

Assemble Your Crew: Automatic Multi-agent Communication Topology Design via Autoregressive Graph Generation

Shiyuan Li¹, Yixin Liu¹, Qingsong Wen², Chengqi Zhang³, Shirui Pan^{1*}

¹Griffith University

²Squirrel Ai Learning

³Hong Kong Polytechnic University

{li.shiy511, qingsonedu}@gmail.com, {yixin.liu, s.pan}@griffith.edu.au, chengqi.zhang@polyu.edu.hk

Abstract

Multi-agent systems (MAS) based on large language models (LLMs) have emerged as a powerful solution for dealing with complex problems across diverse domains. The effectiveness of MAS is critically dependent on its collaboration topology, which has become a focal point for automated design research. However, existing approaches are fundamentally constrained by their reliance on a template graph modification paradigm with a predefined set of agents and hard-coded interaction structures, significantly limiting their adaptability to task-specific requirements. To address these limitations, we reframe MAS design as a conditional autoregressive graph generation task, where both the system composition and structure are designed jointly. We propose **ARG-DESIGNER**, a novel autoregressive model that operationalizes this paradigm by constructing the collaboration graph from scratch. Conditioned on a natural language task query, ARG-DESIGNER sequentially and dynamically determines the required number of agents, selects their appropriate roles from an extensible pool, and establishes the optimal communication links between them. This generative approach creates a customized topology in a flexible and extensible manner, precisely tailored to the unique demands of different tasks. Extensive experiments across six diverse benchmarks demonstrate that ARG-DESIGNER not only achieves state-of-the-art performance but also enjoys significantly greater token efficiency and enhanced extensibility.

Code — <https://github.com/Shiy-Li/ARG-Designer>

Introduction

Agents built on large language models (LLMs) have demonstrated impressive capabilities in tackling complex tasks across domains, including code generation, data analysis, decision-making, and question answering (Zhu et al. 2024; Li et al. 2024a; Song et al. 2023; Wang et al. 2024; Zhong, Wang, and Shang 2024; Tan et al. 2025; Miao et al. 2025). To overcome the limitations of a single agent in tackling more complex tasks, the research interests have increasingly shifted towards multi-agent systems (MAS), which unlock new potential through collaborative interactions among agents with diverse capabilities and roles. Central to MAS is

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

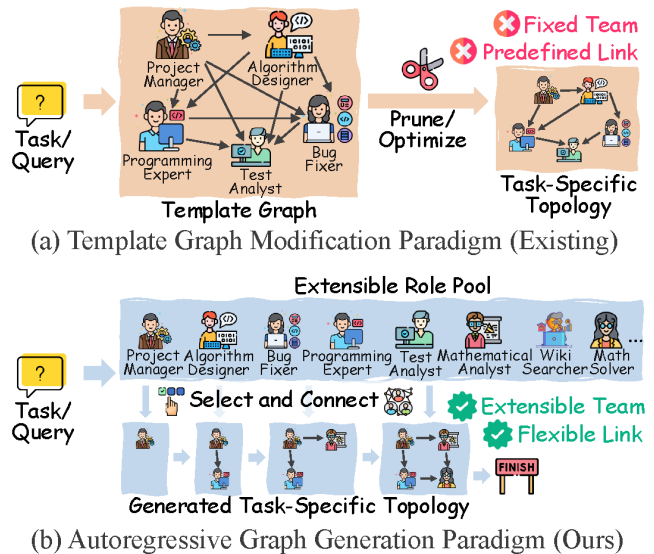


Figure 1: The comparison of two paradigms: (a) template graph modification and (b) autoregressive graph generation.

its **collaboration topology**, a graph that defines how agents with roles are structured and how they exchange information (Liu et al. 2025). A growing body of evidence, spanning from sequential reasoning pipelines to debate-based approaches, demonstrates that MAS performance varies dramatically depending on how inter-agent communication is architected (Zhang et al. 2025c; Zhou et al. 2025). Therefore, designing an effective collaboration graph tailored to specific tasks becomes a critical research challenge.

Early research on MAS topology design focused on static and manually designed graphs, such as chains that enforce sequential workflow (Wei et al. 2022; Hong et al. 2024), trees that enable structured deliberation (Yao et al. 2023), and fully connected graphs that ensure sufficient communication. Although these canonical collaboration topologies can facilitate effective coordination in specific scenarios, the inherent rigidity of these fixed topologies limits their adaptability across diverse tasks, resulting in sub-optimal performance. To enhance flexibility and efficiency, a more recent line of work (Zhuge et al. 2024; Shen et al. 2025) focuses on

adaptively constructing task-specific communication structures using *graph learning models* (Pan et al. 2025b,a; Fu et al. 2025a,b). For example, AgentPrune (Zhang et al. 2025a) and AgentDrop (Wang et al. 2025) learn to create sparse and task-specific graphs by pruning connections or agents from a predefined template topology. More advanced approaches like G-Designer (Zhang et al. 2025b) follow the paradigm of graph structure learning, which uses a graph autoencoder to learn efficient collaboration structures in a task-adaptive manner.

Despite their varied designs, existing graph learning-based methods often follow a shared paradigm: *template graph modification* (Fig. 1a). That is, they typically start from a fixed communication template based on a predefined set of agents and hard-coded interaction structures, and apply learnable adjustments, such as edge reweighting or pruning, to adapt the topology to specific tasks (Zhang et al. 2025a,b). Despite offering reasonable adaptability in constrained settings, this paradigm exhibits two inherent limitations. **Limitation 1: Redundant Composition.** To ensure structural flexibility, template graphs are often initialized with numerous agent roles and densely connected edges, many of which are unnecessary for a specific task. Even with pruning mechanisms, irrelevant agents or connections may be retained in the learned task-specific topology, leading not only to reduced efficiency but also to potential sub-optimal decision-making during execution. **Limitation 2: Limited Extensibility.** In the fast-evolving field of LLM-based agents, a massive number of new agent functionalities are emerging with increasing frequency. However, trained on a fixed template graph, the existing methods struggle to generalize to scenarios with dynamic agent sets or evolving collaboration needs. Meanwhile, it would be prohibitively expensive to build a large-scale template graph that covers all possible agent roles and interaction patterns, and then prune it to a suitable task-specific topology. Given the above limitations, a natural question arises: *Going beyond template graph modification, can we design a more flexible and extensible paradigm for collaboration topology construction?*

To seek the answer to the above question, we draw inspiration from real-world practices of recruiting teams for complex tasks. Rather than starting with a fully staffed team where every possible member is onboarded from the beginning, real-world teams are usually formed incrementally, with members added based on expertise, availability, and evolving task needs. This practical pattern inspires us to explore *autoregressive graph generation* (Fig. 1b) as a more promising paradigm for collaboration topology construction. Unlike pruning from a predefined overcomplete structure, the new generation paradigm constructs the collaboration graph from scratch by progressively selecting appropriate agents. Such an incremental procedure naturally avoids redundant agent-role compositions during the design process, which naturally addresses **Limitation 1**. Moreover, by discarding the fixed template, the generative paradigm enables dynamic expansion of the agent pool, with only linear computational cost during the node generation phase. This merit enhances the extensibility of collaboration graph construction and thus alleviates **Limitation 2**.

Building upon the new paradigm, in this paper, we propose **ARG-DESIGNER**, a novel **AutoRegressive Graph** generation model that acts as a MAS topology **Designer**. Conditioned on a natural language task query, ARG-DESIGNER constructs the entire collaboration graph from scratch by iteratively generating each node (i.e., agent) along with its corresponding edges (i.e., communication links) to previously generated nodes. Compared to prior approaches, ARG-DESIGNER provides enhanced flexibility and scalability with respect to the number of agents, the variety of agent roles, and the richness of potential interactions. To train our generative model, we design a curriculum learning strategy that starts with denser communication topologies to ease the cold-start problem, and gradually transitions to sparser, pruned graphs for fine-tuning, encouraging the model to generalize to minimal yet effective structures. Extensive experiments on six benchmarks demonstrate that our method achieves state-of-the-art effectiveness, communication efficiency, and robustness.

Problem Formulation

In this section, we introduce the graph-based modeling of MAS, and then formulate MAS topology design as an autoregressive graph generation problem.

MAS as a Collaboration Graph. We model a MAS as a *collaboration graph*, a directed acyclic graph (DAG) $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ that outlines the system architecture and the flow of information among its components. The nodes $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$ represent the set of agents, where each agent v_i is an instance of an LLM endowed with a specific role $r_i \in R$ that dictates its function and expertise. It also maintains an internal state $s_i \in S$, which serves as a memory of its past actions and interactions. The edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ define the directed communication pathways. An edge $e_{ji} = (v_j, v_i)$ signifies that agent v_i is a designated recipient of information from agent v_j . The set of direct predecessors of agent v_i is denoted by $\mathcal{N}_{in}(v_i) = \{v_j \mid (v_j, v_i) \in \mathcal{E}\}$.

MAS Collaboration Protocol. Given a collaboration graph \mathcal{G} , the MAS addresses a user query \mathcal{Q} by executing a multi-step collaboration protocol. This protocol governs how information is processed and passed between agents, unfolding over a series of communication rounds (Pan et al. 2026; Zhang et al. 2024; Cai et al. 2024). Unlike traditional GNNs message passing, the operational sequence for agent activation within each round is determined by a topological sort of the nodes, ensuring that an agent is activated only after its prerequisite inputs are available (Zhao et al. 2025; Chen et al. 2025; Zhuang et al. 2025; Li et al. 2024b; Cai et al. 2025). This process can be executed for K rounds to allow for iterative refinement. In each round $k \in \{1, \dots, K\}$, an agent v_i generates its response $m_i^{(k)}$ by invoking its language model with a dynamically constructed prompt $\mathcal{P}_i^{(k)}$:

$$m_i^{(k)} = \text{LLM}_i(\mathcal{P}_i^{(k)}). \quad (1)$$

where the prompt integrates the properties of the agent with outputs of its predecessors from the previous round:

$$\mathcal{P}_i^{(k)} = f(\underbrace{r_i, s_i}_{\text{System}}, \underbrace{\mathcal{Q}, \{m_j^{(k-1)} \mid v_j \in \mathcal{N}_{in}(v_i)\}}_{\text{User}}). \quad (2)$$

After K rounds, the final output \mathcal{O} is obtained by aggregating the final-round responses from some or all agents:

$$\mathcal{O} = \text{Aggregate}(\{m_i^{(K)} \mid v_i \in \mathcal{V}\}), \quad (3)$$

where the aggregation strategy $\text{Aggregate}(\cdot)$ varies across implementations. Common strategies include majority voting, delegating the final decision to a specific terminal agent, or selecting the output from the last agent in the execution order. The number of communication rounds K can be either predefined or adaptively determined via early-stopping. **MAS Topology Design as a Graph Generation Task.** The automatic task-specific construction of MAS topologies is a key challenge and research frontier. Traditional automated approaches that start from a large, predefined template graph, analogous to a fully-staffed team with every possible role, suffer from redundancy and limited extensibility. Drawing inspiration from real-world practices of building expert teams incrementally, we reframe the problem from modifying a fixed template to generating a bespoke graph from scratch. Instead of navigating the enormous graph space \mathbb{G} with an expensive utility function $\phi(\text{Execute}(\mathcal{G}, \mathcal{Q}))$, we propose to learn a **conditional generative model**, i.e., $P(\mathcal{G}|\mathcal{Q}, \mathcal{R})$, where \mathcal{R} is an extensive agent role pool. This model directly captures the relationship between a task query and the principles of effective collaboration, aiming to find the optimal communication topology \mathcal{G}^* that is most probable under this learned distribution:

$$\mathcal{G}^* = \arg \max_{\mathcal{G} \in \mathbb{G}} P(\mathcal{G}|\mathcal{Q}, \mathcal{R}). \quad (4)$$

Compared to modifying a predefined template graph, the generative formulation offers a more flexible, extensible, and scalable approach for constructing high-quality MAS topologies.

Autoregressive Graph Generation. To make this topology generation process more tractable, we formulate it as an autoregressive graph generation problem. This formulation decomposes the intractable joint probability of an entire graph into a tractable sequence of conditional probabilities. The graph is constructed incrementally, where each step involves adding a new node and its corresponding edges, conditioned on the partial graph built so far.

Formally, this factorization is expressed as:

$$P(\mathcal{G}|\mathcal{Q}, \mathcal{R}) = \prod_{i=1}^{|\mathcal{V}|} \underbrace{P(v_i|\mathcal{G}_{<i}, \mathcal{Q}, \mathcal{R})}_{\text{Node Generation}} \cdot \prod_{j=1}^{i-1} \underbrace{P(e_{ji}|v_i, \mathcal{G}_{<i}, \mathcal{Q})}_{\text{Edge Generation}}, \quad (5)$$

where $\mathcal{G}_{<i}$ represents the subgraph of the first $i - 1$ nodes. The generation process at each step i thus involves two key actions: **node generation**, predicting the role of the next agent to add, and **edge generation**, establishing its connections from existing agents. This formulation provides significant flexibility, enabling the model to dynamically determine the total number of agents by learning to sample a special END token, and to model complex structural dependencies by conditioning on the generation history.

Discussion. This generative approach provides several key advantages over traditional template-based methods, as

Method	Task-Adaptive	Variable Size	Flexible Roles
Manual Design	✗	✗	✗
AgentDropout	✓	✗	–
AgentPrune	✓	–	–
G-Designer	✓	✗	–
ARG-DESIGNER (ours)	✓	✓	✓

Table 1: Comparison of degrees of freedom in MAS design paradigms. The icons ✓, –, and ✗ represent full, partial, and no support for each capability, respectively.

summarized in Table 1. **① Task-Adaptive Construction.** By conditioning on the task query, the model constructs a bespoke collaboration graph from scratch, avoiding the rigidity and one-size-fits-all limitations of predefined graphs. **② Dynamic and Extensible Composition.** The model dynamically determines the necessary number of agents and selects their roles from an extensible pool, ensuring the MAS composition is precisely tailored to the task needs and can easily incorporate new agent capabilities. **③ Tractable Generation.** The autoregressive factorization transforms the intractable problem of generating a whole graph into a sequence of simple, conditional steps, making the learning process both manageable and scalable.

ARG-DESIGNER for MAS Topology Design

Based on the autoregressive graph generation paradigm, this section instantiates the proposed method, **ARG-DESIGNER**, which is specifically crafted for MAS topology generation. We first introduce the model architecture designed to implement the sequential generation process, and then describe the training and inference strategies to guide ARG-DESIGNER toward generating both functionally correct and structurally efficient collaboration graphs.

Model Architecture

Following the autoregressive generation paradigm, ARG-DESIGNER constructs collaboration graphs step-by-step. ARG-DESIGNER employs a hierarchical architecture based on gated recurrent units (GRUs), which are well-suited for sequence modeling due to their effectiveness in capturing long-range dependencies while maintaining computational efficiency. The architecture separates the generation model into two sub-components: a *node generator* to select agent roles and an *edge generator* to build communication links. An overview of the model architecture is depicted in Fig. 2a. **Input Representation.** Before generation begins, ARG-DESIGNER encodes all textual conditioning information (i.e., the task query and available agent roles) into dense vector representations. Specifically, the task query \mathcal{Q} is mapped into a fixed-dimensional vector $\mathbf{f}_{\mathcal{Q}} \in \mathbb{R}^d$ by a pre-trained BERT-like sentence encoder followed by a feed-forward network (FFN) with Layer Normalization (LN):

$$\mathbf{f}_{\mathcal{Q}} = \text{FFN}(\text{LN}(\text{SentenceEncoder}(\mathcal{Q}))). \quad (6)$$

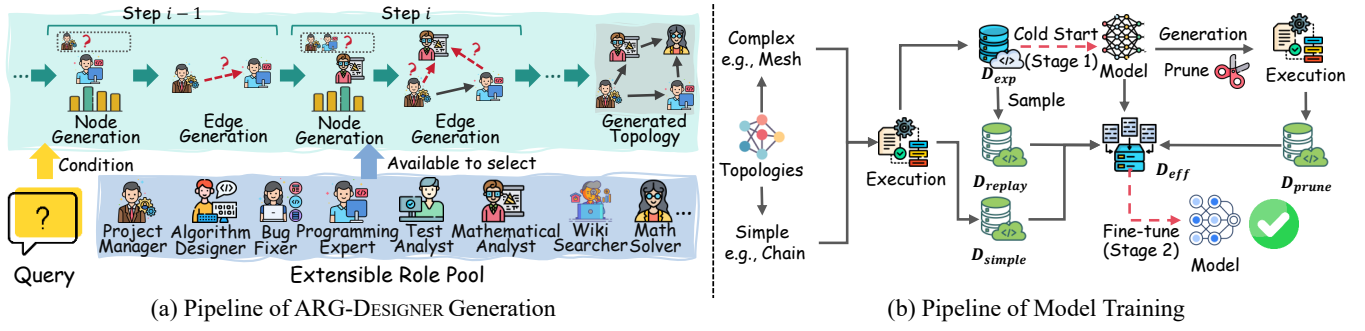


Figure 2: The pipeline of ARG-DESIGNER, including (a) MAS communication topology generation and (b) model training.

Similarly, each available agent role $r_k \in \mathcal{R}$ is converted into an embedding \mathbf{z}_{r_k} . These pre-computed embeddings are collected into a role embedding matrix $\mathbf{Z} \in \mathbb{R}^{|\mathcal{R}| \times d}$, serving as the knowledge base of available agents.

Node Generation. At each step i , the node generator selects the role for the next agent, i.e., v_i . ARG-DESIGNER first models the context information by combining task information with the generation history. A dedicated GRU, GRU_{prev} , is employed to aggregate the role embeddings of all preceding agents to form a historical embedding $\mathbf{f}_{\text{hist}}^{(i)}$:

$$\mathbf{f}_{\text{hist}}^{(i)} = \text{GRU}_{\text{prev}}([\mathbf{z}_{r_1}, \mathbf{z}_{r_2}, \dots, \mathbf{z}_{r_{i-1}}]). \quad (7)$$

Then, we fuse the historical embedding $\mathbf{f}_{\text{hist}}^{(i)}$ with the task embedding $\mathbf{f}_{\mathcal{Q}}$ via a dynamic gate to produce the context embedding $\mathbf{f}_{\text{cont}}^{(i)}$:

$$\mathbf{f}_{\text{cont}}^{(i)} = (1 - \mathbf{g}_i) \cdot \mathbf{f}_{\text{hist}}^{(i)} + \mathbf{g}_i \cdot \mathbf{f}_{\mathcal{Q}}, \quad \mathbf{g}_i = \sigma \left(\frac{\mathbf{f}_{\text{hist}}^{(i)} \cdot \mathbf{f}_{\mathcal{Q}}}{\sqrt{d}} \right), \quad (8)$$

where σ denotes the sigmoid function. This context, along with an edge feature vector $\mathbf{f}_{\text{edge}}^{(i)}$ (i.e., a vector encoding the connectivity pattern of the previously added node v_{i-1}), are concatenated into an input vector $[\mathbf{f}_{\text{cont}}^{(i)}, \mathbf{f}_{\text{edge}}^{(i)}]$. Then, another GRU module, GRU_{node} , updates its hidden state $\mathbf{h}_{\text{node}}^{(i)}$, which captures the full generation condition:

$$\mathbf{h}_{\text{node}}^{(i)} = \text{GRU}_{\text{node}}(\text{MLP}_{\text{node}}([\mathbf{f}_{\text{cont}}^{(i)}, \mathbf{f}_{\text{edge}}^{(i)}]), \mathbf{h}_{\text{node}}^{(i-1)}). \quad (9)$$

To preserve the extensible property of ARG-DESIGNER during agent role selection, we use a metric learning-based module for node generation. Concretely, the hidden state $\mathbf{h}_{\text{node}}^{(i)}$ is projected into a ‘‘node intent’’ embedding. Then, the node prediction scores $s_{\text{node}}^{(i)}$ will be acquired by a dot-product operation with projected role embeddings:

$$s_{\text{node}}^{(i)} = \text{MLP}_{\text{pred.n}}(\mathbf{h}_{\text{node}}^{(i)}) \cdot \text{MLP}_{\text{role}}([\mathbf{Z}, \mathbf{z}_{\text{end}}]), \quad (10)$$

where, \mathbf{z}_{end} is a learnable embedding for ending token END, which signals the termination of the generation process. Finally, we can obtain the predicted probability as follows:

$$P(v_i | \mathcal{G}_{<i}, \mathcal{Q}, \mathcal{R}) = \text{Softmax}(s_{\text{node}}^{(i)}), \quad (11)$$

where the $\text{Softmax}(\cdot)$ function converts the scores into a probability distribution.

Discussion of Extensibility: The design of the node generator in ARG-DESIGNER allows new agent roles to be added at inference time without retraining. When new roles are introduced, we can extend the role embedding matrix \mathbf{Z} by appending new role-specific embedding rows. Since $s_{\text{node}}^{(i)}$ is produced by a metric learning-based retrieval mechanism rather than a fixed-dimensional classifier, the model can flexibly select from an expanded set of roles based on similarity in the shared embedding space. This design ensures that ARG-DESIGNER remains extensible and adaptable to evolving agent pools, which well fits the real-world scenarios where new agents with novel functionalities are frequently introduced to meet emerging task demands.

Edge Generation. Once agent node v_i is chosen, the edge generator determines its incoming connections from existing agents $\{v_1, \dots, v_{i-1}\}$. Here, we use a dedicated GRU, GRU_{edge} to model this sequential process. Its hidden state is initialized from the final state of the node-level GRU, $\mathbf{h}_{\text{node}}^{(i)}$, serving as the condition of edge prediction:

$$\mathbf{h}_{\text{edge}}^{(i,0)} = \text{MLP}_{\text{node2edge}}(\mathbf{h}_{\text{node}}^{(i)}). \quad (12)$$

After that, the model iterates through previously predicted nodes v_j ($j = 1, \dots, i-1$). At each sub-step, the edge GRU updates its hidden state based on the embedding of the previous edge decision:

$$\mathbf{h}_{\text{edge}}^{(i,j)} = \text{GRU}_{\text{edge}}(\text{MLP}_{\text{edge}}(\mathbf{e}^{(j-1,i)}), \mathbf{h}_{\text{edge}}^{(i,j-1)}), \quad (13)$$

where $\mathbf{e}^{(j-1,i)}$ is a one-hot vector representing the previous decision on whether to form an edge from node v_{j-1} to v_i . Following that, the updated state $\mathbf{h}_{\text{edge}}^{(i,j)}$ is passed through an output MLP to predict the score:

$$s_{\text{edge}}^{(i,j)} = \text{MLP}_{\text{pred.e}}(\mathbf{h}_{\text{edge}}^{(i,j)}). \quad (14)$$

We can then obtain the probability of edge $e_{j,i}$ by:

$$P(e_{j,i} = 1 | v_i, \mathcal{G}_{<i}, \mathcal{Q}) = \text{Sigmoid}(s_{\text{edge}}^{(i,j)}). \quad (15)$$

Training and Inference Strategy

Data Construction of Curriculum Learning. To build a powerful graph generator for MAS communication topol-

ogy, we set up two key objectives: ❶ **Functional correctness**, which ensures that the generated topology enables agents to collaboratively complete the given task, with all necessary roles and interactions properly instantiated; ❷ **Communicational efficiency**, which encourages the generated topology to be lightweight, sparse, and compact, by minimizing redundant links or agents. To reach these goals, we design a **curriculum learning** strategy that constructs training data for two-stage training. In the first stage, we create an *exploration dataset* to teach the model to generate correct and diverse topologies; then, an *efficiency dataset* is built to guide the model to design simple yet communication-efficient topologies.

Formally, we define a dataset as $\mathcal{D} = \{(\mathcal{G}_k, \mathcal{Q}_k)\}_{k=1}^M$ that provides strong supervision on what constitutes an effective collaboration graph \mathcal{G}_k for a given task query \mathcal{Q}_k . Since manually authoring such optimal task-graph pairs is infeasible, we propose to construct high-quality datasets in an automatically synthetic manner.

The first phase focuses on creating an exploration dataset (\mathcal{D}_{exp}) for the cold start training of the model, which aims to teach the model to create effective communication graphs. This dataset is formed by pairing tasks from a base set \mathbb{Q}_{base} with resource-rich, complex configurations from a configuration space $\mathbb{C}_{\text{complex}}$, and retaining only empirically successful instances. Formally, this process is defined as:

$$\mathcal{D}_{\text{exp}} = \{(G(c), \mathcal{Q}) \mid S(\mathcal{Q}, G(c)) = 1\}, \quad (16)$$

where $\mathcal{Q} \in \mathbb{Q}_{\text{base}}$, $c \in \mathbb{C}_{\text{complex}}$ is a configuration blueprint specifying the high-level attributes of a graph, such as topology and agent count. For example, a configuration could be defined as $c = (\text{'star'}, 6, \mathcal{R})$. Note that the ‘‘agent_num’’ parameter here does not limit the core capability of ARG-DESIGNER; rather, it constrains the data generation process to ensure a rich diversity of graph structures in the training data, from which the model learns generalizable collaborative patterns, not fixed sizes. $G(\cdot)$ is a deterministic function that maps a configuration c to a specific graph instance \mathcal{G} , and $S(\mathcal{Q}, \mathcal{G}) \in \{0, 1\}$ is an indicator function that verifies the empirical success of the graph for the given task. This initial phase allows the model to learn fundamental collaborative patterns in an unconstrained and resource-abundant environment, which ensures the model has generalizable graph construction abilities.

In the second phase, an efficiency dataset \mathcal{D}_{eff} is built to teach the model to generate more economical graphs. \mathcal{D}_{eff} is a heterogeneous mixture composed of three sources:

$$\mathcal{D}_{\text{eff}} = \mathcal{D}_{\text{simple}} \cup \mathcal{D}_{\text{pruned}} \cup \mathcal{D}_{\text{replay}}, \quad (17)$$

\mathcal{D}_{eff} includes natively efficient graphs from simple configurations $\mathcal{D}_{\text{simple}}$, successful graphs derived from pruning the dense structures generated by the Phase 1 model $\mathcal{D}_{\text{pruned}}$, and a subset of the original exploration data $\mathcal{D}_{\text{replay}}$ to prevent catastrophic forgetting. More specifically: ❶ $\mathcal{D}_{\text{simple}}$, which contains task-graph pairs generated from a predefined set of minimal, manually-verified configurations known to be efficient; ❷ $\mathcal{D}_{\text{pruned}}$, created by taking the overly complex but functional graphs from \mathcal{D}_{exp} , systematically removing individual nodes or edges, and retaining any pruned versions that

still successfully complete the task; ❸ $\mathcal{D}_{\text{replay}}$, a random subset of the initial \mathcal{D}_{exp} dataset, included to prevent the model from forgetting the fundamental patterns learned in the first phase. With the carefully designed dataset for the second training stage, ARG-DESIGNER learns to strike a desirable balance between correctness and simplicity, producing high-quality topologies with minimal redundancy.

Model Training. The training objective is to maximize the conditional log-likelihood of the ground-truth graphs in a given dataset \mathcal{D} . The model parameters θ are optimized by minimizing the negative log-likelihood (NLL) loss:

$$\mathcal{L}(\theta) = - \sum_{(\mathcal{G}, \mathcal{Q}) \in \mathcal{D}} \log P_{\theta}(\mathcal{G}|\mathcal{Q}). \quad (18)$$

Following the autoregressive factorization, the above loss is decomposed into a node generation term and an edge generation term. The final training loss is a weighted sum of these two terms:

$$\mathcal{L}_{\text{total}} = \alpha \cdot \mathcal{L}_{\text{node}} + (1 - \alpha) \cdot \mathcal{L}_{\text{edge}}, \quad (19)$$

where $\alpha \in [0, 1]$ is a hyperparameter balancing the two objectives. The individual loss terms are defined as the NLL over all node and edge generation steps, respectively:

$$\mathcal{L}_{\text{node}} = - \sum_{(\mathcal{G}, \mathcal{Q}) \in \mathcal{D}} \sum_{i=1}^{|\mathcal{V}|} \log P_{\theta}(v_i | \mathcal{G}_{<i}, \mathcal{Q}, \mathcal{R}), \quad (20)$$

$$\mathcal{L}_{\text{edge}} = - \sum_{(\mathcal{G}, \mathcal{Q}) \in \mathcal{D}} \sum_{i=1}^{|\mathcal{V}|} \sum_{j=1}^{i-1} \log P_{\theta}(e_{ji} | v_i, \mathcal{G}_{<i}, \mathcal{Q}). \quad (21)$$

Throughout training, we employ a teacher forcing strategy, feeding the model ground-truth structures at each step to stabilize and accelerate learning. We train ARG-DESIGNER following a two-phase process, see Fig. 2b. We begin with a **cold start** on \mathcal{D}_{exp} , followed by **efficiency fine-tuning** on \mathcal{D}_{eff} with a lower learning rate.

Inference. During inference, given a new task query \mathcal{Q} , the trained ARG-DESIGNER model generates a collaboration graph autoregressively without any ground-truth guidance. The process begins by initializing an empty graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. It then enters a generation loop that iteratively builds the graph node by node. At each step i , the model’s node generator first samples a role r_i for the new agent v_i from the probability distribution $P_{\theta}(v_i | \mathcal{G}_{<i}, \mathcal{Q})$. If the sampled role is the special END token, or if a predefined maximum number of agents N_{max} is reached, the generation process terminates. Otherwise, the new node v_i is added to the vertex set \mathcal{V} . Subsequently, the model’s edge generator is invoked. It sequentially considers each existing node $v_j \in \mathcal{G}_{<i}$ and samples the existence of an incoming edge e_{ji} from v_j to v_i . This entire process of node and edge generation is repeated until a termination condition is met, at which point the final graph \mathcal{G} is returned.

Experiments

Experimental Setting

Datasets and Metrics. Following (Zhang et al. 2025b), we evaluated ARG-DESIGNER on three categories of datasets:

Method	MMLU	GSM8K	AQuA	MultiArith	SVAMP	HumanEval	Average
Vanilla	80.39	82.30	71.06	93.09	86.55	71.39	80.80
CoT	81.69 \uparrow 1.30	86.50 \uparrow 4.20	73.58 \uparrow 2.52	93.25 \uparrow 0.16	87.36 \uparrow 0.81	74.67 \uparrow 3.28	82.84 \uparrow 2.04
SC (CoT)	83.66 \uparrow 3.27	81.60 \downarrow 0.70	75.63 \uparrow 4.57	94.12 \uparrow 1.03	88.59 \uparrow 2.04	79.83 \uparrow 8.44	83.91 \uparrow 3.11
Chain	83.01 \uparrow 2.62	88.30 \uparrow 6.00	74.05 \uparrow 2.99	93.27 \uparrow 0.18	87.17 \uparrow 0.62	81.37 \uparrow 9.98	84.53 \uparrow 3.73
Tree	81.04 \uparrow 0.65	85.20 \uparrow 2.90	71.23 \uparrow 0.17	93.68 \uparrow 0.59	88.91 \uparrow 2.36	80.53 \uparrow 9.14	83.43 \uparrow 2.63
Complete	82.35 \uparrow 1.96	80.10 \downarrow 2.20	72.95 \uparrow 1.89	94.53 \uparrow 1.44	84.01 \downarrow 2.54	79.03 \uparrow 7.64	82.16 \uparrow 1.36
Random	84.31 \uparrow 3.92	86.90 \uparrow 4.60	76.48 \uparrow 5.42	94.08 \uparrow 0.99	87.54 \uparrow 0.99	82.66 \uparrow 11.27	85.33 \uparrow 4.53
LLM-Debate	84.96 \uparrow 4.57	91.40 \uparrow 9.10	77.65 \uparrow 6.59	96.36 \uparrow 3.27	90.11 \uparrow 3.56	84.70 \uparrow 13.31	87.53 \uparrow 6.73
AgentPrune	85.07 \uparrow 4.57	91.10 \uparrow 8.80	80.51 \uparrow 9.45	94.65 \uparrow 1.56	90.58 \uparrow 4.03	86.75 \uparrow 15.36	88.09 \uparrow 7.29
AgentDropout	85.62 \uparrow 5.23	91.70 \uparrow 9.40	80.94 \uparrow 9.88	95.60 \uparrow 2.51	91.04 \uparrow 4.49	85.98 \uparrow 14.59	88.48 \uparrow 7.68
G-designer	86.92 \uparrow 6.53	93.80 \uparrow 11.50	81.60 \uparrow 10.54	96.50 \uparrow 3.41	93.10 \uparrow 6.55	88.33 \uparrow 16.94	90.04 \uparrow 9.24
ARG-DESIGNER	89.54 \uparrow 9.15	94.40 \uparrow 12.10	86.45 \uparrow 15.39	98.93 \uparrow 5.84	95.63 \uparrow 9.08	91.74 \uparrow 20.35	92.78 \uparrow 11.98

Table 2: Performance comparison (%) on six benchmarks. The best results are highlighted in bold.

❶ **General Reasoning:** MMLU (Hendrycks et al. 2021);
❷ **Mathematical Reasoning:** GSM8K (Cobbe et al. 2021), MultiArith (Roy and Roth 2016), SVAMP (Patel, Bhatamishra, and Goyal 2021), and AQuA (Ling et al. 2017);
❸ **Code Generation:** HumanEval (Chen et al. 2021).

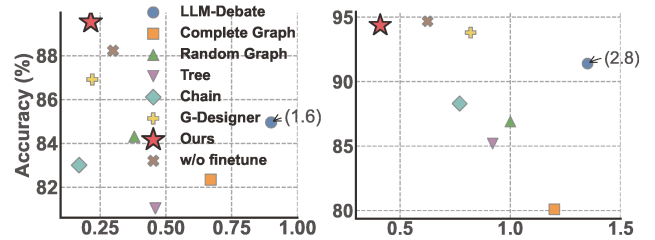
Baselines. We compare ARG-DESIGNER against various baselines, which can be grouped into four main categories:

❶ **Single-agent methods**, including CoT (Wei et al. 2022) and Self-Consistency (Wang et al. 2023);
❷ **MAS with fixed topologies**, such as Chain, Tree, Complete Graph, and Random Graph (Qian et al. 2025);
❸ **MAS with Debate** like LLM-Debate (Du et al. 2023), where multiple agents iteratively critique and refine responses in a structured process;
❹ **MAS with Learnable topologies**, which include Agent-Prune (Zhang et al. 2025a), AgentDropout (Wang et al. 2025), and G-Designer (Zhang et al. 2025b).

Implementation Details. We access GPT models via the OpenAI API, primarily using `gpt-4o-2024-08-06` (GPT-4o). We employ a summarizer agent to aggregate the history of dialogue and produce the final solution $a^{(K)}$, with $K = 3$ for all baselines across all experiments. The node encoder is implemented using `all-MiniLM-L6-v2` (Wang et al. 2020), with the embedding dimension set to $D = 384$. The hyperparameter α is set to 0.2 for all experiments. Following classical configurations in LLM-MAS (Zhuge et al. 2024; Yin et al. 2023; Zhang et al. 2025b), we provide explicit agent profiles for multi-agent methods and use GPT-4 to generate these profile pools. For all datasets, we use $B \in \{40, 60\}$ queries for model training.

Experimental Results

Performance Comparison. The comparison results are presented in Table 2, from which we have the following observations. ❶ ARG-DESIGNER achieves the best performance across all six benchmarks, consistently outperforming a wide range of baselines. The superior performance demonstrates the effectiveness of the autoregressive graph generation paradigm in MAS topology design. ❷ Compared to the strongest learning-based baseline, G-Designer, ARG-DESIGNER shows a significant performance gain. For in-



(a) Token cost of MMLU (10^6) (b) Token cost of GSM8K (10^7)

Figure 3: The prompt token cost comparison.

stance, on AQuA, ARG-DESIGNER achieves an accuracy of 86.45%, surpassing G-Designer by a substantial margin of 4.85%. ❸ When compared to debate-based methods like LLM-Debate, ARG-DESIGNER demonstrates a remarkable improvement of 8.8% on AQuA and 2.66% on GSM8K. This highlights the inefficiency of the fixed and all-to-all communication protocol.

Token Efficiency. A key benefit of ARG-DESIGNER is its ability to generate tailored topologies for different tasks, which prevents unnecessary complexity and thus minimizes token consumption. Fig. 3a and 3b illustrate the trade-off between performance and token cost. We can observe that: ❶ ARG-DESIGNER elegantly balances efficiency and performance. On the GSM8K dataset, ARG-DESIGNER is the most token-efficient method, using only 4.1e6 tokens, while achieving a top-tier accuracy of 94.40%. It surpasses the strong G-Designer baseline in accuracy while using approximately 50% fewer tokens. ❷ A general trend where more complex communication structures, such as the dialogue-based LLM-Debate, achieve relatively good performance but at an extremely high token cost. ❸ A direct comparison between ARG-DESIGNER and its **w/o fine-tune** variant further underscores the value of the efficiency fine-tuning phase. For instance, on MMLU, fine-tuning improves accuracy from 88.23% to 89.54% while simultaneously cutting token usage by nearly 30%. On GSM8K, it reduces token




Figure 4: The robustness, extensibility of ARG-DESIGNER and case studies.

Method	MMLU	GSM8K	HumanEval	Average
Vanilla	80.39	82.30	71.39	78.02
ARG-DESIGNER	89.54	94.40	91.74	91.89
w/o fine-tune	88.23	94.70	90.91	91.28
w/o task emb.	86.93	93.10	89.26	89.76
w/o hist. emb.	88.23	93.60	90.08	90.64

Table 3: Results of ablation study.

consumption by a massive 34% (from 6.25e6 to 4.1e6). This demonstrates that our two-phase training strategy is highly effective at optimizing for both performance and efficiency.

Ablation Study. To validate the effectiveness of key components in ARG-DESIGNER, we investigate three major variants of our model: ❶ **w/o fine-tune**, where efficiency fine-tuning phase is removed. ❷ **w/o task emb.**, where the influence of task embedding f_Q is set to be 0. ❸ **w/o hist. emb.**, where the historical embeddings are removed from the generation model. The results in Table 3 demonstrate the contribution of each component: ❶ ARG-DESIGNER achieves the highest average score of 91.89, significantly outperforming a vanilla baseline (78.02) that lacks these sophisticated mechanisms. ❷ The absence of task-specific guidance (**w/o task emb.**) causes the most significant performance degradation, which highlights that conditioning on the task is crucial for generating a bespoke and effective collaboration topology. ❸ The removal of historical embeddings (**w/o hist. emb.**) also leads to a noticeable decline to 90.64, confirming the value of modeling dependencies between agents. ❹ Interestingly, the performance of the **w/o fine-tune** variant is highly competitive. This indicates that the model can learn fundamental collaboration patterns from the first stage. Nevertheless, the fine-tuning bring significant efficiency benefits (see Fig. 3), while also slightly improving performance.

Robustness Analysis. Following Zhuge et al. (2024), we evaluate the robustness of ARG-DESIGNER by simulating a system prompt attack, where an adversarial prompt is injected into a single agent to disrupt its function. As illustrated in Fig. 4a, this attack causes significant performance degradation in MAS with fixed and naive topologies. In contrast, ARG-DESIGNER shows excellent robustness against attacks with the least performance degradation (2.15%). This resilience emerges from our training objective, which discourages brittle structures and guides the model to construct fault-tolerant topologies with distributed risk and re-

dundant communication paths.

Extensibility Analysis. Furthermore, we examine the model’s extensibility. As depicted in Fig. 4b, we introduce several new roles, to the pre-trained model without any re-training. When presented with a legal question from the MMLU regarding a contract embargo, ARG-DESIGNER demonstrates its adaptability. Correctly identifies the high relevance of the newly added “Lawyer” role and dynamically generates a collaboration graph placing the lawyer at its core, coordinating with other experts. This case vividly illustrates that ARG-DESIGNER can seamlessly scale its capabilities by integrating new knowledge, creating effective specialized team structures on the fly.

Case Study. To further illustrate the advantages of ARG-DESIGNER over learning-based baselines (e.g., G-Designer (Zhang et al. 2025c)), we conduct a comparative case study on representative cases in HumanEval and GSM8K. As shown in Fig. 4c. The key difference lies in the flexibility of composition versus the static design. G-Designer requires a predefined and fixed set of agents and a fixed agent count. Its solution graphs remain within this rigid template, regardless of task complexity. ARG-DESIGNER, in contrast, dynamically generates both the roles and their communication links from an extensible role pool. It adapts the number of agents and connections based on task needs. Therefore, ARG-DESIGNER constructs more efficient collaboration graphs with fewer agents and messages, cutting token usage without sacrificing accuracy.

Conclusion

In this work, we addressed redundant composition and limited extensibility in template graph modification approaches by reformulating collaboration topology design as autoregressive graph generation. We introduced ARG-DESIGNER, a novel autoregressive model that constructs collaboration graphs from scratch, conditioned on natural language task queries. Our approach dynamically determines agent numbers, selects roles from an extensible pool, and establishes optimal communication links, creating bespoke MAS topologies tailored to specific task demands. Extensive experiments across six benchmarks demonstrate that ARG-DESIGNER consistently outperforms existing methods, achieving state-of-the-art performance while maintaining superior token efficiency.

Acknowledgments

The work of S. Pan was partially supported by the Australian Research Council (ARC) under Grant Nos. DP240101547 and FT210100097.

References

- Cai, J.; Zhang, Y.; Liu, F.; and Ng, S.-K. 2025. Leveraging Diffusion Model as Pseudo-Anomalous Graph Generator for Graph-Level Anomaly Detection. In *Proceedings of the International Conference on Machine Learning*.
- Cai, J.; Zhang, Y.; Lu, Z.; Guo, W.; and Ng, S.-K. 2024. Towards Effective Federated Graph Anomaly Detection via Self-boosted Knowledge Distillation. In *Proceedings of the ACM International Conference on Multimedia*, 5537–5546.
- Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; de Oliveira Pinto, H. P.; Kaplan, J.; Edwards, H.; Burda, Y.; Joseph, N.; Brockman, G.; et al. 2021. Evaluating Large Language Models Trained on Code. arXiv:2107.03374.
- Chen, Q.; Li, S.; Liu, Y.; Pan, S.; Webb, G. I.; and Zhang, S. 2025. Uncertainty-Aware Graph Neural Networks: A Multi-hop Evidence Fusion Approach. *IEEE Transactions on Neural Networks and Learning Systems*.
- 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. In *International Conference on Machine Learning*.
- Fu, L.; Deng, B.; Huang, S.; Liao, T.; Pan, S.; and Chen, C. 2025a. Less is More: Federated Graph Learning with Alleviating Topology Heterogeneity from A Causal Perspective. In *Forty-second International Conference on Machine Learning*.
- Fu, L.; Deng, B.; Huang, S.; Liao, T.; Zhang, C.; and Chen, C. 2025b. Learn from Global Rather Than Local: Consistent Context-Aware Representation Learning for Multi-View Graph Clustering. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI*, 16–22.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2021. Measuring Massive Multitask Language Understanding. arXiv:2009.03300.
- Hong, S.; Zhuge, M.; Chen, J.; Zheng, X.; Cheng, Y.; Zhang, C.; Wang, J.; Wang, Z.; Yau, S. K. S.; Lin, Z.; Zhou, L.; Ran, C.; Xiao, L.; Wu, C.; and Schmidhuber, J. 2024. MetaGPT: Meta Programming for A Multi-Agent Collaborative Framework. arXiv:2308.00352.
- Li, B.; Luo, Y.; Chai, C.; Li, G.; and Tang, N. 2024a. The Dawn of Natural Language to SQL: Are We Fully Ready? *Proceedings of the VLDB Endowment*, 17(11): 3318–3331.
- Li, S.; Liu, Y.; Chen, Q.; Webb, G. I.; and Pan, S. 2024b. Noise-resilient unsupervised graph representation learning via multi-hop feature quality estimation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 1255–1265.
- Ling, W.; Yogatama, D.; Dyer, C.; and Blunsom, P. 2017. Program Induction by Rationale Generation: Learning to Solve and Explain Algebraic Word Problems. arXiv:1705.04146.
- Liu, Y.; Zhang, G.; Wang, K.; Li, S.; and Pan, S. 2025. Graph-augmented large language model agents: Current progress and future prospects. *IEEE Intelligent Systems*.
- Miao, R.; Liu, Y.; Wang, Y.; Shen, X.; Tan, Y.; Dai, Y.; Pan, S.; and Wang, X. 2025. Blindguard: Safeguarding llm-based multi-agent systems under unknown attacks. *arXiv preprint arXiv:2508.08127*.
- Pan, J.; Liu, Y.; Zheng, X.; Zheng, Y.; Liew, A. W.-C.; Li, F.; and Pan, S. 2025a. A label-free heterophily-guided approach for unsupervised graph fraud detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 12443–12451.
- Pan, J.; Liu, Y.; Zhou, C.; Xiong, F.; Liew, A. W.-C.; and Pan, S. 2026. Correcting False Alarms from Unseen: Adapting Graph Anomaly Detectors at Test Time. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Pan, J.; Zheng, Y.; Tan, Y.; and Liu, Y. 2025b. A Survey of Generalization of Graph Anomaly Detection: From Transfer Learning to Foundation Models. In *The 16th IEEE International Conference on Knowledge Graphs*.
- 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. arXiv:2406.07155.
- Roy, S.; and Roth, D. 2016. Solving General Arithmetic Word Problems. arXiv:1608.01413.
- Shen, X.; Liu, Y.; Dai, Y.; Wang, Y.; Miao, R.; Tan, Y.; Pan, S.; and Wang, X. 2025. Understanding the Information Propagation Effects of Communication Topologies in LLM-based Multi-Agent Systems. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*.
- Song, C. H.; Wu, J.; Washington, C.; Sadler, B. M.; Chao, W.-L.; and Su, Y. 2023. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2998–3009.
- Tan, Y.; Hu, X.; Xue, H.; De Melo, C.; and Salim, F. D. 2025. Biscle: Binding and Separation in Continual Learning for Video Language Understanding. In *Advances in Neural Information Processing Systems*.
- Wang, W.; Wei, F.; Dong, L.; Bao, H.; Yang, N.; and Zhou, M. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in neural information processing systems*, 33: 5776–5788.
- 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.

- Wang, Z.; Wang, Y.; Liu, X.; Ding, L.; Zhang, M.; Liu, J.; and Zhang, M. 2025. AgentDropout: Dynamic Agent Elimination for Token-Efficient and High-Performance LLM-Based Multi-Agent Collaboration. arXiv:2503.18891.
- Wang, Z.; Zhang, H.; Li, C.-L.; Eisenschlos, J. M.; Perot, V.; Wang, Z.; Miculicich, L.; Fujii, Y.; Shang, J.; Lee, C.-Y.; and Pfister, T. 2024. Chain-of-Table: Evolving Tables in the Reasoning Chain for Table Understanding. arXiv:2401.04398.
- 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.
- Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T.; Cao, Y.; and Narasimhan, K. 2023. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36: 11809–11822.
- Yin, Z.; Sun, Q.; Chang, C.; Guo, Q.; Dai, J.; Huang, X.; and Qiu, X. 2023. Exchange-of-Thought: Enhancing Large Language Model Capabilities through Cross-Model Communication. arXiv:2312.01823.
- Zhang, G.; Yue, Y.; Li, Z.; Yun, S.; Wan, G.; Wang, K.; Cheng, D.; Yu, J. X.; and Chen, T. 2025a. Cut the crap: An economical communication pipeline for llm-based multi-agent systems. In *International Conference on Learning Representations*.
- Zhang, G.; Yue, Y.; Sun, X.; Wan, G.; Yu, M.; Fang, J.; Wang, K.; Chen, T.; and Cheng, D. 2025b. G-designer: Architecting multi-agent communication topologies via graph neural networks. In *International Conference on Machine Learning*.
- Zhang, H.; Wu, B.; Yuan, X.; Pan, S.; Tong, H.; and Pei, J. 2024. Trustworthy graph neural networks: Aspects, methods, and trends. *Proceedings of the IEEE*, 112(2): 97–139.
- Zhang, J.; Xiang, J.; Yu, Z.; Teng, F.; Chen, X.; Chen, J.; Zhuge, M.; Cheng, X.; Hong, S.; Wang, J.; Zheng, B.; Liu, B.; Luo, Y.; and Wu, C. 2025c. AFlow: Automating Agentic Workflow Generation. arXiv:2410.10762.
- Zhao, Y.; Liu, Y.; Li, S.; Chen, Q.; Zheng, Y.; and Pan, S. 2025. Freegad: A training-free yet effective approach for graph anomaly detection. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*.
- Zhong, L.; Wang, Z.; and Shang, J. 2024. Debug like a Human: A Large Language Model Debugger via Verifying Runtime Execution Step by Step. In *Findings of the Association for Computational Linguistics ACL 2024*, 851–870.
- Zhou, H.; Wan, X.; Sun, R.; Palangi, H.; Iqbal, S.; Vulić, I.; Korhonen, A.; and Arik, S. Ö. 2025. Multi-Agent Design: Optimizing Agents with Better Prompts and Topologies. arXiv:2502.02533.
- Zhu, Y.; Du, S.; Li, B.; Luo, Y.; and Tang, N. 2024. Are Large Language Models Good Statisticians? arXiv:2406.07815.
- Zhuang, S.; Wu, Z.; Chen, Z.; Dai, H.-N.; and Liu, X. 2025. Refine then Classify: Robust Graph Neural Networks with Reliable Neighborhood Contrastive Refinement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 13473–13482.
- 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*.