

MASFactory: A Graph-Centric Framework for Orchestrating LLM-Based Multi-Agent Systems with Vibe Graphing

Yang Liu¹, Jinxuan Cai¹, Yishen Li¹, Qi Meng¹, Zedi Liu¹,
Xin Li¹, Chen Qian², Chuan Shi¹ and Cheng Yang^{1*}

¹Beijing University of Posts and Telecommunications

²Shanghai Jiao Tong University

{liuyang1999, yangcheng}@bupt.edu.cn

Abstract

Large language model-based (LLM-based) multi-agent systems (MAS) are increasingly used to extend agentic problem solving via role specialization and collaboration. MAS workflows can be naturally modeled as directed computation graphs, where nodes execute agents/sub-workflows and edges encode dependencies and message passing. However, implementing complex graph workflows in current frameworks still requires substantial manual effort, offers limited reuse, and makes it difficult to integrate heterogeneous external context sources. To overcome these limitations, we present MASFactory, a graph-centric framework for orchestrating LLM-based MAS. It introduces Vibe Graphing, a human-in-the-loop approach that compiles natural-language intent into an editable workflow specification and then into an executable graph. In addition, the framework provides reusable components, skill support, multimodal message handling, and pluggable context integration, as well as a visualizer for topology preview, runtime tracing, and human-in-the-loop interaction. We evaluate MASFactory on seven public benchmarks, validating both reproduction consistency for representative MAS methods and the effectiveness of Vibe Graphing. Our code¹ and video² are publicly available.

1 Introduction

With the introduction of tool use and feedback-driven correction, large language models (LLMs) are often packaged as agents that follow a *perception–reasoning–action* loop, enabling long-horizon task execution in external environments (Yao et al., 2023; Schick et al., 2023; Shinn et al., 2023). As tasks grow in complexity, a

* Corresponding author.

¹<https://github.com/BUPT-GAMMA/MASFactory>. Licensed under Apache-2.0; use, modification, and distribution are permitted within its terms.

²<https://youtu.be/ANynzVfy32k>

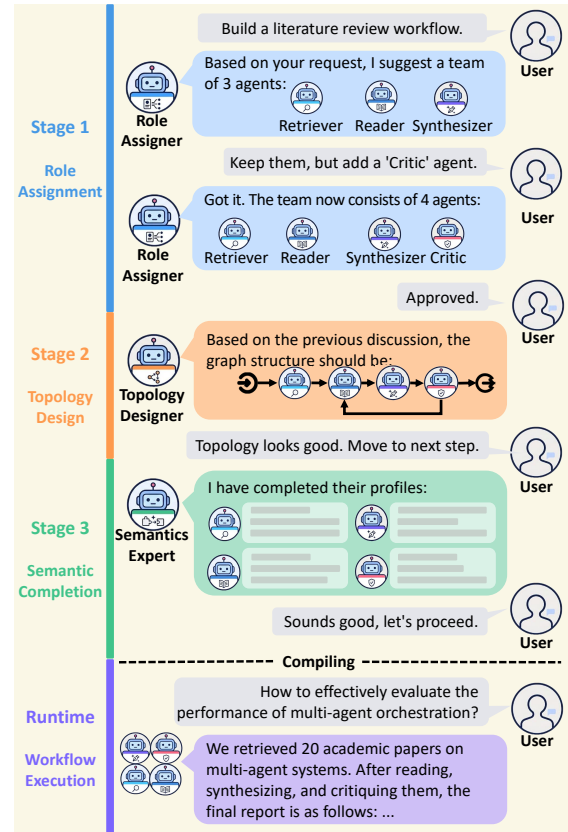


Figure 1: **Vibe Graphing in MASFactory.** MASFactory turns a user’s natural-language intent into an executable multi-agent workflow via a three-stage, human-in-the-loop process, then compiles and executes the resulting workflow at runtime.

single agent often struggles to simultaneously achieve end-to-end coverage, robust error recovery, and multi-step coordination. Consequently, LLM-based multi-agent systems (MAS) have become a widely adopted approach to extend agentic problem solving through role specialization, cross-checking, and iterative collaboration (Li et al., 2023; Wu et al., 2023; Hong et al., 2024; Qian et al., 2024; Chen et al., 2024).

For MAS modeling, an increasingly common orchestration abstraction is directed computation

graphs. A graph consists of nodes that host computation units (agents, tools, or sub-workflows) and edges that encode execution dependencies and message-passing directions. Pure directed acyclic graph (DAG) workflows are well-suited for pipeline-style collaboration, while cyclic structures are crucial for iterative processes such as reflection, revision, and retry. Recent frameworks have begun to explicitly represent multi-agent workflows as graphs. For instance, LangGraph models workflow as a *stateful graph* with explicit execution semantics (LangChain, 2024), and Dify provides a workflow canvas where DAG dependencies define both execution order and data flow (LangGenius, 2024).

However, implementing complex MAS remains engineering-intensive. For example, developers have to manually craft role prompts for agent nodes, wire the routing logic between nodes, and establish inter-agent communication protocols. Also, real-world applications rely on heterogeneous context sources, including reusable skill packages for task-specific guidance and auxiliary resources (Anthropic, 2025a), memory layers (Chhikara et al., 2025; Kang et al., 2025; Packer et al., 2023), retrieval-augmented generation (RAG) (Lewis et al., 2020; Edge et al., 2024), and standardized tool/context integration protocols such as the model context protocol (MCP) (Anthropic, 2025b). Current frameworks often integrate them via workflow-specific glue code, making implementation difficult to port and reuse across environments. Finally, MAS development frequently involves repeated subgraphs that are globally similar but differ slightly in local configurations. Existing frameworks provide limited support for version-controlled, templated reuse of such subgraphs.

To address these challenges, we present **MASFactory**, a composable orchestration framework for multi-agent systems with **Vibe Graphing** (illustrated in Figure 1). Firstly, MASFactory introduces Vibe Graphing: users describe design intent in natural language, and the system compiles it into a readable, editable, and version-controlled structured intermediate representation, which is further compiled into an executable workflow. This process incorporates *human-in-the-loop* interaction, allowing users to review, modify, and provide feedback through a visual interface. Secondly, MASFactory improves interoperability via a Context Adapter that hides heterogeneity across different

context sources, and further supports multimodal message handling by representing image and PDF inputs as structured media assets. Thirdly, MASFactory improves reusability by providing configurable, reusable graph modules for common multi-agent collaboration patterns, enabling developers to quickly assemble workflows without repeatedly implementing structurally similar subgraphs. Beyond these, MASFactory provides a visualizer that supports visual inspection of workflow topology, runtime trace visualization, and human-in-the-loop interaction.

Contributions. (1) We propose **MASFactory**, a graph-centric MAS orchestration framework with reusable components, skill support, multimodal message handling, pluggable context management, delivering both natural-language workflow generation and consistent reproducible results. (2) We introduce **Vibe Graphing**, a human-in-the-loop approach that compiles natural-language intent into executable graphs, reducing implementation effort while achieving competitive performance with manually implemented workflows. (3) We reproduce five representative MASs with MASFactory, and demonstrate competitive performance on seven public benchmarks.

2 Related Work

LLM-based Multi-Agent Systems. Recent work has demonstrated that coordinating multiple LLM agents can improve robustness and coverage via role specialization, cross-checking, and iterative collaboration (Li et al., 2023; Chen et al., 2024; Shen et al., 2023). For example, AutoGen popularizes programmable multi-agent conversation patterns and group coordination (Wu et al., 2023). MetaGPT structures collaboration with role definitions and SOP-like procedures for complex tasks (Hong et al., 2024). ChatDev organizes software-development collaboration into staged role interactions (Qian et al., 2024). However, the implementations of these works vary significantly, and their codebases often span from thousands to tens of thousands of lines, making it difficult for developers to build upon them.

Frameworks for MAS Development. Beyond research prototypes, a growing ecosystem targets practical MAS implementation. Google’s Agent Development Kit (ADK) (Google, 2024) and CrewAI (CrewAI Inc., 2024) emphasize a code-

first, process-driven methodology: the former provides a toolkit towards deploying agentic applications, whereas the latter focuses on the programmatic orchestration of collaborative agent crews. Alternatively, graph-centric frameworks such as LangGraph (LangChain, 2024) and Dify (LangGenius, 2024) represent agentic workflows as explicit node/edge structures to provide fine-grained control over the execution flow. However, building complex MAS remains engineering-intensive, requiring heavy manual configuration and tightly coupled context integrations.

To simplify orchestration, MASFactory complements graph-centric frameworks such as LangGraph and Dify by focusing on reusable and intent-driven MAS construction. Unlike frameworks that mainly emphasize execution control or manual workflow editing, MASFactory represents workflows as editable graph specifications that can be generated via Vibe Graphing, inspected through the visualizer, and instantiated through reusable components. This reduces the effort of constructing, adapting, and reusing MAS workflows.

3 System Design

In this section, we present the architecture of MASFactory. As shown in Figure 2, at the bottom layer the system adopts a backbone composed of Node and Edge for modeling collaboration graphs. Built on top of this backbone, MASFactory provides composable components and modules for flexible collaboration process implementation. Besides, MASFactory unifies MAS-critical mechanisms, including communication protocols, skill packaging, multimodal input handling, and context management, and builds them as pluggable modules. As a result, the system can conveniently integrate external frameworks such as Mem0 (Chhikara et al., 2025) and LlamaIndex (LlamaIndex, 2024) for better user experience. MASFactory also provides three orchestration interfaces: automated orchestration via Vibe Graphing, and manual workflow implementation via declarative or imperative programming. To facilitate development and debugging, MASFactory also provides a visualizer for topology preview, runtime tracing, and human-in-the-loop interactions throughout the workflow.

3.1 Basic Components

Graph-based Organization. MASFactory models collaboration as a directed graph composed of Node

and Edge (Zhuge et al., 2024). A Node is the basic computation unit and can be extended into freely composable components, including Graph, Loop, Agent, CustomNode, Interaction, and Switch. These components expose a consistent interface, but differ in their execution logic. An Edge expresses inter-node dependencies and serves as the carrier of message passing.

Collaboration Flow. MASFactory makes collaboration signals explicit by separating them into three flows. **Control flow** propagates along edges to advance scheduling and dependencies, ensuring causal constraints of execution. **Message flow** propagates horizontally along edges to carry node outputs to downstream nodes. **State flow** propagates along the hierarchy between graphs and sub-graphs to synchronize graph-level context and runtime state.

Multimodal Message Representation. MASFactory extends message flow beyond plain text by representing images, PDFs, and textual content as structured message fields with explicit modality declarations. This design allows multimodal inputs and intermediate outputs to be propagated through the same graph-level message-passing mechanism, while leaving model-specific serialization to downstream adapters.

Node Lifecycle. At runtime, each node follows a unified lifecycle. A node first aggregates incoming messages along its in edges into its input, then reads states from its parent Graph. It subsequently runs node-specific execution logic, e.g., invoking the LLM for Agent, scheduling internal nodes in topological order for Graph and Loop. After execution, the node dispatches produced messages via its out edges to downstream nodes, and finally writes updated states back to the parent Graph to synchronize shared states and support downstream execution.

Runtime Scheduling. MASFactory adopts a readiness-based scheduling strategy that allows multiple ready nodes to execute concurrently. This unified mechanism supports sequential, parallel, branching, and cyclic control structures.

3.2 Key Components

Graph and Loop. Graph and Loop are responsible for topology representation and scheduling of internal nodes. Graph expresses and schedules directed acyclic graph (DAG) workflows. Loop expresses and schedules cyclic structures, and is commonly used for iterative collaboration patterns such as

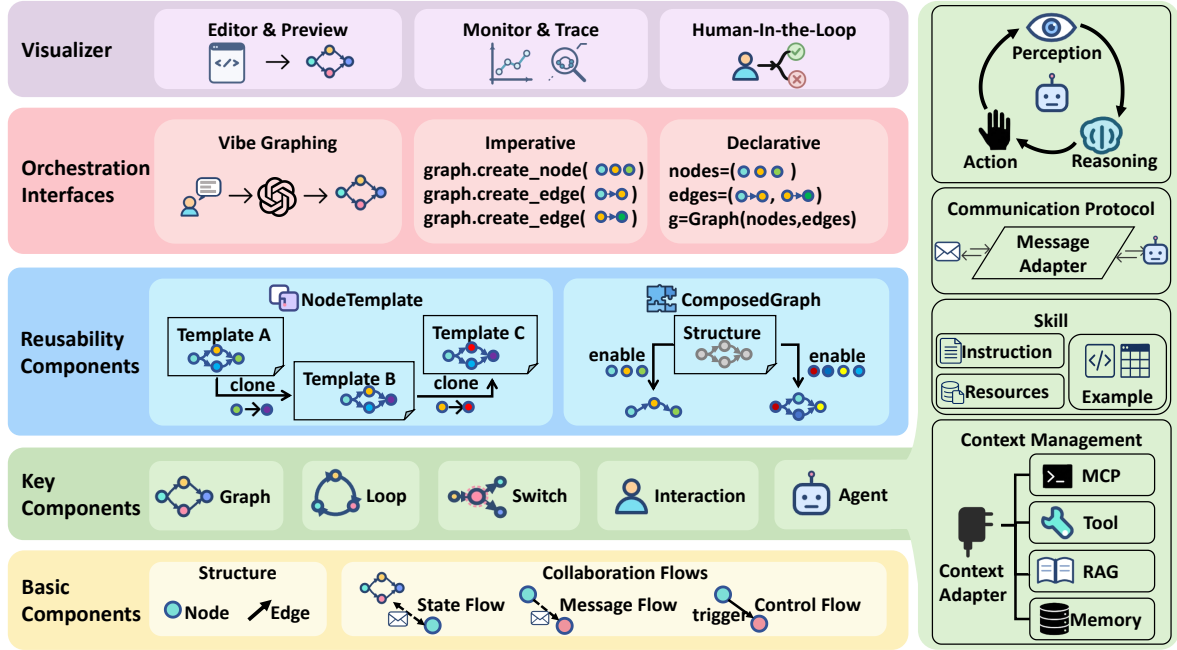


Figure 2: Architecture overview of MASFactory framework.

reflection, revision, and retry.

Switch. A Switch implements control-flow routing inside the graph. Unlike ordinary nodes that broadcast messages and trigger signals to all downstream nodes by default, Switch dynamically selects and activates one or multiple downstream paths based on runtime state.

Interaction. An Interaction node serves as the entry point of the human-in-the-loop mechanism. It can actively query users during execution, collect feedback, and inject user inputs back into the workflow.

3.3 Agent Component

Perception–Reasoning–Action. Agent adopts the classic Perception–Reasoning–Action paradigm and modularizes key steps such as communication and context through pluggable designs. As indicated in Figure 2, Agent relies on a pluggable *Message Adapter* for message processing and a pluggable *Context Adapter* for context management.

Message Adapter. For communication, the Message Adapter formats agent inputs and outputs according to a given communication protocol.

MASFactory provides a set of commonly used adapters, including protocols based on JSON schema, structured Markdown segments, and plaintext paragraph formats. It also exposes interfaces for user-defined protocols. By introducing Message Adapters, MASFactory decouples collabora-

tion graph from protocol definitions, enabling protocol extension or replacement without modifying the collaboration topology.

Context Adapter. For context management, to shield heterogeneity across information sources such as Memory, MCP, and RAG (Packer et al., 2023; Kang et al., 2025; Anthropic, 2025b; Lewis et al., 2020; Edge et al., 2024), MASFactory utilizes a Context Adapter to provide a standardized interface. This adapter segments diverse external contexts into standardized units, thereby enabling unified seamless integration with different context frameworks such as Mem0 and LlamaIndex.

Skill Support. Inspired by recent Agent Skills practice, MASFactory supports lightweight skill packages that encapsulate reusable task guidance and auxiliary resources, allowing agents to reuse task-specific capabilities without changing the workflow topology.

3.4 Reusability and Templated Prototyping

NodeTemplate. NodeTemplate allows users to declare a structural template first and instantiate it into a concrete graph later. This decoupling between declaration and instantiation enables users to clone templates to build multiple graphs that share a similar global structure but differ in local configurations, supporting branch-style reuse and versioned management. Moreover, developers can reuse a graph-level template while adjusting node-

level settings when instantiating the graph.

ComposedGraph. ComposedGraph is a specialized form of Graph that represents a class of predefined structures. It can instantiate a concrete graph by filling node configurations or activating specific branches according to user parameters, hiding low-level construction details. With ComposedGraph, users can package their designs as reusable composite components. Meanwhile, MASFactory uses ComposedGraph to encapsulate structures that are difficult to express with static graphs, such as DyLan-style dynamic scheduling patterns (Liu et al., 2024), as well as commonly used collaboration subgraphs. This supports broader application scenarios while reducing development cost.

3.5 Orchestration Interfaces

Vibe Graphing. Under Vibe Graphing, MASFactory compiles natural-language intent into an executable multi-agent workflow through staged compilation, while providing human-in-the-loop review and revision at each stage. During this process, the system produces a readable, editable and structured intermediate representation, which is then compiled into an executable workflow. As illustrated in Figure 1, the pipeline completes three core tasks. *Role Assignment* maps task intent into a set of candidate agents with clear responsibility boundaries. *Structure Design* generates a directed-graph topology skeleton based on inter-role information dependencies and control constraints, determining connectivity as well as the directions of message and control propagation. *Semantic Completion* performs parameterized instantiation over the skeleton by configuring prompts and tools for each node, producing a workflow that can be compiled and executed directly. This intent–structure–instantiation compilation chain elevates MAS construction from manual workflow configuration to an iterative design process, reducing implementation cost while retaining control over topology and semantics.

Imperative Interface. The imperative interface is code-centric: developers construct graphs by programmatically instantiating nodes and edges and wiring them with explicit control logic. This style offers high flexibility and precise control over topology, parameters, and runtime behaviors, making it suitable for carefully engineered workflows and scenarios that require tight coupling with application logic.

Declarative Interface. The declarative interface specifies a workflow as a structured configu-

ration: developers declare the graph topology and node properties, and MASFactory constructs the executable graph accordingly. This style keeps programming concise and easy to review, and is recommended for fixed workflows and lightly dynamic structures.

3.6 Visualizer

The Visualizer is a visual integrated environment implemented as a VS Code extension. It aligns static workflow topology with runtime traces in a single view, and supports the following features:

Editor & Preview. Editor & Preview supports real-time topology preview and structural inspection during development.

Monitor & Trace. The visualizer tracks node state evolution and message propagation during execution to support debugging and diagnosis.

Human-in-the-Loop. Human-in-the-Loop works with Interaction nodes to visualize runtime user interactions and to incorporate external feedback or inputs into the Vibe Graphing workflow, enabling interactive intervention and iteration.

4 Evaluation and Analysis

We conduct extensive experiments to answer the following questions: (1) whether MASFactory can reproduce representative MAS methods with consistent effectiveness; (2) whether Vibe Graphing can produce competitive workflows compared with manually designed ones; and (3) whether MASFactory can reduce implementation cost through component reuse and intent-driven orchestration.

4.1 Experimental Setup

Benchmarks. We evaluate on coding-oriented benchmarks (HumanEval, MBPP, BigCodeBench, SRDD) (Chen et al., 2021; Austin et al., 2021; Zhuo et al., 2025; Qian et al., 2024), and general reasoning/tool-use benchmarks (MMLU-Pro, GAIA, GPQA) (Wang et al., 2024; Mialon et al., 2024; Rein et al., 2024). All scores are reported in percentage scale (0–100). Coding benchmarks except SRDD report pass@1; SRDD reports an aggregate quality score (SRDD); MMLU-Pro/GAIA/GPQA report accuracy.

LLM Backbones. For Vibe Graphing, the *workflow construction* stage uses gpt-5.2 to compile natural-language intent into an executable workflow. For workflow execution of all methods, we use gpt-4o-mini.

Table 1: Main results are reported on a percentage scale (0-100). The symbol “-” indicates not applicable due to the incompatibility between programming-focused MASs and general-purpose reasoning benchmarks.

Method	HumanEval	MBPP	BigCodeBench	SRDD	MMLU-Pro	GAIA	GPQA
ChatDev (original)	82.50	71.40	50.70	82.91	-	-	-
ChatDev (MASFactory)	81.30	74.20	53.30	84.23	-	-	-
MetaGPT (original)	67.07	36.03	50.10	78.19	-	-	-
MetaGPT (MASFactory)	89.02	59.14	51.70	72.77	-	-	-
AgentVerse (original)	85.00	74.54	65.92	87.55	64.64	12.12	38.39
AgentVerse (MASFactory)	85.00	75.15	64.12	91.06	64.16	12.73	37.50
CAMEL (original)	62.20	60.60	63.51	89.42	50.08	9.70	32.59
CAMEL (MASFactory)	71.85	57.80	78.16	89.69	63.04	12.73	24.78
HuggingGPT (original)	82.32	68.60	28.42	87.96	65.59	9.09	56.67
HuggingGPT (MASFactory)	80.49	64.40	29.91	83.26	63.66	10.91	47.32
Vibe Graphing-ChatDev	83.50	74.20	45.30	88.13	-	-	-
Vibe Graphing-Task Specific	84.76	72.37	51.67	90.71	51.73	12.12	39.51

4.2 Performance of Reproduced Workflows

Table 1 compares five representative MAS methods implemented in MASFactory against their original implementations, including ChatDev (Qian et al., 2024), MetaGPT (Hong et al., 2024), AgentVerse (Chen et al., 2024), CAMEL (Li et al., 2023), and HuggingGPT (Shen et al., 2023). Overall, MASFactory reproductions achieve results broadly consistent or even better with the originals across different benchmarks. This indicates that MASFactory can cover diverse multi-agent architectures and collaboration designs without systematic regressions.

4.3 Performance of Vibe Graphing

We report two Vibe Graphing settings in Table 1. In *Vibe Graphing-ChatDev*, we replace each key phase of ChatDev with an individual VibeGraph component, and connect these phases into a workflow using lightweight glue code. In contrast, *Vibe Graphing-Task Specific* is a task-driven variant: for each benchmark dataset, a developer writes a natural language description of a workflow, and passes the instruction to a single VibeGraph component to compile it into an executable workflow. The last two rows of Table 1 reports the results for workflows produced with Vibe Graphing. Despite using a different model only at the workflow construction stage (gpt-5.2), the resulting workflows execute with gpt-4o-mini and achieve competitive performance on coding benchmarks, as well as on general reasoning/tool-use tasks where applicable. These results suggest that staged intent-to-graph compilation can generate viable multi-agent workflows

Table 2: Comparison of monetary cost between Vibe Graphing (VG) and Vibe Coding (VC). VC-L and VC-M denote the low and medium reasoning lengths of VC.

Metric	ChatDev			AgentVerse		
	VG	VC-L	VC-M	VG	VC-L	VC-M
Cost (\$)	0.26	3.49	3.02	0.59	4.43	6.08

that approach manually designed baselines, while substantially reducing low-level graph wiring and framework-specific engineering effort.

4.4 Case Study of Implementation Cost

The original implementation of ChatDev contains 1,511 lines of Python code for workflow definition. With ComposedGraph-based reuse, our MASFactory reproduction reduces the total implementation to 1,114 lines while maintaining comparable performance as shown in Table 1. If we further use Vibe Graphing to make each ChatDev stage interactively generated (i.e., Vibe Graphing-ChatDev), the implementation only requires 203 lines to connect these stages. When we fully rely on Vibe Graphing to generate the workflow end-to-end (i.e., Vibe Graphing-Task Specific), the workflow specification shrinks to only 45 lines of code.

Furthermore, under the same gpt-5.2 backend, Vibe Graphing reduces API costs by roughly an order of magnitude compared to Vibe Coding, as shown in Table 2. Workflows generated via Vibe Coding frequently exhibit logical flaws in the constructed graphs, failing to return correct execution results. Consequently, we restrict our comparison

with Vibe Coding to cost analysis and exclude it from the performance evaluation.

To summarize, composed components enhance reusability, while Vibe Graphing significantly lowers development barriers by minimizing boilerplate and manual configuration. This reduction is particularly important for MAS research prototyping, where developers often need to compare multiple collaboration patterns and repeatedly modify role definitions, communication paths, and context configurations. By combining declarative specifications with reusable composed components, MASFactory shifts effort from writing orchestration code to designing and refining the collaboration logic itself. We provide a Vibe Graphing case study in Appendix A, which walks through the workflow construction pipeline from natural-language intent to a structured workflow specification and executable workflow definitions.

5 Conclusion

We present MASFactory, a graph-centric framework for orchestrating LLM-based multi-agent systems. MASFactory models multi-agent workflows as executable directed graphs and integrates Vibe Graphing to compile natural-language intent into MAS workflows with human-in-the-loop refinement. It also provides reusable components, plug-gable context and a visualizer for topology preview, runtime tracing, and visual human-in-the-loop interaction. Experiments on seven public benchmarks show that MASFactory consistently reproduces representative MASs and that Vibe Graphing produces competitive workflows while substantially reducing implementation overhead.

Limitations

MASFactory currently does not provide built-in checkpointing for resuming execution from intermediate states after interruptions. In addition, we will continue to enrich and refine the built-in composed-component library in future updates.

Acknowledgments

This work is supported in part by the National Natural Science Foundation of China (No. 62550138, 62192784, 62572064, 62472329), and the Beijing Natural Science Foundation (No. 253004).

References

- Anthropic. 2025a. Agent skills. <https://platform.claude.com/docs/en/agents-and-tools/agent-skills/overview>. Accessed: 2026-05-15.
- Anthropic. 2025b. *Model context protocol (MCP) specification*. Accessed: 2026-02-23.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. 2021. *Program synthesis with large language models*. *arXiv preprint arXiv:2108.07732*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, and 39 others. 2021. *Evaluating large language models trained on code*. *Preprint*, arXiv:2107.03374.
- Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. 2024. *AgentVerse: Facilitating multi-agent collaboration and exploring emergent behaviors*. In *The Twelfth International Conference on Learning Representations (ICLR)*. OpenReview.net.
- Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. 2025. *Mem0: Building production-ready AI agents with scalable long-term memory*. *Preprint*, arXiv:2504.19413.
- CrewAI Inc. 2024. *CrewAI*. Accessed: 2026-02-23.
- Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitanaky, Robert Osazuwa Ness, and Jonathan Larson. 2024. *From local to global: A graph RAG approach to query-focused summarization*. *Preprint*, arXiv:2404.16130.
- Google. 2024. *Agent development kit (ADK)*. Accessed: 2026-02-23.
- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. 2024. *MetaGPT: Meta programming for a multi-agent collaborative framework*. In *The Twelfth International Conference on Learning Representations (ICLR)*. OpenReview.net.
- Jiazheng Kang, Mingming Ji, Zhe Zhao, and Ting Bai. 2025. *Memory OS of AI agent*. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- LangChain. 2024. *LangGraph*. Accessed: 2026-02-23.

- LangGenius. 2024. [Dify: An LLM app development platform](#). Accessed: 2026-02-23.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. [CAMEL: Communicative agents for "mind" exploration of large language model society](#). In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. 2024. [A dynamic LLM-powered agent network for task-oriented agent collaboration](#). In *First Conference on Language Modeling (COLM)*.
- LlamaIndex. 2024. [LlamaIndex](#). Accessed: 2026-02-23.
- Grégoire Mialon, Clémentine Fourrier, Thomas Wolf, Yann LeCun, and Thomas Scialom. 2024. [GAIA: a benchmark for general AI assistants](#). In *The Twelfth International Conference on Learning Representations (ICLR)*.
- Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. 2023. [MemGPT: Towards LLMs as operating systems](#). *Preprint*, arXiv:2310.08560.
- Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024. [ChatDev: Communicative agents for software development](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2024. [GPQA: A graduate-level google-proof q&a benchmark](#). In *First Conference on Language Modeling (COLM)*.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language models can teach themselves to use tools](#). In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. [Hugging-gpt: Solving ai tasks with chatgpt and its friends in hugging face](#). In *Advances in Neural Information Processing Systems*, pages 38154–38180. Curran Associates, Inc.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. [Reflection: Language agents with verbal reinforcement learning](#). In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. 2024. [MMLU-Pro: A more robust and challenging multi-task language understanding benchmark](#). *Preprint*, arXiv:2406.01574.
- Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and Chi Wang. 2023. [AutoGen: Enabling next-gen LLM applications via multi-agent conversation](#). *Preprint*, arXiv:2308.08155.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R. Narasimhan, and Yuan Cao. 2023. [ReAct: Synergizing reasoning and acting in language models](#). In *The Eleventh International Conference on Learning Representations (ICLR)*. OpenReview.net.
- Mingchen Zhuge, Wenyi Wang, Louis Kirsch, Francesco Faccio, Dmitrii Khizbullin, and Jürgen Schmidhuber. 2024. [GPTSwarm: Language agents as optimizable graphs](#). In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, pages 62743–62767.
- Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widayarsi, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, Simon Brunner, Chen Gong, Thong Hoang, Armel Randy Zebaze, Xiaoheng Hong, Wen-Ding Li, Jean Kadour, Ming Xu, Zhihan Zhang, and 14 others. 2025. [BigCodeBench: Benchmarking code generation with diverse function calls and complex instructions](#). In *The Thirteenth International Conference on Learning Representations (ICLR)*.

A Vibe Graphing Case Study

Here we present a concrete Vibe Graphing example that illustrates the end-to-end pipeline from natural-language intent to a structured workflow specification and executable workflow definitions.

A.1 Intent and Structural Constraint

We start from a natural-language build instruction that specifies both the task intent (weekly report writing) and an explicit structural constraint: $START \rightarrow A, B, C \rightarrow D \rightarrow END$, where A , B , and C are parallel drafting agents and D is an evaluator/selector agent that produces the final output.

The build instruction used in this case study is *Design a workflow for writing my weekly report. I will provide what I worked on this week at the beginning. Then run three agents in parallel to draft separate reports, and pass all drafts to a fourth agent to evaluate and select the best one as the final output. The expected workflow structure is: $START \rightarrow A, B, C \rightarrow D \rightarrow END$.*

A.2 Interactive Workflow Construction

As shown in Figure 5, we write a short program that constructs a minimal wrapper graph with the topology $ENTRY \rightarrow VibeGraph \rightarrow EXIT$. Here, `VibeGraph` is a built-in composed graph in `MAS-Factory` that activates the Vibe Graphing pipeline upon graph construction. As illustrated in Figure 1, the pipeline comprises three stages. Each stage is encapsulated by a `Loop` component and centers around an `Agent` node augmented with auxiliary mechanisms for correction, review, and human interaction through `Interaction`. After each stage produces an intermediate design, the system solicits user feedback and iteratively refines the result until the user accepts it, at which point the pipeline proceeds to the next stage. During the human-in-the-loop interactions in Stage 2 and Stage 3, users may either (i) edit the structured intermediate representation directly via the visualizer, or (ii) provide feedback in the interaction panel. Both manual edits and textual feedback are recorded and fed back to the `Agent` as references for subsequent revisions. Figure 3 shows an example interaction trace in the visualizer, and Figure 4(a) shows the workflow preview and editing interface.

A.3 Workflow Specification Snapshot

As shown in Figure 6, we provide a snapshot of the structured intermediate representation gener-

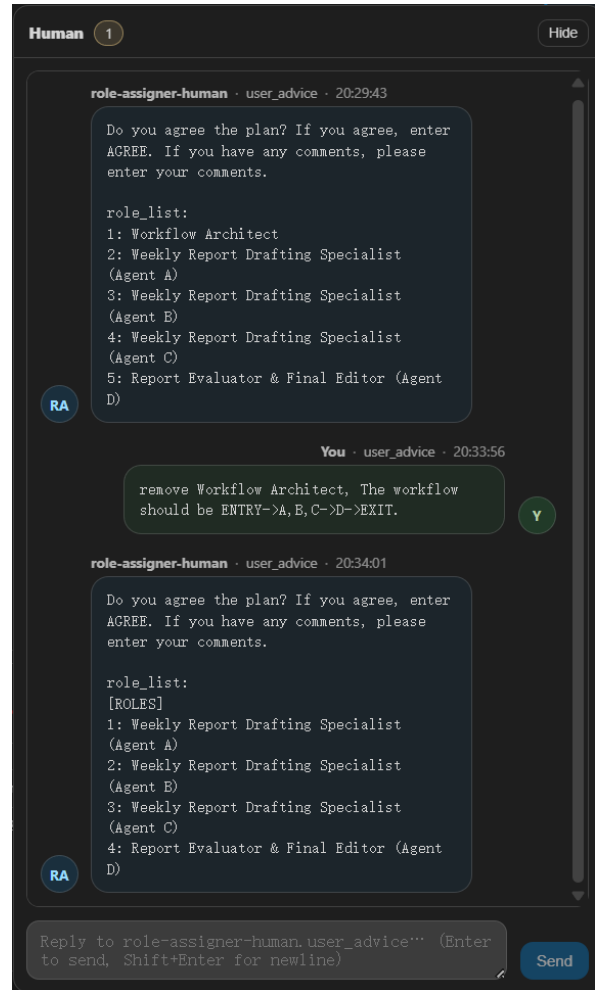
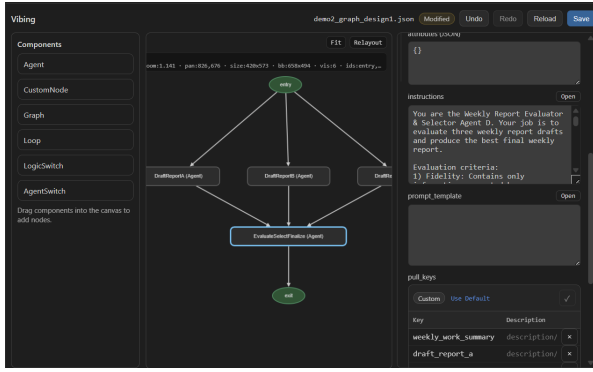


Figure 3: Human-in-the-loop interaction during Vibe Graphing in the visualizer. The user provides feedback to refine the role assignment and structure across stages.

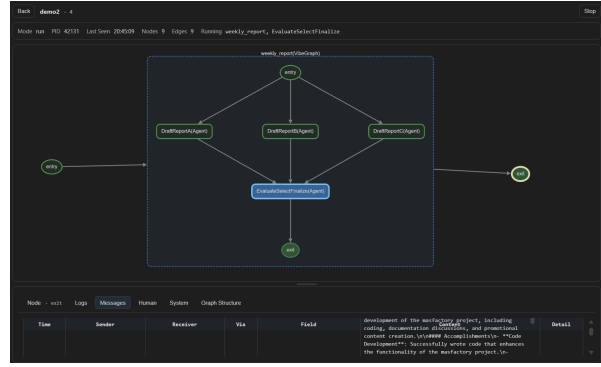
ated by Vibe Graphing. The representation explicitly encodes: (i) node semantics (e.g., label and instructions), (ii) input/output contracts (e.g., `input_fields` and `output_fields`), and (iii) directed dependencies among nodes (i.e., edges).

A.4 Workflow Execution

At runtime, the visualizer provides a synchronized view of execution states and message traces, as shown in Figure 4(b). The workflow first collects user input and normalizes it into `weekly_work_summary`, then executes three drafting branches in parallel, and finally aggregates drafts for selection and light editing.



(a) IR preview and editing in the MASFactory Visualizer.



(b) Runtime preview and trace alignment in the MASFactory Visualizer.

Figure 4: Visualizer views used in the Vibe Graphing.

```

from masfactory import VibeGraph, NodeTemplate, OpenAIModel, RootGraph
invoke_model = OpenAIModel(api_key=os.environ.get("OPENAI_API_KEY", ""),
    base_url=os.environ.get("OPENAI_BASE_URL", ""), model_name="gpt-4o-mini")
build_model = OpenAIModel(api_key=os.environ.get("OPENAI_API_KEY", ""),
    base_url=os.environ.get("OPENAI_BASE_URL", ""), model_name="gpt-5.2")
weekly_report = NodeTemplate(
    VibeGraph,
    invoke_model=invoke_model,
    build_model=build_model,
    build_instructions="... START->A,B,C->D->END.",
    build_cache_path=GRAPH_DESIGN_CACHE_PATH,
    pull_keys={"my_works": "what I worked on this week"},
    push_keys={"final_weekly_report": "final weekly report"},
)
root = RootGraph(name="demo2", nodes=[("weekly_report", weekly_report)],
    edges=[("ENTRY", "weekly_report", {}), ("weekly_report", "EXIT", {})])
root.build()
msg, attr = root.invoke({}, {"my_works": my_works})
print(attr["final_weekly_report"])

```

Figure 5: MASFactory programming code for the Vibe Graphing.

```

{
  "edges": [
    {"source": "ENTRY", "target": "DrafterA"}, {"source": "ENTRY", "target": "DrafterB"},
    {"source": "ENTRY", "target": "DrafterC"}, {"source": "DrafterA", "target": "Finalizer"},
    {"source": "DrafterB", "target": "Finalizer"}, {"source": "DrafterC", "target": "Finalizer"},
    {"source": "Finalizer", "target": "EXIT"}
  ],
  "nodes": [
    {"id": "DrafterA", "type": "Action", "input_fields": ["my_work"],
      "output_fields": ["draft_report_a"], "instructions": "You are Weekly Report Drafting Agent A ..."},
    {"id": "DrafterB", "type": "Action", "input_fields": ["my_work"],
      "output_fields": ["draft_report_b"], "instructions": "You are Weekly Report Drafting Agent B ..."},
    {"id": "DrafterC", "type": "Action", "input_fields": ["my_work"],
      "output_fields": ["draft_report_c"], "instructions": "You are Weekly Report Drafting Agent C ..."},
    {"id": "Finalizer", "type": "Action",
      "input_fields": ["my_work", "draft_report_a", "draft_report_b", "draft_report_c"],
      "output_fields": ["final_weekly_report", "selection_rationale"],
      "instructions": "You are the Weekly Report Evaluator ..."}
  ]
}

```

Figure 6: Workflow specification produced by Vibe Graphing.