

ResMAS: Resilience Optimization in LLM-based Multi-agent Systems

Anonymous submission

Abstract

Large Language Model-based Multi-Agent Systems (LLM-based MAS), where multiple LLM agents collaborate to solve complex tasks, have shown impressive performance in many areas. However, MAS are typically distributed across different devices or environments, making them vulnerable to perturbations such as agent failures. While existing works have studied the adversarial attacks and corresponding defense strategies, they mainly focus on reactively detecting and mitigating attacks after they occur rather than proactively designing inherently resilient systems. In this work, we study the resilience of LLM-based MAS under perturbations and find that both the communication topology and prompt design significantly influence system resilience. Motivated by these findings, we propose ResMAS: a two-stage framework for enhancing MAS resilience. First, we train a reward model to predict the MAS's resilience, based on which we train a topology generator to automatically design resilient topology for specific tasks through reinforcement learning. Second, we introduce a topology-aware prompt optimization method that refines each agent's prompt based on its connections and interactions with other agents. Extensive experiments across a range of tasks show that our approach substantially improves MAS resilience under various constraints. Moreover, our framework demonstrates strong generalization ability to new tasks and models, highlighting its potential for building resilient MASs. Our code is open-source¹ for reproducibility.

Introduction

Large language model-based multi-agent systems (LLM-based MAS; hereafter referred to as **MAS**) have exhibited impressive capabilities across a wide range of domains, including software development (Qian et al. 2024; Hong et al. 2023), mathematical problem solving (Li et al. 2023), and operations research (Xiao et al. 2023). In an MAS, multiple LLM agents collaborate to solve complex tasks through certain role-play and communication patterns, often achieving superior performance compared to single-agent approaches. Despite these advances, agents within an MAS are typically deployed across different devices or environments, making the MAS vulnerable to perturbations and attacks. As a result, the resilience, defined as the ability of a system to maintain its functionality under perturbations (Cohen et al. 2000;

Gao, Barzel, and Barabási 2016), has become an important issue for MAS.

In recent years, there have been extensive studies about resilience of MAS under adversarial attacks. Several studies propose attack strategies such as prompt injection (Yu et al. 2024; Huang et al. 2025; Tian et al. 2023) and knowledge editing (Ju et al. 2025, 2024), and find that harmful information can rapidly spread in MAS (Gu et al. 2024; Lee and Tiwari 2024). Other works have focused on detecting and defending against attacks using techniques like psychological testing (Zhang et al. 2024b) and graph-based anomaly detection (Wang et al. 2025b). However, existing studies mainly focus on the safety issue of MAS caused by adversarial attacks, while largely overlooking more common, non-malicious errors in MAS such as failures or miscommunication. Furthermore, prior works concentrated on detecting and mitigating attacks, with little focus on proactively designing resilient MAS structures that can better tolerate errors.

In this work, we study the resilience of MAS under random agent failures. We find that MAS exhibits significantly higher resilience than a single agent, and the resilience increases with the number of agents and communication links. Moreover, we find that both the topology and prompts of an MAS have a strong impact on its resilience. Motivated by these, we propose a two-stage framework to automatically optimize the topology and prompts to improve the resilience of MAS, named **ResMAS**. In the first stage, we fine-tune an open-source LLM serve as a topology generator capable of producing resilient MAS structures. Specifically, to avoid the high time cost of evaluating an MAS's resilience, we train a GNN-based reward model to predict the resilience of MAS for various tasks. This reward model is then used to fine-tune the LLM using the Group Relative Policy Optimization (GRPO) (Shao et al. 2024) algorithm, enabling it to automatically generate resilient MAS topologies. In the second stage, we propose a topology-aware prompt optimization method to optimize the system prompt for each agent in the MAS based on the generated topology. Specifically, we apply the MAS to a training set of the target task, and refine the system prompt for each agent based on the prompts of its neighbors as well as its interaction history with neighbors.

We conduct experiments on three benchmarks, including commonsense reasoning, math, and game. The results

¹https://anonymous.4open.science/r/LLM_MAS_resilience-891B

demonstrate that our framework can generate more resilient MAS compared to existing approaches for prompt and topology optimization. Moreover, we highlight the versatility of our framework by applying it to optimize the accuracy of MAS on certain tasks, where it achieves the Pareto-optimal performance compared with baselines. In addition, our framework shows strong generalization ability to unseen tasks such as code generation, and to new agents built upon different backbone models.

In summary, our contributions are as follows:

- We systematically study the resilience of MAS under random agent failures, revealing that both the topology and agent prompts have significant influence on system resilience.
- We propose ResMAS: a two-stage framework that automatically optimizes the topology and agent prompts of MAS to enhance resilience. Specifically, we fine-tune an LLM to generate robust topologies, and further refine the prompts in a topology-aware manner.
- Extensive experiments on various tasks show that our framework is able to generate MAS with superior resilience and demonstrates strong generalization ability across tasks and models.

Related Work

Resilience of MAS

The distributed architecture of MASs inherently increases their susceptibility to perturbations. Prior research has focused on adversarial attack and defense methods for MAS. Specifically, some studies propose attack methods such as prompt injection (Yu et al. 2024; Huang et al. 2025; Tian et al. 2023) and knowledge editing (Ju et al. 2025, 2024). Additionally, it has been demonstrated that harmful information can spread extremely fast in MAS, where a single jailbreak agent can undermine the whole MAS (Gu et al. 2024; Lee and Tiwari 2024). In response, some studies focus on detecting and defending against attacks through methods like psychological test (Zhang et al. 2024b) and graph-based anomaly detection (Wang et al. 2025b). However, existing studies mainly focus on the safety issue of MAS caused by adversarial attack, but overlook the resilience to more common, non-malicious agent errors. Furthermore, prior works emphasize detecting and mitigating attacks, with little focus on proactively designing resilient MAS structures that can tolerate errors by design.

Automated Design of MAS

The architecture of MAS, including the agents' roles and prompts as well as the communication patterns between agents, significantly affects the performance of MAS on tasks. Early studies often manually design MAS structures for specific tasks, usually based on human collaboration mechanisms (Du et al. 2024; Hong et al. 2023). Recently, there have been more studies on automated design and optimization of MAS. For instance, Agentverse lets LLM generate and adjust the agent composition based on the status

of the task (Chen et al. 2023). G-designer proposes to optimize the communication network of agents through a variational graph auto-encoder (Zhang et al. 2024a). GPTSwarm represents multi-agent systems as composite graphs and optimizes node-level prompts as well as edges between agents (Zhuge et al. 2024). However, all of the existing methods focus on optimizing the performance of MAS on certain tasks rather than resilience.

Preliminaries

Motivation

In this study, we use a widely used MAS framework (Du et al. 2023; Wang et al. 2025a; Yu et al. 2024), where multiple LLM agents discuss for several rounds to answer a question.

Definition 1 (LLM-based Multi-Agent System). *Specifically, the LLM agents can be modeled as a directed graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, \mathcal{P}\}$, where $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$ is the set of nodes, each node is an LLM agent, \mathcal{E} is the set of edges, and $\mathcal{P} = \{P_1, P_2, \dots, P_N\}$ represents the prompts for agents. We also refer to the graph \mathcal{G} as the communication topology of LLM agent groups. Given a query q , each agent $v_i \in \mathcal{V}$ independently generates an initial response $r_i^{(1)} = v_i(q)$. Then in round t ($t \geq 2$), each agent observes the previous answers of its in-neighbors, and updates its own answer:*

$$r_i^{(t+1)} = v_i(\{r_j^{(t)} | j \in \mathcal{N}_{in}(v_i)\}), \quad (1)$$

where $\mathcal{N}_{in}(v_i)$ denotes the in-neighboring nodes of v_i . After T rounds, the final answer is obtained by aggregating the responses of all agents

$$r^{(T)} = \text{Aggregate}(r_1^{(T)}, r_2^{(T)}, \dots, r_N^{(T)}). \quad (2)$$

In practice, different agents in an MAS are often deployed in a distributed way, i.e., they may be on different devices and environments. As a result, each agent is vulnerable to external perturbations. In this work, we define the perturbation as random agent failure, where in each round of the collaboration, each agent independently has a probability p to output a random response. We call the probability p the error rate.

We first study the behavior of MAS under perturbation. Specifically, we generate 30 random graphs with different agent numbers and edge numbers, and use them as the communication topology to construct 30 MASs. We then test the accuracy of these MASs on the Chess move validity tasks (details in the dataset section) under different error rates ($p = 0, 0.2, 0.4, 0.6, 0.8, 1$). We compare the results of MASs with single agent in Figure 1(a). It can be observed that as the error rate increases, both the performance of single agent and MAS drops. However, the performance of MASs remains much higher than that of a single agent under different error rates, indicating that MASs are more resilient to perturbations.

The resilience of a system is defined as its capability to maintain its functionality under perturbations (Cohen et al. 2000; Gao, Barzel, and Barabási 2016). Generally, the functionality of an MAS is reflected by its performance on certain tasks. Therefore, we quantitatively define the resilience of MAS as follows:

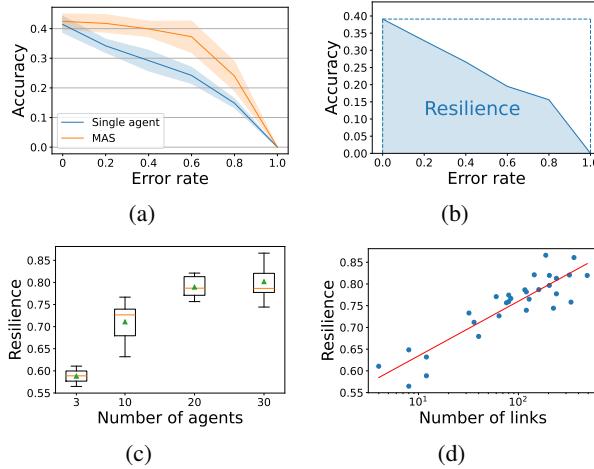


Figure 1: (a) Performance of single agent and MAS under perturbation. (b) Definition of resilience. (c)(d) Resilience generally increases with the number of agents and links in MAS.

Definition 2 (Resilience of MAS). Given an MAS \mathcal{G} and a target task consisting of K problems $\mathcal{T} = \{q_1, \dots, q_K\}$. Let $F(p)$ denote the accuracy (or other performance metrics) of the MAS on the task under error rate p , then the resilience of the MAS on task \mathcal{T} is defined as:

$$R(\mathcal{G}) = \frac{1}{F(0)} \int_0^1 F(p) dp, \quad (3)$$

where $F(0)$ represents the original functionality under no perturbation.

We visualize this definition in Figure 1(b), where resilience is defined as the proportion of area under the curve. In this study, we approximate the resilience by testing the accuracy of an MAS under error rates $p = 0, 0.2, 0.4, 0.6, 0.8, 1$ as follows:

$$R(\mathcal{G}) = \frac{1}{10F(0)} (F(0) + 2F(0.2) + 2F(0.4) + 2F(0.6) + 2F(0.8) + F(1)). \quad (4)$$

Based on the definition above, we further study the factors that affect MAS's resilience. Intuitively, more agents and more communication links between agents will bring more redundancy to the system, resulting in higher error tolerance. To verify this, we visualize the relationship between MAS resilience and the number of agents and links in Figure 1(c)(d), and find that more agents and links do lead to higher resilience, which indicates that we can improve MAS resilience simply by increasing the number of agents and edges. However, since more agents and edges will also bring higher cost, a natural question is whether we can optimize the MAS given the constraints of agent and edge numbers.

To answer this question, we conduct a case study on MASs with the same number of agents and edges, but with different graph structures and agent prompts. As shown in Figure 2, we find that when we refine the prompt for each agent (by calling GPT-4o one time), the resilience increases

significantly. Moreover, even under the same agent prompts, the hierarchical topology (right) demonstrates higher resilience than the centralized topology (left). These findings suggest both topology and prompt have a high impact on the resilience of MAS.

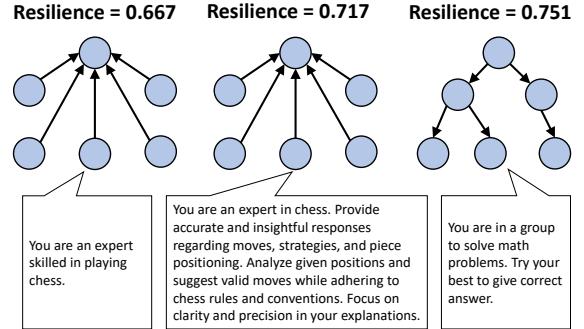


Figure 2: Effect of topology and prompt on the resilience of MAS.

Motivated by these, in this work, we aim to design more resilient MAS by optimizing the topology and prompts under the constraints of agent and edge numbers. We formally define the problem as follows:

Definition 3 (MAS Resilience Optimization Problem). Given the number of agents n and number of edges m , generate an MAS $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, \mathcal{P}\}$ with better resilience on a certain task, i.e.,

$$\begin{aligned} & \arg \max_{\mathcal{E}, \mathcal{P}} R(\mathcal{G}) \\ & \text{s.t. } |\mathcal{V}| = n, |\mathcal{E}| \leq m. \end{aligned} \quad (5)$$

Methods

Framework Overview

As mentioned above, the optimization targets of the problem include both the topology and agent prompts of MAS, which leads to a vast search space. To reduce the search space, we propose a two-stage framework that first designs the topology of MAS, and then optimizes the prompt with the topology fixed.

In the first stage, we focus on designing a more resilient MAS topology with identical prompts for all agents. Existing studies have proposed various methods to optimize the topology of MAS to improve its performance (Zhang et al. 2024a; Zhuge et al. 2024). However, these methods are mainly designed for optimizing the accuracy of MAS on certain tasks, and they face two challenges in our resilience optimization problem. First, all existing methods need to evaluate the performance of the proposed MAS on specific tasks, which is highly time-consuming. This issue is further exacerbated in resilience evaluation, as it requires assessing system performance under varying error rates. Second, existing methods typically conduct optimization independently for different tasks, limiting generalization and incurring additional computational cost. To deal with the first challenge, we train a task-aware reward model to predict the resilience

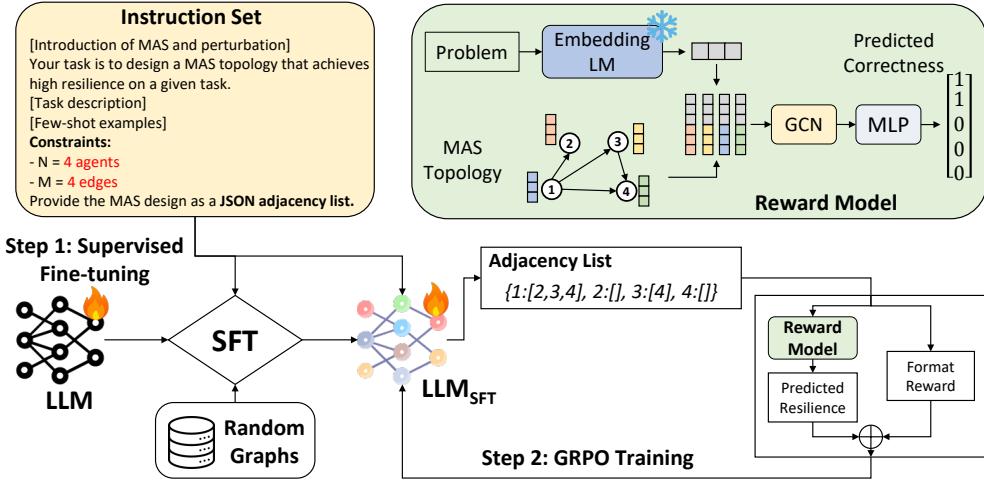


Figure 3: The topology optimization framework in ResMAS.

of MAS for different tasks. Moreover, inspired by the generalization ability of LLMs, we fine-tune an LLM with reinforcement learning guided by the reward model to serve as the topology generator. The generator can produce resilient topologies for diverse tasks and under various constraints specified via prompts. In the second stage, we fix the topology from the first stage and optimize the prompts for each agent. Since the prompt of an agent is tightly coupled with its position in the graph and its connections with other agents, we propose a topology-aware prompt optimization method, where each agent’s prompt is updated based on its connections and interactions with neighbors.

Topology Optimization

In the first stage, we aim to optimize the topology of MAS under given node and edge constraints. To make the optimization process generalizable to different tasks and constraints, we train a generator to directly generate the topology instead of searching in the topology space. Inspired by the outstanding instruction following and generalization ability of LLM, we fine-tune a small-sized LLM as the generator (Qwen2.5-7B-Instruct (Team 2024)). In this way, a single model can be used for different scenarios by designating the task and constraints in the prompt. The framework of topology optimization is shown in Figure 3.

Supervised Fine-tuning on Random Graphs We first construct the instruction set for training. In each instruction, we provide the introduction of MAS and perturbation, the target task, some few-shot examples, node and edge constraints, and ask the model to output an MAS topology with high resilience in adjacency list format. We vary the agent number from $\{10, 15, 20\}$, edge number from 10 to 80, and tasks including MATH, MMLU-Pro, and Chess (details in the dataset section), resulting in 2544 instructions in total. We present an example of instructions in Appendix.

In the experiments, we find that it is difficult for LLM to generate valid adjacency lists that adhere to the constraints. Therefore, we first conduct supervised fine-tuning (SFT) to make the model learn the basic required format. Since there

is no ground truth for the instructions, we generate random graphs as the labels. For example, if the constraint in an instruction is “10 agents, 10 edges”, we generate a random graph with 10 nodes and 10 edges as the label for training.

GRPO Training with Reward Model After teaching the model basic format through SFT, we further leverage reinforcement learning to enhance its ability to generate more resilient topologies. However, it is highly time-consuming to evaluate a generated topology since we need to test it under different error rates. For example, calculating the resilience of an MAS with 10 agents on the MATH dataset may take a few hours. To tackle this challenge, we train a reward model to predict the resilience of an MAS on a given task.

Motivated by the outstanding ability of Graph Neural Network to capture characteristics of graphs, we leverage GCN (Kipf 2016) as the backbone of reward model. Moreover, to enable prediction for different tasks, we use an embedding LM as the task encoder to obtain a low-dimensional embedding for the task, and concatenate it with node features (Zhang et al. 2024a), as shown in Figure 3. Here we initialize the node features with the in-degree and out-degree, and use Sentence-BERT (Reimers and Gurevych 2019) as the task encoder. After information propagation on the graph through two GCN layers, we apply a mean pooling operation to obtain a graph-level embedding, followed by an MLP to generate the final prediction. Here, we do not directly predict the resilience on the whole dataset, as it would be hard to collect a large amount of training data. Instead, we train the reward model to predict the correctness of each problem in the dataset, i.e., whether the MAS can correctly answer the problem under different error rates. Therefore, the output of the reward model is a 5-dimensional 0-1 vector, representing the correctness under $p=0, 0.2, 0.4, 0.6, 0.8$. For the training data, we use the random graphs from previous analysis (motivation section) to test on three datasets, resulting in 7936 [topology, task, correctness] tuples, which are split into train:test=8:2. Experiments show that the trained reward model can achieve an accuracy of 0.86 on the test set in predicting the correctness of an MAS on a problem. In

the inference stage, we calculate the predicted resilience \hat{R} on a dataset by aggregating the predicted correctness for all problems in the dataset with Equation 4.

After obtaining the reward model that can predict the resilience of any MAS topology on any task, we leverage the GRPO (Shao et al. 2024) algorithm to train the LLM to generate more resilient topologies. To ensure the generated topologies satisfy the constraints, we further introduce a format reward that punishes those that are not valid adjacency lists or do not meet the node and edge requirements. The overall reward function is defined as follows:

$$\text{Reward}(\mathcal{G}) = \begin{cases} -1, & \text{if wrong format or } |\mathcal{V}| \neq n \\ -\frac{|\mathcal{E}|-m}{m}, & \text{if } |\mathcal{E}| > m \\ \hat{R}(\mathcal{G}) + |\mathcal{E}|/m, & \text{if } n = N \text{ and } |\mathcal{E}| \leq m \end{cases} \quad (6)$$

In both the SFT and GRPO training, the model is trained using LoRA (Hu et al. 2022).

Topology-aware Prompt Optimization

Agents with different positions in the graph and different connections may play different roles in the MAS. Therefore, it is crucial to consider the topology when optimizing prompts. Previous prompt optimization methods for MAS often update the prompt for each agent sequentially (Zhuge et al. 2024). Moreover, they are primarily designed for improving the accuracy rather than resilience. In this work, we propose a topology-aware prompt optimization method to address such limitations, as shown in Figure 4.

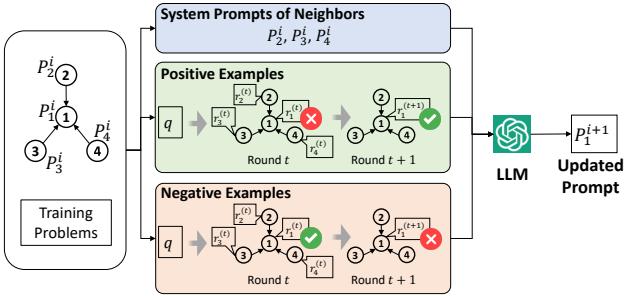


Figure 4: Topology-aware prompt optimization framework in ResMAS.

The intuition behind our method is that the resilience of MAS is largely affected by the interaction between agents. Simply improving the accuracy of each agent may not bring better resilience, as they may be easily affected by perturbed answers from other agents. Therefore, we hope that the agents can learn when to refer to responses from other agents and when to adhere to their own answers.

Instead of using a simple resilience metric as feedback, we make full use of every problem in the training set. Specifically, we run the MAS on the training set and calculate the correctness of each agent on each round of discussion for each problem. If an agent produces an incorrect response in one round but subsequently corrects its answer in the following round after receiving responses from its neighboring agents, this instance is treated as a positive example.

Similarly, we construct negative examples when an agent is misled by responses from neighbors and changes its correct answer to a wrong one. Based on the positive examples, negative examples, and the system prompts of neighbors, we leverage a stronger LLM to update the prompt for each agent. In this way, agents can learn from both the prompts of their neighbors as well as their interaction history with neighbors, enabling topology-aware prompt optimization. The prompt used for optimization is presented in Appendix.

Experiments

Experiment Settings

Datasets We evaluate the performance of our method on three datasets encompassing a wide range of tasks, including commonsense reasoning, mathematical reasoning, and game, which provides a diverse range of tasks to thoroughly assess our approach.

- **MMLU-Pro** (Wang et al. 2024): This dataset is more challenging version of MMLU (Hendrycks et al. 2021a) dataset containing multiple-choice questions with four to ten options. It contains problems from various disciplines and serves as a benchmark to test the general knowledge and commonsense reasoning ability of LLMs.
- **MATH** (Hendrycks et al. 2021b): It contains math problems to test the mathematical reasoning ability of LLMs. We choose the hardest level (level-5) in our experiments.
- **Chess** (Srivastava et al. 2023): It is a subset of the BIG-Bench Benchmark containing Chess move validity tasks, where the problems are to provide a valid move of a piece given the chess move history.

During the training of the reward model and topology generator, we only use data from the three datasets above. Moreover, we also evaluate the generalization ability of our method to unseen tasks on HumanEval (Chen et al. 2021) dataset, which measures the code generation ability of LLM.

Baselines We compare our method with three kinds of baselines that optimize topology, prompt, and both of them.

Topology optimization baselines.

- **G-Designer** (Zhang et al. 2024a): It models MAS as a graph, and trains a variational graph auto-encoder to generate the topology for MAS.

Prompt optimization baselines.

- **OPRO** (Yang et al. 2024): It uses LLM as optimizer to iteratively generate new prompts based on the performance of existing prompts on the training set.
- **TextGrad** (Yuksekgonul et al. 2025): It also leverages LLM as optimizer, and optimizes the prompt of the agent by backpropagating feedback, which is similar to training of neural networks. Following the original paper, we use the performance on the training set as the feedback.

Since the prompt optimization baselines are designed for a single agent, we use a random graph as the topology of MAS, and use the same optimized prompt for all agents.

Method	MATH						MMLU-Pro											
	10 Agents			15 Agents			20 Agents			10 Agents			15 Agents			20 Agents		
	10E	20E	30E	15E	30E	45E	20E	40E	60E	10E	20E	30E	15E	30E	45E	20E	40E	60E
G-Designer	0.490	0.573	0.531	0.531	0.540	0.480	0.546	0.572	0.558	0.456	0.498	0.502	0.459	0.422	0.454	0.441	0.458	0.383
OPRO	0.785	0.718	0.752	0.803	0.827	0.821	0.838	0.820	0.845	0.755	0.758	0.752	0.771	0.774	0.812	0.795	0.802	0.815
TextGrad	0.756	0.746	0.746	0.785	0.800	0.820	0.841	0.816	0.820	0.758	0.731	0.774	0.784	0.765	0.813	0.782	0.806	0.810
GPTSwarm	0.752	0.760	0.786	0.797	0.810	0.786	0.817	0.811	0.846	0.702	0.739	0.702	0.712	0.777	0.789	0.796	0.810	0.808
ResMAS	0.813	0.807	0.806	0.847	0.847	0.837	0.849	0.856	0.853	0.766	0.808	0.780	0.808	0.799	0.815	0.817	0.849	0.839

Table 1: Resilience comparison with baselines on MATH and MMLU-Pro datasets. "E" represents the number of edges. Best results are presented in bold, and the second best results are underlined.

Method	Chess								
	10 Agents			15 Agents			20 Agents		
	10E	20E	30E	15E	30E	45E	20E	40E	60E
G-Designer	0.707	0.731	<u>0.767</u>	0.773	0.717	0.665	0.763	0.776	0.747
OPRO	0.661	0.652	0.730	0.764	0.753	0.764	0.733	0.767	0.810
TextGrad	0.762	0.700	0.738	0.693	0.765	0.750	0.690	0.743	<u>0.811</u>
GPTSwarm	0.768	0.712	0.716	0.755	0.767	0.769	0.745	0.716	0.810
ResMAS	0.798	0.822	0.814	0.810	0.778	0.787	0.790	0.823	0.823

Table 2: Resilience comparison with baselines on Chess dataset. "E" represents the number of edges. Best results are presented in bold, and the second best results are underlined.

Topology and prompt optimization baselines.

- **GPTSwarm** (Zhuge et al. 2024): It optimizes the connection between agents through updating the probabilistic distribution of edges based on performance. Moreover, it sequentially optimizes the prompt for each agent in the MAS with LLM as optimizer.

We evaluate our method and baselines under different node and edge constraints, including 10-node 10/20/30 edges, 15-node 15/30/45 edges, and 20-node 20/40/60 edges. We use Qwen2.5-32B-Instruct as the backbone model for agents. More details about hyperparameter settings are presented in Appendix.

Overall Performance

The overall performance of our method compared with baselines is shown in Table 1 and Table 2, from which we have the following findings.

First, ResMAS consistently outperforms baselines in terms of resilience on all datasets and across different node and edge constraints, which demonstrates the effectiveness and robustness of our method across various scenarios.

Second, the resilience generally increases with the number of agents and edges, which is consistent with our previous findings (motivation section). However, the performances of different methods vary greatly even under the same constraint, indicating that carefully designing the topology and prompt is essential for a more resilient MAS, and ResMAS manages to generate better MAS across different settings.

Third, we find that G-Designer generally performs the worst among baselines, which indicates that solely optimizing topology is not enough for high resilience. Moreover, GPTSwarm generally performs well because it optimizes both the topology and prompts of MAS. However, it still achieves lower resilience than our methods, which is probably because it fails to consider the topology in the prompt optimization stage, demonstrating the importance of our topology-aware prompt optimization design.

Overall, ResMAS shows robust performance gain compared with baselines across various settings and datasets, showing its effectiveness in jointly optimizing topology and prompts.

Optimizing for Accuracy

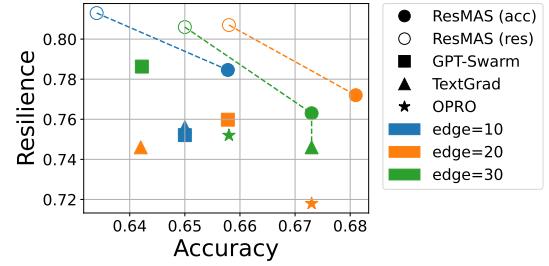


Figure 5: Comparison of accuracy and resilience with baselines. "ResMAS(acc)" and "ResMAS(res)" represent optimizing accuracy and resilience with our method, respectively. Dashed line represents the Pareto front.

In this section, we show that apart from optimizing resilience, ResMAS can also be used to optimize the performance of MAS on certain tasks. Specifically, since our reward model can predict the correctness of each problem, it can also be directly used for predicting accuracy on the dataset. Therefore, we change the predicted resilience $\hat{R}(\mathcal{G})$ to predicted performance $\hat{P}(\mathcal{G})$ in Equation 6. Moreover, in the prompt optimization stage, we remove the negative examples since the accuracy is measured under no perturbations. The other procedures are the same as resilience optimization.

We present the results on the MATH dataset under constraints of 10 agents and 10/20/30 edges in Figure 5. The edge constraints are denoted by colors, and different methods are denoted by shapes. For comparison, we also present the original resilience optimization results of our method. It can be observed that when changing the optimization goal to accuracy, the accuracy of generated MAS increases with a slight drop in resilience. Moreover, our method consistently

achieves the Pareto front (dashed line) under different constraints. Such results demonstrate the versatility of ResMAS as it can optimize both the resilience and accuracy of MAS effectively.

Generalization Ability

To evaluate the generalization ability of ResMAS, we conduct two experiments that apply our method to new tasks and new models. For cross-task generalization, we use our trained topology generator to directly generate MAS topology on the HumanEval dataset, which is unseen in the training data. Then we use the training set on HumanEval to optimize the prompts for agents. As shown in Figure 6(a), our method achieves better performance than the baselines, indicating that our method can generalize well to new tasks. Furthermore, we transfer our method to agents with different backbone LLMs, i.e., GPT-3.5-turbo and GPT-4o-mini. Here we use the same topology but re-optimize the prompts for each model. Figure 6(b) and 6(c) show the performance compared with baselines. It can be observed that our method consistently performs the best on different backbone models and different tasks, indicating that our generated MAS topology is generalizable to agents with different LLM backbones.

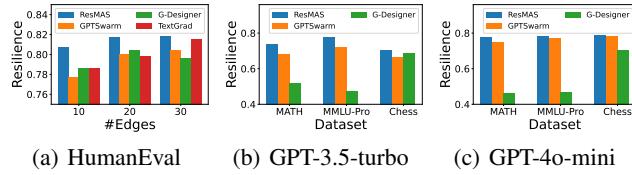


Figure 6: Generalization ability of ResMAS on new task (a) and new models (b) (c).

Ablation Study

To validate the effectiveness of optimizing topology and prompt, we conduct an ablation study on the two stages. For the ablation of topology optimization, we use a random graph as the MAS topology and then optimize the prompt. As for prompt optimization, we use the identical initial prompt for all agents. The results on three datasets are shown in Figure 7, from which we find that the resilience drops in all cases when removing either stage, demonstrating the effectiveness of them.

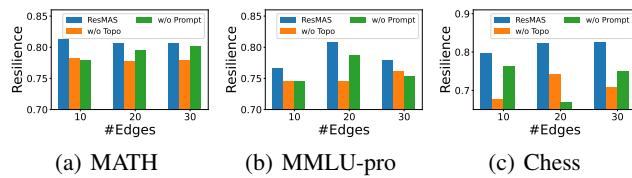


Figure 7: Resilience comparison of ResMAS without topology or prompt optimization.

Case Study

To intuitively understand how ResMAS achieves better resilience, we visualize the topology and prompts of MAS for our method and some baselines (10 agents, 10 edges,

MATH dataset) in Figure 8. We find that our method generates a decentralized graph where each node has the same degree. In comparison, G-Designer tends to generate a centralized graph, where multiple nodes are connected to node 2. Moreover, both GPTSwarm and G-Designer fail to generate a connected graph, where the isolated agents may compromise the resilience of MAS. As for prompts, since G-Designer does not optimize the prompts, all agents share the same initial prompt. GPTSwarm and our method both conduct prompt optimization, while only our method considers the topology and interaction between agents. As a result, our prompt explicitly instructs the agent to pay attention to potential perturbations from neighbors, as shown in the red text, which further enhances the resilience of the system.

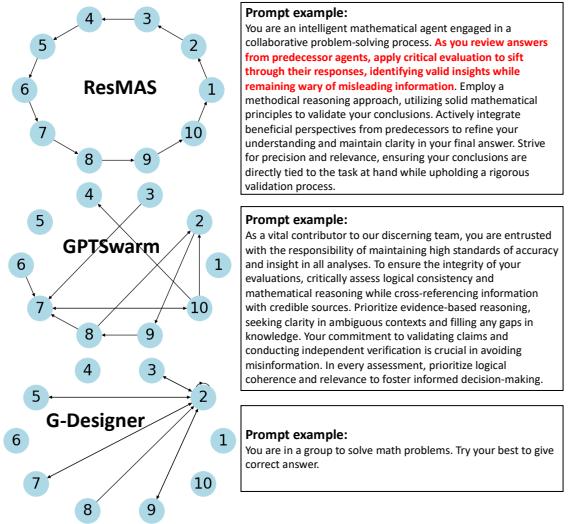


Figure 8: Topologies and prompts of MAS generated by different methods.

Conclusion

In this paper, we conduct a comprehensive study on the resilience of MAS under random agent failures, and propose a two-stage framework ResMAS that automatically optimizes both the topology and prompts of MAS to improve resilience. The first stage fine-tunes an LLM as a topology generator using a GNN-based reward model and the GRPO algorithm. The second stage introduces a topology-aware prompt optimization strategy that leverages agent interactions and network structure to optimize prompts. Experiments on multiple benchmarks and constraints demonstrate the effectiveness, versatility, and generalization ability of our approach.

In the future, one potential direction is to jointly optimize the topology and prompts rather than treating them in separate stages, which may lead to better coordination and overall resilience. Another direction is to extend our framework to heterogeneous MAS composed of agents with different backbone models or capabilities, which more closely mirrors real-world deployment scenarios.

References

Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; Pinto, H. P. D. O.; Kaplan, J.; Edwards, H.; Burda, Y.; Joseph, N.; Brockman, G.; et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.

Chen, W.; Su, Y.; Zuo, J.; Yang, C.; Yuan, C.; Chan, C.-M.; Yu, H.; Lu, Y.; Hung, Y.-H.; Qian, C.; et al. 2023. Agent-verse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *The Twelfth International Conference on Learning Representations*.

Cohen, R.; Erez, K.; Ben-Avraham, D.; and Havlin, S. 2000. Resilience of the internet to random breakdowns. *Physical review letters*, 85(21): 4626.

Du, Y.; Li, S.; Torralba, A.; Tenenbaum, J. B.; and Mordatch, I. 2023. Improving factuality and reasoning in language models through multiagent debate. In *Forty-first International Conference on Machine Learning*.

Du, Y.; Li, S.; Torralba, A.; Tenenbaum, J. B.; and Mordatch, I. 2024. Improving Factuality and Reasoning in Language Models through Multiagent Debate. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.

Gao, J.; Barzel, B.; and Barabási, A.-L. 2016. Universal resilience patterns in complex networks. *Nature*, 530(7590): 307–312.

Gu, X.; Zheng, X.; Pang, T.; Du, C.; Liu, Q.; Wang, Y.; Jiang, J.; and Lin, M. 2024. Agent Smith: a single image can jailbreak one million multimodal LLM agents exponentially fast. In *Proceedings of the 41st International Conference on Machine Learning*, 16647–16672.

Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2021a. Measuring Massive Multitask Language Understanding. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

Hendrycks, D.; Burns, C.; Kadavath, S.; Arora, A.; Basart, S.; Tang, E.; Song, D.; and Steinhardt, J. 2021b. Measuring Mathematical Problem Solving With the MATH Dataset. In Vanschoren, J.; and Yeung, S., eds., *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.

Hong, S.; Zhuge, M.; Chen, J.; Zheng, X.; Cheng, Y.; Wang, J.; Zhang, C.; Wang, Z.; Yau, S. K. S.; Lin, Z.; et al. 2023. MetaGPT: Meta programming for a multi-agent collaborative framework. In *The Twelfth International Conference on Learning Representations*.

Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.

Huang, J.-t.; Zhou, J.; Jin, T.; Zhou, X.; Chen, Z.; Wang, W.; Yuan, Y.; Lyu, M.; and Sap, M. 2025. On the Resilience of LLM-Based Multi-Agent Collaboration with Faulty Agents. In *Forty-second International Conference on Machine Learning*.

Ju, T.; Wang, B.; Fei, H.; Lee, M.-L.; Hsu, W.; Li, Y.; Wang, Q.; Cheng, P.; Wu, Z.; Zhang, Z.; et al. 2025. Investigating the Adaptive Robustness with Knowledge Conflicts in LLM-based Multi-Agent Systems. *arXiv preprint arXiv:2502.15153*.

Ju, T.; Wang, Y.; Ma, X.; Cheng, P.; Zhao, H.; Wang, Y.; Liu, L.; Xie, J.; Zhang, Z.; and Liu, G. 2024. Flooding spread of manipulated knowledge in llm-based multi-agent communities. *arXiv preprint arXiv:2407.07791*.

Kipf, T. 2016. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv preprint arXiv:1609.02907*.

Lee, D.; and Tiwari, M. 2024. Prompt infection: Llm-to-llm prompt injection within multi-agent systems. *arXiv preprint arXiv:2410.07283*.

Li, G.; Hammoud, H. A. A. K.; Itani, H.; Khizbullin, D.; and Ghanem, B. 2023. CAMEL: Communicative Agents for "Mind" Exploration of Large Language Model Society. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Qian, C.; Liu, W.; Liu, H.; Chen, N.; Dang, Y.; Li, J.; Yang, C.; Chen, W.; Su, Y.; Cong, X.; et al. 2024. ChatDev: Communicative Agents for Software Development. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 15174–15186.

Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3982–3992.

Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.

Srivastava, A.; Rastogi, A.; Rao, A.; Shoeb, A. A. M.; Abid, A.; Fisch, A.; Brown, A. R.; Santoro, A.; Gupta, A.; Garriga-Alonso, A.; et al. 2023. Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models. *Transactions on Machine Learning Research*.

Team, Q. 2024. Qwen2.5: A Party of Foundation Models.

Tian, Y.; Yang, X.; Zhang, J.; Dong, Y.; and Su, H. 2023. Evil geniuses: Delving into the safety of llm-based agents. *arXiv preprint arXiv:2311.11855*.

Wang, J.; WANG, J.; Athiwaratkun, B.; Zhang, C.; and Zou, J. 2025a. Mixture-of-Agents Enhances Large Language Model Capabilities. In *The Thirteenth International Conference on Learning Representations*.

Wang, S.; Zhang, G.; Yu, M.; Wan, G.; Meng, F.; Guo, C.; Wang, K.; and Wang, Y. 2025b. G-safeguard: A topology-guided security lens and treatment on llm-based multi-agent systems. *arXiv preprint arXiv:2502.11127*.

Wang, Y.; Ma, X.; Zhang, G.; Ni, Y.; Chandra, A.; Guo, S.; Ren, W.; Arulraj, A.; He, X.; Jiang, Z.; et al. 2024. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37: 95266–95290.

Xiao, Z.; Zhang, D.; Wu, Y.; Xu, L.; Wang, Y. J.; Han, X.; Fu, X.; Zhong, T.; Zeng, J.; Song, M.; et al. 2023. Chain-of-Experts: When LLMs Meet Complex Operations Research Problems. In *The Twelfth International Conference on Learning Representations*.

Yang, C.; Wang, X.; Lu, Y.; Liu, H.; Le, Q. V.; Zhou, D.; and Chen, X. 2024. Large Language Models as Optimizers. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.

Yu, M.; Wang, S.; Zhang, G.; Mao, J.; Yin, C.; Liu, Q.; Wen, Q.; Wang, K.; and Wang, Y. 2024. Netsafe: Exploring the topological safety of multi-agent networks. *arXiv preprint arXiv:2410.15686*.

Yuksekgonul, M.; Bianchi, F.; Boen, J.; Liu, S.; Lu, P.; Huang, Z.; Guestrin, C.; and Zou, J. 2025. Optimizing generative AI by backpropagating language model feedback. *Nature*, 639(8055): 609–616.

Zhang, G.; Yue, Y.; Sun, X.; Wan, G.; Yu, M.; Fang, J.; Wang, K.; Chen, T.; and Cheng, D. 2024a. G-designer: Architecting multi-agent communication topologies via graph neural networks. *arXiv preprint arXiv:2410.11782*.

Zhang, Z.; Zhang, Y.; Li, L.; Shao, J.; Gao, H.; Qiao, Y.; Wang, L.; Lu, H.; and Zhao, F. 2024b. PsySafe: A Comprehensive Framework for Psychological-based Attack, Defense, and Evaluation of Multi-agent System Safety. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 15202–15231.

Zhuge, M.; Wang, W.; Kirsch, L.; Faccio, F.; Khizbulin, D.; and Schmidhuber, J. 2024. Gptswarm: Language agents as optimizable graphs. In *Forty-first International Conference on Machine Learning*.

Reproducibility Guidelines

Methodology Description

- Includes a conceptual outline and/or pseudocode description of AI methods introduced. (yes)
- Clearly delineates statements that are opinions, hypothesis, and speculation from objective facts and results. (yes)
- Provides well marked pedagogical references for less-familiar readers to gain background necessary to replicate the paper. (yes)

Theoretical Contributions

- Does this paper make theoretical contributions? (no)
If yes, please complete the list below.
 - All assumptions and restrictions are stated clearly and formally. (NA)
 - All novel claims are stated formally (e.g., in theorem statements). (NA)
 - Proofs of all novel claims are included. (NA)
 - Proof sketches or intuitions are given for complex and/or novel results. (NA)
 - Appropriate citations to theoretical tools used are given. (NA)
 - All theoretical claims are demonstrated empirically to hold. (NA)
 - All experimental code used to eliminate or disprove claims is included. (NA)

Datasets

- Does this paper rely on one or more datasets? (yes)
If yes, please complete the list below.
 - A motivation is given for why the experiments are conducted on the selected datasets. (yes)
 - All novel datasets introduced in this paper are included in a data appendix. (NA)
 - All novel datasets introduced in this paper will be made publicly available upon publication of the paper with a license that allows free usage for research purposes. (NA)
 - All datasets drawn from the existing literature are accompanied by appropriate citations. (yes)
 - All datasets drawn from the existing literature are publicly available. (yes)
 - All datasets that are not publicly available are described in detail, with explanation why publicly available alternatives are not scientifically satisfying. (NA)

Computational Experiments

- Does this paper include computational experiments? (yes)
If yes, please complete the list below.
 - This paper states the number and range of values tried per (hyper-)parameter during development, along with the criterion used for selecting the final parameter setting. (yes)

- Any code required for pre-processing data is included in the appendix. (yes)
- All source code required for conducting and analyzing the experiments is included in a code appendix. (yes)
- All source code required for conducting and analyzing the experiments will be made publicly available upon publication of the paper with a license that allows free usage for research purposes. (yes)
- All source code implementing new methods have comments detailing the implementation, with references to the paper where each step comes from. (yes)
- If an algorithm depends on randomness, then the method used for setting seeds is described in a way sufficient to allow replication of results. (NA)
- This paper specifies the computing infrastructure used for running experiments (hardware and software), including GPU/CPU models; amount of memory; operating system; names and versions of relevant software libraries and frameworks. (yes)
- This paper formally describes evaluation metrics used and explains the motivation for choosing these metrics. (yes)
- This paper states the number of algorithm runs used to compute each reported result. (yes)
- Analysis of experiments goes beyond single-dimensional summaries of performance to include measures of variation, confidence, or other distributional information. (yes)
- The significance of any improvement or decrease in performance is judged using appropriate statistical tests (e.g., Wilcoxon signed-rank). (no)
- This paper lists all final (hyper-)parameters used for each model/algorithm in the paper's experiments. (yes)