ASMAD: Adaptive Sparse Communication Topology Multi-Agent Debate Framework with Opinion Dynamics

Anonymous ACL submission

Abstract

Large language models (LLMs) still face challenges in complex reasoning within multi-agent debate (MAD) systems due to high computational costs in fully-connected structures. While existing methods use static sparse topologies to reduce computation, they neglect semantic relationships and dynamic opinion evolution. To solve this challenge, we propose ASMAD, an adaptive sparse topology framework that synergizes sociophysical opinion dynamics with LLMs through two innovations: (1) probabilistic semantic-guided attention gates for dynamic opinion visibility control; (2) a hybrid paradigm combining adaptive trust-boundary regulation and opinion synchronization. Experiments show ASMAD reduces token costs to around 33% across GSM8K and MMLU benchmarks while maintaining competitive accuracy with 4-bit quantized 7-9B size models.

1 Introduction

In recent years, the rapid development of large language models (LLM) has greatly promoted the progress of several natural language processing (NLP) tasks (Touvron et al., 2023; Zhao et al., 2023; Naveed et al., 2023; Jiang et al., 2024; Achiam et al., 2023; GLM et al., 2024; Guo et al., 2025). However, performance of LLM in reasoning and logical reasoning tasks is still limited (Zhu et al., 2022; Gou et al., 2023).

To address complex reasoning challenges, various approaches has been developed, including Chain-of-Thought (CoT) (Wei et al., 2022), self-consistency (SC) mechanisms (Wang et al., 2022) with self-correction strategies (Liang et al., 2023). Recent advances in multi-agent debate (MAD) systems have demonstrated superior performance in complex reasoning tasks (Liang et al., 2023). Inspired by the human discussion mechanism (Hill et al., 2015; Liang et al., 2023), MAD systems

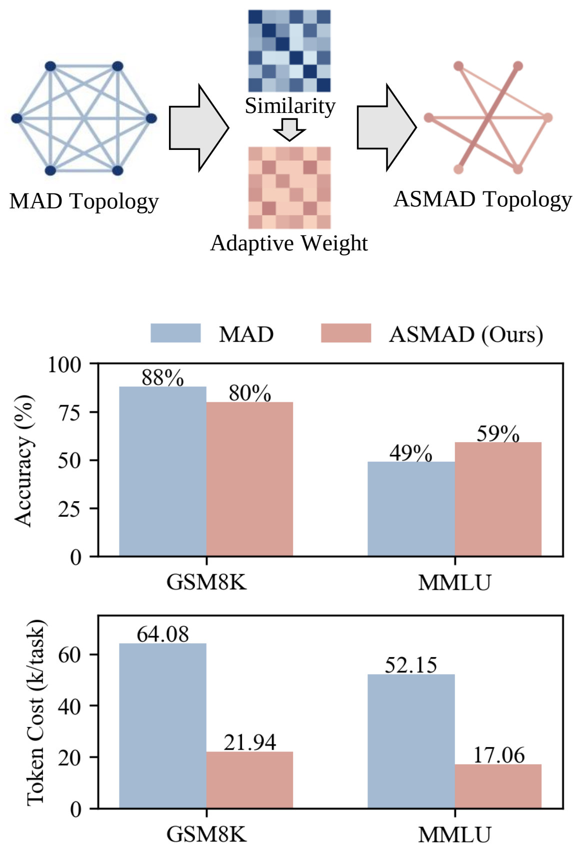

Figure 1: Adaptive topology of ASMAD (Top) and comparison of accuracy (Middle) and token consumption (Bottom)

employ multiple LLM agents to communicate and iteratively argue with each other in a structured debate. However, MAD systems face computation cost problem due to fully-connected communication topology, where every agent interacts with all peers, which incurs quadratic computational complexity that becomes prohibitively expensive for real-world applications (Du et al., 2023).

Existing attempts to address this efficiency challenge focus on either static sparse topologies (e.g., ring or star structures) that reduce token costs through predetermined connection patterns (Du et al., 2023; Sun et al., 2023; Li et al., 2024) or group discussion method that adopts a hierarchical structure by clustering agents into smaller debate groups to exchange intermediate results (Liu et al., 2024; Zeng et al., 2025). However, existing approaches face two fundamental limitations: (1) Task-semantic blindness: fixed topologies cannot adapt to problem difficulty, potentially pruning critical debate pathways; (2) Coarse adaptation granularity: fixed grouping patterns cannot capture nuanced opinion evolution dynamics.

To address these limitations, we propose a adaptive sparse topology framework (ASMAD) that synergies sociophysical opinion dynamics with modern LLM architectures. Our key insight stems from two observations: First, human consensus formation naturally evolves communication networks through confidence-bound adaptation, suggesting that artificial debate systems should similarly adjust interaction patterns based on semantic convergence states. Second, semantic similarity between textual opinions provides a more reliable signal for trust boundary calculation than numerical difference metrics. Building upon this foundation, we propose a dual-regulation debate mechanism that hybridizes two classical models: The Hegselmann-Krause model (Rainer and Krause, 2002) inspired adaptive trust boundary allows agents to dynamically adjust their openness to divergent views based on real-time semantic proximity, while the Deffuant model (Deffuant et al., 2000) derived synchronization protocol coordinates opinion aggregation through gradient descent in the semantic space. The system's core innovation lies in its visibility control module, which implements selective opinion exposure through attention-based gates. By projecting discrete argument exchanges onto the semantic manifold, the module prioritizes information flow along dimensions of highest convergence potential. We evaluate ASMAD across GSM8K (Cobbe et al., 2021) and MMLU (Hendrycks et al., 2021) benchmarks[1] using 4-bit quantized versions of LLaMA-8B (Touvron et al., 2023), ChatGLM-9B (GLM et al., 2024) and Deepseek-7B (Guo et al., 2025). Experiments show ASMAD reduces token costs by 65.7-67.3% across GSM8K and MMLU benchmarks while maintaining competitive accuracy (8% decreasing on GSM8K, 10% improving on MMLU).

In summary, our work contributes as following:

- We developed dynamic visibility control mechanisms for agent opinions, decreasing the communication cost and accelerating consensus formation in MAD.
- We extended classical opinion dynamics to LLM-based MAD systems where a tunable debate paradigm was proposed by integrating Deffuant model's adaptive trust-boundary regulation with Hegselmann-Krause model's synchronized opinion aggregation.
- We introduced a methodology replacing conventional numerical difference metrics with SentenceTransformer-based semantic vector encoding and similarity matrix construction as a potential workaround of effective handling unstructured textual opinions in LLM multi-agent systems.

[1]MIT License

2 Methodology

2.1 Dynamic Opinion Exchange Framework

Multi-agent debate (MAD) with large language models presents unique challenges that traditional frameworks struggle to address. This work reframes the MAD process through the theoretical lens of opinion dynamics, treating each LLM as an agent with bounded rationality, whose willingness to incorporate external viewpoints varies dynamically based on semantic proximity and confidence levels. Drawing from both HK and Deffuant models, we implement: **Simultaneous Updates**: All agents update their states based on visible information, **Probabilistic Interaction**: Probabilities and strength of pairwise interaction determined by adaptive weights.

Unlike classical opinion dynamics that operate in numerical spaces, our framework extends into rich semantic embeddings where agent states comprise both reasoning processes and discrete conclusions. We introduce the agent state as $s_i^t = (r_i^t, c_i^t)$, where $r_i^t \in \mathbb{R}^d$ represents the semantic embedding of agent i's reasoning at time t, and c_i^t denotes its conclusion. This richer state space enables more nuanced modeling of debate dynamics while preserving the mathematical tractability of opinion evolution.

2.2 Adaptive Debate Protocol

As detailed in Figure 3, the proposed protocol orchestrates multi-agent debate through distinct phases that progressively refine agent opinions while maintaining diversity and efficiency.

Independent Initialization Each agent independently generates its initial response to the given problem without access to other agents' outputs. Formally, at $t = 0$, agent i produces state $s_i^0 = (r_i^0, c_i^0)$, where r_i^0 represents its reasoning embedding and c_i^0 its initial conclusion. This independence in initialization is crucial for establishing diverse starting points in the solution space.

Confidence Boundary Determination Following initialization, we adopt the bounded confidence mechanism from classical opinion dynamics models (Deffuant et al., 2000; Rainer and Krause, 2002). A confidence radius $R(t) = R_0 + \lambda\left(\frac{t}{T}\right)$ determines whether agents can consider opinions from each other, where R_0 is the initial radius and λ controls its temporal evolution. Two agents i and j can potentially interact only if their semantic distance falls within this radius: $E_{ij}^t = \mathbb{I}(d(s_i^t, s_j^t) \leq R(t))$, where $d(s_i^t, s_j^t)$ denotes the distance between agents' state and $\mathbb{I}(\cdot)$ is the indicator function. This bounded confidence mechanism helps prevent premature convergence while allowing the interaction scope to gradually expand as the debate progresses.

Weighted Opinion Exchange For agent pairs within confidence bounds, we compute influence weights based on both semantic similarity and answer conclusion consistency (See C.3). The overall influence weight incorporates this similarity measure along with agent-specific attributes:

$$w_{ij}^t = \beta_0 + \beta_1 \left(\frac{t}{T}\right)(1 + \gamma\sigma_i^t), \cdot \mathrm{sim}(s_i^t, s_j^t) \quad (1)$$

where β_0 is the base confidence level, β_1 is the growth rate corresponding to debate progress, γ is the stability influence factor, σ_i^t denotes the agent's stability score and $\mathrm{sim}(s_i^t, s_j^t)$ is the similarity score betweem agents' state.

These weights serve both topology and influence strength in regulating inter-agent interactions. Visibility of agent j's response to i is sampled according to the weight w_{ij}^t (w_{ji}^t if i to j), acting as the probability. Such adaptive directional topology effectively reduces communication token overhead while preserving essential information flow paths.

Construction of agent prompts with varies with degrees of interaction strength, as practical workaround of opinion dynamics model in MAD scenarios. LLMs are prompted with *Critical*, *Reference* and *Background* categories according to w if satisfied various thresholds (See C.1).

Consensus Formation The consensus formation emerges through iterative debate rounds where agents continuously refine their positions through structured interactions:

$$s_i^{t+1} = f_{\mathrm{LLM}}(s_i^t, \{(w_{ij}^t, s_j^t) | j \in \mathcal{N}_i^t\}) \quad (2)$$

where $\mathcal{N}_i^t$ represents the set of visible agents to i at time t, and f_{LLM} denotes the language model's reasoning process. After sufficient rounds of debate, the final conclusion is determined through majority voting.

3 Experiments

3.1 Tasks and Datasets

We evaluate our framework on two benchmark datasets: GSM8K (Cobbe et al., 2021) and MMLU (Hendrycks et al., 2021),that either require multi-step reasoning or admit multiple valid solution paths while maintaining unambiguous answers. GSM8K presents grade school math word problems requiring step-by-step numerical reasoning. MMLU covers multiple-choice questions across various domains, where the challenge lies not only in answer format but in the diversity of valid reasoning approaches. We sampled 100 tasks from each dataset for agents to debate for 5 rounds as benchmark.

3.2 Model Configuration

To thoroughly evaluate the dynamic aspects and diversity benefits of our framework, we construct a heterogeneous agent population using three different LLM architectures: **LLaMA-3.1-8B-Instruct** (Touvron et al., 2023), **ChatGLM-4-9B-chat-abliterated** (GLM et al., 2024) and **Deepseek-math-7b-Instruct** (Guo et al., 2025). Each model type contributes 2 agents, resulting in a debate group of 6 participants. This configuration ensures sufficient diversity in reasoning approaches while maintaining manageable computational requirements. For practical deployment considerations, all deployed models leverage 4-bit block-wise quantization with mixed precision (Q4_K_M), enabling execution on a single NVIDIA GeForce RTX 3090 GPU. This implementation detail is particularly noteworthy as it demonstrates the framework's viability in resource-constrained environments.

3.3 Baseline and Evaluation Protocol

The primary baseline for comparison is a fully-connected debate protocol without visibility control or prompt structuring. This baseline main-

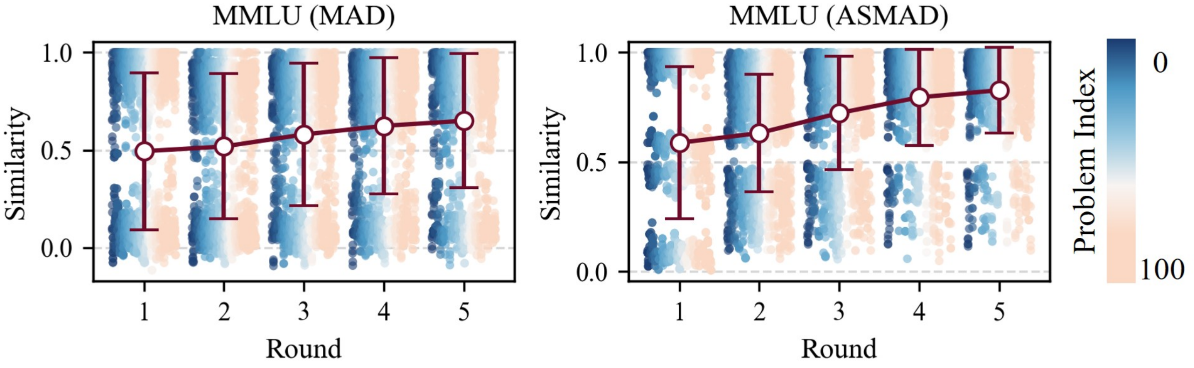

Figure 2: Similarity of agents vary toward consensus with increasing debate rounds where ASMAD provides better consensus rate (demonstrated in mean value and standard variance of similarity among agents) and speed

tains complete information exchange between all agents throughout the debate process, representing the most straightforward implementation of multi-agent debate. All experiments maintain consistent hyperparameters for fair comparison. Key evaluation metrics include: (1) Solution accuracy across different problem types; (2) Computational efficiency measured by token consumption.

3.4 Main Results

See Table 1, which presents our experimental results comparing ASMAD with the baseline MAD method. On GSM8K, although ASMAD shows a moderate accuracy drop of 8 percentage points compared to MAD (80% vs 88%), it achieves a substantial 65.8% reduction in token cost. For MMLU, ASMAD demonstrates superior performance by improving accuracy by 10 percentage points (from 49% to 59%) while simultaneously reducing token consumption by 67.3%. Figure 2 and 4 shows ASMAD accelerates consensus than MAD in each benchmark, with higher mean value and lower standard variance in similarity. The structured interaction framework intrinsic to ASMAD facilitates more comprehensive reasoning processes than mere "majority voting" mechanisms. The significant reduction in token consumption while maintaining or improving performance indicates that ASMAD successfully optimizes the debate process, eliminating redundant exchanges while preserving crucial reasoning steps. This efficiency gain suggests that adaptive structured debate mechanisms can effectively enhance reasoning capabilities, through interactions among even performance-limited quantized models.

Task	Method	ACC	Token Cost (k/task)	Cost Saving
GSM8K	MAD	**88%**	64.08	**-65.8%**
	ASMAD (Ours)	80%	**21.94**	
MMLU	MAD	49%	52.15	**-67.3%**
	ASMAD (Ours)	**59%**	**17.06**	

Table 1: Performance of MAD and ASMAD (our proposed method) across three tasks. Token cost is calculated as average of each topic debated. The results show that while ASMAD achieves comparable or improved accuracy compared to MAD, it significantly reduces token cost.

4 Conclusion

This work introduces ASMAD, a novel framework that synergizes sociophysical opinion dynamics with multi-agent debate systems through two key innovations: (1) probabilistic semantic-guided attention gates that dynamically regulate opinion visibility based on textual reasoning similarity, and (2) a hybrid paradigm integrating adaptive trust-boundary regulation with opinion synchronization mechanisms. By adaptively compute numerical semantic similarity and topology, ASMAD enables efficient consensus formation through structured sparse interactions. The framework establishes a principled bridge between opinion dynamics theory and practical LLM coordination, demonstrating that semantic-aware topology adaptation can simultaneously optimize communication efficiency and reasoning quality. Future work will explore extensions to larger-scale debates and automated parameter adaptation strategies.

Limitations

Our work, while demonstrating promising results, has several limitations worth acknowledging. Computational resource constraints led us to conduct experiments using relatively modest-sized language models (parameters < 10B) with 4-bit quantization. Though this choice enables practical deployment in resource-constrained settings, it inevitably faces an upper bound on the reasoning capabilities our agents can achieve. The potential of our framework when powered by more advanced language models remains to be explored. The effectiveness of our adaptive debate protocol currently hinges on several key hyperparameters, including confidence radius, growth rates, and similarity thresholds. Reinforcement learning approaches could potentially tune these parameters dynamically, adapting them to the specific context and demands of each debate scenario. Beyond these technical constraints, our initial validation of adaptive control mechanisms have promised in moderate-sized agent groups (6 agents). The dynamics and efficacy of our framework in larger debate clusters - particularly the interplay between maintaining diverse perspectives and achieving efficient consensus - represents an intriguing direction for future investigation. We also acknowledge potential risks associated with our work. While our adaptive debate framework aims to enhance reasoning capabilities, it could potentially amplify biases present in individual language models through the consensus formation process. The selective information exchange mechanism, though efficient, might inadvertently create echo chambers where agents reinforce each other's misconceptions. Additionally, the framework's ability to generate more convincing outputs through structured debate could be misused to produce more persuasive misinformation.

Ethical Considerations

In this research, Claude 3.5 Sonnet and Deepseek-R1 models are used as copilot, partially engaging in writing (sentence-level generations and grammar checking) and coding (fuzzing test and code-style polishing).

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.

Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2023. Chateval: Towards better llm-based evaluators through multi-agent debate. *arXiv preprint arXiv:2308.07201*.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.

Guillaume Deffuant, Frédéric Amblard, Gérard Weisbuch, and Thierry Faure. 2002. How can extremism prevail? a study based on the relative agreement interaction model. *Journal of artificial societies and social simulation*, 5(4).

Guillaume Deffuant, David Neau, Frederic Amblard, and Gérard Weisbuch. 2000. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3:87–98.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. 2023. Improving factuality and reasoning in language models through multiagent debate. *arXiv preprint arXiv:2305.14325*.

Seyed Rasoul Etesami and Tamer Başar. 2015. Game-theoretic analysis of the hegselmann-krause model for opinion dynamics in finite dimensions. *IEEE Transactions on Automatic Control*, 60(7):1886–1897.

Seyed Rasoul Etesami, Tamer Başar, Angelia Nedić, and Behrouz Touri. 2013. Termination time of multidimensional hegselmann-krause opinion dynamics. In *2013 American Control Conference*, pages 1255–1260. IEEE.

Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, et al. 2024. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *arXiv preprint arXiv:2406.12793*.

Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Nan Duan, and Weizhu Chen. 2023. Critic: Large language models can self-correct with tool-interactive critiquing. *arXiv preprint arXiv:2305.11738*.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.

Jennifer Hill, W Randolph Ford, and Ingrid G Farreras. 2015. Real conversations with artificial intelligence: A comparison between human–human online conversations and human–chatbot conversations. *Computers in human behavior*, 49:245–250.

Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. 2024. A survey on large language models for code generation. *arXiv preprint arXiv:2406.00515*.

Yunxuan Li, Yibing Du, Jiageng Zhang, Le Hou, Peter Grabowski, Yeqing Li, and Eugene Ie. 2024. Improving multi-agent debate with sparse communication topology. *arXiv preprint arXiv:2406.11776*.

Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2023. Encouraging divergent thinking in large language models through multi-agent debate. *arXiv preprint arXiv:2305.19118*.

Tongxuan Liu, Xingyu Wang, Weizhe Huang, Wenjiang Xu, Yuting Zeng, Lei Jiang, Hailong Yang, and Jing Li. 2024. Groupdebate: Enhancing the efficiency of multi-agent debate using group discussion. *arXiv preprint arXiv:2409.14051*.

Jan Lorenz. 2007. Continuous opinion dynamics under bounded confidence: A survey. *International Journal of Modern Physics C*, 18(12):1819–1838.

Luca Marconi and Federico Cecconi. 2020. Opinion dynamics and consensus formation in a deffuant model with extremists and moderates. *arXiv preprint arXiv:2010.01534*.

Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. 2023. A comprehensive overview of large language models. *arXiv preprint arXiv:2307.06435*.

Hegselmann Rainer and Ulrich Krause. 2002. Opinion dynamics and bounded confidence: Models, analysis and simulation. *Artif. Societies Social Simul.*, 5:1–33.

Qiushi Sun, Zhangyue Yin, Xiang Li, Zhiyong Wu, Xipeng Qiu, and Lingpeng Kong. 2023. Corex: Pushing the boundaries of complex reasoning through multi-model collaboration. *arXiv preprint arXiv:2310.00280*.

Behrouz Touri and Cedric Langbort. 2014. On endogenous random consensus and averaging dynamics. *IEEE Transactions on Control of Network Systems*, 1(3):241–248.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

Fatemeh Zarei, Yerali Gandica, and Luis Enrique Correa Rocha. 2023. Fast but multi-partisan: Bursts of communication increase opinion diversity in the temporal deffuant model. *arXiv preprint arXiv:2307.15614*.

Yuting Zeng, Weizhe Huang, Lei Jiang, Tongxuan Liu, Xitai Jin, Chen Tianying Tiana, Jing Li, and Xiaohua Xu. 2025. S^2-MAD: Breaking the token barrier to enhance multi-agent debate efficiency. *arXiv preprint arXiv:2502.04790*.

Jintian Zhang, Xin Xu, Ningyu Zhang, Ruibo Liu, Bryan Hooi, and Shumin Deng. 2023. Exploring collaboration mechanisms for LLM agents: A social psychology view. *arXiv preprint arXiv:2310.02124*.

YunHong Zhang, QiPeng Liu, and SiYing Zhang. 2017. Opinion formation with time-varying bounded confidence. *PloS one*, 12(3):e0172982.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.

Xinyu Zhu, Junjie Wang, Lin Zhang, Yuxiang Zhang, Ruyi Gan, Jiaxing Zhang, and Yujiu Yang. 2022. Solving math word problems via cooperative reasoning induced language models. *arXiv preprint arXiv:2210.16257*.

A Related Works

Topology in MAD Due to the diversity of human discussion strategies (Liang et al., 2023; Chan et al., 2023; Du et al., 2023), researchers adjust the visibility of interactions between agents and their historical records as well as among the agents themselves, by employing different multi-agent topologies, ultimately reducing token cost or enabling operation in resource-constrained environments (Li et al., 2024; Liu et al., 2024).

Regarding historical records, Du et al. (2023) process information from a centralized topology by summarizing agent outputs at the end of each round, whereas Sun et al. (2023) introduces a forgetting mechanism in which agents can only see the outputs from the previous round. In addition, Zhang et al. (2023) proposes a debate–reflection mechanism in which agents can only review their own past outputs during reflection.

Several studies focus on the topology of inter-agent information exchange. For instance, S-MAD

(Li et al., 2024) employs a sparse topology, limiting information exchange to adjacent agents. GroupDebate (GD) (Liu et al., 2024) adopts a hierarchical structure by clustering agents into smaller debate groups to exchange intermediate results. Furthermore, S^2-MAD (Zeng et al., 2025) utilizes a sparse topology based on grouping and a decision mechanism: agents initially generate independent opinions within groups, and only engage in information exchange within and between groups if a decision mechanism identifies differences in opinions.

Opinion Dynamics In the study of opinion dynamics, the Deffuant model and Hegselmann-Krause (HK) dynamics (Deffuant et al., 2000; Rainer and Krause, 2002) serve as foundational consensus models where a group of agents strive to reach the same objective. The Deffuant model posits that agents update their opinions based on a bounded confidence mechanism: two agents adjust their opinions only when their difference falls below a predefined threshold (Deffuant et al., 2000, 2002; Lorenz, 2007). This model has been extensively applied to investigate opinion convergence and polarization phenomena in social networks (Zhang et al., 2017; Marconi and Cecconi, 2020; Zarei et al., 2023).

The Hegselmann-Krause (HK) dynamics assumes that agents interact exclusively with peers whose opinions lie within their confidence bounds (Rainer and Krause, 2002; Etesami and Başar, 2015). In its synchronous variant, agents simultaneously update opinions by averaging those of neighbors within their confidence interval (Rainer and Krause, 2002; Etesami et al., 2013; Etesami and Başar, 2015), whereas the asynchronous version updates one agent at a time (Rainer and Krause, 2002; Touri and Langbort, 2014; Etesami and Başar, 2015). These consensus models provide critical frameworks for understanding opinion formation and evolution in social systems, particularly in analyzing how local interactions drive collective behaviors.

B Process Pipeline Figure

As is shown in Figure 3, the processing pipeline of proposed ASMAD is shown, which includes three stages. Out key innovation is in the Stage 2, which adopts a opinion dynamics based sparse topology generation mechanism. In detail, the sparse topology generation includes three sub-steps.

C Implementation Details

C.1 Structured Information Exchange

The computed weights determine not only the influence strength but also how information is presented to each agent. We implement a three-tier prompt structure:

$$P_{ij}^t = \begin{cases} \text{[Critical]} & \text{if } w_{ij}^t > 0.40 \\ \text{[Reference]} & \text{if } w_{ij}^t > 0.25 \\ \text{[Background]} & \text{if } w_{ij}^t > 0.10 \end{cases} \quad (3)$$

This structured presentation helps agents prioritize information based on computed influence weights, while maintaining the natural language interaction paradigm of LLMs.

C.2 Self-confidence Evolution

The self-confidence of each agent evolves according to:

$$w_{ij}^t = \text{clip}(\beta_0 + \beta_1 \left(\frac{t}{T}\right)(1 + \gamma\sigma_i^t) \cdot \text{sim}(s_i^t, s_j^t), w_{min}, w_{max})$$

where:

- β_0: base confidence level (0.3 in our implementation)
- β_1: growth rate (0.5)
- γ: stability influence factor (0.2)
- σ_i^t: agent's stability score

C.3 Hybrid Similarity Computation

We introduce a novel similarity measure that combines reasoning process similarity and answer agreement:

$$\text{sim}(i, j) = \lambda \cdot \cos(r_i, r_j) + (1 - \lambda) \cdot \mathbb{I}(c_i = c_j) \quad (4)$$

where:

- $\cos(r_i, r_j)$: cosine similarity between reasoning embeddings
- $\mathbb{I}(c_i = c_j)$: indicator function for answer agreement
- λ: balancing parameter (0.5)

C.4 Stability Mechanism

The stability score for agent i at round t is:

$$\sigma_i^t = 1 - \frac{\sum_{k=2}^{t} \mathbb{I}_{c_i^k \neq c_i^{k-1}}}{t-1} \quad (5)$$

This score influences both self-confidence and inter-agent weights through the mechanisms described above.

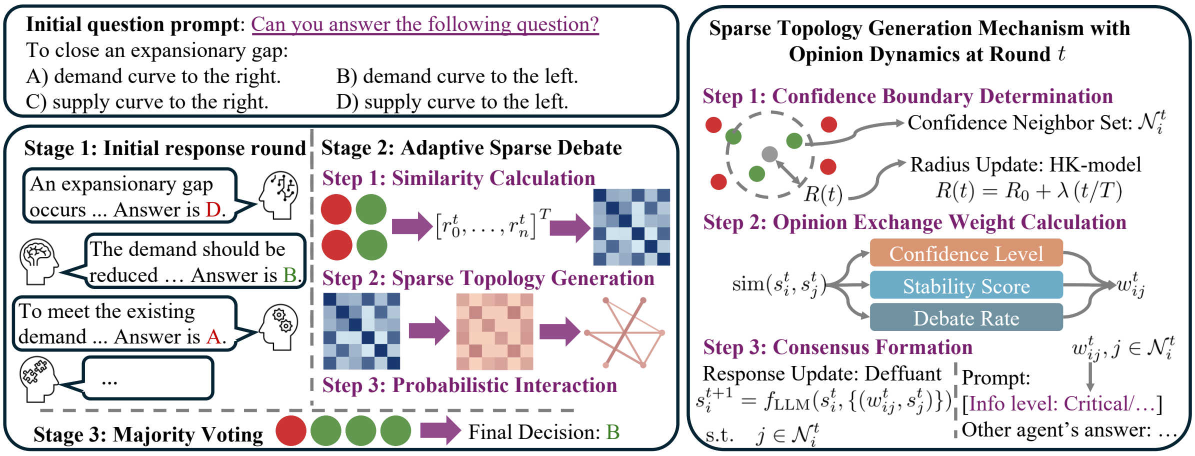

Figure 3: The process pipeline of ASMAD. Following S²-MAD (Zeng et al., 2025), we adopts three stages in total. In the first stage, all agents gives the initial response. In the second stage, with proposed sparse topology generation mechanism, the agents are organized to debeta with each other. In the last stage, the final decision is obtained via majority voting.

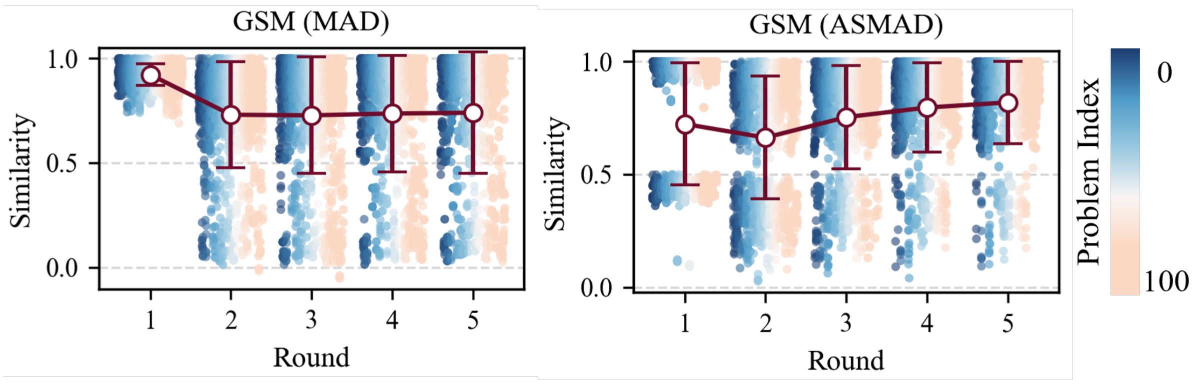

Figure 4: Similarity of agents vary toward consensus with increasing debate rounds where ASMAD provides better consensus rate (demonstrated in mean value and standard variance of similarity among agents) and speed

C.5 Row Normalization

To ensure balanced influence distribution, we apply row normalization to the weight matrix:

$$\hat{w}_{ij}^t = \frac{w_{ij}^t}{\sum_k w_{ik}^t} \tag{6}$$

This normalized weight matrix $\hat{W}^t$ governs the information flow and influence dynamics in each round of debate.

C.6 Consensus formation

ASMAD enables agents to arrive at consensus faster. Figure 2 and Figure 4 show the dynamics of agent opinions through metrics of similarity.