
MermaidFlow: Redefining Agentic Workflow Generation via Safety-Constrained Evolutionary Programming

Chengqi Zheng¹ Jianda Chen² Yueming Lyu¹ Wen Zheng Terence Ng² Haopeng Zhang³ Yew-Soon Ong^{1,2}
Ivor Tsang^{1,2} Haiyan Yin¹

Abstract

Despite the promise of autonomous agentic reasoning, existing workflow generation methods frequently produce fragile, unexecutable plans due to unconstrained LLM-driven construction. We propose **MermaidFlow**, a framework that redefines the agentic search space through safety-constrained graph evolution. At its core, MermaidFlow represent workflows as a verifiable intermediate representation using Mermaid, a structured and human-interpretable graph language. We formulate domain-aware evolutionary operators, i.e., *crossover*, *mutation*, *insertion*, and *deletion*, to preserve semantic correctness while promoting structural diversity, enabling efficient exploration of a high-quality, statically verifiable workflow space. Without modifying task settings or evaluation protocols, MermaidFlow achieves consistent improvements in success rates and faster convergence to executable plans on the agent reasoning benchmark. The experimental results demonstrate that safety-constrained graph evolution offers a scalable, modular foundation for robust and interpretable autonomous decision-making systems.

1. Introduction

Large language models (LLMs) are increasingly instantiated as modular agents that collaborate to solve complex tasks through structured workflows (Guo et al., 2024; Li, 2025a;b). These agentic workflows decompose problems into subtasks, assign them to specialized agents, and inte-

*Equal contribution ¹CFAR and IHPC, Agency for Science, Technology and Research (A*STAR), Singapore ²Nanyang Technological University (NTU), Singapore ³University of Hawaii at Mānoa, USA. Correspondence to: Haiyan Yin <Yin_Haiyan@cfar.a-star.edu.sg>.

grate intermediate outputs toward a shared goal. Moving beyond single-agent prompting, this multi-agent setting requires coherent planning and execution across agents with distinct roles and responsibilities. Designing such workflows involves reasoning over compositional graph structures that represent inter-agent dependencies, data flow, and semantic constraints, forming the foundation for scalable and adaptive multi-agent systems (Zhou et al., 2025).

The lifecycle of agentic workflow is naturally structured into three layers: (1) *workflow planning*, which defines the structure of subtasks, agent roles, and information flow; (2) *code realization*, where the plan is translated into executable programs; and (3) *runtime execution*, where agents are instantiated and carry out their assigned behaviors. In many systems, these layers are collapsed (e.g., (Hu et al., 2024; Zhang et al., 2024c)): workflows are directly generated as *Python code* or serialized *JSON trees*, where planning decisions are entangled with implementation (i.e., through *prompting*-based generation of code or execution traces). As a result, workflows are often encoded in *low-level* formats where **structure** is implicit, **semantics** are entangled with imperative logic, and **validity** can only be assessed at runtime. This implicit representation hinders verifiability, reuse, and search, limiting the robustness and scalability of multi-agent systems.

Indeed, recent studies reveal that multi-agent LLM systems frequently fail due to brittle workflow logic and coordination breakdowns (Cemri et al., 2025; Zhang et al., 2024a; 2025c). These failures typically arise not from deficiencies in language models themselves but emerge from workflows that cannot be reasoned about, verified, or adapted. Without a structured representation of agent roles, task flow, and dependencies, systems struggle to detect errors before execution or to generalize behaviors across tasks. This points to a core limitation: existing workflows lack the abstraction needed for reliable planning.

To address the limitations of implicit, code-bound workflows, we introduce **MermaidFlow**, a declarative representation for agentic planning inspired by the **Mermaid**

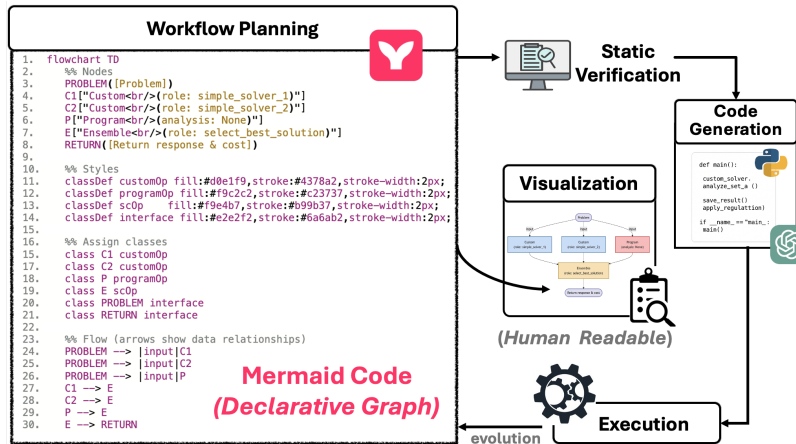


Figure 1. An illustration of the **workflow lifecycle in MermaidFlow**. The workflow is modeled as a declarative graph using Mermaid code, where nodes $\mathcal{V}_{[\tau,\alpha]}$ and edges $\mathcal{E}_{[\rho]}$ are explicitly defined with annotated prompts and roles (lines 3–8), styled and typed (lines 11–21), and connected via directed edges (lines 24–30). This results in a statically verifiable, semantically typed, and structurally interpretable representation that serves as a unified interface for visualization, validation, and code generation.

graph markup language¹. MermaidFlow defines workflows as **declarative graphs**, where nodes represent prompting agents and edges specify information flow (see Figure 1 for an illustrative example of the declarative graph encoded in Mermaid language, which is declarative, structurally explicit, and highly human-interpretable). This high-level representation enables structural and semantics properties with static verification, e.g., *role consistency*, and *type-safe connections*, can be enforced at the graph level. For instance, a simple workflow can be written as `input` \rightarrow `reasoner` \rightarrow `summarizer`, offering a clear plan that is both human-readable and programmatically analyzable. By exposing explicit semantics and structure, MermaidFlow yields substantial downstream benefits in both code generation and workflow evaluation, when LLMs are extensively employed to realize and evaluate workflows. These properties ultimately yield a more robust, verifiable, and high-reward search space for agentic workflow planning. MermaidFlow enforces a clear separation between symbolic planning and executable code, ensuring that workflow structures remain statically verifiable by design.

Building on this foundation, we further propose a novel **evolutionary programming** (EP) framework tailored specifically to explore MermaidFlow’s structured graph space. Our EP approach employs safety-constrained operations, including node replacement, subgraph rewiring, and role-consistent insertions, to maintain workflow correctness throughout the search process. Furthermore, historical workflows generated during search accumulate as structured experience, enabling efficient reuse and adaptation across tasks. Together, MermaidFlow’s declarative representation and EP search framework constitute a programmable and task-agnostic programming layer for agentic workflow generation, enabling efficient search with improved correctness, generalization, and adaptability. To our knowledge, this is the first agentic workflow framework to guarantee static graph-level correctness across the entire generation process.

In summary, our contributions are threefold: (1) We introduce MermaidFlow, a declarative, verifiable graph representation for agentic workflow planning that cleanly separates planning from execution; (2) We develop a novel EP-based search framework leveraging structured mutation operators and workflow experience accumulation; and (3) We empirically demonstrate that MermaidFlow significantly outperforms existing code-based methods on standard agentic reasoning benchmarks, improving success rates, search efficiency, and interpretability.

2. Related Works

Agentic Workflows with LLMs Recent advances in multi-agent LLM systems have enabled structured collaboration among specialized agents to tackle complex, multi-step tasks. AFLOW (Zhang et al., 2024c), MaAS (Zhang et al., 2025b), and MASS (Zhou et al., 2025) formalize agent workflows using execution graphs and message-passing protocols to model multi-step reasoning. MetaGPT (Hong et al., 2024) and MAS-GPT (Ye et al., 2025) implement role-based orchestration by assigning domain-specific functions (e.g., product manager, engineer) and encoding Standard Operating Procedures (SOPs) to reduce cascading errors. Debate-based frameworks such as Multi-Agent Debate (Liang et al., 2024; Du et al., 2024) and DebFlow (Su et al., 2025) introduce structured critique among agents to promote output reliability. While these systems improve modularity and scalability, their workflows are typically encoded in *imperative code* or *loosely structured prompts*, i.e., formats that lack semantic abstraction and resist verification. Recent studies (Cemri et al., 2025; Zhang et al., 2024a; 2025c) identify fragile workflows, rather than model errors, as the primary source of failure in multi-agent systems. Our proposed MermaidFlow addresses this bottleneck by introducing a typed, declarative workflow space that supports safe construction, static validation, and structured exploration, advancing agentic reasoning beyond brittle prompt chaining.

¹<https://mermaid.js.org>

Workflow Representation The representation of agentic workflows governs not only how agents are composed, but also whether they can be verified, reused, or optimized. Natural language-based prompting methods, such as Chain-of-Thought (Kojima et al., 2022), ReAct (Yao et al., 2023), and Self-Refine (Madaan et al., 2023), are expressive but underspecified, lacking formal structure for validation. In contrast, code-centric approaches like AFLOW (Zhang et al., 2024c), ADAS (Hu et al., 2024), and ScoreFlow (Wang et al., 2025) generate executable Python or JSON trees directly, offering precision at the cost of brittleness and poor editability due to tightly entangled logic and implementation. Recent efforts explore more structured workflow abstractions. GPTSwarm (Zhuge et al., 2024) and FlowReasoner (Gao et al., 2025) organize workflows as agent interaction graphs, but lack formal semantics, e.g., no type enforcement, role validation, or support for systematic search. MetaGPT (Hong et al., 2024) and MAS-GPT (Ye et al., 2025) encode workflows through SOP-style templates and DSLs, but rely on rigid decomposition patterns that restrict flexibility. MermaidFlow departs from these by introducing a typed, declarative graph representation grounded in the Mermaid markup language. It makes role semantics and data flow explicit, allowing crucial graph-level structural constraints, such as role consistency, and type safety, to be enforced pre-execution, enabling safe reuse, adaptation, and search.

Workflow Search and Optimization The structure of the search space fundamentally shapes how workflows are generated and optimized. AFlow (Zhang et al., 2024c) applies Monte Carlo Tree Search over executable graphs, while ADAS (Hu et al., 2024) explores code-level candidates via heuristics-guided expansion. Though systematic, both approaches operate over brittle code-centric representations, where small mutations often break correctness, necessitating expensive filtering. ScoreFlow (Wang et al., 2025) and G-Designer (Zhang et al., 2024b) adopt learned or continuous optimization strategies, adjusting prompt topologies or agent graphs via gradient-based tuning or neural controllers. However, these methods require differentiable feedback or training signals and offer limited support for enforcing structural validity. A complementary direction leverages evolutionary and population-based search. DebFlow (Su et al., 2025) refines workflows through iterative agent debates, while EvoFlow (Zhang et al., 2025a) evolves diverse workflows using task complexity-conditioned genetic search. Yet both operate in loosely defined or weakly constrained spaces, where mutations can produce semantically invalid workflows. MermaidFlow addresses this gap by defining a structured, verifiable graph space with safety-aware mutation operators. This design ensures that every candidate is valid by construction, enabling scalable and principled workflow optimization.

3. A Novel Declarative Graph Representation for Agentic Workflows

This section introduces a declarative graph representation for agentic workflows, built on **Mermaid** that is a structured, human-readable language with built-in static verifiability and graph render to help human directly observe the workflow. Departing from unstructured or token-level workflow representations, our workflow formalism leverages Mermaid’s type-aware syntax and explicit control-flow semantics to enable correctness by construction, symbolic manipulation, and modular workflow composition.

3.1. Declarative Workflow Graphs with Mermaid

We model each agentic workflow as a declarative computation graph with explicit typing, annotations, and semantic structure. Formally, we define a workflow graph as:

$$G(\mathcal{V}_{[\tau, \alpha]}, \mathcal{E}_{[\rho]}), \quad (1)$$

where $\mathcal{V}_{[\tau, \alpha]}$ is a set of typed and annotated nodes, $\mathcal{E}_{[\rho]}$ is a set of directed, role-labeled edges.

This structure is instantiated using the **Mermaid** graph language, a lightweight, human-readable syntax for specifying typed graphs with semantically annotated components. Furthermore, Mermaid provides a declarative interface that supports symbolic manipulation and static validation. A real example is illustrated in Figure 1. Each node defines a symbolic identifier and type signature, while edges carry semantic labels (e.g., `inputs`) that describe data-flow.

Following, we will define each component of the workflow in Mermaid.

Nodes. Each node $v \in \mathcal{V}_{[\tau, \alpha]}$ is a tuple $(id, \tau(v), \alpha(v))$, where `id` is a symbolic identifier, $\tau(v) = T_{in} \rightarrow T_{out}$ denotes the type of that node, and attribute $\alpha(v)$ would provide some necessary information according to the type. Nodes represent typed declarative units that are interpretable and statically verifiable, and it can be easily understood by human.

Edges. Each edge $(u, v) \in \mathcal{E}_{[\rho]}$ denotes a dependency annotated with a semantic label $\rho(u, v)$, indicating how information or control flows from u to v e.g., “input”, “problem”. Mermaid syntax supports these semantics with labeled arrows (e.g., `A -> |input| B`).

Types. All graph nodes are explicitly typed and semantically annotated, with types governing interface compatibility and ensuring valid workflow construction. By defining these types up front, we guarantee a consistent translation from Mermaid diagrams to Python code. For each task domain, we introduce dedicated node types for the operators and tools that have proven most effective. Some operators would need to assign parameters when using in Python code,

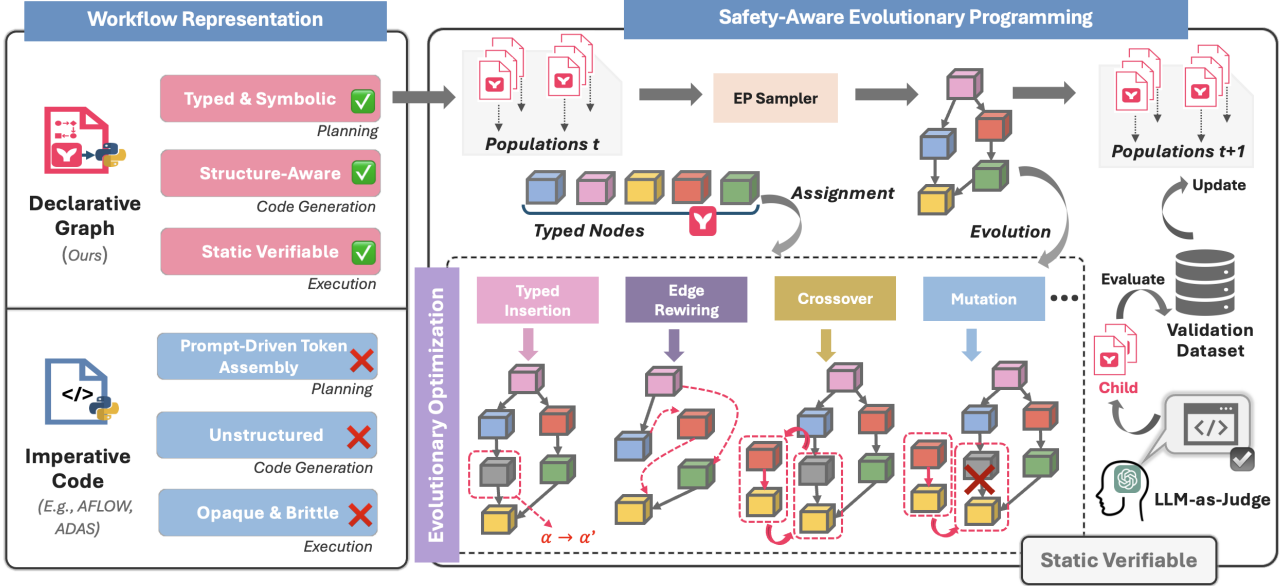


Figure 2. Overview of the MermaidFlow framework. Left: Comparison between imperative (Python-based) and declarative (Mermaid-based) workflow representations. MermaidFlow models workflows as statically typed, verifiable graphs, enabling interpretable planning and structure-aware code generation. Right: Illustration of the safety-aware evolutionary programming process. Given historical Mermaid workflows, the EP sampler selects parent candidates and applies EP operators. Resulting workflows are evaluated by the LLM-as-Judge to update the workflow population.

which is also cleared in the Mermaid field, for example, a CustomOp node may include a role attribute, while a programOp node can carry an analysis attribute. During code generation, these attributes direct the translator to emit the correct Python calls, making sure that all tool arguments are clearly specified. We provide a detailed type description in Appendix A.1.

Each node explicitly defines its symbolic identity, type signature, and related attribution, while edges govern execution flow subject to type constraints. Unlike flat or post-validated node-link DAGs (e.g., JSON plans or token-generated programs), our declarative graph formalism introduces an abstraction layer that bridges symbolic reasoning and execution-level safety. This structure supports (static) correctness-preserving mutation and compositional reuse, which are the key properties we exploit in the constrained optimization process detailed next. To the best of our knowledge, this is the first agentic workflow representation that leverages a graph-oriented abstract coding language to enable more natural graph definition and manipulation. In the next section, we will formalize graph-manipulation actions in Mermaid and present a workflow-optimization method, further illustrating the advantages of using Mermaid for workflow representation.

3.2. Agentic Workflow Search Space

The declarative graph formalism introduced above induces a constrained search space over agentic workflows. We

denote this space as:

$$\mathcal{S} = \{G(\mathcal{V}_{[\tau, \alpha]}, \mathcal{E}_{[\rho]}, \mathcal{C}) \in \mathcal{G}_{\text{Mermaid}} \mid G \models \mathcal{C}_{\text{static}}\}, \quad (2)$$

where $\mathcal{G}_{\text{Mermaid}}$ denotes the set of workflows expressible in the Mermaid graph language, and $\mathcal{C}_{\text{static}}$ captures structural constraints, such as type compatibility, role-consistent edges, and connectivity, automatically enforced by Mermaid’s parser and extended structural schema. This built-in verifiability arises from Mermaid’s declarative syntax and ensures that all elements of \mathcal{S} are valid and executable by construction.

To enable optimization, we parameterize each node $v \in \mathcal{V}_{[\tau, \alpha]}$ as a tuple $v = (m, p(\tau, \alpha), f(\tau))$, following conventions from multi-agent systems (MAS) definition, where $m \in M$ specifies the LLM configuration (e.g., model_name, temperature), $p(\tau, \alpha) \in P$ is the prompt template determined by the node type τ and its argument α , and $f(\tau) \in F$ denotes the input/output format associated with type τ .

$$\mathcal{V}_{[\tau, \alpha]} = \{(m, p(\tau, \alpha), f(\tau)) \mid m \in M, p \in P, f \in F\}. \quad (3)$$

The formula above demonstrates that, when interpreting a Mermaid workflow, each node can be directly mapped to a standard LLM agent instance. By assigning a type τ and parameter/attribute α , one can associate the node with a specific prompt $p(\tau, \alpha)$ and input/output format $f(\tau)$. **This formulation emphasizes that every LLM agent can**

be consistently defined both within the Mermaid representation and in the general context of LLM agent configuration.

The declarative formulation of \mathcal{S} defines a graph space that is statically typed, compositionally valid, and closed under structure-preserving rewrites. Each graph $G \in \mathcal{S}$ admits local edits, such as node substitution, edge relabeling, or typed insertion, without violating global correctness, due to the compositional enforcement of type and role constraints. Crucially, these transformations are decidable and stateless under the constraint that mutations operate on type-consistent boundaries, requiring no forward propagation or re-evaluation, which enables efficient mutation during evolutionary search.

The space is also **inductively closed**: type-compatible subgraphs can be composed without revalidation, while disconnected or cyclic fragments are excluded by construction. This closure property is non-trivial: prior workflow representations, especially those based on imperative or token-level programs, lack structural guarantees, and mutations frequently yield invalid or unrecoverable states. In contrast, our declarative design ensures that local edits remain within the valid region of \mathcal{S} , enabling reliable and efficient search. Though intentionally bounded, \mathcal{S} remains expressive enough to capture planning motifs such as hierarchical refinement and dataflow composition. By unifying generation, mutation, and verification within a compiler-verifiable substrate, it provides a semantically grounded foundation for structure-aware and safe workflow optimization.

4. Constraint-Aware Evolutionary Workflow Optimization

We introduce an evolutionary programming (EP) framework that operates directly over declarative Mermaid workflows. Leveraging Mermaid’s typed and verifiable graph structure, we define correctness-preserving operators that enable safe, modular workflow evolution. Unlike prior approaches over unstructured or token-based spaces, all candidates in MermaidFlow are valid by construction, ensuring safe, compiler-checkable optimization throughout the search process.

4.1. Constraint-Preserving EP Operators for Declarative Workflow Graphs

We define a set of atomic graph-level operators that drive workflow evolution within MermaidFlow. Each operator acts over a candidate graph $G(\mathcal{V}, \mathcal{E}) \in \mathcal{G}_{\text{Mermaid}} \subseteq \mathcal{S}$, and is designed to be *locally scoped*, *type-consistent*, and *statically verifiable*, enabling every candidate to be validated by the Mermaid compiler at each step of the search. Below are the definitions of the operations, which will be used to verify the correctness of the newly generated workflow.

Node Substitution. Changing the attributes of a specific agent $v(\tau, \alpha) \in \mathcal{V}$ to $v(\tau, \alpha')$. Like changing the corresponding role prompt or instruction.

Node Addition. Given an edge $(v_a, v_b) \in \mathcal{E}$, connecting from node v_a to node v_b , insert a new node v' to form (v_a, v') and (v', v_b) and disconnect (v_a, v_b) if: $T_{\text{out}}(v_a) = T_{\text{in}}(v')$, $T_{\text{out}}(v') = T_{\text{in}}(v_b)$ according to node type τ .

Edge Rewiring. Given nodes $\{v_a, v_b, v_c\} \subseteq \mathcal{V}$ and $(v_a, v_b) \in \mathcal{E}$ in the original graph G , rewire to (v_a, v_c) or (v_c, v_b) and disconnect (v_a, v_b) if: $T_{\text{out}}(v_a) = T_{\text{in}}(v_c)$ or $T_{\text{out}}(v_c) = T_{\text{in}}(v_b)$.

Node Deletion. Given a linear path $v_a \rightarrow v_b \rightarrow v_c$, delete v_b and insert an edge (v_a, v_c) if $T_{\text{out}}(v_a) = T_{\text{in}}(v_c)$.

Subgraph Mutation. Let $G_1 \in \mathcal{G}_{\text{Mermaid}}$ be a subgraph of the graph $G \in \mathcal{S}$. Denote the input and output node set of G_1 as I_1 and O_1 , respectively. Let $G_2 \in \mathcal{G}_{\text{Mermaid}}$ be a feasible graph with input and output node set I_2 and O_2 . Replace G_1 in G with G_2 such that $T_{\text{in}}(I_1) = T_{\text{in}}(I_2)$ and $T_{\text{out}}(O_1) = T_{\text{out}}(O_2)$.

Crossover. Given $\{G_1, G_2\} \subseteq \mathcal{S}$ share a common interface node v (e.g., an ensemble node), swap subgraphs rooted at v to yield $\{G'_1, G'_2\}$ such that the type and interface constraints are preserved, i.e., $\{G'_1, G'_2\} \subseteq \mathcal{S}$. Each operator is applied at the level of Mermaid syntax, enabling compiler-level validation of every candidate graph. By constraining transformations to preserve type and role integrity, MermaidFlow ensures that evolutionary search remains within the semantically valid subspace of workflows.

Lemma 4.1 (MermaidFlow Transformation Invariance). *Let \mathcal{S} denote the declarative workflow space defined in Section 3.2. For any workflow graph $G \in \mathcal{S}$ and any atomic transformation operator \mathcal{O} defined above, the resulting graph $G' = \mathcal{O}(G)$ also belongs to \mathcal{S} :*

$$\forall G \in \mathcal{S}, \forall \mathcal{O} \in \mathbb{O}, \quad \mathcal{O}(G) \in \mathcal{S} \quad (4)$$

where \mathbb{O} is the set of constraint-preserving operators over MermaidFlow graphs. That is, \mathcal{S} is closed under all valid EP operations.

Definition 4.2. Let \mathcal{G} denote the space of all candidate workflows, each $G \in \mathcal{G}$ represented as a directed graph $(\mathcal{V}, \mathcal{E})$. We define a static validator function $Q : \mathcal{G} \rightarrow \{0, 1\}$, implemented by a Mermaid parser/compiler, such that:

$$Q(G) = \begin{cases} 1 & \text{if } G \in \mathcal{S} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

Here, $\mathcal{S} \subset \mathcal{G}$ is the subset of workflows satisfying verifiability constraints such as workflow structure, well-typed I/O, role validity, and full connectivity.

By using EP operators above, from Lemma 4.1, given a $G_t \in \mathcal{S}$, each change $\mathcal{O}(G_t)$ at step t leads to a graph

$G_{t+1} = \mathcal{O}(G_t) \in \mathcal{S}$. Given an initial graph $G_0 \in \mathcal{S}$, by induction, we know $\forall t \in \mathbb{N}_+$, $G_{t+1} = \mathcal{O}_t \circ \mathcal{O}_{t-1} \cdots \circ \mathcal{O}_0(G_0) \in \mathcal{S}$. Thus, the evolution in the static Mermaid graph space remains the safe subspace.

In MermaidFlow, when using an LLM to generate a new Mermaid graph, the resulting Mermaid code may sometimes violate predefined safety constraints. To address this, we implement a checker to verify whether the newly generated candidates conform to the defined workflow and operation rules. If any violations are detected, new workflows are regenerated. Thanks to Mermaid’s simple and clear syntax, the code can be treated as structured text. This allows us to easily build a text-based analysis tool and incorporate custom rules into the checker. More implementation details can be found in Appendix A.2.

4.2. Evaluation and Selection in Workflow Populations

We frame each declarative workflow graph as an *experience* and maintain a population of scored experiences over time. At each optimization step t , the system tracks a history buffer: $W_{\text{history},t} = \{(s_i, \text{score}_i)\}_{i=1}^t$, where $s_i \in \mathcal{G}_{\text{Mermaid}}$ denotes a structurally valid workflow, and score_i reflects its estimated performance.

At each optimization cycle, two parent workflows s_a, s_b are sampled from $W_{\text{history},t}$, typically via temperature-scaled softmax sampling according to following distribution: $P_{\text{mixed}}(i) = \lambda \cdot \frac{1}{t} + (1 - \lambda) \cdot \frac{\exp(\alpha \cdot \text{score}_i)}{\sum_{j=1}^n \exp(\alpha \cdot \text{score}_j)}$, where t is the number of workflows in the history buffer, score_i is the validation score of the i -th workflow, and the parameters α and λ control the influence of the scores, and balances exploration-exploitation, respectively. After sample two different parent workflows s_a, s_b where $s_a \neq s_b$. These are used to generate a candidate set through the evolutionary process:

$$S_{\text{candidates}} = \{s_i \mid s_i = \mathcal{O}(s_a, s_b), \mathcal{O} \in \mathbb{O}\}_{i=1}^N, \quad (6)$$

where \mathbb{O} denotes the set of correctness-preserving operators (Section 4.1), for some operator only s_a involved and N is the candidate pool size.

To avoid expensive rollout-based evaluation over the full population, we adopt an *LLM-as-judge* model that scores each candidate $s \in S_{\text{candidates}}$ based on semantic fit, structure, and task relevance. Since all candidates in $S_{\text{candidates}}$ are statically verified by the Mermaid compiler, they are guaranteed to be syntactically valid, type-safe, and structurally executable, dramatically reducing failure cases and increasing effective sample quality.

We then select the highest-scoring candidate and update the history buffer:

$$W_{\text{history},t+1} \leftarrow W_{\text{history},t} \cup \{(s_{\text{child}}^*, \text{Validate}(s_{\text{child}}^*))\}, \quad (7)$$

where $s_{\text{child}}^* = \arg \max_{s \in S_{\text{candidates}}} \text{LLM_as_Judge}(s)$. This experience-centric design, enabled by the declarative and verifiable structure of MermaidFlow, supports efficient, low-cost population evolution without compromising safety, correctness, or search quality. See Appendix A.3 for algorithmic details.

5. Experiments

5.1. Experiment Setup

Baseline. We choose threefold of agentic baselines: (1) **Single-agent execution methods**, including CoT (Kojima et al., 2022), ComplexCoT (Fu et al., 2023), and Self-Consistency (Wang et al., 2023). (2) **Hand-crafted multi-agent systems**, such as LLM-Debate (Du et al., 2024), LLM-Blender (Jiang et al., 2023), DyLAN (Liu et al., 2024), and MAcNet (Qian et al., 2024). (3) **Autonomous multi-agent systems**, including GPTSwarm (Zhuge et al., 2024), MaAS (Zhang et al., 2025b), AutoAgents (Chen et al., 2024), ADAS (Hu et al., 2024), and AFlow (Zhang et al., 2024c). Among them, GPTSwarm and MaAS incorporate trainable modules for assigning workflow structures, while AutoAgents, ADAS, and AFlow rely on an LLM to design the structure, consistent with our setting. More details on baseline setups are provided in Appendix A.4.

Task and Benchmarks. We evaluate MermaidFlow on four public benchmarks covering two domains: (1) **math reasoning**, GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021); (2) **code generation**, HumanEval (Chen et al., 2021), and MBPP (Austin et al., 2021). For MATH benchmark, we follow AFlow (Zhang et al., 2024c) and MaAS (Zhang et al., 2025b) in using the same selected problems from four typical problem types in level 5. The dataset statistics are provided in Appendix A.5.

Implementation details. We use a closed-source LLM (gpt-4o-mini-0718) as both the Optimization and Execution LLM, consistent with the setup in MaAS (Zhang et al., 2025b). The Optimization LLM is responsible for tasks such as generating promising workflows in Mermaid code, selecting from sampled workflows, evolving to new workflows, and translating Mermaid code into Python code. All models are accessed via API with the temperature set to 0. In each round, we generate four different s_{child} candidates. To ensure experimental stability, complex operations such as crossover are applied with only a 10% probability. We set the number of iteration rounds to 20 for both Mermaid and AFlow, and to 30 for ADAS. The evaluation metrics are kept consistent with those used in AFlow and MaAS: for GSM8K and MATH, we report the Solve Rate (%) as the primary metric, while for HumanEval and MBPP, we report the pass@1 score.

Table 1. Performance comparison with single-agent, hand-crafted multi-agent systems, and automated agentic workflows. All methods use `gpt-4o-mini` as the base LLM and are evaluated on the test split, with results averaged over three runs. **Bold** indicates the best result; underline denotes the runner-up. MermaidFlow demonstrates consistent improvement across all datasets.

Method	GSM8K	MATH	HumanEval	MBPP	Avg.
Vanilla	87.57	46.29	87.49	70.29	72.91
CoT (Kojima et al., 2022)	87.45	46.40	88.13	71.83	73.45
ComplexCoT (Fu et al., 2023)	86.89	46.40	87.49	72.36	73.29
SC (CoT×5) (Wang et al., 2023)	87.57	47.91	88.60	73.60	74.42
LLM-Debate (Du et al., 2024)	89.47	48.63	88.80	70.29	74.30
LLM-Blender (Jiang et al., 2023)	88.35	46.92	88.68	77.05	75.25
DyLAN (Liu et al., 2024)	89.98	48.54	90.42	77.30	76.56
MacNet (Qian et al., 2024)	87.95	45.18	84.57	65.28	70.75
GPTswarm (Zhuge et al., 2024)	89.14	47.88	89.32	77.43	75.94
MaAS (Zhang et al., 2025b)	<u>91.47</u>	52.19	<u>91.57</u>	<u>82.17</u>	<u>79.35</u>
AutoAgents (Chen et al., 2024)	87.69	45.32	87.64	71.95	73.15
ADAS (Hu et al., 2024)	88.35	43.18	84.19	77.05	73.69
AFlow (Zhang et al., 2024c)	90.11	<u>52.81</u>	90.08	81.67	78.67
MermaidFlow (Ours)	92.39	55.42	92.87	82.31	80.75

5.2. Experimental Results

We compare MermaidFlow against 13 baselines on the GSM8K, MATH, HumanEval, and MBPP benchmarks, as shown in Table 1. The results demonstrate that MermaidFlow consistently achieves the best performance across all tasks. Compared to methods that search for the next workflow in the Python field, such as ADAS and AFlow, our approach outperforms them by an average margin of 2.08% to 5.54%. On the MATH benchmark specifically, MermaidFlow exceeds the second-best method AFlow by 2.61%. For certain benchmarks, performance is primarily limited by the capabilities of the Execution LLM. For example, in HumanEval and GSM8K, the baseline (Vanilla) performance is already high, so improvements from architectural optimization are less significant. In contrast, for benchmarks where the baseline performance is relatively low, such as MATH and MBPP, our method demonstrates a more substantial impact. Overall, MermaidFlow’s average score across these tasks is 80.75%, which is 1.40% higher than the highest average among all baselines (79.35% by MaAS), fully demonstrating the robustness and superiority of our approach across different problems.

5.3. Ablation Study

Evolution Efficiency To evaluate the effectiveness of our approach, we compare the learning curves of MermaidFlow and AFlow on the MATH dataset, as shown in Figure 5.3. MermaidFlow demonstrates a more consistent improvement in workflow quality during training and better generalization to the test set.

The core difference between MermaidFlow and AFlow lies in the search space. AFlow operates directly on Python code, applying textual edits with prompts constraints. This approach often leads to invalid and non-functional programs, with only a 50% success rate in generating executable code. In contrast, MermaidFlow evolves workflows at the graph level using Mermaid, a domain-specific language that enables structured manipulation (e.g., adding, deleting, or mutating nodes). This representation is better suited for LLM-based optimization (e.g., `gpt-4o-mini`), consistently yielding > 90% success rate in producing valid Python code. This reliability enables more effective exploration and optimization of workflow space.

Table 2. Comparison of different optimization LLMs on HumanEval and GSM8K datasets.

Dataset	Claude 3.5	GPT-4o	GPT-4o-mini
HumanEval	93.13	94.66	92.87
GSM8K	93.83	93.94	92.39

Impact of Optimization LLM Scale We investigate how the choice of Optimization LLM influences the quality of workflows in MermaidFlow by evaluating more capable models on the HumanEval and GSM8K benchmark. Specifically, we compare the effectiveness of larger models, such as Claude 3.5 and GPT-4o, in generating new workflows, while keeping GPT-4o-mini fixed as the Execution LLM. Results are summarized in Table 2. As a result, higher-capacity optimization LLMs consistently yield better performance across both HumanEval and GSM8K. This consistent trend underscores a core strength of Mermaid-

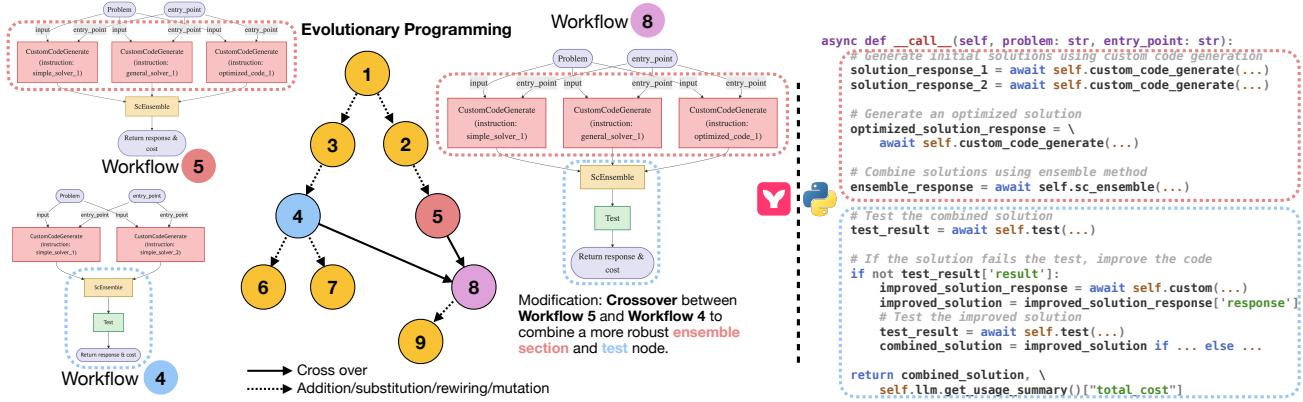


Figure 3. A case study on the HumanEval dataset showcasing how MermaidFlow evolves structured agentic workflows through evolutionary programming (with a detailed example of the crossover operator). The declarative graph representation also enables reliable translation of workflow graphs into executable python code (zoom-in view recommended).

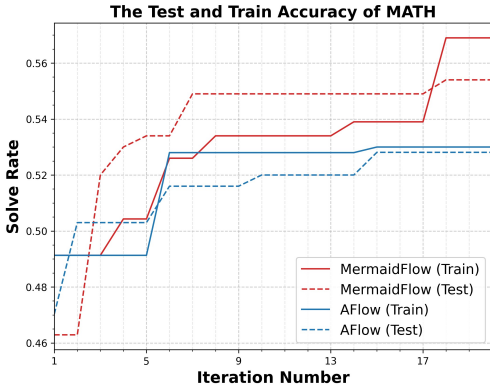


Figure 4. An illustrative figure comparing the highest solve rates on the MATH dataset between MermaidFlow and AFlow on the training set (119 problems) and test set (486 problems) across optimization iterations.

Flow: its well-structured, statically verifiable search space enables even modest improvements in optimization quality to translate directly into more functional, high-reward workflows.

5.4. Case Study: Workflow Evolution with MermaidFlow

In this case study, we present an example of how to generate a new workflow from given parent workflows using the Mermaid representation. During the evolutionary process, a new workflow can be derived from either a single parent workflow or two parent workflows, depending on the type of update operator applied. In the specific example shown in Figure 3, which is based on solving problems from the HumanEval benchmark, Workflow_8 is generated based on Workflow_4 and Workflow_5. Each parent contributes distinct advantages: Workflow_4 is the first to introduce a test node, while Workflow_5 contains a more diverse ensemble section with agents covering different rea-

soning aspects. MermaidFlow combines these strengths to synthesize a new and improved Workflow_8.

This generation process occurs in the Mermaid Field, where all workflows are defined in a structured syntax that can be directly rendered as visual diagrams. Once a new Mermaid workflow is generated, we use gpt-4o-mini to translate the Mermaid code into executable Python code. Due to Mermaid’s well-structured nature, this translation can be both straightforward and reliable. As demonstrated in Figure 3, the generated Python code perfectly resemble Mermaid Workflow_8, consisting of a diverse ensemble section and a test function.

This case study not only demonstrates the efficiency of searching for new high-quality workflow populations in the Mermaid field but also provides a detailed illustration of MermaidFlow’s stable and composable workflow lifecycle.

6. Conclusion

We propose MermaidFlow, a framework that transforms agentic workflow generation by encoding workflows as statically typed, semantically annotated, and compiler-verifiable graphs using the Mermaid language. Our proposed workflow formulation defines a well-structured, declaratively defined search space that supports safety-constrained rewrites and modular composition. Building on this space, we develop a safety-constrained evolutionary programming framework that enables efficient, verifiable, and high-quality workflow synthesis. MermaidFlow offers a principled step toward structurally safer and more interpretable agentic systems, introducing the first workflow optimization framework built atop a statically verifiable workflow representation. While MermaidFlow is evaluated in controlled agentic reasoning settings, its integration with real-world multi-agent systems and user-in-the-loop workflows introduces nuances that merit further exploration.

Acknowledgments

This research is supported by the National Research Foundation, Singapore and Infocomm Media Development Authority under its Trust Tech Funding Initiative, Career Development Fund (CDF) of the Agency for Science, Technology and Research (A*STAR) (No: C233312007, No: C243512014), and the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG-NMLP-2024-003). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of the National Research Foundation, Singapore, and Infocomm Media Development Authority.

References

- Austin, J., Odena, A., Nye, M., Bosma, M., Michalewski, H., Dohan, D., Jiang, E., Cai, C., Terry, M., Le, Q., et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- Cemri, M., Pan, M. Z., Yang, S., Agrawal, L. A., Chopra, B., Tiwari, R., Keutzer, K., Parameswaran, A. G., Klein, D., Ramchandran, K., Zaharia, M., Gonzalez, J. E., and Stoica, I. Why do multi-agent LLM systems fail? *CoRR*, abs/2503.13657, 2025. URL <https://doi.org/10.48550/arXiv.2503.13657>.
- Chen, G., Dong, S., Shu, Y., Zhang, G., Sesay, J., Karlsson, B. F., Fu, J., and Shi, Y. Autoagents: A framework for automatic agent generation. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, 2024*, pp. 22–30, 2024.
- Chen, M., Tworek, J., Jun, H., Yuan, Q., Pinto, H. P. D. O., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano, R., et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Du, Y., Li, S., Torralba, A., Tenenbaum, J. B., and Mordatch, I. Improving factuality and reasoning in language models through multiagent debate. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=zj7YuTE4t8>.
- Fu, Y., Peng, H., Sabharwal, A., Clark, P., and Khot, T. Complexity-based prompting for multi-step reasoning. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, 2023*. OpenReview.net, 2023.
- Gao, H., Liu, Y., He, Y., Dou, L., Du, C., Deng, Z., Hooi, B., Lin, M., and Pang, T. Flowreasoner: Reinforcing query-level meta-agents. *arXiv preprint arXiv:2504.15257*, 2025.
- Guo, T., Chen, X., Wang, Y., Chang, R., Pei, S., Chawla, N. V., Wiest, O., and Zhang, X. Large language model based multi-agents: A survey of progress and challenges. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, 2024*, pp. 8048–8057. ijcai.org, 2024.
- 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. *arXiv preprint arXiv:2103.03874*, 2021.
- Hong, S., Zhuge, M., Chen, J., Zheng, X., Cheng, Y., Wang, J., Zhang, C., 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 *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=VtmBAGCN7o>.
- Hu, S., Lu, C., and Clune, J. Automated design of agentic systems. *CoRR*, abs/2408.08435, 2024. URL <https://doi.org/10.48550/arXiv.2408.08435>.
- Jiang, D., Ren, X., and Lin, B. Y. Llm-blender: Ensembling large language models with pairwise comparison and generative fusion. In *Proceedings of the 61th Annual Meeting of the Association for Computational Linguistics (ACL 2023)*, 2023.
- Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., and Iwasawa, Y. Large language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, 2022*, 2022.
- Li, X. A review of prominent paradigms for llm-based agents: Tool use, planning (including rag), and feedback learning. In *Proceedings of the 31st International Conference on Computational Linguistics, COLING 2025, Abu Dhabi, UAE, 2025*, pp. 9760–9779, 2025a.
- Li, X. A review of prominent paradigms for llm-based agents: Tool use, planning (including rag), and feedback learning. In *Proceedings of the 31st International Conference on Computational Linguistics, COLING 2025, Abu Dhabi, UAE, 2025*, pp. 9760–9779, 2025b.

- Liang, T., He, Z., Jiao, W., Wang, X., Wang, Y., Wang, R., Yang, Y., Shi, S., and Tu, Z. Encouraging divergent thinking in large language models through multi-agent debate. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, 2024*, pp. 17889–17904. Association for Computational Linguistics, 2024. URL <https://aclanthology.org/2024.emnlp-main.992>.
- Liu, Z., Zhang, Y., Li, P., Liu, Y., and Yang, D. A dynamic llm-powered agent network for task-oriented agent collaboration. In *First Conference on Language Modeling, 2024*.
- Madaan, A., Tandon, N., Gupta, P., Hallinan, S., Gao, L., Wiegrefe, S., Alon, U., Dziri, N., Prabhumoye, S., Yang, Y., Gupta, S., Majumder, B. P., Hermann, K., Welleck, S., Yazdanbakhsh, A., and Clark, P. Self-refine: Iterative refinement with self-feedback. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, 2023*, 2023.
- Qian, C., Xie, Z., Wang, Y., Liu, W., Dang, Y., Du, Z., Chen, W., Yang, C., Liu, Z., and Sun, M. Scaling large-language-model-based multi-agent collaboration. *CoRR*, abs/2406.07155, 2024.
- Su, J., Xia, Y., Shi, R., Wang, J., Huang, J., Wang, Y., Shi, T., Yang, J., and He, L. Debflow: Automating agent creation via agent debate. *CoRR*, abs/2503.23781, 2025. URL <https://doi.org/10.48550/arXiv.2503.23781>.
- Wang, X., Wei, J., Schuurmans, D., Le, Q. V., Chi, E. H., Narang, S., Chowdhery, A., and Zhou, D. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, 2023*. OpenReview.net, 2023. URL <https://openreview.net/forum?id=1PL1NIMMrw>.
- Wang, Y., Yang, L., Li, G., Wang, M., and Aragam, B. Scoreflow: Mastering LLM agent workflows via score-based preference optimization. *CoRR*, abs/2502.04306, 2025. URL <https://doi.org/10.48550/arXiv.2502.04306>.
- Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K. R., and Cao, Y. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, 2023*, 2023. URL https://openreview.net/forum?id=WE_vluYUL-X.
- Ye, R., Tang, S., Ge, R., Du, Y., Yin, Z., Chen, S., and Shao, J. MAS-GPT: training llms to build llm-based multi-agent systems. *CoRR*, abs/2503.03686, 2025. URL <https://doi.org/10.48550/arXiv.2503.03686>.
- Zhang, B., Tan, Y., Shen, Y., Salem, A., Backes, M., Zannettou, S., and Zhang, Y. Breaking agents: Compromising autonomous LLM agents through malfunction amplification. *CoRR*, abs/2407.20859, 2024a. URL <https://doi.org/10.48550/arXiv.2407.20859>.
- Zhang, G., Yue, Y., Sun, X., Wan, G., Yu, M., Fang, J., Wang, K., and Cheng, D. G-designer: Architecting multi-agent communication topologies via graph neural networks. *CoRR*, abs/2410.11782, 2024b. URL <https://doi.org/10.48550/arXiv.2410.11782>.
- Zhang, G., Chen, K., Wan, G., Chang, H., Cheng, H., Wang, K., Hu, S., and Bai, L. Evoflow: Evolving diverse agentic workflows on the fly. *CoRR*, abs/2502.07373, 2025a. URL <https://doi.org/10.48550/arXiv.2502.07373>.
- Zhang, G., Niu, L., Fang, J., Wang, K., Bai, L., and Wang, X. Multi-agent architecture search via agentic supernet. *CoRR*, abs/2502.04180, 2025b. URL <https://doi.org/10.48550/arXiv.2502.04180>.
- Zhang, J., Xiang, J., Yu, Z., Teng, F., Chen, X., Chen, J., Zhuge, M., Cheng, X., Hong, S., Wang, J., Zheng, B., Liu, B., Luo, Y., and Wu, C. Aflow: Automating agentic workflow generation. *CoRR*, abs/2410.10762, 2024c. doi: 10.48550/ARXIV.2410.10762. URL <https://doi.org/10.48550/arXiv.2410.10762>.
- Zhang, S., Yin, M., Zhang, J., Liu, J., Han, Z., Zhang, J., Li, B., Wang, C., Wang, H., Chen, Y., et al. Which agent causes task failures and when? on automated failure attribution of llm multi-agent systems. *arXiv preprint arXiv:2505.00212*, 2025c.
- Zhou, H., Wan, X., Sun, R., Palangi, H., Iqbal, S., Vulic, I., Korhonen, A., and Arik, S. Ö. Multi-agent design: Optimizing agents with better prompts and topologies. *CoRR*, abs/2502.02533, 2025. URL <https://doi.org/10.48550/arXiv.2502.02533>.
- Zhuge, M., Wang, W., Kirsch, L., Faccio, F., Khizbullin, D., and Schmidhuber, J. Gptswarm: Language agents as optimizable graphs. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=uTC9AFXIhg>.

A. Implementation details

A.1. Details of Mermaid

Mermaid is a text-based language for creating diagrams and flowcharts, which we use to represent workflows. In our system, we extend Mermaid with custom node types:

1. **CustomOp** - Used for specialized reasoning strategies, executing specific tasks through defined roles. Example: `K["Custom
(role: validate_1)"]`, styled with blue fill.
2. **Interface** - Entry and exit nodes for workflows, including problem input points and return output points. Example: `PROBLEM([Problem])`, styled with light purple fill.
3. **ProgrammerOp** - Nodes that execute code generation and computation, particularly suitable for mathematical problems. Example: `P["Programmer
(analysis: 'Calculate step by step')"]`, styled with red fill. Only math problems use this node type.
4. **ScEnsembleOp** - Nodes that combine multiple solutions into one result, requiring multiple inputs to function properly. Example: `ENSEMBLE["ScEnsemble
"]`, styled with yellow fill.
5. **TestOp** - Nodes for validating solutions, used to verify whether code passes predefined test cases. Example: `T["Test
"]`, styled with green fill. Only code problems use this node type.
6. **CustomCodeGenerateOp** - Nodes specifically designed to generate code based on problem descriptions. Example: `CODE_GEN["CustomCodeGenerate
(instruction: xxx)"]`, styled with red fill. Only code problems use this node type.

Each node type has specific connection rules and usage patterns to ensure data flows correctly through the workflow. Connections between nodes are typically represented by arrows, for example: `PROBLEM -> |problem| P`, indicating that the problem input is passed to the programmer node.

We provide detailed descriptions of each node type in our prompts to guide the LLM in generating structurally valid Mermaid workflows. Additionally, we implement a Mermaid checker to verify the correctness of the generated Mermaid code.

For different types of problems (math or code), we use different prompt templates to guide the LLM in generating appropriate Mermaid code, as each problem type involves different node types and workflow patterns.

For math problems, our templates emphasize clarity in computational steps and rigor in mathematical reasoning. For example, we encourage using `Programmer` nodes for precise calculations and `ScEnsemble` nodes to combine multiple solutions for improved accuracy. The prompt templates include guidelines for LaTeX mathematical expression formatting to ensure outputs conform to mathematical standards.

For code problems, our templates focus on code generation, testing, and validation. We particularly emphasize using `Test` nodes to validate generated code solutions and `CustomCodeGenerate` nodes to produce code based on problem descriptions. The prompt templates include guidelines for code formatting and testing requirements to ensure the generated code executes correctly and passes test cases.

We maintain a dedicated `prompt.py` module to store any node prompts that are too verbose for the Mermaid diagram, keeping the workflow clean and readable.

A.2. Mermaid Checker Implementation

We implemented a comprehensive Mermaid validation system that performs both soft and hard checks on workflow diagrams. The soft check analyzes the structure using regex pattern matching to extract different components of the Mermaid code:

- **Node initialization statements** - Extracts node definitions and their types
- **Class definition statements** - Identifies styling classes for nodes
- **Node class assignments** - Maps nodes to their respective classes

- **Node connection statements** - Extracts source, target, and label information from connections

Our validation system performs several critical checks:

1. **W1** - Verifies the presence of required PROBLEM and RETURN interface nodes
2. **W2** - Ensures each node is properly connected in the workflow with paths from PROBLEM and to RETURN
3. **W3** - Validates that PROBLEM and RETURN nodes are correctly classified as Interface types
4. **W4** - Checks that all nodes have valid types according to configuration
5. **W5** - Ensures ScEnsembleOp nodes have at least two incoming connections

We also implement rigorous hard checks by compiling the Mermaid code directly through the Mermaid CLI², which thoroughly verifies syntactic correctness and catches any compilation errors. This comprehensive two-level validation approach ensures both structural integrity through soft checks and strict syntactic validity through hard checks, guaranteeing that all workflow diagrams are executable and error-free.

A.3. Algorithms

In this section, we define the key algorithms underlying the MermaidFlow system. These algorithms form the core components of our workflow optimization and execution process, including the end-to-end MermaidFlow pipeline and the dedicated Mermaid workflow optimization routines that operate over Mermaid graphs. The following algorithms demonstrate how we generate, validate, and optimize workflow, as well as how we transform the Mermaid code into executable Python code.³

Algorithm 1 MermaidFlow End-to-End Pipeline

```
Input: val_dataset, dataset_type, mmd_checker, max_rounds, num_tries, candi_num
initial_workflow ← LoadInitialWorkflow(dataset_type)
mmd_history ← [initial_workflow]
for round = 1 to max_rounds do
    parent_info ← SampleParent(mmd_history)
    mmd_graph, related_prompt ← OptimizeMermaidWorkflow(parent_info, num_tries, candi_num, mmd_checker)
    python_code ← GeneratePythonCode(parent_info, mmd_graph, related_prompt)
    val_score ← EvaluateWorkflow(python_code, validation_dataset)
    mmd_history ← Update(mmd_history, mmd_graph, python_code, val_score)
end for
return mmd_history
```

²<https://github.com/mermaid-js/mermaid-cli>

³We use “mmd” as the abbreviation for Mermaid.

Algorithm 2 OptimizeMermaidWorkflow Function

```

Input: parent_info, num_tries, candi_num, mermaid_checker
last_mmd, last_errors ← null
prev_attempt ← {last_mmd, last_errors}
for  $i = 1$  to num_tries do
  mmd_graph, related_prompt ← GetNewMermaidGraph(parent_info, prev_attempt, candi_num)
  if mermaid_checker exists then
    mmd_path ← mermaid_checker.TransferMmdCodeToFile(mmd_graph)
    (hard_pass, hard_info) ← mermaid_checker.HardCheck(mmd_path)
    (soft_pass, soft_info) ← mermaid_checker.SoftCheck(mmd_path)
    last_mmd ← mmd_graph
    last_errors ← hard_info + soft_info
    prev_attempt ← {last_mmd, last_errors}
    if hard_pass AND soft_pass then
      break ▷ The new Mermaid Code pass soft and hard check
    end if
  else
    break ▷ Accept first response if no checker
  end if
end for
return mmd_graph, related_prompt

```

Algorithm 3 GetNewMermaidGraph Function

```

Input: parent_info, prev_attempt, candi_num
mmd_candidates ← GetMermaidCandidates(parent_info, prev_attempt, candi_num)
selected_mmd ← LLMsJudge(mmd_candidates)
new_graph ← selected_mmd.graph
new_prompt ← selected_mmd.prompt
return new_graph, new_prompt

```

A.3.1. PROMPT TEMPLATES

We design structured prompt templates to guide the LLM in generating, evaluating, and transforming workflows throughout the optimization process. Each template encodes precise instructions, including output format, semantic constraints, and evaluation objectives, ensuring that the model adheres to the structural and functional requirements of MermaidFlow.

Below, we summarize the purpose of each template and provide the full prompt text for reference.

- **UPDATE_MERMAID_WORKFLOW** - Used to guide the LLM in optimizing mermaid code based on parent nodes information. This prompt includes experience data, score information, the current graph in Mermaid format, the corresponding role prompt, operator descriptions, and log data from previous runs. It instructs the LLM to make specific optimizations by adding, modifying, or deleting nodes while clearly describing each modification. The prompt also emphasizes the importance of proper formatting, critical thinking methods, and effective use of specialized operators like Program for solving math problems.
- **MERMAID_GUIDANCE** - Used to guide the generation of correct Mermaid code according to our node type definitions. This prompt provides specific guidance for generating workflow in Mermaid code following the specification in A.1.
- **GENERATE_PYTHON_CODE** - Used to generate python code from given mermaid code and agents prompt.
- **LLM_AS_JUDGE** - Used to select the most promising graph from multiple Mermaid candidates. It evaluates each candidate graph based on criteria including workflow coherence, innovation, complexity balance, prompt quality, and

modification rationale, providing detailed scoring justifications. The judging process references the performance of historical graphs and avoids selecting graphs with structural defects (such as incorrect connections or improper use of ensemble nodes).

UPDATE_MERMAID_WORKFLOW : Depending on whether a single or two parent workflows are used, we adopt different prompt templates to generate the new Mermaid code.

Single Parent

```
UPDATE_MERMAID_WORKFLOW_MATH = """
You are building a Graph and corresponding Prompt to jointly solve {type} problems. Here is a graph written
↳ in Mermaid and the corresponding prompt that performed excellently in a previous iteration (maximum
↳ score is 1). There are also some special operators defined in the operator_description section.
↳ Referring to this given graph and prompt, which forms a basic example of a code solution approach,
↳ please reconstruct and optimize them.
You should make further optimizations based on this graph. The modified graph should differ from the
↳ provided example or have the same architecture but with prompts fine-tuned based on logs. The specific
↳ differences should be noted within the <modification>xxx</modification> section.
<sample>
  <experience>{experience}</experience>
  <modification>(such as: add/delete/modify/...)</modification>
  <score>{score}</score>
  <graph>{graph_mermaid}</graph>
  <role_prompt>{prompt}</role_prompt>(only prompt_custom)
  <operator_description>{operator_description}</operator_description>
</sample>
Below are the logs of some results with the aforementioned Graph that performed well but encountered errors,
↳ which can be used as references for optimization:
{log}

You can add, modify, or delete nodes, parameters, or connections in the workflow. For each change, clearly
↳ describe your single modification within XML tags using the format: <modification>your detailed
↳ modification description</modification>.

Your optimizations should be one of (this limitation is strict! you can only pick one action!):
1. Expanding the graph by adding new single (operators). Note: adding nodes may require modifying prompts
↳ for related nodes, you should also update the corresponding prompts if needed.
2. Deleting unnecessary nodes from the graph.
3. Modifying existing nodes or their connections or their prompt.

Prompt is very important for performance here are some examples you can learn from it:
{prompt_few_shot}

Ensure all necessary prompts are properly prepared and defined. The generated custom node's role is defined
↳ by their prompt. Use custom methods with proper formatting to ensure your output matches the expected
↳ structure. The system will extract answers based on specific rules for scoring, so maintaining the
↳ correct output format is critical.
The output format is critical for proper evaluation. Analyze the log data carefully to identify patterns in
↳ successful solutions and common errors. Extract specific formatting requirements and incorporate them as
↳ clear guidance in your prompts.
When optimizing, you can incorporate critical thinking methods like review, revise, ensemble (generating
↳ multiple answers through different/similar prompts, then voting/integrating/checking the majority to
↳ obtain a final answer), selfAsk, etc.

NOTE: You should also try to use different operators to enhance workflow capabilities.
NOTE: If you are trying to answer from multiple solutions, you should try to use ensemble operator then to
↳ get the final answer.
NOTE: Each output of nodes except the interface node should be connected to the next node, ensuring data
↳ flows through the entire workflow. This guarantees all nodes properly receive inputs and pass outputs.
NOTE: If you are trying to add an ensemble node, you need to guarantee that there are multiple solution
↳ inputs available for the ensemble to work effectively.
NOTE: Program is a very powerful tool for solving MATH problems, you should try it! The Program operator can
↳ execute Python code to perform calculations and solve complex mathematical problems.
NOTE: Think big! Be bold in innovation and exploration. Don't be afraid to push boundaries and discover new
↳ approaches to problem-solving. Additionally, keep prompts concise - no more than 100 words per prompt.
↳ Shorter, focused prompts often perform better than lengthy ones.
"""

UPDATE_MERMAID_WORKFLOW_CODE = """
You are building a Graph and corresponding Prompt to jointly solve {type} problems. Here is a graph written
↳ in Mermaid and the corresponding prompt that performed excellently in a previous iteration (maximum
↳ score is 1). There are also some special operators defined in the operator_description section.
↳ Referring to this given graph and prompt, which forms a basic example of a code solution approach,
↳ please reconstruct and optimize them.
```

You should make further optimizations based on this graph. The modified graph should differ from the provided example or have the same architecture but with prompts fine-tuned based on logs. The specific differences should be noted within the `<modification>xxx</modification>` section.

```
<sample>
  <experience>{experience}</experience>
  <modification>(such as: add/delete/modify/...)</modification>
  <score>{score}</score>
  <graph>{graph_mermaid}</graph>
  <role_prompt>{prompt}</role_prompt>(only prompt_custom)
  <operator_description>{operator_description}</operator_description>
</sample>
```

Below are the logs of some results with the aforementioned Graph that performed well but encountered errors, which can be used as references for optimization:

```
{log}
```

You can add, modify, or delete nodes, parameters, or connections in the workflow. For each change, clearly describe your single modification within XML tags using the format: `<modification>your detailed modification description</modification>`.

Your optimizations should be one of (this limitation is strict! you can only pick one action!):

1. Expanding the graph by adding new single (operators). Note: adding nodes may require modifying prompts for related nodes, you should also update the corresponding prompts if needed.
2. Deleting unnecessary nodes from the graph.
3. Modifying existing nodes or their connections or their prompt.

Prompt is very important for performance here are some examples you can learn from it:

```
{prompt_few_shot}
```

Ensure all necessary prompts are properly prepared and defined. The generated custom node's role is defined by their prompt. Use custom methods with proper formatting to ensure your output matches the expected structure. The system will extract answers based on specific rules for scoring, so maintaining the correct output format is critical.

Review the log data carefully to understand the expected answer format and ensure your implementation produces compatible output. Proper formatting is essential for accurate evaluation of your solution, and you should emphasize this in the prompt.

When optimizing, you can incorporate critical thinking methods like review, revise, ensemble (generating multiple answers through different/similar prompts, then voting/integrating/checking the majority to obtain a final answer), selfAsk, etc.

You should also try to use different operators to enhance workflow capabilities. Each operator should have the corresponding Mermaid class defined in the graph code.

NOTE: If you are trying to add an ensemble node, you need to guarantee that there are multiple solution inputs available for the ensemble to work effectively.

NOTE: In the code problems, ensemble different kinds of solutions is a good idea, and you should also try to use Test operator to validate the solutions.

NOTE: For code-related tasks, the final output must be either `test_result["solution"]` or `custom_code_generate_result["response"]`, as these represent valid Python code. Make sure your workflow produces one of these formats as the final output to ensure proper evaluation.

NOTE: Think big! Be bold in innovation and exploration. Don't be afraid to push boundaries and discover new approaches to problem-solving. Additionally, keep prompts concise - no more than 100 words per prompt. Shorter, focused prompts often perform better than lengthy ones.

```
"""
```

Two Parents

```
UPDATE_MERMAID_WORKFLOW = """
# Evolution-Based Graph Generation
```

You are tasked with generating a new workflow graph by combining elements from two parent graphs. This process is inspired by genetic algorithms where offspring inherit traits from both parents. This evolutionary approach allows us to create new solutions that may be more effective than either parent graph.

```
<parent_graph>
  <graph_A>
    <graph>{graph_A}</graph>
    <prompt>{prompt_A}</prompt>
    <score>{score_A}</score>
  </graph_A>
  <graph_B>
    <graph>{graph_B}</graph>
    <prompt>{prompt_B}</prompt>
    <score>{score_B}</score>
  </graph_B>
</parent_graph>
```

```
## Your Task
1. Analyze both **parent graphs** carefully (structure, nodes, connections, prompts, and purpose)
2. Create a new graph that:
  - Combines strengths from both parents
  - Introduces strategic innovations for improved performance
  - Applies evolutionary operations such as:
    * Crossover: Merging effective sections(nodes) from both parents, you should separate components in a
      ↳ reasonable and logical way, typically at interface or ensemble nodes
    * Mutation: Making targeted modifications
    * Insertion: Adding beneficial new nodes
    * Deletion: Removing inefficient components

## Guidelines
- Focus on the strengths of both parent graphs, especially the higher-scoring parent, to inform your design
↳ decisions
- Ensure your new graph is structurally correct and follows proper workflow patterns
- If creating a significantly different structure proves challenging, you may maintain a similar structure
↳ to the parents but with improved prompts, or, you can simply add one extra node to ensemble
- Pay careful attention to the connections between nodes to ensure proper data flow - for example, ensemble
↳ nodes should have multiple inputs
- Verify that your graph has correct input/output relationships and follows the operator usage guidelines
- Use custom methods to restrict your output format, rather than using code (outside of the code, the system
↳ will extract answers based on certain rules and score them)

## Prompt Guidance
- **Prompt engineering is crucial for performance**: Carefully analyze and learn from the prompts used in
↳ **high-scoring** historical workflows. Pay special attention to their structure, specificity, and
↳ instruction clarity. Here are some exemplary prompts you can use as reference:
  {prompt_few_shot}

Your response should include:

1. A detailed explanation of your modifications in the <modification> section
2. The complete Mermaid code for the new graph
3. The updated prompts for any custom nodes

NOTE: Ensure the new graph is valid Mermaid syntax and represents a complete workflow solution. The
↳ following section contains critical rules and guidance for using operators in Mermaid that you MUST
↳ follow to create an effective workflow:
{mermaid_custom_prompt}
"""
```

MERMAID GUIDANCE

Mermaid Code Guidance

```
MERMAID_CUSTOM_MATH = """
# Mermaid Graph Style Guide for MATH Problems
# This comprehensive guide defines styling and structure for creating consistent workflow diagrams

# Node Style Definitions
classDef CustomOp fill:#d0e1f9,stroke:#4378a2,stroke-width:2px;
classDef ProgrammerOp fill:#f9c2c2,stroke:#c23737,stroke-width:2px;
classDef ScEnSembleOp fill:#f9e4b7,stroke:#b99b37,stroke-width:2px;
classDef Interface fill:#e2e2f2,stroke:#6a6ab2,stroke-width:2px;

# ===== OPERATOR USAGE GUIDE =====

# 1. Interface Nodes (Entry/Exit Points)
# Every workflow diagram must include these two standard interface nodes:
# PROBLEM([Problem]) - The entry point providing initial input
# RETURN([Return response & cost]) - The exit point receiving final output
#
# Example:
# PROBLEM([Problem])
# RETURN([Return response & cost])
#
# class PROBLEM Interface
# class RETURN Interface
#
# Connection rules:
# - PROBLEM node: Provides input to other nodes but never receives input
# - RETURN node: Receives output from the final node but never produces output

# 2. Custom Operator Nodes
```

```

# Format: NodeName["Custom<br/>(role: role_name)"]
#
# Example:
# K["Custom<br/>(role: validate_1)"]
# class K CustomOp
#
# For multiple nodes with similar roles, use numbered suffixes:
# R1["Custom<br/>(role: review_solution_1)"]
# R2["Custom<br/>(role: review_solution_2)"]
# class R1 CustomOp
# class R2 CustomOp
#
# Connection example:
# PROBLEM --> |input| R1

# 3. Programmer Operator Nodes
# Format: P["Programmer<br/>(analysis: 'your analysis text')"]
#
# The Programmer operator requires two inputs:
# - problem: The math problem to solve
# - analysis: Instructions on how to approach the problem
#
# Examples:
# P["Programmer<br/>(analysis: 'Calculate step by step')"]
# class P ProgrammerOp
#
# Connection rules:
# 1. For the problem input:
# PROBLEM --> |problem|P # The problem must come from the PROBLEM node
#
# 2. For the analysis input (two options):
# C --> |analysis|P # You can use another node's output as analysis content
# # OR
# # You can specify the analysis directly in the node definition
#
# Complete example:
# PROBLEM --> |problem|P
# C --> |analysis|P
# class P ProgrammerOp

# 4. ScEnsemble Operator Nodes
# Format: ENSEMBLE["ScEnsemble<br/>"]
#
# Example:
# ENSEMBLE["ScEnsemble<br/>"]
# class ENSEMBLE ScEnSembleOp
#
# CRITICAL: ScEnsemble nodes MUST have multiple inputs to function correctly
# Example connections:
# SOLUTION_1 --> ENSEMBLE
# SOLUTION_2 --> ENSEMBLE
# ENSEMBLE --> NEXT_NODE

# ===== WORKFLOW PATTERNS =====

# Example Workflow Pattern
# This pattern demonstrates how to effectively combine multiple solution approaches:
# 1. Problem input flows to multiple solution-generating nodes (Custom and/or Programmer)
# 2. All solutions are combined using ScEnsemble
# 3. The ensemble result is returned as the final output
#
# Connection pattern:
# PROBLEM --> |input| SOLUTION_1
# PROBLEM --> |input| SOLUTION_2
# PROBLEM --> |problem| P
# SOLUTION_1 --> ENSEMBLE
# SOLUTION_2 --> ENSEMBLE
# P --> ENSEMBLE
# ENSEMBLE --> RETURN

# ===== GENERAL RULES =====

# Connection Rules
# 1. Direct connections from PROBLEM to any node that needs the initial problem as input
# 2. Direct connections between nodes where one node requires another's output
# 3. When using ProgrammerOp, the edges should have name on it
# 4. All connections must follow logical data flow and maintain workflow coherence

```

```

# Prompt Definition Requirements
# All prompts used in the graph must be defined in this format:
# <prompt>
# PROMPT_NAME_1="Your first prompt text here"
# PROMPT_NAME_2="Your second prompt text here"
# </prompt>
#
# Each prompt must be on a separate line with its own unique variable name

# IMPORTANT: Do not create new Operator types!
# Only use the predefined operators available in the system. Creating custom operators
# will cause workflow execution failures. Stick to the operators documented in this guide.

# Best Practices:
# - Use %% for comments in separate lines for clarity, don't add it to the end of the line because it is not
↳ a valide mermaid code
# - Always include the style block (classDef section) even if not all classes are used
# - Always ensure multiple inputs to ScEnsemble nodes (this is a common mistake)
# - Maintain consistent naming conventions for all nodes and classes
# - Use descriptive role names that indicate the node's purpose
# - Label connections with |input| where appropriate for clarity

"""
MERMAID_CUSTOM_CODE = """
# Operator Style Definitions
# Each operator has a specific style for consistent visualization in the Mermaid graph
classDef CustomOp fill:#d0e1f9,stroke:#4378a2,stroke-width:2px;
classDef CustomCodeGenerateOp fill:#f9c2c2,stroke:#c23737,stroke-width:2px;
classDef ScEnSembleOp fill:#f9e4b7,stroke:#b99b37,stroke-width:2px;
classDef TestOp fill:#d8f0d8,stroke:#2e8b57,stroke-width:2px;
classDef Interface fill:#e2e2f2,stroke:#6a6ab2,stroke-width:2px;

# ===== OPERATOR USAGE GUIDE =====

# 1. Interface Nodes
# Interface nodes represent the entry and exit points of your workflow
# Required in every graph:
#   PROBLEM({Problem})           - Starting point that provides input
#   RETURN({Return response & cost}) - Endpoint that receives final output
#
# Example:
#   class PROBLEM Interface
#   class RETURN Interface
#
# Connection rules:
#   - PROBLEM node provides input but doesn't receive any
#   - RETURN node receives output but doesn't produce any

# 2. Custom Operator
# Used for specialized reasoning strategies with defined roles
#
# Example:
#   K["Custom<br/>(role: xxx)"]
#   class K CustomOp
#
# For multiple instances with same prompt type:
#   R1["Custom<br/>(role: xxx_1)"]
#   R2["Custom<br/>(role: xxx_2)"]
#   class R1 CustomOp
#   class R2 CustomOp
#
# Connection rules:
#   - Connect directly to PROBLEM if it needs the initial problem
#   - Connect to other nodes if it needs their output

# 3. CustomCodeGenerate Operator
# Specialized for code generation tasks based on problem descriptions
#
# Example:
#   CODE_GEN["CustomCodeGenerate<br/>(instruction: xxx)"]
#   class CODE_GEN CustomCodeGenerateOp
#
# Multiple approaches:
#   CODE_GEN_1["CustomCodeGenerate<br/>(instruction: aaa)"]
#   CODE_GEN_2["CustomCodeGenerate<br/>(instruction: bbb)"]
#   class CODE_GEN_1 CustomCodeGenerateOp
#   class CODE_GEN_2 CustomCodeGenerateOp

```

```
#
# 4. ScEnsemble Operator
# Combines multiple solutions into one cohesive result
#
# Example:
# ENSEMBLE["ScEnsemble<br/>"]
# class ENSEMBLE ScEnsembleOp
#
# CRITICAL: Must have multiple inputs to function correctly
# SOLUTION_1 --> ENSEMBLE
# SOLUTION_2 --> ENSEMBLE
# ENSEMBLE --> other_node

# 5. Test Operator
# Validates solutions against test cases
#
# Example:
# T["Test<br/>"]
# class T TestOp
#
# Typical connection pattern:
# CODE_GEN --> |solution|T
# ENTRY_POINT --> |entry_point|T
# PROBLEM --> |problem|T
# T --> DECISION_NODE
#
# Decision node for test results:
# CHECK_TEST{test_result<br/>passed?}
# class CHECK_TEST DecisionOp
# CHECK_TEST -- Failed --> IMPROVE_SOLUTION
# CHECK_TEST -- Passed --> RETURN

# ===== IMPORTANT NOTES =====
# - You cannot create other class operations
# - Always ensure multiple inputs to ScEnsemble nodes
# - All prompts must be defined in the prompt section
# - Format prompts as:
# <prompt>
# PROMPT_NAME_1="Your first prompt text here"
# PROMPT_NAME_2="Your second prompt text here"
# </prompt>

# IMPORTANT: Do not create new Operator types!
# Only use the predefined operators available in the system. Creating custom operators
# will cause workflow execution failures. Stick to the operators documented in this guide.
"""
```

GENERATE PYTHON CODE

Generate Python Code

```
GENERATE_PYTHON_CODE = """
Below is a graph, the corresponding Python code, and the prompt.
<old_workflow>
  <old_graph>{old_mermaid}</old_graph>
  <old_code>{old_code}</old_code>
  <old_role_prompt>{old_prompt}</old_role_prompt> (only prompt_custom)
</old_workflow>

Based on this example of old graph and code, you need to generate new Python code according to the new graph
↔ and given prompt.

<information_for_new_workflow>
  <new_graph>{new_mermaid}</new_graph>
  <new_role_prompt>{new_prompt}</new_role_prompt> (only prompt_custom)
  <operator_description>{operator_description}</operator_description>
</information_for_new_workflow>

The output format should be:

<code>New generated python code</code>

Carefully analyze the new_graph to generate corresponding code. Pay attention to:
1. The connections between each node
2. The role and function of each operator

```

3. The input/output relationships

For each node in the graph, implement the appropriate code and use the corresponding prompts from `new_prompt`.
→ If a node or operator doesn't have an explicit prompt provided, DO NOT create one - use empty strings
→ instead. For example, Program operators may have empty analysis fields.
Every prompt referenced in your Python code must be defined in the `prompt_custom` module. When adding new
→ functionality to your Python code, make sure to import any necessary libraries or modules, except for
→ operator, `prompt_custom`, `create_llm_instance`, and `CostManage` which are automatically imported.
Your implementation must be robust - ensure all methods return appropriate values and **never** return `None`
→ for any field**. Pay special attention to error handling and edge cases.
Use custom methods with proper formatting to ensure your output matches the expected structure. The system
→ will extract answers based on specific rules for scoring, so maintaining the correct output format is
→ critical.

NOTE: especially for the final output, you should ensure that the output format is correct, you can learn it
→ from old code.
"""

LLM AS JUDGE

LLM As Judge

```
LLM_AS_JUDGER = """
```

Your task is to select the most promising graph from the candidates. Here is the content of each graph,
→ written in Mermaid code, along with explanations of how each new graph was generated from parent graphs:
Each new graph is derived from two parent graphs, with parent graph information as follows:

```
<parent_graph>
  <parent_graph_A>{parent_graph_A}</parent_graph_A>
  <parent_prompt_A>{parent_prompt_A}</parent_prompt_A>

  <parent_graph_B>{parent_graph_B}</parent_graph_B>
  <parent_prompt_B>{parent_prompt_B}</parent_prompt_B>
</parent_graph>
```

```
<graph_candidates>
  <graph_A>{graph_A}</graph_A>
  <modification_A>{modification_A}</modification_A>
  <prompt_A>{prompt_A}</prompt_A>

  <graph_B>{graph_B}</graph_B>
  <modification_B>{modification_B}</modification_B>
  <prompt_B>{prompt_B}</prompt_B>

  <graph_C>{graph_C}</graph_C>
  <modification_C>{modification_C}</modification_C>
  <prompt_C>{prompt_C}</prompt_C>

  <graph_D>{graph_D}</graph_D>
  <modification_D>{modification_D}</modification_D>
  <prompt_D>{prompt_D}</prompt_D>
</graph_candidates>
```

Please evaluate the graph candidates and select the most promising one based on these criteria:

1. **Workflow Coherence**: Assess how well the nodes connect and form a logical workflow
2. **Innovation**: Evaluate how the new graph improves upon the parent graphs
3. **Complexity Balance**: Check if the graph has appropriate complexity (neither too simple nor
→ unnecessarily complex)
4. **Prompt Quality**: Examine the quality and specificity of the node prompts
5. **Modification Rationale**: Consider the thoughtfulness of the explanation provided for the changes

Here are some history of previous graphs and their corresponding score:

```
<history>
  {elites_history}
</history>
```

For each candidate graph, provide a score from 1-10 for each criterion and explain your reasoning, you can
→ learn from the history graphs.

You should avoid selecting graphs with structural defects, such as:

1. Incorrectly connecting nodes (e.g., `CustomOp` should not directly feed into `ProgrammerOp`)
2. Not properly using the ensemble node (all solution-generating nodes should feed into it)

3. Missing critical connections between nodes
4. Creating circular dependencies

Here are the specific structural rules for this type of workflow:
{mermaid_usage}

Additionally, consider how well the graph follows established patterns from successful historical examples.

Then select the graph with the highest total score as the most promising candidate.

```
<evaluation>
  <graph_A_score>
    <workflow_coherence>Score (1-10)</workflow_coherence>
    <innovation>Score (1-10)</innovation>
    <complexity_balance>Score (1-10)</complexity_balance>
    <prompt_quality>Score (1-10)</prompt_quality>
    <modification_rationale>Score (1-10)</modification_rationale>
    <total_score>Sum of all scores</total_score>
    <explanation>Detailed explanation of your evaluation</explanation>
  </graph_A_score>

  <graph_B_score>
    <workflow_coherence>Score (1-10)</workflow_coherence>
    <innovation>Score (1-10)</innovation>
    <complexity_balance>Score (1-10)</complexity_balance>
    <prompt_quality>Score (1-10)</prompt_quality>
    <modification_rationale>Score (1-10)</modification_rationale>
    <total_score>Sum of all scores</total_score>
    <explanation>Detailed explanation of your evaluation</explanation>
  </graph_B_score>

  <graph_C_score>
    <workflow_coherence>Score (1-10)</workflow_coherence>
    <innovation>Score (1-10)</innovation>
    <complexity_balance>Score (1-10)</complexity_balance>
    <prompt_quality>Score (1-10)</prompt_quality>
    <modification_rationale>Score (1-10)</modification_rationale>
    <total_score>Sum of all scores</total_score>
    <explanation>Detailed explanation of your evaluation</explanation>
  </graph_C_score>

  <graph_D_score>
    <workflow_coherence>Score (1-10)</workflow_coherence>
    <innovation>Score (1-10)</innovation>
    <complexity_balance>Score (1-10)</complexity_balance>
    <prompt_quality>Score (1-10)</prompt_quality>
    <modification_rationale>Score (1-10)</modification_rationale>
    <total_score>Sum of all scores</total_score>
    <explanation>Detailed explanation of your evaluation</explanation>
  </graph_D_score>
</evaluation>

<selected_graph>[A/B/C/D]</selected_graph>
<justification>
  Please provide a comprehensive justification for your selection, highlighting the key strengths of the
  ↪ chosen graph and how it represents the most effective approach to solving the problem.
</justification>
"""
```

A.4. Details of Baseline Methods

In this section, we provide a detailed description of the configurations for baseline methods:

- **CoT (Kojima et al., 2022)**. Chain-of-Thought prompting guides the LLM to decompose complex reasoning into a sequence of intermediate steps, rather than generating a direct answer.
- **ComplexCoT (Fu et al., 2023)**. This method builds on CoT by explicitly modeling and controlling reasoning complexity across multiple steps. We use the official implementation from <https://github.com/FranxYao/Complexity-Based-Prompting/tree/main>.
- **Self-Consistency (Wang et al., 2023)**. This method generates five independent chain-of-thought trajectories by sampling the LLM multiple times, then consolidates them via majority voting to improve the reliability of the final answer.

- **LLM-Debate** (Du et al., 2024). A panel of five role-specific LLM agents engage in up to two rounds of structured debate; the final answer is selected by majority vote.
- **DyLAN** (Liu et al., 2024). DyLAN employs a dynamic multi-agent framework in which agents iteratively share and refine intermediate reasoning to converge on high-quality solutions.
- **MacNet** (Qian et al., 2024). MacNet structures its agent network as a fully meshed graph, enabling every agent to exchange and integrate information with all others at each reasoning step.
- **GPTSwarm** (Zhuge et al., 2024). GPTSwarm leverages a swarm-intelligence paradigm by orchestrating multiple LLM agents in parallel; agents independently propose solutions and iteratively refine them through shared feedback, following the protocol outlined in the original paper.
- **MaAS** (Zhang et al., 2025b). MaAS integrates a trainable module for dynamic workflow assignment across multiple agents; we replicate the authors’ original architecture using their official code and retain their specified hyperparameters without modification.
- **AutoAgents** (Chen et al., 2024). AutoAgents orchestrates a pipeline of specialized LLM agents that autonomously schedule, execute, and refine subtasks through coordinated interactions; we replicate the authors’ default configuration and parameter settings from the official repository.
- **ADAS** (Hu et al., 2024). ADAS employs an adaptive debate-and-selection mechanism among multiple LLM agents, dynamically refining candidate solutions based on quality criteria.
- **AFlow** (Zhang et al., 2024c). In the original AFlow study, both gpt-4o-mini and claude-3.5-sonnet are used. For a fair comparison, we constrain AFlow to gpt-4o-mini only and set MAX_ITERATION to 20.

A.5. Details of Tasks and Benchmarks

Following established methodologies in workflow automation (Zhang et al., 2024c; 2025b), we partition each dataset into training and test sets with a TRAIN:TEST ratio of 1:4. For the MATH benchmark, we follow the approach in (Zhang et al., 2025b), selecting a subset of 617 challenging problems across four representative categories: Combinatorics & Probability, Number Theory, Pre-algebra, and Pre-calculus, all at difficulty level 5. The dataset statistics are summarized in Table 3.

Table 3. Statistics of datasets used in our experiments.

Domain	Dataset	Training	Testing	Evaluation Metric
Code Generation	HumanEval	33	131	pass@1
	MBPP	86	341	pass@1
Math Reasoning	GSM8K	264	1,055	Accuracy
	MATH	119	486	Accuracy

B. Case Study

In this section, we present examples of the optimal workflows discovered by our approach for each of the four datasets. For each dataset, we provide the **Mermaid code**, corresponding **Python code**, and the rendered **workflow diagram**.

The declarative Mermaid representation serves as the backbone of structured workflow generation, enabling several crucial capabilities:

1. **Typed, Semantically Aligned Representation.** Mermaid encodes workflows as statically typed and semantically annotated graphs, allowing prompt semantics and operator roles to be aligned explicitly within the workflow structure.
2. **Human-Readable and Verifiable Syntax.** Unlike imperative representations, Mermaid provides a format that is both visually interpretable and statically verifiable, facilitating intuitive debugging, planning, and workflow validation.

3. **Reliable Code Translation.** The structured nature of Mermaid graphs enables seamless compilation into executable Python code, with consistent type and role guarantees that reduce runtime errors and improve reliability.

These properties make MermaidFlow not just a search mechanism, but a principled framework for building modular, interpretable, and safer agentic workflows. The case studies illustrate how these benefits translate into practical gains across crucial domains such as programming, math, and symbolic reasoning.

B.1. GSM8K

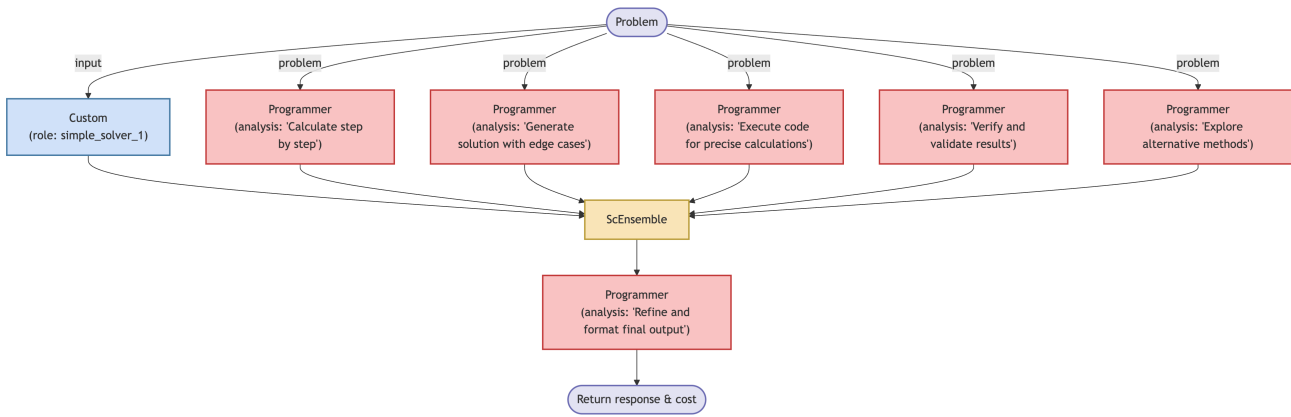


Figure 5. Mermaid diagram for GSM8K.

Mermaid Code

```

flowchart TD
%% Nodes
PROBLEM([Problem])
C[Custom<br/>(role: simple_solver_1)]
P1["Programmer<br/>(analysis: 'Calculate step  
by step')"]
P2["Programmer<br/>(analysis: 'Generate  
solution with edge cases')"]
P3["Programmer<br/>(analysis: 'Execute code  
for precise calculations')"]
P4["Programmer<br/>(analysis: 'Verify and  
validate results')"]
P5["Programmer<br/>(analysis: 'Explore  
alternative methods')"]
P6["Programmer<br/>(analysis: 'Refine and  
format final output')"]
ENSEMBLE[ScEnsemble]
RETURN([Return response & cost])

C --> ENSEMBLE
P1 --> ENSEMBLE
P2 --> ENSEMBLE
P3 --> ENSEMBLE
P4 --> ENSEMBLE
P5 --> ENSEMBLE
ENSEMBLE --> P6
P6 --> RETURN
  
```

```

%% Styles
classDef CustomOp fill:#d0e1f9,stroke:#4378a2,stroke-width:2px;
classDef ProgrammerOp fill:#f9c2c2,stroke:#c23737,stroke-width:2px;
classDef ScEnsembleOp fill:#f9e4b7,stroke:#b99b37,stroke-width:2px;
classDef Interface fill:#e2e2f2,stroke:#6a6ab2,stroke-width:2px;

%% Assign classes
class C CustomOp
  
```

```

class P1 ProgrammerOp
class P2 ProgrammerOp
class P3 ProgrammerOp
class P4 ProgrammerOp
class P5 ProgrammerOp
class P6 ProgrammerOp
class ENSEMBLE ScEnsembleOp
class PROBLEM Interface
class RETURN Interface

%% Flow (arrows show data relationships)
PROBLEM --> |input|C
PROBLEM --> |problem|P1
PROBLEM --> |problem|P2
PROBLEM --> |problem|P3
PROBLEM --> |problem|P4
PROBLEM --> |problem|P5
C --> ENSEMBLE
P1 --> ENSEMBLE
P2 --> ENSEMBLE
P3 --> ENSEMBLE
P4 --> ENSEMBLE
P5 --> ENSEMBLE
ENSEMBLE --> P6
P6 --> RETURN

```

Python Code

```

from typing import Literal
import workspace.GSM8K.workflows.template.operator as operator
import workspace.GSM8K.workflows.round_16.prompt as prompt_custom
from scripts.async_llm import create_llm_instance
import weave

DatasetType = Literal["HumanEval", "MBPP", "GSM8K", "MATH", "HotpotQA", "DROP"]

class Workflow:
    def __init__(
        self,
        name: str,
        llm_config,
        dataset: DatasetType,
    ) -> None:
        self.name = name
        self.dataset = dataset
        self.llm = create_llm_instance(llm_config)
        self.sc_ensemble = operator.ScEnsemble(self.llm)
        self.custom = operator.Custom(self.llm)
        self.programmer = operator.Programmer(self.llm)

    @weave.op()
    async def __call__(self, problem: str):
        """
        Implementation of the workflow
        Each operator is callable, you can call it directly.
        """
        # Step 1: Use the Custom operator to generate a detailed solution
        custom_response = await self.custom(input=problem, instruction=prompt_custom.SIMPLE_SOLVER_1,
        ↪ role="simple_solver_1")

        # Step 2: Use the Programmer operator to analyze the problem and provide a code solution
        programmer_response_1 = await self.programmer(problem=problem, analysis="Calculate step by step")

        # Step 3: Use the Programmer operator to generate a solution considering edge cases
        programmer_response_2 = await self.programmer(problem=problem, analysis="Generate solution with edge
        ↪ cases")

        # Step 4: Use the Programmer operator to execute code for precise calculations
        programmer_response_3 = await self.programmer(problem=problem, analysis="Execute code for precise
        ↪ calculations")

        # Step 5: Use the Programmer operator to verify and validate results
        programmer_response_4 = await self.programmer(problem=problem, analysis="Verify and validate
        ↪ results")

```

```
# Step 6: Use the Programmer operator to explore alternative methods
programmer_response_5 = await self.programmer(problem=problem, analysis="Explore alternative
↳ methods")

# Step 7: Combine the responses from Custom and all Programmer responses for the ScEnsemble
solutions = [
    custom_response['response'],
    programmer_response_1['output'],
    programmer_response_2['output'],
    programmer_response_3['output'],
    programmer_response_4['output'],
    programmer_response_5['output']
]

# Step 8: Use the ScEnsemble operator to select the best solution
ensemble_response = await self.sc_ensemble(solutions=solutions, problem=problem)

# Step 9: Refine and format the final output
final_output = await self.programmer(problem=ensemble_response['response'], analysis="Refine and
↳ format final output")

return final_output['output'], self.llm.get_usage_summary()["total_cost"]
```

Corresponding Prompt

```
SIMPLE_SOLVER_1 = """Please solve the given mathematical problem step by step. Follow these guidelines:
```

1. State the problem clearly.
2. Outline the approach and any relevant formulas or concepts.
3. Provide detailed calculations, using LaTeX notation for mathematical expressions.
4. Explain each step of your reasoning.
5. Verify and validate your results to ensure accuracy.
6. Present the final answer enclosed in `\boxed{}` LaTeX notation.
7. Ensure all mathematical notation is in LaTeX format.

```
Your solution should be thorough, mathematically sound, and easy to understand."""
```

B.2. MATH

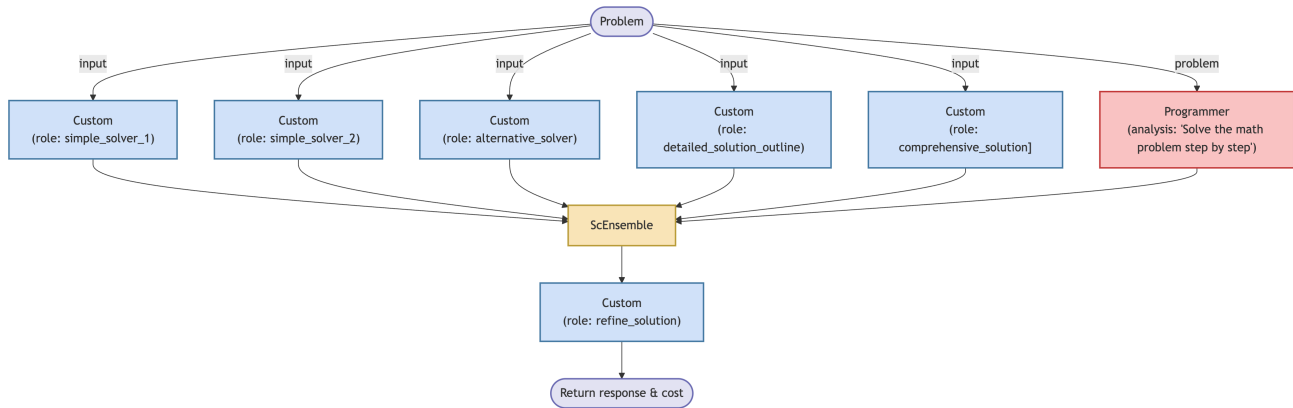


Figure 6. Mermaid diagram for MATH.

Mermaid Code

```

flowchart TD
    %% Nodes
    PROBLEM([Problem])
    C1["Custom<br/>(role: simple_solver_1)"]
    C2["Custom<br/>(role: simple_solver_2)"]
    C3["Custom<br/>(role: alternative_solver)"]
    C4["Custom<br/>(role: detailed_solution_outline)"]
    C5["Custom<br/>(role: comprehensive_solution)"]
    P["Programmer<br/>(analysis: 'Solve the math  
problem step by step')"]
    REFINER["Custom<br/>(role: refine_solution)"]
    ENSEMBLE["ScEnsemble<br/>"]
    RETURN([Return response & cost])

    %% Styles
    classDef CustomOp fill:#d0elf9,stroke:#4378a2,stroke-width:2px;
    classDef ProgrammerOp fill:#f9c2c2,stroke:#c23737,stroke-width:2px;
    classDef ScEnsembleOp fill:#f9e4b7,stroke:#b99b37,stroke-width:2px;
    classDef Interface fill:#e2e2f2,stroke:#6a6ab2,stroke-width:2px;

    %% Assign classes
    class C1 CustomOp
    class C2 CustomOp
    class C3 CustomOp
    class C4 CustomOp
    class C5 CustomOp
    class P ProgrammerOp
    class REFINER CustomOp
    class PROBLEM Interface
    class RETURN Interface
    class ENSEMBLE ScEnsembleOp

    %% Flow (arrows show data relationships)
    PROBLEM -- input --> C1
    PROBLEM -- input --> C2
    PROBLEM -- input --> C3
    PROBLEM -- input --> C4
    PROBLEM -- input --> C5
    PROBLEM -- problem --> P
    C1 --> ENSEMBLE
    C2 --> ENSEMBLE
    C3 --> ENSEMBLE
    C4 --> ENSEMBLE
    C5 --> ENSEMBLE
    P --> ENSEMBLE
    ENSEMBLE --> REFINER
    REFINER --> RETURN
    
```

```

PROBLEM --> |input|C1
PROBLEM --> |input|C2
PROBLEM --> |input|C3
PROBLEM --> |input|C4
PROBLEM --> |input|C5
PROBLEM --> |problem|P
C1 --> ENSEMBLE
C2 --> ENSEMBLE
C3 --> ENSEMBLE
C4 --> ENSEMBLE
C5 --> ENSEMBLE
P --> ENSEMBLE
ENSEMBLE --> REFINED
REFINED --> RETURN
    
```

Python Code

```

from typing import Literal
import workspace.MATH.workflows.template.operator as operator
import workspace.MATH.workflows.round_16.prompt as prompt_custom
from scripts.async_llm import create_llm_instance
import weave

DatasetType = Literal["HumanEval", "MBPP", "GSM8K", "MATH", "HotpotQA", "DROP"]

class Workflow:
    def __init__(
        self,
        name: str,
        llm_config,
        dataset: DatasetType,
    ) -> None:
        self.name = name
        self.dataset = dataset
        self.llm = create_llm_instance(llm_config)
        self.custom1 = operator.Custom(self.llm)
        self.custom2 = operator.Custom(self.llm)
        self.custom3 = operator.Custom(self.llm)
        self.custom4 = operator.Custom(self.llm)
        self.custom5 = operator.Custom(self.llm)
        self.programmer = operator.Programmer(self.llm)
        self.refine = operator.Custom(self.llm)
        self.ensemble = operator.ScEnsemble(self.llm)

    @weave.op()
    async def __call__(self, problem: str):
        """Implementation of the workflow"""
        # Use the first custom operator to solve the problem step by step
        custom_response_1 = await self.custom1(input=problem, instruction=prompt_custom.SIMPLE_SOLVER,
        ↪ role="simple_solver_1")

        # Use the second custom operator to solve the problem step by step
        custom_response_2 = await self.custom2(input=problem, instruction=prompt_custom.SIMPLE_SOLVER,
        ↪ role="simple_solver_2")

        # Use the third custom operator to provide an alternative approach to the problem
        custom_response_3 = await self.custom3(input=problem, instruction=prompt_custom.ALTERNATIVE_SOLVER,
        ↪ role="alternative_solver")

        # Use the fourth custom operator to provide a detailed solution outline
        custom_response_4 = await self.custom4(input=problem,
        ↪ instruction=prompt_custom.DETAILED_SOLUTION_OUTLINE, role="detailed_solution_outline")

        # Use the fifth custom operator to provide a comprehensive solution
        custom_response_5 = await self.custom5(input=problem,
        ↪ instruction=prompt_custom.COMPREHENSIVE_SOLUTION, role="comprehensive_solution")

        # Use the programmer operator to analyze the problem and provide a detailed solution
        programmer_response = await self.programmer(problem=problem, analysis="Solve the math problem step
        ↪ by step")

        # Combine all responses into a list for ensemble processing
        solutions = [
            custom_response_1['response'],
            custom_response_2['response'],
    
```

```
        custom_response_3['response'],
        custom_response_4['response'],
        custom_response_5['response'],
        programmer_response['output']
    ]

    # Use the ensemble operator to select the best solution
    ensemble_response = await self.ensemble(solutions=solutions, problem=problem)

    # Use the refine operator to ensure clarity and correctness of the output
    refined_response = await self.refine(input=ensemble_response['response'],
    ↪ instruction=prompt_custom.REFINE_SOLUTION, role="refine_solution")

    # Return the final output and cost
    return refined_response['response'], self.llm.get_usage_summary()["total_cost"]
```

Corresponding Prompt

```
SIMPLE_SOLVER = """Please solve the given mathematical problem step by step. Follow these guidelines:

1. State the problem clearly.
2. Outline the approach and any relevant formulas or concepts.
3. Provide detailed calculations, using LaTeX notation for mathematical expressions.
4. Explain each step of your reasoning.
5. Present the final answer enclosed in \boxed{} LaTeX notation.
6. Ensure all mathematical notation is in LaTeX format.
"""

ALTERNATIVE_SOLVER = """Please provide an alternative approach to solving the given mathematical problem.
↪ Follow these guidelines:

1. Clearly restate the problem.
2. Identify any different methods or perspectives that could be applied.
3. Provide calculations and reasoning, using LaTeX notation for mathematical expressions.
4. Ensure clarity and correctness in your explanation.
5. Present the final answer enclosed in \boxed{} LaTeX notation.
"""

DETAILED_SOLUTION_OUTLINE = """Please provide a detailed outline for solving the given mathematical problem.
↪ Follow these guidelines:

1. Clearly restate the problem.
2. Identify key concepts and theorems relevant to the problem.
3. Outline the steps needed to solve the problem, including any necessary calculations.
4. Ensure that the outline is structured logically and is easy to follow.
5. Use LaTeX notation for any mathematical expressions.
"""

COMPREHENSIVE_SOLUTION = """Please provide a comprehensive solution to the given mathematical problem.
↪ Follow these guidelines:

1. Clearly restate the problem.
2. Explain the mathematical concepts and theorems involved.
3. Provide a detailed, logical progression of steps leading to the solution.
4. Show all calculations using LaTeX notation for mathematical expressions.
5. Present the final answer clearly marked and enclosed in \boxed{} LaTeX notation.
"""

REFINE_SOLUTION = """Given the mathematical problem and the solutions generated, please refine the output to
↪ ensure clarity and correctness. Follow these guidelines:

1. Review the solutions provided.
2. Ensure all calculations are accurate and clearly presented.
3. Summarize the findings and present the final answer in a clear format.
4. Use LaTeX notation for any mathematical expressions.
5. Ensure the final answer is enclosed in \boxed{} LaTeX notation.
"""
```

B.3. HumanEval

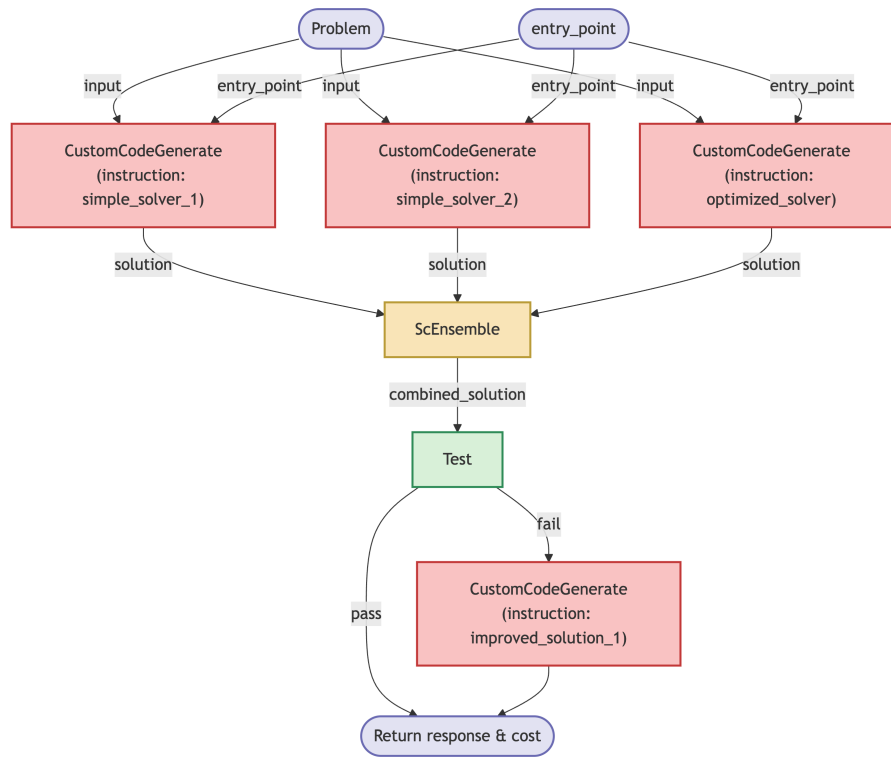


Figure 7. Mermaid diagram for HumanEval.

Mermaid Code

```

flowchart TD
    %% Nodes
    PROBLEM([Problem])
    ENTRY_POINT([entry_point])
    C1["CustomCodeGenerate<br/>(instruction: simple_solver_1)"]
    C2["CustomCodeGenerate<br/>(instruction: simple_solver_2)"]
    C3["CustomCodeGenerate<br/>(instruction: optimized_solver)"]
    C4["CustomCodeGenerate<br/>(instruction: improved_solution_1)"]
    ENSEMBLE["ScEnsemble<br/>"]
    T["Test<br/>"]
    RETURN([Return response & cost])

    %% Styles
    classDef CustomOp fill:#d0e1f9,stroke:#4378a2,stroke-width:2px;
    classDef CustomCodeGenerateOp fill:#f9c2c2,stroke:#c23737,stroke-width:2px;
    classDef ScEnsembleOp fill:#f9e4b7,stroke:#b99b37,stroke-width:2px;
    classDef TestOp fill:#d8f0d8,stroke:#2e8b57,stroke-width:2px;
    classDef Interface fill:#e2e2f2,stroke:#6a6ab2,stroke-width:2px;

    %% Assign classes
    class PROBLEM Interface
    class ENTRY_POINT Interface
    class C1 CustomCodeGenerateOp
    class C2 CustomCodeGenerateOp
    class C3 CustomCodeGenerateOp
    class C4 CustomCodeGenerateOp
    class ENSEMBLE ScEnsembleOp
    class T TestOp
    class RETURN Interface

    %% Flow (arrows show data relationships)
    PROBLEM --> |input| C1
    
```

```

PROBLEM --> |input|C2
PROBLEM --> |input|C3
ENTRY_POINT --> |entry_point|C1
ENTRY_POINT --> |entry_point|C2
ENTRY_POINT --> |entry_point|C3
C1 --> |solution|ENSEMBLE
C2 --> |solution|ENSEMBLE
C3 --> |solution|ENSEMBLE
ENSEMBLE --> |combined_solution|T
T --> |fail|C4
T --> |pass|RETURN
C4 --> RETURN

```

Python Code

```

from typing import Literal
import workspace.HumanEval.workflows.template.operator as operator
import workspace.HumanEval.workflows.round_5.prompt as prompt_custom
from scripts.async_llm import create_llm_instance
import weave

DatasetType = Literal["HumanEval", "MBPP", "GSM8K", "MATH", "HotpotQA", "DROP"]

class Workflow:
    def __init__(
        self,
        name: str,
        llm_config,
        dataset: DatasetType,
    ) -> None:
        self.name = name
        self.dataset = dataset
        self.llm = create_llm_instance(llm_config)
        self.custom = operator.Custom(self.llm)
        self.custom_code_generate = operator.CustomCodeGenerate(self.llm)
        self.sc_ensemble = operator.ScEnsemble(self.llm)
        self.test = operator.Test(self.llm)

    async def __call__(self, problem: str, entry_point: str):
        """
        Implementation of the workflow
        Custom operator to generate multiple solutions and select the best one.
        """
        # Generate initial solutions using custom code generation
        solution_response_1 = await self.custom_code_generate(problem=problem, entry_point=entry_point,
        ↪ instruction=prompt_custom.CODE_GENERATE_PROMPT)
        solution_1 = solution_response_1['response']

        solution_response_2 = await self.custom_code_generate(problem=problem, entry_point=entry_point,
        ↪ instruction=prompt_custom.CODE_GENERATE_PROMPT)
        solution_2 = solution_response_2['response']

        # Generate an optimized solution
        optimized_solution_response = await self.custom_code_generate(problem=problem,
        ↪ entry_point=entry_point, instruction=prompt_custom.OPTIMIZED_CODE_GENERATE_PROMPT)
        optimized_solution = optimized_solution_response['response']

        # Combine solutions using ensemble method
        ensemble_response = await self.sc_ensemble(solutions=[solution_1, solution_2, optimized_solution],
        ↪ problem=problem)
        combined_solution = ensemble_response['response']

        # Test the combined solution
        test_result = await self.test(problem=problem, solution=combined_solution, entry_point=entry_point)

        # If the solution fails the test, improve the code
        if not test_result['result']:
            improved_solution_response = await self.custom(input=problem,
            ↪ instruction=prompt_custom.IMPROVE_CODE_PROMPT)
            improved_solution = improved_solution_response['response']
            # Test the improved solution
            test_result = await self.test(problem=problem, solution=improved_solution,
            ↪ entry_point=entry_point)
            combined_solution = improved_solution if test_result['result'] else combined_solution # Use
            ↪ improved solution if it passes tests

```

```
return combined_solution, self.llm.get_usage_summary()["total_cost"]
```

Corresponding Prompt

```
IMPROVE_CODE_PROMPT = """
The previous solution failed some test cases in the HumanEval benchmark. Please conduct a thorough analysis
↳ of the problem statement, identifying all edge cases and potential pitfalls. Then, provide an improved
↳ solution that not only fixes the issues but also optimizes performance and adheres to industry-standard
↳ coding practices. Ensure your revised code includes clear, concise comments that explain your logic and
↳ design choices, and that it robustly handles all specified requirements.
"""

CODE_GENERATE_PROMPT = """
Generate a Python function to solve the given problem. Ensure the function name matches the one specified in
↳ the problem. Include necessary imports. Use clear variable names and add comments for clarity.

Problem:
{problem}

Function signature:
{entry_point}

Generate the complete function below:
"""

OPTIMIZED_CODE_GENERATE_PROMPT = """
Based on previous attempts and their outcomes, generate an optimized Python function to solve the given
↳ problem. Focus on improving efficiency, readability, and robustness. Ensure the function name matches
↳ the one specified in the problem, include necessary imports, and use clear variable names with comments
↳ for clarity.

Problem:
{problem}

Function signature:
{entry_point}

Generate the complete function below:
"""

ENSEMBLE_PROMPT = """
You have multiple solutions generated for the problem. Your task is to analyze these solutions and select
↳ the one that is most likely to be correct based on their similarities and performance. Ensure that the
↳ selected solution is robust and handles all edge cases effectively.
"""
```

B.4. MBPP

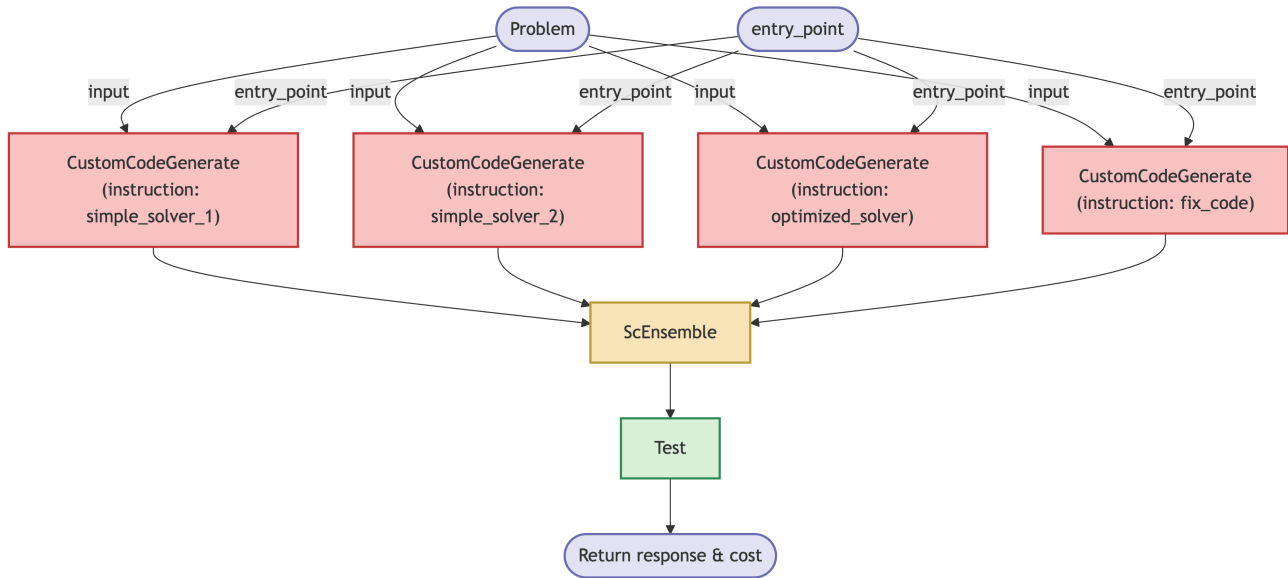


Figure 8. Mermaid diagram for MBPP.

Mermaid Code

```

flowchart TD
%% Nodes
PROBLEM([Problem])
ENTRY_POINT([entry_point])
C1["CustomCodeGenerate<br/>(instruction: simple_solver_1)"]
C2["CustomCodeGenerate<br/>(instruction: simple_solver_2)"]
C3["CustomCodeGenerate<br/>(instruction: optimized_solver)"]
C4["CustomCodeGenerate<br/>(instruction: fix_code)"]
ENSEMBLE["ScEnsemble<br/>"]
T["Test<br/>"]
RETURN([Return response & cost])

%% Styles
classDef CustomOp fill:#d0e1f9,stroke:#4378a2,stroke-width:2px;
classDef CustomCodeGenerateOp fill:#f9c2c2,stroke:#c23737,stroke-width:2px;
classDef ScEnsembleOp fill:#f9e4b7,stroke:#b99b37,stroke-width:2px;
classDef DecisionOp fill:#ffffff,stroke:#444444,stroke-width:1px,stroke-dasharray:2 2;
classDef TestOp fill:#d8f0d8,stroke:#2e8b57,stroke-width:2px;
classDef Interface fill:#e2e2f2,stroke:#6a6ab2,stroke-width:2px;

%% Assign classes
class PROBLEM Interface
class ENTRY_POINT Interface
class C1 CustomCodeGenerateOp
class C2 CustomCodeGenerateOp
class C3 CustomCodeGenerateOp
class C4 CustomCodeGenerateOp
class ENSEMBLE ScEnsembleOp
class T TestOp
class RETURN Interface

%% Flow (arrows show data relationships)
    
```

```

PROBLEM --> |input|C1
PROBLEM --> |input|C2
PROBLEM --> |input|C3
PROBLEM --> |input|C4
ENTRY_POINT --> |entry_point|C1
ENTRY_POINT --> |entry_point|C2
ENTRY_POINT --> |entry_point|C3
ENTRY_POINT --> |entry_point|C4
C1 --> ENSEMBLE
C2 --> ENSEMBLE
C3 --> ENSEMBLE
C4 --> ENSEMBLE
ENSEMBLE --> T
T --> RETURN

```

Python Code

```

from typing import Literal
import workspace.MBPP.workflows.template.operator as operator
import workspace.MBPP.workflows.round.8.prompt as prompt_custom
from scripts.async_llm import create_llm_instance
import weave

DatasetType = Literal["HumanEval", "MBPP", "GSM8K", "MATH", "HotpotQA", "DROP"]

class Workflow:
    def __init__(
        self,
        name: str,
        llm_config,
        dataset: DatasetType,
    ) -> None:
        self.name = name
        self.dataset = dataset
        self.llm = create_llm_instance(llm_config)
        self.custom = operator.Custom(self.llm)
        self.custom_code_generate = operator.CustomCodeGenerate(self.llm)
        self.sc_ensemble = operator.ScEnsemble(self.llm)
        self.test = operator.Test(self.llm)

@weave.op()
async def __call__(self, problem: str, entry_point: str):
    """
    Implementation of the workflow
    Custom operator to generate anything you want.
    But when you want to get standard code, you should use custom_code_generate operator.
    """
    # Generate two solutions using different instructions
    solution_1 = await self.custom_code_generate(problem=problem, entry_point=entry_point,
        ↪ instruction=prompt_custom.CODE_GENERATE_PROMPT)
    solution_2 = await self.custom_code_generate(problem=problem, entry_point=entry_point,
        ↪ instruction=prompt_custom.CODE_GENERATE_PROMPT)
    # Generate an optimized solution
    optimized_solution = await self.custom_code_generate(problem=problem, entry_point=entry_point,
        ↪ instruction=prompt_custom.OPTIMIZED_CODE_PROMPT)
    # Attempt to fix the code if necessary
    fix_attempt = await self.custom_code_generate(problem=problem, entry_point=entry_point,
        ↪ instruction=prompt_custom.FIX_CODE_PROMPT)

    # Use ensemble to select the best solution
    ensemble_response = await self.sc_ensemble(solutions=[solution_1['response'], solution_2['response']],
        ↪ optimized_solution['response'], fix_attempt['response'], problem=problem)

    # Test the selected solution
    test_response = await self.test(problem=problem, solution=ensemble_response['response'],
        ↪ entry_point=entry_point)

    # Return the final response and cost
    return test_response['solution'], self.llm.get_usage_summary()["total_cost"]

```

Corresponding Prompt

```
CODE_GENERATE_PROMPT = """
Generate a Python function to solve the given problem. Ensure the function name matches the one specified in
↳ the problem. Include necessary imports. Use clear variable names and add comments for clarity.

Problem:
{problem}

Function signature:
{entry_point}

Generate the complete function below:
"""

OPTIMIZED_CODE_PROMPT = """
Based on previous attempts, generate an optimized Python function to solve the given problem. Ensure the
↳ function name matches the one specified in the problem. Include necessary imports, and focus on
↳ improving performance and readability. Use clear variable names and add comments for clarity.

Problem:
{problem}

Function signature:
{entry_point}

Generate the complete function below:
"""

FIX_CODE_PROMPT = """
The provided solution failed to pass the tests. Please analyze the error and fix the code. Ensure the
↳ function name and signature remain unchanged. If necessary, add or modify imports, correct logical
↳ errors, and improve the implementation.

Problem:
{input}

Provide the corrected function below:
"""
```