

S-DAG: A Subject-Based Directed Acyclic Graph for Multi-Agent Heterogeneous Reasoning

Jiangwen Dong^{*1}, Zehui Lin^{*1}, Wanyu Lin^{†1,2}, Mingjin Zhang²

¹Department of Data Science and Artificial Intelligence, The Hong Kong Polytechnic University, Hong Kong SAR, China

²Department of Computing, The Hong Kong Polytechnic University, Hong Kong SAR, China

Abstract

Large Language Models (LLMs) have achieved impressive performance in complex reasoning problems. Their effectiveness highly depends on the specific nature of the task, especially the required domain knowledge. Existing approaches, such as mixture-of-experts, typically operate at the task level; they are too coarse to effectively solve the heterogeneous problems involving multiple subjects. This work proposes a novel framework that performs fine-grained analysis at the subject level equipped with a designated multi-agent collaboration strategy for addressing heterogeneous problem reasoning. Specifically, given an input query, we first employ a Graph Neural Network to identify the relevant subjects and infer their interdependencies to generate an *Subject-based Directed Acyclic Graph* (S-DAG), where nodes represent subjects and edges encode information flow. Then, we profile the LLM models by assigning each model a subject-specific expertise score, and select the top-performing one for matching the corresponding subject of the S-DAG. Such subject-model matching enables graph-structured multi-agent collaboration where information flows from the starting model to the ending model over S-DAG. We curate and release multi-subject subsets of standard benchmarks (MMLU-Pro, GPQA, MedMCQA) to better reflect complex, real-world reasoning tasks. Extensive experiments show that our approach significantly outperforms existing task-level model selection and multi-agent collaboration baselines in accuracy and efficiency. These results highlight the effectiveness of subject-aware reasoning and structured collaboration in addressing complex and multi-subject problems.

Code —

https://github.com/WanyuGroup/AAAI2026_S-DAG

Extended version — <https://arxiv.org/abs/2511.06727>

1 Introduction

In recent years, intelligent agents based on large language models (LLMs) have developed rapidly and achieved significant advancements across various fields, ranging from question answering (Yue 2025; Zhuang et al. 2023) to text generation (Huang et al. 2023; Wu et al. 2024) and complex reasoning tasks (Ke et al. 2025; Zhang et al. 2024b). While the

development of a general-purpose LLM is promising (Mumuni and Mumuni 2025; Kojima et al. 2022), it becomes evident that a single LLM often struggles to handle complex reasoning tasks, especially when these problems span multiple disciplines (Feng et al. 2025a). This limitation raises higher demands for model training and fine-tuning (Hoffmann et al. 2022). In this context, multi-agent systems based on LLMs have emerged, aiming to leverage the collective intelligence and specialized expertise of multiple agents to tackle complex and multidisciplinary problems (Guo et al. 2024; Du et al. 2023; Talebirad and Nadiri 2023; Han et al. 2024; Gu et al. 2025; Liang et al. 2024; Yao et al. 2025)¹.

Existing research has explored the mixture-of-experts (MoE) framework that dynamically selects the most suitable LLMs for a given problem (Masoudnia and Ebrahimpour 2014; Zhou et al. 2022; Cai et al. 2024). Subsequently, the mixture-of-agents (MoA) paradigm leveraging multi-agent collaboration is proposed to deal with more complex problems by combining the strengths of diverse LLMs (Du et al. 2023; Zhang et al. 2024c; Wang et al. 2024; Li et al. 2024a). The most relevant to us is Symbolic-MoE, which analyzes the required subject knowledge for a task and then utilizes a set of top-k expert/subject models to solve the heterogeneous problem (Chen et al. 2025). For clarity, we organize the existing heterogeneous reasoning paradigms in Table 1. These prior works often assume queries belong to a single knowledge domain or simply rely on a single “best” model or agent for reasoning (Chen et al. 2025; Feng et al. 2024, 2025b; Feng, Shen, and You 2025). Very rare work considers the fine-grained subject-specific information of the problem, not to mention multi-agent collaboration at the subject level, as shown in Figure 1. Such limitation hinders their applicability of prior works in heterogeneous reasoning tasks, where seamless integration of cross-domain knowledge is critical. Therefore, this paper aims to address the following research problem: *How can we optimally select and coordinate expert LLMs at subject level for complex, multi-subject problems to achieve both high accuracy and efficiency?*

In this work, we propose Subject-based Directed Acyclic Graph (S-DAG), as shown in Figure 2, a novel framework for addressing heterogeneous reasoning problems that re-

^{*}These authors contributed equally.

[†]Corresponding Author: Wanyu Lin, wan-yu.lin@polyu.edu.hk
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹For simplicity, agents, LLMs, and models are used interchangeably.

| Category | Example Work | Subject-Level Analysis | MAS Collaboration | Subject-Specific Collaboration |
|--------------------------|---|------------------------|-------------------|--------------------------------|
| Routing | GraphRouter (Feng, Shen, and You 2025) | ✗ | ✗ | ✗ |
| | SymbolicMoE (Chen et al. 2025) | ✓ | ✗ | ✗ |
| Multi-Agent System (MAS) | Heterogeneous Swarm (Feng et al. 2025b) | ✗ | ✓ | ✗ |
| | Knowledge Card (Feng et al. 2024) | ✓ | ✗ | ✗ |
| Routing for MAS | S-DAG (Ours) | ✓ | ✓ | ✓ |

Table 1: Comparison of our proposed S-DAG and prior methods for heterogeneous reasoning. Unlike prior methods that either perform single-domain routing or lack fine-grained problem understanding, S-DAG supports detailed subject-level analysis and enables dynamic, subject-specific multi-agent collaboration for more effective reasoning.

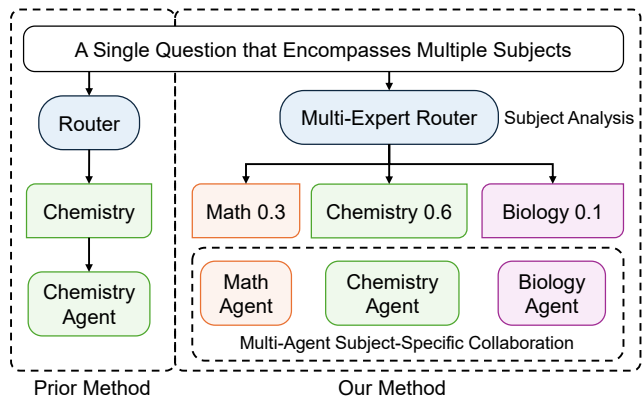


Figure 1: Comparison of the prior method with single agent and the proposed S-DAG approach with multi-agent collaboration. The prior method routes the problem to a single agent based on a coarse domain label, while our S-DAG approach conducts fine-grained subject analysis, identifying multiple relevant domains with associated relevant weights.

quire knowledge across multiple subject domains. The S-DAG identifies the relevant subjects for a given problem and defines the graph-structured information flow for multi-agent collaboration. We begin by modeling the complete set of subjects as a fully connected graph. To extract fine-grained subject-level structure, we introduce a specialized graph neural network that learns node embeddings to capture the relevant subjects and their interdependencies with respect to a given problem. From this, we derive a subject-based directed acyclic graph that reflects the essential subjects and reasoning flow for the problem. Based on the constructed S-DAG, we perform subject-LLM matching by profiling LLMs according to their subject-specific capabilities. The constructed S-DAG and LLM profile guide a structured multi-agent collaboration mechanism, where domain-specialized LLMs are assigned to subject nodes and communicate according to the DAG topology. Through this design, our S-DAG enables efficient and subject-level multi-agent reasoning. In summary, the main contributions of this work are as follows:

- To the best of our knowledge, we are the first to study the heterogeneous reasoning problem that a single complex problem covers multiple subject knowledge. Our novel

framework, S-DAG, enables fine-grained subject-specific decomposition and graph-structured multi-agent collaboration mechanism for the multi-subject problem.

- We develop a fine-grained subject-LLM matching strategy by profiling LLMs according to subject-specific capabilities, enabling precise assignment of expert agents and efficient coordination via the S-DAG.
- We curate multi-subject evaluation datasets by manually selecting samples that require multi-subject knowledge from three challenging benchmarks—MMLU-Pro, GPQA, and MedMCQA. Extensive experiments demonstrate that our approach substantially outperforms both single-model and multi-model baselines in terms of accuracy and computational efficiency.

2 Related Work

Multi-Agent Systems. Recent advances in Multi-Agent Systems (MAS) have introduced diverse collaboration mechanisms to tackle complex tasks (Guo et al. 2024; Talebirad and Nadiri 2023; Han et al. 2024), broadly categorized into fixed and dynamic paradigms. (1) *Fixed Multi-Agent Systems* rely on manually designed architectures, such as LLM debates (Du et al. 2023; Liang et al. 2024), chain-of-agents (Gu et al. 2025; Zhang et al. 2024c; Tao, Zhao, and Feng 2025), and graph-based systems (Yin et al. 2023; Li et al. 2024b; Lin et al. 2025). These approaches enable collaboration among agents with predefined roles and structures, making them effective for well-structured problems, but often lacking adaptability to dynamic tasks. (2) *Dynamic Multi-Agent Systems* adapt their structure in response to real-time task demands. Notable examples include GPTSwarm (Zhuge et al. 2024), GHG (Li et al. 2025) and Heterogeneous Swarms (Feng et al. 2025b), which optimize collaboration through dynamic graph structures. AgentPrune (Zhang et al. 2024a) improves communication efficiency by pruning redundant links, while DyLAN (Liu et al. 2024) dynamically selects agents and communication paths based on task context. MasRouter (Yue et al. 2025) introduces a cascaded controller for mode selection, role assignment, and LLM routing, enabling efficient and adaptive MAS construction.

Heterogeneous Reasoning. Heterogeneous reasoning focuses on solving problems that require knowledge across multiple domains (Xin et al. 2024; Chen et al. 2025; Feng et al. 2024). Existing approaches can be grouped by the

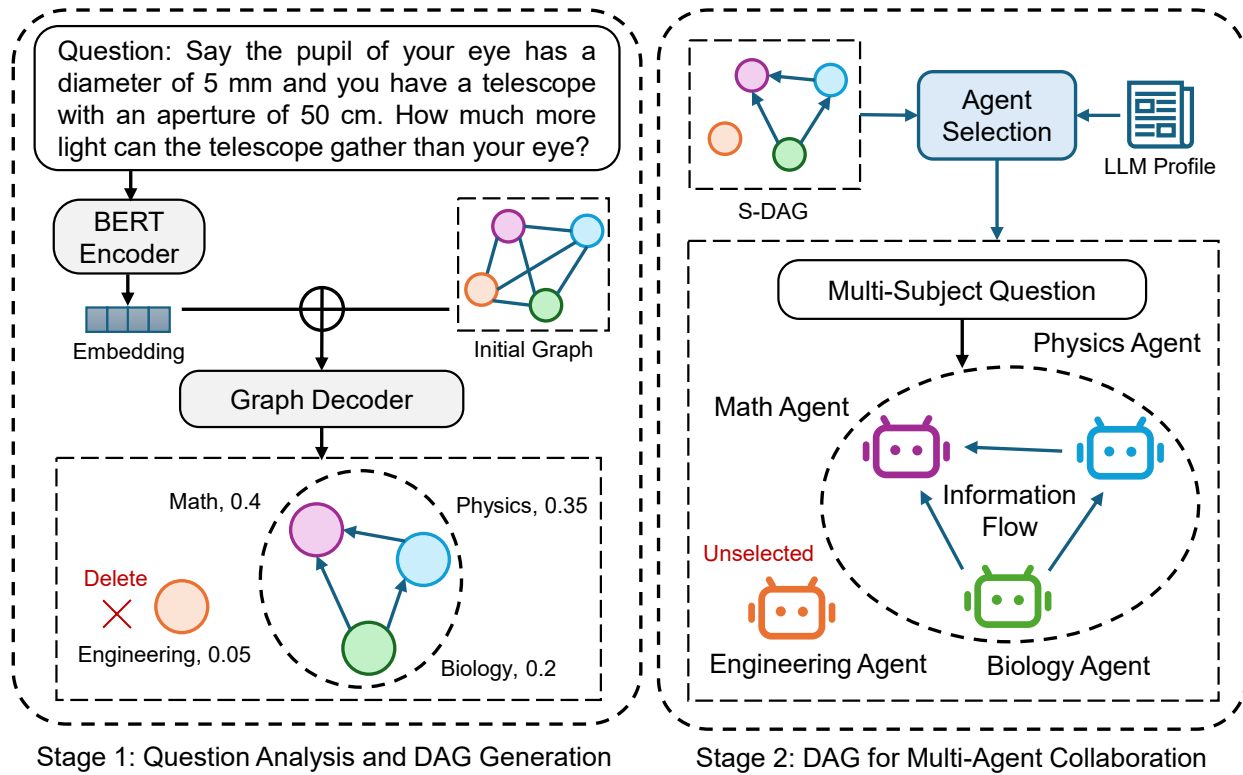


Figure 2: The Overview of the S-DAG Framework. The framework operates in two stages. In Stage 1, the input question is encoded using a BERT encoder, and a Graph Decoder generates the S-DAG, capturing subject dependencies and pruning irrelevant subjects. In Stage 2, expert LLMs are selected based on their subject-specific expertise and organized according to the S-DAG, with directed edges defining the information flow for multi-subject collaborative reasoning.

number of models involved: (1) *Single-Model Approaches* select one expert model per query, often using routing mechanisms. MoE techniques (Masoudnia and Ebrahimpour 2014; Zhou et al. 2022; Cai et al. 2024) specialize models over input space, while FrugalGPT (Chen, Zaharia, and Zou 2024) uses a reliability predictor, and GraphRouter (Feng, Shen, and You 2025) employs a graph neural network to frame selection as edge prediction. Though effective, these methods fall short when queries require multi-domain expertise. (2) *Multi-Model Approaches* enable collaboration across multiple expert models. Mixture-of-Agents (MoA) (Wang et al. 2024; Li et al. 2024a) allows coordinated reasoning, while SymbolicMoE (Chen et al. 2025) aggregates top-k responses based on skill relevance. Knowledge-Card (Feng et al. 2024) dynamically selects smaller fine-tuned agents, and Heterogeneous Swarm (Feng et al. 2025b) optimizes reasoning via a DAG structure. In contrast to prior work focused on dataset-level heterogeneity, our method targets a more granular challenge: each individual question requires reasoning across multiple subject domains.

3 Methodology

Method Overview. As illustrated in Figure 2, our proposed S-DAG framework enables subject-aware multi-agent reasoning for complex, multi-subject questions through a two-

stage process. In Stage 1, the input question is encoded using a BERT encoder, and a Graph Neural Network predicts relevant subjects and their dependencies to construct a Subject-based Directed Acyclic Graph (S-DAG), filtering out irrelevant domains. In Stage 2, expert LLMs are selected based on subject-specific capability profiles and assigned to the S-DAG nodes. These agents collaborate according to the graph’s structure, with directed edges guiding information flow from supporting to dominant subjects, enabling efficient and accurate multi-subject reasoning.

3.1 Problem Setup

Let Q denotes a natural language question that spans multiple subject areas $\mathcal{S} = \{s_1, s_2, \dots, s_K\}$, and let $\mathcal{M} = \{M_1, M_2, \dots, M_n\}$ represent a pool of domain-specific expert LLMs. The objective is to solve Q by determining a small set of relevant subject domains $\mathcal{S}_Q \subseteq \mathcal{S}$ (typically $|\mathcal{S}_Q| \leq 5$), identifying the interdependencies among them, and assigning each subject $s_i \in \mathcal{S}_Q$ to a corresponding expert model $M_j \in \mathcal{M}$ that is most proficient in that domain. More specifically, we model the relationships between the selected subjects by constructing a DAG over \mathcal{S}_Q , i.e., S-DAG, where a directed edge $s_i \rightarrow s_j$ indicates that subject s_i provides auxiliary support for reasoning in subject s_j in solving Q . This structure reflects the compositional nature of multi-domain reasoning and determines how dif-

ferent experts should collaborate. To obtain the S-DAG, we define a fully connected directed graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where each node $v_i \in \mathcal{V}$ corresponds to a subject $s_i \in \mathcal{S}$, and each directed edge $(v_i, v_j) \in \mathcal{E}$ encodes a potential dependency between subject pairs. From \mathcal{G} , we derive the pruned S-DAG $\mathcal{G}_Q = \{\mathcal{S}_Q, \mathcal{A}_Q\}$, which serves as a high-level reasoning blueprint. It guides the selection of a subset of expert LLMs, $\mathcal{M}_Q \subseteq \mathcal{M}$, and defines the collaboration topology among them, enabling effective multi-agent reasoning over complex, interdisciplinary queries.

3.2 Preprocessing

GNN Training. To effectively solve the multi-subject questions, we employ a set of expert agents, each specializing in a distinct domain. While LLMs are capable of identifying relevant subjects via prompting, they often struggle to capture fine-grained inter-subject dependencies and may produce outputs that are noisy, inconsistent, or lacking in structural coherence. To address this limitation, we introduce a trainable GNN module that learns to model subject dependencies through iterative message-passing over a subject-level graph. The resulting subject graph, or S-DAG, serves as a robust structural prior that guides information flow across agents.

Before training the GNN for S-DAG generation, we preprocess the dataset to construct ground-truth subject graphs for supervision. Given a question Q , we prompt a LLM to extract a set of relevant subject domains $\mathcal{S}_Q = \{s_1, s_2, \dots, s_k\}$. Each subject s_i is assigned a relevance weight $\{w_i\}_{i=1}^k \in [0, 1]$, indicating its relative importance for solving Q . Using these subjects and weights, we can construct a ground-truth subject graph $\mathcal{G}_Q = \{\mathcal{S}_Q, \mathcal{A}_Q\}$ for question Q , where $\mathcal{A}_Q \in \{0, 1\}^{k \times k}$ is the adjacency matrix. Specifically, a directed edge $a_{ij}^Q = 1$ indicates that subject s_i (with lower weights) supports subject s_j (with higher weight), reflecting the support-to-dominant subject relationship essential for multi-agent subject-specific reasoning. To ensure consistency, we use `qwen-turbo-0919` (Yang et al. 2024) as the *Subject LLM* and perform three rounds of processing for each question, only retaining subjects that appear consistently. Further details on dataset preprocessing and subject graph construction are provided in appendix. To ensure comprehensive modeling of potential subject interactions, we define a static, fully connected directed graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where each node corresponds to a candidate subject and each edge represents a possible dependency. This graph serves as the structural input for the GNN during S-DAG generation.

Given a question Q , a pretrained transformer encoder, BERT (Devlin et al. 2019), encodes it into a dense vector $\mathbf{h}_Q \in \mathbb{R}^d$, capturing the semantic intent of the input. Each subject node representation $v \in \mathcal{V}$ is initialized with a fused features of its subject embedding and question embedding via an MLP:

$$\mathbf{x}_i^{(0)} = \text{MLP}_{\text{init}}([\mathbf{h}_i; \mathbf{h}_Q]). \quad (1)$$

The initialized node features are then updated through layers of directional message passing within the GNN.

The final node representations are used for joint node and edge prediction, yielding a predicted subject graph $\mathcal{G}_Q = \{\hat{\mathcal{S}}_Q, \hat{\mathcal{A}}_Q\}$, where $\hat{s}_Q^i \in \hat{\mathcal{S}}_Q$ is the predicted relevance score for subject s_i , and \hat{a}_{ij}^Q is the predicted edge score for the dependency from s_i to s_j . Formally, the model is trained to minimize a multi-task binary cross-entropy loss between the predicted and ground-truth node and edge labels:

$$\mathcal{L} = \lambda_{\text{node}} \cdot \sum_{i=1}^K \text{BCE}(\hat{s}_Q^i, s_Q^i) + \lambda_{\text{edge}} \cdot \sum_{i \neq j} \text{BCE}(\hat{a}_{ij}^Q, a_{ij}^Q), \quad (2)$$

where λ_{node} and λ_{edge} control the weighting between node-level and edge-level supervision. To avoid penalizing irrelevant subjects, edge loss terms are masked when both $s_i = 0$ and $s_j = 0$. This objective encourages the model to learn subjects and their interaction patterns reflecting the reasoning dependencies in multi-domain tasks. Once trained, the GNN is used at inference time to construct a S-DAG for any new question, guiding the structure and flow of LLM-based agent collaboration tailored to the problem’s subject composition.

LLMs Subject Capability Profile. To optimize multi-agent collaboration, we construct a capability profile for each LLM based on its performance across various subject domains. This profile captures the subject-specific strengths of each model and serves as the foundation for expert selection within the heterogeneous multi-agent system. By leveraging these profiles, we ensure that queries are routed to the most competent models, improving both accuracy and efficiency in multi-domain reasoning.

Unlike prior work (Chen et al. 2025), which assigns a single subject label to each question, our approach captures the multi-domain nature by assigning weights to all relevant subjects. These weights reflect the relative importance of each subject in solving the question, enabling a more fine-grained and accurate assessment of each LLM’s performance across different domains. For instance, if model M_i answers a question spanning math, physics, and biology with weights $\{\text{'math'} : 0.5, \text{'physics'} : 0.3, \text{'biology'} : 0.2\}$ correctly, it would accordingly obtain the performance score $\text{Score}_{M_i} = \{\text{'math'} : +0.5, \text{'physics'} : +0.3, \text{'biology'} : +0.2\}$. By aggregating such weighted scores across a diverse set of questions, each model builds a subject capability profile that accurately reflects its domain expertise.

To construct these profiles, we randomly select 200 questions from the test set to assess LLM performance across each subject domain. The performance scores are then normalized to ensure comparability across domains and models. The normalized capability scores C_{ij} for model M_i and subject s_j are given by:

$$C_{ij} = \frac{\text{Score}_{M_i, s_j}}{\sum_{s_k \in \mathcal{S}} \text{Score}_{M_i, s_k}}, \quad (3)$$

where S is the set of all subject domains. Normalization ensures that the sum of scores across all domains for each

Algorithm 1: S-DAG Generation

Input: Input question \mathcal{Q} . Randomly initialized node embedding $\{\mathbf{h}_i\}_{i \in \mathcal{V}}$. An initial fully connected graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$.
Output: Generated S-DAG $\mathcal{G}_{\mathcal{Q}}$.

- 1: **for** each question \mathcal{Q} **do**
- 2: **Step 1: Question Embedding**
- 3: Embed the question: $\mathbf{h}_{\mathcal{Q}} \leftarrow \text{BERT}(\mathcal{Q})$.
- 4: Initialize node features: $\mathbf{x}_i \leftarrow \text{MLP}_{\text{init}}([\mathbf{h}_i; \mathbf{h}_{\mathcal{Q}}]), i \in \mathcal{V}$.
- 5: **Step 2: S-DAG Generation**
- 6: Update the node features: $\mathbf{x}_i = f_{\theta}(\mathcal{G}, \mathbf{x}_i^{(0)}), i \in \mathcal{V}$.
- 7: Node prediction: $s_{i, i \in \mathcal{V}} \leftarrow \text{MLP}_{\text{node}}(\mathbf{x}_i)$.
- 8: Edge prediction: $a_{\mathcal{Q}}^{ij, (i,j) \in \mathcal{E}} \leftarrow \text{MLP}_{\text{node}}(\mathbf{x}_i; \mathbf{x}_j; \mathbf{h}_{\mathcal{Q}})$.
- 9: S-DAG Construction: $\mathcal{G}_{\mathcal{Q}} = \{s_{\mathcal{Q}}^i, a_{\mathcal{Q}}^{ij}\}_{i=1}^K$.
- 10: **end for**

model equals 1, preventing skewed evaluations and enabling fair comparisons. This profiling system enables dynamic, context-aware model selection based on the subject composition of each query, ensuring that the most suitable expert LLM is invoked during inference.

3.3 Multi-Agent Collaboration

S-DAG Generation. During inference, given an input question, we employ the trained GNN model to refine both the question embedding and subject node features. The node and edge classifiers (MLP modules) simultaneously predict relevant subject nodes and their dependencies. This inference procedure mirrors the training phase described in Section 3.2. The resulting S-DAG $\mathcal{G}_{\mathcal{Q}}$ is a pruned subgraph of the fully connected subject graph \mathcal{G} , retaining only the most relevant subjects and directed relationships. The complete S-DAG generation process is presented in Alg. 1.

Why is S-DAG suitable for guiding multi-agent subject-specific reasoning? The S-DAG captures both hierarchical and interdependent relationships among subjects. Dominant subjects represent the central focus of reasoning, while supporting subjects—linked via directed edges—provide complementary knowledge. This structured representation naturally defines a collaboration mechanism among expert agents: support agents supply contextual input that enriches the reasoning of dominant agents. Further discussion of the theoretical motivation and construction principles is provided in appendix.

Multi-Agent Information Flow over S-DAG. Given the S-DAG generated in the previous step, the next is to select the appropriate expert LLMs based on their subject proficiency. Based on the LLM subject capability profile process, we match each subject node in the S-DAG to an expert LLM specializing in that domain. If C_{ij} represents the performance score of LLM M_i on subject s_j , the LLM selection for a particular subject s_j could be expressed as:

$$M_j = \arg \max_{M_i} C_{ij}. \quad (4)$$

The multi-agent collaboration mechanism is defined by the directed relationships encoded in the S-DAG, as illustrated in Figure 2. Each edge in the graph represents an information flow dependency between two subject domains,

| Dataset | Train Set | Test Set | Avg. Subject/Q |
|----------|-----------|----------|----------------|
| MMLU-Pro | 1173 | 503 | 4.4 |
| GPQA | 302 | 129 | 3.8 |
| MedMCQA | 396 | 169 | 3.5 |

Table 2: Overview of the curated multi-subject datasets used in our experiments. We manually select questions that span multiple subject areas to better evaluate heterogeneous reasoning capabilities. “Avg. Subject/Q” denotes the average number of distinct subjects involved per question, reflecting the interdisciplinary complexity of each dataset.

guiding how the associated LLM agents should collaborate. Specifically, if a query involves subjects s_1 and s_2 with a directed edge from s_1 to s_2 , the output of the agent corresponding to s_1 and the original query jointly serve as the prompt input for the agent corresponding to s_2 , thereby forming a collaborative reasoning pipeline. The prompting strategy that enables this information flow is detailed in appendix. This process can be formalized as:

$$y_j^{\text{out}} = M_j(\{y_i^{\text{out}} | a_{ij} = 1\}, \mathcal{Q}), \quad (5)$$

where y_i^{out} denotes the output of agent M_i associated with subject s_i , $a_{ij} = 1$ indicates subject s_i has a directed edge to s_j in the S-DAG, and \mathcal{Q} represents the original question, which is included as a shared input to all agents. This dynamic and dependency-driven collaboration enables the system to aggregate reasoning results and progressively refine the final answer through multi-agent cooperation.

4 Experiments

4.1 Experiment Setup

Datasets. We evaluate our proposed method on three benchmarks. MMLU-Pro (Wang et al. 2025) is a challenging extension of the MMLU benchmark, covering 14 college-level subjects. GPQA (Rein et al. 2024) is a dataset of graduate-level science questions designed to be difficult. MedMCQA (Pal, Umapathi, and Sankarasubbu 2022) is a collection of medical entrance exam questions across 21 subdomains. To better reflect real-world heterogeneous reasoning scenarios, we preprocess each dataset to select samples involve multiple subject areas. We construct a dedicated dataset, as shown in Table 2. We also construct a profiling set with 200 samples used to evaluate the subject-specific capabilities of LLMs. This dataset preprocessing ensures that our evaluation aligns with the core challenge addressed in this paper: selecting and coordinating multiple expert agents to solve complex, multi-subject reasoning problems. The details of the dataset curation are shown in appendix.

LLM Pool with Various Experts. To enable subject-aware reasoning and fine-grained agent specialization, we construct a pool of domain-specific expert LLMs. Each model is either pretrained or fine-tuned on data aligned with a specific academic or professional domain, such as mathematics, medicine, law, or economics. These models, typically ranging from 7B to 13B parameters, are computation-

| Category | Method | Model | MMLU-Pro | GPQA | MedMCQA | Avg. |
|----------------------------|----------------------------------|--------------|---------------------|---------------------|---------------------|--------------|
| Closed-Source Single Model | CoT (Wei et al. 2022) | GPT4o-mini | 49.42 ± 0.27 | 47.31 ± 0.52 | 78.82 ± 0.35 | 58.52 |
| Open-Source Single Model | CoT | Qwen2.5 72b | 50.81 ± 0.46 | 48.98 ± 0.35 | 80.44 ± 0.47 | 60.08 |
| | CoT | Llama3.3 70b | 51.92 ± 0.39 | 48.83 ± 0.41 | 79.36 ± 0.61 | 60.04 |
| Single-Model | CoT | Qwen2.5 7b | 41.86 ± 0.29 | 44.51 ± 0.58 | 72.07 ± 0.39 | 52.81 |
| | Self-Refine (Madaan et al. 2023) | Qwen2.5 7b | 44.92 ± 0.34 | 43.83 ± 0.12 | 74.58 ± 0.28 | 54.44 |
| | MoE (Zhou et al. 2022) | LLM Pool | 42.57 ± 0.55 | 45.67 ± 0.34 | 75.45 ± 0.49 | 54.56 |
| | GraphRouter (Feng et al. 2025) | LLM Pool | 44.94 ± 0.94 | 46.23 ± 0.82 | 76.92 ± 0.72 | 56.03 |
| Multi-Model | MAD (Du et al. 2023) | Qwen2.5 7b | 45.82 ± 0.13 | 46.81 ± 0.21 | 76.55 ± 0.15 | 56.39 |
| | Symbolic-MoE (Chen et al. 2025) | LLM Pool | 48.13 ± 0.62 | 45.92 ± 0.51 | 78.55 ± 0.61 | 57.53 |
| | S-DAG (Ours) | LLM Pool | 50.98 ± 0.19 | 49.82 ± 0.24 | 78.38 ± 0.35 | 59.73 |

Table 3: Performance comparison of single-model and multi-model approaches on MMLU-Pro, GPQA, and MedMCQA. We compare various baselines across closed-source, open-source, single-agent and multi-agent settings. We **bold** the best results and underline the second-best (excluding methods using bigger or proprietary models).

ally efficient and well-suited for multi-agent composition. In total, we select 14 expert LLMs spanning a broad range of disciplines. Appendix provides details on the expert models, corresponding subject domains, and Hugging Face links. For instance, DeepseekMath is used for mathematics, while BioMistral is assigned to biology. These expert LLMs form the foundation of our multi-agent system, where each agent is instantiated from the most suitable LLM based on the subject assignments derived from the S-DAG and LLM identification process.

Baselines. Our selection of baselines is guided by the goal of evaluating the challenge of heterogeneous reasoning, where a single complex question spans multiple subject domains. We consider two primary categories. First, we evaluate single-model methods to test whether a single, high-performance, general-purpose LLM can effectively handle multi-subject reasoning. Specifically, we include models such as closed-source GPT4o-mini (Hurst et al. 2024), open-source Qwen2.5-72B (Yang et al. 2024) and Llama3.3-70B (Grattafiori et al. 2024). Then, the Mixture-of-Experts (MoE) approaches that dynamically select one expert model; GraphRouter that utilize a GNN to select expert model based on contextual information given a problem. Second, we explore multi-model methods, which leverage specialized expertise through collaborative or modular strategies. These include Multi-Agent Debate (MAD), where multiple agents reason through dialogue (Liang et al. 2024); and Symbolic-MoE (Chen et al. 2025), which conducts skill-level expert selection. These baselines provide a comprehensive framework to assess both generalist and specialist strategies for tackling complex, interdisciplinary reasoning tasks.

Implementation Details. All experiments, including baselines and our proposed method, are conducted using A100 GPUs with 40 GB memory. For large models such as the 70B open-source LLMs and Qwen variants, we use API-based inference, while smaller expert models in the LLM pool are deployed locally. The decoding temperature is set to 0.7, and the maximum output length is fixed at 4096 tokens across all LLMs. Adam optimizer is used to train GNN and MLP models, and the seed is fixed. Results are

averaged across three trials, and we compute the standard deviations as the statistical indicator. Details on model selection and the LLM pool are provided in appendix.

4.2 Results

The main empirical findings are shown in Table 3. We evaluate the performance of various reasoning frameworks across selected samples from MMLU-Pro, GPQA, and MedMCQA, including both single-model and multi-model settings, as well as open-source and closed-source configurations. The average accuracy is acquired by averaging results across the three benchmarks. Below, we highlight key insights derived from this comparative analysis.

Superior Accuracy of S-DAG. Our S-DAG framework achieves the highest average accuracy (59.73%). It consistently outperforms both single-model and multi-model baselines. In particular, compared to single-model approaches that rely on expert selection—such as MoE (54.56%) and GraphRouter (56.03%)—S-DAG demonstrates a significant improvement. This result underscores the benefit of explicitly modeling inter-subject dependencies, rather than relying solely on selection-based mechanisms.

Robustness over Multi-Model Baselines. Our S-DAG also outperforms several competitive multi-agent systems. It exceeds the performance of Symbolic-MoE (57.53%) and MAD (56.39%) in average accuracy. Although Symbolic-MoE shows strong performance on MedMCQA (78.55%), its effectiveness does not consistently translate across other benchmarks. In contrast, S-DAG maintains a well-balanced performance across all three tasks, demonstrating its robustness and adaptability in coordinating specialized agents across diverse reasoning scenarios.

Competitiveness with Large LLMs. Despite being composed of smaller, domain-specific expert models, S-DAG achieves performance competitive with large-scale, monolithic LLMs. It surpasses the closed-source GPT4o-mini (58.52%) and closely matches the performance of open-source leaders such as Qwen2.5 72B (60.08%) and Llama3.3 70B (60.04%). This indicates that structured coordination

| Variant | GNN Coord. | Model Selection | Graph Structure | Avg. Accuracy (%) | Inf. Time (s) | # LLM Calls |
|-------------------------|------------|-----------------|-----------------|-------------------|---------------|-------------|
| w/o GNN, random model | ✗ | ✗ | S-DAG | 41.12 | 14.21 | 5.1 |
| w/ GNN, random model | ✓ | ✗ | S-DAG | 42.19 | 14.82 | 4.1 |
| w/o GNN, profiled model | ✗ | ✓ | S-DAG | 53.51 | 14.53 | 5.1 |
| Fully Connected Graph | ✓ | ✓ | Fully-connected | 57.29 | 38.45 | 8.2 |
| S-DAG (Ours) | ✓ | ✓ | S-DAG | 59.73 | 15.02 | 4.1 |

Table 4: Ablation study on the effects of coordination (GNN), model selection (LLM profiling), and graph structure. Accuracy is averaged across MMLU-Pro, GPQA, and MedMCQA. Inference efficiency is measured via average inference time and the number of LLM calls per instance.

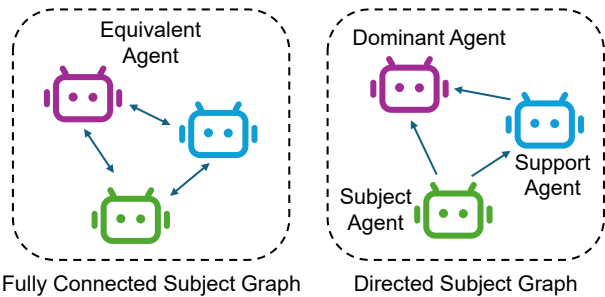


Figure 3: Comparison between the Fully Connected Graph and our S-DAG.

among lightweight experts can rival or exceed the capabilities of significantly larger models, offering a more efficient and cost-effective solution for complex reasoning tasks.

4.3 Ablation Study

To quantify the contribution of individual components in our S-DAG framework, we conduct a structured ablation study, as presented in Table 4. This analysis isolates three key design factors: (1) the presence or absence of GNN-based coordination, (2) the use of subject-aware model selection via LLM profiling versus random assignment, and (3) the impact of graph topology, comparing our sparse S-DAG structure to a fully connected alternative. We evaluate each configuration across two core dimensions: task performance and computational efficiency. Inference time denotes the average wall-clock latency required to process a single multi-subject question, encompassing decomposition, agent invocation, message passing, and response synthesis. LLM call count indicates the average number of distinct language model calls per instance, serving as a proxy for computational and monetary cost.

Effectiveness of GNN Module. Rather than directly prompting the LLM and construct the S-DAG manually, we introduce a GNN module that learns subject interdependencies from data and automatically produces the S-DAG. To evaluate the effectiveness of this design choice, we conduct an ablation study comparing our learned S-DAG with one derived purely from LLM-generated subject weights. The results show that using the LLM-generated S-DAG yields a significantly lower accuracy compared to 59.73% achieved by our GNN-based approach. This performance gap high-

lights the limitations of relying solely on LLM outputs, which can be noisy and inconsistent. In contrast, our GNN module leverages training data to learn robust and context-sensitive subject relationships, resulting in more accurate and concise graph structures for multi-agent reasoning.

Effectiveness of LLM Profile. To assess the benefit of our LLM profiling strategy, we compare our subject-aware model selection approach against a baseline that randomly selects expert models for each subject node in the S-DAG. The results demonstrate that using our profiled subject-model matching significantly improves performance across all benchmarks. This highlights the importance of aligning subject-specific tasks with LLMs that exhibit strong domain expertise. Without profiling, the system risks assigning questions to suboptimal models, which can lead to degraded reasoning quality and inconsistent outputs.

Effectiveness of Directed Acyclic Graph Structure. To evaluate the impact of graph topology, we compare our S-DAG with a fully connected graph (FCG) variant. As illustrated in Figure 3, the FCG allows unrestricted bidirectional communication among all subject nodes, which leads to over-communication and redundant information flow. In contrast, our S-DAG enforces a sparse, hierarchical structure that streamlines reasoning. Empirically, S-DAG achieves higher accuracy (59.73%) while significantly reducing inference time (15.02s) and the number of LLM calls (4.1 per instance). These results demonstrate that fully connected communication is suboptimal for multi-agent reasoning, and that structured, directional coordination leads to more efficient and effective performance.

5 Conclusion

We present S-DAG, a novel framework for heterogeneous reasoning that leverages fine-grained subject-level analysis to guide multi-agent collaboration. By constructing a subject-based directed acyclic graph via a GNN, our method captures subject interdependencies, enabling targeted coordination among specialized LLMs. This structured approach yields efficient, subject-level reasoning. Experiments on three challenging benchmarks show that S-DAG consistently outperforms both single-model and multi-model baselines, achieving competitive accuracy with large-scale LLMs at significantly lower computational cost. Ablation studies further highlight the benefits of the DAG structure in enhancing performance and efficiency.

Acknowledgments

This research was supported by Project P0049179 under the Innovation and Technology Fund – Guangdong–Hong Kong Technology Cooperation Funding Scheme (ITF-TCFS), funded by the Innovation and Technology Commission (Funding Body Ref. No. GHP/386/23SZ).

References

- Cai, W.; Jiang, J.; Wang, F.; Tang, J.; Kim, S.; and Huang, J. 2024. A Survey on Mixture of Experts. *arXiv preprint arXiv:2407.06204*.
- Chen, J. C.-Y.; Yun, S.; Stengel-Eskin, E.; Chen, T.; and Bansal, M. 2025. Symbolic Mixture-of-Experts: Adaptive Skill-based Routing for Heterogeneous Reasoning. *arXiv preprint arXiv:2503.05641*.
- Chen, L.; Zaharia, M.; and Zou, J. 2024. FrugalGPT: How to Use Large Language Models While Reducing Cost and Improving Performance. *Transactions on Machine Learning Research*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (long And Short Papers)*, 4171–4186.
- Du, Y.; Li, S.; Torralba, A.; Tenenbaum, J. B.; and Mordatch, I. 2023. Improving Factuality And Reasoning in Language Models Through Multiagent Debate. In *Forty-first International Conference on Machine Learning*.
- Feng, S.; Ding, W.; Liu, A.; Wang, Z.; Shi, W.; Wang, Y.; Shen, Z.; Han, X.; Lang, H.; Lee, C.-Y.; et al. 2025a. When One LLM Drools, Multi-LLM Collaboration Rules. *arXiv preprint arXiv:2502.04506*.
- Feng, S.; Shi, W.; Bai, Y.; Balachandran, V.; He, T.; and Tsvetkov, Y. 2024. Knowledge Card: Filling LLMs’ Knowledge Gaps with Plug-in Specialized Language Models. In *The Twelfth International Conference on Learning Representations*.
- Feng, S.; Wang, Z.; Goyal, P.; Wang, Y.; Shi, W.; Xia, H.; Palangi, H.; Zettlemoyer, L.; Tsvetkov, Y.; Lee, C.-Y.; et al. 2025b. Heterogeneous Swarms: Jointly Optimizing Model Roles And Weights for Multi-LLM Systems. *arXiv preprint arXiv:2502.04510*.
- Feng, T.; Shen, Y.; and You, J. 2025. GraphRouter: A Graph-based Router for LLM Selections. In *The Thirteenth International Conference on Learning Representations*.
- Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. 2024. The Llama 3 Herd of Models. *arXiv preprint arXiv:2407.21783*.
- Gu, W.; Han, J.; Wang, H.; Li, X.; and Cheng, B. 2025. Explain-Analyze-Generate: A Sequential Multi-Agent Collaboration Method for Complex Reasoning. In *Proceedings of the 31st International Conference on Computational Linguistics*, 7127–7140.
- Guo, T.; Chen, X.; Wang, Y.; Chang, R.; Pei, S.; Chawla, N. V.; Wiest, O.; and Zhang, X. 2024. Large Language Model Based Multi-agents: A Survey of Progress And Challenges. *arXiv preprint arXiv:2402.01680*.
- Han, S.; Zhang, Q.; Yao, Y.; Jin, W.; Xu, Z.; and He, C. 2024. LLM Multi-agent Systems: Challenges And Open Problems. *arXiv preprint arXiv:2402.03578*.
- Hoffmann, J.; Borgeaud, S.; Mensch, A.; Buchatskaya, E.; Cai, T.; Rutherford, E.; de Las Casas, D.; Hendricks, L. A.; Welbl, J.; Clark, A.; et al. 2022. Training Compute-optimal Large Language Models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, 30016–30030.
- Huang, W.; Xia, F.; Shah, D.; Driess, D.; Zeng, A.; Lu, Y.; Florence, P.; Mordatch, I.; Levine, S.; Hausman, K.; et al. 2023. Grounded Decoding: Guiding Text Generation With Grounded Models for Embodied Agents. *Advances in Neural Information Processing Systems*, 36: 59636–59661.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o System Card. *arXiv preprint arXiv:2410.21276*.
- Ke, Z.; Jiao, F.; Ming, Y.; Nguyen, X.-P.; Xu, A.; Long, D. X.; Li, M.; Qin, C.; Wang, P.; Savarese, S.; et al. 2025. A Survey of Frontiers in LLM Reasoning: Inference Scaling, Learning To Reason, And Agentic Systems. *arXiv preprint arXiv:2504.09037*.
- Kojima, T.; Gu, S. S.; Reid, M.; Matsuo, Y.; and Iwasawa, Y. 2022. Large Language Models Are Zero-shot Reasoners. *Advances in Neural Information Processing Systems*, 35: 22199–22213.
- Li, D.; Tan, Z.; Qian, P.; Li, Y.; Chaudhary, K. S.; Hu, L.; and Shen, J. 2024a. Smoa: Improving Multi-agent Large Language Models With Sparse Mixture-of-agents. *arXiv preprint arXiv:2411.03284*.
- Li, J.; Wang, Y.; Zhu, P.; Lin, W.; Yao, X.; and Hu, Q. 2025. Graphs Help Graphs: Multi-Agent Graph Socialized Learning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Li, Y.; Du, Y.; Zhang, J.; Hou, L.; Grabowski, P.; Li, Y.; and Ie, E. 2024b. Improving Multi-Agent Debate With Sparse Communication Topology. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 7281–7294.
- 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.
- Lin, W.; Lan, H.; He, H.; and Li, B. 2025. Graph-Relational Federated Learning: Enhanced Personalization and Robustness. *IEEE Transactions on Dependable and Secure Computing*.
- 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*.

- Madaan, A.; Tandon, N.; Gupta, P.; Hallinan, S.; Gao, L.; Wiegrefe, S.; Alon, U.; Dziri, N.; Prabhunoye, S.; Yang, Y.; et al. 2023. Self-refine: Iterative Refinement With Self-feedback. *Advances in Neural Information Processing Systems*, 36: 46534–46594.
- Masoudnia, S.; and Ebrahimpour, R. 2014. Mixture of Experts: A Literature Survey. *Artificial Intelligence Review*, 42: 275–293.
- Mumuni, A.; and Mumuni, F. 2025. Large Language Models for Artificial General Intelligence (AGI): A Survey of Foundational Principles And Approaches. *arXiv preprint arXiv:2501.03151*.
- Pal, A.; Umaphathi, L. K.; and Sankarasubbu, M. 2022. Medmcqa: A Large-scale Multi-subject Multi-choice Dataset for Medical Domain Question Answering. In *Conference on Health, Inference, And Learning*, 248–260. PMLR.
- Rein, D.; Hou, B. L.; Stickland, A. C.; Petty, J.; Pang, R. Y.; Dirani, J.; Michael, J.; and Bowman, S. R. 2024. Gpqa: A Graduate-level Google-proof Q&a Benchmark. In *First Conference on Language Modeling*.
- Talebirad, Y.; and Nadiri, A. 2023. Multi-agent Collaboration: Harnessing the Power of Intelligent Llm Agents. *arXiv preprint arXiv:2306.03314*.
- Tao, M.; Zhao, D.; and Feng, Y. 2025. Chain-of-Discussion: A Multi-Model Framework for Complex Evidence-Based Question Answering. In *Proceedings of the 31st International Conference on Computational Linguistics*, 11070–11085.
- Wang, J.; Wang, J.; Athiwaratkun, B.; Zhang, C.; and Zou, J. 2024. Mixture-of-agents Enhances Large Language Model Capabilities. *arXiv preprint arXiv:2406.04692*.
- Wang, Y.; Ma, X.; Zhang, G.; Ni, Y.; Chandra, A.; Guo, S.; Ren, W.; Arulraj, A.; He, X.; Jiang, Z.; et al. 2025. MMLU-Pro: A More Robust And Challenging Multi-Task Language Understanding Benchmark. *Advances in Neural Information Processing Systems*, 37: 95266–95290.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought Prompting Elicits Reasoning in Large Language Models. *Advances in Neural Information Processing Systems*, 35: 24824–24837.
- 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*.
- Xin, A.; Liu, J.; Yao, Z.; Lee, Z.; Cao, S.; Hou, L.; and Li, J. 2024. AtomR: Atomic Operator-Empowered Large Language Models for Heterogeneous Knowledge Reasoning. *arXiv preprint arXiv:2411.16495*.
- Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; Huang, F.; Wei, H.; et al. 2024. Qwen2. 5 Technical Report. *arXiv preprint arXiv:2412.15115*.
- Yao, X.; Wang, Y.; Zhu, P.; Lin, W.; Zhao, R.; Guo, Z.; Li, W.; and Hu, Q. 2025. Socialized Coevolution: Advancing a Better World through Cross-Task Collaboration. In *Forty-second International Conference on Machine Learning*.
- Yin, Z.; Sun, Q.; Chang, C.; Guo, Q.; Dai, J.; Huang, X.-J.; and Qiu, X. 2023. Exchange-of-Thought: Enhancing Large Language Model Capabilities Through Cross-Model Communication. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 15135–15153.
- Yue, M. 2025. A Survey of Large Language Model Agents for Question Answering. *arXiv preprint arXiv:2503.19213*.
- Yue, Y.; Zhang, G.; Liu, B.; Wan, G.; Wang, K.; Cheng, D.; and Qi, Y. 2025. Masrouter: Learning To Route Llms for Multi-agent Systems. *arXiv preprint arXiv:2502.11133*.
- Zhang, G.; Yue, Y.; Li, Z.; Yun, S.; Wan, G.; Wang, K.; Cheng, D.; Yu, J. X.; and Chen, T. 2024a. Cut the Crap: An Economical Communication Pipeline for Llm-based Multi-agent Systems. *arXiv preprint arXiv:2410.02506*.
- Zhang, Y.; Mao, S.; Ge, T.; Wang, X.; de Wynter, A.; Xia, Y.; Wu, W.; Song, T.; Lan, M.; and Wei, F. 2024b. Llm As A Mastermind: A Survey of Strategic Reasoning With Large Language Models. *arXiv preprint arXiv:2404.01230*.
- Zhang, Y.; Sun, R.; Chen, Y.; Pfister, T.; Zhang, R.; and Arik, S. 2024c. Chain of Agents: Large Language Models Collaborating on Long-context Tasks. *Advances in Neural Information Processing Systems*, 37: 132208–132237.
- Zhou, Y.; Lei, T.; Liu, H.; Du, N.; Huang, Y.; Zhao, V.; Dai, A. M.; Le, Q. V.; Laudon, J.; et al. 2022. Mixture-of-experts With Expert Choice Routing. *Advances in Neural Information Processing Systems*, 35: 7103–7114.
- Zhuang, Y.; Yu, Y.; Wang, K.; Sun, H.; and Zhang, C. 2023. Toolqa: A Dataset for Llm Question Answering With External Tools. *Advances in Neural Information Processing Systems*, 36: 50117–50143.
- Zhuge, M.; Wang, W.; Kirsch, L.; Faccio, F.; Khizbullin, D.; and Schmidhuber, J. 2024. Gptswarm: Language Agents As Optimizable Graphs. In *Forty-first International Conference on Machine Learning*.