


# **RADAR: Redundancy-Aware Diffusion for Multi-Agent Communication Structure Generation**

**Zhen Zhang<sup>1</sup>, Wanjing Zhou<sup>1</sup>, Juncheng Li<sup>2</sup>,  
Hao Fei<sup>3</sup>, Jun Wen<sup>4</sup>, Wei Ji<sup>1</sup>**


**<sup>1</sup>Nanjing University, <sup>2</sup>Zhejiang University,  
<sup>3</sup>University of Oxford, <sup>4</sup>MBZUAI**

# Background

- Large Language Model based agents have achieved great success across a wide spectrum of domains, including code generation, web navigation and complex reasoning, etc.
- Despite the impressive capabilities of LLM-based agents, single-agent architectures face limitations in specialization, parallel exploration, and robustness



I have to do  
**EVERYTHING** by  
myself



Planning



Coding



Search



Verification



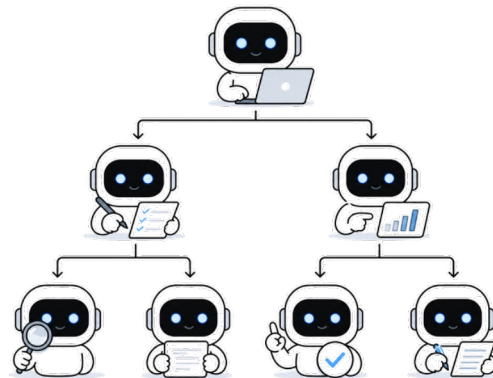
Decision  
Making

# Background

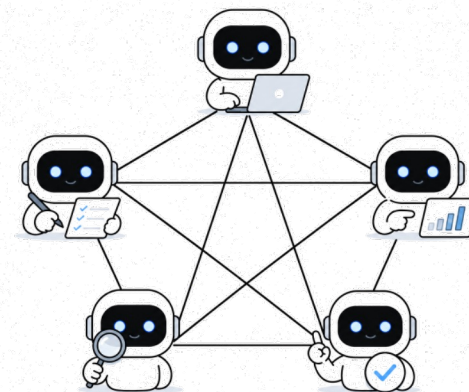
- Recent research has turned to multi-agent systems (MAS), where multiple autonomous agents interact through communication, coordination, and role specialization
- Prior work on multi-agent systems has largely been constrained to static, hand-crafted configurations, such as chain structures, tree architectures, fully connected topologies, etc.



Chain



Tree

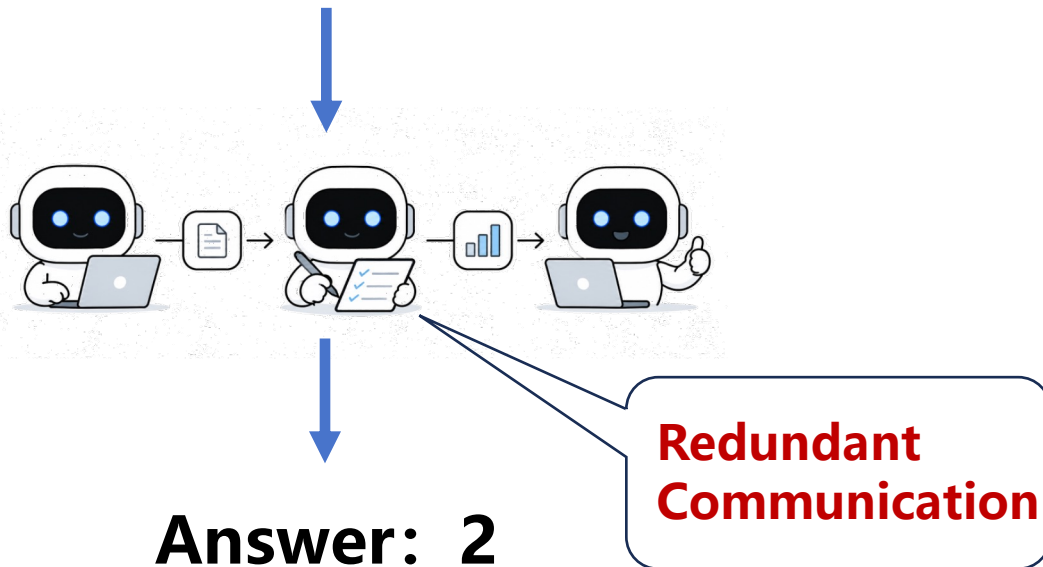


Complete  
Graph

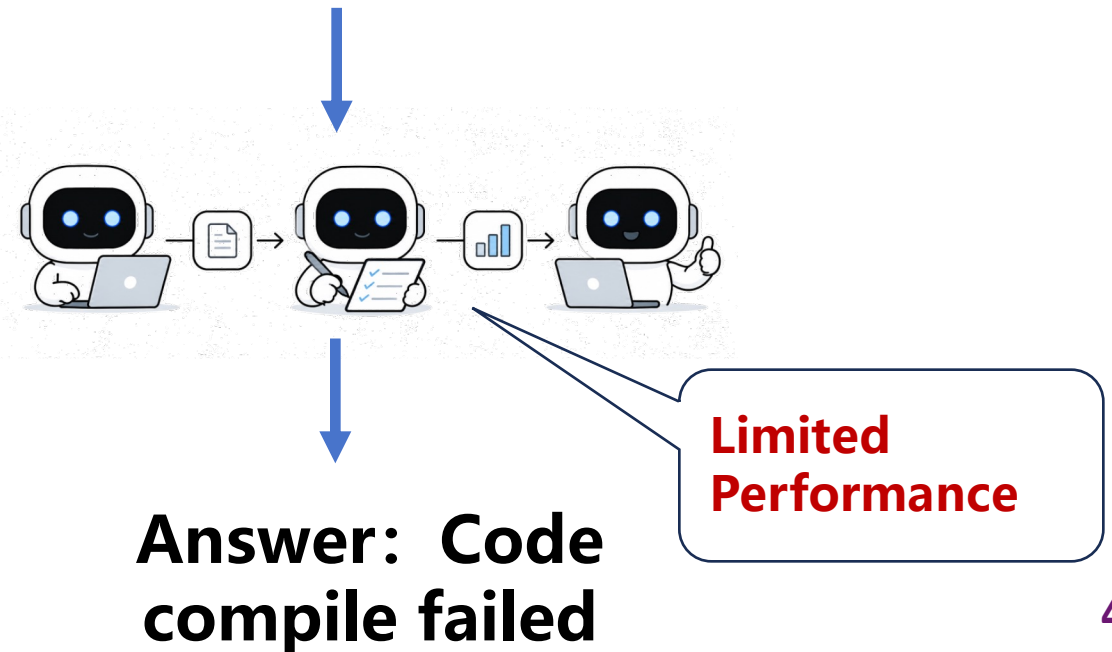
# Background

- Applying a single fixed collaboration pattern across all tasks either incurs redundant communication overhead for simple problems or limits performance on more complicated ones

Query: Compute the value of the expression  $1+1$ ?




Query: Generate code that implements a *balanced binary tree*.



# Background

- In this paper, we propose RADAR, which synthesizes the entire collaboration graph via iterative conditional graph diffusion models and generates collective intelligence tailored to the specific task

Query: Compute the value of the expression  $1+1$ ?

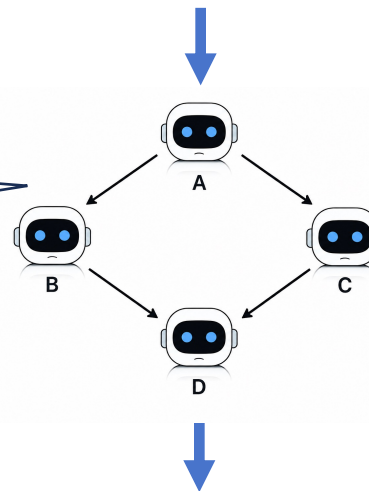


Answer: 2

Task Specific Structure

Reduced Communication

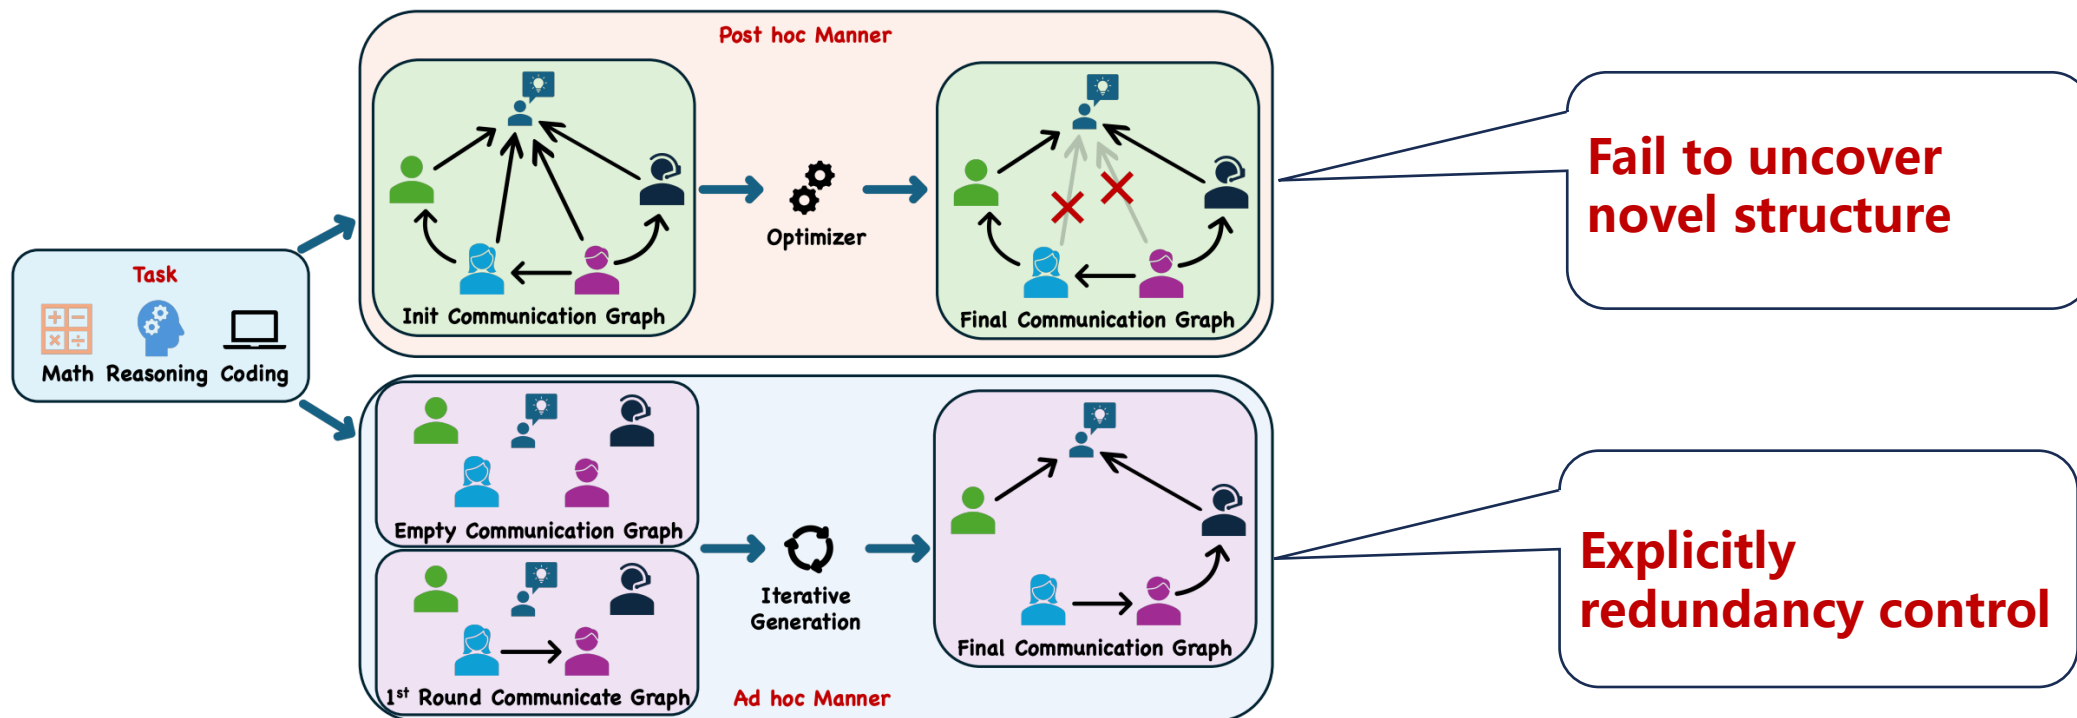
Query: Generate code that implements a *balanced binary tree*.



Answer: Code compile success

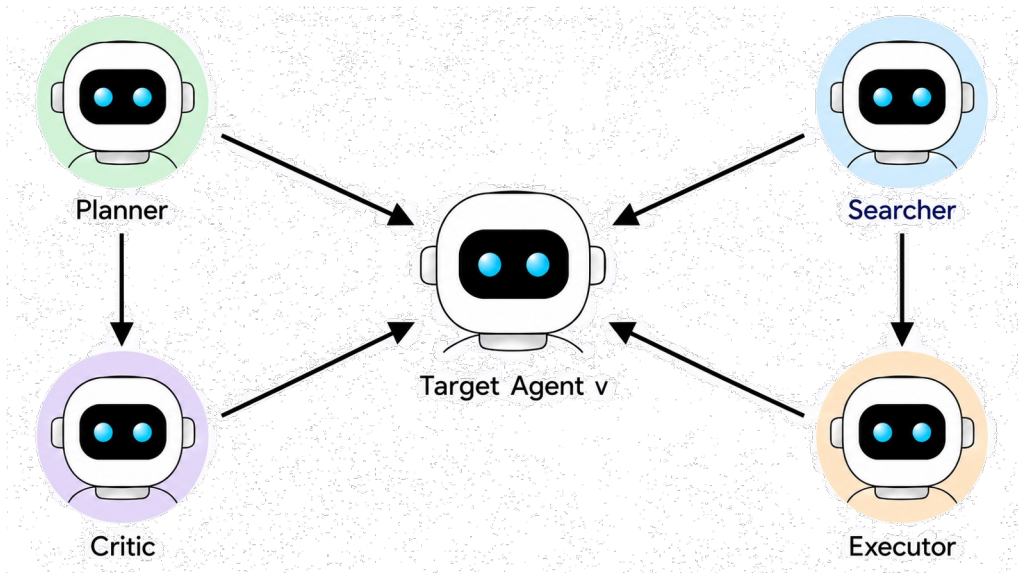
# Related Work

- Recent models generate the communication structure in a **single step**. This often leads to inefficient topologies with redundant links, etc.
- To improve efficiency, they often optimize it in a **post hoc manner**, rather than jointly reasoning about structure formation and redundancy control during topology construction




- We formulate communication topology design as a step-by-step generation process, guided by the *effective size* of the graph

$$\varphi^i(v_k) = |N_i(v_k)| - \frac{\sum_{j,q \in N_i(v_k)} A_{jq} I[r(j) = r(q)]}{|N_i(v_k)|}$$



**High Effective Size** ---> **Low Redundancy**



**Low Effective Size** ---> **High Redundancy**

- We introduce a forward diffusion process that progressively masks nodes together with their connected edges

$$q_{\varphi}(\pi|G_0, \varphi) = \prod_t q_{\varphi}(\pi_t|G_0, \varphi, \pi_{(<t)})$$

- Ordering network  $q_{\varphi}(\pi|G_0, \varphi)$  samples a node  $v_{\pi(t)}$  to be masked at each diffusion step  $t$ , yielding the corresponding partially masked graph  $G_t$

$$q_{\varphi}(\pi_t|G_0, \varphi, \pi_{(<t)}) = \frac{\exp(h_t + \varphi(v_t))}{\sum_{j \notin \pi_{(<t)}} \exp(h_j)}$$

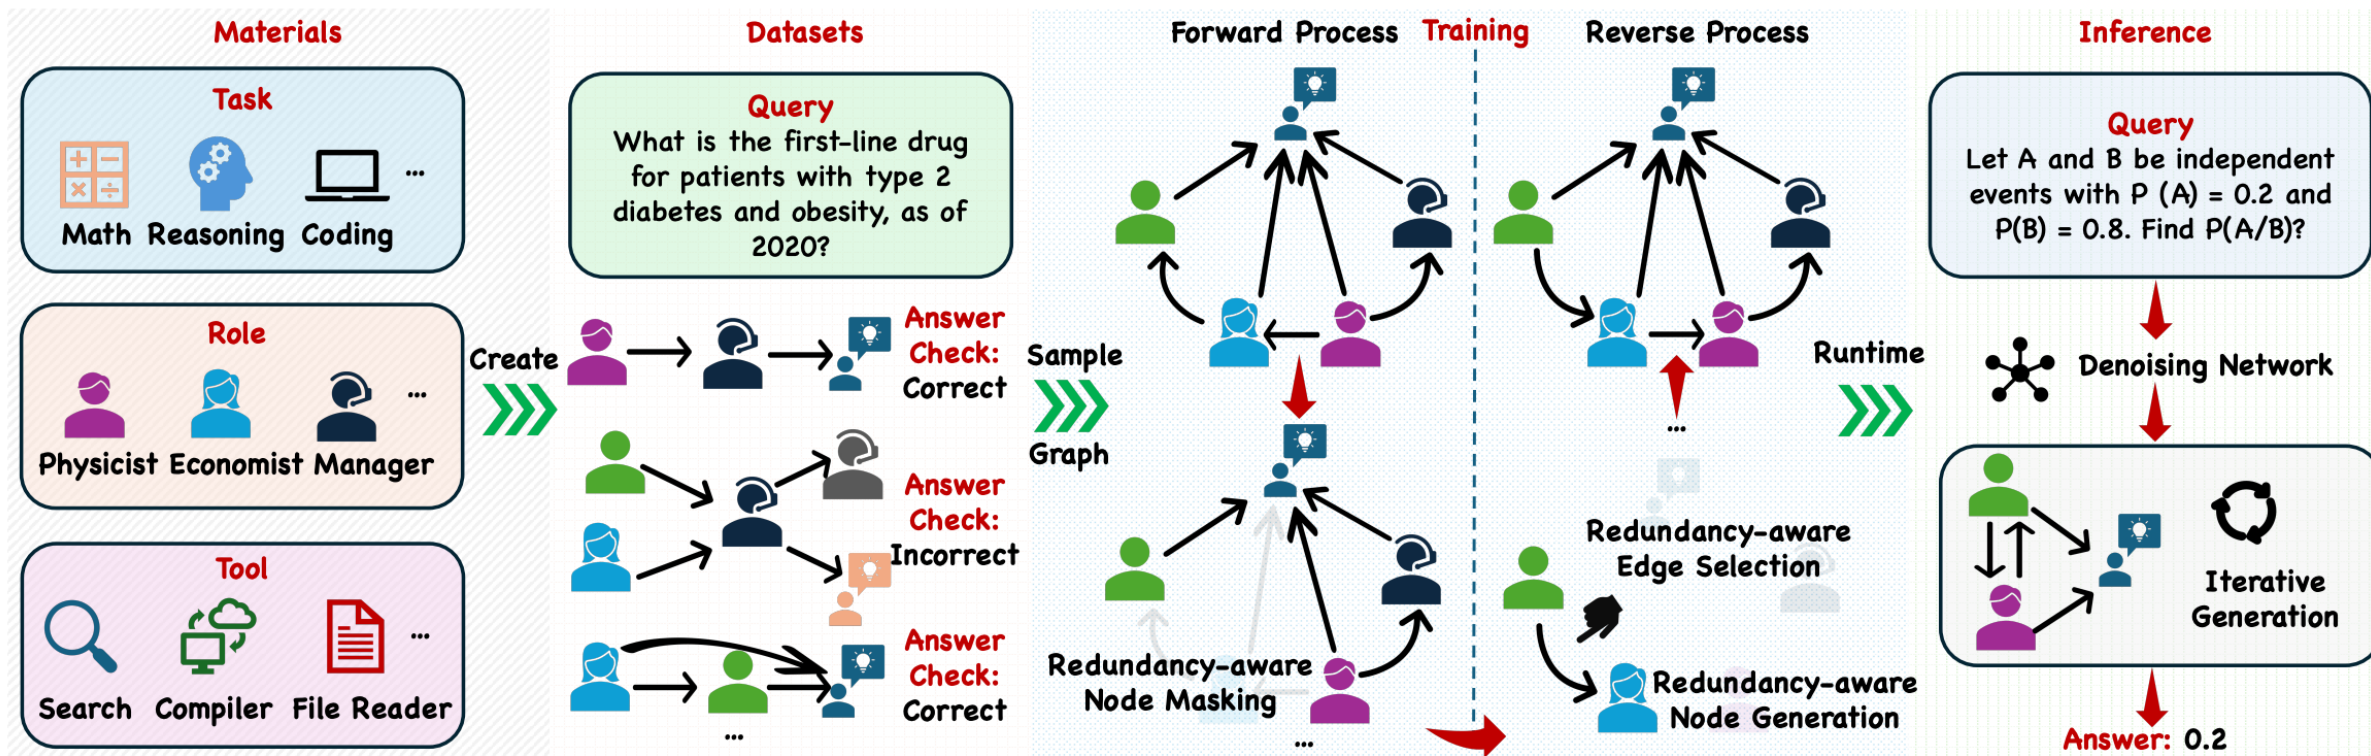
- It outputs the probability of node  $v_t$  being selected at step  $t$

- During the reverse generative phase, a denoising network  $p_{\theta}(G_t|G_{t+1}, Q)$  progressively reconstructs the graph by inverting the forward diffusion process conditioning on the query  $Q$
- At step  $t$ , the denoising network takes the partially masked graph  $G_{t+1}$  as input and maps each node  $v_i$  into a latent space

**Any GNN model can be used!**

- Based on these embeddings, multilayer perceptions are utilized to predict the agent role of the newly recovered node  $v_{\pi_t}$  and its connectivity to the set of previously denoised nodes  $\{v_{\pi_{(>t)}}$

## ➤ We present an overview of the proposed **RADAR** framework



- We use a ***reinforcement learning*** based training strategy to jointly update the diffusion ordering network  $q_\varphi(\pi|G_0, \varphi)$  and the denoising network  $p_\theta(G_t|G_{t+1}, Q)$  using gradient descent
- At each training iteration, for the  $i$ -th training graph  $G_0^{(i)}$ , we generate ***M diffusion trajectories*** by sampling node-masking orderings  $\pi^{i,m}$  from the ordering network
- Each trajectory consists of a sequence of partially masked graphs  $\{G_t^{i,m}\}_{1 \leq t \leq N'}$  where  $N'$  denotes the number of nodes
- Conditioned on these sampled trajectories, the denoising network is trained to minimize the negative ***variational lower bound*** using stochastic gradient descent

- **Datasets:** we conduct experiments using three categories of publicly available datasets
- **Three key category baselines:**
  - *Single Agent Methods*, *Multi-Agent Systems*, *Autonomous Multi-Agent Systems*

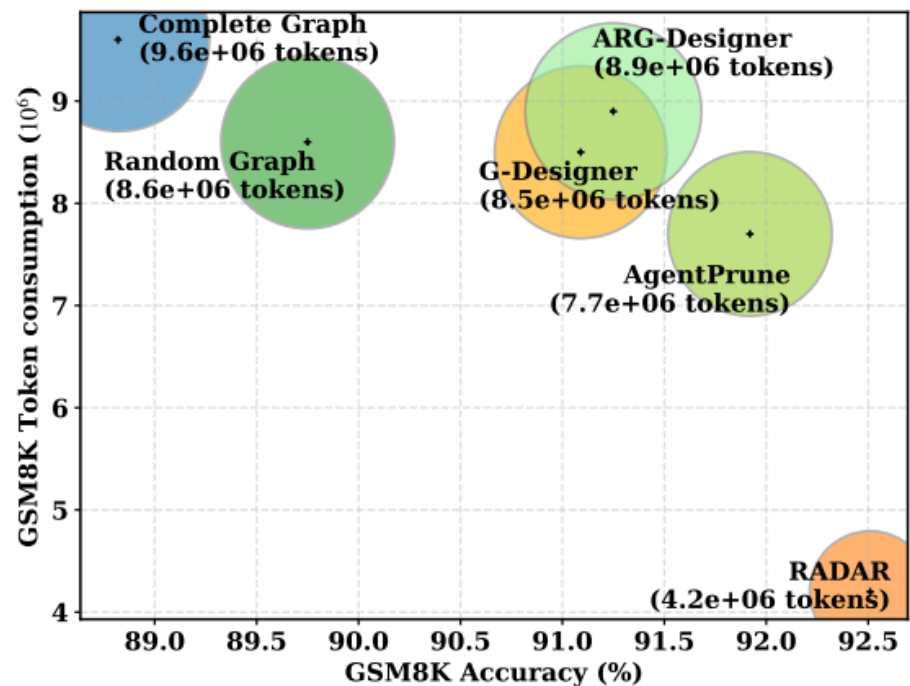
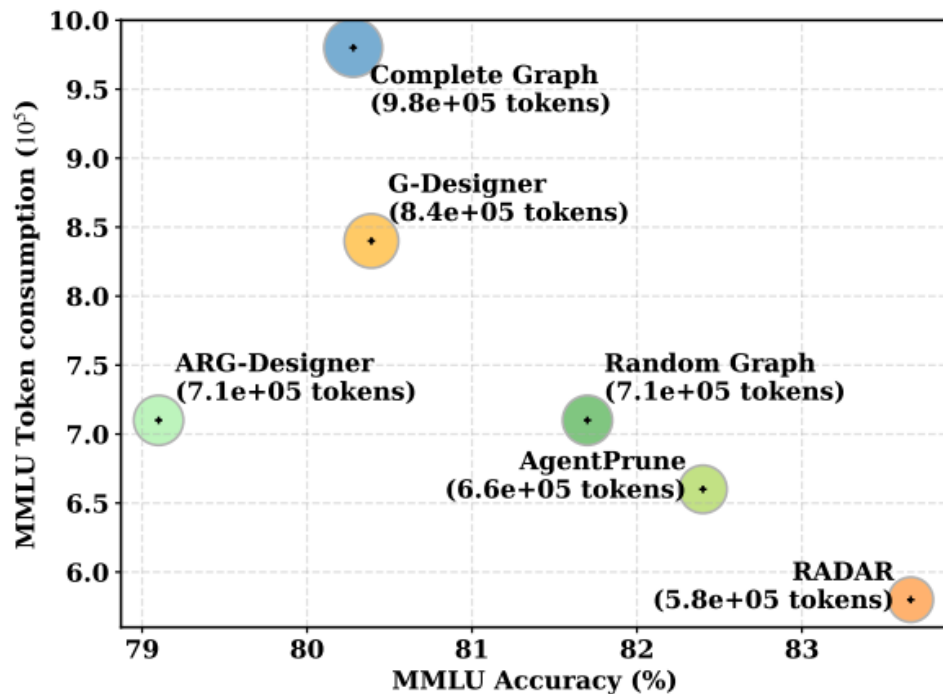
Category	Dataset	Format	Metric	#Samples
General Reasoning	MMLU	Multi-choice	Accuracy	153
	GSM8K	Number	Accuracy	1,319
Math Solving	MultiArith	Number	Accuracy	600
	SVAMP	Number	Accuracy	1,000
	AQuA	Multi-choice	Accuracy	254
Code Generation	HumanEval	Code	Pass@1	164

## ➤ Performance comparison across three categories of baselines

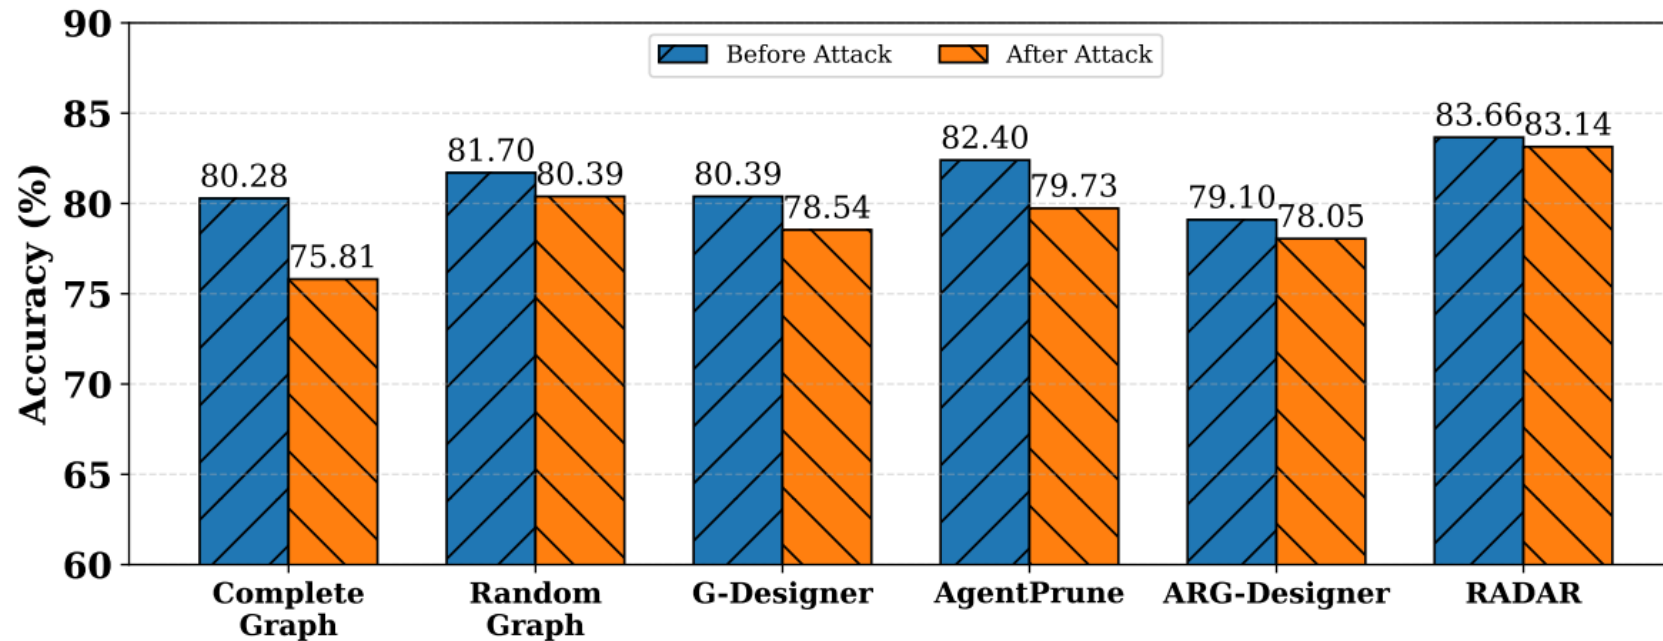
Methods	Auto	MMLU	GSM8K	MultiArith	SVAMP	AQuA	HumanEval	Avg.
Vanilla	✗	78.54	87.45	96.85	86.67	78.92	87.08	85.92
CoT	✗	79.26 <sup>↑0.72</sup>	87.10 <sup>↓0.35</sup>	96.31 <sup>↓0.54</sup>	87.33 <sup>↑0.66</sup>	75.20 <sup>↓3.72</sup>	88.13 <sup>↑1.05</sup>	85.55
ComplexCoT	✗	79.80 <sup>↑1.26</sup>	86.89 <sup>↓0.56</sup>	96.70 <sup>↓0.15</sup>	87.67 <sup>↑1.00</sup>	75.59 <sup>↓3.33</sup>	87.49 <sup>↑0.41</sup>	85.69
SC (COT×5)	✗	80.66 <sup>↑2.12</sup>	87.57 <sup>↑0.12</sup>	96.58 <sup>↓0.27</sup>	88.00 <sup>↑1.33</sup>	82.28 <sup>↑3.36</sup>	88.60 <sup>↑1.52</sup>	87.28
MultiPersona	✗	77.69 <sup>↓0.85</sup>	87.50 <sup>↑0.05</sup>	97.49 <sup>↑0.64</sup>	87.00 <sup>↑0.33</sup>	79.23 <sup>↑0.31</sup>	88.32 <sup>↑1.24</sup>	86.21
LLM-Debate	✗	80.56 <sup>↑2.02</sup>	89.47 <sup>↑2.02</sup>	97.33 <sup>↑0.48</sup>	89.00 <sup>↑2.33</sup>	79.70 <sup>↑0.78</sup>	88.68 <sup>↑1.60</sup>	87.46
LLM-Blender	✗	80.29 <sup>↑1.75</sup>	88.35 <sup>↑0.90</sup>	97.29 <sup>↑0.44</sup>	87.33 <sup>↑0.66</sup>	78.99 <sup>↑0.07</sup>	88.80 <sup>↑1.72</sup>	86.84
DyLAN	✗	79.86 <sup>↑1.32</sup>	89.98 <sup>↑2.53</sup>	97.12 <sup>↑0.27</sup>	88.67 <sup>↑2.00</sup>	79.59 <sup>↑0.67</sup>	90.42 <sup>↑3.34</sup>	87.61
AgentVerse	✗	78.39 <sup>↓0.15</sup>	89.91 <sup>↑2.46</sup>	97.50 <sup>↑0.65</sup>	88.33 <sup>↑1.66</sup>	77.47 <sup>↓1.45</sup>	89.29 <sup>↑2.21</sup>	86.82
MacNet	✗	79.55 <sup>↑1.01</sup>	87.95 <sup>↑0.50</sup>	96.03 <sup>↓0.82</sup>	86.00 <sup>↓0.67</sup>	79.23 <sup>↑0.31</sup>	84.57 <sup>↓2.51</sup>	85.55
AutoAgents	✓	79.59 <sup>↑1.05</sup>	87.69 <sup>↑0.24</sup>	96.42 <sup>↓0.43</sup>	86.34 <sup>↓0.33</sup>	78.65 <sup>↓0.27</sup>	87.64 <sup>↑0.56</sup>	86.05
GPTSwarm	✓	78.36 <sup>↓0.18</sup>	89.14 <sup>↑1.69</sup>	96.79 <sup>↓0.06</sup>	88.67 <sup>↑2.00</sup>	80.71 <sup>↑1.79</sup>	89.32 <sup>↑2.24</sup>	87.17
ADAS	✓	78.39 <sup>↓0.15</sup>	86.12 <sup>↓1.33</sup>	96.02 <sup>↓0.83</sup>	86.33 <sup>↓0.34</sup>	77.71 <sup>↓1.21</sup>	84.19 <sup>↓2.89</sup>	84.79
AgentSquare	✓	79.58 <sup>↑1.04</sup>	87.62 <sup>↑0.17</sup>	97.77 <sup>↑0.92</sup>	88.00 <sup>↑1.33</sup>	81.50 <sup>↑2.58</sup>	89.08 <sup>↑2.00</sup>	87.26
AFlow	✓	81.80 <sup>↑3.26</sup>	91.16 <sup>↑3.71</sup>	96.22 <sup>↓0.63</sup>	88.33 <sup>↑1.66</sup>	80.90 <sup>↑1.98</sup>	90.93 <sup>↑3.85</sup>	88.22
G-Designer	✓	80.39 <sup>↑1.85</sup>	91.09 <sup>↑3.64</sup>	97.78 <sup>↑0.93</sup>	90.00 <sup>↑3.33</sup>	80.75 <sup>↑1.83</sup>	89.37 <sup>↑2.29</sup>	88.23
AgentPrune	✓	82.40 <sup>↑3.86</sup>	91.92 <sup>↑4.47</sup>	97.88 <sup>↑1.03</sup>	90.37 <sup>↑3.70</sup>	80.93 <sup>↑2.01</sup>	87.17 <sup>↑0.09</sup>	88.22
GTD	✓	79.41 <sup>↑0.87</sup>	91.38 <sup>↑3.93</sup>	96.20 <sup>↓0.65</sup>	90.24 <sup>↑3.57</sup>	79.44 <sup>↑0.52</sup>	87.86 <sup>↑0.78</sup>	87.42
MaAS	✓	82.32 <sup>↑3.78</sup>	91.13 <sup>↑3.68</sup>	98.08 <sup>↑1.23</sup>	89.65 <sup>↑2.98</sup>	80.25 <sup>↑1.33</sup>	89.57 <sup>↑2.49</sup>	88.50
ARG-Designer	✓	79.10 <sup>↑0.56</sup>	91.25 <sup>↑3.80</sup>	98.55 <sup>↑1.70</sup>	92.21 <sup>↑5.54</sup>	81.10 <sup>↑2.18</sup>	89.19 <sup>↑2.11</sup>	88.57
RADAR	✓	<b>83.66</b> <sup>↑5.12</sup>	<b>92.51</b> <sup>↑5.06</sup>	<b>98.81</b> <sup>↑1.96</sup>	<b>93.26</b> <sup>↑6.59</sup>	<b>82.84</b> <sup>↑3.92</sup>	<b>91.28</b> <sup>↑4.20</sup>	<b>90.32</b>

# Experiments

- Our proposed RADAR demonstrates a favorable trade-off, maintaining strong performance while exhibiting high token efficiency



- We inject system prompt attacks into two of the five agents in the collaboration framework
- RADAR exhibits exceptional resilience to adversarial perturbations, maintaining nearly identical performance before and after the attack



- We propose an iterative framework for multiagent collaboration topology design, which provides *high accuracy*, *low token consumption*, and *strong robustness* across diverse scenarios
- We employ a graph diffusion model that leverages graph effectiveness signals to iteratively construct collaboration topologies, *explicitly modeling redundancy* as part of the topology generation process
- We conduct extensive experiments on six benchmarks spanning *general reasoning*, *code generation* and *mathematical problem solving*, which collectively demonstrate the effectiveness of our method

**Thanks**  
**Q&A**



**Code & Data**