

REINFORCING QUERY-LEVEL META-AGENTS

Anonymous authors

Paper under double-blind review

ABSTRACT

This paper proposes a query-level meta-agent named FLOWREASONER to automate the design of query-level multi-agent systems, i.e., *one system per user query*. Our core idea is to incentivize a reasoning-based meta-agent via external execution feedback. Concretely, by distilling DeepSeek R1, we first endow the basic reasoning ability regarding the generation of multi-agent systems to FLOWREASONER. Then, we further enhance it via reinforcement learning (RL) with external execution feedback. A multi-purpose reward is designed to guide the RL training from aspects of performance, complexity, and efficiency. In this manner, FLOWREASONER is enabled to generate a personalized multi-agent system for each user query via deliberative reasoning. Experiments on both engineering and competition code benchmarks demonstrate the superiority of FLOWREASONER. Remarkably, it surpasses o1-mini by **10.52%** accuracy across three benchmarks. All the code is included in the Supplemental Material.

1 INTRODUCTION

Large language models (LLMs) (Achiam et al., 2023; Reid et al., 2024; Team, 2024a; Yang et al., 2024; Liu et al., 2024) have exhibited remarkable power in various meaningful yet challenging domains, like chatbots (OpenAI, 2022), code (CognitionAI, 2024), math (OpenAI, 2024), robotics (Kim et al., 2024), etc. LLM-based multi-agent systems (Hong et al., 2023; Wu et al., 2023; Li et al., 2023), which are characterized by planning, reasoning, tool use, and memory, become the foundation of these LLM-driven applications.¹ While effective, most of them are manually designed, increasing human resource costs and limiting scalability.

To mitigate this challenge, early automatic methods are proposed to optimize the prompts (Yuksekgonul et al., 2024; Khattab et al., 2024; Zhou et al., 2024a; Yang et al., 2023) or hyperparameters (Saad-Falcon et al., 2024). But they still rely on the fixed workflow of the multi-agent system, which requires human effort to manually design workflows for each new scenario. From this motivation, various graph-based methods (Zhuge et al., 2024; Liu et al., 2023; Zhang et al., 2024a; Feng et al., 2025) formulate the workflows as graphs or networks and automate the workflow designs. However, the structural complexity of graphs limits their scalability (Hu et al., 2024). To overcome this limitation, state-of-the-art methods represent the multi-agent systems as programming codes (Hu et al., 2024) and prompt a performant LLM, e.g., GPT-4o, as a *meta-agent* to optimize workflows via complex search algorithms on carefully designed search sets (Zhang et al., 2024b; Shang et al., 2024; Zhang et al., 2025).

These previous methods focus on **task-level meta-agents**, generating merely a single **task-specific multi-agent system** that applies to one kind of task, e.g., *code generation task*, as in Figure 1 (a). However, for individual user queries, these one-size-fits-all systems lack the capability for automatic adaptation. To enhance the adaptability of multi-agent systems for individual user queries, this paper aims to design a **query-level meta-agent** to generate a **query-specific multi-agent system** for each user query, e.g., *build a 2048 game*, as shown in Figure 1 (b).

We first identify that the success of task-level meta-agents largely depends on carefully designed search sets, as they rely on complex search algorithms. However, such search sets are unavailable in the setting of query-specific multi-agent systems. To address this issue, instead of relying on search algorithms, we propose to integrate **external execution feedback** of the generated multi-agent system, based on which a **reasoning-driven meta-agent** is leveraged to polish the system.

¹This paper defines a *multi-agent system* as a system consisting of multiple agents operating under a *workflow*.

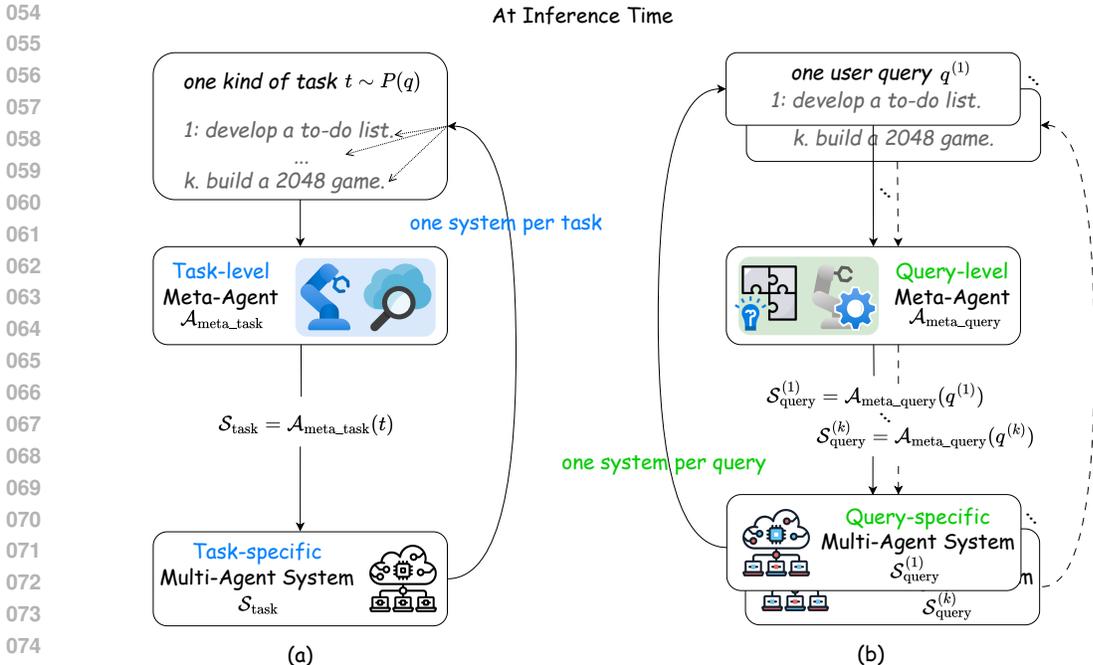


Figure 1: **Task-Level Meta-Agents vs. Query-Level Meta-Agents at Inference Time.** q denotes a user query, e.g., *build a 2048 game*. $t \sim P(q)$ denotes one kind of task, e.g., *code generation task*, which is a distribution of user queries. Given t , previous task-level meta-agent $\mathcal{A}_{\text{meta_task}}$ aims to search a task-specific multi-agent system $\mathcal{S}_{\text{task}}$ to solve all queries sampled from t , i.e., **one system per task**. Differently, given one user query $q^{(i)}$, our query-level meta-agent $\mathcal{A}_{\text{meta_query}}$ conducts reasoning and output a query-specific multi-agent system $\mathcal{S}_{\text{query}}^{(i)}$ for $q^{(i)}$, i.e., **one system per query**.

We dub such a meta-agent as FLOWREASONER. We first synthesize thousands of warm-up SFT data using DeepSeek R1-671B (Team, 2025a) as the meta-agent to generate multi-agent systems and process user queries individually. These synthetic data are then used to finetune DeepSeek-R1-Distill-Qwen-7B, enabling basic reasoning for multi-agent system generation. Furthermore, we enhance its reasoning capabilities for generating novel query-level multi-agent systems through reinforcement learning (RL), incorporating external execution feedback. A **multi-purpose reward** is designed to guide RL training, focusing on *performance*, *complexity*, and *efficiency*. During inference, FLOWREASONER leverages deliberative reasoning to generate a novel query-level multi-agent system for each user query, achieving *one system per user query*. The main contributions are summarized as follows:

- We propose a query-level meta-agent termed FLOWREASONER to automate the designs of query-level multi-agent systems, improving their adaptability in real-world scenarios.
- We train FLOWREASONER to reason from external execution feedback via RL, guided by a multi-purpose reward considering performance, complexity, and efficiency.
- We demonstrate the superiority of FLOWREASONER via extensive experiments and open-source it.

2 RELATED WORK

LLM-Based Multi-Agent Systems. LLM-based multi-agent systems (Guo et al., 2024; Du et al., 2023; Chen et al., 2023b; Park et al., 2023; Huang et al., 2024; Jin et al., 2025) serve as the foundation of various LLM-powered real-world applications, e.g., code intelligence (CognitionAI, 2024), computer use (Anthropic, 2024), and deep research (OpenAI, 2025). LLM-based agents are equipped with planning capabilities, database access, and tool function invocation. These agents collaborate within a multi-agent system, leading to promising performance. Most multi-agent systems are manually designed, increasing the costs of human resources and limiting the scalability. To address this issue, researchers propose automation methods to automate the design of multi-agent

108 systems. Early methods (Yang et al., 2023; Khattab et al., 2024; Zhou et al., 2024a) are proposed to
 109 automate the agentic designs via optimizing prompts (Yuksekgonul et al., 2024) or hyper-parameters
 110 (Saad-Falcon et al., 2024). For example, (Chen et al., 2023a; Yuan et al., 2024; Fernando et al., 2023)
 111 adopt the evolution algorithms to automate the agent profiling.

112 Although effective, they merely optimize the agents but keep the workflow of the multi-agent system
 113 fixed, which still requires human effort to manually design for each new scenario. To address this
 114 problem, several methods (Li et al., 2024; Zhou et al., 2024b; Zhang et al., 2024c; Shang et al., 2024;
 115 Zhuge et al., 2023) are proposed to automate the design of the entire agentic workflow. For example,
 116 researchers (Zhuge et al., 2024; Khattab et al., 2024; Liu et al., 2023; Zhang et al., 2024a; Feng et al.,
 117 2025) formulate the workflows as graphs or networks and then optimize the connections between
 118 nodes. To improve the efficiency, ADAS (Hu et al., 2024) proposes to use programming codes to
 119 represent both agents and workflows. It also introduces a meta-agent to generate these workflows
 120 and presents the meta-agent search to optimize the designs of multi-agent systems. At the same
 121 time, AFLOW (Zhang et al., 2024b) also adopts the code representation but proposes Monte Carlo
 122 Tree Search (MCTS) to optimize it. In addition, MaAS (Zhang et al., 2025) presents the agentic
 123 supernet and then conducts the multi-agent architecture search. Differently, ScoreFlow (Wang et al.,
 124 2025) trains a workflow generator to generate better workflows via direct preference optimization
 125 (DPO) (Rafailov et al., 2023).

126 This paper argues that the previous methods are task-level meta-agents. As shown in Figure 1 (a),
 127 they merely generate a single task-specific multi-agent system for solving one kind of task. However,
 128 these one-size-fits-all systems are rigid and unable to automatically adapt or customize to individual
 129 user queries within a task. From this motivation, we aim to propose a query-level meta-agent to
 130 generate a query-specific multi-agent system for each user query as shown in Figure 1 (b).

131 **Reasoning in LLMs.** The ability to reason is essential for LLMs, enabling them to emulate
 132 human thinking patterns. Pioneering work (Wei et al., 2022; Kojima et al., 2022) has facilitated
 133 this by prompting LLMs to think step by step. Beyond this approach, reasoning capabilities are
 134 further enhanced through frameworks such as self-correction (Kumar et al., 2024), self-critique (Ke
 135 et al., 2023), debate (Liang et al., 2023; Du et al., 2023), and plan-and-solve (Wang et al., 2023).
 136 Additionally, efforts like (Hao et al., 2024; Goyal et al., 2023) seek to transition LLMs’ reasoning
 137 processes into the latent space. OpenAI has advanced reasoning in LLMs by developing the o1 model,
 138 demonstrating the potential for improvement through test-time scaling. Inspired by this, models such
 139 as QwQ (Team, 2024c), QvQ (Team, 2024b), DeepSeek (Team, 2025a), and Kimi (Team, 2025b)
 140 have followed suit, developing o1-like reasoning architectures. Moreover, OpenAI’s o3 model has
 141 been announced to achieve promising results on the ARG-AGI benchmark (ARC-AGI, 2024). LLMs
 142 progressively shift from intuitive processing (System 1) to deliberative reasoning (System 2) (Li et al.,
 143 2025). Besides, researchers demonstrate that reasoning can improve safety (Liu et al., 2025a) and
 144 alleviate hallucination (Gao et al., 2025). However, Chen et al. (2024) examines the overthinking
 145 problem observed in o1-like models. To alleviate this problem, token efficiency methods (Liu et al.,
 146 2025b) are proposed to reduce the token costs while maintaining the reasoning quality. This paper
 147 develops an o1-like reasoning model to serve as a query-level meta-agent, getting rid of complex
 148 search algorithms and the carefully designed search set.

149 3 PROBLEM DEFINITION

150 We denote a user **query** as q , e.g., *build a 2048 game*. Then a user **task** is defined as a distribution of
 151 user queries, denoted as $t = P(q)$, e.g., *code generation task*. A multi-agent system is denoted as
 152 $S = \{\mathcal{A}, \mathcal{W}\}$, where $\mathcal{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_n\}$ denotes the agents in the system and \mathcal{W} is the *workflow* of
 153 collaboration among the agents.

154 As shown in Figure 2 (a), in traditional multi-agent systems (Hong et al., 2023; Wu et al., 2023;
 155 Li et al., 2023), the agents and the workflows are designed manually according to one kind of task
 156 $t = P(q)$, as formulated $S_{\text{task}} = \mathcal{H}(t)$, where \mathcal{H} denotes human experts, and S_{task} denotes the
 157 **task-level multi-agent system** which is fixed for all queries in one task. This kind of manually
 158 designed multi-agent system leads to extensive human costs. Besides, such a one-size-fits-all system
 159 fails to allocate inference resources dynamically for different user queries within the task.

160 As illustrated in Figure 2 (b), search-based automatic multi-agent systems (Zhang et al., 2024b; Hu
 161 et al., 2024; Chen et al., 2023a) are proposed to reduce human effort. Specifically, these approaches

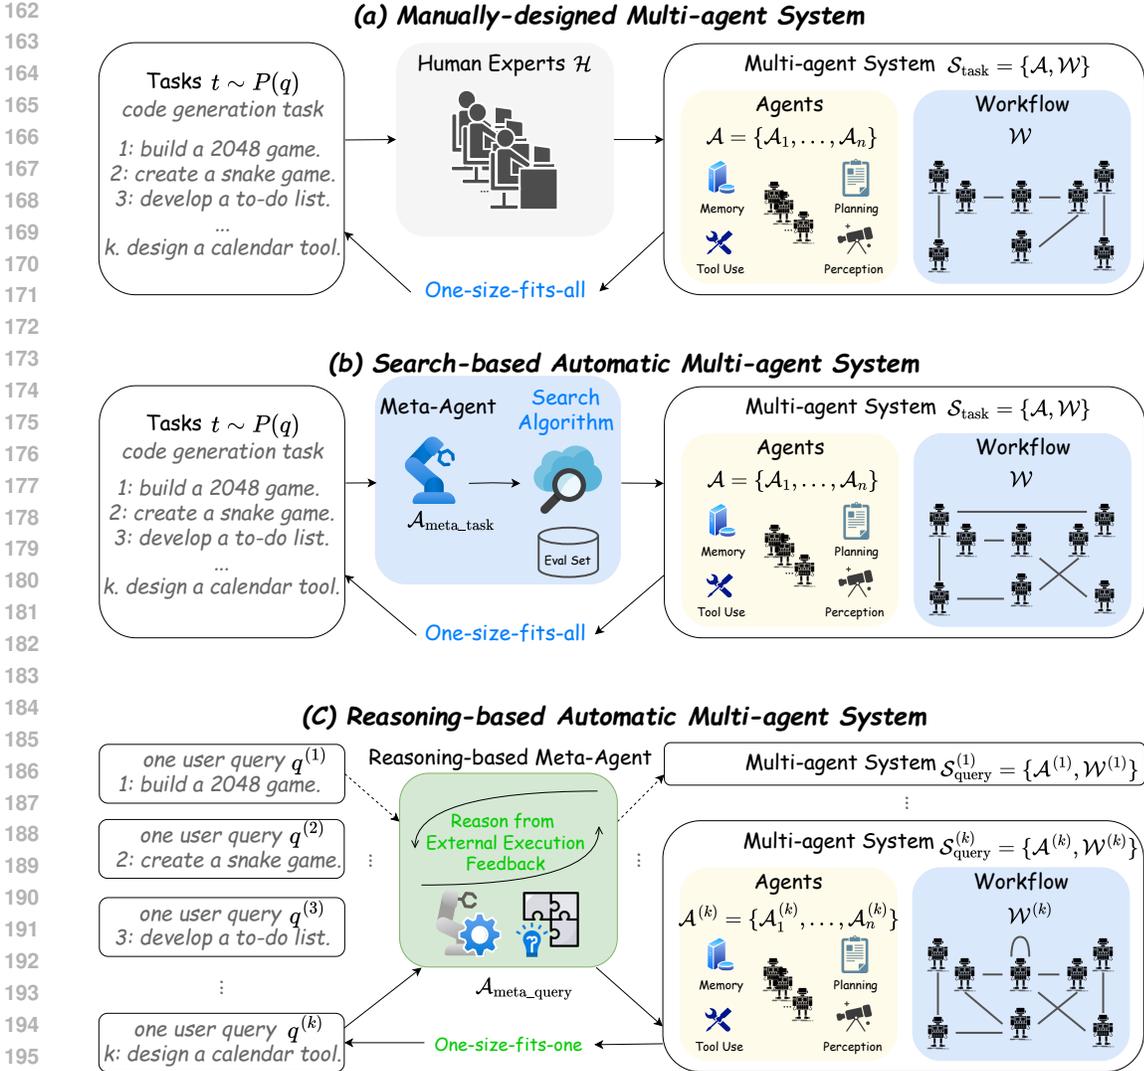


Figure 2: **Architectural Comparison of Three Multi-Agent Systems.** (a) Manually-designed Multi-agent System, (b) Search-based Automatic Multi-agent System, and (c) Reasoning-based Automatic Multi-agent System.

first prompt an LLM, such as GPT-4o, to act as a meta-agent, generating multiple candidate multi-agent system designs. Subsequently, complex search algorithms are employed on the carefully designed search set to identify the optimal system for completing the task. We denote this kind of meta-agent as **task-level meta-agent** $\mathcal{A}_{\text{meta_task}}$ since they can merely generate one **task-level multi-agent system** $\mathcal{S}_{\text{task}}$ to solve one kind of task t . Although effective, they are still one-size-fits-all systems. Besides, the search algorithm is time-consuming and relies on the search set, which is absent in one user query.

4 META-AGENT FLOWREASONER

To solve these problems, we develop a reasoning-based automatic multi-agent system shown in Figure 2 (c). Concretely, by guiding the model to reason from external execution feedback, we train a **query-level meta-agent** denoted as $\mathcal{A}_{\text{meta_query}}$, which can automatically propose a novel **query-level multi-agent system** for each user query q , as formulated $\mathcal{S}_{\text{query}} = \mathcal{A}_{\text{meta_query}}(q)$. Then, $\mathcal{S}_{\text{query}}$ accomplishes the specific user query and obtain the result a , i.e., $a = \mathcal{S}_{\text{query}}(q)$. Subsequently, the evaluator \mathcal{E} evaluates the performance of the proposed multi-agents system $\mathcal{S}_{\text{query}}$ by comparing the result a with ground truth a_{gt} as formulated $\mathcal{E}(a, a_{\text{gt}})$. Our proposed method is more practical

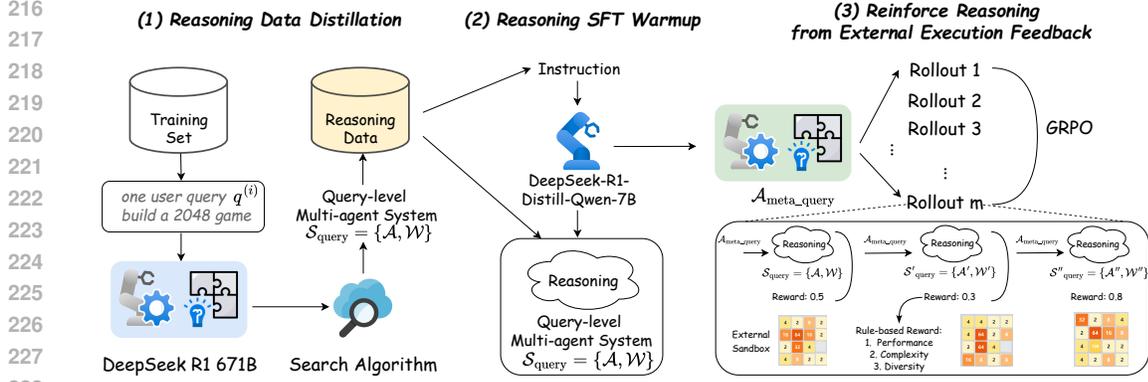


Figure 3: **Training Pipeline of FLOWREASONER.** It consists of (1) Reasoning Data Distillation, (2) Reasoning SFT Warmup, (3) Reinforce Reasoning from external execution feedback.

in real-world scenarios, as it can design an optimal multi-agent system for each specific query. The training pipeline of our method is demonstrated in Figure 3.

4.1 LEARN TO REASON

Reasoning Data Syntheses. To enable our model to learn how to reason workflows based on external execution feedback, we first generate multi-round reasoning data using R1-671B (Team, 2025a). For a given input query q , the R1 model generates l rounds of reasoning $\mathcal{R}_{\text{query}}$ and a multi-agent system $\mathcal{S}_{\text{query}}$ with external execution feedback at each round as follows:

$$\{\{\mathcal{R}_{\text{query}}, \mathcal{S}_{\text{query}}\}^{(1)}, \dots, \{\mathcal{R}_{\text{query}}, \mathcal{S}_{\text{query}}\}^{(l)}\} = \mathcal{A}_{\text{meta_query}}(q). \quad (1)$$

Then, we concatenate the model’s l rounds of reasoning and multi-agent system to form a reasoning process $\hat{\mathcal{R}}$ and final multi-agent system $\hat{\mathcal{S}}$, and pair it with the instruction \mathcal{I} and query q to construct our training sample $\{\mathcal{I}, q, \hat{\mathcal{R}}, \hat{\mathcal{S}}\}$. Based on this, we construct a warmup SFT dataset \mathcal{D} .

Reasoning SFT Warmup. After creating the reasoning training dataset \mathcal{D} , we proceed to perform reasoning SFT for warmup. We input the instruction \mathcal{I} and query q , then guide DeepSeek-R1-Distill-Qwen-7B to output reasoning process $\hat{\mathcal{R}}$ and final multi-agent system $\hat{\mathcal{S}}$. It can be formulated as follows:

$$\mathcal{L}_{\text{SFT}} = -\mathbb{E}_{(\mathcal{I}, q, \hat{\mathcal{R}}, \hat{\mathcal{S}}) \sim \mathcal{D}} \log P_{\theta}(\hat{\mathcal{R}}, \hat{\mathcal{S}} | \mathcal{I}, q), \quad (2)$$

where θ denotes the model parameters. Through SFT, we unlock the model’s reasoning ability regarding workflow generation.

4.2 REINFORCE REASONING FROM EXTERNAL EXECUTION FEEDBACK

After the SFT stage, we use reinforcement learning phase to further enhance the model’s reasoning capabilities through workflows built on external execution feedback. We aim to leverage feedback to improve performance on complex tasks in multi-round multi-agent systems. Following the DeepSeek-R1-Zero (Team, 2025a), we adopt standard GRPO (Grouped Relative Policy Optimization) as our training method. GRPO works by sampling multiple outputs for each query and computing relative advantages based on the rewards these outputs receive. The policy is then updated by maximizing an objective function incorporating these relative advantages. The GRPO objective function can be expressed as follows:

$$\mathcal{L}_{\text{GRPO}}(\theta) = \mathbb{E}_{q \sim t, \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(O|q)} \frac{1}{G} \sum_{i=1}^G \left\{ \min[r_{\text{ratio}}, \text{clip}(r_{\text{ratio}}, 1 - \varepsilon, 1 + \varepsilon)] \cdot \hat{E}_{i,m} - \beta \mathbb{D}_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}) \right\}, \quad (3)$$

where G represents the number of sampled trajectories in each group, o_i denotes the i -th trajectory, $|o_i|$ is the length of trajectory o_i , $r_{\text{ratio}} = \frac{\pi_{\theta}(o_i, m | q, o_i, < m)}{\pi_{\theta_{\text{old}}}(o_i, m | q, o_i, < m)}$ represents the probability ratio between

new and old policies, $\hat{E}_{i,m}$ is the estimated advantage for the m -th token in the i -th trajectory, ε is a small positive clipping parameter that prevents excessively large update steps, and the KL regularization term constrains policy drift to maintain stability during training.

Since we can obtain a reward from the external environment in every round, we use process reward supervision. The reward for each reasoning round is normalized, and the advantage function is calculated as the sum of subsequent normalized rewards.

$$\hat{E}_{i,m} = \sum_{j \geq T} \tilde{r}_j^{(i)}. \quad (4)$$

Here, the normalized reward for the j -th round of the i -th candidate output is defined as $\tilde{r}_j^{(i)} = k^j \cdot \frac{r_j^{(i)} - \text{mean}(\mathbf{R})}{\text{std}(\mathbf{R})}$, where k is a scaling factor and T is a threshold used to exclude the first T items from the calculation. The set \mathbf{R} represents the list of scores for each round across all candidates, and $r_j^{(i)}$ is the score of candidate i in round j . Each score $r_j^{(i)}$ is calculated by the performance of the proposed solution (i.e., pass rate), algorithm complexity (i.e., complexity score of abstract syntax tree) and diversity (i.e., distinctness ratio followed by [Chen et al. \(2024\)](#)) of workflow.

4.3 GENERATE MULTI-AGENT SYSTEMS WITH FLOWREASONER

Constructing a multi-agent system is essentially an optimization problem with the goal of designing an optimized system, $\mathcal{S}_{\text{query}}^*$, that responds to user queries. When a user submits a query q , the system produces an answer $a = \mathcal{S}_{\text{query}}(q)$. The performance of the system in each round is evaluated using an external feedback function $\mathcal{E}(a, a_{\text{gt}})$, where a_{gt} represents the ground truth answer. In this framework, a meta-agent, $\mathcal{A}_{\text{meta_query}}$, is responsible for optimizing the workflow with external execution feedback $\mathcal{E}(a, a_{\text{gt}})$ in every round. The pass rate of the proposed solution for the given query, as the key performance indicator of the system, serves as the external execution feedback. The optimization space is defined by all possible configurations of nodes and edges. Here, nodes represent various parameters (such as language models, prompts, temperature, and output formats), and edges capture the interactions or data flows between these nodes.

By representing both nodes and edges in code and employing predefined operators (such as Ensemble, Review, and Revise) along with a custom operator to combine these elements, FLOWREASONER utilizes an l -round optimization process same as Aflow ([Zhang et al., 2024b](#)) to arrive at the final multi-agent system:

$$\mathcal{S}_{\text{query}}^* = \arg \max_{\mathcal{S}_{\text{query}} \in \mathbf{S}} \mathcal{E}(\mathcal{S}_{\text{query}}(q), a_{\text{gt}}), \quad (5)$$

where $\mathcal{S}_{\text{query}}^*$ is the optimal multi-agent system refined through optimization. By optimizing a multi-agent system through iterative external execution feedback, FLOWREASONER can construct a highly adaptive system that maximizes the accuracy and performance of solution for a query.

5 EXPERIMENTS

Datasets. Given our focus on creating workflows tailored to individual user queries rather than the general task, we restrict our scope to *code generation tasks*, as they can provide test cases as external execution feedback for the workflow construction process. Among various benchmarks, we consider three representative datasets: [BigCodeBench \(Zhuo et al., 2024\)](#), which emphasizes engineering-oriented tasks, and two algorithmically focused benchmarks, [HumanEval \(Chen et al., 2021\)](#) and [MBPP \(Austin et al., 2021\)](#). This selection enables us to comprehensively evaluate workflow discovery across a diverse spectrum of code generation challenges.

Baselines. In line with prior work ([Hu et al., 2024](#); [Zhang et al., 2024b](#)), we evaluate FLOWREASONER against three categories of baselines: (1) single-model direct invocation, where a single LLM is prompted to solve the problem without additional structure; (2) manually designed workflows, including Self-Refine ([Madaan et al., 2023](#)), LLM-Debate ([Du et al., 2023](#)), and LLM-Blender ([Jiang et al., 2023](#)), which incorporate human-crafted reasoning strategies; and (3) automated workflow optimization methods, such as Aflow ([Zhang et al., 2024b](#)), ADAS ([Hu et al., 2024](#)), and MaAS ([Zhang et al., 2025](#)), which construct workflows through search or optimization over possible reasoning

Table 1: **Performance Evaluation.** Accuracy comparison across three code benchmarks for three categories of baselines - individual models, manual workflows, and automated workflow methods - alongside our FLOWREASONER-14B. For manual methods, model names in parentheses indicate the worker model used.

Type	Method	BigCodeBench	HumanEval	MBPP	Overall
Vanilla	o1-mini	57.67	95.42	74.19	71.37
	GPT-4o-mini	56.33	88.55	71.73	68.60
Manual	Self-Refine (4o-mini) (Madaan et al., 2023)	54.78	89.83	69.64	67.29
	LLM-Debate (4o-mini) (Du et al., 2023)	56.88	91.64	70.28	68.69
	LLM-Blender (4o-mini) (Jiang et al., 2023)	57.46	89.44	76.39	71.25
	Self-Refine (o1-mini) (Madaan et al., 2023)	56.68	94.74	73.64	70.63
	LLM-Debate (o1-mini) (Du et al., 2023)	57.25	95.83	74.28	71.33
	LLM-Blender (o1-mini) (Jiang et al., 2023)	59.51	96.37	78.65	74.22
Auto	AutoAgents (Chen et al., 2023a)	56.65	88.91	72.03	68.92
	ADAS (Hu et al., 2024)	53.87	84.26	68.47	65.48
	Aflow (Zhang et al., 2024b)	59.83	94.15	82.40	75.63
	MaAS (Zhang et al., 2025)	60.33	95.42	84.16	76.81
	ScoreFlow (Wang et al., 2025)	60.71	95.67	84.73	77.25
Ours	FLOWREASONER-14B	63.53	97.26	92.15	81.89

Table 2: **Ablation Study on Model Sizes and Training Stages.** Accuracy(%) comparison across three code benchmarks for models of different sizes (7B/14B) at both the Supervised Fine-Tuning (SFT) stage and the combined SFT with Reinforcement Learning (SFT+RL) stage.

Stage	Size	BigCodeBench	HumanEval	MBPP	Overall
SFT	7B	61.79	96.38	87.22	78.89
SFT+RL	7B	62.78	96.95	89.86	80.53
SFT	14B	62.83	97.18	91.91	81.50
SFT+RL	14B	63.53	97.26	92.15	81.89

structures. This comparison enables a comprehensive assessment of FLOWREASONER’s effectiveness relative to both static and adaptive baselines.

Implementation Details. For the manually designed workflow baselines, we employed both o1-mini and GPT-4o-mini as worker models for each method. For the automated workflow optimization baselines, we adopted the original configurations as described in MaAS (Zhang et al., 2025). In our proposed method, FLOWREASONER, we trained two variants of DeepSeek-R1-Distill-Qwen with 7B and 14B parameters, respectively, and used o1-mini as the worker model. We fixed the number of workflow iterations to 10. To assess performance, we used the standard pass@1 metric for code accuracy, consistent with prior work (Chen et al., 2021). More details can be found in Appendix A.

5.1 EXPERIMENT RESULTS

Performance Comparison. Table 1 presents the performance comparison between our proposed method and the baselines. FLOWREASONER-14B consistently outperforms all competing approaches across the three benchmark datasets. Notably, it achieves an overall improvement of 5 percentage points over the strongest baseline, MaAS, and exceeds the performance of its underlying worker model, o1-mini, by a substantial margin of 10%. These results highlight the effectiveness of our workflow-based reasoning framework in enhancing code generation accuracy.

Ablation on Model Size and Training Stages. To investigate the impact of model size and training stages on performance, we conducted an ablation study comparing different configurations of FLOWREASONER in Table 2. We observed that the 14B variant consistently outperformed the 7B counterpart across all benchmarks, indicating a positive correlation between model scale and reasoning effectiveness. Furthermore, within each model size, versions trained with both SFT and RL exhibited notable improvements over those trained with SFT alone, demonstrating the complementary benefits of incorporating RL in enhancing workflow-guided code generation.

Ablation of Meta-agents and Workers. To analyze the impact of meta-agent and worker selection, we conducted an ablation study on the BigCodeBench dataset. Figure 4 presents performance

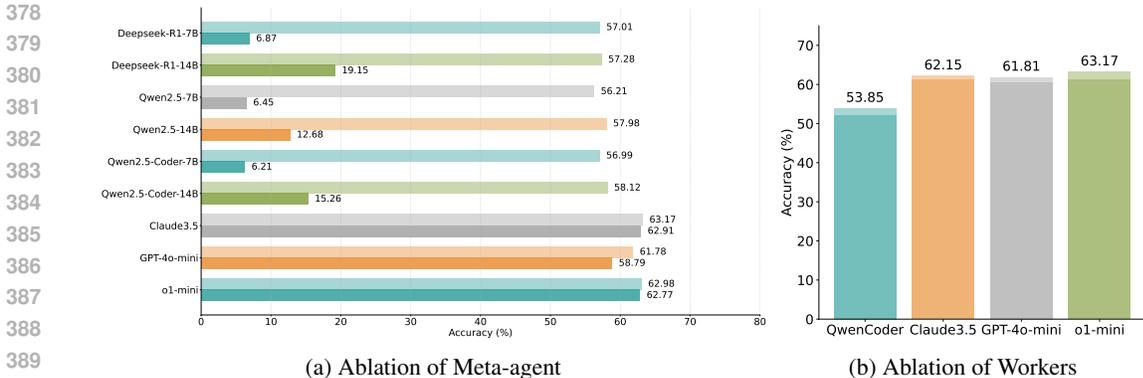


Figure 4: **Ablation of Meta-agent and Workers.** (a) Accuracy of different meta-agents with o1-mini as workers. (b) Accuracy of the generated workflow with different workers. The light bars show results with the initial workflow (multi-round iterations), while the dark bars show results without it.

Table 3: **Generalization Evaluation.** Accuracy of our trained meta-agent FLOWREASONER-7B/14B when paired with alternative workers including Qwen2.5 Coder, Claude, and GPT4o-mini.

Worker/Meta-agent	FLOWREASONER-7B (SFT+RL)			FLOWREASONER-14B (SFT+RL)		
	BigCodeBench	HumanEval	MBPP	BigCodeBench	HumanEval	MBPP
o1-mini	62.77	96.95	89.86	63.53	97.26	92.15
Qwen2.5Coder-32b	50.17	92.89	80.40	52.67	93.69	78.90
Claude3.5	60.67	96.07	87.63	61.12	96.52	89.82
GPT-4o-mini	59.18	94.24	82.19	59.75	94.52	82.27

comparisons under various meta-agent and worker configurations. As shown in Figure 4 (a), open-source models exhibited poor performance when paired with o1-mini as the worker and no initial workflow, frequently generating error-prone workflows. This highlights a limitation of current open-source models, which struggle to produce reliable workflows solely based on instruction prompts and are heavily reliant on predefined, manually crafted workflows. In contrast, API-based models demonstrated stronger performance, likely attributable to their superior instruction-following capabilities. Figure 4 (b) further examines worker effectiveness when using a high-performing meta-agent (Claude 3.5). We compared the open-source model Qwen2.5-Coder-32B and three API-based models, including Claude 3.5, GPT-4o-mini, and o1-mini. Among these, o1-mini achieved the best overall performance, suggesting its suitability as a worker model in the FLOWREASONER framework.

Generalization Capability. To evaluate whether FLOWREASONER, trained with o1-mini as the worker, can generalize its planning capabilities to alternative workers, we conducted experiments by substituting the worker with Qwen2.5-Coder, Claude, and GPT-4o-mini, while keeping the meta-agent fixed as either FLOWREASONER-7B or FLOWREASONER-14B. As shown in Table 3, FLOWREASONER exhibits a notable degree of transferability, maintaining consistent performance across different worker models on the same tasks. These results suggest that the planner is not tightly bound to a specific worker and can adapt its strategies effectively across diverse execution agents.

Case Study. Figure 5 illustrates two example workflows generated by FLOWREASONER-14B for representative tasks from BigCodeBench and HumanEval, respectively. The workflow corresponding to the BigCodeBench task exhibits greater complexity, reflecting the more challenging and engineering-oriented nature of the task. In contrast, the HumanEval workflow is substantially more concise, aligning with the relative simplicity and algorithmic focus of the task. These examples demonstrate FLOWREASONER’s ability to adapt the structure and granularity of workflows based on task complexity. More cases can be found in Appendix B.

6 CONCLUSION

In this paper, we present FLOWREASONER, a query-level meta-agent designed to automate the creation of personalized multi-agent systems for individual user queries. Unlike previous task-level approaches that create fixed, one-size-fits-all systems, FLOWREASONER dynamically generates tai-

```

432 1 async def __call__(self, problem: str, test_cases: list):
433 2     max_retries = 3
434 3     solutions = []
435 4
436 5     for _ in range(max_retries):
437 6         for attempt in range(3):
438 7             # Try generating code up to 3 times
439 8             solution = await self.custom_code_generate(
440 9                 problem=problem, instruction=prompt_custom.GENERATE_CODE_PROMPT
441 10            )
442 11             test_result = await self.test(
443 12                 problem=problem, solution=solution['response'], test_cases=test_cases
444 13            )
445 14
446 15             if test_result['result']:
447 16                 solutions.append(solution['response'])
448 17                 break # Break the inner loop if a correct solution is found
449 18             elif attempt < 2:
450 19                 # If it's not the last attempt, try to improve the solution
451 20                 error_analysis = await self.custom(
452 21                     input=f"Problem: {problem}\nFailed solution: {solution['response']}\n
453 22                     nError: {test_result['solution']}", instruction=prompt_custom.
454 23                     ERROR_ANALYSIS_PROMPT
455 24                )
456 25                 prompt_custom.GENERATE_CODE_PROMPT = f"{prompt_custom.
457 26                     GENERATE_CODE_PROMPT}\n{error_analysis['response']}"
458 27
459 28                 if not solutions: # If no correct solution was found after 3 attempts
460 29                     solutions.append(solution['response']) # Add the last generated solution
461
462             best_solution = await self.sc_ensemble(solutions=solutions, problem=problem)
463             return best_solution['response'], self.llm.cost_manager.total_cost

```

(a) Workflow generated by FLOWREASONER-14B for the BigCodeBench task of *generating traffic data for various vehicle types over a specified number of hours, saving the data to a CSV file with columns, and plotting it in a line chart.*

```

458 1     async def __call__(self, problem: str, test_cases: list):
459 2         solution = await self.custom_code_generate(problem=problem, entry_point=
460 3             entry_point, instruction=prompt_custom.CODE_GENERATE_PROMPT)
461 4
462 5         # Add review to check the solution
463 6         review = await self.custom(
464 7             input=f"Problem: {problem}\nGenerated solution: {solution['response']}",
465 8             instruction=prompt_custom.REVIEW_PROMPT
466 9         )
467 10
468 11         # Improve the solution based on the review information
469 12         improved_solution = await self.custom_code_generate(
470 13             problem=problem, instruction=f"{prompt_custom.CODE_IMPROVE_PROMPT}\n\nConsider
471 14             this review: {review['response']}"
472
473         return improved_solution['response'], self.llm.cost_manager.total_cost

```

(b) Workflow generated by FLOWREASONER-14B for the HumanEval task of *splitting the string into words and returning an array of the words.*

Figure 5: Cases of Workflows generated by FLOWREASONER-14B for the tasks of BigCodeBench and HumanEval.

lored workflows for each specific query through reasoning-based optimization. The system leverages external execution feedback and reinforcement learning with multi-purpose rewards focusing on performance, complexity, and efficiency to generate optimized workflows without relying on complex search algorithms or carefully designed search sets. Experimental results demonstrate that FLOWREASONER-14B outperforms both manually designed workflows and existing automated methods across multiple code generation benchmarks, notably improving o1-mini’s performance by overall 10.52% on three benchmarks, thus proving the effectiveness and adaptability of reasoning-driven workflow generation. Besides, the observation when pairing FLOWREASONER with different worker models further confirms its generalization capabilities. Our approach reduces human resource costs while enhancing scalability by enabling more adaptive and efficient multi-agent systems that dynamically optimize their structure based on specific user queries rather than relying on fixed workflows for entire task categories.

ETHICS STATEMENT

Our research on automated multi-agent system design does not involve human subjects and does not raise concerns regarding discrimination, bias, fairness, privacy, or security. The entire construction process focuses solely on technical methodology for code generation workflow optimization without involving sensitive data or human participation.

REPRODUCIBILITY STATEMENT

To ensure the reproducibility of our research, we have included all the code required to reproduce our experimental results in the supplementary material. Furthermore, we are committed to making our code, data, and trained models publicly available upon publication. This will allow the research community to verify our results and build upon our work.

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Anthropic. Introducing computer use, a new claude 3.5 sonnet, and claude 3.5 haiku. <https://www.anthropic.com/news/3-5-models-and-computer-use>, 2024.
- ARC-AGI. Abstraction and reasoning corpus for artificial general intelligence. <https://github.com/fchollet/ARC-AGI/>, 2024.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- Guangyao Chen, Siwei Dong, Yu Shu, Ge Zhang, Jaward Sesay, Börje F Karlsson, Jie Fu, and Yemin Shi. Autoagents: A framework for automatic agent generation. *arXiv preprint arXiv:2309.17288*, 2023a.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chen Qian, Chi-Min Chan, Yujia Qin, Yaxi Lu, Ruobing Xie, et al. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors in agents. *arXiv preprint arXiv:2308.10848*, 2(4):6, 2023b.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, et al. Do not think that much for $2+3=?$ on the overthinking of o1-like llms. *arXiv preprint arXiv:2412.21187*, 2024.
- CognitionAI. Introducing devin, the first ai software engineer. <https://www.cognition.ai/blog/introducing-devin/>, 2024.
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate, 2023. URL <https://arxiv.org/abs/2305.14325>, 3, 2023.
- Shangbin Feng, Zifeng Wang, Palash Goyal, Yike Wang, Weijia Shi, Huang Xia, Hamid Palangi, Luke Zettlemoyer, Yulia Tsvetkov, Chen-Yu Lee, et al. Heterogeneous swarms: Jointly optimizing model roles and weights for multi-llm systems. *arXiv preprint arXiv:2502.04510*, 2025.
- Chrisantha Fernando, Dylan Banarse, Henryk Michalewski, Simon Osindero, and Tim Rocktäschel. Promptbreeder: Self-referential self-improvement via prompt evolution. *arXiv preprint arXiv:2309.16797*, 2023.

- 540 Hongcheng Gao, Jiashu Qu, Jingyi Tang, Baolong Bi, Yue Liu, Hongyu Chen, Li Liang, Li Su, and
541 Qingming Huang. Exploring hallucination of large multimodal models in video understanding:
542 Benchmark, analysis and mitigation. *arXiv preprint arXiv:2503.19622*, 2025.
- 543
- 544 Sachin Goyal, Ziwei Ji, Ankit Singh Rawat, Aditya Krishna Menon, Sanjiv Kumar, and Vaishnavh
545 Nagarajan. Think before you speak: Training language models with pause tokens. *arXiv preprint*
546 *arXiv:2310.02226*, 2023.
- 547
- 548 Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V Chawla, Olaf Wiest,
549 and Xiangliang Zhang. Large language model based multi-agents: A survey of progress and
550 challenges. *arXiv preprint arXiv:2402.01680*, 2024.
- 551
- 552 Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong
553 Tian. Training large language models to reason in a continuous latent space. *arXiv preprint*
554 *arXiv:2412.06769*, 2024.
- 555
- 556 Sirui Hong, Xiawu Zheng, Jonathan Chen, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang,
557 Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, et al. Metagpt: Meta programming for multi-agent
558 collaborative framework. *arXiv preprint arXiv:2308.00352*, 3(4):6, 2023.
- 559
- 560 Shengran Hu, Cong Lu, and Jeff Clune. Automated design of agentic systems. *arXiv preprint*
561 *arXiv:2408.08435*, 2024.
- 562
- 563 Yichong Huang, Xiaocheng Feng, Baohang Li, Yang Xiang, Hui Wang, Ting Liu, and Bing Qin.
564 Ensemble learning for heterogeneous large language models with deep parallel collaboration.
565 In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.),
566 *Advances in Neural Information Processing Systems*, volume 37, pp. 119838–119860. Curran
567 Associates, Inc., 2024.
- 568
- 569 Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. Llm-blender: Ensembling large language models
570 with pairwise ranking and generative fusion. *arXiv preprint arXiv:2306.02561*, 2023.
- 571
- 572 Can Jin, Hongwu Peng, Qixin Zhang, Yujin Tang, Dimitris N Metaxas, and Tong Che. Two heads
573 are better than one: Test-time scaling of multi-agent collaborative reasoning. *arXiv preprint*
574 *arXiv:2504.09772*, 2025.
- 575
- 576 Pei Ke, Bosi Wen, Zhuoer Feng, Xiao Liu, Xuanyu Lei, Jiale Cheng, Shengyuan Wang, Aohan
577 Zeng, Yuxiao Dong, Hongning Wang, et al. Critiquellm: Scaling llm-as-critic for effective and
578 explainable evaluation of large language model generation. *arXiv preprint arXiv:2311.18702*,
579 2023.
- 580
- 581 Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Saiful
582 Haq, Ashutosh Sharma, Thomas T Joshi, Hanna Moazam, Heather Miller, et al. Dspy: Compiling
583 declarative language model calls into state-of-the-art pipelines. In *The Twelfth International*
584 *Conference on Learning Representations*, 2024.
- 585
- 586 Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair,
587 Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source
588 vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- 589
- 590 Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large
591 language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:
592 22199–22213, 2022.
- 593
- 594 Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli,
595 Shariq Iqbal, Colton Bishop, Rebecca Roelofs, et al. Training language models to self-correct via
596 reinforcement learning. *arXiv preprint arXiv:2409.12917*, 2024.
- 597
- 598 Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. Camel: Com-
599 municative agents for "mind" exploration of large language model society. *Advances in Neural*
600 *Information Processing Systems*, 36:51991–52008, 2023.

- 594 Zelong Li, Shuyuan Xu, Kai Mei, Wenyue Hua, Balaji Rama, Om Raheja, Hao Wang, He Zhu, and
595 Yongfeng Zhang. Autoflow: Automated workflow generation for large language model agents.
596 *arXiv preprint arXiv:2407.12821*, 2024.
- 597
598 Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian
599 Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, et al. From system 1 to system 2: A survey of
600 reasoning large language models. *arXiv preprint arXiv:2502.17419*, 2025.
- 601 Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi,
602 and Zhaopeng Tu. Encouraging divergent thinking in large language models through multi-agent
603 debate. *arXiv preprint arXiv:2305.19118*, 2023.
- 604
605 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,
606 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint*
607 *arXiv:2412.19437*, 2024.
- 608 Yue Liu, Hongcheng Gao, Shengfang Zhai, Jun Xia, Tianyi Wu, Zhiwei Xue, Yulin Chen, Kenji
609 Kawaguchi, Jiaheng Zhang, and Bryan Hooi. Guardreasoner: Towards reasoning-based llm
610 safeguards. *arXiv preprint arXiv:2501.18492*, 2025a.
- 611 Yue Liu, Jiaying Wu, Yufei He, Hongcheng Gao, Hongyu Chen, Baolong Bi, Jiaheng Zhang, Zhiqi
612 Huang, and Bryan Hooi. Efficient inference for large reasoning models: A survey. *arXiv preprint*
613 *arXiv:2503.23077*, 2025b.
- 614
615 Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. Dynamic llm-agent network: An
616 llm-agent collaboration framework with agent team optimization. *arXiv preprint arXiv:2310.02170*,
617 2023.
- 618 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri
619 Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement
620 with self-feedback. *Advances in Neural Information Processing Systems*, 36:46534–46594, 2023.
- 621
622 OpenAI. Introducing chatgpt. <https://openai.com/index/chatgpt/>, 2022.
- 623
624 OpenAI. Learning to reason with llms. <https://openai.com/index/learning-to-reason-with-llms/>, 2024.
- 625
626 OpenAI. Introducing deep research. <https://openai.com/index/introducing-deep-research/>, 2025.
- 627 Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S
628 Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th*
629 *annual acm symposium on user interface software and technology*, pp. 1–22, 2023.
- 630 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
631 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances*
632 *in Neural Information Processing Systems*, 36:53728–53741, 2023.
- 633 Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste
634 Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. Gemini
635 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint*
636 *arXiv:2403.05530*, 2024.
- 637
638 Jon Saad-Falcon, Adrian Gamarra Lafuente, Shlok Natarajan, Nahum Maru, Hristo Todorov, Etash
639 Guha, E Kelly Buchanan, Mayee Chen, Neel Guha, Christopher Ré, et al. Archon: An architecture
640 search framework for inference-time techniques. *arXiv preprint arXiv:2409.15254*, 2024.
- 641 Yu Shang, Yu Li, Keyu Zhao, Likai Ma, Jiahe Liu, Fengli Xu, and Yong Li. Agentsquare: Automatic
642 llm agent search in modular design space. *arXiv preprint arXiv:2410.06153*, 2024.
- 643
644 Anthropic Team. The claude 3 model family: Opus, sonnet, haiku. [https://www-](https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf)
645 [cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf](https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf),
646 2024a.
- 647 Deepseek Team. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning.
arXiv preprint arXiv:2501.12948, 2025a.

- 648 Kimi Team. Kimi k1.5: Scaling reinforcement learning with llms. *arXiv preprint 2501.12599v1*,
649 2025b.
- 650
- 651 Qwen Team. Qvq: To see the world with wisdom. <https://qwenlm.github.io/blog/qvq-72b-preview/>,
652 2024b.
- 653 Qwen Team. Qwq: Reflect deeply on the boundaries of the unknown.
654 <https://qwenlm.github.io/blog/qwq-32b-preview/>, 2024c.
- 655
- 656 Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi Lan, Roy Ka-Wei Lee, and Ee-Peng Lim.
657 Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language
658 models. *arXiv preprint arXiv:2305.04091*, 2023.
- 659
- 660 Yinjie Wang, Ling Yang, Guohao Li, Mengdi Wang, and Bryon Aragam. Scoreflow: Mastering llm
661 agent workflows via score-based preference optimization. *arXiv preprint arXiv:2502.04306*, 2025.
- 662 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
663 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*
664 *neural information processing systems*, 35:24824–24837, 2022.
- 665
- 666 Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang,
667 Shaokun Zhang, Jiale Liu, et al. Autogen: Enabling next-gen llm applications via multi-agent
668 conversation. *arXiv preprint arXiv:2308.08155*, 2023.
- 669
- 670 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
671 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint*
672 *arXiv:2412.15115*, 2024.
- 673 Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen.
674 Large language models as optimizers. *arXiv preprint arXiv:2309.03409*, 2023.
- 675
- 676 Siyu Yuan, Kaitao Song, Jiangjie Chen, Xu Tan, Dongsheng Li, and Deqing Yang. Evoagent: Towards
677 automatic multi-agent generation via evolutionary algorithms. *arXiv preprint arXiv:2406.14228*,
678 2024.
- 679
- 680 Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Zhi Huang, Carlos Guestrin, and
681 James Zou. Textgrad: Automatic "differentiation" via text. *arXiv preprint arXiv:2406.07496*,
682 2024.
- 683 Guibin Zhang, Yanwei Yue, Xiangguo Sun, Guancheng Wan, Miao Yu, Junfeng Fang, Kun Wang,
684 and Dawei Cheng. G-designer: Architecting multi-agent communication topologies via graph
685 neural networks. *arXiv preprint arXiv:2410.11782*, 2024a.
- 686
- 687 Guibin Zhang, Luyang Niu, Junfeng Fang, Kun Wang, Lei Bai, and Xiang Wang. Multi-agent
688 architecture search via agentic supernet. *arXiv preprint arXiv:2502.04180*, 2025.
- 689
- 690 Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xionghui Chen, Jiaqi Chen, Mingchen Zhuge,
691 Xin Cheng, Sirui Hong, Jinlin Wang, et al. Aflow: Automating agentic workflow generation. *arXiv*
692 *preprint arXiv:2410.10762*, 2024b.
- 693
- 694 Shaokun Zhang, Jieyu Zhang, Jiale Liu, Linxin Song, Chi Wang, Ranjay Krishna, and Qingyun
695 Wu. Offline training of language model agents with functions as learnable weights. In *Forty-first*
696 *International Conference on Machine Learning*, 2024c.
- 697
- 698 Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and
699 Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. *arXiv*
700 *preprint arXiv:2403.13372*, 2024.
- 701
- 702
- 703
- 704
- 705
- 706
- 707
- 708
- 709
- 710
- 711
- 712
- 713
- 714
- 715
- 716
- 717
- 718
- 719
- 720
- 721
- 722
- 723
- 724
- 725
- 726
- 727
- 728
- 729
- 730
- 731
- 732
- 733
- 734
- 735
- 736
- 737
- 738
- 739
- 740
- 741
- 742
- 743
- 744
- 745
- 746
- 747
- 748
- 749
- 750
- 751
- 752
- 753
- 754
- 755
- 756
- 757
- 758
- 759
- 760
- 761
- 762
- 763
- 764
- 765
- 766
- 767
- 768
- 769
- 770
- 771
- 772
- 773
- 774
- 775
- 776
- 777
- 778
- 779
- 780
- 781
- 782
- 783
- 784
- 785
- 786
- 787
- 788
- 789
- 790
- 791
- 792
- 793
- 794
- 795
- 796
- 797
- 798
- 799
- 800
- 801
- 802
- 803
- 804
- 805
- 806
- 807
- 808
- 809
- 810
- 811
- 812
- 813
- 814
- 815
- 816
- 817
- 818
- 819
- 820
- 821
- 822
- 823
- 824
- 825
- 826
- 827
- 828
- 829
- 830
- 831
- 832
- 833
- 834
- 835
- 836
- 837
- 838
- 839
- 840
- 841
- 842
- 843
- 844
- 845
- 846
- 847
- 848
- 849
- 850
- 851
- 852
- 853
- 854
- 855
- 856
- 857
- 858
- 859
- 860
- 861
- 862
- 863
- 864
- 865
- 866
- 867
- 868
- 869
- 870
- 871
- 872
- 873
- 874
- 875
- 876
- 877
- 878
- 879
- 880
- 881
- 882
- 883
- 884
- 885
- 886
- 887
- 888
- 889
- 890
- 891
- 892
- 893
- 894
- 895
- 896
- 897
- 898
- 899
- 900
- 901
- 902
- 903
- 904
- 905
- 906
- 907
- 908
- 909
- 910
- 911
- 912
- 913
- 914
- 915
- 916
- 917
- 918
- 919
- 920
- 921
- 922
- 923
- 924
- 925
- 926
- 927
- 928
- 929
- 930
- 931
- 932
- 933
- 934
- 935
- 936
- 937
- 938
- 939
- 940
- 941
- 942
- 943
- 944
- 945
- 946
- 947
- 948
- 949
- 950
- 951
- 952
- 953
- 954
- 955
- 956
- 957
- 958
- 959
- 960
- 961
- 962
- 963
- 964
- 965
- 966
- 967
- 968
- 969
- 970
- 971
- 972
- 973
- 974
- 975
- 976
- 977
- 978
- 979
- 980
- 981
- 982
- 983
- 984
- 985
- 986
- 987
- 988
- 989
- 990
- 991
- 992
- 993
- 994
- 995
- 996
- 997
- 998
- 999
- 1000

702 Wangchunshu Zhou, Yixin Ou, Shengwei Ding, Long Li, Jialong Wu, Tiannan Wang, Jiamin Chen,
703 Shuai Wang, Xiaohua Xu, Ningyu Zhang, et al. Symbolic learning enables self-evolving agents.
704 *arXiv preprint arXiv:2406.18532*, 2024b.
705

706 Mingchen Zhuge, Haozhe Liu, Francesco Faccio, Dylan R Ashley, Róbert Csordás, Anand Gopalakr-
707 ishnan, Abdullah Hamdi, Hasan Abed Al Kader Hammoud, Vincent Herrmann, Kazuki Irie, et al.
708 Mindstorms in natural language-based societies of mind. *arXiv preprint arXiv:2305.17066*, 2023.

709 Mingchen Zhuge, Wenyi Wang, Louis Kirsch, Francesco Faccio, Dmitrii Khizbullin, and Jürgen
710 Schmidhuber. Gptswarm: Language agents as optimizable graphs. In *Forty-first International
711 Conference on Machine Learning*, 2024.
712

713 Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widayarsi, Imam
714 Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, et al. Bigcodebench: Benchmarking code
715 generation with diverse function calls and complex instructions. *arXiv preprint arXiv:2406.15877*,
716 2024.
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

A MORE IMPLEMENTATION DETAILS

In the workflow optimization process of our method, we evaluate each workflow for 3 times as external execution feedback. We use the following 6 operators based on Aflow (Zhang et al., 2024b) as our base operators:

- Code Generator: generates code solutions for a given problem.
- Format Generator: produces formatted answers for a given problem.
- Ensemble Operator: combines multiple solutions or approaches to create a more robust final result.
- Review Operator: evaluates solutions for correctness, efficiency, and adherence to requirements.
- Revise Operator: refines solutions based on feedback from the review process.
- Code Test Operator: executes and validates code solutions against test cases to ensure functionality.

For our supervised fine-tuning (SFT), we utilized approximately 1,400 items sourced from three datasets generated by R1 through our optimization process. We conducted the training using LLaMA-Factory (Zheng et al., 2024), with a per-device training batch size of 1, gradient accumulation over 2 steps, a learning rate of 1e-5, and max train epochs of 3. For the reinforcement learning (RL) phase, we set the scaling factor k to 1.1, the threshold T to 3, rollout number m to 5, and max episodes to 5. All experiments are conducted on NVIDIA A100 GPUs with 80GB of memory. The total cost of training stage is about 150 GPU hours.

The multi-purpose reward in RL is designed to balance correctness, efficiency, and exploration. However, both diversity and complexity rewards showed minimal improvement in accuracy-driven benchmarks. Consequently, we set the coefficients for these two rewards to 0.

B MORE CASES OF WORKFLOW

In this section, we present additional cases of both successful and failed results generated by FLOWREASONER.

B.1 SUCCESSFUL CASES

Fig. 6, Fig. 8 and Fig. 7 demonstrate three successful cases generated by FLOWREASONER for BigCodeBench, HumanEval and MBPP.

B.2 FAILURE CASES

Fig. 10, Fig. 11 and Fig. 9 demonstrate three failure cases generated by FLOWREASONER for BigCodeBench, HumanEval and MBPP.

C LIMITATION

Despite demonstrating superior performance over existing approaches, FLOWREASONER faces the limitation of the need for training, which needs some computational cost. Besides, while some generalization capability is shown when switching worker agent, performance still varies based on the worker agent used, suggesting partial dependency of the trained meta-agent on the underlying execution agent.

D LLM USAGE

We used an OpenAI LLM (GPT-5) as a writing and formatting assistant. In particular, it helped refine grammar and phrasing, improve clarity, and suggest edits to figure/table captions and layout (e.g., column alignment, caption length, placement). The LLM did not contribute to research ideation, experimental design, implementation, data analysis, or technical content beyond surface-level edits. All outputs were reviewed and edited by the authors, who take full responsibility for the final text and visuals.

810

811

812

813

814

815

816

817

818

819

820

821

822

823

824

825

826

827

828

829

830

831

832

833

834

835

836

837

838

839

840

841

```

1 async def __call__(self, problem: str, test_cases: list):
2     # Extract key problem requirements and constraints
3     problem_analysis = await self.custom(
4         input=problem,
5         instruction=prompt_custom.GRAPH_ANALYZE_PROMPT
6     )
7
8     # Generate initial solution with problem analysis context
9     solution = await self.custom_code_generate(
10        problem=problem,
11        entry_point=entry_point,
12        instruction=f"{prompt_custom.GRAPH_GENERATE_PROMPT}\nProblem Analysis: {
13        problem_analysis['response']}"
14    )
15
16    current_solution = solution['response']
17
18    # Loop for max_iterations to improve the solution if tests fail
19    for iteration in range(3):
20        # Test current solution
21        test_results = await self.custom(
22            input=f"Solution: {current_solution}\nTest Cases: {test_cases}",
23            instruction=prompt_custom.GRAPH_TEST_PROMPT
24        )
25
26        # If tests pass, return the solution
27        if "failed" not in test_results['response'].lower():
28            return current_solution, self.llm.cost_manager.total_cost
29
30        # Improve solution with test feedback and problem analysis
31        improved_solution = await self.custom_code_generate(
32            problem=problem,
33            entry_point=entry_point,
34            instruction=f"{prompt_custom.GRAPH_IMPROVE_PROMPT}\nProblem Analysis: {
35            problem_analysis['response']}\nCurrent Solution: {current_solution}\nTest
36            Results: {test_results['response']}\nIteration: {iteration + 1}/{
37            max_iterations}"
38        )
39
40        # Update current solution
41        current_solution = improved_solution['response']
42
43        # Return the last solution attempt
44        return current_solution, self.llm.cost_manager.total_cost

```

842

843

844

845

Figure 6: Successful workflow generated by FLOWREASONER-14B for the BigCodeBench task of *Generating and plot weather data for a specified date range*.

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

```

1 async def __call__(self, problem: str, test_cases: list):
2     # Generate solution
3     solution = await self.custom_code_generate(problem=problem, entry_point=entry_point,
4         instruction=prompt_custom.GRAPH_GENERATE_PROMPT)
5
6     # Test solution
7     test_results = await self.custom(
8         input=f"Solution: {solution['response']}\nTest Cases: {test_cases}",
9         instruction=prompt_custom.GRAPH_TEST_PROMPT
10    )
11
12    # Optimize solution if tests fail
13    if "failed" in test_results['response'].lower():
14        improved_solution = await self.custom_code_generate(
15            problem=problem,
16            instruction=f"{prompt_custom.GRAPH_IMPROVE_PROMPT}\n\nTest Results: {test_results
17            ['response']}"
18        )
19
20        return improved_solution['response'], self.llm.cost_manager.total_cost
21
22    return solution['response'], self.llm.cost_manager.total_cost

```

Figure 7: Successful workflow generated by FLOWREASONER-14B for the MBPP task of *Writing a function to reverse words in a given string*.

```

864
865 1 async def __call__(self, problem: str, test_cases: list):
866 2     # 1. Prepare generation prompts
867 3     gen_prompts = [
868 4         prompt_custom.CODE_GENERATE_PROMPT_A,
869 5         prompt_custom.CODE_GENERATE_PROMPT_B,
870 6         prompt_custom.CODE_GENERATE_PROMPT_C,
871 7     ]
872 8
873 9     # 2. Generate three candidate solutions
874 10    sols = [
875 11        (await self.custom_code_generate(
876 12            problem=problem,
877 13            instruction=inst
878 14        ))['response'] for inst in gen_prompts
879 15    ]
880 16
881 17    # 3. Ensemble: pick or merge the best solution
882 18    ensemble_input = "Solutions:\n" + "\n---\n".join(sols)
883 19
884 20    best = (await self.custom(
885 21        input=ensemble_input,
886 22        instruction=prompt_custom.ENSEMBLE_PROMPT
887 23    ))['response']
888 24
889 25    # 4. Return the best solution and total cost
890 26    return best, self.llm.cost_manager.total_cost

```

Figure 8: Successful workflow generated by FLOWREASONER-14B for the HumanEval task of *Returning list of prime factors of given integer in the order from smallest to largest.*

```

888
889 1 async def __call__(self, problem: str, test_cases: list):
890 2     # 1. Analyze the problem
891 3     analysis = await self.custom(
892 4         input=problem,
893 5         instruction=prompt_custom.ANALYZE_PROMPT
894 6     )
895 7
896 8     # 2. Generate initial solution
897 9     solution = await self.custom_code_generate(
898 10        problem=problem,
899 11        instruction=prompt_custom.GENERATE_PROMPT,
900 12        context=analysis['response']
901 13    )
902 14
903 15    # 3. Test with a subset of test cases
904 16    sample_tests = test_cases[:min(3, len(test_cases))]
905 17    test_result = await self.custom(
906 18        input=f"Solution:\n{solution['response']}\nTests:\n{sample_tests}",
907 19        instruction=prompt_custom.TEST_PROMPT
908 20    )
909 21
910 22    # 4. Handle based on test results
911 23    if "all passed" in test_result['response'].lower():
912 24        # Optimize if all tests pass
913 25        final = await self.custom_code_generate(
914 26            problem=problem,
915 27            instruction=prompt_custom.OPTIMIZE_PROMPT,
916 28            context=f"Solution:\n{solution['response']}"
917 29        )
918 30        return final['response'], self.llm.cost_manager.total_cost
919 31    else:
920 32        # Fix issues if tests fail
921 33        fixed = await self.custom_code_generate(
922 34            problem=problem,
923 35            instruction=prompt_custom.FIX_PROMPT,
924 36            context=f"Solution:\n{solution['response']}\nTest results:\n{test_result['response']}"
925 37        )
926 38        return fixed['response'], self.llm.cost_manager.total_cost

```

Figure 9: Failed workflow generated by FLOWREASONER-14B for the MBPP task of *Writing a function to count those characters which have vowels as their neighbors in the given string.*

918

919

```

1  async def __call__(self, problem: str, test_cases: list):
2      async def __call__(self, problem: str, test_cases: list):
3          # Extract key problem requirements and constraints with detailed specifications
4          problem_analysis = await self.custom(
5              input=problem,
6              instruction=prompt_custom.GRAPH_ANALYZE_PROMPT
7          )
8
9          # Initial approach planning with appropriate algorithm selection
10         algorithm_selection = await self.custom(
11             input=f"Problem: {problem}\nAnalysis: {problem_analysis['response']}",
12             instruction=prompt_custom.GRAPH_ALGORITHM_SELECTION_PROMPT
13         )
14
15         # Generate solution with problem analysis and algorithm selection
16         solution = await self.custom_code_generate(
17             problem=problem,
18             entry_point=entry_point,
19             instruction=f"{prompt_custom.GRAPH_GENERATE_PROMPT}\nProblem Analysis: {
20                 problem_analysis['response']}\nAlgorithm Selection: {algorithm_selection['
21                 response']}"
22         )
23
24         # Self-review for logical errors and optimization opportunities
25         code_review = await self.custom(
26             input=f"Code: {solution['response']}",
27             instruction=prompt_custom.GRAPH_CODE_REVIEW_PROMPT
28         )
29
30         # Generate improved solution with comprehensive context
31         improved_solution = await self.custom_code_generate(
32             problem=problem,
33             entry_point=entry_point,
34             instruction=f"{prompt_custom.GRAPH_IMPROVE_PROMPT}\nProblem Analysis: {
35                 problem_analysis['response']}\nAlgorithm Selection: {algorithm_selection['
36                 response']}\nCode Review: {code_review['response']}"
37         )
38
39         # Add explanation to the solution
40         solution_explanation = await self.custom(
41             input=f"Solution: {improved_solution['response']}\nProblem: {problem}",
42             instruction=prompt_custom.GRAPH_SOLUTION_EXPLANATION_PROMPT
43         )
44
45         final_solution = f"# Solution\n{improved_solution['response']}\n\n# Explanation\n{
46             solution_explanation['response']}"
47         return final_solution, self.llm.cost_manager.total_cost

```

951

952

Figure 10: Failed workflow generated by FLOWREASONER-14B for the BigCodeBench task of *Extracting the text and href attributes of all anchor tags from a given URL's HTML content.*

953

954

955

956

957

```

1  async def __call__(self, problem: str, test_cases: list):
2      # 1. Generate solution
3      solution = (await self.custom_code_generate(
4          problem=problem,
5          instruction=prompt_custom.CODE_GENERATE_PROMPT
6      ))['response']
7
8      # 2. Test and refine solution
9      refined_solution = (await self.custom_code_generate(
10         problem=problem,
11         instruction=prompt_custom.REFINE_PROMPT,
12         context=f"Initial solution:\n{solution}\n\nTest cases:\n{test_cases}"
13     ))['response']
14
15     # 3. Return the solution and total cost
16     return refined_solution, self.llm.cost_manager.total_cost

```

968

969

Figure 11: Failed workflow generated by FLOWREASONER-14B for the HumanEval task of *Evaluating whether the given number n can be written as the sum of exactly 4 positive even numbers.*

970

971