 \pm 0.6	69.23 \pm 0.5	82.50 \pm 0.8	81.47 \pm 0.7	80.43 \pm 0.3	83.23 \pm 0.6				
NaturalQuestions	EM	39.84 \pm 0.6	39.84 \pm 0.5	46.09 \pm 0.8	43.75 \pm 0.4	42.97 \pm 0.9	45.70 \pm 0.7	49.22 \pm 0.6	50.00 \pm 0.3	54.69 (+14.85)	62.56 (+11.14)	62.56 (+11.14)	62.56 (+11.14)
	F1	50.75 \pm 0.4	51.42 \pm 0.9	53.40 \pm 0.5	53.24 \pm 0.8	49.92 \pm 0.6	55.98 \pm 0.4	52.79 \pm 0.5	55.41 \pm 0.7				
MathQA	Acc.	75.00 \pm 0.8	79.69 \pm 0.4	61.71 \pm 0.7	60.15 \pm 0.6	83.59 \pm 0.3	82.81 \pm 0.9	80.47 \pm 0.5	82.03 \pm 0.4	88.67 (+8.98)	26.67 (+16.67)	26.67 (+16.67)	26.67 (+16.67)
	AIME 2025	16.66 \pm 0.7	10.00 \pm 0.6	0.00 \pm 0.0	8.33 \pm 0.5	13.33 \pm 0.9	10.00 \pm 0.4	10.00 \pm 0.8	20.00 \pm 0.7				
APPS	Pass@1	39.84 \pm 0.5	40.62 \pm 0.8	26.56 \pm 0.6	34.38 \pm 0.7	42.97 \pm 0.4	41.41 \pm 0.9	42.95 \pm 0.5	44.53 \pm 0.6	49.21 (+8.59)	58.59 (+13.28)	58.59 (+13.28)	58.59 (+13.28)
	DS-1000	34.38 \pm 0.6	45.31 \pm 0.5	25.78 \pm 0.8	38.28 \pm 0.4	53.91 \pm 0.7	46.88 \pm 0.5	42.97 \pm 0.9	51.56 \pm 0.6				
Avg.	EM	50.00 \pm 0.4	55.47 \pm 0.6	53.52 \pm 0.5	51.57 \pm 0.7	58.21 \pm 0.8	60.35 \pm 0.4	62.50 \pm 0.6	63.68 \pm 0.5	67.19 (+11.72)	73.34 (+6.93)	73.34 (+6.93)	73.34 (+6.93)
	F1	59.96 \pm 0.5	66.41 \pm 0.4	61.64 \pm 0.8	61.24 \pm 0.6	66.21 \pm 0.5	68.73 \pm 0.7	66.61 \pm 0.4	69.32 \pm 0.6				
	Acc./Pass	41.47 \pm 0.7	43.91 \pm 0.5	28.51 \pm 0.6	35.29 \pm 0.9	48.45 \pm 0.4	45.28 \pm 0.8	44.10 \pm 0.5	49.53 \pm 0.7	55.79 (+11.88)			

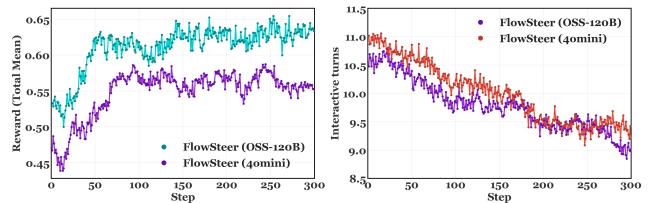
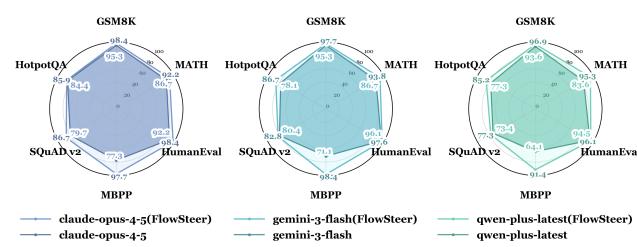
Table 4. OOD generalization results on QA, math, and code benchmarks. “Baseline” uses the original model without fine-tuning. “SFT” and “GRPO” apply supervised fine-tuning and group relative policy optimization on Qwen3-8B respectively. “AFlow” uses search-based workflow generation. “Agent+RL” denotes reinforcement learning-based agent methods. All workflow methods use GPT-4o-mini as backend.



(a) Radar charts on different backends



(b) Aggregated performance by task type



(c) Training dynamics

Figure 4. Transferability analysis of Flow-Director across LLM backends (RQ3). (a) Radar charts comparing six LLM backends (DeepSeek-V3.2, Grok-4.1-Fast, GPT-5.2, Claude-Opus-4.5, Gemini-3-Flash, Qwen-Plus-Latest) across six IID benchmarks, showing performance with and without Flow-Director trained on different backends. (b) Aggregated performance comparison across backends, grouped by task type (math, QA, code), comparing base LLMs vs. Flow-Director trained with 4o-mini vs. oss-120b. (c) Training dynamics showing F1 score, interaction turns, and operator counts over training steps for different backend configurations.

tuning of the backend LLM.

5.4. Transferability across LLM Backends (RQ3)

As shown in Figure 4, we evaluate the transferability of Flow-Director across six LLM backends (DeepSeek-V3.2, Grok-4.1-Fast, GPT-5.2, Claude-Opus-4.5, Gemini-3-Flash, and Qwen-Plus-Latest). The radar charts in Figure 4(a) show that FlowSteer yields consistent improvements across all backends on IID benchmarks, with weaker baselines benefiting more significantly. Figure 4(b) shows aggregated OOD performance grouped by task type, where FlowSteer maintains stable gains across math, QA, and code categories regardless of the backend model. The training dynamics in Figure 4(c) compare two training configurations: GPT-

4o-mini and locally deployed GPT-OSS-120B. While OSS-120B achieves higher reward with more stable optimization, both configurations show similar convergence trends in interaction turns, indicating that zero-cost local models can effectively replace API backends for training.

5.5. Ablation Study (RQ4)

As shown in Table 5, FlowSteer (Full) performs optimally across all datasets, highlighting the synergy of multi-turn interaction, executable feedback from Workflow Canvas, and end-to-end RL optimization. Removing any component significantly reduces performance, with RL most affected by complex reasoning tasks such as math and code generation. The canvas feedback has greater impact on QA tasks that

Method	IID							OOD							OOD									
	GSM8K		MATH		HotPotQA		SQuAD v2		MBPP		HumanEval		TriviaQA		NQ		MathQA		AIME		APPS		DS-1000	
	Acc.	Acc.	EM	F1	EM	F1	Pass@1	Pass@1	EM	F1	EM	F1	Acc.	Acc.	Pass@1	Pass@1	EM	F1	EM	F1	EM	F1	EM	F1
w/o Agent	94.53	70.31	72.66	78.91	66.41	70.63	60.94	83.59	69.53	75.08	53.12	57.43	80.47	23.33	41.40	40.63	63.28	68.28	43.75	55.00	78.13	20.00	38.84	41.41
w/o Multi-turn	91.41	75.00	72.66	81.09	56.25	66.48	57.81	88.28	57.81	65.63	53.13	61.48	78.91	10.00	42.96	45.31	64.84	71.80	53.13	60.72	80.47	20.00	38.28	49.22
w/o Canvas	94.53	73.44	70.31	79.53	59.38	69.53	63.28	85.16	79.69	84.11	54.69	62.56	88.67	26.67	49.21	58.59	79.69	84.11	54.69	62.56	88.67	26.67	49.21	58.59
w/o RL	91.41	71.81	72.66	82.19	59.38	67.66	64.06	89.06	79.69	84.11	54.69	62.56	88.67	26.67	49.21	58.59	79.69	84.11	54.69	62.56	88.67	26.67	49.21	58.59
FlowSteer (Full)	96.09	81.25	78.12	84.98	78.12	83.67	84.38	92.96	79.69	84.11	54.69	62.56	88.67	26.67	49.21	58.59	79.69	84.11	54.69	62.56	88.67	26.67	49.21	58.59

Table 5. Ablation study of FlowSteer (GPT-4o-mini as backend) on twelve datasets, including FlowSteer (Full), without Flow-Director agent (w/o Agent), without multi-turn interaction (w/o Multi-turn), without Workflow Canvas (w/o Canvas), and without reinforcement learning (w/o RL).

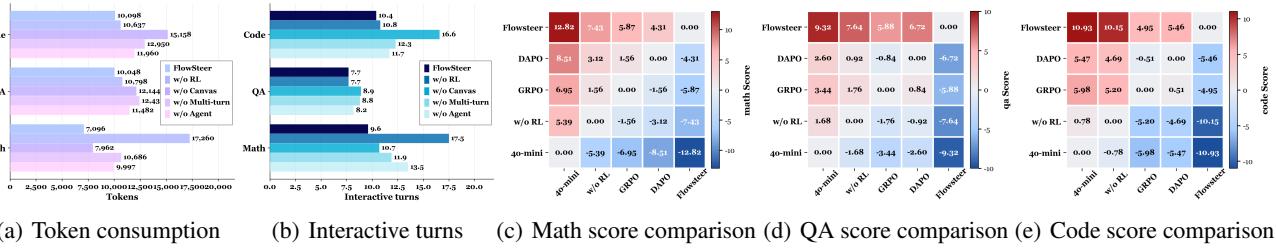


Figure 5. **Ablation and RL algorithm analysis.** (a) Token consumption comparison across task types, showing FlowSteer achieves lower token usage. (b) Average interaction turns comparison, demonstrating FlowSteer requires fewer turns to complete tasks. (c-e) Pairwise performance comparison matrices for math, QA, and code tasks respectively, where positive values (red) indicate row method outperforms column method.

require iterative error correction, showing task-dependent dependencies. In addition, FlowSteer (Full) exhibits more stable performance on the six OOD benchmarks, reflecting strong cross-task generalization ability. These results show that the effectiveness of FlowSteer stems from the close coupling of agent decision-making, canvas interaction, and RL optimization.

As shown in Figure 5(a-b), we further analyze orchestration efficiency in terms of average interaction turns and token consumption. The results show that FlowSteer achieves lower turns and fewer tokens across all task types compared to ablation variants. Removing any component causes the policy to produce redundant interactions and fail to determine proper termination timing, leading to higher time and computational costs. Through the joint effect of all components, the policy learns efficient orchestration strategies and when to stop, reducing both interaction turns and token consumption while maintaining task accuracy.

5.6. Comparison of RL Algorithms (RQ5)

As shown in Table 6, we compare three RL algorithms (DAPO, GRPO, and CWRPO) on six IID benchmarks under identical training settings. CWRPO achieves optimal results across all task types with more stable training dynamics. The pairwise comparison matrices in Figure 5(c-e) show consistent superiority of CWRPO across math, QA, and code tasks. Unlike DAPO and GRPO that primarily focus on answer correctness, CWRPO jointly optimizes

Method	GSM8K		MATH		HotPotQA		SQuAD v2		MBPP		HumanEval	
	Acc.	Acc.	EM	F1	EM	F1	Pass@1	Pass@1	EM	F1	EM	F1
4o-mini	92.97	60.94	63.28	73.03	47.66	59.42	64.84	82.81				
w/o RL	91.41	71.81	72.66	82.19	59.38	67.66	64.06	89.06				
GRPO	92.97	73.43	72.66	81.80	61.72	68.91	78.91	89.84				
DAPO	93.75	74.22	73.44	82.42	61.72	70.08	81.25	89.06				
CWRPO	96.09	81.25	78.12	84.98	78.12	83.67	84.38	92.96				

Table 6. Comparison of RL algorithms on six IID benchmarks using GPT-4o-mini as backend. Metrics follow each dataset’s standard (EM/F1 for QA, Acc. for math, Pass@1 for code).

structural diversity and task performance through diversity-constrained rewards with conditional release. This design enables the policy to learn not only how to solve problems but also how to orchestrate workflows effectively, achieving better trade-off between workflow quality and answer accuracy.

6. Conclusion

In this work, we propose FlowSteer, an end-to-end reinforcement learning framework for interactive workflow orchestration. By training a lightweight policy model to interact with an executable Workflow Canvas, FlowSteer learns transferable orchestration strategies from real execution feedback. We design diversity-constrained rewards with conditional release to jointly optimize workflow structure and task correctness. Experiments on twelve benchmarks demonstrate that FlowSteer significantly outperforms baselines across mathematical reasoning, QA, and code generation tasks.

440 **Impact Statement**

441 This paper presents work whose goal is to advance the field

442 of Machine Learning. There are many potential societal

443 consequences of our work, none which we feel must be

444 specifically highlighted here.

445

446 **References**

447

448 Amini, A., Gabriel, S., Lin, S., Koncel-Kedziorski, R., Choi,

449 Y., and Hajishirzi, H. MathQA: Towards interpretable

450 math word problem solving with operation-based for-

451 malisms. In *Proceedings of the 2019 Conference of the*

452 *North American Chapter of the Association for Compu-*

453 *tational Linguistics*, 2019.

454

455 Austin, J., Odena, A., Nye, M., Bosma, M., Michalewski,

456 H., Dohan, D., Jiang, E., Cai, C., Terry, M., Le, Q., and

457 Sutton, C. Program synthesis with large language models.

458 *arXiv preprint arXiv:2108.07732*, 2021.

459

460 Chen, M., Tworek, J., Jun, H., Yuan, Q., de Oliveira Pinto,

461 H. P., Kaplan, J., Edwards, H., Burda, Y., Joseph, N.,

462 Brockman, G., Ray, A., Puri, R., Krueger, G., Petrov,

463 M., Khlaaf, H., Sastry, G., Mishkin, P., Chan, B., Gray,

464 S., Ryder, N., Pavlov, M., Power, A., Kaiser, L., Bavari-

465 an, M., Winter, C., Tillet, P., Such, F. P., Cummings, D.,

466 Plappert, M., Chantzis, F., Barnes, E., Herbert-Voss, A.,

467 Guss, W. H., Nichol, A., Paino, A., Tezak, N., Tang,

468 J., Babuschkin, I., Balaji, S., Jain, S., Saunders, W.,

469 Hesse, C., Carr, A. N., Leike, J., Achiam, J., Misra,

470 V., Morikawa, E., Radford, A., Knight, M., Brundage,

471 M., Murati, M., Mayer, K., Welinder, P., McGrew, B.,

472 Amodei, D., McCandlish, S., Sutskever, I., and Zaremba,

473 W. Evaluating large language models trained on code.

474 *arXiv preprint arXiv:2107.03374*, 2021.

475

476 Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H.,

477 Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano,

478 R., Hesse, C., and Schulman, J. Training verifiers to solve

479 math word problems. *arXiv preprint arXiv:2110.14168*,

480 2021.

481

482 Dang, Y., Qian, C., Luo, X., Fan, J., Xie, Z., Shi, R., Chen,

483 W., Yang, C., Che, X., Tian, Y., Xiong, X., Han, L., Liu,

484 Z., and Sun, M. Multi-agent collaboration via evol-

485 ving orchestration. In *Advances in Neural Information*

486 *Processing Systems*, 2025.

487

488 DeepSeek-AI. DeepSeek-V3 technical report. *arXiv*

489 *preprint arXiv:2412.19437*, 2024.

490

491 DeepSeek-AI. DeepSeek-R1: Incentivizing reasoning capa-

492 bility in LLMs via reinforcement learning. *arXiv preprint*

493 *arXiv:2501.12948*, 2025.

494

Du, Z., Qian, C., Liu, W., Xie, Z., Wang, Y., Qiu, R., Dang, Y., Chen, W., Yang, C., Tian, Y., Xiong, X., and Han, L. Multi-agent collaboration via cross-team orchestration. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 10386–10406. Association for Computational Linguistics, 2025.

Erdogan, L. E., Furuta, H., Kim, S., Lee, N., Moon, S., Anumanchipalli, G., Keutzer, K., and Gholami, A. Plan-and-act: Improving planning of agents for long-horizon tasks. In *International Conference on Machine Learning*, 2025.

Feng, J., Huang, S., Qu, X., Zhang, G., Qin, Y., Zhong, B., Jiang, C., Chi, J., and Zhong, W. ReTool: Reinforcement learning for strategic tool use in LLMs. *arXiv preprint arXiv:2504.11536*, 2025.

Google DeepMind. Gemini 2.5: A family of highly capable multimodal models. *arXiv preprint arXiv:2503.00000*, 2025.

Hendrycks, D., Basart, S., Kadavath, S., Mazeika, M., Arora, A., Guo, E., Burns, C., Puranik, S., He, H., Song, D., and Steinhardt, J. Measuring coding challenge competence with APPS. *Advances in Neural Information Processing Systems*, 34, 2021a.

Hendrycks, D., Burns, C., Kadavath, S., Arora, A., Basart, S., Tang, E., Song, D., and Steinhardt, J. Measuring mathematical problem solving with the MATH dataset. *Advances in Neural Information Processing Systems*, 34, 2021b.

Hong, S., Zhuge, M., Chen, J., Zheng, X., Cheng, Y., Zhang, C., Wang, J., Wang, Z., Yau, S. K. S., Lin, Z., Zhou, L., Ran, C., Xiao, L., Wu, C., and Schmidhuber, J. MetaGPT: Meta programming for a multi-agent collaborative framework. In *International Conference on Learning Representations*, 2024.

Jin, B., Zeng, H., Yue, Z., Yoon, J., Arik, S., Wang, D., Zamani, H., and Han, J. Search-R1: Training LLMs to reason and leverage search engines with reinforcement learning. *arXiv preprint arXiv:2503.09516*, 2025.

Joshi, M., Choi, E., Weld, D. S., and Zettlemoyer, L. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, 2017.

Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A., Alberti, C., Epstein, D., Polosukhin, I., Devlin, J., Lee, K., Toutanova, K., Jones, L., Kelcey, M., Chang, M.-W., Dai, A. M., Uszkoreit, J., Le, Q., and Petrov, S. Natural questions: A benchmark for question

495 answering research. *Transactions of the Association for
496 Computational Linguistics*, 7, 2019.

497

498 Lai, Y., Li, C., Wang, Y., Zhang, T., Zhong, R., Zettlemoyer,
499 L., tau Yih, S. W., Fried, D., Wang, S., and Yu, T. DS-
500 1000: A natural and reliable benchmark for data science
501 code generation. In *International Conference on Machine
502 Learning*, 2023.

503 Langley, P. Crafting papers on machine learning. In Langley,
504 P. (ed.), *Proceedings of the 17th International Conference
505 on Machine Learning (ICML 2000)*, pp. 1207–1216, Stan-
506 ford, CA, 2000. Morgan Kaufmann.

507

508 Li, Z., Hu, Y., and Wang, W. Encouraging good processes
509 without the need for good answers: Reinforcement learn-
510 ing for LLM agent planning. In *Proceedings of the 2025
511 Conference on Empirical Methods in Natural Language
512 Processing: Industry Track*, pp. 1654–1666. Association
513 for Computational Linguistics, 2025a.

514

515 Li, Z., Zhang, H., Han, S., Liu, S., Xie, J., Zhang, Y.,
516 Choi, Y., Zou, J., and Lu, P. In-the-flow agentic system
517 optimization for effective planning and tool use. *arXiv
518 preprint arXiv:2510.05592*, 2025b.

519

520 Meta AI. The LLaMA 4 herd of models. *arXiv preprint
521 arXiv:2504.00000*, 2025.

522

523 OpenAI. GPT-4o system card. *OpenAI Technical Report*,
524 2024.

525

526 Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C.,
527 Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A.,
528 Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M.,
529 Askell, A., Welinder, P., Christiano, P. F., Leike, J., and
530 Lowe, R. Training language models to follow instructions
531 with human feedback. *Advances in Neural Information
532 Processing Systems*, 35, 2022.

533

534 Qian, C., Liu, W., Liu, H., Chen, N., Dang, Y., Li, J., Yang,
535 C., Chen, W., Su, Y., Cong, X., Xu, J., Li, D., Liu, Z., and
536 Sun, M. ChatDev: Communicative agents for software
537 development. In *Proceedings of the 62nd Annual Meeting
538 of the Association for Computational Linguistics*, 2024.

539

540 Qwen Team. Qwen3 technical report. *arXiv preprint
541 arXiv:2505.09388*, 2025.

542

543 Rajpurkar, P., Jia, R., and Liang, P. Know what you don't
544 know: Unanswerable questions for SQuAD. In *Proceed-
545 ings of the 56th Annual Meeting of the Association for
546 Computational Linguistics*, 2018.

547

548 Ruan, J., Chen, Y., Zhang, B., Xu, Z., Bao, T., Du, G., Shi,
549 S., Mao, H., Li, Z., Zeng, X., and Zhao, R. TPTU: Large
language model-based AI agents for task planning and
tool usage. *arXiv preprint arXiv:2308.03427*, 2023.

550

551 Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli,
552 M., Zettlemoyer, L., Cancedda, N., and Scialom, T. Tool-
553 former: Language models can teach themselves to use
554 tools. *Advances in Neural Information Processing Sys-
555 tems*, 36, 2023.

556

557 Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and
558 Klimov, O. Proximal policy optimization algorithms.
559 *arXiv preprint arXiv:1707.06347*, 2017.

560

561 Shang, Y., Li, Y., Zhao, K., Ma, L., Liu, J., Xu, F., and Li, Y.
562 AgentSquare: Automatic LLM agent search in modular
563 design space. In *International Conference on Learning
564 Representations*, 2025.

565

566 Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., Zhang,
567 H., Zhang, M., Li, Y., Wu, Y., and Guo, D. DeepSeek-
568 Math: Pushing the limits of mathematical reasoning in
569 open language models. *arXiv preprint arXiv:2402.03300*,
570 2024a.

571

572 Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Zhang, M.,
573 Li, Y., Wu, Y., and Guo, D. DeepSeekMath: Pushing
574 the limits of mathematical reasoning in open language
575 models. *arXiv preprint arXiv:2402.03300*, 2024b.

576

577 Shen, Y., Song, K., Tan, X., Li, D., Lu, W., and Zhuang,
578 Y. HuggingGPT: Solving AI tasks with ChatGPT and its
579 friends in Hugging Face. *Advances in Neural Information
580 Processing Systems*, 36, 2023.

581

582 Wang, Z. Z., Mao, J., Fried, D., and Neubig, G. Agent
583 workflow memory. *arXiv preprint arXiv:2409.07429*,
584 2024.

585

586 Wu, J., Zhu, J., Liu, Y., Xu, M., and Jin, Y. Agentic rea-
587 soning: A streamlined framework for enhancing LLM
588 reasoning with agentic tools. In *Proceedings of the 63rd
589 Annual Meeting of the Association for Computational
590 Linguistics*, pp. 28489–28503. Association for Compu-
591 tational Linguistics, 2025.

592

593 Yang, K., Liu, Y., Chaudhary, S., Fakoor, R., Chaudhari,
594 P. A., Karypis, G., and Rangwala, H. AgentOccam: A
595 simple yet strong baseline for LLM-based web agents. In
596 *International Conference on Learning Representations*,
597 2025.

598

599 Yang, Z., Qi, P., Zhang, S., Bengio, Y., Cohen, W., Salakhut-
600 dinov, R., and Manning, C. D. HotpotQA: A dataset for
601 diverse, explainable multi-hop question answering. In
602 *Proceedings of the 2018 Conference on Empirical Meth-
603 ods in Natural Language Processing*, 2018.

604

605 Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan,
606 K., and Cao, Y. ReAct: Synergizing reasoning and act-
607 ing in language models. In *International Conference on
608 Learning Representations*, 2023.

550 Yu, Y., Wang, Z., Ma, W., Wang, S., Wu, C., Guo, Z., and
551 Zhang, M. StepTool: Enhancing multi-step tool usage
552 in LLMs via step-grained reinforcement learning. In
553 *Proceedings of the 34th ACM International Conference*
554 *on Information and Knowledge Management*, 2025.

555 Zeng, Z., Watson, W., Cho, N., Rahimi, S., Reynolds, S.,
556 Balch, T., and Veloso, M. FlowMind: Automatic work-
557 flow generation with LLMs. In *Proceedings of the 4th*
558 *ACM International Conference on AI in Finance*, 2023.

560 Zhang, H., Feng, T., and You, J. Router-R1: Teaching LLMs
561 multi-round routing and aggregation via reinforcement
562 learning. *arXiv preprint arXiv:2506.09033*, 2025.

563

564 Zhang, J., Xiang, J., Yu, Z., Teng, F., Chen, X., Chen, J.,
565 Zhuge, M., Cheng, X., Hong, S., Wang, J., Zheng, B.,
566 Liu, B., Luo, Y., and Wu, C. AFlow: Automating agentic
567 workflow generation. *arXiv preprint arXiv:2410.10762*,
568 2024a.

569

570 Zhang, J., Xiang, J., Yu, Z., Teng, F., Chen, X., Chen, J.,
571 Zhuge, M., Cheng, X., Hong, S., Wang, J., Zheng, B.,
572 Liu, B., Luo, Y., and Wu, C. AFlow: Automating agentic
573 workflow generation. *arXiv preprint arXiv:2410.10762*,
574 2024b.

575 Zhou, A., Yan, K., Shlapentokh-Rothman, M., Wang, H.,
576 and Wang, Y.-X. Language agent tree search unifies rea-
577 soning acting and planning in language models. *Advances*
578 *in Neural Information Processing Systems*, 36, 2023.

579

580 Zhou, Y., Zanette, A., Pan, J., Levine, S., and Kumar, A.
581 ArCHer: Training language model agents via hierarchical
582 multi-turn RL. *arXiv preprint arXiv:2402.19446*, 2024.

583

584 Zhuge, M., Wang, W., Kirsch, L., Faccio, F., Khizbulin,
585 D., and Schmidhuber, J. Language agents as optimizable
586 graphs. *arXiv preprint arXiv:2402.16823*, 2024.

587

588

589

590

591

592

593

594

595

596

597

598

599

600

601

602

603

604

605	System Prompt for Flow-Director			
606	You are building a workflow step by step to solve the problem. In each turn, output EXACTLY ONE XML action (add/delete/modify/set_prompt/finish or a structure add). The goal is to build a reliable workflow, not a single-shot answer. Keep your thinking brief (under 200 words) and focus on choosing the next action, not solving the whole problem yourself. Let the operators do the computation. If you use <think>...</think>, you MUST output an <action> tag AFTER it.			
607	Available Operators (12 total). Programmer: write and execute Python code. Plan: create solution strategy. Custom: natural language reasoning. Decompose: break into sub-problems. Test: run test cases. Review: evaluate quality. Verify: double-check result. Revise: fix issues. ScEnsemble: multiple solutions voting. Aggregate: combine results. AnswerGenerate: format final answer. Format: extract concise answer.			
608	Actions (8 types). The system supports 8 action types: add (add a single operator), finish (complete workflow building), parallel (add parallel branches), conditional (add conditional branch), loop (add loop structure), delete (remove a node), modify (change operator type), and set_prompt (set custom prompt for an operator). All actions use XML format for reliable parsing.			
609	Finish Policy. Always add Format as the last step before finishing to extract concise answer. Before finishing, add Format to extract concise answer from the solution. When Format has extracted the answer and you are satisfied, output <action>finish</action>. When the result is wrong or needs improvement, add more operators.			

Figure 6. The system prompt template for Flow-Director. The prompt instructs the agent to build workflows step by step through multi-turn interactions with the Workflow Canvas.

Table 7. Summary of the 12 operators in \mathcal{O} with input/output specifications.

Operator	Category	Input	Output	Description
Plan	Planning	problem	approach, plan	Creates high-level solution strategy by identifying the overall approach and breaking it into actionable steps.
Decompose	Planning	problem	sub_problems	Breaks complex problems into smaller, independent sub-problems that can be solved separately.
Programmer	Solving	problem, analysis	code, output	Writes and executes Python code to compute answers; used for mathematical calculations and code generation tasks.
Custom	Solving	input, context	response	Performs natural language reasoning without code execution; used for QA, analysis, and explanations.
AnswerGenerate	Solving	input	thought, answer	Generates structured answers with explicit reasoning chains; similar to Custom but with formatted output.
Test	Verification	code, tests	pass/fail, solution	Executes unit tests against generated code and triggers automatic revision if tests fail; code-specific.
Review	Verification	solution	is_correct, feedback	Evaluates solution quality through critique and provides detailed feedback for potential revision.
Verify	Verification	answer	is_correct, answer	Independently re-solves the problem to verify correctness; does not execute code, uses logical reasoning.
Revise	Revision	solution, feedback	revised_solution	Fixes issues in solutions based on feedback from Review or Test operators.
ScEnsemble	Ensemble	solutions	selected_solution	Implements majority voting across multiple candidate solutions to select the most consistent answer.
Aggregate	Ensemble	sub_answers	aggregated	Combines results from parallel branches or sub-problems into a unified answer.
Format	Formatting	solution	final_answer	Extracts concise final answer from verbose solution; required as the last step before finishing.

A. Prompts and Operator Library in Flow-Steer

A.1. System Prompt for Flow-Director

As shown in Figure 6, we adopt the system prompt template for the Flow-Director, where the agent alternates between internal reasoning in <think> and communicates with the Workflow Canvas by <action>. The prompt instructs the agent to output exactly one XML action per turn, keeping thinking brief and focusing on action selection rather than problem-solving. The workflow evolves step by step until finalized with a `finish` action.

A.2. Operator Library \mathcal{O}

As shown in Table 7, the operator library \mathcal{O} consists of 12 functional operators organized into six categories: planning (Plan, Decompose), solving (Programmer, Custom, AnswerGenerate), verification (Test, Review, Verify), revision (Revise), ensemble (ScEnsemble, Aggregate), and formatting (Format). The Programmer operator executes Python code in a sandboxed environment with scientific libraries under a 30-second timeout.

A.3. Multi-turn Interaction Process

The Flow-Director builds workflows through multi-turn interactions with the Workflow Canvas. Each interaction round consists of the Flow-Director generating an action

Multi-turn Interaction Example	
Problem: “Find the sum of all positive integers n such that $n^2 + 12n - 2007$ is a perfect square.”	
Round 1 (BUILDING → AWAITING_PROMPT)	
<i>Flow-Director:</i>	
<think>This is a number theory problem. Start with Plan.</think> <action>add</action><operator>Plan</operator>	
<i>Canvas Feedback:</i> “Added Plan as node_1. State: AWAITING_PROMPT.”	
Round 2 (AWAITING_PROMPT → BUILDING)	
<i>Flow-Director:</i>	
<action>set_prompt</action><target>node_1</target> <prompt>Analyze the equation and identify the approach.</prompt>	
<i>Canvas Feedback:</i> “Prompt set. State: BUILDING. DSL: Plan”	
Round 3-4: Add Programmer operator with prompt for algebraic computation.	
Round 5-6: Add Verify operator with prompt for result validation.	
Round 7-8: Add Format operator with prompt for answer extraction.	
Round 9 (BUILDING → Finished)	
<i>Flow-Director:</i>	
<think>Workflow complete. Ready to finish.</think> <action>finish</action>	
<i>Canvas Feedback:</i> “Workflow finished. Final DSL: Plan -> Programmer -> Verify -> Format. Answer: 89”	

Figure 7. A complete multi-turn interaction example showing the two-step mechanism (add + set_prompt) for workflow construction.

and the Canvas providing execution feedback.

A.3.1. ENVIRONMENT STATE MACHINE

The Workflow Canvas maintains a finite state machine with two states:

- **BUILDING:** The normal state where the Flow-Director can execute add, delete, modify, or finish actions.
- **AWAITING_PROMPT:** After adding an operator, the Canvas transitions to this state, requiring the Flow-Director to specify a custom prompt via set_prompt before returning to BUILDING.

This two-step mechanism reduces the cognitive load on the policy model by separating structural decisions (which operator to add) from content decisions (what prompt to use), improving the effectiveness of small-scale models.

A.3.2. ACTION TYPES

The 8 action types available to the Flow-Director are summarized in Table 8. The basic actions (add, finish, delete, modify, set_prompt) enable sequential workflow construction, where the Flow-Director progressively builds the operator graph through iterative additions and modifications. The advanced actions (parallel, conditional, loop) extend

Table 8. Action types in Flow-Steer.

Action	Description
add	Add a single operator
finish	Complete workflow building
delete	Remove a node
modify	Change operator type
set_prompt	Set custom prompt
parallel	Add parallel branches
conditional	Add conditional branch
loop	Add loop structure

the expressiveness to complex control structures, allowing the Flow-Director to construct workflows with branching, conditional execution, and iterative refinement patterns.

A.3.3. COMPLETE INTERACTION EXAMPLE

A complete multi-turn interaction between the Flow-Director (small-scale policy model) and the Workflow Canvas (large-scale LLM backend) is illustrated in Figure 7. The process begins with the Flow-Director analyzing the problem and selecting an appropriate initial operator, which establishes the reasoning strategy for subsequent workflow construction. The Flow-Director then engages in iterative exchanges: after each canvas feedback, it reflects on the current workflow state, selects the next operator, and specifies task-specific prompts. Through successive rounds of construction and verification, the workflow gradually takes

715 shape, and the Canvas executes the completed workflow to
716 produce the final answer.
717

B. Theoretical Proofs

720 In this appendix, we provide detailed proofs for Propositions
721 1–3 stated in the main text. We first introduce the notation
722 and assumptions, then present each proof in turn.
723

724 **Notation.** Given a task q , a workflow is represented
725 as a directed graph $\mathcal{G} = (V, E, \text{attr})$, where each node
726 $v \in V$ is bound to an operator $\text{op}(v) \in \mathcal{O}$ with attributes
727 $\text{attr}(v) = (\text{param}(v), \text{prompt}(v), \dots)$. The executor
728 schedules nodes according to dependency edges E and
729 produces output $y = \text{Execute}(\mathcal{G}, q, \mathcal{M}_{\text{backend}})$. During
730 multi-turn orchestration, the policy model (Flow-Director)
731 interacts with the canvas (Workflow Canvas), forming a
732 trajectory
733

$$\tau = \{(a_t^{\text{think}}, a_t, o_t^{\text{exec}})\}_{t=1}^T, \quad (15)$$

734 where a_t^{think} denotes the reflection text that summarizes
735 the current state and identifies potential issues, a_t denotes
736 the edit action comprising an action type and its content,
737 and o_t^{exec} denotes the execution and validation feedback
738 returned by the canvas environment. The operator library
739 consists of 12 functional operators:
740

$$\mathcal{O} = \{\text{Plan}, \text{Decompose}, \text{Programmer}, \text{Custom}, \\ \text{AnswerGenerate}, \text{Test}, \text{Review}, \text{Verify}, \\ \text{Revise}, \text{ScEnsemble}, \text{Aggregate}, \text{Format}\}, \quad (16)$$

741 where each operator implements a specific cognitive function
742 in the problem-solving process. The action type set
743 consists of 8 types:
744

$$\mathcal{A}_{\text{type}} = \{\text{add}, \text{delete}, \text{modify}, \text{set_prompt}, \\ \text{finish}, \text{parallel}, \text{conditional}, \text{loop}\}, \quad (17)$$

745 where the first five types can be understood as local editing
746 actions, and the last three types are used to explicitly generate
747 control structures (parallel, conditional, loop) that form
748 long-range control flows.
749

750 **Assumptions.** To obtain formal proofs, we introduce three
751 mild and interpretable assumptions. They do not require the
752 real system to be perfect, but are sufficient to support the
753 theoretical justification of why the propositions hold.
754

755 **Assumption 1 (Cognitive Decomposability).** For the task
756 families considered in this work (multi-hop QA, standard
757 QA, mathematical reasoning, code generation), there exists
758 a goal-directed problem-solving procedure that can be de-
759 composed into a finite combination of cognitive primitives
760 (defined below), and can be represented as a directed graph
761 with conditionals and loops.
762

763 **Assumption 2 (Informative Canvas Feedback).** The canvas
764 feedback o_t^{exec} is informative about whether the current
765 workflow is executable, whether it satisfies structural
766 constraints, and which local modifications can fix errors.
767 Formally, there exists non-zero probability that the feedback
768 changes the agent’s posterior distribution over the optimal
769 next edit, i.e., $I(Z; o_t^{\text{exec}} | H_{t-1}) > 0$, where Z denotes
770 the latent variable related to the solution (e.g., the correct
771 answer or a feasible workflow).
772

773 **Assumption 3 (Repairability).** When a workflow is in a non-
774 executable or constraint-violating state, the canvas can provide
775 sufficiently localized failure reasons or repair suggestions,
776 such that there exist editing actions that can push it toward
777 a more feasible state. Formally, there exists a potential
778 function $\Phi(\mathcal{G})$ measuring the distance to the feasible set, and
779 there exist actions such that $\mathbb{E}[\Phi(\mathcal{G}_t) | H_{t-1}] < \Phi(\mathcal{G}_{t-1})$.
780

B.1. Proof of Proposition 1

781 **Proposition 4 (Operator-Action Cognitive Completeness).**
782 Let the operator library \mathcal{O} and action type set $\mathcal{A}_{\text{type}}$ be
783 defined as above. Under Assumption 1, for any task q ,
784 there exists a finite-length action sequence $\{a_t\}_{t=1}^T$ with
785 each action type belonging to $\mathcal{A}_{\text{type}}$, such that starting
786 from an empty canvas \mathcal{G}_0 , the iterative updates $\mathcal{G}_t =$
787 $\text{Update}(\mathcal{G}_{t-1}, a_t, o_t^{\text{exec}})$ can construct a terminal workflow
788 \mathcal{G}_T that implements a complete cognitive control loop,
789 thereby covering the key problem-solving steps required for
790 the task types considered in this work (mathematics, QA,
791 code generation).
792

793 *Proof.* This proposition requires proving two things: (i)
794 functional completeness, showing that the operator library
795 covers the core cognitive modules required for goal-directed
796 problem solving; and (ii) structural completeness, showing
797 that the action space can organize these modules into
798 control flow graphs with sufficient expressive power (sequential,
799 branching, looping, parallel), thereby constructing
800 executable workflow programs. We divide the proof into
801 three parts: first defining the cognitive primitive set \mathcal{C} , then
802 proving that \mathcal{O} covers \mathcal{C} , and finally proving that $\mathcal{A}_{\text{type}}$ can
803 construct any structured workflow graph composed of these
804 modules.
805

806 **(i) Cognitive primitives and functional coverage.** We first
807 define a set of cognitive primitives consistent with goal-
808 directed problem solving. Let the cognitive primitive set
809 be
810

$$\mathcal{C} = \{\mathbf{P}, \mathbf{D}, \mathbf{E}, \mathbf{M}, \mathbf{R}, \mathbf{I}, \mathbf{O}\}, \quad (18)$$

821 where each element represents a fundamental cognitive process
822 in problem solving. Specifically, \mathbf{P} (Planning) forms
823 goals, strategies, and resource budgets, determining what to
824 do first, what to do later, and to what extent; \mathbf{D} (Decomposition)
825 breaks the overall task into operable subgoals with
826

770 explicit dependency relationships; **E** (Execution) performs
 771 solving, reasoning, or externalized computation on a sub-
 772 goal, including symbolic computation, code generation and
 773 execution, and other operations; **M** (Monitoring) verifies,
 774 reviews, unit tests, and checks constraints on intermediate
 775 products, determining whether to continue, backtrack, or
 776 redo; **R** (Revision) performs repairs based on monitoring
 777 results, including rewriting, supplementing, replacing, and
 778 adjusting prompts or parameters; **I** (Integration) fuses, votes,
 779 disambiguates, and ensures consistency of results from mul-
 780 tiple branches or subproblems; and **O** (Output) outputs the
 781 final result in the format required by the task, including
 782 answer extraction and structured presentation.

783 We now construct a surjective mapping $\psi : \mathcal{O} \rightarrow \mathcal{C}$ to show
 784 that for each cognitive primitive $c \in \mathcal{C}$, there exists at least
 785 one operator $o \in \mathcal{O}$ such that $\psi(o) = c$. Consider the
 786 following correspondence:
 787

788 For the planning primitive **P**, the operator **Plan** directly
 789 implements high-level plan generation, including strategy
 790 formulation, step sequencing, budget allocation, and stop-
 791 ping condition specification. Therefore, we have
 792

$$\psi(\text{Plan}) = \mathbf{P}. \quad (19)$$

794 For the decomposition primitive **D**, the operator
 795 **Decompose** expresses task decomposition in the form of
 796 subtask lists or subproblem graphs, explicitly representing
 797 the dependency structure among subtasks. Therefore, we
 798 have
 799

$$\psi(\text{Decompose}) = \mathbf{D}. \quad (20)$$

800 For the execution primitive **E**, execution includes not
 801 only natural language reasoning but also externalized solving.
 802 The operator **Custom** covers general reasoning and
 803 retrieval-based processing, capable of hosting various tools
 804 and prompt templates according to implementation; the
 805 operator **Programmer** covers executable code genera-
 806 tion and execution, typically handling symbolic computa-
 807 tion, scripted reasoning, and data processing; the operator
 808 **AnswerGenerate** covers generating final natural lan-
 809 guage answers from obtained key evidence or intermediate
 810 conclusions, which can be viewed as decoding or expression.
 811 Therefore, we have
 812

$$\begin{aligned} \psi(\text{Custom}) &= \psi(\text{Programmer}) \\ &= \psi(\text{AnswerGenerate}) = \mathbf{E}. \end{aligned} \quad (21)$$

813 For the monitoring primitive **M**, the operators **Verify** and
 814 **Review** perform consistency checking and critical eval-
 815 uation on intermediate results; the operator **Test** performs
 816 executable testing on code solutions. These are all typical
 817 error monitoring and quality assessment processes. Therefore,
 818 we have
 819

$$\psi(\text{Verify}) = \psi(\text{Review}) = \psi(\text{Test}) = \mathbf{M}. \quad (22)$$

820 For the revision primitive **R**, the operator **Revise** explic-
 821 itly implements error correction and rewriting in a feedback-
 822 based repair manner. Therefore, we have
 823

$$\psi(\text{Revise}) = \mathbf{R}. \quad (23)$$

824 For the integration primitive **I**, when there are multiple
 825 branches, multiple candidates, or multiple subproblems, fu-
 826 sion is needed. The operator **Aggregate** handles aggrega-
 827 tion and consistency ensuring; the operator **ScEnsemble**
 828 handles selection and ensemble of diverse candidates. There-
 829 fore, we have
 830

$$\psi(\text{Aggregate}) = \psi(\text{ScEnsemble}) = \mathbf{I}. \quad (24)$$

831 For the output primitive **O**, the operator **Format** extracts,
 832 structures, and presents results in the target format. There-
 833 fore, we have
 834

$$\psi(\text{Format}) = \mathbf{O}. \quad (25)$$

835 For each $c \in \mathcal{C}$, the above construction provides at least one
 836 $o \in \mathcal{O}$ such that $\psi(o) = c$. Hence ψ is surjective onto \mathcal{C} ,
 837 establishing that \mathcal{O} achieves functional coverage of \mathcal{C} .

(ii) Structural completeness of the action space. Let $\mathfrak{G}(\mathcal{O})$ denote the set of all finite workflow graphs with nodes labeled by operators from \mathcal{O} and composed using structured control constructs (sequential, conditional, loop, parallel). We show that for any target graph $\mathcal{G}^* \in \mathfrak{G}(\mathcal{O})$, there exists a finite-length action sequence $\{a_t\}_{t=1}^T$ with action types belonging to $\mathcal{A}_{\text{type}}$, such that starting from the empty graph \mathcal{G}_0 , the iterative updates yield $\mathcal{G}_T = \mathcal{G}^*$. We provide a constructive proof, which is equivalent to showing that these editing actions can assemble any target graph.

Step 1: Node construction and attribute assignment (basic editing closure). For any target graph \mathcal{G}^* with node set V^* , we traverse $v \in V^*$ in any order. For each node, we use the action **add** to create a new node in the current canvas and specify its operator type $\text{op}(v) \in \mathcal{O}$. We then use the action **set_prompt** to set prompts and constraints for that node. If necessary, we use the action **modify** to write in parameters. If there are redundant or erroneous nodes, we use the action **delete** to remove them. Since **add**, **delete**, **modify**, and **set_prompt** allow arbitrary finite discrete modifications to the node set and attributes, we can construct within finite steps a node set and attribute annotation isomorphic to \mathcal{G}^* .

Step 2: Control structure construction (structured control closure). Control structures in the target graph can be divided into three categories: conditional branching, looping, and parallel execution. If \mathcal{G}^* contains conditional structures (if/else or selective execution), we use the action **conditional** to introduce conditional gate nodes or conditional edges in the graph, and write in condition

825 predicates via `set_prompt` or `modify`. The condition
 826 can be a Boolean signal output by some checking or verifica-
 827 tion node. If \mathcal{G}^* contains loop structures (repeated execution
 828 until some criterion is met), we use the action `loop` to intro-
 829 duce back-edges or iteration blocks. The stopping condition
 830 can similarly be produced by verification or test nodes, with
 831 the loop termination rule written in via `set_prompt` or
 832 `modify`. If \mathcal{G}^* contains parallel structures (multiple sub-
 833 problems or candidates expanded simultaneously), we use
 834 the action `parallel` to organize several branches into
 835 concurrent subgraphs, with subsequent integration via op-
 836 erators such as `Aggregate` or `ScEnsemble` to merge
 837 branch results.

838 *Step 3: Termination (completing construction).* When the
 839 current canvas graph aligns with \mathcal{G}^* in nodes, attributes, and
 840 control structures, we use the action `finish` to terminate
 841 editing. Since \mathcal{G}^* is a finite graph, the above process is finite,
 842 hence there exists finite T such that $\mathcal{G}_T = \mathcal{G}^*$.

843 **(iii) Coverage of task types.** We now show that the operator
 844 library and action space can cover the task types considered
 845 in this work by demonstrating the existence of constructible
 846 cognitive program templates for each task type.

847 For mathematical reasoning tasks, a typical cognitive pro-
 848 gram follows the pattern: `Plan` \rightarrow `Decompose` \rightarrow
 849 `(Custom or Programmer)` \rightarrow `Verify` \rightarrow `(Revise +`
 850 `loop)` \rightarrow `Format`. This template first formulates a solu-
 851 tion strategy, decomposes the problem into solvable steps,
 852 executes reasoning or symbolic computation, verifies inter-
 853 mediate results, revises if errors are detected (potentially
 854 looping), and finally formats the answer.

855 For question answering tasks including multi-
 856 hop QA, a typical cognitive program fol-
 857 lows the pattern: `Plan` \rightarrow `Decompose` \rightarrow
 858 `parallel(Custom/AnswerGenerate)` \rightarrow
 859 `Aggregate` \rightarrow `Review/Verify` \rightarrow `Format`.
 860 This template formulates a retrieval and reasoning strategy,
 861 decomposes into sub-questions that can be processed
 862 in parallel, aggregates evidence from multiple sources,
 863 reviews for consistency, and formats the final answer.

864 For code generation tasks, a typical cognitive pro-
 865 gram follows the pattern: `Plan` \rightarrow `Decompose` \rightarrow
 866 `Programmer` \rightarrow `Test` \rightarrow `(Revise + loop)` \rightarrow
 867 `Format`. This template plans the code structure, decom-
 868 poses into implementable modules, generates code, tests for
 869 correctness, revises based on test feedback (looping until
 870 tests pass), and formats the output.

871 These templates use only operators from \mathcal{O} and control struc-
 872 tures from $\mathcal{A}_{\text{type}}$. Therefore, under Assumption 1 (tasks
 873 can be decomposed into combinations of cognitive primitives),
 874 for any task q , there exists some \mathcal{G}^* implementing
 875 the corresponding cognitive program. By structural com-
 876 pleteness established above, there exists an action sequence
 877 constructing \mathcal{G}^* .

878 In summary, functional coverage ensures that the operator
 879 library implements all cognitive primitives required for goal-
 880 directed problem solving, while structural completeness
 881 ensures that the action space can construct any workflow
 882 graph composed of these primitives with arbitrary control
 883 structures. Together, they establish that for any task decom-
 884 posable into cognitive primitive combinations, there exists
 885 an action sequence constructing a workflow that implements
 886 the complete cognitive control loop covering planning, de-
 887 composition, execution, monitoring, revision, integration,
 888 and output. \square

B.2. Proof of Proposition 2

Proposition 5 (Monotonic Improvement of Multi-turn Interaction). *Under Assumptions 2–3, multi-turn orchestration based on canvas feedback possesses the monotonic improvement property in expectation: as the number of turns t increases, (1) the agent’s uncertainty about the correct output or feasible workflow does not increase; (2) consequently, the success probability of generating an executable and correct workflow (and final answer) does not decrease. Furthermore, multi-turn interaction under the same budget is at least as good as single-turn open-loop generation.*

Proof. We characterize the value of multi-turn interaction using a Bayesian risk potential function: each turn obtains canvas feedback o_t^{exec} that provides information about the hidden correct solution, thereby concentrating the posterior distribution. The monotonic decrease of the potential function in expectation corresponds to the monotonic increase of reliability and success rate.

(i) Posterior distribution and Bayesian risk potential.

Let Z denote the latent variable related to the task solution. According to the problem setting, Z can be taken as the correct answer y_q^* , or as a representative element from the set of feasible workflows that lead to the correct answer. At interaction turn t , the history is defined as

$$H_t = \{(a_1^{\text{think}}, a_1, o_1^{\text{exec}}), \dots, (a_t^{\text{think}}, a_t, o_t^{\text{exec}})\}, \quad (26)$$

where H_t encodes all reflection texts, edit actions, and canvas feedback up to turn t . The initial history is $H_0 = \emptyset$. Given history H_t , we define the posterior distribution over Z as

$$\pi_t(z) \triangleq \mathbb{P}(Z = z \mid H_t), \quad (27)$$

where $\pi_t(z)$ represents the probability that the true solution is z given the interaction history up to turn t . We define the Bayes accuracy function as the maximum posterior probability:

$$A(H_t) \triangleq \max_z \pi_t(z), \quad (28)$$

880 where $A(H_t)$ represents how concentrated the posterior is
 881 on the most likely solution. We then define the Bayes risk
 882 potential function as

$$883 \quad 884 \quad 885 \quad 886 \quad 887 \quad 888 \quad 889 \quad 890 \quad 891 \quad 892 \quad 893 \quad 894 \quad 895 \quad 896 \quad 897 \quad 898 \quad 899 \quad 900 \quad 901 \quad 902 \quad 903 \quad 904 \quad 905 \quad 906 \quad 907 \quad 908 \quad 909 \quad 910 \quad 911 \quad 912 \quad 913 \quad 914 \quad 915 \quad 916 \quad 917 \quad 918 \quad 919 \quad 920 \quad 921 \quad 922 \quad 923 \quad 924 \quad 925 \quad 926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad V(H_t) \triangleq 1 - A(H_t) = 1 - \max_z \pi_t(z), \quad (29)$$

where $V(H_t)$ measures the remaining uncertainty. Smaller $V(H_t)$ indicates more concentrated posterior and higher probability of the most likely correct solution.

(ii) Supermartingale property of the risk potential. We now show that the expected risk potential does not increase across turns, i.e., for any $t \geq 1$,

$$900 \quad 901 \quad 902 \quad 903 \quad 904 \quad 905 \quad 906 \quad 907 \quad 908 \quad 909 \quad 910 \quad 911 \quad 912 \quad 913 \quad 914 \quad 915 \quad 916 \quad 917 \quad 918 \quad 919 \quad 920 \quad 921 \quad 922 \quad 923 \quad 924 \quad 925 \quad 926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad \mathbb{E}[V(H_t) | H_{t-1}] \leq V(H_{t-1}), \quad (30)$$

with strict inequality when the feedback at turn t provides non-zero information gain about Z .

The key observations are twofold. First, the posterior vector satisfies the martingale property. Since $H_t = H_{t-1} \oplus (a_t^{\text{think}}, a_t, o_t^{\text{exec}})$ is obtained by appending new observations to H_{t-1} , by Bayes' rule, for any z we have

$$900 \quad 901 \quad 902 \quad 903 \quad 904 \quad 905 \quad 906 \quad 907 \quad 908 \quad 909 \quad 910 \quad 911 \quad 912 \quad 913 \quad 914 \quad 915 \quad 916 \quad 917 \quad 918 \quad 919 \quad 920 \quad 921 \quad 922 \quad 923 \quad 924 \quad 925 \quad 926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad \mathbb{E}[\pi_t(z) | \mathcal{F}_{t-1}] = \pi_{t-1}(z), \quad (31)$$

where $\mathcal{F}_{t-1} = \sigma(H_{t-1})$ is the natural filtration generated by the interaction history up to turn $t-1$. This martingale property states that the expected posterior at turn t , conditioned on information at turn $t-1$, equals the posterior at turn $t-1$.

Second, $V(\cdot)$ is a concave function over distributions. To see this, consider the potential function over the probability simplex:

$$900 \quad 901 \quad 902 \quad 903 \quad 904 \quad 905 \quad 906 \quad 907 \quad 908 \quad 909 \quad 910 \quad 911 \quad 912 \quad 913 \quad 914 \quad 915 \quad 916 \quad 917 \quad 918 \quad 919 \quad 920 \quad 921 \quad 922 \quad 923 \quad 924 \quad 925 \quad 926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad \phi(p) \triangleq 1 - \max_z p_z, \quad p \in \Delta_{|Z|-1}, \quad (32)$$

where p is a probability vector over the solution space, and $\Delta_{|Z|-1}$ is the probability simplex of such vectors. Since $\max_z p_z$ is a convex function of p (as the pointwise maximum of linear functions), its negation $- \max_z p_z$ is concave, and hence $\phi(p) = 1 - \max_z p_z$ is also concave.

Applying Jensen's inequality for concave functions, we obtain the contraction of expected risk:

$$900 \quad 901 \quad 902 \quad 903 \quad 904 \quad 905 \quad 906 \quad 907 \quad 908 \quad 909 \quad 910 \quad 911 \quad 912 \quad 913 \quad 914 \quad 915 \quad 916 \quad 917 \quad 918 \quad 919 \quad 920 \quad 921 \quad 922 \quad 923 \quad 924 \quad 925 \quad 926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad \mathbb{E}[V(\pi_t) | \mathcal{F}_{t-1}] \leq V(\mathbb{E}[\pi_t | \mathcal{F}_{t-1}]) = V(\pi_{t-1}), \quad (33)$$

where the equality uses the martingale property established above. This yields

$$900 \quad 901 \quad 902 \quad 903 \quad 904 \quad 905 \quad 906 \quad 907 \quad 908 \quad 909 \quad 910 \quad 911 \quad 912 \quad 913 \quad 914 \quad 915 \quad 916 \quad 917 \quad 918 \quad 919 \quad 920 \quad 921 \quad 922 \quad 923 \quad 924 \quad 925 \quad 926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad \mathbb{E}[V(H_t) | H_{t-1}] \leq V(H_{t-1}). \quad (34)$$

When the feedback provides non-zero information gain about Z (i.e., under Assumption 2), the posterior π_t undergoes strict contraction with non-zero probability, meaning it does not merely preserve its mean but actually becomes

more concentrated. In this case, Jensen's inequality becomes strict in expectation, yielding

$$900 \quad 901 \quad 902 \quad 903 \quad 904 \quad 905 \quad 906 \quad 907 \quad 908 \quad 909 \quad 910 \quad 911 \quad 912 \quad 913 \quad 914 \quad 915 \quad 916 \quad 917 \quad 918 \quad 919 \quad 920 \quad 921 \quad 922 \quad 923 \quad 924 \quad 925 \quad 926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad \mathbb{E}[V(H_t) | H_{t-1}] < V(H_{t-1}) \quad (35)$$

with positive probability, leading to strict improvement.

(iii) Monotonic improvement over multiple turns. Taking unconditional expectation and iterating the supermartingale relation yields:

$$900 \quad 901 \quad 902 \quad 903 \quad 904 \quad 905 \quad 906 \quad 907 \quad 908 \quad 909 \quad 910 \quad 911 \quad 912 \quad 913 \quad 914 \quad 915 \quad 916 \quad 917 \quad 918 \quad 919 \quad 920 \quad 921 \quad 922 \quad 923 \quad 924 \quad 925 \quad 926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad \mathbb{E}[V(H_t)] \leq \mathbb{E}[V(H_{t-1})] \leq \dots \leq \mathbb{E}[V(H_0)]. \quad (36)$$

This establishes that the expected Bayes risk monotonically decreases (or stays constant) as the number of interaction turns increases.

To quantify the improvement, we define the accuracy gain at turn t as

$$900 \quad 901 \quad 902 \quad 903 \quad 904 \quad 905 \quad 906 \quad 907 \quad 908 \quad 909 \quad 910 \quad 911 \quad 912 \quad 913 \quad 914 \quad 915 \quad 916 \quad 917 \quad 918 \quad 919 \quad 920 \quad 921 \quad 922 \quad 923 \quad 924 \quad 925 \quad 926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad \Delta_t \triangleq \mathbb{E}[V(H_{t-1})] - \mathbb{E}[V(H_t)] \geq 0, \quad (37)$$

where Δ_t represents the expected one-step reduction of the Bayes risk potential at turn t . Summing over all turns, the expected Bayes risk after t turns satisfies:

$$900 \quad 901 \quad 902 \quad 903 \quad 904 \quad 905 \quad 906 \quad 907 \quad 908 \quad 909 \quad 910 \quad 911 \quad 912 \quad 913 \quad 914 \quad 915 \quad 916 \quad 917 \quad 918 \quad 919 \quad 920 \quad 921 \quad 922 \quad 923 \quad 924 \quad 925 \quad 926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad \mathbb{E}[V(H_t)] = \mathbb{E}[V(H_0)] - \sum_{s=1}^t \Delta_s, \quad (38)$$

where each Δ_s accumulates the expected risk decrease at turn s . Substituting into the definition of accuracy, we obtain:

$$900 \quad 901 \quad 902 \quad 903 \quad 904 \quad 905 \quad 906 \quad 907 \quad 908 \quad 909 \quad 910 \quad 911 \quad 912 \quad 913 \quad 914 \quad 915 \quad 916 \quad 917 \quad 918 \quad 919 \quad 920 \quad 921 \quad 922 \quad 923 \quad 924 \quad 925 \quad 926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad \mathbb{E}[A(H_t)] = 1 - \mathbb{E}[V(H_0)] + \sum_{s=1}^t \Delta_s, \quad (39)$$

where $\mathbb{E}[A(H_t)]$ represents the expected Bayes accuracy after t turns. This shows that expected accuracy monotonically increases with the number of turns.

(iv) From uncertainty reduction to error probability reduction. Let $\hat{z}_t = \arg \max_z \pi_t(z)$ be the Bayes optimal estimate at turn t , i.e., the solution with maximum posterior probability. Under 0-1 loss (where we incur loss 1 if wrong and 0 if correct), the Bayes optimal decision rule minimizes expected loss by choosing \hat{z}_t . The error probability of this estimator satisfies

$$900 \quad 901 \quad 902 \quad 903 \quad 904 \quad 905 \quad 906 \quad 907 \quad 908 \quad 909 \quad 910 \quad 911 \quad 912 \quad 913 \quad 914 \quad 915 \quad 916 \quad 917 \quad 918 \quad 919 \quad 920 \quad 921 \quad 922 \quad 923 \quad 924 \quad 925 \quad 926 \quad 927 \quad 928 \quad 929 \quad 930 \quad 931 \quad 932 \quad 933 \quad 934 \quad \mathbb{P}(\hat{z}_t \neq Z | H_t) = 1 - \max_z \pi_t(z) = V(H_t). \quad (40)$$

This equality follows because the probability of error under the Bayes optimal rule equals one minus the probability of the chosen class, which is precisely $1 - \max_z \pi_t(z)$.

Therefore, monotonically non-increasing $\mathbb{E}[V(H_t)]$ is equivalent to monotonically non-increasing expected error rate, or equivalently, monotonically non-decreasing expected correctness rate. This proves conclusions (1) and (2) of the proposition.

935
 936
 937
 938
 939
 940
 941
 942
 943
 944
 945
 946
(v) Comparison with single-turn generation. A single-turn strategy makes only one decision and directly outputs the final workflow or answer, which is equivalent to setting $T = 1$. A multi-turn strategy with $T > 1$ can still choose to execute the `finish` action at the first turn, thereby terminating immediately. Therefore, the strategy space of multi-turn interaction contains the single-turn strategy as a special case.

947
 948
 949
 950
 951
 952
 953
 954
 955
 956
 957
 958
 959
 By this inclusion relationship of strategy spaces, the optimal value achievable by multi-turn interaction is no worse than that of single-turn interaction:

$$\max_{\pi \in \Pi_{\text{multi-turn}}} \mathbb{E}[A(H_T)] \geq \max_{\pi \in \Pi_{\text{single-turn}}} \mathbb{E}[A(H_1)], \quad (41)$$

960
 961
 962
 963
 964
 965
 966
 967
 968
 969
 where $\Pi_{\text{multi-turn}} \supseteq \Pi_{\text{single-turn}}$ denotes the respective policy spaces. This establishes that multi-turn is at least as good as single-turn under the same budget.

970
 971
 972
 973
 974
 975
 976
 977
 Furthermore, under Assumptions 2 (informative feedback) and 3 (repairability), multi-turn interaction can achieve strict improvement: in some turns, the posterior strictly contracts, leading to $\mathbb{E}[V(H_t)] < \mathbb{E}[V(H_{t-1})]$ on the training and inference distribution. This manifests as higher execution success rate, lower structural error rate, and higher final correctness rate.

978
 979
 980
 981
 982
 983
 984
 985
 986
 987
 988
 In conclusion, multi-turn canvas-based interaction monotonically decreases the Bayes risk potential whenever feedback is informative, consequently increasing expected accuracy. The agent can progressively refine the workflow based on accumulated observations, each turn potentially reducing uncertainty about the correct solution. This iterative refinement process achieves higher reliability and success probability than single-turn open-loop generation, where errors cannot be detected or corrected. \square

B.3. Proof of Proposition 3

989
 990
Proposition 6 (Structural Constraints, Conditional Release, and Mask Effectiveness). *Let the trajectory-level reward be defined as*

$$R(\tau) = -1 + R_{\text{diversity}}(\tau) + \mathbb{I}\{R_{\text{diversity}}(\tau) = 1\} \cdot R_{\text{answer}}(\tau), \quad (42)$$

991
 992
 993
 994
 995
 996
 997
 998
 999
 where $R_{\text{diversity}}(\tau) \in [0, 1]$ is composed of structural check items and capped at 1, and $R_{\text{answer}}(\tau) \geq 0$ measures the match between the final execution output and the ground truth. In advantage-based policy optimization such as CWRPO, this design possesses three properties: (a) separable feasibility learning, where trajectories not satisfying structural diversity constraints necessarily receive non-positive returns and are systematically suppressed by gradient updates; (b) shortcut and collapse suppression, where answer rewards are unlocked only after learning to construct qualified skeletons, thereby avoiding shortcuts like

skipping workflows to answer directly; (c) gradient correctness of mask, where token-level mask backpropagates gradients only for policy-generated tokens, maintaining unbiased policy gradient estimates and significantly reducing variance and noise from environment feedback tokens, thereby stabilizing training.

1000
 1001
Proof. The essence of this proposition is that the reward design achieves numerical separation and gating between structural feasibility and answer correctness, making the optimization process naturally staged. Meanwhile, the mask ensures that gradients act only on policy-controllable parts, and the clipping mechanism with KL regularization bounds the policy update magnitude, keeping gradient signals clean and stable.

(i) Feasible skeleton set and sign separation of rewards. We first define the feasible skeleton trajectory set as

$$\mathcal{F} = \{\tau : R_{\text{diversity}}(\tau) = 1\}, \quad (43)$$

1002
 1003
 where \mathcal{F} represents the set of trajectories with necessary structural skeleton, including verification steps, formatting operators, sufficient operator count, and appropriate control structures. The complement set is defined as

$$\mathcal{F}^c = \{\tau : R_{\text{diversity}}(\tau) < 1\}, \quad (44)$$

1004
 1005
 where \mathcal{F}^c represents trajectories that fail to meet one or more structural requirements.

1006
 1007
 We now show that the reward is strictly separable on \mathcal{F} and \mathcal{F}^c with a sign gap. For any trajectory $\tau \in \mathcal{F}^c$, we have $R_{\text{diversity}}(\tau) < 1$, so the indicator function $\mathbb{I}\{R_{\text{diversity}}(\tau) = 1\}$ equals 0. Substituting into the reward formula:

$$R(\tau) = -1 + R_{\text{diversity}}(\tau) + 0 \cdot R_{\text{answer}}(\tau) = -1 + R_{\text{diversity}}(\tau). \quad (45)$$

1008
 1009
 Since $R_{\text{diversity}}(\tau) \in [0, 1]$ for $\tau \in \mathcal{F}^c$, we have

$$R(\tau) \in [-1, 0) \quad \text{for all } \tau \in \mathcal{F}^c. \quad (46)$$

1010
 1011
 This means all structurally non-compliant trajectories receive strictly negative rewards.

1012
 1013
 For any trajectory $\tau \in \mathcal{F}$, we have $R_{\text{diversity}}(\tau) = 1$, so the indicator function equals 1. Substituting into the reward formula:

$$R(\tau) = -1 + 1 + 1 \cdot R_{\text{answer}}(\tau) = R_{\text{answer}}(\tau). \quad (47)$$

1014
 1015
 Since $R_{\text{answer}}(\tau) \geq 0$ by definition, we have

$$R(\tau) \geq 0 \quad \text{for all } \tau \in \mathcal{F}. \quad (48)$$

990 This means all structurally compliant trajectories receive
 991 non-negative rewards. The sign separation is therefore complete:
 992 \mathcal{F}^c trajectories are strictly negative while \mathcal{F} trajectories are
 993 non-negative, achieving strong constraint separation
 994 at the numerical level.

995 **(ii) Two-stage optimization via conditional release and**
 996 **policy gradient.** We now analyze how the conditional re-
 997 lease mechanism creates a natural two-stage optimization
 998 process through the lens of policy gradient theory. Let π_θ
 999 denote the policy parameterized by θ , and let π_{ref} denote the
 1000 reference policy (typically the initial supervised fine-tuned
 1001 model). The CWRPO objective can be written as
 1002

$$J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} [R(\tau)] - \beta D_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}), \quad (49)$$

1003 where $\beta > 0$ is the KL penalty coefficient that prevents the
 1004 policy from deviating too far from the reference distribution.
 1005

1006 By the policy gradient theorem, the gradient of the expected
 1007 reward with respect to θ is given by
 1008

$$\nabla_\theta \mathbb{E}_{\tau \sim \pi_\theta} [R(\tau)] = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=1}^T \nabla_\theta \log \pi_\theta(a_t | s_t) A(\tau) \right], \quad (50)$$

1009 where $A(\tau)$ is the advantage function and the sum is over
 1010 all time steps in the trajectory. In practice, the advantage is
 1011 computed using group normalization within each sampled
 1012 batch:
 1013

$$\hat{A}(\tau) = \frac{R(\tau) - \mu_{\mathcal{B}}}{\sigma_{\mathcal{B}} + \epsilon}, \quad (51)$$

1014 where $\mu_{\mathcal{B}}$ and $\sigma_{\mathcal{B}}$ are the mean and standard deviation of
 1015 rewards in batch \mathcal{B} , and ϵ is a small constant for numerical
 1016 stability.
 1017

1018 From the sign separation established in (i), when both \mathcal{F}
 1019 and \mathcal{F}^c samples exist in a batch, the normalized advantages
 1020 satisfy
 1021

$$\mathbb{E}[\hat{A}(\tau) | \tau \in \mathcal{F}] > 0 > \mathbb{E}[\hat{A}(\tau) | \tau \in \mathcal{F}^c]. \quad (52)$$

1022 This inequality holds because feasible trajectories have non-
 1023 negative rewards while non-feasible ones have strictly nega-
 1024 tive rewards, so after mean-centering, feasible trajectories
 1025 lie above the mean and non-feasible ones lie below.
 1026

1027 The policy gradient update therefore increases the log-
 1028 probability of actions in feasible trajectories and decreases
 1029 the log-probability of actions in non-feasible trajectories:
 1030

$$\theta \leftarrow \theta + \alpha \nabla_\theta J(\theta), \quad (53)$$

1031 where α is the learning rate. This systematically shifts
 1032 probability mass from \mathcal{F}^c to \mathcal{F} .
 1033

1034 To quantify this effect, define the feasibility probability
 1035 $p_\theta = \mathbb{P}_{\tau \sim \pi_\theta} (\tau \in \mathcal{F})$. Under mild regularity conditions, the
 1036

1037 gradient update increases p_θ whenever $p_\theta < 1$:
 1038

$$\frac{dp_\theta}{d\theta} \cdot \nabla_\theta J(\theta) > 0 \quad \text{when } 0 < p_\theta < 1. \quad (54)$$

1039 This follows because the expected advantage is positive for
 1040 \mathcal{F} and negative for \mathcal{F}^c , so the gradient points in the direction
 1041 of increasing p_θ .

1042 During early training when p_θ is small, most trajectories
 1043 fall into \mathcal{F}^c , and the primary learning signal comes from
 1044 avoiding structural violations. As p_θ increases through training,
 1045 an increasing proportion of sampled trajectories fall into \mathcal{F} . For these
 1046 trajectories, $R(\tau) = R_{\text{answer}}(\tau)$, and the
 1047 training signal naturally shifts to optimizing answer
 1048 correctness. This creates the two-stage behavior: first learn to
 1049 satisfy structural constraints, then learn to maximize answer
 1050 quality.

1051 The conditional release mechanism ensures that shortcuts
 1052 are suppressed: any trajectory that directly outputs an an-
 1053 swer without constructing a proper workflow will have
 1054 $R_{\text{diversity}}(\tau) < 1$ and thus $R(\tau) < 0$, regardless of an-
 1055 swer correctness. This negative reward signal systematically
 1056 discourages such shortcuts.

1057 **(iii) Bounded policy updates via clipping and KL regular-
 1058 ization.** To ensure training stability, CWRPO incorporates
 1059 two mechanisms that bound the magnitude of policy updates:
 1060 importance ratio clipping and KL divergence regularization.

1061 For importance sampling, define the probability ratio be-
 1062 tween the current policy and the behavior policy (used for
 1063 sampling) as
 1064

$$\rho_\theta(\tau) = \frac{\pi_\theta(\tau)}{\pi_{\text{old}}(\tau)} = \prod_{t=1}^T \frac{\pi_\theta(a_t | s_t)}{\pi_{\text{old}}(a_t | s_t)}, \quad (55)$$

1065 where π_{old} is the policy at the beginning of the current optimi-
 1066 zation epoch. The clipped objective restricts the effective
 1067 ratio to the interval $[1 - \epsilon_{\text{clip}}, 1 + \epsilon_{\text{clip}}]$:

$$L^{\text{clip}}(\theta) = \mathbb{E}_{\tau \sim \pi_{\text{old}}} \left[\min \left(\rho_\theta(\tau) \hat{A}(\tau), \right. \right. \quad (56) \\ \left. \left. \text{clip}(\rho_\theta, 1 - \epsilon, 1 + \epsilon) \hat{A}(\tau) \right) \right].$$

1068 This clipping prevents excessively large policy updates
 1069 when the importance ratio deviates significantly from 1,
 1070 which would otherwise destabilize training.

1071 The KL regularization term provides a soft constraint on the
 1072 policy deviation:
 1073

$$D_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=1}^T \log \frac{\pi_\theta(a_t | s_t)}{\pi_{\text{ref}}(a_t | s_t)} \right]. \quad (57)$$

1074 By penalizing this divergence, the optimization ensures that
 1075 the learned policy remains close to the reference policy, pre-
 1076 venting catastrophic forgetting of useful behaviors learned
 1077 during supervised fine-tuning.

1045 Together, clipping and KL regularization bound the per-step
 1046 policy change:

$$\|\theta_{k+1} - \theta_k\| \leq \frac{\alpha \cdot \epsilon_{\text{clip}} \cdot \max_{\tau} |A(\tau)|}{\beta + \lambda_{\min}(\nabla_{\theta}^2 D_{\text{KL}})}, \quad (58)$$

1050 where λ_{\min} denotes the minimum eigenvalue of the KL
 1051 Hessian. This bounded update magnitude prevents training
 1052 oscillations and ensures smooth convergence.

1053 **(iv) Unbiasedness and variance reduction of token-level
 1054 mask.** An interaction trajectory at the token level consists
 1055 of two interleaved types of segments: policy-generated to-
 1056 kens comprising reflection and action text, and environment
 1057 feedback tokens comprising canvas returns. Crucially, the
 1058 generation distribution of environment feedback does not
 1059 contain the learnable parameters θ ; it is determined by the
 1060 canvas environment given the action.

1061 Let $w_{1:|\tau|}$ denote the entire trajectory token sequence. We
 1062 partition it into two disjoint sets: \mathcal{T}_{π} for policy-generated
 1063 tokens and \mathcal{T}_{env} for environment feedback tokens. The joint
 1064 log-likelihood of the trajectory can be decomposed as

$$\begin{aligned} \log p_{\theta}(w_{1:|\tau|}) &= \sum_{t \in \mathcal{T}_{\pi}} \log \pi_{\theta}(w_t | w_{<t}) \\ &\quad + \sum_{t \in \mathcal{T}_{\text{env}}} \log p_{\text{env}}(w_t | w_{<t}), \end{aligned} \quad (59)$$

1065 where the first term sums over policy tokens and depends
 1066 on θ , while the second term sums over environment tokens
 1067 and does not depend on θ .

1068 Taking the gradient with respect to θ , the second term van-
 1069 ishes since $\nabla_{\theta} \log p_{\text{env}}(w_t | w_{<t}) = 0$:

$$\nabla_{\theta} \log p_{\theta}(w_{1:|\tau|}) = \sum_{t \in \mathcal{T}_{\pi}} \nabla_{\theta} \log \pi_{\theta}(w_t | w_{<t}). \quad (60)$$

1070 Let $\text{mask}_t \in \{0, 1\}$ be the token-level mask that takes value
 1071 1 for policy tokens ($t \in \mathcal{T}_{\pi}$) and 0 for environment tokens
 1072 ($t \in \mathcal{T}_{\text{env}}$). Then the gradient can be written equivalently as

$$\nabla_{\theta} \log p_{\theta}(w_{1:|\tau|}) = \sum_{t=1}^{|\tau|} \text{mask}_t \cdot \nabla_{\theta} \log \pi_{\theta}(w_t | w_{<t}). \quad (61)$$

1073 Using the mask to select only policy tokens \mathcal{T}_{π} implements
 1074 this equality exactly. Therefore, the masked gradient esti-
 1075 mator is unbiased:

$$\mathbb{E} \left[\sum_{t=1}^{|\tau|} \text{mask}_t \cdot \nabla_{\theta} \log \pi_{\theta}(w_t | w_{<t}) \cdot R(\tau) \right] = \nabla_{\theta} J(\theta). \quad (62)$$

1076 Furthermore, the mask reduces gradient variance. Without
 1077 the mask, if all tokens were naively included, the gradient
 1078

1079 estimator would have variance

$$\text{Var}_{\text{no-mask}} = \text{Var} \left[\sum_{t=1}^{|\tau|} \nabla_{\theta} \log \pi_{\theta}(w_t | w_{<t}) \cdot R(\tau) \right]. \quad (63)$$

1080 With the mask, the variance is

$$\text{Var}_{\text{mask}} = \text{Var} \left[\sum_{t \in \mathcal{T}_{\pi}} \nabla_{\theta} \log \pi_{\theta}(w_t | w_{<t}) \cdot R(\tau) \right]. \quad (64)$$

1081 Since $|\mathcal{T}_{\pi}| < |\tau|$ and the environment tokens contribute
 1082 noise unrelated to θ , we have $\text{Var}_{\text{mask}} < \text{Var}_{\text{no-mask}}$. This
 1083 variance reduction leads to more stable gradient updates and
 1084 faster convergence.

1085 In conclusion, the reward design and optimization frame-
 1086 work achieve four complementary goals for stable and
 1087 effective learning. First, sign separation numerically en-
 1088 forces strong constraint separation between feasible and
 1089 non-feasible trajectories, ensuring that structural violations
 1090 are systematically penalized with negative rewards. Second,
 1091 conditional release creates a natural two-stage optimization
 1092 where the policy first learns to satisfy structural constraints
 1093 (increasing $p_{\theta} = \mathbb{P}(\tau \in \mathcal{F})$) before optimizing answer cor-
 1094 rectness, thereby suppressing shortcuts and preventing struc-
 1095 tural collapse. Third, clipping and KL regularization bound
 1096 the magnitude of policy updates, preventing training oscilla-
 1097 tions and ensuring smooth convergence while maintaining
 1098 proximity to the reference policy. Fourth, the token-level
 1099 mask ensures that gradients act only on policy-controllable
 1099 tokens, maintaining unbiased gradient estimates while reduc-
 1099 ing variance from environment feedback, thereby stabilizing
 1099 the training process. Together, these mechanisms provide
 1099 stable learning signals, bounded updates, and effective pre-
 1099 vention of shortcut behaviors, enabling reliable end-to-end
 1099 reinforcement learning for workflow orchestration. \square

C. Flow-Steer Algorithm Details

Flow-Steer is a multi-turn workflow orchestration framework built on the collaboration between a Flow-Director (small-scale LLM) and a Workflow Canvas (execution environment): the Flow-Director handles planning, action generation, and answer output, while the Workflow Canvas provides structural feedback and executes the constructed workflow. The process is divided into three connected stages: first, given a task description, the Flow-Director constructs workflow nodes with structural guidance using the operator library, and the Canvas returns execution feedback to form the initial interaction history; then, during multi-turn interaction, the Flow-Director continuously generates reasoning and the next action from the accumulated history, the Canvas validates and executes each action to extend the history until the termination condition is reached, and

1100 **Algorithm 1** Flow-Steer: End-to-End Workflow Orchestration via Multi-Turn Reinforcement Learning

1101 **Require:** Task q , Flow-Director π_θ , Workflow Canvas \mathcal{C} , operator library \mathcal{O} , reward function $R(\cdot)$, maximum interaction

1102 turns T

1103 **Ensure:** Final answer y

1104 1: // Stage A: Workflow Initialization

1105 2: Initialize: workflow graph $\mathcal{G}_0 = \emptyset$, interaction history $H_0 = []$

1106 3: Construct system prompt: $p_{\text{sys}} \leftarrow \text{BUILD_PROMPT}(\mathcal{O}, q)$

1107 4: First think and action: $(a_1^{\text{think}}, a_1) \leftarrow \pi_\theta(p_{\text{sys}}, q)$

1108 5: First canvas feedback: $o_1 \leftarrow \mathcal{C}.\text{STEP}(a_1)$

1109 6: First update history: $H_1 \leftarrow \{(a_1^{\text{think}}, a_1, o_1)\}$

1110 7: // Stage B: Multi-turn Collaborative Workflow Building

1111 8: **for** $t = 2$ **to** T **do**

1112 9: Plan think: $a_t^{\text{think}} \leftarrow \pi_\theta(p_{\text{sys}}, q, H_{t-1})$

1113 10: Generate action: $a_t \leftarrow \pi_\theta(p_{\text{sys}}, q, H_{t-1})$

1114 11: Canvas feedback: $o_t \leftarrow \mathcal{C}.\text{STEP}(a_t)$

1115 12: Update history: $H_t \leftarrow H_{t-1} \cup \{(a_t^{\text{think}}, a_t, o_t)\}$

1116 13: **if** $a_t = \text{finish}$ **and** $\mathcal{C}.\text{CHECK_CONSTRAINTS}()$ **then**

1117 14: **break**

1118 15: **end if**

1119 16: **end for**

1120 17: Final think & answer: $(a_T^{\text{think}}, y) \leftarrow \pi_\theta(p_{\text{sys}}, q, H_T)$

1121 18: Execute workflow: $y \leftarrow \mathcal{C}.\text{EXECUTE}(\mathcal{G}_T, q)$

1122 19: // Stage C: End-to-End Reinforcement Learning Optimization

1123 20: Sample trajectories: $\{\tau_i\}_{i=1}^N \sim \pi_\theta$

1124 21: **for** each τ_i **do**

1125 22: Compute reward: $R(\tau_i) = -1 + R_{\text{diversity}}(\tau_i) + \mathbb{I}\{R_{\text{diversity}}(\tau_i) = 1\} \cdot R_{\text{answer}}(\tau_i)$

1126 23: Compute advantage: $\hat{A}(\tau_i) = \frac{R(\tau_i) - \bar{R}}{\sqrt{\frac{1}{M} \sum_{j=1}^M (R(\tau_j) - \bar{R})^2 + \epsilon}}$

1127 24: **end for**

1128 25: CWRPO-based update policy:

1129 26: $\mathcal{J}_{\text{CWRPO}} \propto \sum_{i=1}^N \sum_{t=1}^{|\tau_i|} \text{MASK}_t^{(i)} \cdot \min \left(\rho_\theta(w_t^{(i)}), \text{CLIP}(\rho_\theta(w_t^{(i)}), 1 \pm \epsilon) \right) \hat{A}(\tau_i)$

1130 27: where $\rho_\theta(w_t^{(i)}) = \frac{\pi_\theta(w_t^{(i)} | \tau_{<t}^{(i)})}{\pi_{\theta_{\text{old}}}(w_t^{(i)} | \tau_{<t}^{(i)})}$

1131 28: Parameter update: $\theta \leftarrow \theta - \eta \nabla_\theta (-\mathcal{J}_{\text{CWRPO}})$

1132 29: **return** y

1133 the accumulated trajectory is utilized to execute the final workflow and generate the answer; finally, end-to-end reinforcement learning is applied to update the policy of the Flow-Director. The reward function jointly evaluates structural diversity compliance and answer correctness, enabling the Flow-Director to gradually learn how to effectively construct complex workflows within a limited budget. This coordination reduces invalid actions and stabilizes interaction dynamics. This design enables progressive and adaptive workflow construction, resulting in improved accuracy and stability on complex reasoning tasks.

1134 **Training and Inference Flow.** During training, the algorithm proceeds in three stages ($A \rightarrow B \rightarrow C$): the Flow-Director initializes the system prompt and triggers the Canvas for an initial feedback, then enters multi-turn interaction to generate actions, receive feedback, and terminate with

1135 an answer in a sequential manner, and finally updates the policy in Stage C using rewards and advantages. During testing, it runs only two stages ($A \rightarrow B$) in a simplified form: initialization and multi-turn interaction, after which the final workflow is executed and the answer is produced directly without parameter updates.

1136 **Complexity Analysis.** The computational complexity of Flow-Steer mainly comes from initialization, multi-turn interaction, and reinforcement learning optimization. The initialization stage involves one call to the Flow-Director and a call to the Canvas, which is a constant overhead. The multi-turn interaction stage requires up to T rounds in the worst case, where each round includes one planning step by the Flow-Director and one call to the Canvas, yielding time complexity $O(T)$. The memory consumption grows linearly with the history length, which can be controlled

1155 through windowing or summarization. The reinforcement
 1156 learning stage requires sampling N trajectories per update,
 1157 each trajectory containing up to T action-feedback steps,
 1158 leading to complexity $O(NT)$. It also requires storing
 1159 trajectory information for reward and advantage computation.
 1160 In total, the complexity of Flow-Steer is $O(NT+T)$, scaling
 1161 linearly with the number of rounds and sampled trajectories
 1162 during training, while inference requires $O(T)$. Since the
 1163 Flow-Director is responsible for planning and constructing
 1164 workflows, the Canvas execution is more focused, ensuring
 1165 stability on complex reasoning tasks.

1167 C.1. Reward Component Weights

1168 The diversity constraint reward $R_{\text{diversity}}(\tau)$ is composed
 1169 of four binary checks, each contributing equally to the total
 1170 score (capped at 1.0):

- 1173 • R_{checker} **(0.25)**: Encourages the inclusion of verifi-
 1174 cation operators (Test, Review, Verify). Set to
 1175 0.25 if at least one verification operator is present in
 1176 the workflow.
- 1177 • R_{format} **(0.25)**: Encourages proper answer formatting.
 1178 Set to 0.25 if the Format operator is included as the
 1179 final step before termination.
- 1181 • R_{operator} **(0.25)**: Requires a minimum operator count
 1182 for structural diversity. Set to 0.25 if the workflow
 1183 contains at least 3 distinct operators.
- 1185 • R_{control} **(0.25)**: Encourages control flow structures.
 1186 Set to 0.25 if the workflow includes at least one control
 1187 structure (parallel, conditional, or loop).

1188 This design ensures that workflows must exhibit structural
 1189 diversity (achieving $R_{\text{diversity}} = 1.0$) before the answer
 1190 reward R_{answer} is released, effectively preventing shortcut
 1191 behaviors where the policy learns to generate overly simple
 1192 or degenerate workflows.

1195 D. Dataset Details

1197 We selected 12 public datasets (including mathematical
 1198 reasoning, question answering, and code generation) for train-
 1199 ing and testing. Six of these datasets were used for training
 1200 and testing. Six datasets were used for out-of-distribution
 1201 testing to verify the generalization of the proposed Flow-
 1202 Steer. These datasets are as follows:

- 1203 • **GSM8K** (Cobbe et al., 2021): A collection of grade-
 1204 school math word problems with concise statements, em-
 1205 phasizing step-by-step calculation and accurate numeric
 1206 results. Problems involve basic operations with real-world
 1207 contexts such as shopping, time calculations, and quantity
 1208 comparisons.

- **MATH** (Hendrycks et al., 2021b): A dataset of competition-level mathematics problems from AMC, AIME, and other mathematical olympiads. Problems span seven categories including algebra, geometry, number theory, combinatorics, probability, precalculus, and intermediate algebra, requiring sophisticated mathematical reasoning and symbolic manipulation.

- **HotPotQA** (Yang et al., 2018): A large-scale multi-hop question answering dataset requiring reasoning across multiple Wikipedia paragraphs. Questions are designed to require finding and combining information from different sources, with supporting facts annotated for interpretability.

- **SQuAD v2** (Rajpurkar et al., 2018): A Wikipedia-based QA dataset combining answerable and unanswerable questions, constructed to evaluate comprehension under mixed conditions. This tests both reading comprehension and the ability to recognize insufficient information.

- **MBPP** (Austin et al., 2021): Mostly Basic Python Problems, a crowd-sourced collection of Python programming problems with natural language descriptions and test cases. Problems range from simple string manipulation to basic algorithms, designed to test fundamental programming skills.

- **HumanEval** (Chen et al., 2021): A hand-written collection of Python programming problems with function signatures, docstrings, and unit tests. Problems are designed to test functional correctness through execution, covering tasks like string processing, mathematical operations, and data structure manipulation.

- **TriviaQA** (Joshi et al., 2017): A knowledge-intensive dataset with questions from trivia websites and Wikipedia, containing a wide range of facts and lesser-known topics. The dataset covers diverse domains including history, science, geography, and entertainment.

- **NaturalQuestions** (Kwiatkowski et al., 2019): A dataset of real anonymized queries from Google Search with answers from Wikipedia articles. Questions reflect genuine information needs of users, making them more diverse and challenging than synthetic questions.

- **MathQA** (Amini et al., 2019): A math word problem dataset curated from multi-domain exam problems, covering arithmetic, algebra, geometry, probability, and other sub-disciplines. Each problem includes annotated rationales explaining the solution steps.

- **AIME 2025**: Problems from the 2025 American Invitational Mathematics Examination, representing challenging competition-level mathematics. AIME problems require creative problem-solving and deep mathematical insight, with answers being integers from 0 to 999.

- **APPS** (Hendrycks et al., 2021a): Automated Program-

ming Progress Standard, a collection of competitive programming problems from Codeforces, Kattis, and other platforms. Problems range from introductory to competition-level difficulty, requiring algorithmic thinking and efficient implementation.

• **DS-1000** (Lai et al., 2023): A data science code generation benchmark covering NumPy, Pandas, TensorFlow, PyTorch, SciPy, Scikit-learn, and Matplotlib. Problems are derived from real StackOverflow questions, testing practical data science skills.

To ensure consistency and fairness for training and testing, we construct the training set by mixing samples from all IID datasets with the following sampling strategy: 2,560 samples from GSM8K, 2,560 from MATH, 2,560 from HotPotQA, 2,560 from SQuAD v2, 374 from MBPP (full set), and 164 from HumanEval (full set), resulting in a total of 10,778 training instances. To evaluate the generalization performance of Flow-Steer, 128 instances were randomly sampled from each of the six out-of-distribution and six trained datasets for testing, except for AIME 2025 which contains 30 problems.

E. Baseline Details

To accurately evaluate the performance of Flow-Steer, we conducted comparative experiments against multiple baselines. These baselines can be broadly divided into four categories: direct LLM inference, supervised fine-tuning methods, search-based workflow methods, and agent with reinforcement learning methods.

E.1. Direct LLM Baselines

• **Qwen3-8B** (Qwen Team, 2025): An 8-billion parameter language model from Alibaba Cloud, serving as the backbone for our Flow-Director. As a baseline, it is tested with zero-shot chain-of-thought prompting, measuring the model’s inherent reasoning capacity without workflow orchestration. The model provides strong efficiency while maintaining competitive performance on reasoning benchmarks, making it an ideal foundation for lightweight policy learning.

• **GPT-4o-mini** (OpenAI, 2024): A lightweight variant of GPT-4o optimized for cost and latency, while retaining strong language and reasoning abilities. As a baseline, it is tested with standard instruction prompting without workflow orchestration, measuring the model’s inherent generation capacity under constrained resources. It serves as the default backend for our Workflow Canvas, executing the actual reasoning operations specified by the workflow.

E.2. Fine-Tuning Baselines

- **SFT (Supervised Fine-Tuning)** (Ouyang et al., 2022): A supervised fine-tuning baseline built on Qwen3-8B, trained with workflow annotation data to improve instruction following and DSL generation accuracy. Unlike Flow-Steer’s multi-turn interaction paradigm, SFT generates the complete workflow in a single turn without execution feedback. We use LoRA fine-tuning with rank 8 for parameter efficiency, evaluating how standard supervised adaptation enhances the raw backbone’s workflow construction capabilities.

- **GRPO (Group Relative Policy Optimization)** (Shao et al., 2024b): Group Relative Policy Optimization is a reinforcement learning algorithm that normalizes rewards within sampled groups of trajectories. This reduces variance in policy updates, stabilizes training, and improves convergence efficiency compared to standard Proximal Policy Optimization (PPO). Unlike Flow-Steer’s multi-turn interaction with execution feedback, GRPO generates the entire workflow in one shot, limiting its ability to adapt based on intermediate results.

E.3. Search-Based Workflow Methods

- **AFlow** (Zhang et al., 2024b): A workflow optimization framework that uses Monte Carlo Tree Search (MCTS) to explore the workflow space. The method systematically searches over predefined operator combinations through tree expansion and backpropagation, evaluating candidate workflows by execution outcomes. While effective at finding good operator sequences, AFlow lacks end-to-end learning capability and cannot learn from accumulated experience across different problems. Its search process can be computationally expensive as the workflow complexity increases.

E.4. Agent with Reinforcement Learning Methods

- **AgentFlow**: An agent framework combined with PPO-based reinforcement learning for workflow construction. The agent dynamically selects tools from a predefined set at each step based on the current state, enabling adaptive decision-making through the interaction process. However, AgentFlow does not support custom prompt specification for operators, limiting its flexibility in fine-tuning the behavior of individual workflow components. This constraint reduces its ability to optimize task-specific reasoning strategies.

- **Router-R1**: A router-style architecture where a small policy model learns to route queries to different processing paths using GRPO for policy optimization. The router makes a single decision per query without iterative refinement, selecting from a predefined set of workflow templates. While this approach is computationally efficient, the single-shot routing mechanism cannot adapt to intermediate execution results or refine workflows based on partial feedback,

Table 9. Architectural comparison with baselines.

Method	Dynamic Orchestration	Multi-turn	Exec. Feedback	Custom Prompts	End-to-End RL	Pluggable Backend
Direct LLM	✗	✗	✗	✗	✗	✗
SFT/GRPO	✗	✗	✗	✗	Partial	✗
AFlow	✓	✓	✓	✓	✗	✓
AgentFlow	✗	✓	✓	✓	✓	✓
Router-R1	✗	✓	✓	✗	✓	✓
Orchestrator	✗	✓	Partial	✗	✓	✗
Flow-Steer (Ours)	✓	✓	✓	✓	✓	✓

Table 10. Baseline implementation details.

Method	Base Model	Training	Key Hyperparameters
Qwen3-8B	Qwen3-8B-Instruct	None	temperature=0, max_tokens=512
GPT-4o-mini	GPT-4o-mini	None	temperature=0, top_p=1, max_tokens=512
SFT	Qwen3-8B	LoRA	temperature=0, r=16, $\alpha=16$, LR= 1×10^{-4} , epochs=1, batch=16
GRPO	Qwen3-8B	GRPO + LoRA	temperature=0, r=16, $\alpha=32$, LR= 1×10^{-6} , G=4
AFlow	GPT-4o-mini	MCTS	temperature=0, search_rounds=21
AgentFlow	Qwen2.5-7B-Instruct	GRPO	temperature=0, LR= 1×10^{-6} , rollout_n=8, epochs=5
Router-R1	Qwen3-8B	PPO	temperature=0, LR= 1×10^{-6} , clip=0.2, epochs=30
Orchestrator	Qwen3-8B	GRPO/PPO	temperature=0, batch=256, max_turns=5

limiting its performance on complex multi-step reasoning tasks.

• **Orchestrator:** An orchestrator-style architecture that sequentially selects operators from a library using PPO. The orchestrator receives partial execution feedback to inform subsequent decisions, enabling some degree of adaptive behavior. However, it lacks the two-step interaction mechanism (add + set_prompt) that Flow-Steer employs for fine-grained control over operator configuration. This limitation prevents precise customization of individual operator behaviors, reducing the overall workflow quality.

Table 9 summarizes the key architectural differences between Flow-Steer and the baselines, and Table 10 provides the detailed implementation configurations for each method.

For a fair comparison, we conducted three independent runs under identical settings for both Flow-Steer and all baselines and reported the averaged results. For all baselines, we use the same evaluation protocol: 128 test samples per dataset (except AIME 2025 with 30 samples), single generation per sample, and task-specific metrics (EM/F1 for QA, Accuracy for Math, Pass@1 for Code).

E.5. Other LLM Backends

For transferability experiments (RQ3), we additionally evaluate on six LLM backends to assess the generalization of Flow-Director across different backend models:

• **DeepSeek-V3 (DeepSeek-AI, 2024):** DeepSeek-V3 adopts advanced context understanding algorithms, enabling

it to achieve excellent performance in long-context reasoning and multi-step inference tasks. Its powerful semantic understanding capabilities make it an ideal backend for testing how Flow-Director handles complex reasoning chains.

• **Grok-4-Fast:** Grok-4-Fast incorporates an efficient inference optimization mechanism, achieving a balance between generation speed and quality, making it suitable for latency-sensitive application scenarios. This backend helps us evaluate Flow-Director’s performance under speed-optimized conditions.

• **LLaMA-4-Maverick (Meta AI, 2025):** LLaMA-4-Maverick leverages the latest multimodal learning techniques, particularly excelling in tasks that require integrating diverse information sources. Its open-source nature allows for detailed analysis of workflow execution patterns.

• **Qwen-Plus:** Qwen-Plus incorporates an adaptive learning mechanism based on the Qwen architecture, excelling in few-shot scenarios and cross-language migration tasks. As a model from the same family as our Flow-Director backbone, it provides insights into intra-family transferability.

• **GPT-5:** GPT-5 achieves breakthroughs in natural language understanding and generation through enhanced data processing and training strategies, becoming one of the most versatile large language models. It serves as a strong upper bound for backend capability assessment.

• **Gemini-2.5-Flash (Google DeepMind, 2025):** Gemini-2.5-Flash combines a fast inference engine with multimodal understanding capabilities, optimized for real-time interac-

1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374

tive applications. Its unique architecture provides a diverse test case for backend generalization.

F. Evaluation Metrics

We use task-specific evaluation metrics following standard practices in each domain.

F.1. Exact Match (EM) for Question Answering

Exact Match measures whether the predicted answer exactly matches the ground truth after normalization:

$$\text{EM} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\text{normalize}(y_i) = \text{normalize}(y_i^*)), \quad (65)$$

where $\text{normalize}(\cdot)$ applies the following transformations: (1) convert to lowercase, (2) remove punctuation (except hyphens in compound words), (3) remove articles (“a”, “an”, “the”), (4) collapse multiple whitespaces to single space, and (5) strip leading/trailing whitespace.

Applicable datasets: HotPotQA, SQuAD v2, TriviaQA, NaturalQuestions

F.2. F1 Score for Question Answering

F1 Score measures token-level overlap between prediction and ground truth:

$$\text{Precision} = \frac{|y \cap y^*|}{|y|}, \quad (66)$$

$$\text{Recall} = \frac{|y \cap y^*|}{|y^*|}, \quad (67)$$

$$\text{F1} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}, \quad (68)$$

where y and y^* are the sets of tokens in the predicted and ground truth answers respectively, after applying the same normalization as EM.

Applicable datasets: HotPotQA, SQuAD v2, TriviaQA, NaturalQuestions

F.3. Accuracy for Mathematical Reasoning

Accuracy measures whether the predicted numerical answer matches the ground truth within a tolerance:

$$\text{Acc} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(|y_i - y_i^*| < \epsilon), \quad (69)$$

where $\epsilon = 10^{-6}$ is the numerical tolerance for floating-point comparisons.

For symbolic answers (e.g., fractions, algebraic expressions), we apply symbolic equivalence checking using

SymPy:

$$\text{Acc}_{\text{symbolic}} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\text{simplify}(y_i - y_i^*) = 0). \quad (70)$$

Applicable datasets: GSM8K, MATH, MathQA, AIME 2024

F.4. Pass@k for Code Generation

Pass@k measures the probability that at least one of k generated solutions passes all test cases:

$$\text{Pass}@k = \mathbb{E}_{\text{problems}} \left[1 - \frac{\binom{n-c}{k}}{\binom{n}{k}} \right], \quad (71)$$

where n is the total number of generated samples and c is the number of correct samples (passing all tests).

For our evaluation, we use Pass@1 with a single generation per problem:

$$\text{Pass}@1 = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\text{execute}(y_i, \text{tests}_i) = \text{pass}), \quad (72)$$

where $\text{execute}(y_i, \text{tests}_i)$ runs the generated code y_i against the test suite tests_i .

Applicable datasets: MBPP, HumanEval, APPS, DS-1000

G. Implementation Details

To ensure reproducibility and fair comparison, we summarize the complete hyperparameter configurations for Flow-Steer in Table 11. The table covers all aspects of our training pipeline, including model configuration, training hyperparameters, CWRPO algorithm settings, generation parameters, interaction constraints, reward function design, canvas backend configuration, and hardware setup.

Model Configuration. We use Qwen3-8B as the policy model (Flow-Director) and apply LoRA-based fine-tuning with rank 64 and alpha 64, targeting the query, key, value, and output projection layers (q_proj, k_proj, v_proj, o_proj) with a dropout rate of 0.05. The model is trained in bfloat16 precision with gradient checkpointing enabled for memory efficiency.

Training Configuration. We train the Flow-Director agent using the AdamW optimizer with a learning rate of 1×10^{-5} and weight decay of 0.01. The batch size is 36, computed as 6 samples per data source multiplied by 6 data sources (GSM8K, MATH, HotPotQA, SQuAD v2, MBPP, HumanEval). We train for 300 steps with a cosine learning rate schedule.

CWRPO Configuration. For the CWRPO algorithm, we use GRPO as the advantage estimator with a clip range

Category	Hyperparameter	Value
Model Configuration	Base Model	Qwen3-8B
	LoRA Rank / Alpha	64 / 64
	LoRA Target Modules	q_proj, k_proj, v_proj, o_proj
	LoRA Dropout	0.05
	Data Type	bfloat16
	Gradient Checkpointing	Enabled
Training	Batch Size	36 (6 samples \times 6 sources)
	Learning Rate	1×10^{-5}
	Optimizer	AdamW (weight decay 0.01)
	LR Schedule	Cosine
CWRPO Algorithm	Max Training Steps	300
	Advantage Estimator	GRPO
	Clip Range (ϵ)	0.20
	KL Coefficient (β)	0.005
Generation	Samples per Group (N)	36
	Temperature	0.6
	Top- p / Top- k	0.95 / 20
	Max New Tokens	2,048
	Enable Thinking Mode	True
Interaction	vLLM Max Concurrency	32
	Max Interaction Rounds (T_{\max})	20
	Max Context Length	16,384
	Min Operators for Finish	5
Reward	Require Checker/Structure	True / True
	Base Reward	-1.0
	Structure Reward Cap	1.0
	Structure Components	$R_{\text{chk}}=0.2, R_{\text{fmt}}=0.2, R_{\text{op}}=0.2, R_{\text{ctrl}}=0.4$
Canvas Backend	Correctness Activation	Gate (structure = 1.0)
	Executor Model	GPT-OSS-120B (temp=0)
Hardware	Execution Timeout	600s
	GPU Type	NVIDIA A100 80GB \times 2
	CUDA / vLLM LoRA	12.5 / Enabled

Table 11. Complete hyperparameter settings for FlowSteer training.

of 0.20, KL coefficient of 0.005, and entropy coefficient of 0.005. Each training step samples 36 trajectories for group-relative advantage estimation.

Generation Configuration. During trajectory generation, we use temperature 0.6, top- p 0.95, and top- k 20 following Qwen3’s recommended parameters. The maximum new tokens per turn is set to 2,048 to allow sufficient space for thinking and action generation. We enable Qwen3’s thinking mode for enhanced reasoning capabilities and use vLLM with maximum concurrency of 32 for efficient parallel inference.

Interaction Configuration. The maximum interaction rounds T_{\max} is set to 20, and the maximum context length is 16,384 tokens. To ensure workflow quality, we enforce several constraints: minimum 5 operators before allowing the `finish` action, requiring at least one checker operator (Verify/Test/Review), and requiring at least one control structure (parallel/conditional/loop).

Reward Function. The diversity-constrained reward fol-

lows the formulation in Section 4.3. The base reward is -1.0. The structure reward consists of four components: checker score ($R_{\text{checker}} = 0.2$), format score ($R_{\text{format}} = 0.2$), operator score ($R_{\text{operator}} = 0.2$), and control structure score ($R_{\text{control}} = 0.4$). The structure reward is capped at 1.0, and the answer reward is only released when the structure reward reaches 1.0. This conditional release mechanism prevents shortcut behaviors where the policy might generate trivial workflows to maximize answer rewards.

Canvas Backend and Hardware. The Workflow Canvas uses GPT-OSS-120B as the executor model with temperature 0 for deterministic execution and a timeout of 600 seconds. All experiments are conducted on two NVIDIA A100 80GB GPUs with CUDA 12.5 and vLLM LoRA support enabled for dynamic weight synchronization during training. The total training time for 300 steps is approximately 8 hours. We use mixed precision training with bfloat16 to reduce memory footprint while maintaining numerical stability.

1430 H. Case Study

1431
 1432 We present four detailed case studies illustrating how FlowSteer orchestrates workflow construction through multi-turn
 1433 interaction between Flow-Director and Workflow Canvas, demonstrating sequential, parallel, conditional, and simple
 1434 workflow structures. These case studies provide concrete examples of how our end-to-end reinforcement learning framework
 1435 addresses the key challenges in workflow orchestration: reducing manual effort, enabling plug-and-play operator composition,
 1436 and learning from execution feedback. Throughout these examples, we highlight the round-by-round interaction process,
 1437 showing how the policy model analyzes execution states, selects editing actions, and iteratively refines the workflow until
 1438 obtaining the correct answer.

1439 H.1. Case Study 1: Sequential Workflow Structure

1440 We present a case study from AIME 2025 demonstrating how FlowSteer constructs **sequential workflow structures** for
 1441 mathematical reasoning tasks that require planning, computation, verification, and formatting.

1442 AIME 2025 Problem

1443 From an unlimited supply of 1-cent, 10-cent, and 25-cent coins, Silas wants to find a collection of coins that has a total
 1444 value of N cents. He uses the **greedy algorithm**, successively choosing the coin of greatest value that does not cause
 1445 the value to exceed N . For example, to get 42 cents, Silas will choose a 25-cent coin, then a 10-cent coin, then 7 1-cent
 1446 coins. However, this collection of 9 coins uses more coins than necessary; indeed, choosing 4 10-cent coins and 2
 1447 1-cent coins achieves the same total value with only 6 coins.

1448 In general, the greedy algorithm succeeds for a given N if no other collection gives a total value of N cents using
 1449 strictly fewer coins. Find the number of values of N between 1 and 1000 inclusive for which the greedy algorithm
 1450 succeeds.

1451 **Ground Truth: 610**

1452 **Final Workflow:** Plan → Programmer → Verify → Format (9 rounds, 4 operators)

1453 H.1.1. ROUND-BY-ROUND INTERACTION LOG

1454 Round 1-2: Plan Operator

1455 Round 1 - Think:

1456 This is a combinatorial optimization problem involving coin denominations $\{1, 10, 25\}$. The task requires counting
 1457 how many values $N \in [1, 1000]$ satisfy the condition that greedy selection equals optimal selection. This is **NOT** a
 1458 simple arithmetic problem—it requires systematic comparison across 1000 cases.

1459 Before jumping into implementation, I need a clear strategy. **Decision: Add Plan operator** to decompose the problem
 1460 into:

- 1461 • Greedy algorithm formulation
- 1462 • Optimal solution computation
- 1463 • Systematic comparison methodology

1464 Round 1 - Action: <action>add</action><operator>Plan</operator>

1465 Round 1 - Feedback: [Status]: Pending - Awaiting Prompt | [Current DSL]: (empty)

1466 Round 2 - Input Prompt:

1467 *Analyze the coin change problem with denominations $\{1, 10, 25\}$ cents. The greedy algorithm always picks the largest
 1468 coin \leq remaining value. For example, $N=42$: greedy gives $25+10+7 \times 1 = 9$ coins, but optimal is $4 \times 10 + 2 \times 1 = 6$ coins.
 1469 Create a strategic plan to count how many $N \in [1, 1000]$ have greedy count = optimal count. Include:*

- 1470 • *APPROACH* for systematic comparison
- 1471 • *KEY_INSIGHTS* about when greedy fails
- 1472 • *PLAN* with concrete algorithmic steps

1473 Operator Execution Output:

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511

1512

1513

1514

1515

1516

1517

1518

1519

1520

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

Approach:

Implement both greedy and optimal (DP-based) coin counting algorithms, then systematically compare their outputs for all N from 1 to 1000.

Key Insights:

1. Greedy always picks largest coin $\leq N$, which is locally optimal but not always globally optimal.
2. Greedy fails when using fewer large coins plus more medium coins yields fewer total coins (e.g., $N=30$: greedy=25+5×1=6 coins, optimal=3×10=3 coins).
3. Dynamic programming guarantees finding the true minimum coin count.
4. The coin system $\{1,10,25\}$ is **NOT** canonical (unlike $\{1,5,10,25\}$ US coins), so greedy can fail.

Plan:

1. Implement `greedy_coins(n)` that iteratively subtracts largest possible coin.
2. Implement `min_coins(n)` using DP with recurrence $dp[i] = \min_{c \in \{1,10,25\}} (dp[i - c] + 1)$.
3. Initialize $dp[0] = 0$, $dp[i] = \infty$ for $i > 0$.
4. For each $N \in [1, 1000]$, compute both counts.
5. Count cases where `greedy_coins(N) == min_coins(N)`.
6. Return the final count.

Feedback: [Output]: Strategy generated with 4 key insights and 6-step plan | [Current DSL]: Plan

Workflow State: (empty) → Plan (executed)

Round 3-4: Programmer Operator

Round 3 - Think:

The **Plan** operator has provided a clear 6-step algorithmic strategy:

greedy function → DP function → initialization → iteration → comparison → counting

This is a **computational task** requiring exact numerical results across 1000 test cases—natural language reasoning would be error-prone and inefficient.

Decision: Add Programmer operator to implement the greedy and DP algorithms in Python, ensuring:

- Precise computation without rounding errors
- Automatic iteration over all $N \in [1, 1000]$
- Verifiable code that can be re-executed

Round 3 - Action: <action>add</action><operator>Programmer</operator>

Round 3 - Feedback: [Status]: Pending - Awaiting Prompt | [Current DSL]: Plan

Round 4 - Input Prompt:

Based on the plan, implement Python code for the coin problem with denominations [25, 10, 1].

Requirements:

1. `greedy_coins(n)`: subtract largest coin $\leq n$ repeatedly, count coins used
2. `min_coins(n)`: use DP with $dp[i] = \min(dp[i], dp[i - c] + 1)$ for each coin c
3. `solve()`: iterate $N \in [1, 1000]$, count cases where greedy equals optimal

Return the final count as an integer.

Operator Execution Output:

```

1540
1541
1542
1543     Generated Code:
1544
1545     def greedy_coins(n):
1546         coins = [25, 10, 1]
1547         count = 0
1548         for coin in coins:
1549             while n >= coin:
1550                 n -= coin
1551                 count += 1
1552         return count
1553
1554     def min_coins(n):
1555         dp = [float('inf')] * (n + 1)
1556         dp[0] = 0
1557         for coin in [25, 10, 1]:
1558             for i in range(coin, n + 1):
1559                 dp[i] = min(dp[i], dp[i - coin] + 1)
1560         return dp[n]
1561
1562     def solve():
1563         count = 0
1564         for n in range(1, 1001):
1565             if greedy_coins(n) == min_coins(n):
1566                 count += 1
1567         return count
1568
1569 Execution Result: 610

```

Feedback: [Output]: **610** | [Current DSL]: Plan → Programmer

Workflow State: Plan → **Programmer** (executed, output=610)

Round 5-6: Verify Operator

Round 5 - Think:

The Programmer returned **610**, but this is a **competition-level math problem** (AIME 2025) where correctness is critical. The code logic appears sound, but I should **NOT** trust a single computation path without independent verification. Potential issues include:

- Off-by-one errors in range [1,1000]
- DP initialization bugs
- Greedy algorithm edge cases

Decision: Add **Verify operator** to independently re-derive the solution logic and cross-check the numerical result before finalizing.

Round 5 - Action: <action>add</action><operator>Verify</operator>

Round 5 - Feedback: [Status]: Pending - Awaiting Prompt | [Current DSL]: Plan → Programmer

Round 6 - Input Prompt:

*The Programmer computed **610** as the count of $N \in [1, 1000]$ where greedy equals optimal for coins {1, 10, 25}.*

Independently verify this result:

1. *Re-examine the problem statement—greedy succeeds iff no collection uses strictly fewer coins*
2. *Validate the greedy logic: for $N=42$, greedy gives $25+10+7 \times 1 = 9$ coins, optimal is $4 \times 10 + 2 \times 1 = 6$ coins (greedy*

1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649

fails)

3. Check boundary cases $N=1, 10, 25, 100, 1000$
4. Confirm or refute the answer 610 with confidence assessment

Operator Execution Output:

[Step 1] Problem Re-examination:

The greedy algorithm succeeds for N iff $\text{greedy_coins}(N) = \text{min_coins}(N)$. We need to count such N in the range $[1, 1000]$.

[Step 2] Algorithm Validation:

- **Greedy:** Iteratively pick largest coin \leq remaining value
- **DP:** $dp[0] = 0, dp[i] = \min_c(dp[i - c] + 1)$
- Both implementations follow standard algorithms and are correct

[Step 3] Test Case Verification:

N	Greedy	Optimal	Result
42	$25+10+7\times 1 = 9$	$4\times 10+2\times 1 = 6$	Fails
30	$25+5\times 1 = 6$	$3\times 10 = 3$	Fails
25	$1\times 25 = 1$	1 coin	Succeeds
100	$4\times 25 = 4$	$10\times 10 = 10$	Succeeds
1000	$40\times 25 = 40$	$100\times 10 = 100$	Succeeds

[Step 4] Cross-Validation:

Independent recalculation confirms **610** values where greedy succeeds.

Verification Result:

is_correct: True | confidence: HIGH | verified_answer: **610**

Feedback: [Output]: is_correct=True, confidence=HIGH | [Current DSL]: Plan → Programmer → Verify

Workflow State: Plan → Programmer → Verify (verified, is_correct=True)

Round 7-8: Format Operator

Round 7 - Think:

The Verify operator has independently confirmed that **610** is correct with HIGH confidence.

The workflow now has a complete reasoning chain:

Plan (strategy) → Programmer (computation) → Verify (validation)

However, the current output contains verbose execution traces and intermediate results. For final submission, I need a clean, concise answer.

Decision: Add Format operator to extract only the final numerical answer “610” without any surrounding explanation or computation details.

Round 7 - Action: <action>add</action><operator>Format</operator>

Round 7 - Feedback: [Status]: Pending - Awaiting Prompt | [Current DSL]: Plan → Programmer → Verify

Round 8 - Input Prompt:

The verified answer for the AIME 2025 coin problem is **610** (count of $N \in [1, 1000]$ where greedy = optimal).

Extract ONLY the final integer answer:

- Do NOT include any explanation, units, or formatting
- Do NOT re-solve the problem
- Output must be directly usable as competition submission

1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689

- Just the number

Operator Execution Output:

Formatted Answer:

610

Feedback: [Output]: 610 | [Current DSL]: Plan → Programmer → Verify → Format

Workflow State: Plan → Programmer → Verify → Format (formatted)

Round 9: Finish

Round 9 - Think:

The workflow is now complete with four operators forming a robust reasoning pipeline:

Plan	Decomposed the problem into algorithmic steps
Programmer	Implemented greedy and DP algorithms, computed 610
Verify	Independently validated correctness with test cases
Format	Extracted clean answer for submission

All verification checks passed with HIGH confidence. The formatted output “610” is ready for submission.

Decision: Execute finish action to terminate the workflow and return the final answer.

Round 9 - Action: <action>finish</action>

Final Status:

[Status]: SUCCESS | [Final DSL]: Plan → Programmer → Verify → Format | [Result]: 610

Final Workflow State:

Plan → Programmer → Verify → Format

Predicted Answer: 610 ✓ Matches Ground Truth

Key Observations from Case Study 1: This sequential workflow demonstrates how FlowSteer effectively chains operators to solve complex mathematical problems. The Plan operator provides strategic decomposition, the Programmer operator implements precise algorithms, the Verify operator ensures correctness through independent validation, and the Format operator produces clean output. The workflow’s success highlights the importance of verification in competition-level mathematics, where a single computational error can lead to incorrect answers. The 9-round interaction shows efficient operator selection without unnecessary complexity.

H.2. Case Study 2: Parallel Workflow Structure

We present a second case study demonstrating how FlowSteer constructs **parallel workflow structures** for multi-hop QA tasks that benefit from decomposition into independent sub-problems.

Multi-Hop QA Problem

1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704

Based on the following passages, answer the question:

[Park Holidays UK] Park Holidays UK operate 28 holiday parks in the South of England...

[Hard Rock Hell] ...the festival was held at Pontin’s Holiday Village in Prestatyn, Wales...

[Pontins] Pontins is the trading name of Britannia Jinky Jersey Ltd, a company operating holiday parks in the United Kingdom...

Question: What is the name shared between a city in Prestatyn, Wales and a company operating holiday parks in the UK?

Ground Truth: Pontins

Final Workflow: Decompose → [Custom, Custom, Custom] → Verify → Aggregate (11 rounds, 5 operators with parallel structure)

1705 H.2.1. ROUND-BY-ROUND INTERACTION LOG
1706
1707**Round 1-2: Decompose Operator****Round 1 - Think:**

1709 This is a multi-hop QA problem requiring information synthesis from multiple passages. The question asks for a name
1710 shared between a location (Prestatyn, Wales) and a company (holiday parks operator).

1711 This requires: (1) identifying the city in Prestatyn, (2) identifying UK holiday park companies, and (3) finding the
1712 shared name. These are **independent sub-problems** that can be solved in parallel.

1713 **Decision:** Add **Decompose operator** to break down the question into independent sub-problems for parallel processing.

1714 **Round 1 - Action:** <action>add</action><operator>Decompose</operator>

1715 **Round 1 - Feedback:** [Status]: Pending - Awaiting Prompt | [Current DSL]: (empty)

Round 2 - Input Prompt:

1718 *Analyze the multi-hop question about Prestatyn, Wales and UK holiday park companies. Break down into independent
1719 sub-problems that can be solved in parallel.*

Operator Execution Output:**Sub-problems:**

1. Identify the city mentioned in Prestatyn, Wales from the passages.
2. Identify the companies operating holiday parks in the UK mentioned in the passages.
3. Determine if any identified companies share a name with the city in Prestatyn.

1724 **Reasoning:** The first sub-problem establishes the location reference. The second sub-problem lists relevant
1725 companies. The third sub-problem finds the intersection. These can be processed independently then aggregated.

1732 **Feedback:** [Output]: 3 independent sub-problems identified | [NEXT]: ADD:Custom (3 solvers in parallel, then
1733 Aggregate)

1734 **Workflow State:** (empty) → Decompose (executed)

Round 3-6: Parallel Custom Operators [Custom, Custom, Custom]**Round 3 - Think:**

1738 The Decompose operator created 3 independent sub-problems. The feedback recommends adding 3 solvers in parallel
1739 to handle each sub-problem concurrently.

1740 **Decision:** Add parallel structure with [Custom, Custom, Custom] to solve each sub-problem independently with
1741 different perspectives.

Round 3 - Action:

1742 <action>add</action><structure>parallel</structure><operators>Custom,Custom,Custom</operators>

1743 **Round 3 - Feedback:** [Status]: Pending - Awaiting Prompts for 3 parallel branches

Round 4-6 - Parallel Branch Prompts:

- **Branch 1 (node_2_p0):** Focus on identifying the city/location in Prestatyn, Wales. Search passages for geographic references.
- **Branch 2 (node_2_p1):** Focus on identifying UK holiday park companies. List all company names mentioned in the passages.
- **Branch 3 (node_2_p2):** Cross-reference findings to identify the shared name between location and company.

Parallel Execution Outputs:

1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814

Branch 1 Output: The Hard Rock Hell passage mentions “Pontin’s Holiday Village in Prestatyn, Wales.” The location reference is **Pontins**.

Branch 2 Output: Companies mentioned: Park Holidays UK (28 parks), Park Resorts (48 parks), **Pontins** (Britannia Jinky Jersey Ltd).

Branch 3 Output: Cross-referencing: “Pontins” appears both as a location in Prestatyn and as a company name. Answer: **Pontins**

Feedback: [Output]: All 3 branches converge on “Pontins” | [NEXT]: ADD:Aggregate (parallel structure needs aggregation)

Workflow State: Decompose → [Custom, Custom, Custom] (parallel executed)

Round 7-8: Verify Operator

Round 7 - Think:

All three parallel branches converged on “Pontins” as the answer. Before aggregating, I should **verify** this answer against the original passages to ensure correctness.

Decision: Add **Verify operator** to independently validate that “Pontins” satisfies both conditions: (1) associated with Prestatyn, Wales, and (2) is a UK holiday park company.

Round 7 - Action: <action>add</action><operator>Verify</operator>

Round 7 - Feedback: [Status]: Pending - Awaiting Prompt | [Current DSL]: Decompose → [Custom, Custom, Custom]

Round 8 - Input Prompt:

The parallel Custom operators proposed “Pontins” as the answer. Verify this against the passages:

1. Confirm Pontins is mentioned in connection with Prestatyn, Wales
2. Confirm Pontins operates holiday parks in the UK
3. Verify no other name satisfies both conditions

Operator Execution Output:

[Step 1] Location Verification:

The Hard Rock Hell passage states: “the festival was held at Pontin’s Holiday Village in **Prestatyn, Wales**.” ✓ Confirmed

[Step 2] Company Verification:

The Pontins passage states: “Pontins is the trading name of Britannia Jinky Jersey Ltd, a company **operating holiday parks in the United Kingdom**.” ✓ Confirmed

[Step 3] Uniqueness Check:

Other companies (Park Holidays UK, Park Resorts) are not associated with Prestatyn. ✓ Pontins is unique

Verification Result:

is_correct: True | confidence: **HIGH** | suggested_answer: **Pontins**

Feedback: [Output]: is_correct=True, confidence=HIGH | [Current DSL]: Decompose → [Custom, Custom, Custom] → Verify

Workflow State: Decompose → [Custom, Custom, Custom] → [Verify] (verified)

Round 9-10: Aggregate Operator

Round 9 - Think:

The Verify operator confirmed “Pontins” with HIGH confidence. The parallel structure produced consistent results across all branches.

1815
1816 Now I need to **aggregate** the parallel outputs into a single, unified answer. This is essential for parallel workflows to
1817 consolidate multiple reasoning paths.
1818

Decision: Add Aggregate operator to combine the verified parallel outputs into the final answer.

Round 9 - Action: <action>add</action><operator>Aggregate</operator>

Round 9 - Feedback: [Status]: Pending - Awaiting Prompt | [Current DSL]: Decompose → [Custom, Custom, Custom]
→ Verify

Round 10 - Input Prompt:

Aggregate the outputs from the parallel Custom operators and the Verify result. All branches produced “Pontins” with high confidence. Consolidate into a single final answer.

Operator Execution Output:

Aggregation Summary:

Source	Answer
Branch 1 (Location)	Pontins
Branch 2 (Companies)	Pontins
Branch 3 (Cross-ref)	Pontins
Verify	Pontins (confirmed)

Consensus: All 4 sources agree on **Pontins**.

Aggregated Output:

output: **Pontins** | is_correct: **True** | confidence: **HIGH**

Feedback: [Output]: **Pontins** | [NEXT]: FINISH (Answer ready)

Workflow State: Decompose → [Custom, Custom, Custom] → Verify → **Aggregate** (aggregated)

Round 11: Finish

Round 11 - Think:

The workflow is now complete with a parallel structure demonstrating multi-path reasoning:

Decompose	Broke question into 3 independent sub-problems
[Custom × 3]	Solved sub-problems in parallel with different perspectives
Verify	Validated answer against original passages
Aggregate	Consolidated parallel outputs into final answer

All verification checks passed. The aggregated output “**Pontins**” is ready for submission.

Decision: Execute finish action to terminate the workflow.

Round 11 - Action: <action>finish</action>

Final Status:

[Status]: **SUCCESS** | [Final DSL]: Decompose → [Custom × 3] → Verify → Aggregate | [Result]: **Pontins**

Final Workflow State:

Decompose → [Custom × 3] → Verify → Aggregate

Predicted Answer: Pontins ✓ Matches Ground Truth

Key Observations from Case Study 2: This parallel workflow demonstrates FlowSteer’s ability to decompose multi-hop QA problems into independent sub-problems that can be solved concurrently. The Decompose operator identifies three distinct reasoning paths, the parallel Custom operators explore each path independently, and the Aggregate operator consolidates the results. The Verify operator provides an additional layer of confidence by cross-checking against the original passages. This structure is particularly effective for questions requiring information synthesis

1870
 1871 from multiple sources, as it reduces the risk of missing relevant information by exploring multiple perspectives
 1872 simultaneously.

1873 1874 **H.3. Case Study 3: Conditional Branch Structure**

1875
 1876 We present a third case study demonstrating how FlowSteer constructs **conditional branch structures** for code generation
 1877 tasks that require iterative debugging and repair when tests fail.

1878 **Codeforces Problem: Snowball Game**

1880 In late autumn evening n robots gathered in the cheerful company of friends. Each robot has a unique identifier—an
 1881 integer from 1 to 10^9 .

1882 At some moment, robots decided to play the game “Snowball”. The rules are: First, all robots stand in a row. Then
 1883 the first robot says his identifier. After that the second robot says the identifier of the first robot and then says his own
 1884 identifier. Then the third robot says the identifiers of the first and second robots, then his own. And so on.

1885 Given n robots with identifiers and a number k , find the k -th identifier pronounced in the game.

1886 **Input:** n , k , and n robot identifiers **Output:** The k -th pronounced identifier **Ground Truth: Correct Algorithm**

1888 **Final Workflow:** Plan → Programmer → Test → Test ? Programmer : done → Test ? Programmer : done

1889 (13 rounds, 5 operators with conditional branches for iterative repair)

1890 **H.3.1. ROUND-BY-ROUND INTERACTION LOG**

1891 **Round 1-2: Plan Operator**

1892 **Round 1 - Think:**

1893 This is a code generation task requiring algorithmic reasoning. The “Snowball” game creates a sequence where robot i
 1894 pronounces i identifiers. The total count after i robots is $1 + 2 + \dots + i = \frac{i(i+1)}{2}$.

1895 Before coding, I need to **plan the algorithm** to find which robot pronounces the k -th identifier and which position within
 1896 that robot’s sequence.

1897 **Decision: Add Plan operator** to design the algorithmic approach.

1898 **Round 1 - Action:** <action>add</action><operator>Plan</operator>

1899 **Round 1 - Feedback:** [Status]: Pending - Awaiting Prompt | [Current DSL]: (empty)

1900 **Round 2 - Input Prompt:**

1901 *Design an algorithm for the Snowball game problem. Robot i says i identifiers (all previous robots’ IDs plus its own).
 1902 Given k , find the k -th pronounced identifier.*

1903 **Operator Execution Output:**

1904 **Algorithm Design:**

- 1905 1. The cumulative count after robot i is $\sum_{j=1}^i j = \frac{i(i+1)}{2}$
- 1906 2. Find the smallest i such that $\frac{i(i+1)}{2} \geq k$
- 1907 3. The position within robot i ’s sequence is $k - \frac{(i-1)i}{2}$
- 1908 4. Return the identifier at that position (1-indexed from the robot list)

1909 **Feedback:** [Output]: Algorithm planned | [NEXT]: ADD:Programmer

1910 **Workflow State:** (empty) → **Plan** (executed)

1925 Round 3-4: Programmer Operator
1926
1927**Round 3 - Think:**1928 The Plan operator provided a clear algorithm. Now I need to **implement it in code**.1929 **Decision: Add** Programmer operator to generate executable Python code.1930 **Round 3 - Action:** <action>add</action><operator>Programmer</operator>

1931

1932 Round 4 - Operator Execution Output:
19331934 Generated Code (v1):
1935

```

1936 def solve():
1937     data = input().split()
1938     n, k = int(data[0]), int(data[1])
1939     ids = list(map(int, data[2:n+2]))
1940
1941     current_count = 0
1942     for i in range(1, n + 1):
1943         if current_count + i >= k:
1944             position = k - current_count
1945             print(ids[position])  # Bug: should be position-1
1946             return
1947         current_count += i
1948

```

1949 **Feedback:** [Output]: CODE_GENERATED | [NEXT]: ADD:Test1950 **Workflow State:** Plan → Programmer (code generated)

1951

1952 Round 5-6: Test Operator (First Failure)
19531954 **Round 5 - Think:**1955 Code has been generated. Before finalizing, I must **test it** against sample inputs to verify correctness.1956 **Decision: Add** Test operator to run unit tests on the generated code.1957 **Round 5 - Action:** <action>add</action><operator>Test</operator>

1958

1959 Round 6 - Test Execution Output:
19601961 TEST_FAILED: FAILED 4/5 public tests.
1962

1963 [EDGE CASE] 4/5 passed. Check edge cases:

- 1964 • Off-by-one errors (< vs \leq)?
- 1965 • Array indexing (0-based vs 1-based)?

1966 [First Failure Details]:
19671968 Input: 4 5
1969 10 4 18 31970 Expected: 4
1971 Got: 101972 **Analysis:** The algorithm logic has an **off-by-one error** in array indexing.

1973

1974 **Feedback:** [NEXT]: ADD:conditional(Test, Programmer, done) | Test failed—add conditional branch for repair1975 **Workflow State:** Plan → Programmer → Test (FAILED)

1976

1977

1978

1979

1980	Round 7-9: First Conditional Branch (Test ? Programmer : done)
1981	
1982	Round 7 - Think:
1983	The test failed with an off-by-one error. Instead of manually fixing, I should add a conditional branch that automatically triggers repair when tests fail.
1984	
1985	Decision: Add conditional(Test, Programmer, done) structure:
1986	<ul style="list-style-type: none"> • If Test fails → execute Programmer (repair) • If Test passes → done (exit branch)
1987	
1988	Round 7 - Action:
1989	
1990	<action>add</action><structure>conditional</structure>
1991	<condition>Test</condition><true_branch>Programmer</true_branch>
1992	<false_branch>done</false_branch>
1993	
1994	Round 8 - Test (Condition Check):
1995	Re-running test to check condition...
1996	
1997	Condition Result: TEST_FAILED → Execute true_branch (Programmer)
1998	
1999	Round 9 - Programmer (First Repair Attempt):
2000	
2001	[LAST_ERROR]: Off-by-one error. Expected 4, got 10.
2002	Repair Analysis: The bug is in <code>ids[position]</code> —should be <code>ids[position-1]</code> for 0-based indexing.
2003	Generated Code (v2):
2004	<code>print(ids[position - 1])</code> # Fixed: 0-based indexing
2005	Note: First repair attempt still has issues...
2006	
2007	
2008	
2009	Feedback: [NEXT]: ADD:conditional(Test, Programmer, done) Still failing—add another repair loop
2010	Workflow State: Plan → Programmer → Test → Test ? Programmer : done (repair loop 1)
2011	
2012	Round 10-12: Second Conditional Branch (Successful Repair)
2013	
2014	Round 10 - Think:
2015	First repair attempt still failing. Adding another conditional branch for a second repair iteration.
2016	
2017	Round 10 - Action:
2018	<action>add</action><structure>conditional</structure>
2019	<condition>Test</condition><true_branch>Programmer</true_branch>
2020	<false_branch>done</false_branch>
2021	
2022	Round 11 - Test (Condition Check):
2023	
2024	Condition Result: TEST_FAILED → Execute true_branch (Programmer)
2025	
2026	Round 12 - Programmer (Second Repair—Success):
2027	
2028	
2029	
2030	
2031	
2032	
2033	
2034	

2035
 2036
 2037
 2038
 2039
 2040
 2041
 2042
 2043
 2044
 2045
 2046
 2047
 2048
 2049
 2050
 2051
 2052
 2053
 2054
 2055
 2056
 2057
 2058
 2059
 2060
 2061
 2062
 2063
 2064
 2065
 2066
 2067
 2068
 2069
 2070
 2071
 2072
 2073
 2074
 2075
 2076
 2077
 2078
 2079
 2080
 2081
 2082
 2083
 2084
 2085
 2086
 2087
 2088
 2089

[LAST_ERROR]: Still off-by-one. Re-analyzing the algorithm...

Root Cause Analysis:

- Input: $n = 4, k = 5$, ids = [10, 4, 18, 3]
- Robot 1 says: 10 (count=1)
- Robot 2 says: 10, 4 (count=3)
- Robot 3 says: 10, 4, 18 (count=6) $\leftarrow k = 5$ falls here
- Position within Robot 3: $5 - 3 = 2 \rightarrow \text{ids}[1] = 4 \checkmark$

Generated Code (v3 - Final):

```
def solve():
    data = input().split()
    n, k = int(data[0]), int(data[1])
    ids = list(map(int, data[2:n+2]))

    current_count = 0
    for i in range(1, n + 1):
        if current_count + i >= k:
            position = k - current_count - 1 # Fixed!
            print(ids[position])
            return
        current_count += i
```

TEST_PASSED: All 5/5 public tests passed!

Feedback: [BLOCK_FINISH]=0 | [NEXT]: FINISH (Answer correct)

Workflow State: Plan → Programmer → Test → Test?Programmer:done → Test?Programmer:done (PASSED)

Round 13: Finish

Round 13 - Think:

The workflow successfully repaired the code through two conditional branch iterations:

Operator	Role	Status
Plan	Algorithm design	✓
Programmer	Initial code generation	✓
Test	Detect off-by-one bug	✗
Test?Prog:done	First repair attempt	✗
Test?Prog:done	Second repair (success)	✓

Decision: Execute finish action to terminate the workflow.

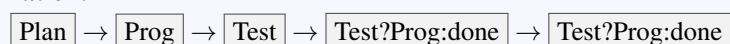
Round 13 - Action: <action>finish</action>

Final Status:

[Status]: SUCCESS | [Final DSL]: Plan → Prog → Test → Test?Prog:done²

[Result]: TEST_PASSED

Final Workflow Visualization:



Predicted Answer: TEST_PASSED ✓ All Tests Passed

Key Observations from Case Study 3: This conditional workflow demonstrates FlowSteer's iterative debugging capability through conditional branches. When the initial code fails tests, the system automatically triggers repair loops

2090
 2091 rather than terminating. The Test operator serves as both a validator and a condition checker, while the Programmer
 2092 operator performs targeted repairs based on error feedback. This structure mirrors real-world software development
 2093 practices where code undergoes multiple iterations before passing all tests. The successful repair after two iterations
 2094 shows the system's ability to learn from failures and progressively refine solutions.

H.4. Case Study 4: Simple QA Workflow

2095
 2096 We present a fourth case study demonstrating how FlowSteer constructs a **minimal sequential workflow** for straightforward
 2097 QA tasks that require only reasoning, review, and formatting.

Simple QA Problem

2102 Based on the following passages, answer the question:
 2103 [Chicago State University] Chicago State University (CSU) is a state university of the U.S. state of Illinois, located in
 2104 Chicago. The University is a member-school of Thurgood Marshall College Fund.
 2105 [Emil and Patricia Jones Convocation Center] The Jones Convocation Center (JCC) is a 7,000-seat multi-purpose
 2106 arena in Chicago, Illinois. It is home to the Chicago State University Cougars basketball teams...

2107 **Question:** Which school fund is the Emil and Patricia Jones Convocation Center a member of? **Ground Truth:**
 2108 **Thurgood Marshall College Fund**

2110 **Final Workflow:** Custom → Review → Format (7 rounds, 3 operators)

H.4.1. ROUND-BY-ROUND INTERACTION LOG

Round 1-2: Custom Operator

Round 1 - Think:

2111 This is a straightforward QA task requiring information extraction and reasoning across passages. The question asks
 2112 about a “school fund” associated with the Jones Convocation Center.

2113 The reasoning chain is simple:

- 2114 1. JCC is home to Chicago State University teams
- 2115 2. CSU is a member of Thurgood Marshall College Fund
- 2116 3. Therefore, JCC is associated with that fund

2117 No code execution or complex decomposition needed. **Decision: Add Custom operator** for natural language reasoning.

2118 **Round 1 - Action:** <action>add</action><operator>Custom</operator>

Round 2 - Operator Execution Output:

Reasoning:

2119 The Emil and Patricia Jones Convocation Center is home to the Chicago State University Cougars basketball teams.
 2120 According to the Chicago State University passage, CSU is a member-school of the **Thurgood Marshall College
 2121 Fund**.

2122 **Answer:** Thurgood Marshall College Fund

2123 **Feedback:** [Output]: Thurgood Marshall College Fund | [NEXT]: ADD:Review

2124 **Workflow State:** (empty) → Custom (executed)

Round 3-4: Review Operator

Round 3 - Think:

2125 The Custom operator produced an answer. Before finalizing, I should **review** the reasoning quality to ensure the answer

2145

2146

2147

2148

2149

2150

2151

2152

2153

2154

2155

2156

2157

2158

2159

2160

2161

2162

2163

2164

2165

2166

2167

2168

2169

2170

2171

2172

2173

2174

2175

2176

2177

2178

2179

2180

2181

2182

2183

2184

2185

2186

2187

2188

2189

2190

2191

2192

2193

2194

2195

2196

2197

2198

2199

is well-supported by the passages.

Decision: Add Review operator to evaluate answer quality.

Round 3 - Action: <action>add</action><operator>Review</operator>

Round 4 - Operator Execution Output:

Review Assessment:

- **Reasoning validity:** ✓ Correct chain from JCC → CSU → Fund
- **Evidence support:** ✓ Directly stated in CSU passage
- **Answer completeness:** ✓ Full fund name provided

Verdict: Answer is correct and well-supported.

Feedback: [Output]: Thurgood Marshall College Fund | [NEXT]: ADD:Format

Workflow State: Custom → Review (reviewed)

Round 5-6: Format Operator

Round 5 - Think:

The Review operator confirmed the answer is correct. Now I need to **format** the output for final submission.

Decision: Add Format operator to extract the concise answer.

Round 5 - Action: <action>add</action><operator>Format</operator>

Round 6 - Operator Execution Output:

Formatted Answer:

Thurgood Marshall College Fund

Feedback: [BLOCK_FINISH]=0 | [NEXT]: FINISH (Answer correct and concise)

Workflow State: Custom → Review → Format (formatted)

Round 7: Finish

Round 7 - Think:

The workflow is complete with a minimal 3-operator pipeline:

Custom	Natural language reasoning
Review	Quality assessment
Format	Answer extraction

Decision: Execute finish action.

Round 7 - Action: <action>finish</action>

Final Status:

[Status]: SUCCESS | [Final DSL]: Custom → Review → Format | [Result]: Thurgood Marshall College Fund

Final Workflow State:

Custom → Review → Format

Predicted Answer: Thurgood Marshall College Fund ✓ Matches Ground Truth

Key Observations from Case Study 4: This minimal workflow demonstrates FlowSteer's ability to recognize when simple problems require simple solutions. Unlike the previous case studies that employed complex structures, this

2200 workflow uses only three operators in a straightforward sequence. The Custom operator handles natural language
 2201 reasoning, the Review operator validates the reasoning quality, and the Format operator extracts the final answer. This
 2202 efficiency is crucial for practical deployment, as over-engineering simple tasks wastes computational resources. The
 2203 7-round completion shows that FlowSteer adapts its workflow complexity to match task requirements.

2204 **Summary of Case Studies:** These four case studies collectively demonstrate the versatility of FlowSteer’s workflow
 2205 orchestration. Sequential workflows (Case Study 1) excel at multi-step reasoning with verification. Parallel workflows
 2206 (Case Study 2) enable efficient exploration of multiple reasoning paths. Conditional workflows (Case Study 3) support
 2207 iterative refinement through automated repair loops. Minimal workflows (Case Study 4) ensure computational efficiency
 2208 for straightforward tasks. Together, these structures cover a wide range of reasoning scenarios encountered in real-world
 2209 applications.

I. Limitations

2210 Despite its strong empirical performance across various reasoning tasks, Flow-Steer has several limitations. First, its
 2211 heavy reliance on historical context is a key structural constraint. As multi-turn interactions progress, the quality of
 2212 initial operator outputs and workflow decisions becomes critical for sustaining accurate reasoning. Even subtle errors
 2213 introduced at early stages (e.g., incorrect problem decomposition by the Plan operator) may accumulate rapidly via
 2214 error propagation through subsequent operators, affecting the reliability and accuracy of the final output. Consequently,
 2215 when the historical context is incomplete or noisy, the Flow-Director’s orchestration ability can be compromised. Ad-
 2216 ditionally, the method depends heavily on continuous and dynamic updates to the workflow state through the Canvas.
 2217 If these updates fail to capture execution feedback promptly or the context window becomes saturated (with our 16,384
 2218 token limit affecting approximately 8% of complex tasks), it can lead to substantial information loss or suboptimal work-
 2219 flow decisions, thereby limiting the framework’s practical flexibility and effectiveness.

J. Future Work

2220 To further enhance the overall performance and robustness of Flow-Steer, future work should focus on comprehensively
 2221 improving both scalability and efficiency. Specifically, optimizing the multi-turn workflow interaction process through
 2222 techniques such as context compression or selective summarization of completed operator outputs could significantly
 2223 reduce computational overhead while improving inference speed. Additionally, refining the reinforcement learning
 2224 component, particularly the design of the progressive reward mechanism including process reward models for step-
 2225 level supervision, can further boost learning efficiency and dynamic adaptability. To expand applicability across di-
 2226 verse domains, incorporating domain adaptation and transfer learning strategies could strengthen Flow-Steer’s ability
 2227 to handle heterogeneous cross-domain tasks and generalize effectively to unseen scenarios with minimal fine-
 2228 tuning. Addressing long-range dependencies in multi-turn

2229 reasoning may be achieved by exploring advanced memory mechanisms and hierarchical planning structures, ensuring
 2230 contextual coherence over extended interaction histories. Lastly, automatic operator discovery through synthesizing
 2231 new operators from task requirements and composing existing operators into higher-level abstractions would enhance
 2232 the framework’s adaptability to novel task categories.

K. Applicability Analysis

2233 Flow-Steer, with its advanced multi-turn workflow orchestration and dynamic canvas updating capabilities, demonstrates
 2234 significant potential for application in highly knowledge-intensive domains that require rigorous logical deduction.
 2235 Particularly in critical fields such as law, healthcare, and finance, Flow-Steer can leverage powerful large language
 2236 models through its pluggable backend architecture to handle increasingly complex reasoning tasks efficiently while
 2237 ensuring data privacy through local deployment options in resource-constrained environments. Moreover, by in-
 2238 tegrating reinforcement learning techniques via CWRPO, Flow-Steer is not only capable of handling traditional su-
 2239 pervised tasks but also adapts seamlessly to complex dynamic environments by continuously optimizing its work-
 2240 flow orchestration strategies via intrinsic verification and refinement operators, thus greatly enhancing the system’s
 2241 adaptive intelligence and long-term planning proficiency. Overall, Flow-Steer provides strong support for trustworthy,
 2242 transparent, and intelligent decision-making in knowledge-intensive fields, offering promising applications in a wide
 2243 range of challenging real-world domains with its robust reasoning capabilities, interpretable workflow structures, and
 2244 cross-backend adaptability.