

Augmented Runtime Collaboration for Self-Organizing Multi-Agent Systems: A Hybrid Bi-Criteria Routing Approach

Qingwen Yang^{1*†}, Feiyu Qu^{2*‡}, Tiezheng Guo¹, Yanyi Liu¹, Yingyou Wen¹

¹School of Computer Science and Engineering, Northeastern University, Shenyang 110819, China

²Department of Computer Science, Dartmouth College, Hanover, NH 03755, USA

2290182@stu.neu.edu.cn, feiyu.qu.gr@dartmouth.edu, 2310725@stu.neu.edu.cn, 2290175@stu.neu.edu.cn, wenyinyou@mail.neu.edu.cn

Abstract

LLM-based multi-agent systems have demonstrated significant capabilities across diverse domains. However, the task performance and efficiency are fundamentally constrained by their collaboration strategies. Prevailing approaches rely on static topologies and centralized global planning, a paradigm that limits their scalability and adaptability in open, decentralized networks. Effective collaboration planning in distributed systems using only local information thus remains a formidable challenge. To address this, we propose BiRouter, a novel dual-criteria routing method for Self-Organizing Multi-Agent Systems (SO-MAS). This method enables each agent to autonomously execute “next-hop” task routing at runtime, relying solely on local information. Its core decision-making mechanism is predicated on balancing two metrics: (1) the ImpScore, which evaluates a candidate agent’s long-term importance to the overall goal, and (2) the GapScore, which assesses its contextual continuity for the current task state. Furthermore, we introduce a dynamically updated reputation mechanism to bolster system robustness in untrustworthy environments and have developed a large-scale, cross-domain dataset, comprising thousands of annotated task-routing paths, to enhance the model’s generalization. Extensive experiments demonstrate that BiRouter achieves superior performance and token efficiency over existing baselines, while maintaining strong robustness and effectiveness in information-limited, decentralized, and untrustworthy settings.

1 Introduction

Recent breakthroughs in Large Language Models (LLMs) such as GPT-4 (Achiam et al. 2023) have unlocked unprece-

*These authors contributed equally.

† Author Contributions: Yang led the creation of the GapScore dataset and was responsible for training BiRouter.

‡ Author Contributions: Qu designed the overall network framework and the BiRouter forwarding mechanism, and contributed to building the ImpScore dataset. Qu and Yang jointly executed all experiments and collaborated on writing the manuscript.

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

ented capabilities in autonomous reasoning and planning (Wei et al. 2022; Wang et al. 2023; Yao et al. 2023b; Zhao et al. 2025). When embodied as agents (Li et al. 2023a; Sun et al. 2021), these models can comprehend complex instructions, decompose tasks, and execute actions to solve problems in diverse domains (Wang et al. 2024; Xi et al. 2023). This success has catalyzed a shift from single-agent systems to the more ambitious paradigm of Multi-Agent Systems (MAS) (Chan et al. 2024; Chen et al. 2025). By orchestrating agents with distinct roles and specialized skills, MAS can achieve a sophisticated division of labor, mirroring human societal structures to tackle tasks of far greater complexity (Tran et al. 2025). However, orchestrating this collaboration, especially in open and decentralized environments, remains a formidable scientific challenge.

The prevailing paradigm for MAS development has been the closed-world system, which relies on global optimization to coordinate agents (Zhang et al. 2025b). This approach inherently limits scalability and adaptability, prompting a community-wide effort towards open agent networks that integrate heterogeneous agents from different systems. Foundational protocols such as Agent-to-Agent (A2A) (Google 2025) and Agent Network Protocol (ANP) (Chang 2024) have emerged from this effort. However, these protocols provide mechanisms for agent discovery and invocation, but lack the strategies for dynamic, task-oriented collaboration. This absence leads to sub-optimal, or even chaotic, agent interactions, resulting in wasted computational resources, cascading failures, and an inability to reliably solve complex, multi-step tasks. Existing planning methods like GPTSwarm (Zhuge et al. 2024) and G-Designer (Zhang et al. 2025c) rely on static, global plans that pre-compute a complete execution path. Such an approach is not only brittle to dynamic changes but fundamentally incompatible with decentralized settings where no single agent possesses a global view. Furthermore, their generalization capabilities are often limited as open environments demand adaptability across a wide array of tasks and domains. This exposes a critical research gap: **How can an agent, relying solely on local in-**

formation, dynamically make “next hop” routing decisions that contribute to an emergent, globally efficient solution across diverse tasks?

To address this challenge, we situate our work within the framework of a **Self-Organizing Multi-Agent System (SO-MAS)**. We define this as a decentralized network where individual agents autonomously make task-forwarding decisions using only local information, obviating the need for a central scheduler. In such a paradigm, we introduce *BiRouter*, a novel routing method that empowers each agent to dynamically plan task pathways at runtime. Inspired by A* search algorithm (Hart, Nilsson, and Raphael 1968), *BiRouter* makes intelligent, localized routing decisions through a bi-criteria heuristic that evaluates potential next-hop agents. This heuristic balances two metrics: (1) an **Importance (Imp) score**, which serves as a long-term heuristic by quantifying an agent’s capability relevance to the overall task goal (analogous to A*’s heuristic cost, $h(n)$), and (2) a **Gap score**, which assesses an agent’s contextual continuity to bridge the gap between the task’s current state and the next logical step (akin to A*’s path cost, $g(n)$). By integrating these scores, *BiRouter* enables the emergence of globally coherent problem-solving pathways from simple, local rules. To enhance robustness in open, untrusted environments, the model further incorporates an iteratively updated agent reputation score into its decision-making process. Furthermore, to train our scoring functions and ensure their generalization, we constructed a large-scale, domain-rich dataset comprising thousands of annotated task-routing paths for diverse multi-step problems. Our main contributions can be summarized as follows:

- We introduce *BiRouter*, a novel bi-criteria routing method for dynamic and decentralized task planning that balances long-term relevance with short-term state-awareness.
- We develop and release a large-scale and diverse dataset specifically designed for training multi-agent collaboration models, providing a valuable resource for the research community.
- We provide extensive experimental validation demonstrating that *BiRouter* significantly improves collaboration efficiency and robustness in both centralized and distributed settings, outperforming established baselines.

2 Related Work

LLM-driven Multi-Agent Systems. The rise of LLMs has enabled Multi-Agent Systems (MAS) that outperform single agents on complex tasks (Li et al. 2023b; Chen et al. 2024a; Du et al. 2023; Liang et al. 2024). Early closed systems relied on predefined roles (Chen et al. 2024b; Wu et al. 2024) or aggregating outputs (Jiang, Ren, and Lin 2023; Wang et al. 2025), inherently limiting extensibility. To address this, the community has shifted towards open networks, developing foundational protocols like A2A (Google 2025), ANP (Chang 2024) and Agora (Marro et al. 2025) for dynamic discovery and decentralized communication. While these protocols establish the infrastructure for interoperability, they primarily solve for connectivity, leaving

a critical gap in the task-oriented, intelligent collaboration strategies needed for effective coordination.

Agent Collaboration Planning Mechanisms. Effective agent collaboration is crucial for task success. A dominant paradigm is static planning, which pre-computes a fixed workflow before task execution (Zhang et al. 2023; Zhuge et al. 2024). Methods include pruning connections (Agent-Prune (Zhang et al. 2025b)), generating task-specific structures (G-Designer (Zhang et al. 2025c), AFlow (Zhang et al. 2025d)), or determining a dependency-based order (MacNet (Qian et al. 2025)). However, these static approaches lack adaptability to dynamic changes or failures. To address this rigidity, dynamic approaches like DyLAN (Liu et al. 2024) and MaAS (Zhang et al. 2025a) adapt team structures or select executors during the task. Despite their flexibility, both static and dynamic methods typically rely on a centralized planner with a global view of the system. This centralized model is not applicable in large-scale, decentralized environments where information is inherently local. Our work addresses this gap by proposing a mechanism for agents to make effective runtime decisions using only local information.

Heuristic Search Methods. Planning a collaborative workflow can be framed as a pathfinding problem in a complex state space, making heuristic search highly relevant (Meng et al. 2024; Veerapaneni, Saleem, and Likhachev 2023). Recent methods have leveraged this, such as Tree-of-Thought (Yao et al. 2023a) using tree search for reasoning paths, and ToolChain* (Zhuang et al. 2024) and AFlow applying A* (Hart, Nilsson, and Raphael 1968) and MCTS (Browne et al. 2012) to optimize workflows. However, A*’s reliance on a static heuristic makes it ill-suited for dynamic environments, while the more adaptable MCTS incurs significant computational overhead from extensive simulations. Inspired by the goal-oriented nature of A*, we introduce *BiRouter*: a novel, hybrid bi-criteria routing method. It navigates these dynamic spaces by learning its own heuristic functions, enabling efficient and adaptive task routing.

3 Methodology

Figure 1 presents the architecture of our method. Our system operates as a self-organizing agent network where any agent can initiate a query. It then employs its built-in routing mechanism, *BiRouter*, to dynamically forward the task. This distributed routing process progressively approaches an optimal execution path, forming a runtime adaptive agent chain. To enhance system robustness, we employ a reputation mechanism that dynamically updates each agent’s credit score upon task completion. In the following sections, Section 3.1 introduces definitions and optimization objectives, Section 3.2 details the forwarding criteria, Section 3.3 details how selection policy of *BiRouter* helps to generate agent chain, and Section 3.4 describes the generation strategy of our *BiRouter* training dataset.

3.1 Preliminaries

System and Agent Chain Definition. We define a Self-organizing Multi-agent System \mathcal{S} as a decentralized network

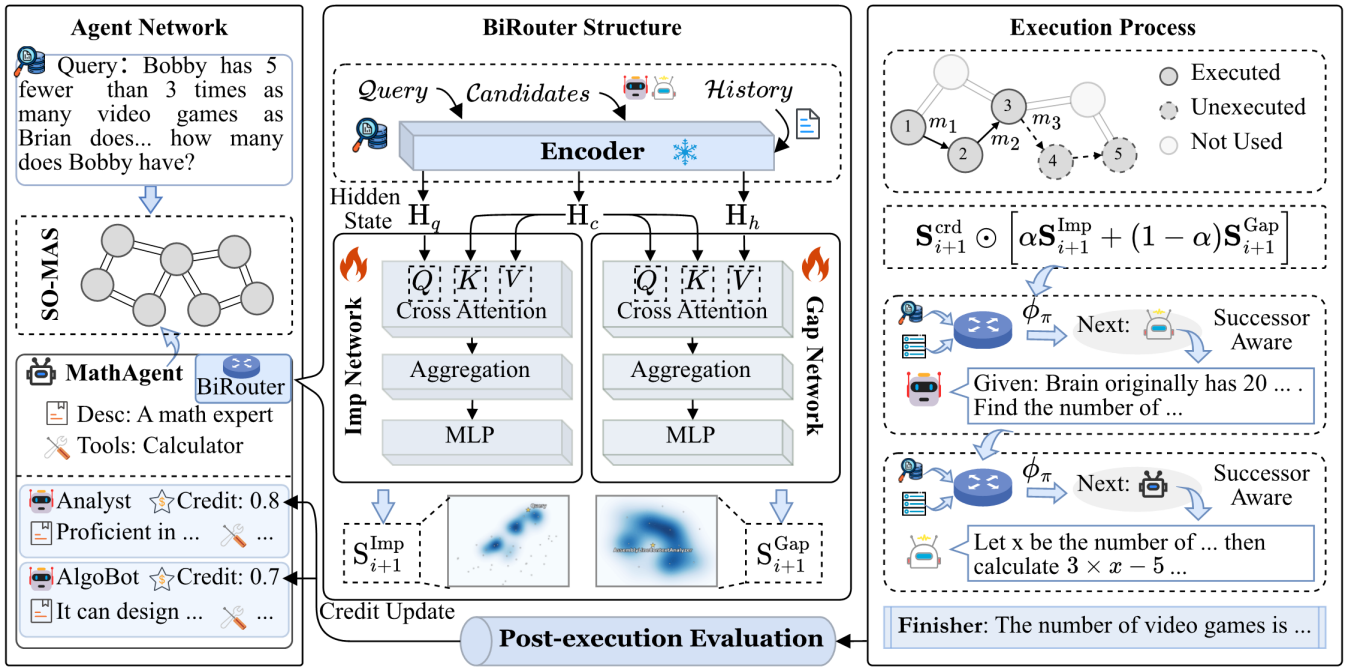


Figure 1: The architecture of BiRouter and its workflow in self-organizing agent systems.

operating without a central coordinator or hierarchical structure:

$$\mathcal{S} = \langle \mathcal{A}, \mathcal{C}^{\mathcal{A}}, \phi_{\pi} \rangle \quad (1)$$

where $\mathcal{A} = \{a_i | i = 1, \dots, n\}$ is the set of agents, each with a description $\text{Desc}(a_i)$; $\mathcal{C}^{\mathcal{A}} = \{\mathcal{C}^{a_i} | a_i \in \mathcal{A}\}$ defines the set of known successor candidates for each agent a_i ; and ϕ_{π} is the successor selection policy. In this system, each agent can partially execute a query and autonomously forward the task based on its local policy.

Given a query q , the system's objective is to form an *agent chain*, $C_{agt}(q) = \{x_1, x_2, \dots, x_L\}$, which is an ordered sequence of agents that collectively solve the task. The formation of this chain, denoted as $C_{agt} = \Pi(\mathcal{A}, q)$, results from this decentralized decision-making process.

Problem Formulation. In contrast to centralized approaches that pre-plan the entire agent chain based on global knowledge, our system operates without a central scheduler. The agent chain emerges dynamically from a sequence of local, autonomous decisions. Each agent x_i makes its decision based on a strictly local observation $o^{x_i}(q)$, which represents the agent's limited, egocentric view of the system:

$$o^{x_i}(q) = \underbrace{(\mathcal{H}_{i-1}^{(q)})}_{\text{Past}} \cup \underbrace{\text{Desc}(x_i)}_{\text{Present}} \cup \underbrace{\{\text{Desc}(x_j) | x_j \in \mathcal{C}^{x_i}\}}_{\text{Future}} \quad (2)$$

where $\mathcal{H}_{i-1}^{(q)} = (q, x_1, m_1, \dots, x_{i-1}, m_{i-1})$ is the chain's history up to step $i-1$, and m is the runtime message.

Our objective is to learn an optimal local policy π^* that selects the best successor $c^{x_i^*}$ from the local candidate set \mathcal{C}^{x_i} , rather than the global set \mathcal{A} :

$$c^{x_i^*} = \phi_{\pi^*}(\mathcal{C}^{x_i} | o^{x_i}(q)) \quad (3)$$

The optimal policy π^* favors successors based on a local decision score \mathbf{S}^{x_i} :

$$\pi^*(\mathcal{C}^{x_i} | o^{x_i}(q)) \propto \mathbf{S}^{x_i} \quad (4)$$

This single-hop forwarding mechanism allows each agent to make the best possible local decision, progressively constructing a globally effective agent chain C_{agt}^* .

3.2 Forwarding Criteria

The forwarding decision is governed by three key criteria:

Importance (ImpScore). This score quantifies an agent's capability relevance to the overall task query. It is computed by a function $\text{ImpScore} : \mathcal{A} \times \mathcal{Q} \rightarrow \mathbb{R}^{|\mathcal{A}| \times |\mathcal{Q}|}$, which outputs a matrix where each element (i, j) represents the importance of agent a_i for query q_j , where \mathcal{Q} is the query set and $q_j \in \mathcal{Q}$. The score prioritizes agents with higher task-criticality, thereby enhancing execution efficiency.

Cohesion (GapScore). This score evaluates how well a candidate agent fits into the emerging chain, ensuring logical and contextual continuity. The function $\text{GapScore} : \mathcal{C}^{x_i} \rightarrow \mathbb{R}^{|\mathcal{C}^{x_i}|}$ computes a vector of compatibility scores for the set of candidate successors \mathcal{C}^{x_i} of the current agent x_i . A higher score indicates a smoother, more logical transition.

The heatmaps in Figure 1 show the semantic distribution of ImpScore and GapScore, where darker colors indicate higher scores. Agents with a high ImpScore form a path from the query to the goal, while those with a high GapScore form a cluster around the current agent.

Agent Reputation. To ensure system reliability, the policy incorporates agent reputation, a dynamic value represented

by a credit score $\mathbf{S}^{\text{crd}} \in \mathbb{R}^{|A|}$, reflecting each agent’s historical trustworthiness and performance. This score is updated post-execution, where an LLM evaluates the agent’s performance to determine a multiplicative update factor.

3.3 The Successor Selection Policy

At each step i , the current agent x_i selects a successor x_{i+1} from its candidates \mathcal{C}^{x_i} . This decision is executed by its built-in **BiRouter** module, which implements our policy ϕ_π . The BiRouter consists of two branches that process inputs through a shared encoder, followed by distinct Cross Attention and MLP layers to compute the ImpScore and GapScore, respectively:

- **Importance Branch** (\mathcal{R}_{Imp}): Computes \mathbf{S}^{Imp} , assessing the semantic relevance of each candidate to the original query q .
- **Cohesion Branch** (\mathcal{R}_{Gap}): Computes \mathbf{S}^{Gap} , evaluating the contextual fit of each candidate as a successor to x_i .

Formally, the score vectors for the next step ($i + 1$) are computed as:

$$\begin{aligned} \mathbf{S}_{i+1}^{\text{Imp}}(q) &= \mathcal{R}_{\text{Imp}}(q, \text{Desc}(\mathcal{C}^{x_i})) \\ \mathbf{S}_{i+1}^{\text{Gap}}(q) &= \mathcal{R}_{\text{Gap}}(o^{x_i}(q)) \end{aligned} \quad (5)$$

where $\mathcal{R}_{\text{Imp}}(\cdot)$ and $\mathcal{R}_{\text{Gap}}(\cdot)$ are the neural network functions of the BiRouter, and the outputs $\mathbf{S}_{i+1}^{\text{Imp}}, \mathbf{S}_{i+1}^{\text{Gap}} \in \mathbb{R}^{|\mathcal{C}^{x_i}|}$.

Probabilistic Synthesis. These dynamic scores are integrated with the static agent reputation scores (\mathbf{S}^{crd}) to produce a final probability distribution for successor selection. The local policy π_i is defined as:

$$\begin{aligned} \text{Logits}_{i+1} &= \mathbf{S}_{i+1}^{\text{crd}} \odot \left[\alpha \mathbf{S}_{i+1}^{\text{Imp}} + (1 - \alpha) \mathbf{S}_{i+1}^{\text{Gap}} \right] \\ \pi_i(\mathcal{C}^{x_i} | x_i, q) &= \text{Softmax}(\text{Logits}_{i+1}) \end{aligned} \quad (6)$$

where $\mathbf{S}_{i+1}^{\text{crd}}$ is the vector of credit scores for the candidates, \odot denotes the element-wise product, and $\alpha \in [0, 1]$ is a hyperparameter balancing importance and cohesion. This formulation allows an agent’s reputation to act as a multiplicative gate on the combined task-relevance scores.

A key feature of our design is that the transition message m_i is generated **after** the successor x_{i+1} has been selected. This enables the current agent x_i to craft a **successor-aware message**, tailoring its content to the specific capabilities of x_{i+1} to ensure a seamless and efficient task handoff:

$$m_i = x_i(\mathcal{H}_{i-1}^{(q)}, \text{Desc}(x_{i+1})) \quad (7)$$

Overall Chain Probability. The formation of a complete agent chain $C_{\text{agt}}(q)$ is initiated from a user-specified starting agent x_1 . The probability of a specific chain, given the starting agent, is the product of the individual transition probabilities determined by our local policy π :

$$\mathcal{P}(C_{\text{agt}} | x_1, q) = \prod_{i=1}^{L-1} \pi(\mathcal{C}^{x_i} | x_i, q) \quad (8)$$

Ours
 GSM8K
 MultiArith
 MMLU
 HumanEval
 MBPP

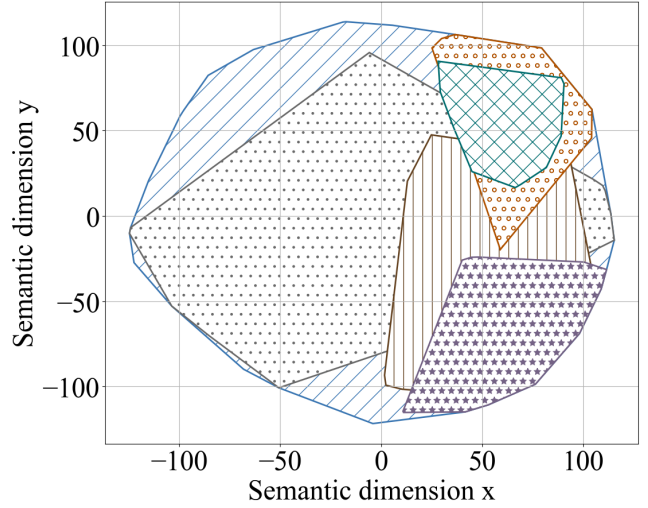


Figure 2: Distribution of MARS against others in a 2D semantic space projected by t-SNE.

Adaptive Task Termination. We introduce a special **Finisher** agent to enable adaptive task completion. As a potential candidate at each step, its selection concludes tasks based on the task state rather than a fixed length, thus reducing unnecessary token consumption.

3.4 Training Data Generation

To foster the BiRouter’s cross-domain generalization, we constructed Multi-agent Routing Dataset (MARS), a diverse dataset for BiRouter training. MARS consists of two main components: $\mathcal{D}_I = \{(q^{(i)}, a^{(i)}; s_I^{(i)})\}$ for ImpScore, which pairs queries and agents with importance scores, and $\mathcal{D}_C = \{(c^{(j)}, a^{(j)}; s_C^{(j)})\}$ for GapScore, which pairs execution contexts and agents with cohesion scores.

We initiated by prompting an LLM to produce a broad set of domain keywords, which were manually curated into 115 representative domains. Then, we prompted the LLM to generate unique queries within one domain or cross domains. To ensure comprehensive semantic coverage, we densified the query embedding space by identifying and populating sparse regions. This process was guided by a semantic density function using an RBF kernel:

$$\rho(x) = \frac{1}{N_q} \sum_{i=1}^{N_q} \exp\left(-\frac{\|x - q_i\|^2}{2\sigma^2}\right) \quad (9)$$

where the LLM was prompted to generate new queries in regions of low density $\rho(x)$.

For each query, an LLM generated multiple candidate execution paths (i.e., agent sequences). Within each path, the LLM ranked agents by their criticality; these ranks were averaged across multiple iterations for robustness. To promote efficiency, a penalty factor γ was applied to agents in long




Figure 3: BiRouter’s performance and token consumption on GSM8K and MBPP across different α values.

execution paths known for high token consumption. The final ImpScore is calculated using a scaled Sigmoid function to amplify the distinction between critical and non-critical agents within a target interval $[l, u]$:

$$s_I(a_i) = [l + (u - l) \cdot \frac{1}{1 + e^{-\beta(N_r - r_i)}}] \cdot \gamma \quad (10)$$

where N_r is the number of agents, r_i is the agent’s rank, l and u are set to 0.3 and 1, and $\beta = 2$ is a scaling parameter.

The context for GapScore is derived from the generated agent sequences. An agent a_i at step k receives a GapScore based on its position in the full path:

$$s_G(a_i, k) = \begin{cases} 1/(i - k + 1) & \text{if } i \geq k \\ 0 & \text{if } i < k \end{cases} \quad (11)$$

where i is the agent’s index. Finally, a dedicated “Finisher” agent was appended to each path to signal task completion. As shown in Figure 2, a t-SNE visualization of the dataset reveals a rich semantic distribution, which is crucial for enhancing the BiRouter’s cross-domain generalization.

4 Experiments

4.1 Experimental Setup

Benchmarks. We evaluate BiRouter on six public benchmarks: MMLU (Hendrycks et al. 2021), GSM8K (Cobbe et al. 2021), MultiArith (Roy and Roth 2016), SVAMP (Patel, Bhattamishra, and Goyal 2021), HumanEval (Chen et al. 2021) and MBPP (Austin et al. 2021).

Baselines. We compare BiRouter with three series of agentic baselines: (1) **Single-agent approach:** Single-agent LLM, CoT (Wei et al. 2022), ComplexCoT (Fu et al. 2023), and CoT-SC (Wang et al. 2023). (2) **Static coordination MAS:** MacNet (Qian et al. 2025), LLM-Blender (Jiang, Ren, and Lin 2023), LLM-Debate (Du et al. 2023), GPTSwarm (Zhuge et al. 2024), G-Designer (Zhang et al. 2025c). (3) **Dynamic coordination MAS:** AgentVerse (Chen et al. 2024b), DyLAN (Liu et al. 2024), MaAS (Zhang et al. 2025a).

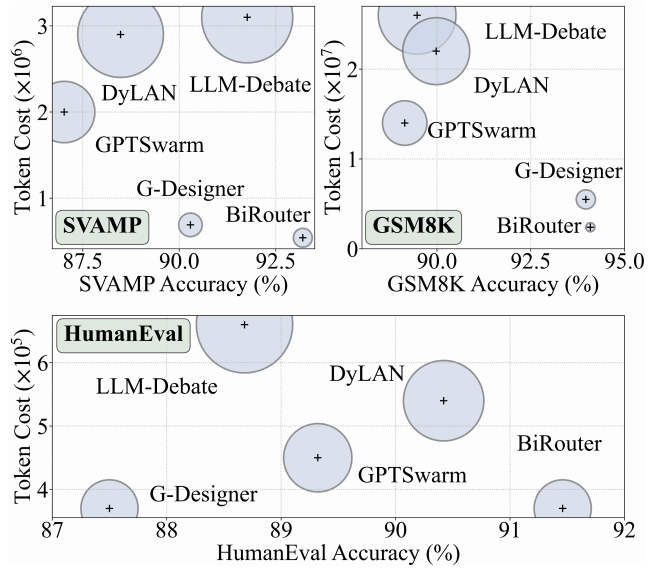


Figure 4: Visualization of the performance metrics and token consumption. The diameter of each point is proportional to its y -axis value.

Implementation Details. We use `qwen3-embedding-0.6b` as the frozen encoder for BiRouter, training only the Cross-Attention and scoring head on our synthetically generated data. All agents use `gpt-4o-mini` via the OpenAI API. We set the temperature to 0 for all single-agent methods and 1 for all multi-agent methods. We use `gpt-4.1` to generate agent profiles to build the agent network, including agents specialized for the benchmarks and one fixed “Finisher” agent for task exit and answer summarization. We conduct each experiment twice and take the average as the final result. The optimal hyperparameter α was tuned using $Q \in \{100, 200\}$ queries randomly sampled from GSM8K and MBPP. As illustrated in Figure 3, we evaluated BiRouter’s performance across various α values. Balancing performance against token consumption, we set the final value to $\alpha = 0.3$.

4.2 Results in Centralized MAS

We first conducted experiments in a centralized MAS environment where all agents are available at each stage of the task. We compared our method with 12 other baselines, with all agents powered by `gpt-4o-mini`. The evaluation results demonstrate that BiRouter is:

High-performing. As summarized in Table 1, BiRouter achieves state-of-the-art or highly competitive results across all benchmarks. The average accuracy confirms that BiRouter surpasses all other baseline categories, demonstrating performance gains of 5.86% to 7.06% over single-agent methods, 4.65% to 7.58% over static coordination methods, and 4.52% to 5.23% over dynamic coordination methods. BiRouter achieved optimal results on 4 out of 6 benchmarks, demonstrating strong task performance.

Method	Dyn.	Dis.	MMLU	GSM8K	SVAMP	HumanEval	MultiArith	MBPP	Avg.
Single-agent	–	–	77.81	87.45	88.26	87.08	96.85	71.83	84.88
CoT	–	–	78.43 _{+0.62}	87.10 _{−0.35}	86.24 _{−2.02}	88.13 _{+1.05}	96.31 _{−0.54}	71.83 _{+0.00}	84.67
ComplexCoT	–	–	81.05 _{+3.24}	86.89 _{−0.56}	90.53 _{+2.27}	87.49 _{+0.41}	96.70 _{−0.15}	72.36 _{+0.53}	85.84
SC (CoT×5)	–	–	80.96 _{+3.15}	87.57 _{+0.12}	87.92 _{−0.34}	88.60 _{+1.52}	96.58 _{−0.27}	73.60 _{+1.77}	85.87
MacNet	×	×	82.98 _{+5.17}	87.95 _{+0.50}	88.06 _{−0.20}	84.57 _{−2.51}	96.03 _{−0.82}	65.28 _{−6.55}	84.15
LLM-Blender	×	×	81.22 _{+3.41}	88.35 _{+0.90}	89.52 _{+1.26}	88.80 _{+1.72}	97.29 _{+0.44}	77.05 _{+5.22}	87.04
LLM-Debate	×	×	81.04 _{+3.23}	89.47 _{+2.02}	91.76 _{+3.50}	88.68 _{+1.60}	97.33 _{+0.48}	70.29 _{−1.54}	86.43
GPTSwarm	×	×	82.80 _{+4.99}	89.14 _{+1.69}	87.02 _{−1.24}	89.32 _{+2.44}	96.79 _{−0.06}	77.43 _{+5.60}	87.08
G-Designer	×	×	87.20 _{+9.39}	<u>93.97</u> _{+6.52}	90.29 _{+2.03}	87.50 _{+0.42}	98.33 _{+1.48}	–	–
AgentVerse	✓	×	78.36 _{+0.55}	89.91 _{+2.46}	89.64 _{+1.38}	89.29 _{+2.21}	97.50 _{+0.65}	74.28 _{+2.45}	86.50
DyLAN	✓	×	79.96 _{+2.15}	89.98 _{+2.53}	88.48 _{+0.22}	90.42 _{+3.34}	97.12 _{+0.27}	77.30 _{+5.47}	87.21
MaAS	✓	×	–	92.30 _{+4.85}	–	92.85 _{+5.77}	<u>98.80</u> _{+1.95}	<u>82.17</u> _{+10.34}	–
BiRouter (Ours)	✓	✓	<u>86.80</u> _{+8.99}	94.09 _{+6.64}	93.20 _{+4.04}	<u>91.46</u> _{+4.38}	100.00 _{+3.15}	84.82 _{+12.99}	91.73

Table 1: Performance comparison with BiRouter and baselines in centralized MAS settings. “Dyn.” and “Dis.” indicate whether a method supports dynamic topology and distributed execution (✓: yes, ×: no).

Method	GSM8K		HumanEval	
	Acc.(%)	Cost	Pass@1(%)	Cost
DyLAN	87.95	6.3×10^6	68.29	6.2×10^5
MaAS	86.43	3.8×10^6	79.27	4.9×10^5
BiRouter	91.99	2.8×10^6	89.63	4.6×10^5

Table 2: Performance and token consumption comparison in SO-MAS settings. “Acc.” denotes accuracy, and “Cost” represents the number of tokens used.

Token-economical. Figure 4 visualizes the trade-off between token consumption (both prompt and completion tokens) and performance. Across both SVAMP and GSM8K, BiRouter not only surpasses all other baselines in performance but does so with the lowest token expenditure. On HumanEval, it delivers a 3.96% performance boost over G-Designer while maintaining a comparable token count. This stands in sharp contrast to methods like GPTSwarm and DyLAN, whose complex, iterative approaches inflate token costs. By comparison, BiRouter exemplifies an exceptional balance between high performance and resource efficiency.

Generalizable. The consistent high performance across diverse domains (general reasoning, math, and code) demonstrates BiRouter’s strong generalization capabilities, especially given it was trained exclusively on synthetically generated data. Its task-adaptive and dynamically adaptive routing mechanism allows it to adjust agent selection in real-time, showing its scalability in dynamic environments.

4.3 Results in SO-MAS

In this section, we simulate a self-organizing MAS (SO-MAS) environment where each agent can only communicate with three other randomly selected agents and must choose




Figure 5: Agent credit score dynamics on GSM8K and HumanEval benchmarks.

its successor from this limited set. We compare BiRouter with two other dynamic methods capable of operating in such a setting, DyLAN and MaAS, with results shown in Table 2. The localized information in this decentralized setting leads to a general performance degradation across all methods. However, BiRouter exhibits the smallest performance drop, achieving 91.99 on GSM8K and 89.63 on HumanEval, respectively. This resilience stems from its native design for local decision-making. In contrast, DyLAN and MaAS rely heavily on the availability of global information, leading to a significant performance decline. BiRouter’s strong performance in both centralized and decentralized architectures underscores its cross-architecture adaptability.

4.4 Robustness

To evaluate BiRouter’s robustness, we introduced an equal number of unreliable agents into the network, created by using `gpt-4.1` to rewrite agent prompts to intentionally produce incorrect outputs. As shown in Figure 5, which plots agent credit scores during experiments on GSM8K and HumanEval, BiRouter rapidly identifies and down-weights

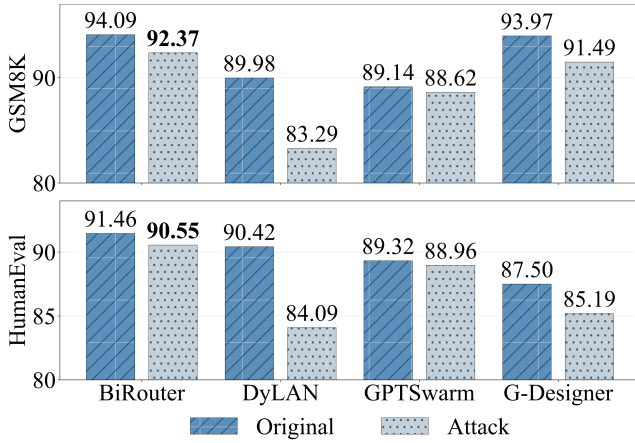


Figure 6: Performance comparison between before and after unreliable agent attack.

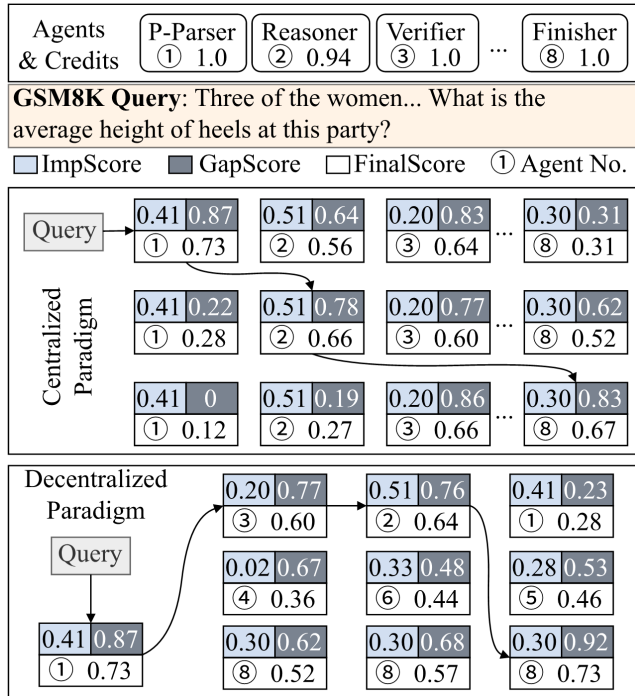


Figure 7: Visualization of BiRouter’s routing mechanism across centralized and decentralized paradigms.

these unreliable agents, mitigating their negative influence. Consequently, as seen in Figure 6, while all methods suffer from performance degradation, BiRouter remains the highest performance under unreliable agent attacks. This demonstrates its robustness in open, untrustworthy environments.

4.5 Case Study

We present a case study on GSM8K to illustrate BiRouter’s mechanism. As shown in Figure 7, in the centralized setting, BiRouter efficiently selects agents based on global visibility: ProblemParser (0.73) initiates, followed by MathReasoner

Variant	GSM8K				HumanEval			
	Origin Acc. (%)	Attack Cost (10 ⁶)	Origin Acc. (%)	Attack Cost (10 ⁶)	Origin Pass (%)	Attack Cost (10 ⁵)	Origin Pass (%)	Attack Cost (10 ⁵)
BiRouter	94.09	2.1	92.37	2.5	91.46	3.5	90.55	3.7
w/o S ^{crd}	92.31	2.5	84.79	3.2	90.24	3.7	82.93	4.9
w/o Succ.	90.84	2.5	90.35	3.0	87.80	3.8	87.20	3.9
w/o Fin.	90.08	4.1	88.73	4.2	88.72	5.4	88.10	5.4

Table 3: Ablation study of BiRouter on GSM8K and HumanEval.

(0.66) for reasoning and computation, and finally Finisher (the highest GapScore of 0.83) upon task completion. In the decentralized setting, with limited visibility, BiRouter initially selects the locally optimal Verifier (0.60) but corrects its path by routing to MathReasoner (0.64), still reaching the solution.

4.6 Ablation Study

We conducted ablation studies on GSM8K and HumanEval with the following configurations: (1) **w/o S^{crd}**: Agent credit scores are fixed at 1. (2) **w/o Succ.**: The current agent is not informed of the next-hop agent. (3) **w/o Fin.**: The Finisher agent is removed, and the number of execution steps is fixed at 5. The results in Tables 3 show that all three components contribute significantly. Without the Reputation mechanism, BiRouter fails to recognize unreliable agents, degrading performance in untrustworthy networks. Removing the Successor-aware mechanism weakens inter-agent cooperation and lowers collaborative efficiency. The absence of the Finisher agent eliminates the ability for early task completion, forcing a fixed execution length. This reduces dynamic adaptability and leads to unnecessary token consumption, increasing the token overhead by approximately 69% on GSM8K and 50% on HumanEval.

5 Conclusion

This paper addresses the critical challenge of task planning in open, decentralized multi-agent systems. We introduce BiRouter, a novel bi-criteria routing method designed for self-organizing MAS, which enables individual agents to make autonomous “next-hop” routing decisions based solely on local information. By balancing overall task relevance and contextual continuity of candidate agents at runtime, BiRouter dynamically constructs collaborative pathways, demonstrating high adaptability and scalability. Experiments show that BiRouter holds significant advantages in performance, token efficiency, and robustness, and can seamlessly adapt to both centralized and decentralized architectures. Our work provides an effective solution for achieving emergent intelligent collaboration. Future work will focus on online adaptive optimization and more complex collaboration patterns.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Austin, J.; Odena, A.; Nye, M.; Bosma, M.; Michalewski, H.; Dohan, D.; Jiang, E.; Cai, C.; Terry, M.; Le, Q.; and Sutton, C. 2021. Program Synthesis with Large Language Models. *arXiv:2108.07732*.
- Browne, C. B.; Powley, E.; Whitehouse, D.; Lucas, S. M.; Cowling, P. I.; Rohlfshagen, P.; Tavener, S.; Perez, D.; Samothrakis, S.; and Colton, S. 2012. A Survey of Monte Carlo Tree Search Methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1): 1–43.
- Chan, C.-M.; Chen, W.; Su, Y.; Yu, J.; Xue, W.; Zhang, S.; Fu, J.; and Liu, Z. 2024. ChatEval: Towards Better LLM-based Evaluators through Multi-Agent Debate. In *The Twelfth International Conference on Learning Representations*.
- Chang, G. 2024. Agent Network Protocol: Enabling All Agents to Connect and Communicate.
- Chen, G.; Dong, S.; Shu, Y.; Zhang, G.; Sesay, J.; Karlsson, B. F.; Fu, J.; and Shi, Y. 2024a. AutoAgents: A Framework for Automatic Agent Generation.
- Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; de Oliveira Pinto, H. P.; Kaplan, J.; Edwards, H.; Burda, Y.; Joseph, N.; Brockman, G.; Ray, A.; Puri, R.; Krueger, G.; Petrov, M.; Khlaaf, H.; Sastry, G.; Mishkin, P.; Chan, B.; Gray, S.; Ryder, N.; Pavlov, M.; Power, A.; Kaiser, L.; Bavarian, M.; Winter, C.; Tillet, P.; Such, F. P.; Cummings, D.; Plappert, M.; Chantzis, F.; Barnes, E.; Herbert-Voss, A.; Guss, W. H.; Nichol, A.; Paino, A.; Tezak, N.; Tang, J.; Babuschkin, I.; Balaji, S.; Jain, S.; Saunders, W.; Hesse, C.; Carr, A. N.; Leike, J.; Achiam, J.; Misra, V.; Morikawa, E.; Radford, A.; Knight, M.; Brundage, M.; Murati, M.; Mayer, K.; Welinder, P.; McGrew, B.; Amodei, D.; McCandlish, S.; Sutskever, I.; and Zaremba, W. 2021. Evaluating Large Language Models Trained on Code. *arXiv:2107.03374*.
- Chen, S.; Liu, Y.; Han, W.; Zhang, W.; and Liu, T. 2025. A Survey on LLM-based Multi-Agent System: Recent Advances and New Frontiers in Application. *arXiv:2412.17481*.
- Chen, W.; Su, Y.; Zuo, J.; Yang, C.; Yuan, C.; Chan, C.-M.; Yu, H.; Lu, Y.; Hung, Y.-H.; Qian, C.; Qin, Y.; Cong, X.; Xie, R.; Liu, Z.; Sun, M.; and Zhou, J. 2024b. AgentVerse: Facilitating Multi-Agent Collaboration and Exploring Emergent Behaviors. In *The Twelfth International Conference on Learning Representations*.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; Hesse, C.; and Schulman, J. 2021. Training Verifiers to Solve Math Word Problems. *arXiv:2110.14168*.
- Du, Y.; Li, S.; Torralba, A.; Tenenbaum, J. B.; and Mordatch, I. 2023. Improving Factuality and Reasoning in Language Models through Multiagent Debate. *arXiv:2305.14325*.
- Fu, Y.; Peng, H.; Sabharwal, A.; Clark, P.; and Khot, T. 2023. Complexity-Based Prompting for Multi-step Reasoning. In *The Eleventh International Conference on Learning Representations*.
- Google. 2025. A2A: An open protocol enabling communication and interoperability between opaque agentic applications.
- Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2): 100–107.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2021. Measuring Massive Multitask Language Understanding. In *International Conference on Learning Representations*.
- Jiang, D.; Ren, X.; and Lin, B. Y. 2023. LLM-Blender: Ensembling Large Language Models with Pairwise Ranking and Generative Fusion. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 14165–14178. Toronto, Canada: Association for Computational Linguistics.
- Li, C.; Chen, H.; Yan, M.; Shen, W.; Xu, H.; Wu, Z.; Zhang, Z.; Zhou, W.; Chen, Y.; Cheng, C.; Shi, H.; Zhang, J.; Huang, F.; and Zhou, J. 2023a. ModelScope-Agent: Building Your Customizable Agent System with Open-source Large Language Models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 566–578. Singapore: Association for Computational Linguistics.
- Li, G.; Hammoud, H.; Itani, H.; Khizbullin, D.; and Ghanem, B. 2023b. CAMEL: Communicative Agents for "Mind" Exploration of Large Language Model Society. In *Advances in Neural Information Processing Systems*, volume 36, 51991–52008. Curran Associates, Inc.
- Liang, T.; He, Z.; Jiao, W.; Wang, X.; Wang, Y.; Wang, R.; Yang, Y.; Shi, S.; and Tu, Z. 2024. Encouraging Divergent Thinking in Large Language Models through Multi-Agent Debate. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 17889–17904. Miami, Florida, USA: Association for Computational Linguistics.
- Liu, Z.; Zhang, Y.; Li, P.; Liu, Y.; and Yang, D. 2024. A Dynamic LLM-Powered Agent Network for Task-Oriented Agent Collaboration. In *First Conference on Language Modeling*.
- Marro, S.; Malfa, E. L.; Wright, J.; Li, G.; Shadbolt, N.; Wooldridge, M. J.; and Torr, P. 2025. A Scalable Communication Protocol for Networks of Large Language Models.
- Meng, S.; Wang, Y.; Yang, C.-F.; Peng, N.; and Chang, K.-W. 2024. LLM-A*: Large Language Model Enhanced Incremental Heuristic Search on Path Planning. *Research Square*.
- Patel, A.; Bhattamishra, S.; and Goyal, N. 2021. Are NLP Models really able to Solve Simple Math Word Problems? *arXiv:2103.07191*.
- Qian, C.; Xie, Z.; Wang, Y.; Liu, W.; Zhu, K.; Xia, H.; Dang, Y.; Du, Z.; Chen, W.; Yang, C.; Liu, Z.; and Sun, M.

2025. Scaling Large Language Model-based Multi-Agent Collaboration. In *The Thirteenth International Conference on Learning Representations*.
- Roy, S.; and Roth, D. 2016. Solving General Arithmetic Word Problems. arXiv:1608.01413.
- Sun, D.; Vlasic, D.; Herrmann, C.; Jampani, V.; Krainin, M.; Chang, H.; Zabih, R.; Freeman, W. T.; and Liu, C. 2021. AutoFlow: Learning a Better Training Set for Optical Flow. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10088–10097.
- Tran, K.-T.; Dao, D.; Nguyen, M.-D.; Pham, Q.-V.; O’Sullivan, B.; and Nguyen, H. D. 2025. Multi-Agent Collaboration Mechanisms: A Survey of LLMs. arXiv:2501.06322.
- Veerapaneni, R.; Saleem, M. S.; and Likhachev, M. 2023. Learning Local Heuristics for Search-Based Navigation Planning. *Proceedings of the International Conference on Automated Planning and Scheduling*, 33(1): 634–638.
- Wang, J.; WANG, J.; Athiwaratkun, B.; Zhang, C.; and Zou, J. 2025. Mixture-of-Agents Enhances Large Language Model Capabilities. In *The Thirteenth International Conference on Learning Representations*.
- Wang, L.; Ma, C.; Feng, X.; Zhang, Z.; Yang, H.; Zhang, J.; Chen, Z.; Tang, J.; Chen, X.; Lin, Y.; Zhao, W. X.; Wei, Z.; and Wen, J. 2024. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18: 186345.
- Wang, X.; Wei, J.; Schuurmans, D.; Le, Q.; Chi, E.; Narang, S.; Chowdhery, A.; and Zhou, D. 2023. Self-Consistency Improves Chain of Thought Reasoning in Language Models. arXiv:2203.11171.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E.; Le, Q. V.; and Zhou, D. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems*, volume 35, 24824–24837. Curran Associates, Inc.
- Wu, Q.; Bansal, G.; Zhang, J.; Wu, Y.; Li, B.; Zhu, E.; Jiang, L.; Zhang, X.; Zhang, S.; Liu, J.; Awadallah, A. H.; White, R. W.; Burger, D.; and Wang, C. 2024. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversations. In *First Conference on Language Modeling*.
- Xi, Z.; Chen, W.; Guo, X.; He, W.; Ding, Y.; Hong, B.; Zhang, M.; Wang, J.; Jin, S.; Zhou, E.; Zheng, R.; Fan, X.; Wang, X.; Xiong, L.; Zhou, Y.; Wang, W.; Jiang, C.; Zou, Y.; Liu, X.; Yin, Z.; Dou, S.; Weng, R.; Cheng, W.; Zhang, Q.; Qin, W.; Zheng, Y.; Qiu, X.; Huang, X.; and Gui, T. 2023. The Rise and Potential of Large Language Model Based Agents: A Survey. arXiv:2309.07864.
- Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T.; Cao, Y.; and Narasimhan, K. 2023a. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. In *Advances in Neural Information Processing Systems*, volume 36, 11809–11822. Curran Associates, Inc.
- Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K. R.; and Cao, Y. 2023b. ReAct: Synergizing Reasoning and Acting in Language Models. In *The Eleventh International Conference on Learning Representations*.
- Zhang, G.; Niu, L.; Fang, J.; Wang, K.; BAI, L.; and Wang, X. 2025a. Multi-agent Architecture Search via Agentic Supernet. In *Forty-second International Conference on Machine Learning*.
- Zhang, G.; Yue, Y.; Li, Z.; Yun, S.; Wan, G.; Wang, K.; Cheng, D.; Yu, J. X.; and Chen, T. 2025b. Cut the Crap: An Economical Communication Pipeline for LLM-based Multi-Agent Systems. In *The Thirteenth International Conference on Learning Representations*.
- Zhang, G.; Yue, Y.; Sun, X.; Wan, G.; Yu, M.; Fang, J.; Wang, K.; Chen, T.; and Cheng, D. 2025c. G-Designer: Architecting Multi-agent Communication Topologies via Graph Neural Networks. In *Forty-second International Conference on Machine Learning*.
- Zhang, J.; Xiang, J.; Yu, Z.; Teng, F.; Chen, X.-H.; Chen, J.; Zhuge, M.; Cheng, X.; Hong, S.; Wang, J.; Zheng, B.; Liu, B.; Luo, Y.; and Wu, C. 2025d. AFlow: Automating Agentic Workflow Generation. In *The Thirteenth International Conference on Learning Representations*.
- Zhang, J.; Xu, X.; Zhang, N.; Liu, R.; Hooi, B.; and Deng, S. 2023. Exploring Collaboration Mechanisms for LLM Agents: A Social Psychology View. arXiv:2310.02124.
- Zhao, W. X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.; Du, Y.; Yang, C.; Chen, Y.; Chen, Z.; Jiang, J.; Ren, R.; Li, Y.; Tang, X.; Liu, Z.; Liu, P.; Nie, J.-Y.; and Wen, J.-R. 2025. A Survey of Large Language Models. arXiv:2303.18223.
- Zhuang, Y.; Chen, X.; Yu, T.; Mitra, S.; Bursztyjn, V.; Rossi, R. A.; Sarkhel, S.; and Zhang, C. 2024. ToolChain*: Efficient Action Space Navigation in Large Language Models with A* Search. In *The Twelfth International Conference on Learning Representations*.
- Zhuce, M.; Wang, W.; Kirsch, L.; Faccio, F.; Khizbullin, D.; and Schmidhuber, J. 2024. GPTSwarm: Language Agents as Optimizable Graphs. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 62743–62767. PMLR.