

Research Article

AgentNet: Decentralized Evolutionary Coordination for LLM-based Multi-Agent Systems

Yingxuan Yang¹, Huacan Chai¹, Weinan Zhang^{1,2}

1. Shanghai Jiao Tong University, Shanghai, China; 2. SII

The rapid advancement of Large Language Models (LLMs) has catalyzed the development of multi-agent systems, where multiple LLM-based agents collaborate to solve complex tasks. However, existing systems predominantly rely on centralized coordination, which introduces scalability bottlenecks, limits adaptability, and creates single points of failure. Additionally, concerns over privacy and proprietary knowledge sharing hinder cross-organizational collaboration, leading to siloed expertise. To address these challenges, we propose AgentNet, a decentralized, Retrieval-Augmented Generation (RAG)-based framework that enables LLM-based agents to autonomously evolve their capabilities and collaborate efficiently in a Directed Acyclic Graph (DAG)-structured network. Unlike traditional multi-agent systems that depend on static role assignments or centralized control, AgentNet allows agents to specialize dynamically, adjust their connectivity, and route tasks without relying on predefined workflows. AgentNet's core design is built upon several key innovations: (1) Fully Decentralized Paradigm: Removing the central orchestrator, allowing agents to coordinate and specialize autonomously, fostering fault tolerance and emergent collective intelligence. (2) Dynamically Evolving Graph Topology: Real-time adaptation of agent connections based on task demands, ensuring scalability and resilience. (3) Adaptive Learning for Expertise Refinement: A retrieval-based memory system that enables agents to continuously update and refine their specialized skills. By eliminating centralized control, AgentNet enhances fault tolerance, promotes scalable specialization, and enables privacy-preserving collaboration across organizations. Through decentralized coordination and minimal data exchange, agents can leverage diverse knowledge sources while safeguarding sensitive information. Experimental results demonstrate that AgentNet outperforms traditional centralized multi-agent systems, significantly improving efficiency,

adaptability, and scalability in dynamic environments, making it a promising foundation for next-generation autonomous, privacy-respecting multi-agent ecosystems.

Yingxuan Yang and Huacan Chai equally contributed to this work.

Corresponding author: Weinan Zhang, wnzhang@sjtu.edu.cn

1. Introduction

Recently, large language models (LLMs) have demonstrated remarkable capabilities in various domains, ranging from basic text understanding to complex reasoning and multimodal integration^{[1][2][3]}. Consequently, LLM-based agents have exhibited exceptional performance in numerous tasks, including scientific discovery^[4], automated reasoning^[5], and website operations^[6]. However, due to the lack of collective intelligence and collaboration, LLM-based Single Agents struggle to address the complex challenges encountered in the real world. By leveraging collective intelligence through parallel decision-making or workflow collaboration, LLM-based Multi-Agent Systems (MAS) have emerged as a promising framework for tackling complex real-world problems^{[7][8][9]}. However, most MAS following the workflow collaboration paradigm rely heavily on a centralized controller or a static, predefined workflow to allocate tasks among agents with fixed roles^{[10][11][12][13][14]}. While such designs simplify orchestration, they also introduce inherent constraints—including limited scalability, a single point of failure, and challenges to cross-organizational collaboration due to privacy and proprietary knowledge concerns.

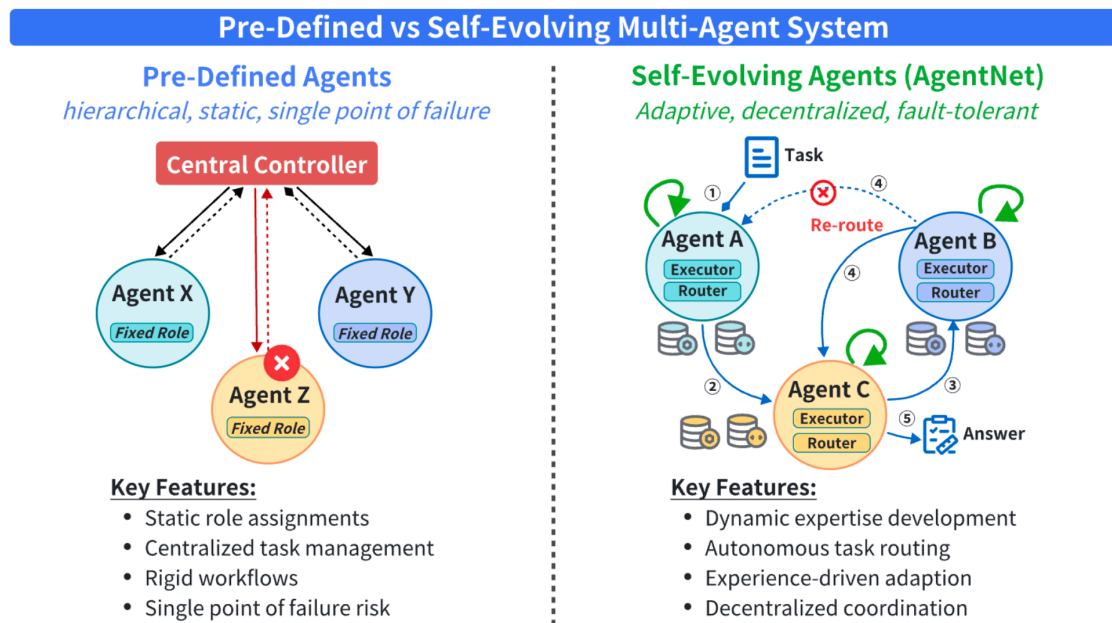


Figure 1. The illustration contrasts Pre-Defined Multi-Agents (hierarchical, static, with centralized control and single point of failure) against Self-Evolving Agents/AgentNet (adaptive, decentralized, and fault-tolerant with dynamic expertise development).

A more critical drawback arises from the inability of these systems to adapt to real-time fluctuations in agent performance or rapidly changing task requirements. Relying on a central controller inflates deployment complexity and restricts dynamic role reassignment, rendering the system vulnerable when the controller fails or becomes overloaded. Furthermore, rigid role definitions prevent agents from flexibly leveraging their full expertise in dynamic environments, ultimately undermining both efficiency and scalability. Taken together, these limitations highlight the need for more decentralized, fault-tolerant approaches that support dynamic task allocation, enhance adaptability, and safeguard privacy across organizational boundaries.

Beyond the scalability and failure-tolerance issues previously discussed, centralized architectures become even more problematic when organizations attempt to collaborate at scale^{[9][15]}. Each institution—be it an enterprise, research lab, or government agency—typically holds proprietary expertise, sensitive data, or both. In a centralized setup, concerns over data ownership, privacy regulations, and inconsistent governance often create barriers that prevent free exchange of knowledge. As a result, LLM-based agents contributed by multiple organizations remain siloed, unable to fully capitalize on each other’s specialized capabilities or datasets. This fragmentation not only hampers collective intelligence but also highlights

the urgency of developing secure, decentralized collaboration mechanisms. By enabling each participant to maintain and share only the minimal necessary information, these mechanisms address data confidentiality requirements while still allowing for a richer, more collaborative multi-agent ecosystem.

To address these challenges in multi-agent systems, we propose AgentNet, a novel framework designed to foster adaptive agent evolution, optimize task coordination, and preserve privacy. By eliminating the reliance on a central orchestrator, AgentNet enables agents to dynamically reconfigure their connections and redistribute tasks, forming a self-organizing, fault-tolerant architecture. Within this architecture, tasks are efficiently routed via a Directed Acyclic Graph (DAG)^{[16][17]}, which supports flexible collaboration and prevents cyclic dependencies.

Unlike traditional MAS frameworks that fix each agent's role, AgentNet incorporates a retrieval-based RAG^{[18][19][6]} memory mechanism to refine agent expertise over time. Each agent maintains a limited-capacity pool of successful task trajectories; when a new task arises, it retrieves the most relevant trajectories through few-shot learning, thus improving decision-making. To prevent memory overflow, agents autonomously prune less pertinent trajectories, ensuring the retention of valuable knowledge. This dynamic specialization strategy not only streamlines task allocation and agent adaptation but also supports a highly scalable and privacy-respecting environment for multi-agent collaboration.

AgentNet's core design is built upon several key innovations:

- **Fully Decentralized Paradigm:** By removing the need for a central orchestrator, AgentNet fosters emergent collective intelligence. Decision-making authority is distributed across all agents, thereby eliminating single points of failure and allowing each agent to coordinate, delegate, and specialize as conditions evolve. This approach leads to a self-organizing and fault-tolerant architecture that can rapidly respond to new tasks and unforeseen challenges. This decentralized setup also encourages emergent collective intelligence—in other words, agents can collectively discover and refine optimal strategies rather than waiting for instructions from a central controller.
- **Dynamically Evolving Graph Topology:** AgentNet employs a network structure in which both nodes (agents) and edges (agent-to-agent connections) adapt in real time based on task demands and agent performance. Rather than relying on fixed workflows, the system continuously reconfigures its topology to optimize information flow and task distribution, ensuring scalability and resilience in complex, changing environments.

- **Adaptive Learning Mechanism for Expertise Refinement:** AgentNet’s third innovation is its retrieval-based memory system, enabling agents to capture and update knowledge from successful task trajectories. This mechanism continuously refines each agent’s specialized skills without altering the network’s topology, allowing agents to avoid over-reliance on outdated information and sustain high performance in dynamic scenarios.

Moreover, each of these three innovations inherently enhances data privacy. By eliminating a central orchestrator, every agent stores and processes knowledge locally, sharing only minimal task-relevant metadata. The dynamic graph topology further confines data flow to necessary agent-to-agent interactions, reducing the exposure of sensitive information. Meanwhile, the retrieval-based memory mechanism restricts how much and how long data is retained, pruning outdated trajectories so that only high-value knowledge persists. Together, these design choices safeguard privacy and intellectual property, particularly crucial for cross-organizational collaborations.

Our experimental evaluation shows that AgentNet significantly outperforms traditional LLM-based multi-agent frameworks in dynamic environments, demonstrating improved task efficiency, specialization stability, and adaptive learning speed. These results highlight the effectiveness of decentralized evolutionary coordination in large-scale AI ecosystems.

2. Related Work

2.1. LLM-based Multi-Agent Systems

The development of LLM-based multi-agent systems (LaMAS)^[9] has advanced rapidly in recent years. Early frameworks, such as AutoGen^[11] and MetaGPT^[10], made significant strides in establishing foundational architectures for orchestrating multiple LLM agents through structured workflows. AutoGen provided a flexible framework for defining agent interactions, while MetaGPT incorporated software development principles to enhance collaboration. These centralized frameworks proved effective for managing multi-agent interactions. However, they also faced inherent challenges, including limited scalability, single points of failure, and difficulty in dynamically adapting to evolving tasks or incorporating new expertise.

In response to these limitations, more recent frameworks such as AgentScope^[20] and MegaAgent^[12] have focused on improving robustness and scalability. AgentScope introduced modular

design patterns to enhance system reliability, while MegaAgent employed hierarchical structures to scale agent interactions. Although these frameworks offer improvements, they still operate under centralized control paradigms, with a master agent delegating tasks, which continues to lead to scalability bottlenecks and single points of failure. Moreover, existing LaMAS implementations predominantly utilize single-source LLMs, lacking the integration of heterogeneous models. Their workflows are typically static, unable to dynamically allocate resources based on task complexity, further constraining adaptability.

In contrast, AgentNet introduces a novel decentralized approach, addressing these challenges by enabling agents to autonomously refine their expertise and dynamically allocate resources. AgentNet supports scalable, fault-tolerant collaboration without reliance on a central orchestrator, overcoming the limitations of centralized frameworks.

2.2. Evolutionary Agent Systems

Inspired by natural evolution, recent researchers have explored evolutionary approaches to automate and optimize agent behaviors and workflows in LaMAS. Existing efforts can be broadly categorized into the following areas:

- **Prompt Evolution and Optimization** – Techniques such as PromptBreeder^[21], DsPy^[22] and AgentPrune^[23] apply evolutionary algorithms to iteratively refine prompt generation, improving task performance through better input design.
- **Inter-Agent Topology Optimization** – Systems like GPTSwarm^[24], DyLAN^[25], and G-Designer^[26] focus on evolving the structural organization of agent interactions. These works aim to optimize communication patterns, task allocation, and collaboration efficiency within multi-agent networks.
- **Agent Role and Persona Specialization** – Frameworks such as AgentVerse and MorphAgent^[27]^[28] refine agent roles and profiles, enabling more effective specialization and coordination among agents in complex tasks.

While these evolutionary approaches have shown promise, they primarily focus on individual agent adaptation rather than collective coordination. Additionally, they still tend to operate within centralized control structures, which limits their scalability and dynamic adaptability. Recent frameworks like AgentSquare^[29] and AFlow^[30] have begun to formalize automated design processes for agentic systems,

improving system-level orchestration and workflow automation. Another key direction is self-adaptive agent architectures, where agents adjust their strategies in real-time based on feedback and accumulated experience. For example, EvoMAC^[31] combines reinforcement learning with evolutionary algorithms to optimize agent decision-making and policy updates.

However, these approaches are often limited to single-agent adaptation and lack mechanisms for decentralized specialization and coordination across large-scale agent collectives. While EvoMAC and other systems focus on optimizing individual agents, they are not designed for scalable, multi-agent, decentralized collaboration. In contrast, AgentNet integrates evolutionary learning with decentralized control, enabling heterogeneous agents to dynamically evolve their roles, adapt their strategies in real-time, and collaborate flexibly across a large-scale multi-agent system. This integration of evolutionary learning with decentralized control makes AgentNet a more suitable framework for real-time, adaptive, and scalable multi-agent collaboration.

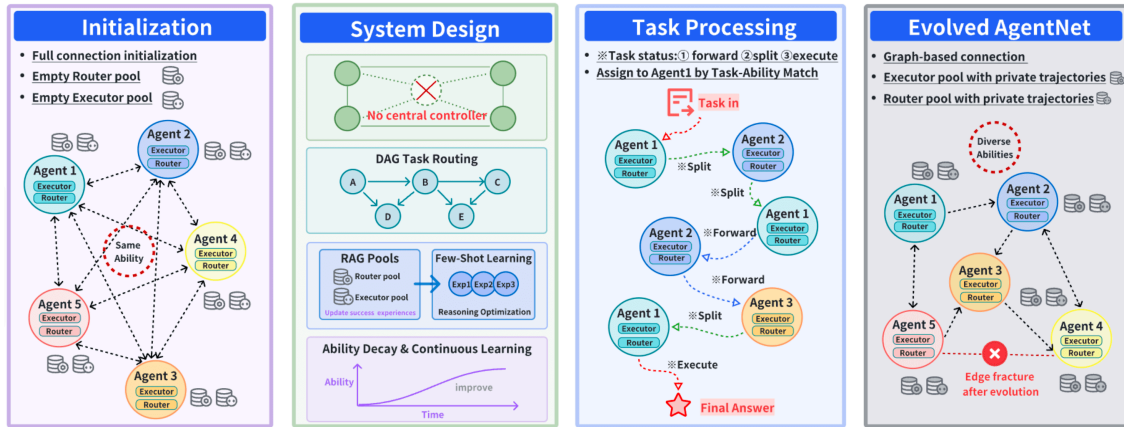


Figure 2. Illustration of AgentNet. Initially, agents are fully connected and equipped with executors and routers. The system eliminates the need for a central controller, using a DAG for dynamic task routing and agents leveraging RAG pools and few-shot learning. In the evolved phase, the network adapts with agents developing private trajectories and diversified abilities, ensuring scalability, fault tolerance, and continuous evolution of expertise..

3. Methodology

3.1. Overview of AgentNet Architecture

Unlike traditional MAS frameworks with fixed agent roles and rigid workflows using central coordinators, AgentNet creates a privacy-preserving, collective intelligence multi-agent system with high scalability and failure-tolerance by leveraging an innovative framework, consisting of a fully decentralized network architecture, a dynamic task allocation mechanism, and an adaptive agent learning method.

We begin with a brief introduction of AgentNet, including notation and basic architectures of agents employed.

Formally, we define AgentNet as a tuple $\mathcal{G} = (\mathcal{A}, \mathcal{E})$, where $\mathcal{A} = \{a_1, a_2, \dots, a_n\}$ represents the set of autonomous agents, $\mathcal{C} = \{c_1, c_2, \dots, c_n\}$ represents each agent's ability, and $\mathcal{E} \subseteq \mathcal{A} \times \mathcal{A}$ represents the communication connections between agents, specifically $e_{i,j} \in \mathcal{E}$ referring to a unidirectional connection from Agent a_i to Agent a_j .

For each agent $a_i \in \mathcal{A}$ contains two key components. rou_i is an agent router, responsible for analyzing received routing queries and making routing decisions. exe_i is an agent executor, responsible for responding to executing queries through operations and tools.

The two components mentioned above are underpinned by a substantial LLM that leverages its extensive knowledge and understanding to solve specific problems. Furthermore, both rou_i and exe_i in a_i maintain fixed-size memory modules M_i^{rou} and M_i^{exe} , respectively, providing a_i with powerful adaptive evolutionary capabilities by storing and utilizing the agent's experiences through the RAG mechanism.

For optimization, AgentNet will be given a series of tasks denoted as $\mathbf{T} = \{t_1, t_2, \dots, t_M\}$ to resolve, along with an evaluation function $Eval(\cdot)$. The optimization goal of AgentNet is to maximize the evaluated score by $Eval(\cdot)$ for the solution output by AgentNet, specifically optimizing \mathcal{A} and \mathcal{E} , as the following formula:

$$\mathcal{G}^* = (\mathcal{A}^*, \mathcal{E}^*) = \arg \max_{\mathcal{A}, \mathcal{E}} Eval(\mathcal{G}, \mathbf{T}). \quad (1)$$

The innovation of AgentNet emerges from the synergistic integration of three key mechanisms: (1) AgentNet realized a fully decentralized network architecture by distributing decision-making authority

across all agents, leading to a high failure tolerance and privacy-preserving MAS. (2) By using a dynamic task allocation mechanism, AgentNet can optimize workload distribution based on agent capabilities and the current system state flexibly. (3) The adaptive learning in AgentNet can achieve continuous specialization of agents, making the whole MAS more scalable and adaptive. Therefore, we can create a self-organizing system capable of handling complex tasks while preserving privacy and adapting to changing environments.

Algorithm 1 AgentNet System

Require: Task set $T = \{t_1, t_2, \dots, t_M\}$
Ensure: Optimized network $\mathcal{G}^* = (\mathcal{A}^*, \mathcal{E}^*)$

- 1: Initialize $\mathcal{A} = \{a_1, a_2, \dots, a_n\}$, $\mathcal{E}, \mathcal{C} = \{c_1, c_2, \dots, c_n\}$, w_0, M_i^{rou} and $M_i^{exe} \forall a_i \in \mathcal{A}$
- 2: **for** each task $t_{m+1} \in T$ **do**
- 3: // 1. Task allocation and processing
- 4: $c_{t_{m+1}} \leftarrow \Phi(o_{t_{m+1}})$, $a_{curr} \leftarrow \arg \max_{a_i \in \mathcal{A}_m} \text{sim}(c_{t_{m+1}}, c_i^m)$
- 5: $task_state \leftarrow (o_{t_{m+1}}, \emptyset, p_{t_{m+1}})$, $visited \leftarrow \emptyset$, $finished \leftarrow false$
- 6: **while** not $finished$ and $a_{curr} \notin visited$ **do**
- 7: $visited \leftarrow visited \cup \{a_{curr}\}$
- 8: $fragments^{rou} \leftarrow \text{Select}(M_{curr}^{rou}, t_{m+1}, k)$
- 9: $action \leftarrow \mathcal{F}_{act}(o_{t_{m+1}}, c_{t_{m+1}}, \mathcal{F}_{reason}(o_{t_{m+1}}, c_{t_{m+1}}, fragments^{rou}), fragments^{rou})$
- 10: **if** $action = \mathcal{O}_{fwd}$ **then**
- 11: $a_{curr} \leftarrow \arg \max_{a_k \in \mathcal{A}_m \setminus \{a_{curr}\}} \text{sim}(c_{t_{m+1}}, c_k^m)$
- 12: **else if** $action = \mathcal{O}_{split}$ **then**
- 13: $subtasks \leftarrow \text{DecomposeTask}(t_{m+1})$
- 14: $task_state.context \leftarrow task_state.context \oplus \text{ProcessSubtasks}(subtasks, a_{curr}, \mathcal{A}_m)$
- 15: $finished \leftarrow \text{AllSubtasksCompleted}(subtasks)$
- 16: **else**
- 17: $task_state.context \leftarrow task_state.context \oplus \text{ExecuteTask}(a_{curr}, t_{m+1}, \text{Select}(M_{curr}^{exe}, t_{m+1}, k))$
- 18: $finished \leftarrow true$
- 19: **end if**
- 20: **end while**
- 21: // 2. Network update
- 22: **for** each interacting pair (a_i, a_j) **do**
- 23: $w_{m+1}(i, j) \leftarrow \alpha \cdot w_m(i, j) + (1 - \alpha) \cdot S(a_i^{m+1}, a_j^{m+1}, t_{m+1})$
- 24: **end for**
- 25: $\mathcal{E}_{m+1} \leftarrow \{(a_i^{m+1}, a_j^{m+1}) \mid w_{m+1}(i, j) > \theta_w\}$
- 26: // 3. Agent capability and memory update
- 27: **for** each participating agent a_i **do**
- 28: $c_i^{m+1} \leftarrow \beta \cdot c_i^m + (1 - \beta) \cdot \Delta c_i^{m+1}$, Update M_i^{rou} and M_i^{exe}
- 29: **end for**
- 30: **end for**
- 31: **return** $\mathcal{G}^* = (\mathcal{A}, \mathcal{E})$

3.2. Decentralized Network Topology

As illustrated in Figure 3, AgentNet employs a dual-role design. Each agent a_i is equipped with a router rou_i to facilitate routing decisions and an executor exe_i to execute specific tasks. We will introduce details of the router and executor in the following sections. In essence, the router within each agent endows AgentNet with a fully decentralized network structure because the routing decisions are made independently by each agent, without relying on a central authority or coordinator.

This contrasts with traditional LLM based Multi-Agent Systems that typically depend on a centralized controller to manage the coordination and allocation of tasks. Each agent in AgentNet autonomously determines how to route tasks to other agents based on its local knowledge and the task requirements, ensuring that decision-making is distributed across the network and that there is no single point of control, thus achieving full decentralization.

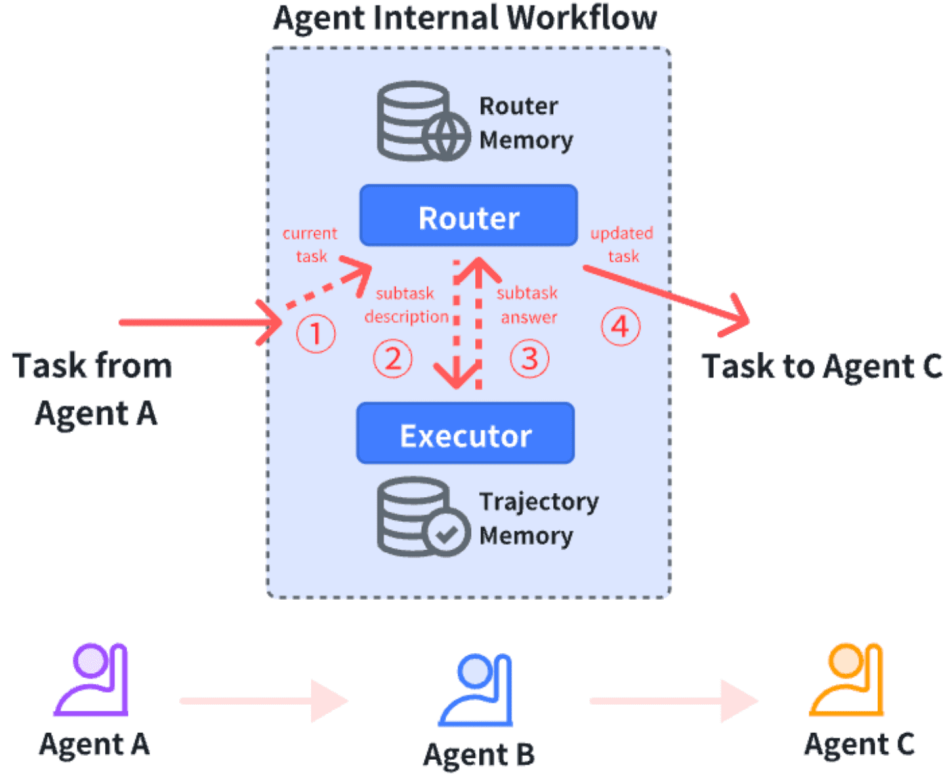


Figure 3. Dual-role agent architecture.

Mathematically, we represent the architecture of AgentNet as $\mathcal{G}_m = (\mathcal{A}_m, \mathcal{E}_m)$ when given the $m + 1$ -th task t_{m+1} after completing task t_m , where $\mathcal{A}_m = \{a_1^m, a_2^m, \dots, a_n^m\}$ represents the states of agents after task t_m and $\mathcal{E}_m \subseteq \mathcal{A}_m \times \mathcal{A}_m$ represents the set of directed edges between agents and each edge $e_{i,j}^m$ means a directed edge from a_i^m to a_j^m . A weight matrix w_m will be maintained throughout all the tasks before t_{m+1} to weight the connection between agents, namely $w_m(i, j)$. After completing t_{m+1} , w_{m+1} can be updated using the following formula from w_m :

$$w_{m+1}(i, j) = \alpha \cdot w_m(i, j) + (1 - \alpha) \cdot S(a_i^{m+1}, a_j^{m+1}, t_{m+1}), \quad (2)$$

where $\alpha \in [0, 1]$ is a decay factor that balances historical performance with recent interactions, and $S(a_i^{m+1}, a_j^{m+1}, t_{m+1})$ is a success metric for task t_{m+1} routed from agent a_i^{m+1} to a_j^{m+1} . This adaptive weighting mechanism ensures that the network continuously refines its structure based on operational experience.

Over tasks, the weight matrix w_m will evolve based on collaborative success, and the edges with a lower weight than a hyper-parameter threshold θ_w are periodically pruned:

$$\mathcal{E}_{m+1} = \{(a_i^{m+1}, a_j^{m+1}) \mid w_{m+1}(i, j) > \theta_w\}. \quad (3)$$

This pruning mechanism ensures that the network maintains efficient pathways while eliminating unproductive connections, optimizing both communication overhead and routing efficiency.

3.3. Adaptive Learning and Specialization

AgentNet's adaptive learning mechanism facilitates continuous improvement and specialization of agents based on their task experiences, without the need for explicit role assignment. This process enables agents to gradually develop expertise in specific domains, differentiating AgentNet from static multi-agent systems and allowing it to adapt to evolving requirements over time.

Agents in AgentNet follow the *ReAct* (Reasoning + Acting) framework Yao et al.^[32]; Zhou et al.^[6], which empowers agents to reason about a given query and its context before deciding appropriate actions for the executor modules. In addition to the given query and its context, the agent also retrieves relevant trajectory fragments from its memory modules to enhance reasoning and acting. The retrieval process is performed using a Retrieval-Augmented Generation (RAG) mechanism^{[18][19][6]}, which allows the agent to leverage past experiences to generate informed decisions and actions for new tasks.

In the AgentNet, Each agent $a_i \in \mathcal{A}$ maintains two memory modules M_i^{rou} and M_i^{exe} for its router module rou_i and exe_i , which store local trajectory fragments from prior tasks corresponding to the specific steps where a_i was actively involved instead of storing the whole task trajectories cooperated by all agents.

Formally, each entry in M_i^{rou} and M_i^{exe} is the local step fragment represented as: $f^r = (o^r, c^r, a^r)$, where r represents this entry belongs to rou_i (when $r = rou$) or exe_i (when $r = exe$). o^r denotes the observation, namely the query of the corresponding task, and c^r represents the context of the corresponding task so far (i.e., partial trajectory before this step), and a^r is the action or response of the

agent. These step fragments are collected from different tasks in which the agent has participated and serve as experiential knowledge for future reasoning.

When agent a_i^m receives a new task t_{m+1} to solve, it retrieves the k most relevant fragments from both memory modules. For each module type $r \in \{rou, exe\}$, the retrieval process is defined as:

$$\begin{aligned} \text{Select}(M_i^r, t_{m+1}, k) &= \{f_1^r, f_2^r, \dots, f_k^r\} \subset M_i^r, \text{ such that} \\ \forall f_j^r \in \text{Select}(M_i^r, t_{m+1}, k) \text{ and } \forall f \in M_i^r \setminus \text{Select}(M_i^r, t_{m+1}, k) : \\ \text{sim}(\text{embed}(o_{f_j^r}^r, c_{f_j^r}^r), \text{embed}(o_{t_{m+1}}^r, c_{t_{m+1}}^r)) &\geq \text{sim}(\text{embed}(o_f^r, c_f^r), \text{embed}(o_{t_{m+1}}^r, c_{t_{m+1}}^r)). \end{aligned} \quad (4)$$

Here, $\text{embed}(\cdot)$ is a semantic embedding function that projects the input context into a high-dimensional vector space, and the fragments with the highest relevance are retrieved to inform the agent's reasoning or action for both routing and execution processes.

Both the reasoning and acting processes are enhanced by the retrieval of historical task fragments, allowing the agent to make better decisions based on prior experiences. The reasoning function for each module type is modeled as:

$$\mathcal{R}_{a_i}(t_{m+1}, r) = \mathcal{F}_{\text{reason}}(o_{t_{m+1}}, c_{t_{m+1}}, \{f_j^r\}_{j=1}^k), \quad (5)$$

where $\mathcal{F}_{\text{reason}}$ represents the large language model that serves as the backbone of the LLM Agent, processing the inputs to generate reasoned decisions. The reasoning function \mathcal{R}_{a_i} takes the current observation and question $o_{t_{m+1}}$, the historical context $c_{t_{m+1}}$ representing the partial task trajectory and interactions up to the current point, and the retrieved fragments $\{f_j^r\}_{j=1}^k$ as input to generate the reasoning output. The fragments allow the agent to reason based on prior experiences that are most relevant to the current situation.

Once the reasoning process has been completed, the agent executes the chosen action. The action is informed by the reasoning output, which can be expressed as:

$$\mathcal{A}_{a_i}(t_{m+1}, r) = \mathcal{F}_{\text{act}}(o_{t_{m+1}}, c_{t_{m+1}}, \mathcal{R}_{a_i}(t_{m+1}, r), \{f_j^r\}_{j=1}^k), \quad (6)$$

where \mathcal{F}_{act} represents the large language model that serves as the backbone of the LLM Agent, translating reasoning into concrete operations. The $\mathcal{A}_{a_i}(t_{m+1}, r)$ function utilizes the reasoning output $\mathcal{R}_{a_i}(t_{m+1}, r)$ along with the retrieved memory fragments to determine the appropriate action. The specific action depends on the module type: for $r = rou$, the router module may produce actions such as forwarding the task to another agent or splitting it into subtasks; for $r = exe$, the executor module generates a single-step operation or response to directly address the final answer.

To manage memory effectively, each agent employs a dynamic memory management strategy. As the agent receives new tasks, it evaluates the trajectories stored in its memory modules and decides which to retain or prune. This evaluation is based on reasoning about the task context, historical usage patterns, and the relevance of each trajectory to future tasks. Several factors influence the decision-making process, including the frequency of use, recency of tasks, and the uniqueness of the trajectory. The agent assesses these factors through a prompt-based reasoning process that helps determine the utility of each trajectory. When a memory module reaches its capacity limit C_{\max} , the agent compares the new trajectory with the existing ones in that memory module and selects the least useful trajectory to remove, thus ensuring the memory pool remains focused on high-quality and relevant knowledge.

Through this adaptive and memory-driven learning process, agents in AgentNet continuously refine their expertise and specialize in the areas where they excel. This specialization occurs naturally over time, allowing the system to self-organize and adapt to the demands of a wide variety of tasks.

3.4. Dynamic Task Allocation

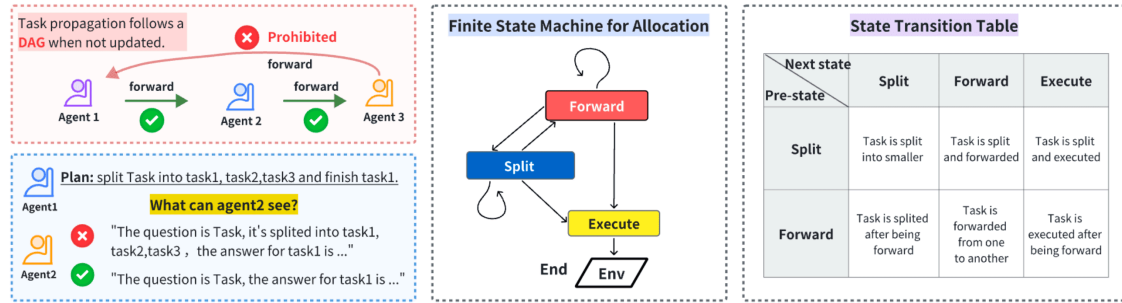


Figure 4. Details of Dynamic Task Allocation.

The dynamic task allocation mechanism in AgentNet enables efficient distribution of tasks without centralized coordination, creating a responsive system that optimizes both performance and load balancing. This decentralized approach to task routing represents a significant advancement over static assignment strategies employed in traditional multi-agent frameworks.

Each task $t \in T$ is formally represented as a tuple $t = (o_t, c_t, p_t)$, where o_t contains the task description in natural language, c_t is a vector of capability requirements, and p_t denotes the priority level. To efficiently process a new task t_{m+1} after completing task t_m , AgentNet employs a sophisticated mechanism to select the most suitable initial agent.

Agent capability representation and matching form the foundation of task allocation. Each agent a_i^m , after completing task t_m , possesses a capability vector c_i^m that is dynamically updated through task performance during system operation. In the initial allocation phase, the system selects an entry agent for t_{m+1} using the following formula:

$$a_{initial} = \underset{a_i \in \mathcal{A}_m}{\operatorname{argmax}} \{ \operatorname{sim}(c_{t_{m+1}}, c_i^m) \}, \quad (7)$$

where $c_{t_{m+1}} = \Phi(o_{t_{m+1}})$ represents the capability requirements of task t_{m+1} , c_i^m denotes the capability vector of agent a_i , and $\operatorname{sim}(\cdot, \cdot)$ is a similarity function measuring the match between task requirements and agent capabilities. The capability requirements are determined through different methodologies depending on task complexity:

$$c_{t_{m+1}} = \begin{cases} \Phi_{atomic}(t_{m+1}), & \text{for atomic tasks} \\ \Phi_{compound}(t_{m+1}), & \text{for compound tasks.} \end{cases} \quad (8)$$

For atomic tasks, the system employs function Φ_{atomic} that maps task properties to capability requirements based on predefined heuristics. For compound tasks, function $\Phi_{compound}$ leverages an instruction set containing carefully crafted prompts that guide the large language model in analyzing task descriptions and inferring the required capability vectors. The system subsequently ranks all agents according to capability matching scores and selects the highest-scoring agent as the initial agent.

Once a task is assigned to the initial agent, this agent determines how to process the task based on the reasoning results from its router module rou_i . As illustrated in Figure 4, the agent can perform three operations:

1. **Forward** (\mathcal{O}_{fwd}): Transfer the task unchanged to another more suitable agent, maintaining the task's original state and preserving the Directed Acyclic Graph (DAG) property of the routing path. Forwarding decisions are based on analyzing the gap between the current agent's capabilities and the task requirements, as well as evaluating the capability vectors of other agents in the network.
2. **Split** (\mathcal{O}_{split}): Decompose the task into subtasks, execute portions matching the agent's expertise, and route the remaining subtasks to an appropriate agents. Subtask routing follows this formula:

$$a_{next} = \underset{a_k \in \mathcal{A}_m \setminus \{a_i\}}{\operatorname{argmax}} \{ \operatorname{sim}(\Phi(o_{t_{m+1}}), c_k^m) \}, \quad (9)$$

where $\Phi(o_{t_{m+1}})$ represents the capability requirements derived from the observation of subtask j , determined through the current agent's task decomposition reasoning, and $\mathcal{A}_m \setminus \{a_i\}$ denotes the set of all agents excluding the current one.

3. **Execute** (\mathcal{O}_{exec}): Complete the entire task without further delegation.

A key design feature in the system is that when an agent chooses to split a task, it only forwards the results of the subtasks it has completed, and not the reasoning behind the decomposition. This prevents the transfer of unnecessary information and ensures that task decomposition errors made by one agent do not propagate to other agents in the network.

The agent capability vector c_i^m is updated based on task execution history and success rates, using the following formula:

$$c_i^{m+1} = \beta \cdot c_i^m + (1 - \beta) \cdot \Delta c_i^{m+1}, \quad (10)$$

where $\beta \in [0, 1]$ is a decay factor balancing historical capabilities with newly acquired ones, and Δc_i^{m+1} represents the new capability contribution demonstrated by the agent in task t_{m+1} , calculated by analyzing the types of operations successfully executed by the agent and the quality of results.

Furthermore, the task's state is updated only when an agent completes a part of the task (whether by executing or splitting it). When the agent completes a subtask, it updates the context and forwards it to the next agent:

$$\text{context}_{\text{updated}} = \text{context}_{\text{original}} \oplus \text{result}(a_j, t_i). \quad (11)$$

While the task is only being forwarded from one agent to another, its state remains unchanged, preserving the Directed Acyclic Graph (DAG) structure of the task routing path. This ensures that the task's progression avoids being trapped in an infinite loop during task forwarding, maintaining a consistent and effective routing process across different agents.

Through this dynamic task allocation mechanism, AgentNet can adaptively optimize task flow based on task characteristics and changes in agent capabilities, achieving improved overall system performance and efficient resource utilization.

4. Experiment

4.1. Experimental Setup

Tasks and Benchmarks

We evaluate methods using several benchmarks across three task categories, along with custom constructed training and test sets for each benchmark:

- **Mathematics:** This task involves mathematical problem and is evaluated using MATH Hendrycks et al.^[33], which includes problems with 7 different types. The training set consists of 100 examples per type (total of 700 problems), while the test set consists of 20 examples per type (total of 140 problems).
- **Logical Question Answering:** This task tests reasoning and logical question answering abilities using the BBH (Big-Bench Hard) benchmark Suzgun et al.^[34]. The training set follows the MorphAgent setup, selecting 627 examples from 20 tasks. For testing, each task has 5 examples of varying difficulty, totaling 100 test problems.
- **Function-Calling:** This benchmark evaluates the agent’s ability to perform tool-augmented task planning and API usage, based on the API-Bank dataset Li et al.^[35]. We construct a training set of 100 tasks and a test set of 100 tasks, randomly sampled from the full API-Bank corpus. Since the original dataset does not include category labels, we annotate each task using GPT-4o-mini to assign one of the seven task types: health, account, schedule, information, housework, finance, and others. Each task is further categorized into one of three difficulty levels, determined by prompt complexity and required toolchain length.

Baselines

We compare AgentNet with two categories of baselines: single-agent and multi-agent frameworks:

- **Single-agent frameworks:** These methods involve a single agent solving tasks independently without collaboration or coordination with other agents.
 - **Direct:** A baseline approach where the LLM directly generates outputs.
 - **React:** A prompting technique that elicits step-by-step reasoning from language models Wei et al.^[36]; Yao et al.^[32].
 - **Synapse:** A trajectory-as-exemplar prompting method, which prompts the LLM with complete trajectories of the abstracted states and actions to improve multi-step decision-making. Zheng et al.^[37]
 - **Self-Consistency:** A decoding strategy that samples multiple reasoning paths and selects the most consistent answer through majority voting, enhancing reliability Wang et al.^[38].
 - **Self-Refinement:** An iterative approach where models critically evaluate and improve their own solutions over multiple passes, progressively enhancing solution quality Madaan et al.^[39].
- **Multi-agent frameworks:** These methods involve multiple agents working collaboratively to solve tasks, each contributing to different aspects of the task-solving process.

- **MorphAgent:** A framework featuring self- evolving agent profiles that dynamically optimize individual expertise in the profile through three metrics Lu et al.^[27].
- **MetaGPT:** A software development framework where specialized agents (like product manager, architect, engineer) collaborate in a waterfall workflow to complete complex engineering tasks Hong et al.^[10].
- **AFLOW:** A framework that optimizes agent workflows using Monte Carlo Tree Search over code- represented workflows with execution feedback Zhang et al.^[30].
- **GPTSwarm:** A framework modeling agents as computational graphs with automatic optimization of both prompts and agent collaboration patterns Zhuge et al.^[24].

Parameter Configuration

In our implementation, we configure the LLM API with a temperature of 0.0, a maximum token limit of 2048, and a top-p value of 1.0, ensuring consistent results throughout our experiments and enabling reliable comparisons and analysis. For the memory pool experiment, we utilize the "BAAI/bge-large-en-v1.5" model to compute the similarity between task queries and database trajectories.

4.2. Main Results

Backbone	Category	Method	MATH(Acc/%)	BBH(Acc/%)	API-Bank(Acc/%)
DeepSeek-V3	Single Agent	Direct	47.86	69.00	26.00
		React	77.14	88.00	29.00
		Synapse	89.28	92.00	28.00
		Self-Consistency	88.00	85.00	29.00
		Self-Refinement	87.14	84.00	25.00
	Multi-Agent	MorphAgent	39.29	56.00	16.00
		MetaGPT	92.14	64.00	22.00
		AFLOW	<u>91.67</u>	88.00	28.00
		GPTSwarm	72.14	90.00	21.00
		AgentNet	92.86	94.00	30.00
GPT-4o-mini	Single Agent	Direct	31.43	59.00	15.00
		React	55.71	80.00	24.00
		Synapse	77.14	79.00	22.00
		Self-Consistency	54.28	85.00	22.00
		Self-Refinement	68.57	81.00	23.00
	Multi-Agent	MorphAgent	80.71	56.00	16.00
		MetaGPT	73.57	53.00	19.00
		AFLOW	85.00	75.00	21.00
		GPTSwarm	85.00	86.00	13.00
		AgentNet	85.00	86.00	29.00
Qwen-turbo	Single Agent	Direct	37.85	57.00	27.00
		React	53.57	69.00	23.00
		Synapse	67.14	68.00	24.00

Backbone	Category	Method	MATH(Acc/%)	BBH(Acc/%)	API-Bank(Acc/%)
		Self-Consistency	64.28	70.00	28.00
		Self-Refinement	76.43	74.00	23.00
	Multi-Agent	MorghAgent	16.43	56.00	9.00
		MetaGPT	63.57	51.00	20.00
		AFLOW	82.14	57.00	22.00
		GPTSwarm	79.29	<u>75.00</u>	30.00
		AgentNet	81.43	92.00	32.00

Table 1. Performance comparison of different methods across various tasks. In all multi-agent methods, we set 3 agents for each method to ensure a fair comparison. The best result is in bold, while the second is underlined.

Table 1 summarizes performance across Math, Logical QA tasks and API-Calling tasks. For Math, Logical QAs and API-Calling tasks, accuracy is reported. Compared to single-agent methods (e.g., Synapse, ReAct), AgentNet achieves competitive or superior performance across all tasks. While ReAct performs well on Math and Logic, its static prompting strategy limits generalization to more complex tasks. Against multi-agent baselines, AgentNet consistently outperforms centralized frameworks such as MetaGPT, which suffers from limited scalability—e.g., only 53.00% accuracy on Logical QA. AgentNet’s decentralized coordination and retrieval-augmented memory contribute to its robustness across domains, particularly in tasks requiring contextual understanding and adaptive role specialization.

4.3. Experiments on Heterogeneous Agents

To investigate the impact of agent diversity on performance, we designed a heterogeneity experiment across different settings on the BBH task. Agents were tested under four configurations: fully homogeneous (identical models and capabilities), LLM heterogeneity (different language models, same capabilities), skill heterogeneity (same model but varied capabilities), and a combination of both. This design allows us to isolate and analyze how model-level and capability-level diversity influence multi-agent collaboration.

Setting	Fully Homogeneous	Skill Hetero.	LLM Hetero.	Both Hetero.
3 Agents	0.86	0.84	0.81	0.81
5 Agents	0.79	0.86	0.85	0.85

Table 2. BBH Accuracy under Different Heterogeneity Settings (Acc%)

The results show that the impact of heterogeneity on multi-agent performance depends on team size. With 3 agents, the fully homogeneous setting performs best, while introducing either model or skill diversity reduces accuracy, suggesting uniform reasoning is more effective in small teams. However, with 5 agents, heterogeneous configurations outperform the homogeneous one, indicating that diversity enhances collaboration and complementary reasoning in larger teams. Overall, heterogeneity may introduce coordination overhead in small groups but offers clear benefits at larger scales.

4.4. Ablation Study

Router Effectiveness in AgentNet

To evaluate AgentNet’s decentralized router, experiments were conducted comparing AgentNet with ablation configurations: "Totally Random", "Random Operations", "Random Next Agent ID" and a centralized "Global Router". Each router manages external routing (selecting the next agent) and internal routing (deciding to forward, split, or execute).

Performance was tested on the BBH task (training: 627 problems, testing: 100 problems), with results in Figure 5. AgentNet outperforms randomized methods, achieving 82.14% accuracy during training and 86.00% during testing. Randomizing operations (forward/split/execute) affects task execution more directly, randomizing next agent ID primarily results in suboptimal task delegation but does not disrupt task completion as severely. These results underscore the critical role of effective routing and suggest that optimizing routing decisions can significantly enhance multi-agent system performance.




Figure 5. AgentNet's Router Performance on the BBH (Backbone: gpt-4o-mini)

Impact of Evolution Phase

Results in Table 3 clearly indicate that AgentNet significantly improves performance compared to the non-evolution baseline. On the MATH task, AgentNet achieves a score of 85.00 versus 77.86. For the function-calling task, performance improves notably from 23.00 to 32.00. On the BBH task, accuracy rises from 76% to 86%, demonstrating the impact of evolution phase for AgentNet.

3 Agents	MATH	API-Bank	BBH
	Acc(%)	Acc(%)	Acc(%)
w/o evolution	77.86	23.00	76.00
AgentNet	85.00	32.00	86.00

Table 3. Performance Comparison of AgentNet vs. Without evolution (Backbone: gpt-4o-mini)

These results confirm that AgentNet’s adaptive learning during the evolution phase effectively enhances agent specialization and task performance, demonstrating its essential role in the system’s optimization and overall efficiency.

4.5. Case Studies

This case study is presented to illustrate the differences between the two methods, based on results obtained using GPT-4o-mini on the BBH dataset. The left image shows the trajectory produced by the ReAct method, while the right image illustrates the trajectory generated by AgentNet. In the case of ReAct, the lack of collective reasoning results in an incorrect response after a single-step inference, highlighting the limitations of the method in handling tasks that require more complex reasoning. In contrast, AgentNet uses a multi-step workflow where agents without the necessary expertise are bypassed, while those with the relevant skills divide the task into smaller steps, leading to a more accurate final solution.

Question:

Which sentence has the correct adjective order?

Options:

A. old-fashioned circular green cardboard exercise computer

B. old-fashioned green exercise circular cardboard computer

ReAct

Reasoning: (B) was selected because it follows a common pattern in English, where opinion adjectives (e.g., "old-fashioned") precede color adjectives (e.g., "green") and shape adjectives (e.g., "circular").

Action: (B)

Ground Truth:(A)

Figure 6. ReAct Response with Reasoning

AgentNet (5 Agents)

Agent 0: **Decision:** Forward Next Agent: Agent 1

Agent 1: **Decision:** Forward Next Agent: Agent 2

Agent 2: **Decision:** Forward Next Agent: Agent 3

Agent 3: **Decision:** Split

Executable: Adjective order analysis

Delegate: Final answer to another agent

Findings:

Correct order: Quantity, opinion, size, age, shape, color, origin, material, purpose Correct Option: (A)

Next Agent: Agent 4

Agent 4: **Decision:** Execute Action: Confirmed option (A)

Final Outcome: (A)

Ground Truth:(A)

Figure 7. AgentNet Task Breakdown

5. Analysis

5.1. Scalability and Robustness of the System

Based on the experimental results as illustrated in Figure 8, we observed that both training and testing performance improve slightly as the number of agents and executor pool limit increase. However, the improvements are incremental, with diminishing returns as the system scales up. Specifically, performance in the training phase increased from 80.38 for 3 agents and 30 executors to 81.18 for 9 agents and 40 executors. In the testing phase, performance fluctuated between 80 and 86, with the highest performance seen in configurations with 40 executors.




Figure 8. AgentNet Performance with Different Net Parameters. Experiments were conducted with routers without pool limit, where (A, B) represents A as the number of agents and B as the upper limit of the executor pool, with performance evaluated on the BBH.

These results suggest that AgentNet’s decentralized coordination model allows for gradual performance improvement as resources are added. This indicates that AgentNet can scale effectively, enhancing performance without the dramatic bottlenecks commonly seen in centralized systems. Additionally, while performance benefits from increased resources, the marginal gains suggest there may be an optimal point where resource allocation reaches its most efficient balance. The experiment demonstrates that AgentNet’s design, which dynamically adjusts agent connections and executor pool sizes, effectively supports scalable and adaptable multi-agent systems with a high degree of fault tolerance.

5.2. Autonomous Specialization of Agents

Based on the observed results in Figure 9, the experiments demonstrate that AgentNet’s multi-agent system can naturally specialize agents in a decentralized environment. With varying numbers of agents and a fixed executor pool of 40 pieces, the ability scores across different tasks such as reasoning, language, knowledge, and sequence showed significant variation. As the number of agents increased, specialization became more evident, particularly in complex tasks, with certain agents excelling in specific areas while others focused on different abilities. This highlights AgentNet’s capacity to dynamically refine agent expertise and optimize performance in a decentralized, task-driven system.




Figure 9. Autonomous Specialization under Different Agent Sets. The upper limit of the executor pool is fixed at 40, regardless of the number of agents. The experiment was run on the BBH, and the images show the final ability scores of the agents after training.

Notably, with 3 agents, the abilities were more evenly distributed across the agents, but as the number of agents increased, specialization became more evident, especially in tasks like knowledge and sequence, where specific agents showed notable proficiency. In the (5, 40) configuration, the specialization became clearer, with some agents excelling in certain abilities, while others lagged behind in different tasks. As we scaled up to 7 and 9 agents, the system displayed even greater specialization, especially in complex tasks, demonstrating AgentNet’s ability to allow for dynamic expertise refinement. This confirms that AgentNet supports the hypothesis of autonomous specialization within a decentralized, task-driven multi-agent system, where agents evolve to optimize task performance independently, without the need for a central controller.

5.3. Evolution of Agents Networks

The evolution of the agent network in our experiment is illustrated in Figure 10, which demonstrates the transition of a multi-agent system composed of 5 agents running on the BBH (627 pieces) benchmark. The figure captures the network at three key stages: the initial state, an intermediate state, and the final evolved state.

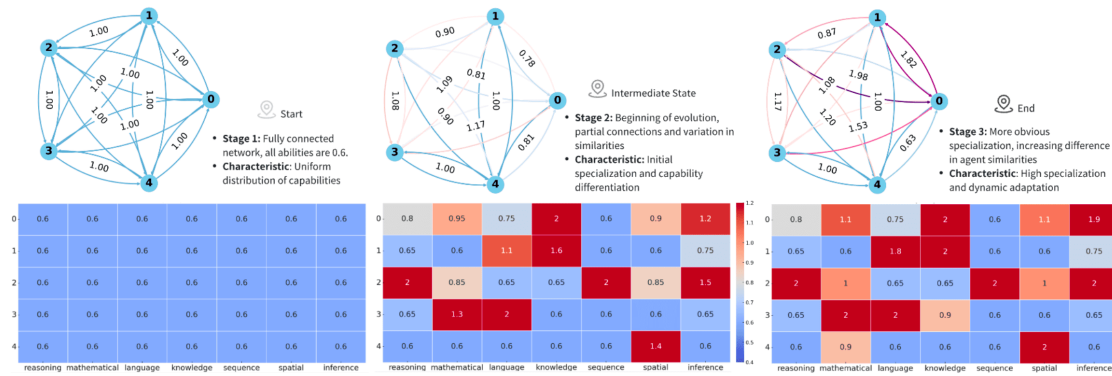


Figure 10. Evolution Example of Agents Networks.

In the initial state, the network is fully connected with uniform connection values of 1.00, indicating equal capabilities among all agents. At this stage, there is no specialization, and all agents are equally equipped to handle tasks.

As the network evolves, agents begin to specialize, and connection values vary, reflecting the strength of collaboration. Stronger connections indicate tighter cooperation, while weaker ones suggest less interaction. This evolution shows how agents naturally adapt and form more efficient collaboration patterns.

By the final stage, the network exhibits clear specialization, with agents taking on distinct roles. The connection values further emphasize the growing cooperation between specialized agents, improving task performance. This progression demonstrates the effectiveness of decentralized coordination, where evolving collaboration enhances task allocation, scalability, and fault tolerance.

6. Limitations and Future Work

Despite AgentNet implementing a fully distributed, adaptive learning multi-agent system (MAS) with dynamic task allocation, several important limitations remain that require further exploration in future work.

One key challenge is how to improve task performance in heterogeneous agent environments. In real-world applications, agents often vary significantly in terms of model capabilities, workflow structures, tools, and available data. The impact of such heterogeneity on AgentNet's performance, especially in terms of task coordination and resource allocation, remains an open question. Understanding how to

adapt the system to handle such variations efficiently will be crucial for its scalability and effectiveness in complex environments.

Secondly, the decision-making process of the router within each agent, particularly in relation to exploration and discovery, requires more in-depth study. Currently, the router selects agents from a relatively small pool of predefined candidates. However, in larger-scale systems involving hundreds or potentially thousands of agents, the challenge of accurately identifying the most suitable agent for task delegation becomes significantly more complex. This problem is further compounded in heterogeneous settings, where different agents may possess distinct strengths and weaknesses. To address this, future research could focus on developing more sophisticated routing mechanisms that can autonomously identify and delegate tasks to the most appropriate agents, even in large and diverse agent pools.

Additionally, a promising direction for future work involves designing incentives that encourage the router to explore agents beyond the predefined candidate set. By enabling AgentNet to dynamically discover new agents or specialized capabilities, such an approach would enhance its adaptability and scalability, ultimately improving the system's overall performance and autonomy.

7. Conclusion

In conclusion, AgentNet provides an effective approach to addressing the limitations of traditional centralized multi-agent systems. With its decentralized architecture, dynamic task allocation, and adaptive learning mechanisms, AgentNet improves scalability, fault tolerance, and task efficiency in collaborative environments. Its privacy-preserving features further ensure secure cooperation across organizations. Our experimental results highlight the advantages of this approach, demonstrating improvements in task efficiency, adaptability, and specialization. AgentNet offers a practical framework for developing more flexible and secure multi-agent systems in dynamic, real-world settings.

Statements and Declarations

Funding

This research was supported by Bytedance through a sponsored project that facilitated the execution and completion of this work.

Ethics Statement

This study did not involve human participants, animal subjects, or the use of personal data. All datasets and benchmarks employed are publicly available and used in accordance with their respective licenses. Therefore, no ethics approval was required.

Acknowledgements

The authors gratefully acknowledge the generous support provided by Bytedance, which has been instrumental in enabling the successful execution of this research. Their sponsorship significantly contributed to the advancement and completion of this work.

References

1. [△]OpenAI, Achiam J, Adler S, Agarwal S, et al. (2024). "GPT-4 Technical Report." arXiv. [2303.08774](https://arxiv.org/abs/2303.08774).
2. [△]Touvron H, Martin L, Stone K, et al. (2023). "Llama 2: Open Foundation and Fine-Tuned Chat Models." arXiv. [2307.09288](https://arxiv.org/abs/2307.09288).
3. [△]Yang Y, Huang B, Qi S, Feng C, Hu H, Zhu Y, Hu J, Zhao H, He Z, Liu X, Wang Z, Qiu L, Cao X, Cai X, Yu Y, Zhang W (2025). "Who's the MVP? A Game-Theoretic Evaluation Benchmark for Modular Attribution in LLM Agents." arXiv. [2502.00510](https://arxiv.org/abs/2502.00510).
4. [△]Gottweis J, Weng W, Daryin A, Tu T, Palepu A, Sirkovic P, Myaskovsky A, Weissenberger F, Rong K, Tanno R, Saab K, Popovici D, Blum J, Zhang F, Chou K, Hassidim A, Gokturk B, Vahdat A, Kohli P, Matias Y, Carroll A, Kulkarni K, Tomasev N, Guan Y, Dhillon V, Vaishnav ED, Lee B, Costa TRD, Penadés JR, Peltz G, Xu Y, Pawlosky A, Karthikesalingam A, Natarajan V (2025). "Towards an AI Co-Scientist." arXiv. [2502.18864](https://arxiv.org/abs/2502.18864).
5. [△]Putta P, Mills E, Garg N, Motwani S, Finn C, Garg D, Rafailov R (2024). "Agent Q: Advanced Reasoning and Learning for Autonomous AI Agents." arXiv. [2408.07199](https://arxiv.org/abs/2408.07199).
6. ^{a, b, c, d}Zhou R, Yang Y, Wen M, Wen Y, Wang W, Xi C, Xu G, Yu Y, Zhang W (2024). "TRAD: Enhancing LLM Agents With Step-Wise Thought Retrieval and Aligned Decision." arXiv. [2403.06221](https://arxiv.org/abs/2403.06221).
7. [△]Guo T, Chen X, Wang Y, Chang R, Pei S, Chawla N, Wiest O, Zhang X (2024). "Large Language Model Based Multi-Agents: A Survey of Progress and Challenges." In International Joint Conference on Artificial Intelligence. S2CID [267412980](https://doi.org/10.26434/chemrxiv-2024-26741).
8. [△]Sun C, Huang S, Pompili D (2024). "LLM-Based Multi-Agent Reinforcement Learning: Current and Future Directions." arXiv. [abs/2405.11106](https://arxiv.org/abs/2405.11106). S2CID [269921354](https://doi.org/10.26434/chemrxiv-2024-26992).

9. ^a ^b Yang Y, Peng Q, Wang J, Wen Y, Zhang W (2024). "LLM-Based Multi-Agent Systems: Techniques and Business Perspectives." arXiv. [2411.14033](#).
10. ^a ^b Hong S, Zhuge M, Chen J, Zheng X, Cheng Y, Zhang C, Wang J, Wang Z, Yau SKS, Lin Z, Zhou L, Ran C, Xiao L, Wu C, Schmidhuber J (2024). "MetaGPT: Meta Programming for a Multi-Agent Collaborative Framework." arXiv. [2308.00352](#).
11. ^a ^b Wu Q, Bansal G, Zhang J, Wu Y, Li B, Zhu E, Jiang L, Zhang X, Zhang S, Liu J, Awadallah AH, White RW, Burger D, Wang C (2023). "AutoGen: Enabling Next-Gen LLM Applications Via Multi-Agent Conversation." arXiv. [2308.08155](#).
12. ^a ^b Wang Q, Wang T, Li Q, Liang J, He B (2024). "MegaAgent: A Practical Framework for Autonomous Cooperation in Large-Scale LLM Agent Systems." arXiv preprint arXiv:2408.09955.
13. ^Δ Ye R, Tang S, Ge R, Du Y, Yin Z, Chen S, Shao J (2025). "MAS-GPT: Training LLMs to Build LLM-Based Multi-Agent Systems." arXiv. [2503.03686](#).
14. ^Δ Chen G, Dong S, Shu Y, Zhang G, Jaward S, Börje K, Fu J, Shi Y (2023). "AutoAgents: The Automatic Agents Generation Framework." arXiv preprint.
15. ^Δ Shi Z, Wan G, Huang W, Zhang G, Shao J, Ye M, Yang C (2025). "Privacy-Enhancing Paradigms Within Federated Multi-Agent Systems." arXiv. [2503.08175](#).
16. ^Δ Kahn AB (1962). "Topological Sorting of Large Networks." Commun ACM. 5(11):558–562.
17. ^Δ Ahuja RK, Magnanti TL, Orlin JB, et al. (1993). Network Flows: Theory, Algorithms, and Applications. Vol. 1. Englewood Cliffs, NJ: Prentice Hall.
18. ^a ^b Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, Küttler H, Lewis M, Yih W, Rocktäschel T, et al. (2020). "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." Adv Neural Inf Process Syst. 33:9459–9474.
19. ^a ^b Gao Y, Xiong Y, Gao X, Jia K, Pan J, Bi Y, Dai Y, Sun J, Wang H, Wang H (2023). "Retrieval-Augmented Generation for Large Language Models: A Survey." arXiv preprint arXiv:2312.10997.
20. ^Δ Gao D, Li Z, Pan X, Kuang W, Ma Z, Qian B, Wei F, Zhang W, Xie Y, Chen D, Yao L, Peng H, Zhang Z, Zhu L, Cheng C, Shi H, Li Y, Ding B, Zhou J (2024). "AgentScope: A Flexible Yet Robust Multi-Agent Platform." arXiv. [2402.14034](#).
21. ^Δ Fernando C, Banarse D, Michalewski H, Osindero S, Rocktäschel T (2023). "Promptbreeder: Self-Referential Self-Improvement Via Prompt Evolution." arXiv preprint arXiv:2309.16797.
22. ^Δ Khattab O, Singhvi A, Maheshwari P, Zhang Z, Santhanam K, Vardhamanan S, Haq S, Sharma A, Joshi TT, Moazam H, et al. (2023). "DSPy: Compiling Declarative Language Model Calls Into Self-Improving Pipeline

- s." arXiv preprint arXiv:2310.03714.
23. [△]Zhang G, Yue Y, Li Z, Yun S, Wan G, Wang K, Cheng D, Yu JX, Chen T (2024). "Cut the Crap: An Economical Communication Pipeline for LLM-Based Multi-Agent Systems."
 24. [△]Zhuge M, Wang W, Kirsch L, Faccio F, Khizbullin D, Schmidhuber J (2024). "GPTSwarm: Language Agents as Optimizable Graphs." In International Conference on Machine Learning.
 25. [△]Liu Z, Zhang Y, Li P, Liu Y, Yang D (2024). "A Dynamic LLM-Powered Agent Network for Task-Oriented Agent Collaboration." arXiv. [2310.02170](#).
 26. [△]Zhang G, Yue Y, Sun X, Wan G, Yu M, Fang J, Wang K, Chen T, Cheng D (2024). "G-Designer: Architecting Multi-Agent Communication Topologies Via Graph Neural Networks." arXiv preprint arXiv:2410.11782.
 27. [△]Lu S, Shao J, Luo B, Lin T (2024). "MorphAgent: Empowering Agents Through Self-Evolving Profiles and Decentralized Collaboration." arXiv preprint arXiv:2410.15048.
 28. [△]Chen W, Su Y, Zuo J, Yang C, Yuan C, Qian C, Chan C, Qin Y, Lu Y, Xie R, et al. (2023). "Agentverse: Facilitating Multi-Agent Collaboration and Exploring Emergent Behaviors in Agents." arXiv preprint arXiv:2308.10848 2 (4), pp. 6.
 29. [△]Shang Y, Li Y, Zhao K, Ma L, Liu J, Xu F, Li Y (2024). "AgentSquare: Automatic LLM Agent Search in Modular Design Space." arXiv preprint arXiv:2410.06153.
 30. [△]Zhang J, Xiang J, Yu Z, Teng F, Chen X, Chen J, Zhuge M, Cheng X, Hong S, Wang J, et al. (2024). "Aflow: Automating Agentic Workflow Generation." arXiv preprint arXiv:2410.10762.
 31. [△]Hu Y, Cai Y, Du Y, Zhu X, Liu X, Yu Z, Hou Y, Tang S, Chen S (2024). "Self-Evolving Multi-Agent Collaboration Networks for Software Development." arXiv preprint arXiv:2410.16946.
 32. [△]Yao S, Zhao J, Yu D, Du N, Shafran I, Narasimhan K, Cao Y (2023). "ReAct: Synergizing Reasoning and Acting in Language Models." arXiv. [2210.03629](#).
 33. [△]Hendrycks D, Burns C, Kadavath S, Arora A, Basart S, Tang E, Song D, Steinhardt J (2021). "Measuring Mathematical Problem Solving With the MATH Dataset." NeurIPS.
 34. [△]Suzgun M, Scales N, Schärli N, Gehrmann S, Tay Y, Chung HW, Chowdhery A, Le QV, Chi EH, Zhou D, Wei J (2022). "Challenging Big-Bench Tasks and Whether Chain-of-Thought Can Solve Them." arXiv. [2210.09261](#).
 35. [△]Li M, Zhao Y, Yu B, Song F, Li H, Yu H, Li Z, Huang F, Li Y (2023). "API-Bank: A Comprehensive Benchmark for Tool-Augmented LLMs." arXiv. [2304.08244](#).
 36. [△]Wei J, Wang X, Schuurmans D, Bosma M, Ichter B, Xia F, Chi E, Le Q, Zhou D (2023). "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models." arXiv. [2201.11903](#).

37. [△]Zheng L, Wang R, Wang X, An B (2024). "Synapse: Trajectory-as-Exemplar Prompting With Memory for Computer Control." arXiv. [2306.07863](https://arxiv.org/abs/2306.07863).
38. [△]Wang X, Wei J, Schuurmans D, Le Q, Chi E, Narang S, Chowdhery A, Zhou D (2023). "Self-Consistency Improves Chain of Thought Reasoning in Language Models." arXiv. [2203.11171](https://arxiv.org/abs/2203.11171).
39. [△]Madaan A, Tandon N, Gupta P, Hallinan S, Gao L, Wiegrefe S, Alon U, Dziri N, Prabhumoye S, Yang Y, Gupta S, Majumder BP, Hermann K, Welleck S, Yazdanbakhsh A, Clark P (2023). "Self-Refine: Iterative Refinement With Self-Feedback." arXiv. [2303.17651](https://arxiv.org/abs/2303.17651).

Declarations

Funding: This research was supported by Bytedance through a sponsored project that facilitated the execution and completion of this work.

Potential competing interests: No potential competing interests to declare.