

HyperAgent: Leveraging Hypergraphs for Topology Optimization in Multi-Agent Communication

Heng Zhang
South China Normal University
Guangzhou, China
2024025450@m.scnu.edu.cn

Yuling Shi
Shanghai Jiao Tong University
Shanghai, China
yuling.shi@sjtu.edu.cn

Xiaodong Gu
Shanghai Jiao Tong University
Shanghai, China
xiaodong.gu@sjtu.edu.cn

Zijian Zhang
University of Pennsylvania
Philadelphia, USA
zzjharry@alumni.upenn.edu

Haochen You
Columbia University
New York, USA
hy2854@columbia.edu

Lubin Gan
University of Science and Technology
of China
Hefei, China
ganlubin@mail.ustc.edu.cn

Yilei Yuan
University of Michigan
Ann Arbor, USA
yiliey@umich.edu

Jin Huang*
South China Normal University
Guangzhou, China
huangjin@m.scnu.edu.cn

ABSTRACT

Recent advances in large language model-powered multi-agent systems have demonstrated remarkable collective intelligence through effective communication. However, existing approaches face two primary challenges: (i) *Ineffective group collaboration modeling*, as they rely on pairwise edge representations in graph structures, limiting their ability to capture relationships among multiple agents; and (ii) *Limited task-adaptiveness in communication topology design*, leading to excessive communication cost for simple tasks and insufficient coordination for complex scenarios. These issues restrict the scalability and practical deployment of adaptive collaboration frameworks. To address these challenges, we propose **HyperAgent**, a hypergraph-based framework that optimizes communication topologies and effectively captures group collaboration patterns using direct hyperedge representations. Unlike edge-based approaches, HyperAgent uses hyperedges to link multiple agents within the same subtask and employs hypergraph convolutional layers to achieve one-step information aggregation in collaboration groups. Additionally, it incorporates a variational autoencoder framework with sparsity regularization to dynamically adjust hypergraph topologies based on task complexity. Experiments highlight the superiority of HyperAgent in both performance and efficiency. For instance, on GSM8K, HyperAgent achieves 95.07% accuracy while reducing token consumption by 25.33%, demonstrating the potential of hypergraph-based optimization for multi-agent communication.

*Corresponding author.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/QTVF9552>

KEYWORDS

Large Language Model, Multi-agent Systems, Multi-agent Communication, Graph Neural Networks

ACM Reference Format:

Heng Zhang, Yuling Shi, Xiaodong Gu, Zijian Zhang, Haochen You, Lubin Gan, Yilei Yuan, and Jin Huang. 2026. HyperAgent: Leveraging Hypergraphs for Topology Optimization in Multi-Agent Communication. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 10 pages. <https://doi.org/10.65109/QTVF9552>

1 INTRODUCTION

Large language model-powered agents have demonstrated exceptional capabilities across diverse tasks[6, 7, 52]. When multiple agents collaborate, their collective intelligence can surpass the performance of individual agents[12, 16, 33]. The effectiveness of this collaboration depends on the design of communication topologies[44, 72], which govern how information is exchanged and actions are coordinated among agents. Well-designed topologies allow information to flow efficiently, enabling agents to integrate their efforts and handle complex tasks such as problem solving and decision making[53, 68]. Conversely, ineffective topologies increase communication inefficiency and hinder coordination[21, 43]. For example, in software development, architects, programmers, and testers must interact through structured communication to ensure smooth workflows. As a result, optimizing communication topologies to suit specific tasks has become a core focus in multi-agent system research[67].

Current research has explored various approaches to multi-agent topology design, resulting in several well-established paradigms. Static topology methods rely on predefined structures, such as chain, star, tree, and complete graph configurations, to ensure consistent task execution[44, 60, 63]. Template-based approaches modify existing graph structures by adding or removing edges to accommodate different scenarios[12, 72]. Optimization-driven methods use search

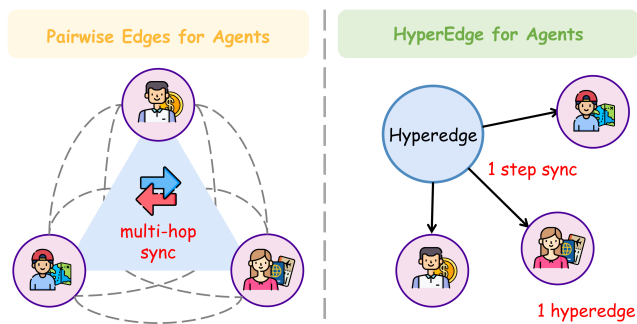


Figure 1: Comparison of communication topologies for multi-agent collaboration. (a) Pairwise edges require multiple connections and multi-hop information propagation among agents. (b) Hyperedges enable direct one-step synchronization by connecting all collaborating agents within a single structure.

algorithms to explore topology spaces and identify more effective configurations through iterative refinement[4]. Recent developments have introduced dynamic approaches that adjust communication structures during task execution, aiming to balance performance and efficiency[14, 37, 70]. These methods share a common foundation: agents are represented as graph nodes, and communication relationships are modeled as pairwise edges[5, 42]. Using graph neural networks, these approaches learn connection patterns directly from task data[44], enabling significant performance improvements in benchmarks such as code generation[10, 23], mathematical reasoning[54, 60], and question answering[1, 68, 73].

Despite their architectural variations, recent topology design methods share a common foundation that limits scalability and efficiency. Approaches like G-Designer[32], GPTSwarm[72], and DyLAN[38] represent multi-agent systems as graphs, where agents are nodes and communication channels are modeled as edges. Our analysis of these methods on tasks with varying complexity in the MMLU benchmark reveals a consistent trend: for simple tasks, such as basic multiple-choice questions, they naturally converge to sparse topologies, reducing unnecessary communication. However, as tasks grow more complex, these methods require denser topologies to maintain performance. For advanced reasoning tasks, effective configurations approach near-complete graphs to ensure sufficient information exchange, but this densification drives communication costs to scale quadratically with agent count, as every agent establishes edges with all others. More critically, graph-based representations can only model pairwise relationships, making it difficult to capture multi-agent collaborations as unified units. For example, three agents collaborating on the same subtask must be connected by multiple pairwise edges. This limitation forces topology design into a tradeoff: sparse graphs reduce costs but fragment coordination, while dense graphs enable better communication but introduce high overhead. Existing methods tackle this balance differently, but the issue persists as a consequence of graph-based modeling, not algorithmic choices.

The core problem lies in how graph representations decompose collaborative units. Consider a scenario in multi-agent problem-solving where a mathematician derives a formula, a programmer

implements it, and a validator checks its correctness. Instead of treating this group as a whole, graph models decompose these interactions into three pairwise edges, creating inefficiencies. When the mathematician updates the formula, the message must follow a sequential path. The programmer receives the update first, but the validator must wait for the programmer to process and relay it. This multi-hop propagation increases latency and risks information degradation over intermediate steps. Additionally, this decomposition treats group interactions as emerging from separate edges, rather than unified structures. Algorithms then must infer these units, searching through an optimization space that grows quadratically with agent count. Hypergraphs overcome this limitation by treating collaborative units as first-class structures. A hyperedge directly connects all agents in a group, enabling unified information aggregation. All agents contribute to and receive from a shared representation in a single step, eliminating multi-hop delays and preserving semantic consistency. Moreover, hypergraphs reduce the optimization space from quadratic pairwise edges to linear collaborative units, allowing algorithms to directly work with meaningful structures rather than inferring them indirectly.

Based on the above observations, we propose HyperAgent, a hypergraph-based framework for optimizing communication topologies in multi-agent systems. HyperAgent models the system as a hypergraph, where each hyperedge connects agents working on the same subtask, directly capturing group-level interactions. To process this structure, we introduce hypergraph convolutional layers that enable efficient information aggregation within collaboration units. For dynamic topology generation, HyperAgent uses a variational autoencoder. The encoder transforms agent features and task information into latent representations, and the decoder constructs task-specific hypergraphs. Sparsity regularization controls the number of collaboration units, maintaining communication efficiency. This design allows HyperAgent to adapt topologies to task demands, creating sparse structures for simple tasks and denser ones for complex coordination. By balancing adaptability and efficiency, HyperAgent delivers a robust solution for multi-agent communication. Our contributions can be concluded as:

- We observe that pairwise graph representations fail to capture multi-agent collaboration effectively. Existing methods force a tradeoff between sparse topologies, which fragment coordination, and dense topologies, which increase communication overhead.
- We propose **HyperAgent**, a hypergraph-based framework that uses hyperedges to represent group collaboration units. This representation enables direct synchronization within teams through unified node-edge-node transformations.
- We develop a variational hypergraph autoencoder for task-adaptive topology generation. It encodes agent features and task semantics into latent representations and generates sparse hypergraphs, adjusting complexity according to task difficulty.
- Extensive experiments demonstrate HyperAgent outperforms state-of-the-art methods, achieving **88.50%** accuracy on MMLU and **92.90%** pass@1 on HumanEval, with up to **25.33%** reduction in communication token consumption.

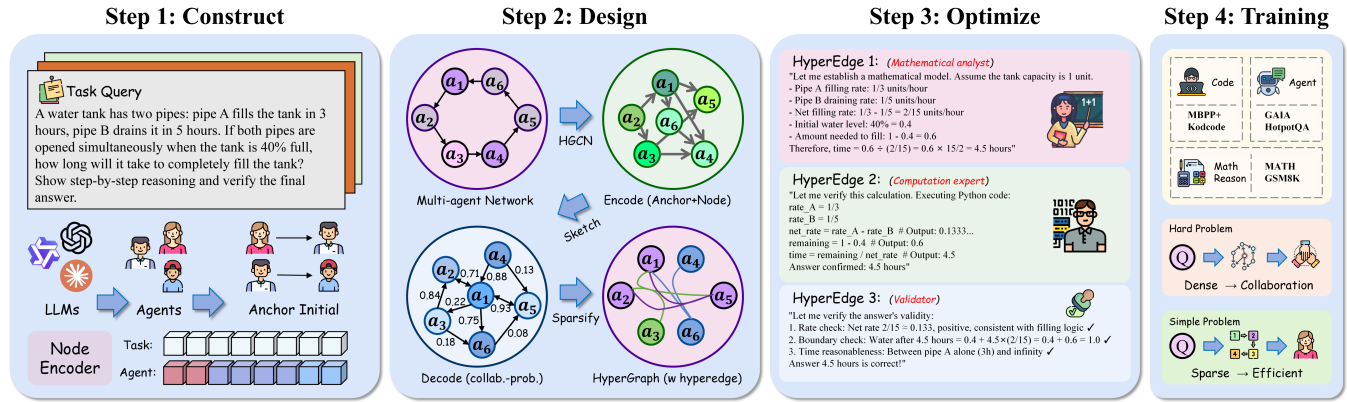


Figure 2: The HyperAgent pipeline. We encode agents and tasks into a hypergraph, then apply a variational autoencoder with sparsity regularization to generate task-adaptive communication topologies. Agents interact through hyperedge-based collaboration for multiple rounds. The VAE is trained via policy gradients to maximize task performance while minimizing communication overhead.

2 RELATED WORK

2.1 LLMs-agent Collaboration

Large language models have driven the shift from single-agent to collaborative multi-agent systems[6, 52]. Early works demonstrated that multiple LLM-based agents working together can outperform individual agents through role specialization[33] and coordinated interactions[12]. Collaboration takes various architectural forms. Sequential architectures organize agents in chains where each refines the output of its predecessor[21, 43]. Hierarchical systems employ star topologies with a central coordinator directing subordinates[57, 62]. Debate-based approaches let agents iteratively argue to enhance reasoning and factuality[9, 16, 26]. Recent advances focus less on static structures and more on collaboration mechanisms, distinguishing cooperative, competitive, and hybrid frameworks[29, 42, 49]. Studies confirm that effective collaboration arises from well-designed communication protocols rather than mere aggregation of agents[22, 61], and multi-agent systems further exhibit human-like collective behaviors such as consensus formation and adaptive negotiation[2, 56]. In software engineering domains, hierarchical debugging frameworks decompose complex code for systematic error resolution[47], competitive debates enable diverse reasoning along fault propagation traces[30], and experience-driven approaches accumulate repair knowledge from historical trajectories[11]. These efforts are complemented by research on LLM-generated code patterns[48], long-context compression[17, 46, 65], reinforcement learning-based reasoning[36], and cross-language translation[55].

2.2 Graphs for Multi-agents

Graphs naturally model relationships and communication structures in multi-agent systems, a perspective rooted in multi-agent reinforcement learning[34, 39]. The advent of LLMs extended this approach to language-driven systems, where early frameworks implicitly used graph structures without defining explicit topologies[8, 27]. Modern research explicitly represents multi-agent organizations

as directed graphs, with nodes as agents and edges as communication links[37, 44, 72]. Predefined topologies like complete graphs allow unrestricted communication but incur high costs, while chain and tree structures support hierarchical flows useful for sequential reasoning[5, 44, 60, 63, 71]. To move beyond fixed designs, graph neural networks now enable dynamic topology learning that generates task-adaptive structures[31, 66]. Pruning techniques further simplify dense graphs by removing redundant edges[28, 59]. Practical applications demonstrate task-aware coordination: repository-level code understanding leverages dependency graphs[41], issue resolution employs fault propagation graphs for collaborative diagnosis[30], and experience banks enable knowledge reuse across problem instances[11]. However, graph-based methods share a common limitation: they represent only pairwise relationships, failing to model collaborative units where multiple agents work jointly on shared tasks[19, 25]. This gap highlights the need for richer frameworks to capture group-level interactions.

3 PRELIMINARY

3.1 Problem Formulation

We model the multi-agent system as a hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E}, \mathbf{W})$. The node set $\mathcal{V} = \{v_1, \dots, v_N\}$ represents N agents in the system. Each node $v_i \in \mathcal{V}$ corresponds to an agent, formalized as:

$$v_i = \{\text{Base}_i, \text{Role}_i, \text{State}_i, \text{Plugin}_i\} \quad (1)$$

Each agent v_i is composed of four key elements: (1) Base_i , the language model instance powering v_i ; (2) Role_i , the agent’s pre-assigned role or function; (3) State_i , representing the agent’s accumulated knowledge and interaction history; and (4) Plugin_i , a set of external tools available to v_i , such as web searchers, code compilers, or file readers. The hyperedge set \mathcal{E} contains collaboration units. Each hyperedge $e \in \mathcal{E}$ connects multiple agents participating in the same subtask. The weight matrix $\mathbf{W} \in \mathbb{R}^{|\mathcal{E}| \times |\mathcal{E}|}$ assigns importance scores to different collaboration units. We set $\mathbf{W} = \mathbb{I}_{|\mathcal{E}| \times |\mathcal{E}|}$ by default, treating all collaboration units equally.

The hypergraph structure can be represented by an incidence matrix $\mathbf{H} \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{E}|}$. Each element is defined as:

$$h_{v,e} = \begin{cases} 1, & \text{if } v \in e \\ 0, & \text{if } v \notin e \end{cases} \quad (2)$$

This binary encoding directly captures group-wise collaboration patterns. The degree of node v_i is $d_i = \sum_{e \in \mathcal{E}} \mathbf{W}_{ee} h_{i,e}$. The degree of hyperedge e is $\delta_e = \sum_{v \in \mathcal{V}} h_{v,e}$. These statistics form diagonal matrices $\mathbf{D} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ and $\mathbf{B} \in \mathbb{R}^{|\mathcal{E}| \times |\mathcal{E}|}$ respectively.

Each LLM-based agent v_i receives prompt \mathcal{P} and generates response \mathcal{R}_i :

$$\mathcal{R}_i = v_i(\mathcal{P}) = v_i(\mathcal{P}_{\text{sys}}, \mathcal{P}_{\text{usr}}) \quad (3)$$

where $\mathcal{P}_{\text{sys}} = \{\text{Role}_i, \text{State}_i\}$ represents the system prompt, and \mathcal{P}_{usr} denotes the user prompt including the given task, responses from other agents, and externally retrieved knowledge.

3.2 Hypergraph Convolution

Given node feature matrix $\mathbf{X}^{(l)} \in \mathbb{R}^{|\mathcal{V}| \times D}$ at layer l , the hypergraph convolutional layer computes updated features through:

$$\mathbf{X}^{(l+1)} = \sigma(\mathbf{D}^{-1/2} \mathbf{H} \mathbf{W} \mathbf{B}^{-1} \mathbf{H}^T \mathbf{D}^{-1/2} \mathbf{X}^{(l)} \Theta^{(l)}) \quad (4)$$

The term $\mathbf{H}^T \mathbf{X}^{(l)}$ performs node-to-edge aggregation. Each row represents one hyperedge’s aggregated feature computed as the sum of features from participating agents. The term \mathbf{B}^{-1} normalizes by dividing each hyperedge feature by the number of participating agents. The term $\mathbf{H}(\mathbf{B}^{-1} \mathbf{H}^T \mathbf{X}^{(l)})$ performs edge-to-node propagation. Each node receives messages from all hyperedges it participates in through summation. The degree normalization $\mathbf{D}^{-1/2}$ ensures balanced information flow. The learnable weight $\Theta^{(l)} \in \mathbb{R}^{D \times D'}$ transforms features. The activation function $\sigma(\cdot)$ introduces nonlinearity. We use ReLU by default. This node-edge-node transformation naturally captures group-wise collaboration patterns.

4 METHODOLOGY

4.1 Multi-Agent Hypergraph Construction

Given input query Q and agent set \mathcal{V} , HyperAgent aims to design a task-adaptive communication topology \mathcal{H}_{com} . We begin by assigning each agent a unique role and profile. Previous research has shown that assigning distinct personas to LLM-based agents enhances cognitive synergy. Based on these roles, different external tools are allocated to agents. For example, Mathematica for a math analyst, Python compiler for a programmer. Thus we initialize each agent v_i as $\{\text{Base}_i, \text{Role}_i, \text{State}_i, \text{Plugin}_i\}$.

We construct a structured multi-agent hypergraph as input to HyperAgent, represented as $\mathcal{H} = (\mathbf{X}_{\text{agent}}, \mathbf{H}_{\text{anchor}})$. Here $\mathbf{X}_{\text{agent}} \in \mathbb{R}^{N \times D}$ is the node feature matrix. $\mathbf{H}_{\text{anchor}} \in \{0, 1\}^{N \times |\mathcal{E}_{\text{anchor}}|}$ is the initial incidence matrix. For the feature matrix, we employ a node encoder to transform each agent’s profile into a fixed-length embedding:

$$\mathbf{x}_i \leftarrow \text{NodeEncoder}(\mathcal{T}(\text{Base}_i), \text{Role}_i, \mathcal{T}(\text{Plugin}_i)) \quad (5)$$

The function $\mathcal{T}(\cdot)$ extracts textual descriptions of the agent’s LLM backbone and assigned plugins. NodeEncoder can be realized using small text embedding models such as Sentence-BERT. After encoding individual agents, we ensure the hypergraph incorporates

task-related information. This query-dependent approach enables HyperAgent to be task-aware and adaptive.

We introduce a task-specific virtual global node v_{task} . This node connects bidirectionally to all agent nodes, enabling a global storage mechanism and facilitating smoother information flow among agents. The task node is encoded as:

$$\mathbf{x}_{\text{task}} \leftarrow \text{NodeEncoder}(Q) \quad (6)$$

After obtaining agent node features $\mathbf{X}_{\text{agent}} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]^T$ and task embedding \mathbf{x}_{task} , we provide a simple anchor hypergraph structure $\mathbf{H}_{\text{anchor}}$. This serves as a starting point for topology design. For instance, given a code generation task with three agents (manager, programmer, code reviewer), the anchor could configure a sequential pipeline where each hyperedge connects adjacent agents. The anchor topology can be user-defined or automatically generated by LLMs. It is often simple and sub-optimal but provides foundational reference and prior knowledge.

We incorporate the task-specific vertex v_{task} and obtain $\tilde{\mathbf{H}}_{\text{anchor}} \in \{0, 1\}^{(N+1) \times (|\mathcal{E}_{\text{anchor}}| + N)}$. The additional N hyperedges represent bidirectional connections between v_{task} and each agent. We establish a task-specific multi-agent hypergraph:

$$\tilde{\mathcal{H}} = \left(\begin{bmatrix} \mathbf{X}_{\text{agent}} \\ \mathbf{x}_{\text{task}}^T \end{bmatrix}, \tilde{\mathbf{H}}_{\text{anchor}} \right) = (\tilde{\mathcal{V}}, \tilde{\mathcal{E}}) \quad (7)$$

where $\tilde{\mathcal{V}} = \mathcal{V} \cup \{v_{\text{task}}\}$ and $\begin{bmatrix} \mathbf{X}_{\text{agent}} \\ \mathbf{x}_{\text{task}}^T \end{bmatrix}$ can be denoted as $\tilde{\mathbf{X}}$.

4.2 Designing Communication Hypergraph Topology

Building upon the task-specific hypergraph $\tilde{\mathcal{H}}$, HyperAgent seeks to establish a fine-grained communication topology \mathcal{H}_{com} . Drawing inspiration from the variational graph auto-encoder (VGAE) framework, HyperAgent employs a VGAE-based encoder-decoder f_v to generate the hypergraph topology:

$$\mathcal{H}_{\text{com}} = f_v(\tilde{\mathcal{H}}; \Theta_v) = p(\mathcal{H}_{\text{com}} | \mathbf{H}_{\text{latent}}) q(\mathbf{H}_{\text{latent}} | \tilde{\mathbf{X}}, \tilde{\mathbf{H}}_{\text{anchor}}) \quad (8)$$

Here f_v is the encoder-decoder architecture with parameters Θ_v . The encoder $q(\cdot)$ maps node embeddings to low-dimensional latent representations. The decoder $p(\cdot)$ reconstructs hypergraph structure from these representations.

The encoder consists of two hypergraph convolutional layers followed by sampling operations. Given node features $\tilde{\mathbf{X}}$ and anchor structure $\tilde{\mathbf{H}}_{\text{anchor}}$, the encoder computes mean vectors $\boldsymbol{\mu}$ and variance vectors $\boldsymbol{\sigma}$ through separate paths:

$$\boldsymbol{\mu} = \text{HGNCN}_{\boldsymbol{\mu}}(\tilde{\mathbf{X}}, \tilde{\mathbf{H}}_{\text{anchor}}), \quad \log \boldsymbol{\sigma} = \text{HGNCN}_{\boldsymbol{\sigma}}(\tilde{\mathbf{X}}, \tilde{\mathbf{H}}_{\text{anchor}}) \quad (9)$$

Both $\text{HGNCN}_{\boldsymbol{\mu}}$ and $\text{HGNCN}_{\boldsymbol{\sigma}}$ are two-layer hypergraph convolutional networks with distinct parameters $\Theta_{\boldsymbol{\mu}}$ and $\Theta_{\boldsymbol{\sigma}}$. The encoder outputs latent representation matrix $\mathbf{H}_{\text{latent}} \in \mathbb{R}^{N \times D}$ by sampling:

$$q(\mathbf{H}_{\text{latent}} | \tilde{\mathbf{X}}, \tilde{\mathbf{H}}_{\text{anchor}}) = \prod_{i=1}^N q(\mathbf{h}_i | \tilde{\mathbf{X}}, \tilde{\mathbf{H}}_{\text{anchor}}) \quad (10)$$

$$q(\mathbf{h}_i | \tilde{\mathbf{X}}, \tilde{\mathbf{H}}_{\text{anchor}}) = \mathcal{N}(\mathbf{h}_i | \boldsymbol{\mu}_i, \text{diag}(\boldsymbol{\sigma}_i^2))$$

Here \mathbf{h}_i , $\boldsymbol{\mu}_i$, and $\boldsymbol{\sigma}_i$ denote the i -th row of corresponding matrices. The encoder parameters are $\Theta_e = \{\Theta_{\boldsymbol{\mu}}, \Theta_{\boldsymbol{\sigma}}\}$. This stochastic

encoding enables diverse topology generation while maintaining meaningful structure.

The decoder transforms latent representations into hypergraph structure through a two-phase process. The decoder $p(\cdot) = p_c \circ p_s$ first constructs a sketched pairwise affinity matrix S , then refines it into the final hypergraph topology:

$$p(\mathcal{H}_{\text{com}}|\mathbf{H}_{\text{latent}}) = \int_S p_c(\mathcal{H}_{\text{com}}|S)p_s(S|\mathbf{H}_{\text{latent}})dS \quad (11)$$

At the first step, $p_s(\cdot)$ constructs sketched adjacency matrix $S \in [0, 1]^{N \times N}$ from latent representations:

$$p_s(S|\mathbf{H}_{\text{latent}}) = \prod_{i=1}^N \prod_{j=1}^N p_s(S_{ij}|\mathbf{h}_i, \mathbf{h}_j, \mathbf{h}_{\text{task}}; \Theta_d) \quad (12)$$

The detailed derivation is:

$$\begin{aligned} p_s(S_{ij} = 1|\mathbf{h}_i, \mathbf{h}_j, \mathbf{h}_{\text{task}}) &= g(\mathbf{h}_i, \mathbf{h}_j, \mathbf{h}_{\text{task}}) \\ &= \text{Sigmoid}((\log(\epsilon) - \log(1 - \epsilon) + \omega_{ij})/\tau) \end{aligned} \quad (13)$$

where $\omega_{ij} = \text{FFN}_d([\mathbf{h}_i, \mathbf{h}_j, \mathbf{h}_{\text{task}}])$ with parameters Θ_d . The uniform random variable $\epsilon \sim \text{Uniform}(0, 1)$ introduces stochasticity. The temperature τ controls the sharpness of the sigmoid function. When τ approaches zero, the output becomes increasingly discrete.

The sketched matrix S typically contains many nonzero entries, resulting in a dense structure. The second decoder phase $p_c(\cdot)$ refines S into a sparse hypergraph topology through structured regularization:

$$\begin{aligned} \tilde{S} &= \arg \min_{S' \in \mathbb{S}} \frac{1}{2} \|S - \mathbf{Z}\mathbf{W}\mathbf{Z}^T\|_F^2 + \zeta \|\mathbf{W}\|_* \\ &\quad + \frac{1}{2} \|\mathbf{A}_{\text{anchor}} - \mathbf{Z}\mathbf{W}\mathbf{Z}^T\|_F^2 \quad (14) \\ \text{subject to } \tilde{S} &= \mathbf{Z}\mathbf{W}\mathbf{Z}^T \end{aligned}$$

The matrix $\mathbf{Z} \in \mathbb{R}^{N \times r}$ contains the top r left singular vectors of S . The weight matrix $\mathbf{W} \in \mathbb{R}^{r \times r}$ is optimized to balance three objectives. The first term keeps refined structure \tilde{S} close to original sketch S . The third term maintains similarity to anchor topology $\mathbf{A}_{\text{anchor}}$ (derived by converting $\mathbf{H}_{\text{anchor}}$ to pairwise connections). The second term applies nuclear norm regularization $\|\mathbf{W}\|_* = \sum_i \lambda_i$ where λ_i are singular values of \mathbf{W} . This encourages low-rank structure in \mathbf{W} , translating to sparsity in \tilde{S} since $\|\tilde{S}\|_* = \|\mathbf{W}\|_*$ holds due to $\mathbf{Z}^T \mathbf{Z} = \mathbb{I}_{r \times r}$. The hyperparameter ζ controls sparsification strength.

The refined adjacency matrix \tilde{S} defines pairwise collaboration affinities. We convert this into hyperedge structure by grouping strongly connected agents. Specifically, for each agent i , we identify the k agents with highest values in row $\tilde{S}[i, :]$ and form a hyperedge connecting these $k+1$ agents. This grouping produces the incidence matrix \mathbf{H}_{com} defining the final topology:

$$\mathcal{H}_{\text{com}} = (\mathcal{V}, \mathcal{E}_{\text{com}}), \quad (15)$$

$$\mathcal{E}_{\text{com}} = \{e_i | e_i \text{ formed by top-}k \text{ connections of agent } i\} \quad (16)$$

The resulting structure captures group-wise collaboration patterns while maintaining sparsity for communication efficiency.

4.3 Multi-Round Agent Interaction

The generated hypergraph topology \mathcal{H}_{com} guides information flow during collaboration. At each round t , agents execute according to a topological ordering. An agent v_i can only execute after all agents in its in-neighborhood have produced responses. When agent v_i executes, it receives system prompt $\mathcal{P}_{\text{sys}}^{(t)} = \{\text{Role}_i, \text{State}_i\}$ and user prompt $\mathcal{P}_{\text{usr}}^{(t)} = \{Q, \cup_{v_j \in \mathcal{N}_{\text{in}}(v_i)} \mathcal{R}_j^{(t)}\}$. The in-neighborhood $\mathcal{N}_{\text{in}}(v_i)$ includes agents sharing hyperedges with v_i . The agent generates:

$$\mathcal{R}_i^{(t)} = v_i(\mathcal{P}_{\text{sys}}^{(t)}, \mathcal{P}_{\text{usr}}^{(t)}) \quad (17)$$

After K rounds, an aggregation function produces the final answer:

$$a^{(K)} \leftarrow \text{Aggregate}(\mathcal{R}_1^{(K)}, \mathcal{R}_2^{(K)}, \dots, \mathcal{R}_N^{(K)}) \quad (18)$$

The aggregation can be majority voting, weighted combination, or delegation to a specific agent depending on the task type.

4.4 Training Objective

HyperAgent optimizes topology generation through policy gradient methods. The training objective maximizes expected utility:

$$\mathcal{L}_{\text{utility}} = \mathbb{E}_{\Theta} [u(\mathcal{H}_{\text{com}}(Q))] \quad (19)$$

The utility function $u(\cdot)$ evaluates the quality of final answer $a^{(K)}$. We approximate the gradient using sampled topologies:

$$\nabla_{\Theta} \mathcal{L}_{\text{utility}} \approx \frac{1}{M} \sum_{m=1}^M u(a_m^{(K)}) \nabla_{\Theta} \log P(\mathcal{H}_m) \quad (20)$$

The system samples M different hypergraph topologies $\{\mathcal{H}_m\}$ during training. Each produces answer $a_m^{(K)}$. The complete training loss combines utility maximization with regularization:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{utility}} + \mathcal{L}_{\text{anchor}} + \mathcal{L}_{\text{sparse}} \quad (21)$$

The anchor regularization $\mathcal{L}_{\text{anchor}} = \|\mathbf{A}_{\text{anchor}} - \tilde{S}\|_F^2$ keeps generated topologies grounded in reasonable prior structures. The sparsity regularization $\mathcal{L}_{\text{sparse}} = \zeta \|\mathbf{W}\|_*$ ensures communication efficiency.

5 EXPERIMENTS

5.1 Datasets

We evaluate HyperAgent on three categories of benchmarks spanning diverse reasoning and generation tasks. The general reasoning category includes MMLU[20], a comprehensive benchmark containing multiple-choice questions across 57 subjects. The mathematical reasoning category comprises GSM8K[15] for grade school math problems, MultiArith[45] for arithmetic word problems, SVAMP[40] for structurally diverse math questions, and AQuA[35] for algebraic reasoning. The code generation category uses HumanEval[10], containing 164 programming tasks requiring function implementations. These benchmarks exhibit varying task complexities and collaboration demands, enabling comprehensive evaluation of hypergraph-based topology optimization.

5.2 Baselines

We compare HyperAgent against three categories of baselines. The Single-agent methods include CoT[60] for chain-of-thought prompting, ComplexCoT[18] for complexity-based prompting, Self-Consistency[58] for multiple sampling with voting, PHP[69] for

Table 1: Performance comparison with three types of baselines, including single-agent execution, static multi-agent topologies, and adaptive multi-agent frameworks. The best results are in bold, and the runner-ups are underlined. All multi-agent methods utilize five gpt-4-based agents. “Mul.”, “Ada.”, and “Rob.” indicate whether the method supports a multi-agent setting, whether it is task-adaptive, and whether it is adversarially robust, respectively. ✗, ⚡ and ✓ signifies no/partial/full support in these aspects.

Method	Mul.	Ada.	Rob.	MMLU	GSM8K	MultiArith	SVAMP	AQuA	HumanEval	Avg.
<i>Single-Agent Methods</i>										
Vanilla	✗	✗	✗	82.14	85.40	93.15	87.18	70.34	71.68	81.65
CoT	✗	✗	✗	82.65 _{↑0.51}	87.17 _{↑1.77}	94.79 _{↑1.64}	88.32 _{↑1.14}	73.91 _{↑3.57}	75.52 _{↑3.84}	83.73
ComplexCoT	✗	✗	✗	83.78 _{↑1.64}	87.62 _{↑2.22}	95.86 _{↑2.71}	90.17 _{↑2.99}	77.58 _{↑7.24}	74.94 _{↑3.26}	84.99
SC (CoT)	✗	✗	✗	82.66 _{↑0.52}	87.93 _{↑2.53}	96.88 _{↑3.73}	88.69 _{↑1.51}	75.08 _{↑4.74}	77.30 _{↑5.62}	84.75
SC (ComplexCoT)	✗	✗	✗	83.65 _{↑1.51}	86.14 _{↓0.74}	96.94 _{↑3.79}	89.72 _{↑2.54}	77.69 _{↑7.35}	77.94 _{↑6.26}	85.35
AutoGPT	✗	✗	✗	83.65 _{↑1.51}	86.14 _{↓0.74}	96.94 _{↑3.79}	89.72 _{↑2.54}	77.69 _{↑7.35}	77.94 _{↑6.26}	85.35
PHP	✓	✗	✗	83.45 _{↑1.31}	95.50 _{↑10.1}	<u>98.10</u> _{↑2.84}	90.02 _{↑3.44}	79.00 _{↑8.66}	82.96 _{↑11.36}	88.17
ReAct	✗	✗	✗	83.12 _{↑0.98}	88.24 _{↑2.84}	95.37 _{↑2.22}	89.15 _{↑1.97}	76.42 _{↑6.08}	78.35 _{↑6.67}	85.11
ToT	✗	✗	✗	83.89 _{↑1.75}	89.06 _{↑3.66}	96.52 _{↑3.37}	90.24 _{↑3.06}	77.95 _{↑7.61}	80.12 _{↑8.44}	86.30
GoT	✗	✗	✗	84.01 _{↑1.87}	89.47 _{↑4.07}	96.73 _{↑3.58}	90.38 _{↑3.20}	78.24 _{↑7.90}	81.26 _{↑9.58}	86.68
<i>Static Multi-Agent Topologies</i>										
Chain	✓	✗	✗	82.35 _{↑0.21}	85.57 _{↑0.17}	94.38 _{↑1.23}	83.41 _{↓3.77}	70.94 _{↑0.60}	80.88 _{↑9.20}	82.92
Star	✓	✗	✗	80.79 _{↓1.35}	85.55 _{↑0.15}	93.79 _{↓0.64}	88.09 _{↑0.91}	68.57 _{↓1.77}	75.65 _{↑3.97}	82.07
Tree	✓	✗	✗	81.89 _{↓0.25}	84.56 _{↓0.84}	94.60 _{↑1.45}	89.25 _{↑2.07}	72.84 _{↑2.50}	77.38 _{↑5.70}	83.42
Complete Graph	✓	✗	✗	83.15 _{↑1.01}	86.49 _{↑1.09}	97.20 _{↑4.05}	89.48 _{↑2.30}	<u>79.21</u> _{↑8.87}	83.75 _{↑12.07}	86.55
Random Graph	✓	✗	✗	83.76 _{↑1.62}	86.14 _{↑0.74}	95.46 _{↑2.31}	85.41 _{↓1.77}	74.07 _{↑3.73}	82.66 _{↑10.98}	84.58
<i>Adaptive Multi-Agent Frameworks</i>										
AutoGen	✓	✗	✗	82.13 _{↓0.01}	90.06 _{↑7.92}	93.80 _{↑0.65}	88.44 _{↓1.26}	73.65 _{↑3.31}	85.41 _{↑13.73}	85.58
MetaGPT	✓	✗	✗	83.24 _{↑1.10}	89.84 _{↑4.44}	95.12 _{↑1.97}	89.56 _{↑2.38}	76.18 _{↑5.84}	85.90 _{↑14.22}	86.64
LLM-Blender	✓	✗	✗	81.22 _{↓0.92}	89.17 _{↑3.77}	94.27 _{↑1.12}	88.77 _{↑1.59}	77.05 _{↑6.71}	84.52 _{↑12.84}	85.83
LLM-Debate	✓	✗	✓	83.69 _{↑1.55}	90.23 _{↑4.83}	96.27 _{↑3.12}	90.56 _{↑3.38}	77.52 _{↑7.18}	83.79 _{↑12.11}	87.01
DyLAN	✓	⚡	✓	80.16 _{↓1.98}	88.16 _{↑2.76}	94.27 _{↑1.12}	87.40 _{↑0.22}	74.16 _{↑3.82}	<u>89.70</u> _{↑18.02}	85.64
GPTSwarm	✓	⚡	✓	<u>83.98</u> _{↑1.84}	89.74 _{↑4.34}	97.84 _{↑4.69}	86.42 _{↓0.76}	78.16 _{↑7.82}	88.49 _{↑16.81}	87.44
AgentVerse	✓	⚡	✗	83.52 _{↑1.38}	90.12 _{↑4.72}	96.45 _{↑3.30}	89.87 _{↑2.69}	77.83 _{↑7.49}	86.24 _{↑14.56}	87.34
COPPER	✓	✓	✗	83.76 _{↑1.62}	91.35 _{↑5.95}	96.82 _{↑3.67}	90.18 _{↑3.00}	78.42 _{↑8.08}	87.53 _{↑15.85}	88.01
AutoAgents	✓	✓	✗	83.45 _{↑1.31}	90.58 _{↑5.18}	96.15 _{↑3.00}	89.64 _{↑2.46}	77.29 _{↑6.95}	86.87 _{↑15.19}	87.33
G-Designer	✓	✓	✓	84.25 _{↑2.11}	92.18 _{↑6.78}	97.56 _{↑4.41}	91.02 _{↑3.84}	78.94 _{↑8.60}	88.72 _{↑17.04}	88.78
AgentPrune	✓	✓	✓	84.15 _{↑2.01}	91.86 _{↑6.46}	97.38 _{↑4.23}	90.73 _{↑3.55}	78.65 _{↑8.31}	88.15 _{↑16.47}	88.49
AgentDropout	✓	✓	✓	84.08 _{↑1.94}	91.52 _{↑6.12}	97.21 _{↑4.06}	90.45 _{↑3.27}	78.51 _{↑8.17}	87.68 _{↑15.00}	88.24
HyperAgent (Ours)	✓	✓	✓	86.50 _{↑4.36}	<u>96.57</u> _{↑11.17}	99.30 _{↑6.15}	93.85 _{↑6.67}	81.97 _{↑11.63}	92.40 _{↑20.72}	91.77

progressive-hint prompting, AutoGPT[51] for autonomous task execution, ReAct[64] for synergizing reasoning and acting, ToT[63] for tree-based thought exploration, and GoT[3] for graph-based reasoning. Predefined multi-agent topologies include Chain, Star, Tree, Complete Graph, and Random Graph structures. Adaptive multi-agent frameworks include AutoGen[62] providing conversational coordination, MetaGPT[21] organizing software development agents, LLM-Blender[24] fusing multiple responses, LLM-Debate[16] enabling multi-agent debate, DyLAN[37] constructing dynamic layered networks, GPTSwarm[72] optimizing graph structures, AgentVerse[13] facilitating multi-agent collaboration, and G-Designer[50] using graph-based design. These baselines represent

state-of-the-art approaches in both static and dynamic topology design for multi-agent systems.

5.3 Evaluation Models and Metrics

We conduct experiments using two base language models accessed via OpenAI API: gpt-4-1106-preview and gpt-3.5-turbo-0125. Performance evaluation uses accuracy for multiple-choice questions on MMLU and AQuA, as well as for mathematical reasoning on GSM8K, MultiArith, and SVAMP. Code generation on HumanEval reports pass@1, measuring the percentage of problems solved correctly in the first attempt. All metrics are computed on test sets,

Table 2: Ablation study of different components in HyperAgent. We evaluate the contribution of each key component across six benchmarks. “ Δ ” denotes the performance drop compared to the full model. The results demonstrate that the hypergraph structure is the most critical component, followed by the VAE framework and task node.

Methods	MMLU	GSM8K	MultiArith	SVAMP	AQuA	HumanEval	Avg.	Δ Avg.
HyperAgent (Full)	86.50	96.57	99.30	93.85	81.97	92.40	91.77	-
w/o Hypergraph (Graph-based)	84.80	93.50	97.85	91.50	79.20	89.15	89.33	-2.44
w/o VAE (Fixed Topology)	85.20	94.80	98.50	92.80	80.85	90.50	90.44	-1.33
w/o Sparsity ($\zeta = 0$)	86.30	96.20	99.15	93.60	81.75	91.80	91.47	-0.30
w/o Task Node	85.65	94.95	98.65	92.45	80.50	90.85	90.51	-1.26
w/o Anchor Regularization	85.80	95.30	98.80	92.70	80.72	91.25	90.76	-1.01




Figure 3: Training dynamics of HyperAgent. (a) Loss components over training iterations. The utility loss (blue) steadily decreases while sparsity regularization (green) maintains stable constraint. (b) Validation accuracy improves and plateaus after 50 iterations. (c) Generated hypergraphs become progressively sparser during training, demonstrating the model learns efficient topologies.




Figure 4: (a) Effect of hyperedge size parameter k on performance and communication efficiency. (b) Impact of sparsity regularization coefficient on the performance-efficiency frontier. (c) Information propagation: graphs need multi-hop passing whereas hyperedges enable direct 1-step synchronization. (d) Visualization of the performance metrics and prompt token consumption.

with single-agent baselines using temperature 0 and multi-agent methods using temperature 1 to enable diverse responses.

5.4 Implementation Details

We access GPT via the OpenAI API, and mainly test on gpt-4 and gpt-3.5-turbo. We set temperature to 0 for single execution and single agent baselines and 1 for multi-agent methods. We set a summarizer agent to aggregate the dialogue history and produce the final solution $a^{(K)}$, with $K = 3$ across all experiments. The NodeEncoder(\cdot) is implemented using all-MiniLM-L6-v2, with the embedding dimension set to $D = 384$. The anchor hypergraph

H_{anchor} is predefined as a simple chain structure where each hyperedge connects two adjacent agents. The hypergraph encoder $HGCN_{\mu}$ and $HGCN_{\sigma}$ are two-layer hypergraph convolutional networks with hidden dimension 64. The decoder feedforward network FFN_d has hidden dimension 128. We set the rank $r = 16$ for low-rank approximation in Equation (14), the temperature $\tau = 1e - 2$ for Gumbel-Softmax sampling in Equation (13), and the sparsity coefficient $\zeta = 1e - 1$ for nuclear norm regularization. The hyperedge grouping parameter $k = 2$, meaning each collaboration unit connects 3 agents on average. The sampling times M are set as 10 for policy gradient approximation. We provide explicit agent




Figure 5: Performance vs. number of interaction rounds K. Accuracy improves with more rounds but exhibits diminishing returns after K=3.

profiling for multi-agent methods, following the classical configurations in LLM-MA systems, and use gpt-4 to generate agent profile pools. For all benchmarks, we merely use $B' \in \{40, 80\}$ queries for optimization.

5.5 Main Result

Table 1 presents comprehensive performance comparisons across six benchmarks spanning general reasoning, mathematical problem solving, and code generation tasks. HyperAgent consistently outperforms all baseline methods, achieving an average accuracy of 91.77% across all tasks, which represents substantial improvements over the strongest competitors including G-Designer at 88.78% and AgentDropout at 88.24%. Notably, HyperAgent demonstrates exceptional performance on mathematical reasoning benchmarks, reaching 96.57% on GSM8K and 99.30% on MultiArith, while also excelling in code generation with 92.40% pass@1 on HumanEval. The results highlight HyperAgent’s advantage in capturing group collaboration patterns through hyperedges rather than pairwise edges, enabling more efficient information aggregation within collaboration units. Furthermore, HyperAgent is the only method that fully supports multi-agent settings, task-adaptive topology generation, and adversarial robustness simultaneously, as indicated by the checkmarks in the rightmost columns. The performance gains are particularly pronounced on complex tasks requiring intensive coordination, demonstrating that hypergraph-based topology optimization effectively addresses the limitations of graph-based approaches in modeling multi-agent collaboration.

5.6 Model Analysis

Figure 4 provides a direct comparison between graph-based and hypergraph-based topology learning during training. The hypergraph approach demonstrates superior convergence properties, with training loss decreasing more rapidly and stabilizing at approximately 0.25 compared to the graph-based method’s 0.65 after 1200 training steps. This performance advantage stems from the hypergraph’s ability to directly model collaborative units through single hyperedges connecting multiple agents, eliminating the need for multi-hop information propagation required by pairwise edges in traditional graphs. When agents collaborate on shared subtasks, the hypergraph representation enables one-step synchronization where all participating agents contribute to and receive from a unified hyperedge representation simultaneously. In contrast, graph-based methods must decompose group interactions into multiple

pairwise connections, leading to sequential information flow and potential degradation over intermediate steps. The training dynamics demonstrate that optimizing over hyperedge space rather than quadratic pairwise edge space significantly reduces computational complexity while preserving richer semantic structure, ultimately translating to better task performance.

5.7 Hyper-parameter Analysis

Figure 5 investigates the impact of interaction rounds K on system performance across three representative benchmarks. Accuracy consistently improves as K increases from 1 to 3, with MMLU rising from 82.5% to 86.5%, HumanEval improving from 84.2% to 91.6%, and GSM8K advancing from 90.1% to 96.6%. However, the performance gains exhibit diminishing returns beyond K equals 3, with only marginal improvements observed at K equals 4 and 5. This phenomenon suggests that three interaction rounds provide sufficient capacity for effective collaboration, allowing agents to propose initial solutions, receive feedback, and produce refined outputs. Additional rounds contribute minimal value while incurring increased computational costs through redundant communication. The consistency of this pattern across diverse task types indicates that K equals 3 represents an optimal balance between collaboration effectiveness and efficiency, supporting our design choice to fix K at 3 throughout all experiments.

5.8 Ablation Study

Table 2 systematically evaluates the contribution of each component in HyperAgent through ablation experiments across six benchmarks. Removing the hypergraph structure and reverting to graph-based pairwise edges causes the most significant performance degradation, with average accuracy dropping 2.44 percentage points from 91.77% to 89.33%, demonstrating that direct hyperedge representation of collaboration units constitutes the most critical architectural choice. The VAE framework for dynamic topology generation proves to be the second most important component, as replacing it with fixed topology results in a 1.33 percentage point decrease, highlighting the value of task-adaptive structure learning. Disabling sparsity regularization leads to a 0.30 percentage point drop, while removing the task-specific virtual node reduces accuracy by 1.26 percentage points, confirming its role in facilitating global information flow. Figure 3 provides deeper insights into training dynamics, showing that utility loss steadily decreases while validation accuracy improves rapidly and plateaus after approximately 50 epochs at around 90%, demonstrating that the model learns to construct efficient topologies that balance communication overhead with coordination effectiveness.

6 CONCLUSION

In this work, we propose HyperAgent, a hypergraph-based multi-agent communication framework that connects agents sharing subtasks via hyperedges for efficient single-step information aggregation. A variational autoencoder with sparsity regularization generates task-adaptive topologies. Experiments on multiple benchmarks show that HyperAgent outperforms state-of-the-art methods while significantly reducing communication overhead.

7 ACKNOWLEDGMENTS

This work was supported by the Natural Science Foundation of Guangdong Province, China. “Research on Key Theories and Technologies for Nano-learning”

REFERENCES

- [1] Omar Adjali, Olivier Ferret, Sahar Ghannay, and Hervé Le Borgne. 2024. Multi-Level Information Retrieval Augmented Generation for Knowledge-based Visual Question Answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. 16499–16513.
- [2] Abdullah Almaatouq, Mohammed Alsobay, Ming Yin, and Duncan J Watts. 2021. Task Complexity Moderates Group Synergy. *Proceedings of the National Academy of Sciences* 118, 36 (2021).
- [3] Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Michal Podstawski, Hubert Niewiadomski, Piotr Nyczyk, and Torsten Hoefler. 2024. Graph of Thoughts: Solving Elaborate Problems with Large Language Models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 17682–17690.
- [4] Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, et al. 2024. Graph of Thoughts: Solving Elaborate Problems with Large Language Models. *AAAI Conference on Artificial Intelligence* (2024).
- [5] Maciej Besta, Florim Memedi, Zhenyu Zhang, Robert Gerstenberger, Guangyuan Piao, Nils Blach, Piotr Nyczyk, Marcin Copik, Grzegorz Kwasniewski, Jurgen Muller, et al. 2024. Demystifying Chains, Trees, and Graphs of Thoughts. *arXiv preprint arXiv:2401.14295* (2024).
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems* 33 (2020), 1877–1901.
- [7] Sebastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of Artificial General Intelligence: Early Experiments with GPT-4. *arXiv preprint arXiv:2303.12712* (2023).
- [8] Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2024. ChatEval: Towards Better LLM-based Evaluators through Multi-Agent Debate. In *International Conference on Learning Representations*.
- [9] Justin Chen, Swarnadeep Saha, and Mohit Bansal. 2024. ReConcile: Round-Table Conference Improves Reasoning via Consensus among Diverse LLMs. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*. 6827–6844.
- [10] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating Large Language Models Trained on Code. *arXiv preprint arXiv:2107.03374* (2021).
- [11] Silin Chen, Shaoxin Lin, Xiaodong Gu, Yuling Shi, Heng Lian, Longfei Yun, Dong Chen, Weiguo Sun, Lin Cao, and Qianxiang Wang. 2025. Swe-exp: Experience-driven software issue resolution. *arXiv preprint arXiv:2507.23361* (2025).
- [12] Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chen Qian, Chi-Min Chan, Yujia Qin, Yaxi Lu, Ruobing Xie, et al. 2024. AgentVerse: Facilitating Multi-agent Collaboration and Exploring Emergent Behaviors in Agents. In *International Conference on Learning Representations*.
- [13] Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chen Qian, Chi-Min Chan, Yujia Qin, Yaxi Lu, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. 2023. AgentVerse: Facilitating Multi-Agent Collaboration and Exploring Emergent Behaviors. *arXiv preprint arXiv:2308.10848* (2023).
- [14] Weize Chen, Ziming You, Ran Li, Yitong Guan, Chen Qian, Chenyang Zhao, Cheng Yang, Ruobing Xie, Zhiyuan Liu, and Maosong Sun. 2024. Internet of Agents: Weaving a Web of Heterogeneous Agents for Collaborative Intelligence. *arXiv preprint arXiv:2407.07061* (2024).
- [15] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training Verifiers to Solve Math Word Problems. *arXiv preprint arXiv:2110.14168* (2021).
- [16] Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. 2024. Improving Factuality and Reasoning in Language Models through Multiagent Debate. In *International Conference on Machine Learning*.
- [17] Yixiong Fang, Tianran Sun, Yuling Shi, and Xiaodong Gu. 2025. Attentionrag: Attention-guided context pruning in retrieval-augmented generation. *arXiv preprint arXiv:2503.10720* (2025).
- [18] Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. 2022. Complexity-Based Prompting for Multi-Step Reasoning. *arXiv preprint arXiv:2210.00720* (2022).
- [19] Mustafa Hajji, Ghada Zamzmi, Theodore Papamarkou, Aldo Guzman-Saenz, Tolga Birdal, and Michael T Schaub. 2023. Combinatorial Complexes: Bridging the Gap Between Cell Complexes and Hypergraphs. *57th Asilomar Conference on Signals, Systems, and Computers* (2023), 799–803.
- [20] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring Massive Multitask Language Understanding. In *International Conference on Learning Representations (ICLR)*.
- [21] Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jurgen Schmidhuber. 2024. MetaGPT: Meta Programming for A Multi-Agent Collaborative Framework. In *The Twelfth International Conference on Learning Representations (ICLR)*.
- [22] John J Hopfield. 1982. Neural Networks and Physical Systems with Emergent Computational Abilities. *Proceedings of the National Academy of Sciences* 79, 8 (1982), 2554–2558.
- [23] Yiming Huang, Jianwen Luo, Yan Yu, Yitong Zhang, Fanguy Lei, Yifan Wei, Shizhu He, Lifu Huang, Xiao Liu, Jun Zhao, and Kang Liu. 2024. DA-Code: Agent Data Science Code Generation Benchmark for Large Language Models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. 13487–13521.
- [24] Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. 2023. LLM-Blender: Ensembling Large Language Models with Pairwise Ranking and Generative Fusion. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)*. 14165–14178.
- [25] Wei Ju, Zhengyang Mao, Siyu Yi, Yifang Qin, Yiyang Gu, Zhiping Xiao, Yifan Wang, Xiao Luo, and Ming Zhang. 2024. Hypergraph-enhanced Dual Semi-supervised Graph Classification. *arXiv preprint arXiv:2405.04773* (2024).
- [26] Akbir Khan, John Hughes, Dan Valentine, Laura Ruis, Kshitij Sachan, Ansh Radhakrishnan, Edward Grefenstette, Samuel R Bowman, Tim Rocktaschel, and Ethan Perez. 2024. Debating with More Persuasive LLMs Leads to More Truthful Answers. In *International Conference on Machine Learning*.
- [27] Omar Khattab, Arnab Sintra, Keshav Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Varshney, Mohammadreza Komeili, Nader Moazam, Yuval Kirstain, Matei Zaharia, and Christopher Ré. 2024. DSPy: Compiling Declarative Language Model Calls into State-of-the-Art Pipelines. In *International Conference on Learning Representations*.
- [28] Boyi Li, Zhonghan Zhao, Der-Horng Lee, and Gaoang Wang. 2025. Adaptive Graph Pruning for Multi-Agent Communication. *arXiv preprint arXiv:2506.02951* (2025).
- [29] Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. CAMEL: Communicative Agents for “Mind” Exploration of Large Language Model Society. In *Advances in Neural Information Processing Systems*, Vol. 36. 51991–52008.
- [30] Han Li, Yuling Shi, Shaoxin Lin, Xiaodong Gu, Heng Lian, Xin Wang, Yantao Jia, Tao Huang, and Qianxiang Wang. 2025. Swe-debate: Competitive multi-agent debate for software issue resolution. *arXiv preprint arXiv:2507.23348* (2025).
- [31] Shiyuan Li, Yixin Liu, Qingsong Wen, Chengqi Zhang, and Shirui Pan. 2025. Assemble Your Crew: Automatic Multi-agent Communication Topology Design via Autoregressive Graph Generation. *arXiv preprint arXiv:2507.18224* (2025).
- [32] Yuan Li, Yilei Yao, Dong Li, Huazheng Zhang, and Tong Zhao. 2024. G-Designer: Architecting Multi-agent Communication Topologies via Graph Neural Networks. *arXiv preprint arXiv:2410.11782* (2024).
- [33] Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Rui Wang, Yuju Yang, Zhaopeng Tu, and Shuming Shi. 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*. 16310–16324.
- [34] Bor Jiun Lin and Chun-Yi Lee. 2024. HGAP: Boosting Permutation Invariant and Permutation Equivariant in Multi-Agent Reinforcement Learning via Graph Attention Network. In *International Conference on Machine Learning*.
- [35] Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. Program Induction by Rationale Generation: Learning to Solve and Explain Algebraic Word Problems. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*, Vol. 1. 158–167.
- [36] Runze Liu, Jiakang Wang, Yuling Shi, Zhihui Xie, Chenxin An, Kaiyan Zhang, Jian Zhao, Xiaodong Gu, Lei Lin, Wengping Hu, et al. 2025. Attention as a Compass: Efficient Exploration for Process-Supervised RL in Reasoning Models. *arXiv preprint arXiv:2509.26628* (2025).
- [37] Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. 2023. Dynamic LLM-Agent Network: An LLM-agent Collaboration Framework with Agent Team Optimization. *arXiv preprint arXiv:2310.02170* (2023).
- [38] Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. 2023. A Dynamic LLM-Powered Agent Network for Task-Oriented Agent Collaboration. *arXiv preprint arXiv:2310.02170* (2023).
- [39] Yaru Niu, Rohan R. Paleja, and Matthew Craig Gombolay. 2021. Multi-Agent Graph-Attention Communication and Teaming. In *Adaptive Agents and Multi-Agent Systems*. <https://api.semanticscholar.org/CorpusID:234351960>
- [40] Arkil Patel, Satwik Bhattamishra, and Navin Goyal. 2021. Are NLP Models Really Able to Solve Simple Math Word Problems?. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics*

- Human Language Technologies (NAACL-HLT)*. 2080–2094.
- [41] Weihang Peng, Yuling Shi, Yuhang Wang, Xinyun Zhang, Beijun Shen, and Xiaodong Gu. 2025. SWE-QA: Can Language Models Answer Repository-level Code Questions? *arXiv preprint arXiv:2509.14635* (2025).
- [42] Giorgio Piatti, Zhijing Jin, Max Kleiman-Weiner, Bernhard Scholkopf, Mrinmaya Sachan, and Rada Mihalcea. 2024. Cooperate or Collapse: Emergence of Sustainability Behaviors in a Society of LLM Agents. *arXiv preprint arXiv:2404.16698* (2024).
- [43] Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, et al. 2024. ChatDev: Communicative Agents for Software Development. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*. 14174–14190.
- [44] Chen Qian, Zihao Xie, Yifei Wang, Wei Liu, Kunlun Zhu, Hanchen Xia, Yufan Dang, Zhuoyun Du, Weize Chen, Cheng Yang, Zhiyuan Liu, and Maosong Sun. 2025. Scaling Large Language Model-based Multi-Agent Collaboration. In *International Conference on Learning Representations*.
- [45] Subhro Roy and Dan Roth. 2015. Solving General Arithmetic Word Problems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1743–1752.
- [46] Yuling Shi, Yichun Qian, Hongyu Zhang, Beijun Shen, and Xiaodong Gu. 2025. LongCodeZip: Compress Long Context for Code Language Models. *arXiv preprint arXiv:2510.00446* (2025).
- [47] Yuling Shi, Songsong Wang, Chengcheng Wan, Min Wang, and Xiaodong Gu. 2024. From code to correctness: Closing the last mile of code generation with hierarchical debugging. *arXiv preprint arXiv:2410.01215* (2024).
- [48] Yuling Shi, Hongyu Zhang, Chengcheng Wan, and Xiaodong Gu. 2024. Between Lines of Code: Unraveling the Distinct Patterns of Machine and Human Programmers. In *2025 IEEE/ACM 47th International Conference on Software Engineering (ICSE)*. IEEE Computer Society, 51–62.
- [49] Yashar Talebirad and Amirhossein Nadiri. 2023. Multi-Agent Collaboration: Harnessing the Power of Intelligent LLM Agents. *arXiv preprint arXiv:2306.03314* (2023).
- [50] Xiangru Tang, Anni Zou, Zhuosheng Zhang, Yilun Zhao, Xingyao Zhang, Arman Cohan, and Mark Gerstein. 2024. MedAgents: Large Language Models as Collaborators for Zero-shot Medical Reasoning. *arXiv preprint arXiv:2311.10537* (2024).
- [51] Torantulino and Contributors. 2023. AutoGPT: An Autonomous GPT-4 Experiment. <https://github.com/Significant-Gravitas/AutoGPT>.
- [52] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothee Lacroix, Baptiste Roziere, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971* (2023).
- [53] Khanh-Tung Tran, Dung Dao, Minh-Duong Nguyen, Quoc-Viet Pham, Barry O’Sullivan, and Hoang D Nguyen. 2025. Multi-Agent Collaboration Mechanisms: A Survey of LLMs. *arXiv preprint arXiv:2501.06322* (2025).
- [54] Trieu H Trinh, Yuhuai Wu, Quoc V Le, He He, and Thang Luong. 2024. Solving Olympiad Geometry without Human Demonstrations. *Nature* 625 (2024), 476–482.
- [55] Chaofan Wang, Tingrui Yu, Jie Wang, Dong Chen, Wenrui Zhang, Yuling Shi, Xiaodong Gu, and Beijun Shen. 2025. EVOC2RUST: A Skeleton-guided Framework for Project-Level C-to-Rust Translation. *arXiv preprint arXiv:2508.04295* (2025).
- [56] Haotian Wang, Xiyuan Du, Weijiang Yu, Qianglong Chen, Kun Zhu, Zheng Chu, Lian Yan, and Yi Guan. 2024. Learning to Break: Knowledge-Enhanced Reasoning in Multi-Agent Debate System. In *arXiv preprint arXiv:2312.04854*.
- [57] Qineng Wang, Zihao Wang, Ying Su, Hanghang Tong, and Yangqiu Song. 2024. Rethinking the Bounds of LLM Reasoning: Are Multi-Agent Discussions the Key?. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*. 6024–6041.
- [58] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-Consistency Improves Chain of Thought Reasoning in Language Models. *arXiv preprint arXiv:2203.11171* (2022).
- [59] Zhexuan Wang, Yutong Wang, Xuebo Liu, Liang Ding, Miao Zhang, Jie Liu, and Min Zhang. 2025. AgentDropout: Dynamic Agent Elimination for Token-Efficient and High-Performance LLM-Based Multi-Agent Collaboration. *arXiv preprint arXiv:2503.18891* (2025).
- [60] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *Advances in Neural Information Processing Systems* 35 (2022), 24824–24837.
- [61] Anita Williams Woolley, Christopher F Chabris, Alex Pentland, Nada Hashmi, and Thomas W Malone. 2010. Evidence for a Collective Intelligence Factor in the Performance of Human Groups. *Science* 330, 6004 (2010), 686–688.
- [62] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, et al. 2023. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation. In *ICLR 2024 Workshop on LLM Agents*.
- [63] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. *Advances in Neural Information Processing Systems* 36 (2023), 11809–11822.
- [64] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. ReAct: Synergizing Reasoning and Acting in Language Models. In *International Conference on Learning Representations (ICLR)*.
- [65] Wenhao Zeng, Yaoning Wang, Chao Hu, Yuling Shi, Chengcheng Wan, Hongyu Zhang, and Xiaodong Gu. 2025. Pruning the unsurprising: Efficient code reasoning via first-token surprisal. *arXiv preprint arXiv:2508.05988* (2025).
- [66] Guibin Zhang, Yanwei Yue, Xiangguo Sun, Guancheng Wan, Miao Yu, Junfeng Fang, Kun Wang, Tianlong Chen, and Dawei Cheng. 2025. G-Designer: Architecting Multi-agent Communication Topologies via Graph Neural Networks. In *International Conference on Machine Learning*.
- [67] Jintian Zhang, Xin Xu, Ningyu Zhang, Ruibo Liu, Bryan Hooi, and Shumin Deng. 2024. Exploring Collaboration Mechanisms for LLM Agents: A Social Psychology View. In *ICLR 2024 Workshop on LLM Agents*.
- [68] Jun Zhao, Can Zu, Xu Hao, Yi Lu, Wei He, Yiwen Ding, Tao Gui, Qi Zhang, and Xuanjing Huang. 2024. LONGAGENT: Achieving Question Answering for 128k-Token-Long Documents through Multi-Agent Collaboration. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. 16310–16324.
- [69] Chuanyang Zheng, Zhengying Liu, Enze Xie, Zhenguang Li, and Yu Li. 2023. Progressive-Hint Prompting Improves Reasoning in Large Language Models. *arXiv preprint arXiv:2304.09797* (2023).
- [70] Wangchunshu Zhou, Yixin Ou, Shengwei Ding, Long Li, Jialong Wu, Tiannan Wang, Jiamin Chen, Shuai Wang, Xiaohua Xu, Ningyu Zhang, et al. 2024. Symbolic Learning Enables Self-Evolving Agents. *arXiv preprint arXiv:2406.18532* (2024).
- [71] Yuchen Zhuang, Xiang Chen, Tong Yu, Saayan Mitra, Victor Bursztyn, Ryan A Rossi, Somdeb Sarkhel, and Chao Zhang. 2024. ToolChain*: Efficient Action Space Navigation in Large Language Models with A* Search. In *arXiv preprint arXiv:2310.13227*.
- [72] Mingchen Zhuge, Wenyi Wang, Louis Kirsch, Francesco Faccio, Dmitrii Khizbullin, and Jürgen Schmidhuber. 2024. Language Agents as Optimizable Graphs. In *International Conference on Machine Learning*.
- [73] Chang Zong, Yuchen Yan, Weiming Lu, Jian Shao, Yongfeng Huang, Heng Chang, and Yueting Zhuang. 2024. Triad: A Framework Leveraging a Multi-Role LLM-based Agent to Solve Knowledge Base Question Answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. 1698–1710.