

CONTEXT LEARNING FOR MULTI-AGENT DISCUSSION

Xingyuan Hua¹, Sheng Yue^{2*}, Xinyi Li³, Yizhe Zhao¹, Jinrui Zhang¹, Ju Ren^{1,4}

¹Department of Computer Science and Technology, Tsinghua University,

²School of Cyber Science and Technology, Sun Yat-sen University,

³College of Computer Science, Northwest University,

⁴State Key Laboratory of Internet Architecture, Tsinghua University

ABSTRACT

Multi-Agent Discussion (MAD) has garnered increasing attention very recently, where multiple LLM instances collaboratively solve problems via structured discussion. However, we find that current MAD methods easily suffer from discussion inconsistency—LLMs fail to reach a coherent solution—due to the misalignment between their individual contexts. In this paper, we introduce a multi-LLM context learning method (M2CL) that learns a context generator for each agent, capable of dynamically generating context instructions per discussion round via automatic information organization and refinement. Specifically, inspired by our theoretical insights on the context instruction, M2CL trains the generators to control context coherence and output discrepancies via a carefully crafted self-adaptive mechanism. It enables LLMs to avoid premature convergence on “majority noise” and progressively reach the correct consensus. We evaluate M2CL on challenging tasks, including academic reasoning, embodied tasks, and mobile control. The results show that the performance of M2CL significantly surpasses existing methods by 20%–50%, while enjoying favorable transferability and computational efficiency.¹

1 INTRODUCTION

Large Language Models (LLMs) have demonstrated transformative impact across a large number of real-world domains, including education, healthcare, and scientific research, where an LLM instance is employed to automate the process of content generation, reasoning, or decision-making (Zheng et al., 2023; Imani et al., 2023; Zhang et al., 2024; Goyal et al., 2024; Wu et al., 2025; Yu et al., 2025). Yet, it has been recognized that single-LLM-built systems often struggle in the problems requiring complex multi-step reasoning or multi-tool using, such as complicated proof (Cobbe et al., 2021), large-scale code generation (Wang et al., 2025), and embodied agentic tasks (Ahn et al., 2022; Shen et al., 2024), because its single viewpoint of a problem easily limits the ability to explore multiple reasoning paths, leverage external tools effectively, or adapt to dynamic task requirements (Li et al., 2023; Smit et al., 2024).

Very recently, research has shifted towards *Multi-Agent Discussion* (MAD) (Smit et al., 2024), where multiple LLM instances collaboratively solve problems via structured discussions (Du et al., 2023; Liu et al., 2024). In a typical MAD framework, each LLM instance is pre-assigned a set of crafted *contexts* that represent diverse solution perspectives of the problem to be solved. Equipped with these distinct context instructions, LLMs continue to discuss with each other for a solution consensus (Park et al., 2023; Shanahan et al., 2023; Wei et al., 2023; Lu et al., 2024; Liu et al., 2024). Such ‘society-of-mind’ paradigms are expected to improve reasoning accuracy by enhancing creativity and expanding the search space for possible solutions, and have shown great potential across various complex tasks, including software engineering (Gu, 2023) and scientific discovery (Sprueill et al., 2024).

Albeit achieving improved performance over single-LLM settings, we find that current multi-agent collaboration approaches typically suffer from discussion inconsistency, that is, the majority of LLM instances fail to reach an agreement on a coherent solution (as showcased in Fig. 2), easily making the collaborative decision dominated by noise rather than principled reasoning. The underlying reason

*Correspondence to: yuesh5@mail.sysu.edu.cn

¹Code is available at <https://github.com/HansenHua/M2CL-ICLR26>.

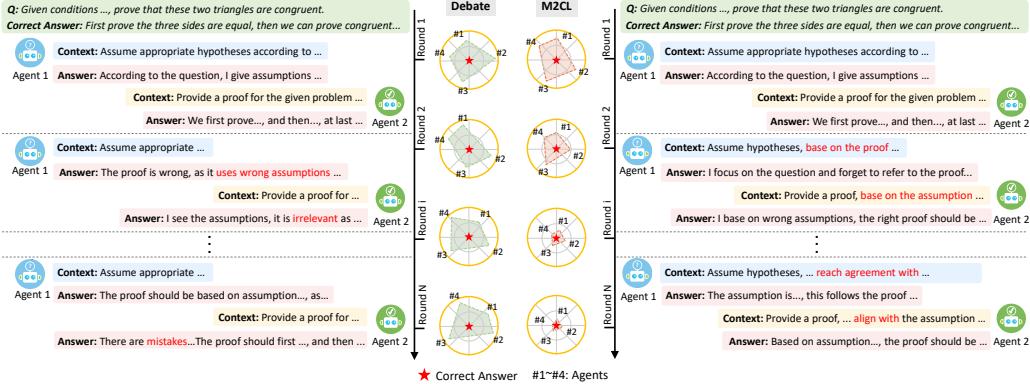


Figure 1: An illustration of context misalignment of an existing method (Debate Du et al. (2023)) on a multi-step proof task. Pre-assigned context instructions (in the blue and yellow boxes of the left part) provide insufficient guidance on information fusion, leading to conflict in reasoning.

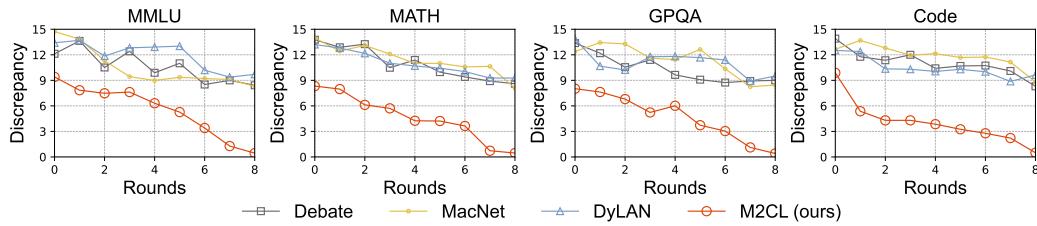


Figure 2: The discrepancy between the answers of participating LLM instances. The discrepancy is characterized by the maximum distance between participating LLMs' output embeddings.

primarily lies in *context misalignment* between LLMs. On one hand, the pre-assigned contexts or role-based instructions lack a nuanced understanding of the task; they are often rigid, incomplete or biased, which would misguide the reasoning of individual LLMs (Jang et al., 2025). On the other hand, these contexts often fall short in effective fusion of information exchanged among LLMs, and thus hardly steer the discussion towards coherent solutions. As illustrated in Fig. 1, for a multi-step mathematical proof task, one LLM agent may correctly derive an intermediate result, yet another LLM—despite receiving this result in the extended context—does not effectively incorporate it into its own reasoning chain. Since the context instruction does not explicitly enforce leveraging conclusions from other LLMs, other LLMs may redundantly re-derive the step or even give an inconsistent argument (see Appendix H for full results).

This paper aims to answer: “*how can we obtain the contexts for MAD that can continually guide multi-LLM discussion towards a correct consensus?*” A straightforward solution is to manually adapt instructions in contexts as the discussion progresses. It, however, is labor-intensive and requires extra expert knowledge, rendering it impractical for complex tasks or large-scale collaboration. Instead, a more reasonable approach is to develop context learning mechanisms that enable evolving the context instructions based on the intermediate discussion results. Although promising, it is highly challenging to evaluate the contribution of LLMs’ contexts to the final solution and control the coherence among inter- and intra-LLM outputs.

To tackle these challenges, we propose a *multi-LLM context learning* (M2CL) method for efficient MAD, which learns a context generator for each agent, capable of dynamically generating context instructions per discussion round via automatic information organization and refinement. First, we characterize the impact of initial and evolving contexts on the discussion performance. Building upon the analytical insights, we train the generators to control context coherence and output discrepancies. To strike the right tradeoff therein, we devise a self-adaptive balancing mechanism, enabling LLMs to progressively align on correct consensus while avoiding premature convergence on “majority noise”. Further, we develop a lightweight context initialization approach, which tends to assign LLMs with diverse initial instructions that are approximately orthogonal in the latent space, enabling sufficient coverage of complementary solution perspectives.

We systematically evaluate the proposed method across 9 challenging benchmarks, including LLM reasoning, embodied agentic tasks, and mobile GUI control. The results demonstrate that M2CL significantly enhances the MAD performance under various numbers of participating LLMs, consistently outperforming existing methods by **20%–50%**, particularly in complex GUI control tasks. M2CL enjoys a more favorable MAD “scaling law” and exhibits great efficiency, where a runtime overhead of at most 10% suffices to achieve more than 20% performance gains. Further, we find that the learned context generators can be migrated to different LLM architectures with consistent performance improvement.

Our main contributions are summarized as follows:

- We propose M2CL, a principled multi-LLM context learning method that learns a context generator for each agent, capable of dynamically generating context instructions per discussion round.
- We devise a lightweight context initialization approach that can assign LLMs with diverse initial instructions, thereby enabling sufficient coverage of complementary solution perspectives.
- We systematically evaluate the proposed method across a range of challenging benchmarks and corroborate the efficacy and efficiency of the proposed method.

2 RELATED WORK

Multi-agent framework. Multi-agent frameworks have been proposed to enhance the reasoning capabilities of single-LLM prompting methods. Several works utilize two LLMs to iteratively generate and evaluate to refine the final answer (Raman et al., 2022; Yao et al., 2023; Paul et al., 2023; Madaan et al., 2024). However, scaling them to a larger number of LLMs to make full use of the wisdom of crowd is challenging due to the pairwise dependency of generation and evaluation. To incorporate more LLMs, several studies have explored another multi-agent framework in which each LLM has access to the history of all preceding LLMs’ responses (Chan et al., 2023; Du et al., 2023; Liang et al., 2023; Smit et al., 2024; Liu et al., 2024; Zhuge et al., 2024). Most existing methods remain constrained by manually defined inter-LLM topologies or workflows. More recently, researches have attempted to overcome this limitation by automatically optimizing workflows (Zhuge et al., 2024; Liu et al., 2024; Zhang et al., 2025b). However, such evolution is often one-shot. For instance, Zhuge et al. (2024) describe language agent systems as optimizable computational graphs and enables automatic improvements of LLM prompts and inter-LLM orchestration. Zhang et al. (2025b) employs Monte Carlo Tree Search to construct complex Multi-LLM systems tailored to a specific task domain. Liu et al. (2024) propose a multi-LLM collaboration method which constructs dynamic communication structure by scoring other LLMs. While these approaches allow for scaling the number of LLMs, designing appropriate contexts for them to collaboratively solve problems remains difficult, as single-viewpoint context instructions fall short in guidance on how to organize and refine information across agents.

Context learning. Context learning has recently attracted significant attention, as it allows LLMs to adjust their behavior at inference time by modifying the input context, without requiring any gradient-based training Von Oswald et al. (2023); Todd et al. (2024); Li et al. (2024a). One line of work focuses on context selection, which identifies the most relevant examples or information to include in the context Zhang et al. (2022); Lu et al. (2023); Xiong et al. (2024); Purohit et al. (2025). To eliminate manual selection, Lu et al. (2023) leverage reinforcement learning to learn a context selection policy. Xiong et al. (2024) query LLM to obtain knowledge and then query a retriever to obtain the final context. Beyond selection, some recent methods have proposed to generate or evolve context (Zhuge et al., 2024; Li et al., 2024b; Zhang et al., 2025a). Madaan et al. (2024) leverage feedback generated by LLMs as extra prompt and iteratively incorporate it into revised drafts, aiming to enhance the coherence of the generated text. Pandita et al. (2025) dynamically refines prompts during inference using textual feedback from prior outputs, aiming to improve contextual alignment and generate more consistent responses. Albeit with promising results, these methods are grounded in single-LLM formulations and struggle with inter-LLM inconsistency during multi-agent discussion, where it is crucial to guide LLMs to make full use of others’ intermediate results.

3 PRELIMINARIES

In this section, we provide necessary backgrounds and definitions of our investigated problem.

Context learning. To formally define context learning, we begin with the standard probabilistic model of an autoregressive LLM. An autoregressive model, parameterized by ϕ , generates an output sequence $X = (x_1, x_2, \dots, x_L)$ given an input context C by maximizing the conditional probability:

$$P_\phi(X|C) = \prod_{l=1}^L P_\phi(x_l|x_{<l}, C). \quad (1)$$

Historically, in the paradigm of prompt engineering, the context C was treated as a static string of text. This view is insufficient for complex agentic tasks which leverage dynamic, structured, and multifaceted information stream. To address this, context learning is redefined as a dynamically organized collection of information, denoted as $\{c_1, c_2, \dots, c_n\}$. Each component c_k can be instructions, external knowledge (Lewis et al., 2020), available external tools (Qin et al., 2023), and memory (Zhang et al., 2025a). These components are sourced, filtered, and formatted by a set of functions f_k , and orchestrated into a coherent representation by a high-level assembly function \mathcal{A} :

$$C = \mathcal{A}(f_1(c_1), f_2(c_2), \dots, f_n(c_n)). \quad (2)$$

Multi-LLM context learning. MAD involves a set of N LLM instances collaboratively solving a task via multi-round inter-LLM discussion (debate). Each LLM i is endowed with an evolving instruction context in terms of the task description, available external tools and current knowledge aggregation. At the t -th round, we use the concatenation function as the assembly function \mathcal{A} and construct the context with three components: (i) the task goal $P \in \mathbb{R}^{d_{model} \times n}$, (ii) the concatenation of responses from all other LLMs in the previous round $\bar{X}_i^{t-1} = [X_j^{t-1}]_{j \neq i} \in \mathbb{R}^{d_{model} \times n}$, where X_j^{t-1} is the response of LLM j at round $t-1$, and (iii) the current instruction context $I_i^t \in \mathbb{R}^{d_{model} \times n}$ (d_{model} refers to the dimension of embedding). The task goal represents the ultimate objective of collaboration and remains invariant across rounds. The second component serves as a dynamic memory that incorporates cross-LLM interaction history. For the third component, in contrast to existing efforts relying on static preassigned roles, we employ an instruction generator \mathcal{G} , parameterized by θ_i , to adaptively refine it into a per-step instruction I_i^t , conditioned on the task goal P , the initial instruction I_i^b , and the concatenated response \bar{X}_i^{t-1} :

$$I_i^t = \mathcal{G}_{\theta_i}([P; I_i^b; \bar{X}_i^{t-1}]). \quad (3)$$

Given context $C_i^t = [I_i^t, \bar{X}_i^{t-1}, P]$, each LLM i , parameterized by ϕ_i , generate its response by:

$$X_i^t = \arg \max_X P_{\phi_i}(X|C_i^t). \quad (4)$$

After T rounds of interaction, the final result is obtained by a majority vote on the LLMs' outputs generated in the final round.

4 MOTIVATION

In this section, we investigate the quantified impact of the contexts on the MAD performance. First, we introduce the formulation of attention activation, denoted as $a(\cdot)$, as follows (Vaswani, 2017):

$$a(C_i^t) \doteq W_V[I_i^t; \bar{X}_i^{t-1}; P] \text{softmax} \left(\frac{(W_K[I_i^t; \bar{X}_i^{t-1}; P])^T W_Q P}{\sqrt{d}} \right) \quad (5)$$

where \sqrt{d} is a scaling factor. W_Q , W_K , and W_V denote the parameter weight matrices in the attention mechanism. $a(C_i^t) \in \mathbb{R}^{d_{model} \times n}$ is a matrix with the same shape as $[P; X; I]$

From this, we have the following theorem characterizing the total distance between the activations of the correct answer a_c and that induced by context C_j^t . Here, we utilize activation distance instead of token embedding distance as attention activation captures deep representational similarity learned through the model's internal reasoning process, making it more robust to superficial linguistic variations.

Theorem 4.1. Assume that the attention activation function is L_a -smooth and define the weight vector as $\omega \doteq [\omega_1, \omega_2, \dots, \omega_N]$, with the initial context $C_i^b \doteq [I_i^b; P]$. Then, the following fact holds:

$$\begin{aligned} \sum_{i=1}^N \|a_c - a(C_i^t)\| &\leq \sum_{i=1}^N \left(\sum_{j=1}^N \|a(C_i^t) - a(C_j^t)\| + (N+1)L_a \|C_i^t - C_i^b\| \right) \\ &\quad + N \min_{\omega} \|a_c - \sum_{i=1}^N \omega_i a(C_i^b)\|. \end{aligned} \quad (6)$$

Proof. For a detailed proof, please refer to Appendix B. \square

Theorem 4.1 corroborates the necessity of multi-LLM context learning, consisting of both initialization and evolution. The first term, $\|a(C_i^t) - a(C_j^t)\| + (N+1)L_a \|C_i^t - C_i^b\|$, captures the divergence among LLMs' activations as well as the deviation from their initial contexts. This indicates contexts must be continuously evolved to reduce inter-LLM discrepancies while keeping coherent, promoting consistency of reasoning chains. The second term, $\min_{\omega} \|a_c - \sum_{i=1}^N \omega_i a(C_i^b)\|$, depends solely on the initial contexts I_i^b . It implies that orthogonality among initial activations provides a comprehensive basis, allowing the contexts to approximate the correct activation more effectively. This highlights the importance of diverse solution perspectives at the initialization stage. Together, the result motivates a two-stage design: initializing contexts to ensure diverse solution perspectives and evolving intermediate contexts to drive a consensus across LLMs.

5 MULTI-LLM CONTEXT LEARNING

In this section, we introduce a multi-LLM context learning method that can select appropriate contexts for each LLM and dynamically adapt the context according to the evolving task completion status.

5.1 CONTEXT INITIALIZATION

The context initialization serves as the foundation of the entire multi-agent interaction process, as it frames the capability scope of LLMs and fundamentally shapes the chain of subsequent information exchange. Therefore, we propose to select initial contexts from a predefined pool $\{I_i^b\}_{i=1}^M$ containing prompts with diverse perspectives.

Motivated by Theorem 4.1, we provide the following context initialization mechanism:

$$\mathbf{I}^b = \{I_1^b, \dots, I_N^b\} = \arg \min_{\mathbf{I}^b} \left\{ \min_{\omega} \left\| \sum_{i=1}^N \omega_i a([I_i^b; P] - a_c) \right\| \right\}. \quad (7)$$

Eq. (7) identifies a subset of contexts \mathbf{I}^b whose activation best reconstruct the target activation a_c . As the dimension of the activation matrix $a([I_i^b; P]) \in \mathbb{R}^{d_{model} \times n}$ is far larger than the number of selected contexts N , a set of matrices that aims to best reconstruct the correct activation a_c naturally tends toward forming a set of basis-like directions. The resulting near-orthogonal activations form a compact basis, ensuring each context contributes unique, non-overlapping information for subsequent discussion.

However, Eq. (7) is impractical since the correct activation a_c is not accessible during initialization. Inspired by prior work (Yang et al., 2024b; 2025), we project the activation into a latent space with function f :

$$\mathbf{I}^b = \{I_1^b, \dots, I_N^b\} = \arg \min_{\mathbf{I}^b} \left\{ \min_{\omega} \left\| \sum_{i=1}^N \omega_i f(a([I_i^b; P])) - f(a_c) \right\| \right\}, \quad (8)$$

Then, we use the problem space as the selected contexts are solely dependent on the problem. This projection preserves the original orthogonality properties of activations, ensuring that diverse perspectives remain distinguishable in other spaces. Therefore, we reformulate the initialization

mechanism:

$$\mathbf{I}^b = \{I_1^b, \dots, I_N^b\} = \arg \min_{\mathbf{I}^b} \left\{ \min_{\omega} \left\| \sum_{i=1}^N \omega_i f(a([I_i^b; P])) - v_P \right\| \right\}, \quad (9)$$

where v_P denotes the sentence vector of the question P . To obtain the projection $f(\cdot)$ (parameterized by ϕ_f), we utilize the activation of inputting the answer A with the question P as the correct activation $a_c = a([A; P])$ and provide the loss function as:

$$L(\phi_f) = \|v_P - f(a([A; P]))\|. \quad (10)$$

However, directly leveraging Eq. (9) to initialize context is computational costly as it requires all the context activations in the context pool. Therefore, we distill a lightweight $\mathcal{F}(\cdot)$ (parameterized by ϕ_F) which directly projects the initial context into the problem space through the following loss function:

$$L(\phi_F) = \|\mathcal{F}([I_i^b; P]) - f(a([I_i^b; P]))\|. \quad (11)$$

Then, we provide the final context initialization as:

$$\mathbf{I}^b = \{I_1^b, \dots, I_N^b\} = \arg \min_{\mathbf{I}^b} \left\{ \min_{\omega} \left\| \sum_{i=1}^N \omega_i \mathcal{F}([I_i^b; P]) - v_P \right\| \right\}. \quad (12)$$

Eq. (12) provides a computational efficient context initialization (as illustrated in Fig. 23) that aligns with Eq. (7) to encourage orthogonality of selected context activations, thereby providing diverse reasoning perspectives and expanding the search space for solutions.

5.2 CONTEXT EVOLUTION

As discussed in Section 5.1, context initialization aims to select contexts with diverse perspectives to solve the problem. Nevertheless, these individual contexts address the problem from a single viewpoint while lacking explicit instructions on how to incorporate perspectives provided by other LLMs. Directly using them often results in inconsistent reasoning across LLMs due to misaligned inter-LLM guidance. To address this, we iteratively refine the instruction to integrate collaborative instructions that guide the LLMs more effectively toward generating the desired solution.

5.2.1 EVALUATE THE CONTRIBUTION OF CONTEXTS

To achieve efficient context evolution, it is crucial to quantify the contribution of different contexts (refer to a utility function). A natural solution is to assign utilities only at the final round. However, this leads to inefficient and unstable training process due to its sparsity Liu et al. (2025). On the other hand, current methods typically use correctness as the only criterion (Zelikman et al., 2022; Rafailov et al., 2023; Guo et al., 2025). Albeit with promising performance in single-LLM question answering, this simple criterion ignores the complex dependency among LLMs with different contexts in MAD. As a result, an LLM that provides crucial insights which enable others to reach the correct answer may still receive a low utility, simply because it fails to produce the correct final answer itself. To tackle these challenges, we introduce a novel round-wise criterion to evaluate the contribution of C_i^t in multi-agent interaction:

$$\max_{j \in [N]} \left\{ -\alpha \|C_i^t - C_i^b\| - \|a(C_i^t) - a(C_j^t)\| \right\}, \quad (13)$$

where $\alpha \geq 0$ is a weighting parameter and $\|a(C_i^t) - a(C_j^t)\|$ represents the activation difference between contexts.

Remarks. Of note, if removing the second term of Eq. (13), the context generation reduces to that of prompt engineering: $\min_{I_{1:N}^t} \sum_{i \in [N]} \alpha \|C_i^t - C_i^b\|$, which utilize fixed instructions in contexts to LLMs throughout the discussion (Liu et al., 2024; Lu et al., 2024). As mentioned in Section 1, this easily results in inconsistency across LLMs. Thus, we introduce $\|a(C_i^t) - a(C_j^t)\|$ into the objective to encourage LLMs to remain aligned with each other.

This design addresses the two aforementioned challenges. First, this round-wise criterion provides denser feedback, mitigating the inefficiency and instability caused by sparse final-round utilities.

Second, the activation-based alignment term captures inter-LLM dependencies, thereby encouraging them to offer guidance on how to integrate others' responses.

Applying Eq. (13) as the criterion for LLM i requires the contexts of all other LLMs. However, since these contexts are simultaneously being optimized, a bias inevitably arises as each LLM is updated based on stale snapshots of others' contexts. Such outdated information introduces a biased utility, which in turn hinders consistency and degrades the quality of the learned contexts. To overcome this issue, we decouple the inter-LLM dependencies by designing an alternative per-LLM criterion:

$$-\alpha \|C_i^t - C_i^b\| - \|a([I_i^t, P]) - a([X_i^{t-1}, P])\|. \quad (14)$$

In Lemma C.1, we can prove that the summation of Eq. (14) over all LLMs serves as an upper bound of the summation of Eq. (13) over all LLMs. Intuitively, this criterion conveys a clear meaning: the first term, $-\alpha \|C_i^t - C_i^b\|$, preserves the fundamental problem-solving capability endowed by the initialization, while the second term, $-\|a([I_i^t, P]) - a([X_i^{t-1}, P])\|$, enforces consistency between an LLM's current instruction and its own previous response. This local consistency constraint implicitly aligns all LLMs, because when every LLM evolves its context in a temporally coherent way, the divergence among them is gradually reduced. Hence, the collective effect of all LLMs following this rule is that their contexts evolve coherently across rounds, progressively shrinking activation differences and ultimately driving the MAD toward consistency in their final answers.

To intuitively show the tightness of the bound, consider the case where $I_i^t = X_i^{t-1}$ holds, i.e., the instruction of each LLM in the next round equals its response in the previous round. That means $a(I_i^t) = a(X_i^{t-1})$. Given that all LLMs receive the same context, the second term of both Eq. (13) and Eq. (14) becomes zero. In this case, the two formulations coincide.

5.2.2 MULTI-ROUND CONTEXT EVOLVING

We slightly abuse notation by denoting $a(I_i^t) = a([I_i^t; P])$ and $a(X_i^{t-1}) = a([X_i^{t-1}; P])$, and provide the multi-round context evolution objective by accumulating contributions as follows:

$$\sum_{t=1}^T \max_{I_i^t} \left\{ -\alpha \|C_i^t - C_i^b\| - \|a(I_i^t) - a(X_i^{t-1})\| \right\}. \quad (15)$$

However, selecting the optimal value for the weight α is non-trivial, as it must be carefully tuned for the evolving discussion process. To alleviate the need for manual tuning, we recast Eq. (15) into a constrained optimization problem, where the context adjustment $\|I_i^t - I_i^b\|$ is treated as a constraint:

$$\begin{aligned} & \min_{I_i^1} \left(\|a(I_i^1) - a(X_i^0)\| + \min_{I_i^2} \left(\|a(I_i^2) - a(X_i^1)\| + \cdots + \min_{I_i^T} \|a(I_i^T) - a(X_i^{T-1})\| \right) \right) \\ & \text{s.t. } \|C_i^t - C_i^b\| \leq \beta, \quad \forall t, i. \end{aligned} \quad (16)$$

According to the detailed derivation in Appendix E, we prove that taking the dual of this problem recovers Eq. (15) and yields an auxiliary update for the dual variable α . Consequently, the optimization of Problem (16) can be implemented using an approximate dual gradient descent procedure, alternating gradient updates of $L(\theta_i)$ and $L(\alpha_i)$.

$$\begin{aligned} L(\theta_i) &= \|a(\mathcal{G}_{\theta_i}(P, I_i^b, \bar{X}_i^{t-1})) - a(X_i^{t-1})\| + \alpha \|C_i^t - C_i^b\| \\ L(\alpha_i) &= \alpha_i (\beta - \|\mathcal{G}_{\theta_i}(P, I_i^b, \bar{X}_i^{t-1}) - C_i^b\|). \end{aligned} \quad (17)$$

At the beginning of the discussion, the initial contexts are intentionally diverse to encourage multi-perspective reasoning of the problem. At this stage, when the answers differ greatly, α decreases rapidly, thereby weakening the constraint on the distance between the generated contexts and their initial contexts. This guides the generation of contexts towards promoting faster convergence on a unified solution. As the discussion progresses and the LLMs gradually reach agreement, α will be kept at a certain level. This adjustment prevents premature consensus and supports a richer, more comprehensive final solution by shifting the contexts' focus toward preserving multiple perspectives and exploring nuanced differences.

Overall, we name our proposed algorithm M2CL, with its pseudocode detailed in Appendix F.4.

Table 1: Accuracy (%) on different datasets. The number of LLMs is 4 for all dataset. We exhibit the performance advantage with BoN and highlight the best result.

Model	Method	MMLU	MATH	GPQA	Code	ALFWorld	SciWorld	GAIA	PDDL
Qwen-7B	Single	61.2 _{±13.0}	12.9 _{±12.0}	20.2 _{±16.2}	51.2 _{±11.3}	23.8 _{±7.7}	25.2 _{±10.1}	15.6 _{±5.5}	21.0 _{±5.3}
	BoN	74.2 _{±0.0}	24.9 _{±0.0}	36.4 _{±0.0}	62.5 _{±0.0}	31.5 _{±0.0}	35.3 _{±0.0}	21.1 _{±0.0}	26.3 _{±0.0}
	Debate	71.1 _{±3.1}	19.9 _{±5.0}	28.7 _{±7.7}	60.0 _{±2.5}	30.6 _{±0.9}	32.2 _{±3.1}	21.0 _{±0.1}	24.8 _{±1.5}
	DyLAN	74.3 _{±0.1}	26.7 _{±1.8}	35.4 _{±1.0}	63.4 _{±0.9}	29.8 _{±1.7}	29.4 _{±5.9}	18.4 _{±2.7}	23.4 _{±2.9}
	GPTSwarm	76.3 _{±2.1}	26.2 _{±1.3}	35.6 _{±0.8}	62.7 _{±0.2}	29.9 _{±1.6}	30.8 _{±4.5}	20.6 _{±0.5}	24.5 _{±1.8}
	MacNet	71.5 _{±2.7}	21.4 _{±3.5}	30.8 _{±5.6}	59.3 _{±3.2}	33.6 _{±2.1}	37.0 _{±1.7}	21.5 _{±0.4}	29.5 _{±3.2}
	M2CL (ours)	92.5 _{±18.3}	47.8 _{±22.9}	66.1 _{±29.7}	80.3 _{±17.8}	39.9 _{±8.4}	45.3 _{±10.0}	33.6 _{±12.5}	34.7 _{±8.4}
Qwen-14B	Single	67.2 _{±12.5}	21.6 _{±6.2}	21.2 _{±11.6}	56.7 _{±12.7}	30.1 _{±7.4}	31.3 _{±10.3}	18.7 _{±6.9}	25.7 _{±2.8}
	BoN	79.7 _{±0.0}	27.8 _{±0.0}	32.8 _{±0.0}	69.4 _{±0.0}	37.5 _{±0.0}	41.6 _{±0.0}	25.6 _{±0.0}	28.5 _{±0.0}
	Debate	77.2 _{±2.5}	27.4 _{±0.4}	30.3 _{±2.5}	66.9 _{±2.5}	36.8 _{±0.7}	39.4 _{±2.2}	26.4 _{±0.8}	30.0 _{±1.5}
	DyLAN	86.8 _{±7.1}	31.2 _{±3.4}	39.4 _{±6.6}	76.6 _{±7.2}	34.8 _{±2.7}	36.0 _{±5.6}	22.8 _{±2.8}	28.0 _{±0.5}
	GPTSwarm	87.0 _{±7.3}	31.3 _{±3.5}	38.7 _{±5.9}	75.9 _{±6.5}	34.9 _{±2.6}	35.9 _{±5.7}	25.0 _{±0.6}	28.4 _{±0.1}
	MacNet	78.9 _{±0.8}	28.7 _{±0.9}	31.9 _{±0.9}	70.4 _{±1.0}	39.2 _{±1.7}	46.0 _{±4.4}	32.2 _{±6.6}	32.7 _{±4.2}
	M2CL (ours)	93.7 _{±14.0}	51.7 _{±23.9}	66.2 _{±33.4}	91.1 _{±21.7}	48.2 _{±10.7}	56.1 _{±14.5}	42.0 _{±16.4}	43.0 _{±14.5}
Qwen-72B	Single	72.5 _{±11.7}	31.6 _{±19.4}	34.9 _{±11.0}	59.1 _{±13.1}	48.2 _{±9.3}	50.4 _{±11.7}	31.1 _{±10.1}	41.0 _{±10.0}
	BoN	84.2 _{±0.0}	51.0 _{±0.0}	45.9 _{±0.0}	72.2 _{±0.0}	57.5 _{±0.0}	62.1 _{±0.0}	41.2 _{±0.0}	51.0 _{±0.0}
	Debate	82.7 _{±1.5}	48.4 _{±2.6}	43.4 _{±2.5}	69.1 _{±3.1}	60.4 _{±2.9}	65.2 _{±3.1}	42.9 _{±1.7}	49.1 _{±1.9}
	DyLAN	91.5 _{±7.3}	63.1 _{±12.1}	51.6 _{±5.7}	80.4 _{±8.2}	55.1 _{±2.4}	58.0 _{±4.1}	40.4 _{±0.8}	45.5 _{±5.5}
	GPTSwarm	91.5 _{±7.3}	64.7 _{±13.7}	52.4 _{±6.5}	79.6 _{±7.4}	56.8 _{±0.7}	60.2 _{±1.9}	40.3 _{±0.9}	49.7 _{±1.3}
	MacNet	83.8 _{±0.4}	52.9 _{±1.9}	46.2 _{±0.3}	70.5 _{±1.7}	61.8 _{±4.3}	68.4 _{±6.3}	46.4 _{±5.2}	53.7 _{±2.7}
	M2CL (ours)	95.1 _{±10.9}	72.5 _{±21.5}	78.9 _{±33.0}	90.7 _{±18.5}	79.0 _{±21.5}	88.9 _{±26.8}	67.2 _{±26.0}	70.5 _{±19.5}

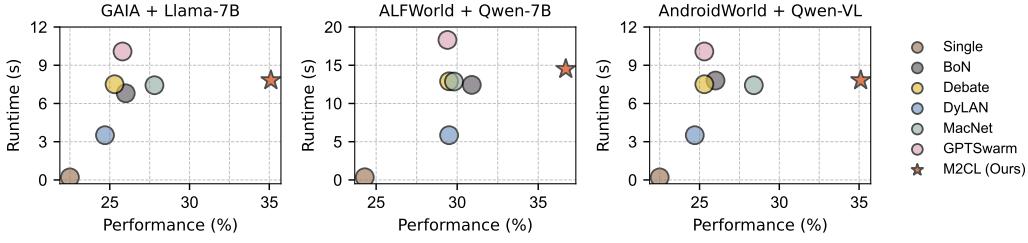


Figure 3: Performance versus runtime under different settings. Circles closer to the lower-left corner indicate higher efficiency.

6 EXPERIMENT

In this section, we conduct experiments to evaluate the performance of M2CL by answering the following research questions:

- Q1. How does M2CL perform compared to existing methods across various benchmarks, especially in complex agentic tasks?
- Q2. How does the performance of MAD scale with the number of LLMs?
- Q3. How is the performance affected by factors such as context constraint and components such as context initialization and context evolution?
- Q4. How do contexts promote consensus, and are they transferable to other models?

6.1 EXPERIMENTAL SETUP

Dataset. We run experiments with 3 domains including 9 datasets: **1) LLM reasoning**, including MMLU (Hendrycks et al., 2021), MATH (Hendrycks et al., 2021), and GPQA (Rein et al., 2023), HumanEval (Chen et al., 2021). **2) Embodied Agentic**, including ALFWorld Shridhar et al. (2021), SciWorld Wang et al. (2022), GAIA (Mialon et al., 2024), and PDDL Chang et al. (2024). **3) Mobile GUI** AndroidWorld Rawles et al. (2025). Details on datasets can be found in Appendix F.1.

Baselines. We evaluate our method against six strong baseline methods: **1) Single execution**, querying a single LLM to solve the task. **2) Best-of-N**, querying a single LLM N times and sampling the most correct answer. **3) Debate** (Du et al., 2023), a multi-agent framework where LLMs discuss their responses and reasoning processes over multiple rounds. **4) DyLAN** (Liu et al., 2024), a multi-agent framework where LLMs score each other and collaborate dynamically. **5) GPTSwarm** (Zhuge et al., 2024), a multi-agent framework that refines LLM prompts and improves

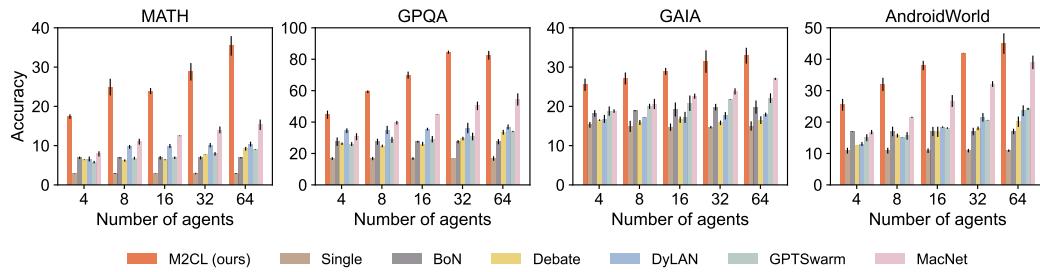


Figure 4: Performance of varying the numbers of LLMs. Uncertainty intervals depict standard deviation over three seeds.

LLM orchestration by changing their connectivity. 6) MacNet (Jiang et al., 2023), a recent multi-LLM framework where LLMs are invoked between LLM interactions to provide actionable instructions to the next LLM based on the previous LLM’s outputs.

Reproducibility. All details of our experiments are provided in the appendices in terms of the tasks, network architectures, hyperparameters, etc. We conduct experiment on two series of LLM (Llama-2 (Touvron et al., 2023) and Qwen-2.5 (Yang et al., 2024a)) and three model sizes each for LLM reasoning, and Qwen2.5-VL (Bai et al., 2025) for GUI reasoning. All the experiments are run on Ubuntu 22.04.4 LTS with 8 NVIDIA H800 GPUs.

6.2 EXPERIMENTAL RESULTS

Comparative results. To answer the first question, we evaluate M2CL’s performance across all datasets, with varying base models and number of LLMs (ranging from 4 to 64). We select Qwen series models and 4 LLMs participating in Table 1 and provide full results in Appendices G.2 and G.3. We find M2CL consistently outperforms baselines in all 9 datasets, often by a significant margin in terms of performance. Of note, BoN outperforms most of the baselines especially in complex multi-round agentic tasks, revealing the drawback of fixed contexts, which, despite expanding the exploration space, do not converge and thus hinder LLMs from achieving true cooperative reasoning. In contrast, M2CL can adapt contexts and enhance the relevance of responses and questions while ensuring creativity, indicating that M2CL well avoids that LLMs with different contexts easily influence each other and successfully brings LLMs into cooperation by reaching a consensus through discussion.

In addition, we evaluate the efficiency of M2CL by visualizing the performance versus runtime. As shown in Fig. 3, M2CL consistently delivers the highest performance improvement (more than 20%) while maintaining a modest increase in runtime (less than 10%). This clearly demonstrates the efficiency of M2CL due to its lightweight context generator.

Multi-agent scaling law. To answer the second question, we run experiments with varying numbers of LLMs (ranging from 4 to 64). The data and parameter setup adhere to that of Table 2. We present selected results in Fig. 4 and full results in Tables 4 to 8 and Figs. 5 to 10 of Appendix G.2. Scaling our method reveals a more efficient scaling law as the performance grows logarithmically before saturation and improves faster than baselines. We speculate this arises because collaborative instructions in our generated contexts enable genuine inter-LLM cooperation, thereby unleashing the multidimensional reasoning capabilities of the MAD.

Context constraint. To answer the third question, we vary the context constraint β from 0 to 10 and run experiments across all datasets. Full results are shown in Figs. 12 to 16 of Appendix G.4. The results clearly indicate that as β increases, there is an initial improvement in performance; once it reaches a sufficiently large value, performance tends to drop. A strict (refers to a smaller β) context constraint leads to generated contexts getting closer to the initial contexts, thereby causing discussion inconsistency. On the other hand, a loose context constraint leads to naive consistency, where LLMs tend to generate the same answers, resulting in insufficient creativity.

Ablation studies. We assess the effect of key components by ablating them on all datasets under the same setting. 1) *Importance of context initialization.* 2) *Importance of tuning α .* 3) *Importance of context evolution.* As illustrated in Tables 11 to 13, LLMs struggle to specialize and coordinate

efficiently, whereas context initialization enables them to acquire high-impact contexts, significantly enhancing the foundational capabilities of MAD. Without tuning α during discussion rounds, LLMs tend to reach an agreement in the first round, leading to responses that lack creativity and diversity, which ultimately reduces problem-solving ability. Without context evolution, LLMs lack collaborative guidance considering previous responses, which prevents them from effectively leveraging the outputs of other LLMs.

Discrepancy intensity. To answer the fourth question, we explored the change of Discrepancy intensity ($\max_{i,j \in [N]} \|a_i - a_j\|^2$) with rounds to verify whether the generated context can help the LLMs gradually reach agreement with the discussion. Complete results are shown in Figs. 17 to 21, of Appendix G.5. We find that discrepancy intensity from M2CL decreases faster than other methods, corroborating its efficacy and superiority in converging the search space of multiple LLMs.

Transferability of contexts. To answer the fourth question, we also conduct experiments on transferring contexts directly into stronger LLMs to verify whether the generated contexts have better interpretation ability and efficacy. Full results are depicted in Table 10 of Appendix G.6. The results show that the transferred contexts deliver consistent improvement, indicating that the trained context generator can directly adapt to a wide range of models without additional retraining.

7 LIMITATION AND DISCUSSION

In this paper, we propose a novel context learning method, designed to expand LLMs’ horizons and reach consensus in MAD. By initializing and adjusting LLMs’ contexts based on the problem and discussion states, our method significantly improves problem-solving capabilities across diverse benchmarks while maintaining computational efficiency. A limitation of M2CL is the MAD framework in which diversity is brought about by the number of LLMs with heterogeneous characteristics which is computationally inefficient. An avenue for future work is to enable LLMs to truly capture the sub-tasks they are interested in or excel at.

ACKNOWLEDGMENTS

This research was supported in part by the National Natural Science Foundation of China under Grants 62572496 and 62402269, the Shenzhen Science and Technology Program under Grant JCYJ20250604175500001, and the Young Elite Scientist Sponsorship Program by CAST under Contract ZB2025-218.

ETHICS STATEMENT

This work advances the field of large language model (LLM) research by introducing a context learning method for multi-agent discussion, demonstrating significant improvements in accuracy, diversity, and consensus across multiple benchmarks.

However, the broader implications of deploying such systems warrant careful consideration. Multi-agent discussion, while powerful, may inadvertently propagate biases introduced during context initialization or amplify errors during consensus-building. These risks are particularly critical in high-stakes domains like legal, financial, and healthcare applications, where erroneous outputs could lead to significant ethical, social, or economic consequences.

REPRODUCIBILITY STATEMENT

The full algorithmic details of M2CL are presented in the main paper, with additional implementation details and hyperparameter settings included in Section 6. We provide source code at <https://github.com/HansenHua/M2CL-ICLR26>, which includes scripts for training, evaluation, and reproducing all reported experiments. We provide checkpoints at <https://huggingface.co/hansenhua/M2CL>. For theoretical results, we include complete proofs in Appendices B to D.

REFERENCES

Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.

Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. Chateval: Towards better llm-based evaluators through multi-agent debate. In *Proceedings of The 11st International Conference on Learning Representations*, 2023.

Ma Chang, Junlei Zhang, Zhihao Zhu, Cheng Yang, Yujiu Yang, Yaohui Jin, Zhenzhong Lan, Lingpeng Kong, and Junxian He. Agentboard: An analytical evaluation board of multi-turn llm agents. In *Proceedings of Advances in Neural Information Processing Systems*, volume 37, pp. 74325–74362, 2024.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

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

Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. In *Proceedings of The 40th International Conference on Machine Learning*, 2023.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Sagar Goyal, Eti Rastogi, Sree Prasanna Rajagopal, Dong Yuan, Fen Zhao, Jai Chintagunta, Gautam Naik, and Jeff Ward. Healai: A healthcare llm for effective medical documentation. In *Proceedings of The 17th ACM International Conference on Web Search and Data Mining*, pp. 1167–1168, 2024.

Qiuhan Gu. Llm-based code generation method for golang compiler testing. In *Proceedings of The 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pp. 2201–2203, 2023.

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

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. In *Proceedings of The 35th Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.

Shima Imani, Liang Du, and Harsh Shrivastava. Mathprompter: Mathematical reasoning using large language models. *arXiv preprint arXiv:2303.05398*, 2023.

Doohyuk Jang, Yoonjeon Kim, Chanjae Park, Hyun Ryu, and Eunho Yang. Reasoning model is stubborn: Diagnosing instruction overriding in reasoning models. *arXiv preprint arXiv:2505.17225*, 2025.

Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. Llm-blender: Ensembling large language models with pairwise ranking and generative fusion. *arXiv preprint arXiv:2306.02561*, 2023.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings of Advances in Neural Information Processing Systems*, volume 33, pp. 9459–9474, 2020.

Dongfang Li, Zhenyu Liu, Xinshuo Hu, Zetian Sun, Baotian Hu, and Min Zhang. In-context learning state vector with inner and momentum optimization. In *Proceedings of Advances in Neural Information Processing Systems*, 2024a.

Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbulin, and Bernard Ghanem. Camel: Communicative agents for "mind" exploration of large language model society. In *Proceedings of Advances in Neural Information Processing Systems*, volume 36, pp. 51991–52008, 2023.

Zhuowei Li, Zihao Xu, Ligong Han, Yunhe Gao, Song Wen, Di Liu, Hao Wang, and Dimitris N Metaxas. Implicit in-context learning. *arXiv preprint arXiv:2405.14660*, 2024b.

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

Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. A dynamic llm-powered agent network for task-oriented agent collaboration. In *Proceedings of The 1st Conference on Language Modeling*, 2024.

Ziru Liu, Cheng Gong, Xinyu Fu, Yaofang Liu, Ran Chen, Shoubo Hu, Suiyun Zhang, Rui Liu, Qingfu Zhang, and Dandan Tu. Ghpo: Adaptive guidance for stable and efficient llm reinforcement learning. *arXiv preprint arXiv:2507.10628*, 2025.

Li-Chun Lu, Shou-Jen Chen, Tsung-Min Pai, Chan-Hung Yu, Hung-yi Lee, and Shao-Hua Sun. Llm discussion: Enhancing the creativity of large language models via discussion framework and role-play. *arXiv preprint arXiv:2405.06373*, 2024.

Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning. In *Proceedings of The 11th International Conference on Learning Representations*, 2023.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. In *Proceedings of Advances in Neural Information Processing Systems*, volume 36, 2024.

Grégoire Mialon, Clémentine Fourrier, Thomas Wolf, Yann LeCun, and Thomas Scialom. Gaia: a benchmark for general ai assistants. In *Proceedings of The 12th International Conference on Learning Representations*, 2024.

Deepak Pandita, Tharindu Cyril Weerasooriya, Ankit Parag Shah, Isabelle Diana May-Xin Ng, Christopher M Homan, and Wei Wei. Prorefine: Inference-time prompt refinement with textual feedback. *arXiv preprint arXiv:2506.05305*, 2025.

Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of The 36th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–22, 2023.

Debjit Paul, Mete Ismayilzada, Maxime Peyrard, Beatriz Borges, Antoine Bosselut, Robert West, and Boi Faltings. Refiner: Reasoning feedback on intermediate representations. *arXiv preprint arXiv:2304.01904*, 2023.

Kiran Purohit, V Venktesh, Sourangshu Bhattacharya, and Avishek Anand. Sample efficient demonstration selection for in-context learning. In *Proceedings of The 42nd International Conference on Machine Learning*, 2025.

Chen Qian, Zihao Xie, YiFei Wang, Wei Liu, Kunlun Zhu, Hanchen Xia, Yufan Dang, Zhuoyun Du, Weize Chen, Cheng Yang, et al. Scaling large language model-based multi-agent collaboration. In *Proceedings of The 13th International Conference on Learning Representations*, 2025.

Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, et al. Toolllm: Facilitating large language models to master 16000+ real-world apis. In *Proceedings of The 12th International Conference on Learning Representations*, 2023.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Proceedings of Advances in Neural Information Processing Systems*, volume 36, pp. 53728–53741, 2023.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020.

Shreyas Sundara Raman, Vanya Cohen, Eric Rosen, Ifrah Idrees, David Paulius, and Stefanie Tellex. Planning with large language models via corrective re-prompting. In *Proceedings of The 36th Conference on Neural Information Processing System Foundation Models for Decision Making Workshop*, 2022.

Christopher Rawles, Sarah Clinckemaillie, Yifan Chang, Jonathan Waltz, Gabrielle Lau, Marybeth Fair, Alice Li, William E Bishop, Wei Li, Folawiyo Campbell-Ajala, et al. Androidworld: A dynamic benchmarking environment for autonomous agents. In *Proceedings of The Thirteenth International Conference on Learning Representations*, 2025.

David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. *arXiv preprint arXiv:2311.12022*, 2023.

Murray Shanahan, Kyle McDonell, and Laria Reynolds. Role play with large language models. *Nature*, 623(7987):493–498, 2023.

Weizhou Shen, Chenliang Li, Hongzhan Chen, Ming Yan, Xiaojun Quan, Hehong Chen, Ji Zhang, and Fei Huang. Small llms are weak tool learners: A multi-lm agent. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 16658–16680, 2024.

Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Cote, Yonatan Bisk, Adam Trischler, and Matthew Hausknecht. Alfworld: Aligning text and embodied environments for interactive learning. In *Proceedings of The 10th International Conference on Learning Representations*, 2021.

Andries Petrus Smit, Nathan Grinsztajn, Paul Duckworth, Thomas D Barrett, and Arnu Pretorius. Should we be going mad? a look at multi-agent debate strategies for llms. In *Proceedings of The 41st International Conference on Machine Learning*, 2024.

Henry W Sprueill, Carl Edwards, Khushbu Agarwal, Mariefel V Olarte, Udishnu Sanyal, Conrad Johnston, Hongbin Liu, Heng Ji, and Sutanay Choudhury. Chemreasoner: Heuristic search over a large language model's knowledge space using quantum-chemical feedback. In *Proceedings of The 41st International Conference on Machine Learning*, 2024.

Eric Todd, Millicent Li, Arnab Sen Sharma, Aaron Mueller, Byron C Wallace, and David Bau. Function vectors in large language models. In *Proceedings of The 12th International Conference on Learning Representations*, 2024.

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

A Vaswani. Attention is all you need. In *Proceedings of Advances in Neural Information Processing Systems*, 2017.

Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. Transformers learn in-context by gradient descent. In *Proceedings of The 41st International Conference on Machine Learning*, pp. 35151–35174. PMLR, 2023.

Ruoyao Wang, Peter Jansen, Marc-Alexandre Côté, and Prithviraj Ammanabrolu. Scienceworld: Is your agent smarter than a 5th grader? In *Proceedings of The 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 11279–11298, 2022.

Sizhe Wang, Zhengren Wang, Dongsheng Ma, Yongan Yu, Rui Ling, Zhiyu Li, Feiyu Xiong, and Wentao Zhang. Codeflowbench: A multi-turn, iterative benchmark for complex code generation. *arXiv preprint arXiv:2504.21751*, 2025.

Jimmy Wei, Kurt Shuster, Arthur Szlam, Jason Weston, Jack Urbanek, and Mojtaba Komeili. Multi-party chat: Conversational agents in group settings with humans and models. *arXiv preprint arXiv:2304.13835*, 2023.

Shiwei Wu, Yu Liu, Yan Gao, Zhi Li, Chao Zhang, Qimeng Wang, Enhong Chen, and You He. Evaluating conversational question answering in multi-instructional documents with large language models. *Chinese Journal of Electronics*, 2025.

Jing Xiong, Zixuan Li, Chuanyang Zheng, Zhijiang Guo, Yichun Yin, Enze Xie, Zhicheng YANG, Qingxing Cao, Haiming Wang, Xiongwei Han, Jing Tang, Chengming Li, and Xiaodan Liang. DQ-lore: Dual queries with low rank approximation re-ranking for in-context learning. In *Proceedings of The 12th International Conference on Learning Representations*, 2024.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024a.

Hao Yang, Qianghua Zhao, and Lei Li. Chain-of-thought in large language models: Decoding, projection, and activation. *arXiv preprint arXiv:2412.03944*, 2024b.

Jingyuan Yang, Rongjun Li, Weixuan Wang, Ziyu Zhou, Zhiyong Feng, and Wei Peng. Lf-steering: Latent feature activation steering for enhancing semantic consistency in large language models. *arXiv preprint arXiv:2501.11036*, 2025.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *Proceedings of The 11th International Conference on Learning Representations*, 2023.

Beiyuan Yu, Jian Mao, and Jianwei Liu. Cookiechecker: Automated cookie setting compliance analysis in real-world web environments. *Chinese Journal of Electronics*, 2025.

Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. Star: Bootstrapping reasoning with reasoning. In *Advances in Neural Information Processing Systems*, volume 35, pp. 15476–15488, 2022.

Guibin Zhang, Muxin Fu, Guancheng Wan, Miao Yu, Kun Wang, and Shuicheng Yan. G-memory: Tracing hierarchical memory for multi-agent systems. *arXiv preprint arXiv:2506.07398*, 2025a.

Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xiong-Hui Chen, Jiaqi Chen, Mingchen Zhuge, Xin Cheng, Sirui Hong, Jinlin Wang, et al. Aflow: Automating agentic workflow generation. In *Proceedings of The 13th International Conference on Learning Representations*, 2025b.

Yiming Zhang, Shi Feng, and Chenhao Tan. Active example selection for in-context learning. In *Proceedings of The 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 9134–9148, 2022.

Zheyuan Zhang, Daniel Zhang-Li, Jifan Yu, Linlu Gong, Jinchang Zhou, Zhanxin Hao, Jianxiao Jiang, Jie Cao, Huiqin Liu, Zhiyuan Liu, et al. Simulating classroom education with llm-empowered agents. *arXiv preprint arXiv:2406.19226*, 2024.

Qinkai Zheng, Xiao Xia, Xu Zou, Yuxiao Dong, Shan Wang, Yufei Xue, Lei Shen, Zihan Wang, Andi Wang, Yang Li, et al. Codegeex: A pre-trained model for code generation with multilingual benchmarking on humaneval-x. In *Proceedings of The 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 5673–5684, 2023.

Mingchen Zhuge, Wenyi Wang, Louis Kirsch, Francesco Faccio, Dmitrii Khizbulin, and Jürgen Schmidhuber. Gptswarm: Language agents as optimizable graphs. In *Proceedings of The 41st International Conference on Machine Learning*, 2024.

A USAGE OF LLMs

We use large language models (LLMs) solely as an assistive tool for polishing the writing and improving clarity of exposition. All content generated by LLMs was carefully reviewed, verified, and, where necessary, revised by the authors. The authors take full responsibility for the correctness and integrity of the final manuscript.

B PROOF OF THEOREM 4.1

Proof. Define the optimal $\omega^* = \arg \min_{\omega} \|a_c - \sum_{i=1}^N \omega_i a(C_i^b)\|$ and its corresponding activation $a = \sum_{i=1}^N \omega_i^* a(C_i^b)$. Then, we can derive the upper bound of the activation difference using the triangle inequality:

$$\sum_{i=1}^N \|a_c - a(C_i^t)\| \leq \sum_{i=1}^N \|a(C_i^t) - a\| + N \|a - a_c\|. \quad (18)$$

Using the triangle inequality, we can bound the first term in Eq. (18):

$$\begin{aligned} \sum_{i=1}^N \|a(C_i^t) - a\| &= \sum_{i=1}^N \left\| \sum_{j=1}^N \omega_j^* (a(C_i^t) - a(C_j^b)) \right\| \\ &\leq \sum_{i=1}^N \sum_{j=1}^N |\omega_j^*| \cdot \|a(C_i^t) - a(C_j^b)\| \leq \sum_{i=1}^N \sum_{j=1}^N \|a(C_i^t) - a(C_j^b)\| \\ &\leq \sum_{i=1}^N \sum_{j=1}^N \|a(C_i^t) - a(C_j^t)\| + \sum_{i=1}^N \sum_{j=1}^N \|a(C_j^t) - a(C_j^b)\| \\ &\leq \sum_{i=1}^N \sum_{j=1}^N \|a(C_i^t) - a(C_j^t)\| + \sum_{i=1}^N \sum_{j=1}^N L_a \|C_j^t - C_j^b\|. \end{aligned} \quad (19)$$

Then, we bound the second term in Eq. (18):

$$\begin{aligned} \|a - a_c\| &\leq \left\| \sum_{i=1}^N \omega_i^* a(C_i^t) - a_c \right\| = \left\| \sum_{i=1}^N \omega_i^* a(C_i^b) + \sum_{i=1}^N \omega_i^* (a(C_i^t) - a(C_i^b)) - a_c \right\| \\ &\leq \min_{\omega} \left\| a_c - \sum_{i=1}^N \omega_i a(C_i^b) \right\| + \sum_{i=1}^N \|\omega_i^* (a(C_i^t) - a(C_i^b))\| \\ &\leq \min_{\omega} \left\| a_c - \sum_{i=1}^N \omega_i a(C_i^b) \right\| + \sum_{i=1}^N \|a(C_i^t) - a(C_i^b)\| \\ &\leq \min_{\omega} \left\| a_c - \sum_{i=1}^N \omega_i a(C_i^b) \right\| + L_a \sum_{i=1}^N \|C_i^t - C_i^b\|. \end{aligned} \quad (20)$$

In the first line, we add and subtract $\sum_{i=1}^N \omega_i^* a(C_i^t)$. In the second line, we use the triangle inequality and the definition of ω^* . In the third line, we scale ω^* to 1. Finally, we utilize the smoothness property of the activation function to bound the second term. Substituting into Eq. (18), we can derive:

$$\begin{aligned} \sum_{i=1}^N \|a_c - a(C_i^t)\| &\leq \sum_{i=1}^N \left(\sum_{j=1}^N \|a(C_i^t) - a(C_j^t)\| + (N+1)L_a \|C_i^t - C_i^b\| \right) \\ &\quad + N \min_{\omega} \left\| a_c - \sum_{i=1}^N \omega_i a(C_i^b) \right\|, \end{aligned} \quad (21)$$

□

C DECOUPLED CRITERION FUCNTION

Lemma C.1. *Under the assumption of one-block transformer, LLMs' activation diversity can be bound by the activation difference between their instructions and responses.*

$$\|a(C_i^t) - a(C_j^t)\| \leq \|a(I_i^t) - a(X_i^t)\| + \|a(I_j^t) - a(X_j^t)\| + 18L_V n N \exp(2\rho^2). \quad (22)$$

Proof. We first decompose the activation of context into a combination of its components:

$$\begin{aligned} a(C_i^t) &= a(C_i^t) - [a(I_i^t) + \sum_{k \neq i}^N a(X_k^{t-1}) - (N-1)a(P)] + [a(I_i^t) + \sum_{k \neq i}^N a(X_k^{t-1}) - (N-1)a(P)] \\ a(C_j^t) &= a(C_j^t) - [a(I_j^t) + \sum_{k \neq j}^N a(X_k^{t-1}) - (N-1)a(P)] + [a(I_j^t) + \sum_{k \neq j}^N a(X_k^{t-1}) - (N-1)a(P)]. \end{aligned} \quad (23)$$

By using Lemma D.3, we can derive the upper bound of the activation difference as:

$$\begin{aligned} &\|a(C_i^t) - a(C_j^t)\| \\ &\leq \|a(C_i^t) - [a(I_i^t) + \sum_{k \neq i}^N a(X_k^{t-1}) - (N-1)a(P)]\| + \|a(I_i^t) - a(X_i^{t-1})\| \\ &\quad + \|a(C_j^t) - [a(I_j^t) + \sum_{k \neq j}^N a(X_k^{t-1}) - (N-1)a(P)]\| + \|a(I_j^t) - a(X_j^{t-1})\| \\ &\leq \|a(I_i^t) - a(X_i^{t-1})\| + \|a(I_j^t) - a(X_j^{t-1})\| + 18L_V n N \exp(2\rho^2). \end{aligned} \quad (24)$$

□

D USEFUL LEMMA

Lemma D.1. *Suppose $\|W_V X\| \leq L_V$ and $\|(W_K X)^T W_Q X\| \leq \rho^2$, the difference between the activation of softmax attention $a(X)$ and linear attention $a'(X)$ can be bounded by:*

$$\|a(X) - a'(X)\| \leq 3L_V n \exp(2\rho^2) \quad (25)$$

where $a(X) = W_V X \text{softmax}\left(\frac{(W_K X)^T W_Q X}{\sqrt{d}}\right)$ and $a'(X) = W_V X (W_K X)^T W_Q X$.

Proof. Define $\Delta = D^{-1} \exp(S) - \exp(S)$, where $S = (W_K X)^T W_Q X$ and $D = \text{diag}(\langle \exp(S), \mathbf{1}_n \rangle)$. Then, we can derive the upper bound of each element of Δ .

$$\Delta_{ij} = \frac{\exp(S_{ij})}{D_{ii}} - \exp(S_{ij}) \leq \exp(S_{ij}) \left(\left| \frac{1}{D_{ii}} \right| + 1 \right) \leq \exp(\rho^2) (1 + \exp(\rho^2)) \leq 2 \exp(2\rho^2) \quad (26)$$

Summing them up, we can derive:

$$\|D^{-1} \exp(S) - \exp(S)\| \leq \|D^{-1} \exp(S) - \exp(S)\|_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^n \Delta_{ij}^2} \leq 2n \exp(2\rho^2) \quad (27)$$

Define $\delta = \exp(S) - S$, where $S = (W_K X)^T W_Q X$. Then, we derive the upper bound of δ by finding the upper bound of each element of δ .

$$\delta_{ij} = \exp(S_{ij}) - S_{ij} \leq \max\{\exp(\rho^2) - \rho^2, \rho^2 + \exp(-\rho^2)\} \leq \exp(\rho^2) \quad (28)$$

Summing them up, we can derive:

$$\|\exp(S) - S\| \leq \|\exp(S) - S\|_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^n \exp(S_{ij}) - S_{ij}} \leq n \exp(\rho^2) \quad (29)$$

After bounding $\|D^{-1} \exp(S) - \exp(S)\|$ and $\|\exp(S) - S\|$, we can derive the upper bound of $\|D^{-1} \exp(S) - S\|$ by basic algebra as:

$$\begin{aligned} & \|D^{-1} \exp(S) - S\| \\ &= \|D^{-1} \exp(S) - \exp(S) + \exp(S) - S\| \\ &\leq \|D^{-1} \exp(S) - \exp(S)\| + \|\exp(S) - S\| \\ &\leq 2n \exp(2\rho^2) + n \exp(\rho^2) \\ &\leq 3n \exp(2\rho^2) \end{aligned} \quad (30)$$

Therefore, we derive the difference between the activation of softmax attention and linear attention as:

$$\begin{aligned} & \|W_V X \text{softmax}\left(\frac{(W_K X)^T W_Q X}{\sqrt{d}}\right) - W_V X (W_K X)^T W_Q X\| \\ &\leq \|W_V X\| \cdot \|\text{softmax}\left(\frac{(W_K X)^T W_Q X}{\sqrt{d}}\right) - (W_K X)^T W_Q X\| \\ &\leq L_V \cdot \|D^{-1} \exp(S) - S\| \\ &\leq 3L_V n \exp(2\rho^2) \end{aligned} \quad (31)$$

□

Lemma D.2. Define the activation of linear attention:

$$\begin{aligned} a'(Y) &\doteq W_V [Y, P] (W_K [Y, P])^T W_Q Y \\ a'(P) &\doteq W_V P (W_K P)^T W_Q P, \end{aligned} \quad (32)$$

then the activation of a long prompt $Y = [Y_1, Y_2, \dots, Y_N]$ can be derived by the activation of its component:

$$a'(Y) = \sum_{i=1}^N a'(Y_i) - (N-1)a'(P). \quad (33)$$

Proof.

$$\begin{aligned} a'(Y) &= W_V [Y_1, Y_2, \dots, Y_N, P] \left(W_K [Y_1, Y_2, \dots, Y_N, P] \right)^T W_Q P \\ &= W_V [Y_1, Y_2, \dots, Y_N, P] [Y_1, Y_2, \dots, Y_N, P]^T W_K^T W_Q P \\ &= W_V \left(\sum_{i=1}^N Y_i Y_i^T + P P^T \right) W_K^T W_Q P \\ &= W_V \left[\sum_{i=1}^N (Y_i Y_i^T + P P^T) - (N-1) P P^T \right] W_K^T W_Q P \\ &= \sum_{i=1}^N W_V (Y_i Y_i^T + P P^T) W_K^T W_Q P - W_V (N-1) P P^T W_K^T W_Q P \\ &= \sum_{i=1}^N W_V [Y_i, P] (W_K [Y_i, P])^T W_Q P - (N-1) W_V P (W_K P)^T W_Q P \\ &= \sum_{i=1}^N a'(Y_i) - (N-1)a'(P). \end{aligned} \quad (34)$$

□

Lemma D.3. *We derive the difference between the activation of a long prompt Y and a combination of the activations of its components $Y = [Y_1, Y_2, \dots, Y_N, P]$ as:*

$$\|a(Y) - \sum_{i=1}^N a(Y_i) + (N-1)a(P)\| \leq 9L_V n N \exp(\rho^2). \quad (35)$$

Proof. Denote the activation of linear attention as a' , we can complete the proof by using Lemma D.2 to simplify the relationship of linear attention and Lemma D.1 to bound the difference:

$$\begin{aligned} & \|a(Y) - \sum_{i=1}^N a(Y_i) + (N-1)a(P)\| \\ &= \|a(Y) - a'(Y) + a'(Y) - [\sum_{i=1}^N a'(Y_i) - (N-1)a'(P)] \\ &\quad + [\sum_{i=1}^N a'(Y_i) - (N-1)a'(P)] - [\sum_{i=1}^N a(Y_i) - (N-1)a(P)]\| \\ &\leq \|a(Y) - a'(Y)\| + \sum_{i=1}^N \|a'(Y_i) - a(Y_i)\| + (N-1)\|a'(P) - a(P)\| \\ &\leq \exp(2\rho^2)[3L_V n N + 3L_V N n + 3L_V (N-1)n] \\ &\leq 9L_V n N \exp(2\rho^2) \end{aligned} \quad (36)$$

□

Lemma D.4. *If $\|S\| \leq \rho$, we have $\langle \exp(S), \mathbf{1}_n \rangle \geq \exp(-\rho)$*

Proof.

$$\begin{aligned} \langle \exp(S), \mathbf{1}_n \rangle &= \sum_{i=1}^n \exp(S_i) \geq \min_{i \in [n]} \exp(S_i) \geq \min_{i \in [n]} \exp(-|S_i|) \\ &= \exp(-\max_{i \in [n]} |S_i|) = \exp(-\|S\|_\infty) \geq \exp(-\|S\|_2) \geq \exp(-\rho) \end{aligned} \quad (37)$$

□

E MULTI-ROUND CONTEXT LEARNING

First, we define:

$$h(I_i^t) \doteq \beta - \|I_i^t - I_i^b\| \quad (38)$$

$$f(I_i^t) \doteq \begin{cases} \|a(I_i^t) - a(X_i^{t-1})\|, & \text{if } h(I_i^t) \geq 0 \\ +\infty, & \text{otherwise.} \end{cases} \quad (39)$$

To solve the minimization optimization with inequality constraint, we can construct a Lagrangian expression with a Lagrange multiplier α_T :

$$\text{minimize } f(I_i^t) \text{ s.t. } h(I_i^t) \geq 0 \Leftrightarrow L(I_i^t, \alpha_t) = f(I_i^t) - \alpha_t h(I_i^t). \quad (40)$$

The optimization changes to:

$$\min_{I_i^t} f(I_i^t) = \max_{\alpha_t \geq 0} \min_{I_i^t} L(I_i^t, \alpha_t). \quad (41)$$

Therefore, to minimize $f(I_i^t)$, the dual problem is listed as below. Note that to make sure $\min_{I_i^t} f(I_i^t)$ is properly minimized and would not become $+\infty$, the constraint has to be satisfied.

$$\begin{aligned}
& \min_{I_i^t} \left[\|a(I_i^t) - a(X_i^{t-1})\| \right] \\
&= \min_{I_i^t} f(I_i^t) \\
&= \max_{\alpha_t \geq 0} \min_{I_i^t} L(I_i^t, \alpha_t) \\
&= \max_{\alpha_t \geq 0} \min_{I_i^t} f(I_i^t) - \alpha_t h(I_i^t) \\
&= \max_{\alpha_t \geq 0} \min_{I_i^t} \left[\|a(I_i^t) - a(X_i^{t-1})\| \right] - \alpha_t \left(\beta - \|I_i^t - I_i^b\| \right) \\
&= \max_{\alpha_t \geq 0} \min_{I_i^t} \left[\|a(I_i^t) - a(X_i^{t-1})\| - \alpha_t \beta + \alpha_t \|I_i^t - I_i^b\| \right]. \tag{42}
\end{aligned}$$

Define the Q-function for this multi-round multi-agent collaboration as:

$$Q(a(I_i^t), a(X_i^{t-1})) = \|a(I_i^t) - a(X_i^{t-1})\| + \min_{I_i^t} \|a(I_i^{t+1}) - a(X_i^t)\| + \alpha_{t+1} \|I_i^{t+1,*} - I_i^b\|. \tag{43}$$

Here, $I_i^{t+1,*}$ denotes the optimal I_i^{t+1} . Therefore, the expected return is as follows, when we take one round further back to the round $T-1$:

$$\begin{aligned}
& \min_{I_i^{T-1}} \left(\|a(I_i^{T-1}) - a(X_i^{T-2})\| + \min_{I_i^t} \|a(I_i^t) - a(X_i^{T-1})\| \right) \\
&= \min_{I_i^{T-1}} \left(Q(a(I_i^{T-1}), a(X_i^{T-2})) - \alpha_T^* \|I_i^{T,*} - I_i^b\| \right) \\
&= \max_{\alpha_{T-1} > 0} \min_{I_i^{T-1}} \left(Q(a(I_i^{T-1}), a(X_i^{T-2})) - \alpha_{T-1} (\beta - \|I_i^{T-1} - I_i^b\|) - \alpha_T^* \|I_i^{T,*} - I_i^b\| \right) \\
&= \max_{\alpha_{T-1} > 0} \min_{I_i^{T-1}} \left(Q(a(I_i^{T-1}), a(X_i^{T-2})) - \alpha_{T-1} \beta + \alpha_{T-1} \|I_i^{T-1} - I_i^b\| \right) - \alpha_T^* \|I_i^{T,*} - I_i^b\|. \tag{44}
\end{aligned}$$

Similar to the previous round,

$$\begin{aligned}
I_i^t &= \arg \min_{I_i^t} \left[\|a(I_i^t) - a(X_i^{t-1})\| - \alpha_t \beta + \alpha_t \|I_i^t - I_i^b\| \right] \\
\alpha_t^* &= \arg \max_{\alpha_t \geq 0} \left[\alpha_t \|I_i^t - I_i^b\| - \alpha_t \beta \right]. \tag{45}
\end{aligned}$$

By repeating this process, we can learn the optimal temperature parameter in every round. Hence, the loss functions for θ and α are as follows:

$$\begin{aligned}
L(\theta_i^t) &= \|a(I_i^t) - a(X_i^{t-1})\| + \alpha_t \|I_i^t - I_i^b\| \\
L(\alpha) &= \alpha (\beta - \|I_i^t - I_i^b\|). \tag{46}
\end{aligned}$$

F EXPERIMENT SETUP

F.1 DATASETS

We evaluate our method on three areas with 7 datasets which are widely used in prior studies (Dubey et al., 2024; Qian et al., 2025). We elaborate on what follows.

- MMLU (Hendrycks et al., 2021), a comprehensive benchmark covering diverse subjects and difficulty levels, designed to test world knowledge and logical reasoning through multiple-choice questions.
- MATH (Hendrycks et al., 2021), a dataset of challenging competition-level math problems requiring multi-step symbolic reasoning and advanced problem-solving ability.
- GPQA (Rein et al., 2023), a benchmark of graduate-level multiple-choice science questions that assesses deep domain knowledge and reasoning under uncertainty.
- HumanEval (Chen et al., 2021), a widely recognized benchmark for function-level code generation, designed to evaluate fundamental programming skills.
- ALFWorld (Shridhar et al., 2021), a text-based embodied environment featuring household tasks, where agents navigate and interact with objects via natural language commands.
- SciWorld (Wang et al., 2022), a text-based embodied environment for interactive science tasks, requiring agents to navigate rooms, conduct experiments, and perform procedural reasoning.
- GAIA (Mialon et al., 2024), a benchmark of real-world question answering tasks that integrate knowledge retrieval, reasoning, and multi-step tool use.
- PDDL (Chang et al., 2024), an environment comprising diverse strategic games, where agents must employ PDDL expressions to plan and execute complex tasks.
- AndroidWorld (Rawles et al., 2025), an environment with 116 dynamic tasks across 20 real-world Android apps, designed to evaluate mobile agents’ capabilities in app navigation and system-level control.

We use 20% of the questions to construct the training dataset and the rest as testing dataset.

F.2 BASELINES

We test our method against six baselines. We implement them based on their publicly available implementations.

- *Single Execution* (`Single`), querying a single LLM to solve the task.
- *Best-of-N sampling* (`BoN`), querying a single LLM N times and sampling the most correct answer. We set N to 32, as increasing N further does not yield additional performance benefits and aligns closely with the number of LLM calls used for multi-agent methods.
- *Discussion* (Du et al., 2023), a multi-agent framework in which LLMs are assigned predefined distinct contexts and iteratively exchange their reasoning processes as additional prompts over multiple rounds, before producing a final answer via majority voting.
- *Dynamic LLM-Powered Agent Network* (Liu et al., 2024) (`DyLAN`), a discussion-style framework which incorporates an LLM selection algorithm based on an unsupervised metric, namely the Agent Importance Score, which identifies the most contributive LLMs through a preliminary trial tailored to the specific task.
- *GPTSwarm* (Zhuge et al., 2024), formalizing a swarm of autonomous agents as computational graphs, with nodes as manually-customized functions and edges facilitating information flow, adaptively optimizing node prompts and modifying graph connectivity during collective reasoning.
- *Multi-LLM Collaboration Network* (Qian et al., 2025) (`MacNet`), a representative framework for decentralized and scalable multi-LLM systems. It introduces edge agents that mediate interactions by generating actionable instructions for the next agent based on the outputs of the previous one.

F.3 IMPLEMENTATION DETAILS

The context pool is constructed using GPT-4o, where we prompt it to generate a large collection of high-quality initial contexts across diverse domains, including mathematics, science, coding, and

embodied reasoning, and various domain-specific sub-contexts are included in each domain. This ensures that the pool provides a broad coverage of reasoning perspectives and is shared across all tasks.

The context generators are implemented by T5-small model (Raffel et al., 2020) for all the tasks. The dimension of the generated context vectors is 512. The learning rates for the context generators and α are both $1e-4$. We use 20% of the questions to train the context initialization and generator.

Table 2: Hyperparameters (identical across datasets).

Hyperparameter	Value
Sentence vector dimension	512
Optimizer	Adam
Batchsize	32
Learning rate of context	1e-4
Learning rate of α	1e-4
Maximum rounds for discussion	8
Training rounds	100
Size of context pool	100

We implement our code using Pytorch 2.3.0, built upon the open-source parameters of the llama2, Qwen2.5, and Qwen2.5-VL. Models provided at <https://huggingface.co/meta-llama> and <https://huggingface.co/Qwen>. All the experiments are run on Ubuntu 22.04.4 LTS with 8 NVIDIA H800 GPUs.

F.4 PSEUDOCODE OF M2CL

We present the pseudocode of training M2CL in Algorithm 1. It begins with training the context initialization, then the context generators are trained along with the weight α during the discussion.

Algorithm 1: Pseudocode of M2CL

```

1 Initialize parameters  $\phi_f$ ,  $\phi_F$ ,  $\{\theta_i\}_{i=1}^N$ , and  $\{\alpha_i\}_{i=1}^N$ ;
2 for each question do
3    $\phi_f \leftarrow \phi_f - \eta_f \nabla L(\phi_f)$ ,  $\phi_F \leftarrow \phi_F - \eta_F \nabla L(\phi_F)$ ;
4 end
5 for each epoch do
6   Obtain initial contexts  $\{I_i^b\}_{i=1}^N$  via Eq. (12);
7   for each round  $t$  do
8     for  $i = 1$  to  $N$  do
9       Generate instructions  $I_i^t$  via Eq. (3) and obtain responses  $X_i^t$  via Eq. (4);
10       $\theta_i \leftarrow \theta_i - \eta_\theta \nabla L(\theta_i)$ ,  $\alpha_i \leftarrow \alpha_i - \eta_\alpha \nabla L(\alpha_i)$ ;
11    end
12  end
13 end

```

G ADDITIONAL RESULTS

G.1 COMPARISON WITH MORE MODELS

To further investigate M2CL’s performance, we compare it with `best-of-N` using stronger close-sourced LLMs, including Qwen2.5-Max and GPT-4. As shown in Table 3, llama series models perform worse when fewer LLMs participate but achieve higher accuracy as the number of LLMs increases, demonstrating that M2CL can collaborate weaker models to achieve comparable performance.

Table 3: Accuracy using different base models on different datasets. We exhibit the performance advantage with Qwen-max and highlight the best result.

	Model	MMLU	MATH	GPQA	Code	ALFWorld	SciWorld	GAIA	PDDL
n=1	Qwen-max	86.5 ^{±0.0}	43.6 ^{±0.0}	35.7 ^{±0.0}	65.9 ^{±0.0}	73.1 ^{±0.0}	68.7 ^{±0.0}	61.5 ^{±0.0}	76.4 ^{±0.0}
	GPT-4	84.3 ^{±2.2}	41.1 ^{±2.5}	33.3 ^{±2.4}	62.7 ^{±3.2}	70.2 ^{±2.9}	65.4 ^{±3.3}	58.7 ^{±2.8}	73.5 ^{±2.9}
	Llama-7B	45.3 ^{±41.2}	2.9 ^{±40.7}	16.8 ^{±18.9}	15.8 ^{±50.1}	22.5 ^{±50.6}	20.1 ^{±48.6}	15.3 ^{±46.2}	24.2 ^{±52.2}
	Llama-13B	54.8 ^{±31.7}	3.9 ^{±39.7}	25.0 ^{±10.7}	23.7 ^{±42.2}	28.5 ^{±44.6}	24.5 ^{±44.2}	18.6 ^{±42.9}	30.4 ^{±46.0}
	Llama-70B	68.9 ^{±17.6}	31.6 ^{±12.0}	27.6 ^{±8.1}	35.5 ^{±30.4}	46.1 ^{±27.0}	40.2 ^{±28.5}	29.9 ^{±31.6}	48.4 ^{±28.0}
n=4	Qwen-max	87.5 ^{±0.0}	44.2 ^{±0.0}	39.2 ^{±0.0}	67.5 ^{±0.0}	74.0 ^{±0.0}	69.8 ^{±0.0}	62.5 ^{±0.0}	77.5 ^{±0.0}
	GPT-4	86.4 ^{±1.1}	43.9 ^{±0.3}	38.8 ^{±0.4}	67.0 ^{±0.5}	72.8 ^{±1.2}	68.0 ^{±1.8}	61.0 ^{±1.5}	76.0 ^{±1.5}
	Llama-7B	59.1 ^{±28.4}	17.4 ^{±26.8}	44.7 ^{±5.5}	32.9 ^{±34.6}	35.1 ^{±38.9}	33.0 ^{±36.8}	25.5 ^{±37.0}	33.1 ^{±44.4}
	Llama-13B	86.9 ^{±0.6}	23.4 ^{±20.8}	58.9 ^{±19.7}	47.5 ^{±20.0}	42.9 ^{±31.1}	39.9 ^{±29.9}	30.5 ^{±32.0}	40.9 ^{±36.6}
	Llama-70B	95.6 ^{±8.1}	40.3 ^{±3.9}	78.7 ^{±39.5}	70.6 ^{±3.1}	69.5 ^{±4.5}	65.5 ^{±4.3}	49.6 ^{±13.0}	66.5 ^{±11.0}
n=8	Qwen-max	88.5 ^{±0.0}	45.0 ^{±0.0}	40.5 ^{±0.0}	68.8 ^{±0.0}	75.0 ^{±0.0}	71.0 ^{±0.0}	63.5 ^{±0.0}	79.0 ^{±0.0}
	GPT-4	86.5 ^{±2.0}	44.6 ^{±0.4}	38.9 ^{±1.6}	67.1 ^{±1.7}	73.2 ^{±1.8}	69.5 ^{±1.5}	61.0 ^{±2.5}	76.0 ^{±3.0}
	Llama-7B	63.8 ^{±24.7}	24.9 ^{±20.1}	59.3 ^{±18.8}	38.8 ^{±30.0}	37.2 ^{±37.8}	36.1 ^{±34.9}	27.1 ^{±36.4}	34.7 ^{±44.3}
	Llama-13B	92.7 ^{±4.2}	23.0 ^{±22.0}	70.7 ^{±30.2}	55.4 ^{±13.4}	46.1 ^{±28.9}	44.7 ^{±26.3}	33.6 ^{±29.9}	43.8 ^{±35.2}
	Llama-70B	95.4 ^{±6.9}	43.5 ^{±1.5}	87.6 ^{±47.1}	81.0 ^{±12.2}	78.9 ^{±3.9}	76.2 ^{±3.2}	58.7 ^{±4.8}	73.6 ^{±5.4}
n=16	Qwen-max	88.5 ^{±0.0}	45.0 ^{±0.0}	40.5 ^{±0.0}	68.8 ^{±0.0}	75.0 ^{±0.0}	71.0 ^{±0.0}	63.5 ^{±0.0}	79.0 ^{±0.0}
	GPT-4	87.9 ^{±0.6}	44.8 ^{±0.2}	40.1 ^{±0.4}	68.5 ^{±0.3}	74.3 ^{±0.7}	70.5 ^{±0.5}	63.0 ^{±0.5}	78.2 ^{±0.8}
	Llama-7B	71.5 ^{±17.0}	23.9 ^{±21.1}	69.9 ^{±29.4}	48.3 ^{±20.5}	39.8 ^{±35.2}	38.4 ^{±32.6}	28.9 ^{±34.6}	37.0 ^{±42.0}
	Llama-13B	94.5 ^{±6.0}	32.7 ^{±12.3}	88.5 ^{±48.0}	66.1 ^{±2.7}	48.7 ^{±26.3}	47.1 ^{±23.9}	36.2 ^{±27.3}	45.5 ^{±33.5}
	Llama-70B	93.9 ^{±5.4}	51.5 ^{±6.5}	91.3 ^{±50.8}	97.2 ^{±28.4}	84.9 ^{±9.9}	81.7 ^{±10.7}	61.5 ^{±2.0}	78.1 ^{±0.9}
n=32	Qwen-max	89.0 ^{±0.0}	45.5 ^{±0.0}	41.0 ^{±0.0}	69.5 ^{±0.0}	75.5 ^{±0.0}	71.5 ^{±0.0}	64.0 ^{±0.0}	79.5 ^{±0.0}
	GPT-4	89.1 ^{±0.1}	45.4 ^{±0.1}	41.5 ^{±0.5}	69.7 ^{±0.2}	75.0 ^{±0.5}	71.0 ^{±0.5}	63.5 ^{±0.5}	78.5 ^{±1.0}
	Llama-7B	79.1 ^{±9.9}	28.8 ^{±16.7}	84.5 ^{±43.5}	56.5 ^{±13.0}	42.5 ^{±33.0}	40.7 ^{±30.8}	31.4 ^{±32.6}	38.9 ^{±40.6}
	Llama-13B	95.8 ^{±6.8}	41.4 ^{±4.1}	94.7 ^{±53.7}	75.6 ^{±6.1}	52.2 ^{±23.3}	50.7 ^{±20.8}	38.5 ^{±25.5}	48.0 ^{±31.5}
	Llama-70B	97.0 ^{±8.0}	54.5 ^{±9.0}	95.1 ^{±54.1}	93.7 ^{±24.2}	88.5 ^{±13.0}	86.0 ^{±14.5}	65.7 ^{±1.7}	79.8 ^{±0.3}
n=64	Qwen-max	89.5 ^{±0.0}	46.0 ^{±0.0}	41.5 ^{±0.0}	70.0 ^{±0.0}	76.0 ^{±0.0}	72.0 ^{±0.0}	64.5 ^{±0.0}	80.0 ^{±0.0}
	GPT-4	88.5 ^{±1.0}	45.4 ^{±0.6}	41.8 ^{±0.3}	70.5 ^{±0.5}	75.2 ^{±0.8}	71.5 ^{±0.5}	64.0 ^{±0.5}	79.2 ^{±0.8}
	Llama-7B	81.5 ^{±8.0}	35.4 ^{±10.6}	82.5 ^{±41.0}	60.6 ^{±9.4}	44.2 ^{±31.8}	42.9 ^{±29.1}	32.9 ^{±31.6}	40.9 ^{±39.1}
	Llama-13B	96.8 ^{±7.3}	40.0 ^{±6.0}	93.1 ^{±51.6}	82.5 ^{±12.5}	54.8 ^{±21.2}	52.8 ^{±19.2}	40.3 ^{±24.2}	49.8 ^{±30.2}
	Llama-70B	96.3 ^{±6.8}	58.0 ^{±12.0}	96.9 ^{±55.4}	92.0 ^{±22.0}	90.8 ^{±14.8}	88.9 ^{±16.9}	68.2 ^{±3.7}	82.2 ^{±2.2}

G.2 NUMBER OF LLMs

We evaluate M2CL’s performance with varying number of LLMs, ranging from 4 to 64. The comparative results are shown in Tables 4 to 8 and Figs. 5 to 9.

Summary of key findings. The results show that M2CL significantly improves MAD performance, consistently surpassing existing methods by 20% – 50%, particularly in complex tasks like math and tool-using. This highlights its ability to tackle intricate reasoning tasks. M2CL also exhibits a more effective multi-agent scaling law, where performance consistently improves as the number of LLMs increases to 64, especially in agentic domains where a larger amount of LLMs enhances problem-solving accuracy and efficiency.

Table 4: Accuracy with varying number of LLMs on different datasets. The number of LLMs is 4 for all datasets. We exhibit the performance advantage with BoN and highlight the best result.

Model	Method	MMLU	MATH	GPQA	Code	ALFWorld	SciWorld	GAIA	PDDL
Llama-7B	Single	45.3 _{±5.2}	2.9 _{±4.0}	16.8 _{±10.7}	15.8 _{±6.0}	22.5 _{±2.9}	20.1 _{±4.8}	15.3 _{±2.9}	24.2 _{±2.0}
	BoN	50.5 _{±0.0}	6.9 _{±0.0}	27.5 _{±0.0}	21.8 _{±0.0}	25.4 _{±0.0}	24.9 _{±0.0}	18.2 _{±0.0}	26.2 _{±0.0}
	Debate	47.3 _{±3.2}	6.4 _{±0.5}	26.2 _{±1.3}	19.8 _{±2.0}	25.3 _{±0.1}	23.8 _{±1.1}	16.5 _{±1.7}	26.2 _{±0.0}
	DyLAN	48.1 _{±2.4}	6.6 _{±0.3}	34.5 _{±7.0}	22.2 _{±0.4}	24.7 _{±0.7}	21.8 _{±3.1}	16.8 _{±1.4}	25.5 _{±0.7}
	GPTSwarm	48.0 _{±2.5}	5.8 _{±1.1}	26.0 _{±1.5}	18.9 _{±2.9}	25.8 _{±0.4}	25.1 _{±0.2}	18.8 _{±0.6}	27.9 _{±1.7}
	MacNet	50.8 _{±0.3}	7.9 _{±1.0}	30.7 _{±3.2}	24.0 _{±2.2}	27.8 _{±2.4}	26.5 _{±1.6}	18.8 _{±0.6}	29.3 _{±3.1}
	M2CL (ours)	59.1 _{±8.6}	17.4 _{±10.5}	44.7 _{±17.2}	32.9 _{±11.1}	35.1 _{±9.7}	33.0 _{±8.1}	25.5 _{±7.3}	33.1 _{±6.9}
Llama-14B	Single	54.8 _{±16.2}	3.9 _{±4.9}	25.0 _{±12.7}	23.7 _{±8.3}	28.5 _{±2.7}	24.5 _{±7.2}	18.6 _{±3.2}	30.4 _{±2.8}
	BoN	71.0 _{±0.0}	8.8 _{±0.0}	37.7 _{±0.0}	32.0 _{±0.0}	31.2 _{±0.0}	31.7 _{±0.0}	21.8 _{±0.0}	33.2 _{±0.0}
	Debate	64.5 _{±6.5}	8.4 _{±0.4}	36.3 _{±1.4}	27.9 _{±4.1}	30.8 _{±0.4}	29.3 _{±2.4}	19.8 _{±2.0}	31.9 _{±1.3}
	DyLAN	69.7 _{±1.3}	10.3 _{±1.5}	46.5 _{±8.8}	32.5 _{±0.5}	31.3 _{±0.1}	28.7 _{±3.0}	19.4 _{±2.4}	30.4 _{±2.8}
	MacNet	70.3 _{±0.7}	9.2 _{±0.4}	40.9 _{±3.2}	35.2 _{±3.2}	31.5 _{±0.3}	31.8 _{±0.1}	23.4 _{±1.6}	35.3 _{±2.1}
	MacNet	70.3 _{±0.7}	9.2 _{±0.4}	40.9 _{±3.2}	35.2 _{±3.2}	31.5 _{±0.3}	31.8 _{±0.1}	23.4 _{±1.6}	35.3 _{±2.1}
	M2CL (ours)	86.9 _{±15.9}	23.4 _{±14.6}	58.9 _{±21.2}	47.5 _{±15.5}	42.9 _{±11.7}	39.9 _{±8.2}	30.5 _{±8.7}	40.9 _{±7.7}
Llama-70B	Single	68.9 _{±14.1}	31.6 _{±3.2}	27.6 _{±32.0}	35.5 _{±12.5}	46.1 _{±4.8}	40.2 _{±6.5}	29.9 _{±2.8}	48.4 _{±4.4}
	BoN	83.0 _{±0.0}	34.8 _{±0.0}	59.6 _{±0.0}	48.0 _{±0.0}	50.9 _{±0.0}	46.7 _{±0.0}	32.7 _{±0.0}	52.8 _{±0.0}
	Debate	74.8 _{±8.2}	29.8 _{±5.0}	50.2 _{±9.4}	43.2 _{±4.8}	50.4 _{±0.5}	47.8 _{±1.1}	32.4 _{±0.3}	52.9 _{±0.1}
	DyLAN	82.1 _{±0.9}	31.5 _{±3.3}	53.8 _{±5.8}	44.5 _{±3.5}	49.3 _{±1.6}	45.1 _{±1.6}	31.5 _{±1.2}	50.3 _{±2.5}
	GPTSwarm	81.4 _{±1.6}	31.5 _{±3.3}	54.8 _{±4.8}	43.7 _{±4.3}	51.0 _{±0.1}	45.2 _{±1.5}	32.5 _{±0.2}	51.5 _{±1.3}
	MacNet	72.9 _{±10.1}	32.7 _{±2.1}	49.1 _{±10.5}	40.8 _{±4.2}	52.7 _{±1.8}	55.9 _{±9.2}	36.5 _{±3.8}	57.0 _{±4.2}
	M2CL (ours)	95.6 _{±12.6}	40.3 _{±5.5}	78.7 _{±19.1}	70.6 _{±22.6}	69.5 _{±18.6}	65.5 _{±18.8}	49.6 _{±16.9}	66.5 _{±13.7}
Qwen-7B	Single	61.2 _{±13.0}	12.9 _{±12.0}	20.2 _{±16.2}	51.2 _{±11.3}	24.3 _{±6.6}	25.9 _{±6.6}	15.3 _{±4.7}	21.1 _{±4.1}
	BoN	74.2 _{±0.0}	24.9 _{±0.0}	36.4 _{±0.0}	62.5 _{±0.0}	30.9 _{±0.0}	32.5 _{±0.0}	20.0 _{±0.0}	25.2 _{±0.0}
	Debate	69.3 _{±4.9}	17.9 _{±7.0}	28.7 _{±7.7}	54.5 _{±8.0}	29.5 _{±1.4}	31.4 _{±1.1}	20.6 _{±0.6}	24.3 _{±0.9}
	DyLAN	74.2 _{±0.0}	23.3 _{±1.6}	34.2 _{±2.2}	62.5 _{±0.0}	29.5 _{±1.4}	28.3 _{±4.2}	19.0 _{±1.0}	23.0 _{±2.2}
	GPTSwarm	74.5 _{±0.3}	24.1 _{±0.8}	37.2 _{±0.8}	60.4 _{±2.1}	29.4 _{±1.5}	29.5 _{±3.0}	20.0 _{±0.0}	23.9 _{±1.3}
	MacNet	68.0 _{±6.2}	20.3 _{±4.6}	25.0 _{±11.4}	57.1 _{±5.4}	29.8 _{±1.1}	33.8 _{±1.3}	20.5 _{±0.5}	26.6 _{±1.4}
	M2CL (ours)	88.7 _{±14.5}	40.8 _{±15.9}	58.4 _{±22.0}	76.8 _{±14.3}	36.7 _{±5.8}	41.7 _{±9.2}	30.2 _{±10.2}	33.0 _{±7.8}
Qwen-14B	Single	67.2 _{±12.5}	21.6 _{±6.2}	21.2 _{±11.6}	56.7 _{±12.7}	30.2 _{±6.0}	31.7 _{±7.7}	19.2 _{±4.6}	26.4 _{±2.4}
	BoN	79.7 _{±0.0}	27.8 _{±0.0}	32.8 _{±0.0}	69.4 _{±0.0}	36.2 _{±0.0}	39.4 _{±0.0}	23.8 _{±0.0}	28.8 _{±0.0}
	Debate	71.7 _{±8.0}	24.7 _{±3.1}	24.8 _{±8.0}	62.7 _{±4.7}	35.0 _{±0.6}	38.5 _{±0.9}	25.5 _{±1.7}	29.9 _{±1.1}
	DyLAN	77.5 _{±2.2}	26.9 _{±0.9}	31.9 _{±0.9}	67.9 _{±1.5}	33.6 _{±2.6}	34.9 _{±4.5}	22.7 _{±1.1}	27.8 _{±1.0}
	GPTSwarm	80.1 _{±0.4}	26.9 _{±0.9}	33.0 _{±0.2}	68.1 _{±1.3}	34.6 _{±0.6}	34.8 _{±4.6}	24.5 _{±0.7}	28.3 _{±0.5}
	MacNet	71.2 _{±8.5}	24.6 _{±3.2}	27.6 _{±5.2}	62.9 _{±6.5}	36.1 _{±0.1}	41.4 _{±2.0}	28.6 _{±4.8}	29.4 _{±0.6}
	M2CL (ours)	91.4 _{±11.7}	43.6 _{±15.8}	64.6 _{±31.8}	86.5 _{±17.1}	45.0 _{±8.8}	51.3 _{±11.9}	37.2 _{±13.4}	39.7 _{±10.9}
Qwen-72B	Single	72.5 _{±11.7}	31.6 _{±19.4}	34.9 _{±11.0}	59.1 _{±13.1}	49.3 _{±7.3}	51.1 _{±8.7}	31.4 _{±7.2}	41.8 _{±8.0}
	BoN	84.2 _{±0.0}	51.0 _{±0.0}	45.9 _{±0.0}	72.2 _{±0.0}	56.6 _{±0.0}	59.8 _{±0.0}	38.6 _{±0.0}	49.8 _{±0.0}
	Debate	76.6 _{±7.6}	48.0 _{±3.0}	39.6 _{±6.3}	66.7 _{±5.5}	58.3 _{±1.7}	62.5 _{±2.7}	40.7 _{±2.1}	47.8 _{±2.0}
	DyLAN	84.6 _{±0.4}	61.5 _{±10.5}	45.5 _{±0.4}	71.3 _{±0.9}	54.2 _{±2.4}	57.8 _{±2.0}	38.2 _{±0.4}	45.9 _{±3.9}
	GPTSwarm	84.9 _{±0.7}	60.9 _{±9.9}	46.6 _{±0.7}	69.7 _{±2.5}	55.2 _{±1.4}	59.1 _{±0.7}	39.0 _{±0.4}	48.8 _{±1.0}
	MacNet	80.3 _{±3.9}	48.1 _{±12.9}	41.2 _{±4.7}	68.1 _{±4.1}	57.2 _{±0.6}	64.1 _{±4.3}	42.0 _{±3.4}	49.8 _{±0.0}
	M2CL (ours)	93.5 _{±9.3}	69.3 _{±18.3}	73.1 _{±27.2}	88.3 _{±16.1}	73.9 _{±17.3}	81.9 _{±22.1}	60.3 _{±21.7}	65.3 _{±15.5}

Table 5: Accuracy with varying number of LLMs on different datasets. The number of LLMs is 8 for all datasets. We exhibit the performance advantage with BoN and highlight the best result.

Model	Method	MMLU	MATH	GPQA	Code	ALFWorld	SciWorld	GAIA	PDDL
Llama-7B	Single	45.3 _{±5.2}	2.9 _{±4.0}	16.8 _{±10.7}	15.8 _{±6.0}	22.6 _{±3.6}	19.7 _{±5.3}	14.9 _{±4.0}	23.6 _{±3.1}
	BoN	50.5 _{±0.0}	6.9 _{±0.0}	27.5 _{±0.0}	21.8 _{±0.0}	26.2 _{±0.0}	25.0 _{±0.0}	18.9 _{±0.0}	26.7 _{±0.0}
	Debate	49.5 _{±1.0}	6.2 _{±0.7}	24.8 _{±2.7}	19.2 _{±2.6}	25.9 _{±0.3}	24.7 _{±0.3}	15.9 _{±3.0}	26.9 _{±0.2}
	DyLAN	54.2 _{±3.7}	9.7 _{±2.8}	35 _{±7.5}	26 _{±4.2}	25.0 _{±1.2}	22.0 _{±3.0}	17.2 _{±1.7}	25.4 _{±1.3}
	GPTSwarm	49.2 _{±1.3}	6.8 _{±0.1}	28.7 _{±1.2}	20.7 _{±1.1}	26.2 _{±0.0}	26.4 _{±1.4}	20.0 _{±1.1}	29.3 _{±2.6}
	MacNet	56.5 _{±6.0}	11.0 _{±4.1}	39.6 _{±12.1}	30.0 _{±8.2}	30.4 _{±4.2}	30.2 _{±5.2}	20.6 _{±1.7}	31.7 _{±5.0}
	M2CL (ours)	63.8 _{±13.3}	24.9 _{±18.0}	59.3 _{±31.8}	38.8 _{±17.0}	37.2 _{±11.0}	36.1 _{±11.1}	27.1 _{±8.2}	34.7 _{±8.0}
Llama-14B	Single	54.8 _{±16.2}	3.9 _{±4.9}	25.0 _{±12.7}	23.7 _{±8.3}	27.8 _{±4.1}	24.1 _{±8.2}	18.6 _{±3.9}	29.4 _{±5.4}
	BoN	71.0 _{±0.0}	8.8 _{±0.0}	37.7 _{±0.0}	32.0 _{±0.0}	31.9 _{±0.0}	32.3 _{±0.0}	22.5 _{±0.0}	34.8 _{±0.0}
	Debate	66.9 _{±4.1}	8.0 _{±0.8}	34.5 _{±3.2}	30.8 _{±4.2}	31.7 _{±0.2}	30.3 _{±2.0}	20.0 _{±2.5}	32.0 _{±2.8}
	DyLAN	82.1 _{±11.1}	12.4 _{±3.6}	46.7 _{±9.0}	37.5 _{±5.5}	31.7 _{±0.2}	29.6 _{±2.7}	19.4 _{±3.1}	31.3 _{±3.5}
	MacNet	81.7 _{±10.7}	14.2 _{±5.4}	53.9 _{±16.2}	43.3 _{±11.3}	33.4 _{±1.5}	35.6 _{±3.3}	24.9 _{±2.4}	37.0 _{±2.2}
	MacNet	81.7 _{±10.7}	14.2 _{±5.4}	53.9 _{±16.2}	43.3 _{±11.3}	33.4 _{±1.5}	35.6 _{±3.3}	24.9 _{±2.4}	37.0 _{±2.2}
	M2CL (ours)	92.7 _{±21.7}	23.0 _{±14.2}	70.7 _{±33.0}	55.4 _{±23.4}	46.1 _{±14.2}	44.7 _{±12.4}	33.6 _{±11.1}	43.8 _{±9.0}
Llama-70B	Single	68.9 _{±14.1}	31.6 _{±3.2}	27.6 _{±32.0}	35.5 _{±12.5}	44.1 _{±8.3}	39.1 _{±9.7}	29.1 _{±3.8}	47.1 _{±5.7}
	BoN	83.0 _{±0.0}	34.8 _{±0.0}	59.6 _{±0.0}	48.0 _{±0.0}	52.4 _{±0.0}	48.8 _{±0.0}	32.9 _{±0.0}	52.8 _{±0.0}
	Debate	82.2 _{±0.8}	34.3 _{±0.5}	58.1 _{±1.5}	47.1 _{±0.9}	51.7 _{±0.7}	50.1 _{±1.3}	32.2 _{±0.7}	53.3 _{±0.5}
	DyLAN	93.5 _{±10.5}	37.1 _{±2.3}	72.8 _{±13.2}	56.8 _{±8.8}	50.2 _{±2.2}	46.4 _{±2.4}	31.7 _{±1.2}	49.7 _{±3.1}
	GPTSwarm	93.6 _{±10.6}	36.6 _{±1.8}	71.2 _{±11.6}	57.1 _{±9.1}	51.4 _{±1.0}	46.5 _{±2.3}	32.2 _{±0.7}	52.2 _{±0.6}
	MacNet	82.0 _{±1.0}	36.9 _{±2.1}	62.2 _{±2.6}	48.8 _{±0.8}	57.3 _{±4.9}	64.0 _{±15.2}	40.7 _{±7.8}	62.6 _{±9.8}
	M2CL (ours)	95.4 _{±12.4}	43.5 _{±8.7}	87.6 _{±28.0}	81.0 _{±33.0}	78.9 _{±26.5}	76.2 _{±27.4}	58.7 _{±25.8}	73.6 _{±20.8}
Qwen-7B	Single	61.2 _{±13.0}	12.9 _{±12.0}	20.2 _{±16.2}	51.2 _{±11.3}	23.8 _{±7.7}	25.2 _{±10.1}	15.6 _{±5.5}	21.0 _{±5.3}
	BoN	74.2 _{±0.0}	24.9 _{±0.0}	36.4 _{±0.0}	62.5 _{±0.0}	31.5 _{±0.0}	35.3 _{±0.0}	21.1 _{±0.0}	26.3 _{±0.0}
	Debate	71.1 _{±3.1}	19.9 _{±5.0}	28.7 _{±7.7}	60.0 _{±2.5}	30.6 _{±0.9}	32.2 _{±3.1}	21.0 _{±0.1}	24.8 _{±1.5}
	DyLAN	74.3 _{±0.1}	26.7 _{±1.8}	35.4 _{±1.0}	63.4 _{±0.9}	29.8 _{±1.7}	29.4 _{±5.9}	18.4 _{±2.7}	23.4 _{±2.9}
	GPTSwarm	76.3 _{±2.1}	26.2 _{±1.3}	35.6 _{±0.8}	62.7 _{±0.2}	29.9 _{±1.6}	30.8 _{±4.5}	20.6 _{±0.5}	24.5 _{±1.8}
	MacNet	71.5 _{±2.7}	21.4 _{±3.5}	30.8 _{±5.6}	59.3 _{±3.2}	33.6 _{±2.1}	37.0 _{±1.7}	21.5 _{±0.4}	29.5 _{±3.2}
	M2CL (ours)	92.5 _{±18.3}	47.8 _{±22.9}	66.1 _{±29.7}	80.3 _{±17.8}	39.9 _{±8.4}	45.3 _{±10.0}	33.6 _{±12.5}	34.7 _{±8.4}
Qwen-14B	Single	67.2 _{±12.5}	21.6 _{±6.2}	21.2 _{±11.6}	56.7 _{±12.7}	30.1 _{±7.4}	31.3 _{±10.3}	18.7 _{±6.9}	25.7 _{±2.8}
	BoN	79.7 _{±0.0}	27.8 _{±0.0}	32.8 _{±0.0}	69.4 _{±0.0}	37.5 _{±0.0}	41.6 _{±0.0}	25.6 _{±0.0}	28.5 _{±0.0}
	Debate	77.2 _{±2.5}	27.4 _{±4.0}	30.3 _{±2.5}	66.9 _{±2.5}	36.8 _{±0.7}	39.4 _{±2.2}	26.4 _{±0.8}	30.0 _{±1.5}
	DyLAN	86.8 _{±7.1}	31.2 _{±3.4}	39.4 _{±6.6}	76.6 _{±7.2}	34.8 _{±2.7}	36.0 _{±5.6}	22.8 _{±2.8}	28.0 _{±0.5}
	GPTSwarm	87.0 _{±7.3}	31.3 _{±3.5}	38.7 _{±5.9}	75.9 _{±6.5}	34.9 _{±2.6}	35.9 _{±5.7}	25.0 _{±0.6}	28.4 _{±0.1}
	MacNet	78.9 _{±0.8}	28.7 _{±0.9}	31.9 _{±0.9}	70.4 _{±1.0}	39.2 _{±1.7}	46.0 _{±4.4}	32.2 _{±6.6}	32.7 _{±4.2}
	M2CL (ours)	93.7 _{±14.0}	51.7 _{±23.9}	66.2 _{±33.4}	91.1 _{±21.7}	48.2 _{±10.7}	56.1 _{±14.5}	42.0 _{±16.4}	43.0 _{±14.5}
Qwen-72B	Single	72.5 _{±11.7}	31.6 _{±19.4}	34.9 _{±11.0}	59.1 _{±13.1}	48.2 _{±9.3}	50.4 _{±11.7}	31.1 _{±10.1}	41.0 _{±10.0}
	BoN	84.2 _{±0.0}	51.0 _{±0.0}	45.9 _{±0.0}	72.2 _{±0.0}	57.5 _{±0.0}	62.1 _{±0.0}	41.2 _{±0.0}	51.0 _{±0.0}
	Debate	82.7 _{±1.5}	48.4 _{±2.6}	43.4 _{±2.5}	69.1 _{±3.1}	60.4 _{±2.9}	65.2 _{±3.1}	42.9 _{±1.7}	49.1 _{±1.9}
	DyLAN	91.5 _{±7.3}	63.1 _{±12.1}	51.6 _{±5.7}	80.4 _{±8.2}	55.1 _{±2.4}	58.0 _{±4.1}	40.4 _{±0.8}	45.5 _{±5.5}
	GPTSwarm	91.5 _{±7.3}	64.7 _{±13.7}	52.4 _{±6.5}	79.6 _{±7.4}	56.8 _{±0.7}	60.2 _{±1.9}	40.3 _{±0.9}	49.7 _{±1.3}
	MacNet	83.8 _{±0.4}	52.9 _{±1.9}	46.2 _{±0.3}	70.5 _{±1.7}	61.8 _{±4.3}	68.4 _{±6.3}	46.4 _{±5.2}	53.7 _{±2.7}
	M2CL (ours)	95.1 _{±10.9}	72.5 _{±21.5}	78.9 _{±33.0}	90.7 _{±18.5}	79.0 _{±21.5}	88.9 _{±26.8}	67.2 _{±26.0}	70.5 _{±19.5}

Table 6: Accuracy with varying number of LLMs on different datasets. The number of LLMs is 16 for all datasets. We exhibit the performance advantage with BoN and highlight the best result.

Model	Method	MMLU	MATH	GPQA	Code	ALFWorld	SciWorld	GAIA	PDDL
Llama-7B	Single	45.3 _{±5.2}	2.9 _{±4.0}	16.8 _{±10.7}	15.8 _{±6.0}	22.4 _{±4.3}	19.6 _{±6.1}	14.7 _{±4.5}	24.2 _{±2.1}
	BoN	50.5 _{±0.0}	6.9 _{±0.0}	27.5 _{±0.0}	21.8 _{±0.0}	26.7 _{±0.0}	25.7 _{±0.0}	19.2 _{±0.0}	26.3 _{±0.0}
	Debate	51 _{±0.5}	6.4 _{±0.5}	26.1 _{±1.4}	21.3 _{±0.5}	25.7 _{±1.0}	25.3 _{±0.4}	16.6 _{±2.6}	26.4 _{±0.1}
	DyLAN	54.5 _{±4.0}	9.9 _{±3.0}	35.5 _{±8.0}	26.2 _{±4.4}	24.6 _{±2.1}	22.7 _{±3.0}	17.2 _{±2.0}	25.8 _{±0.5}
	GPTSwarm	51.6 _{±1.1}	6.9 _{±0.0}	29.1 _{±1.6}	23.0 _{±1.2}	27.2 _{±0.5}	26.8 _{±1.1}	20.8 _{±1.6}	29.7 _{±3.4}
	MacNet	59.3 _{±8.8}	12.5 _{±5.6}	45.0 _{±17.5}	33.7 _{±11.9}	32.8 _{±6.1}	33.4 _{±7.7}	22.6 _{±3.4}	34.2 _{±7.9}
	M2CL (ours)	71.5 _{±21.0}	23.9 _{±17.0}	69.9 _{±42.4}	48.3 _{±26.5}	39.8 _{±13.1}	38.4 _{±12.7}	28.9 _{±9.7}	37.0 _{±10.7}
Llama-14B	Single	54.8 _{±16.2}	3.9 _{±4.9}	25.0 _{±12.7}	23.7 _{±8.3}	27.4 _{±5.6}	24.0 _{±9.5}	18.3 _{±5.0}	29.2 _{±4.6}
	BoN	71.0 _{±0.0}	8.8 _{±0.0}	37.7 _{±0.0}	32.0 _{±0.0}	33.0 _{±0.0}	33.5 _{±0.0}	23.3 _{±0.0}	33.8 _{±0.0}
	Debate	72.2 _{±1.2}	8.2 _{±0.6}	37.0 _{±0.7}	32.0 _{±0.0}	32.0 _{±1.0}	30.9 _{±2.6}	19.6 _{±3.7}	32.7 _{±1.1}
	DyLAN	83.2 _{±12.2}	12.4 _{±3.6}	47.0 _{±9.3}	38.1 _{±6.1}	31.9 _{±1.1}	30.2 _{±3.3}	20.0 _{±3.3}	31.2 _{±2.6}
	MacNet	86.8 _{±15.8}	16.0 _{±7.2}	59.7 _{±22.0}	49.1 _{±17.1}	34.8 _{±1.8}	37.8 _{±4.3}	26.7 _{±3.4}	39.9 _{±6.1}
	MacNet	86.8 _{±15.8}	16.0 _{±7.2}	59.7 _{±22.0}	49.1 _{±17.1}	34.8 _{±1.8}	37.8 _{±4.3}	26.7 _{±3.4}	39.9 _{±6.1}
	M2CL (ours)	94.5 _{±23.5}	32.7 _{±23.9}	88.5 _{±50.8}	66.1 _{±34.1}	48.7 _{±15.7}	47.1 _{±13.6}	36.2 _{±12.9}	45.5 _{±11.7}
Llama-70B	Single	68.9 _{±14.1}	31.6 _{±3.2}	27.6 _{±32.0}	35.5 _{±12.5}	43.7 _{±10.2}	38.3 _{±11.5}	28.6 _{±4.6}	46.8 _{±6.2}
	BoN	83.0 _{±0.0}	34.8 _{±0.0}	59.6 _{±0.0}	48.0 _{±0.0}	53.9 _{±0.0}	49.8 _{±0.0}	33.2 _{±0.0}	53.0 _{±0.0}
	Debate	82.9 _{±0.1}	34.6 _{±0.2}	58.5 _{±4.1}	46.8 _{±1.2}	52.2 _{±1.7}	51.2 _{±1.4}	33.0 _{±0.2}	52.8 _{±0.2}
	DyLAN	93.2 _{±10.2}	37.2 _{±2.4}	73.2 _{±13.6}	57.2 _{±9.2}	50.5 _{±3.0}	46.8 _{±3.0}	31.1 _{±2.1}	49.7 _{±3.3}
	GPTSwarm	93.0 _{±10.0}	37.4 _{±2.6}	69.7 _{±10.1}	55.5 _{±7.5}	51.7 _{±2.2}	47.5 _{±2.3}	32.4 _{±0.8}	51.6 _{±1.4}
	MacNet	86.7 _{±3.7}	38.9 _{±4.1}	68.3 _{±8.7}	49.8 _{±1.8}	61.6 _{±7.7}	70.8 _{±21.0}	43.8 _{±10.6}	67.1 _{±14.1}
	M2CL (ours)	93.9 _{±10.9}	51.5 _{±16.7}	91.3 _{±31.7}	97.2 _{±49.2}	84.9 _{±31.0}	81.7 _{±31.9}	61.5 _{±28.3}	78.1 _{±25.1}
Qwen-7B	Single	61.2 _{±13.0}	12.9 _{±12.0}	20.2 _{±16.2}	51.2 _{±11.3}	23.4 _{±9.5}	24.9 _{±10.5}	15.6 _{±6.4}	20.6 _{±6.5}
	BoN	74.2 _{±0.0}	24.9 _{±0.0}	36.4 _{±0.0}	62.5 _{±0.0}	32.9 _{±0.0}	35.4 _{±0.0}	22.0 _{±0.0}	27.1 _{±0.0}
	Debate	73.4 _{±0.8}	24.1 _{±0.8}	35.6 _{±0.8}	62.1 _{±0.4}	31.7 _{±1.2}	32.6 _{±2.8}	22.1 _{±0.1}	24.2 _{±2.9}
	DyLAN	77.4 _{±3.2}	26.0 _{±1.1}	38.9 _{±2.5}	64.8 _{±2.3}	30.4 _{±2.5}	28.7 _{±6.7}	19.2 _{±2.8}	22.7 _{±4.4}
	GPTSwarm	78.0 _{±3.8}	27.6 _{±2.7}	41.1 _{±4.7}	66.2 _{±3.7}	30.1 _{±2.8}	30.4 _{±5.0}	22.1 _{±0.1}	24.4 _{±2.7}
	MacNet	76.8 _{±2.6}	27.8 _{±2.9}	40.3 _{±3.9}	69.4 _{±6.9}	35.4 _{±2.5}	39.7 _{±4.3}	24.6 _{±2.6}	30.7 _{±3.6}
	M2CL (ours)	94.6 _{±20.4}	47.4 _{±22.5}	72.8 _{±36.4}	84.9 _{±22.4}	41.8 _{±8.9}	47.6 _{±12.2}	36.5 _{±14.5}	37.4 _{±10.3}
Qwen-14B	Single	67.2 _{±12.5}	21.6 _{±6.2}	21.2 _{±11.6}	56.7 _{±12.7}	28.3 _{±9.8}	30.9 _{±11.4}	18.5 _{±7.6}	25.5 _{±3.8}
	BoN	79.7 _{±0.0}	27.8 _{±0.0}	32.8 _{±0.0}	69.4 _{±0.0}	38.1 _{±0.0}	42.3 _{±0.0}	26.1 _{±0.0}	29.3 _{±0.0}
	Debate	78.7 _{±1.0}	26.7 _{±1.1}	33.4 _{±0.6}	70.1 _{±0.7}	37.7 _{±0.4}	40.5 _{±1.8}	26.7 _{±0.6}	30.4 _{±1.1}
	DyLAN	88.0 _{±8.3}	31.9 _{±4.1}	41.0 _{±8.2}	78.2 _{±8.8}	34.9 _{±3.2}	35.7 _{±6.6}	23.9 _{±2.2}	28.3 _{±1.0}
	GPTSwarm	88.0 _{±8.3}	31.8 _{±4.0}	40.4 _{±7.6}	78.4 _{±9.0}	35.5 _{±2.6}	36.8 _{±5.5}	25.1 _{±1.0}	28.7 _{±0.6}
	MacNet	83.1 _{±3.4}	30.8 _{±3.0}	35.1 _{±2.3}	74.8 _{±5.4}	42.1 _{±4.0}	49.3 _{±7.0}	35.5 _{±9.4}	33.8 _{±4.5}
	M2CL (ours)	95.9 _{±16.2}	52.7 _{±24.9}	69.0 _{±36.2}	93.3 _{±23.9}	51.2 _{±13.1}	59.4 _{±17.1}	43.9 _{±17.8}	46.0 _{±16.7}
Qwen-72B	Single	72.5 _{±11.7}	31.6 _{±19.4}	34.9 _{±11.0}	59.1 _{±13.1}	47.7 _{±10.8}	50.1 _{±12.8}	31.2 _{±12.1}	40.8 _{±11.5}
	BoN	84.2 _{±0.0}	51.0 _{±0.0}	45.9 _{±0.0}	72.2 _{±0.0}	58.5 _{±0.0}	62.9 _{±0.0}	43.3 _{±0.0}	52.3 _{±0.0}
	Debate	83.9 _{±0.3}	56.1 _{±5.1}	44.1 _{±1.8}	71.6 _{±0.6}	60.8 _{±2.3}	66.7 _{±3.8}	44.1 _{±0.8}	49.3 _{±3.0}
	DyLAN	92.7 _{±8.5}	65.8 _{±14.8}	53.6 _{±7.7}	81.4 _{±9.2}	56.1 _{±2.4}	59.1 _{±3.8}	41.1 _{±2.2}	46.3 _{±6.0}
	GPTSwarm	92.4 _{±8.2}	66.0 _{±15.0}	53.3 _{±7.4}	81.1 _{±8.9}	57.3 _{±1.2}	59.8 _{±3.1}	41.3 _{±2.0}	50.6 _{±1.7}
	MacNet	88.1 _{±3.9}	60.6 _{±9.6}	49.2 _{±3.3}	72.6 _{±0.4}	65.8 _{±7.3}	74.7 _{±11.8}	50.3 _{±7.0}	58.1 _{±5.8}
	M2CL (ours)	96.6 _{±12.4}	74.3 _{±23.3}	80.0 _{±34.1}	94.4 _{±22.2}	82.9 _{±24.4}	94.6 _{±31.7}	71.9 _{±28.6}	74.3 _{±22.0}

Table 7: Accuracy with varying number of LLMs on different datasets. The number of LLMs is 32 for all datasets. We exhibit the performance advantage with BoN and highlight the best result.

Model	Method	MMLU	MATH	GPQA	Code	ALFWorld	SciWorld	GAIA	PDDL
Llama-7B	Single	45.3 _{±5.2}	2.9 _{±4.0}	16.8 _{±10.7}	15.8 _{±6.0}	22.4 _{±4.4}	18.8 _{±8.1}	14.7 _{±5.1}	23.2 _{±4.0}
	BoN	50.5 _{±0.0}	6.9 _{±0.0}	27.5 _{±0.0}	21.8 _{±0.0}	26.8 _{±0.0}	26.9 _{±0.0}	19.8 _{±0.0}	27.2 _{±0.0}
	Debate	51.8 _{±1.3}	7.7 _{±0.8}	29.5 _{±2.0}	23.5 _{±1.7}	26.4 _{±0.4}	25.6 _{±1.3}	15.8 _{±4.0}	26.2 _{±1.0}
	DyLAN	54.7 _{±4.2}	10.1 _{±3.2}	36.2 _{±8.7}	26.7 _{±4.9}	24.9 _{±1.9}	23.1 _{±3.8}	17.6 _{±2.2}	26.0 _{±1.2}
	GPTSwarm	52.4 _{±1.9}	7.9 _{±1.0}	30.8 _{±3.3}	23.7 _{±1.9}	27.3 _{±0.5}	27.9 _{±1.0}	21.8 _{±2.0}	30.4 _{±3.2}
	MacNet	62.7 _{±12.2}	14.0 _{±7.1}	50.3 _{±22.8}	37.2 _{±15.4}	35.7 _{±8.9}	37.5 _{±10.6}	23.8 _{±4.0}	37.5 _{±10.3}
	M2CL (ours)	79.1 _{±28.6}	28.8 _{±21.9}	84.5 _{±57.0}	56.5 _{±34.7}	42.5 _{±15.7}	40.7 _{±13.8}	31.4 _{±11.6}	38.9 _{±11.7}
Llama-14B	Single	54.8 _{±16.2}	3.9 _{±4.9}	25.0 _{±12.7}	23.7 _{±8.3}	27.0 _{±5.8}	23.5 _{±11.3}	18.1 _{±6.1}	28.9 _{±5.4}
	BoN	71.0 _{±0.0}	8.8 _{±0.0}	37.7 _{±0.0}	32.0 _{±0.0}	32.8 _{±0.0}	34.8 _{±0.0}	24.2 _{±0.0}	34.3 _{±0.0}
	Debate	73.1 _{±2.1}	9.5 _{±0.7}	40.7 _{±3.0}	34.3 _{±2.3}	31.9 _{±0.9}	31.7 _{±3.1}	19.9 _{±4.3}	33.0 _{±1.3}
	DyLAN	83.6 _{±12.6}	12.7 _{±3.9}	47.7 _{±10.0}	38.7 _{±6.7}	32.1 _{±0.7}	30.8 _{±4.0}	19.8 _{±4.4}	31.0 _{±3.3}
	MacNet	90.7 _{±19.7}	17.6 _{±8.8}	67.2 _{±29.5}	53.9 _{±21.9}	37.1 _{±4.3}	41.7 _{±6.9}	29.1 _{±4.9}	42.3 _{±8.0}
	MacNet	90.7 _{±19.7}	17.6 _{±8.8}	67.2 _{±29.5}	53.9 _{±21.9}	37.1 _{±4.3}	41.7 _{±6.9}	29.1 _{±4.9}	42.3 _{±8.0}
	M2CL (ours)	95.8 _{±24.8}	41.4 _{±32.6}	94.7 _{±57.0}	75.6 _{±43.6}	52.2 _{±19.4}	50.7 _{±15.9}	38.5 _{±14.3}	48.0 _{±13.7}
Llama-70B	Single	68.9 _{±14.1}	31.6 _{±3.2}	27.6 _{±32.0}	35.5 _{±12.5}	43.1 _{±10.0}	38.6 _{±11.5}	28.1 _{±4.7}	46.0 _{±7.2}
	BoN	83.0 _{±0.0}	34.8 _{±0.0}	59.6 _{±0.0}	48.0 _{±0.0}	53.1 _{±0.0}	50.1 _{±0.0}	32.8 _{±0.0}	53.2 _{±0.0}
	Debate	85.7 _{±2.7}	35.3 _{±0.5}	63.4 _{±3.8}	51.2 _{±3.2}	52.6 _{±0.5}	52.1 _{±2.0}	32.9 _{±0.1}	53.7 _{±0.5}
	DyLAN	94.5 _{±11.5}	37.3 _{±2.5}	74.6 _{±15.0}	57.8 _{±9.8}	50.1 _{±3.0}	47.2 _{±2.9}	31.4 _{±1.4}	49.7 _{±3.5}
	GPTSwarm	96.4 _{±13.4}	37.8 _{±3.0}	73.4 _{±13.8}	56.4 _{±8.4}	52.2 _{±0.9}	48.4 _{±1.7}	32.5 _{±10.3}	51.6 _{±1.6}
	MacNet	90.0 _{±7.0}	41.2 _{±6.4}	73.9 _{±14.3}	54.1 _{±6.1}	64.5 _{±11.4}	76.6 _{±26.5}	46.5 _{±13.7}	71.6 _{±18.4}
	M2CL (ours)	97.0 _{±14.0}	54.5 _{±19.7}	95.1 _{±35.5}	93.7 _{±45.7}	88.5 _{±35.4}	86.0 _{±35.9}	65.7 _{±32.9}	79.8 _{±26.6}
Qwen-7B	Single	61.2 _{±13.0}	12.9 _{±12.0}	20.2 _{±16.2}	51.2 _{±11.3}	24.2 _{±8.6}	24.5 _{±13.3}	14.6 _{±7.7}	20.4 _{±7.3}
	BoN	74.2 _{±0.0}	24.9 _{±0.0}	36.4 _{±0.0}	62.5 _{±0.0}	32.8 _{±0.0}	37.8 _{±0.0}	22.3 _{±0.0}	27.7 _{±0.0}
	Debate	76.8 _{±2.6}	26.8 _{±1.9}	40.0 _{±3.6}	64.3 _{±1.8}	31.7 _{±1.1}	33.4 _{±4.4}	23.0 _{±0.7}	25.3 _{±2.4}
	DyLAN	78.2 _{±4.0}	29.5 _{±4.6}	40.3 _{±3.9}	66.5 _{±4.0}	30.7 _{±2.1}	29.2 _{±8.6}	19.0 _{±3.3}	22.6 _{±5.1}
	GPTSwarm	79.6 _{±5.4}	29.9 _{±5.0}	41.8 _{±5.4}	67.8 _{±5.3}	31.2 _{±1.6}	31.3 _{±6.5}	22.6 _{±0.3}	25.0 _{±2.7}
	MacNet	80.6 _{±6.4}	31.5 _{±6.6}	45.8 _{±9.4}	73.7 _{±11.2}	38.6 _{±5.8}	43.1 _{±5.3}	26.6 _{±4.3}	34.3 _{±6.6}
	M2CL (ours)	96.0 _{±21.8}	47.1 _{±22.2}	74.9 _{±38.5}	87.3 _{±24.8}	44.4 _{±11.6}	50.4 _{±12.6}	39.7 _{±17.4}	40.2 _{±12.5}
Qwen-14B	Single	67.2 _{±12.5}	21.6 _{±6.2}	21.2 _{±11.6}	56.7 _{±12.7}	29.4 _{±8.0}	29.5 _{±15.6}	18.9 _{±8.2}	24.5 _{±4.5}
	BoN	79.7 _{±0.0}	27.8 _{±0.0}	32.8 _{±0.0}	69.4 _{±0.0}	37.4 _{±0.0}	45.1 _{±0.0}	27.1 _{±0.0}	29.0 _{±0.0}
	Debate	82.7 _{±3.0}	29.1 _{±1.3}	35.5 _{±2.7}	71.8 _{±2.4}	39.0 _{±1.6}	41.7 _{±3.4}	28.7 _{±1.6}	30.8 _{±1.8}
	DyLAN	89.3 _{±9.6}	32.6 _{±4.8}	42.2 _{±9.4}	78.9 _{±9.5}	35.4 _{±2.0}	36.2 _{±8.9}	23.2 _{±3.9}	29.3 _{±0.3}
	GPTSwarm	89.0 _{±9.3}	32.6 _{±4.8}	42.2 _{±9.4}	78.9 _{±9.5}	36.4 _{±1.0}	37.5 _{±7.6}	26.6 _{±0.5}	28.9 _{±0.1}
	MacNet	87.8 _{±8.1}	33.7 _{±5.9}	40.1 _{±7.3}	83.0 _{±13.6}	44.8 _{±7.4}	54.2 _{±9.1}	39.8 _{±12.7}	36.2 _{±7.2}
	M2CL (ours)	96.1 _{±16.4}	58.0 _{±30.2}	70.5 _{±37.7}	94.9 _{±25.5}	54.9 _{±17.5}	62.1 _{±17.0}	49.0 _{±21.9}	49.6 _{±20.6}
Qwen-72B	Single	72.5 _{±11.7}	31.6 _{±19.4}	34.9 _{±11.0}	59.1 _{±13.1}	46.7 _{±12.6}	49.0 _{±14.2}	30.4 _{±13.5}	40.1 _{±13.2}
	BoN	84.2 _{±0.0}	51.0 _{±0.0}	45.9 _{±0.0}	72.2 _{±0.0}	59.3 _{±0.0}	63.2 _{±0.0}	43.9 _{±0.0}	53.3 _{±0.0}
	Debate	86.5 _{±2.3}	62.9 _{±11.9}	48.3 _{±2.4}	74.5 _{±2.3}	62.8 _{±3.5}	67.8 _{±4.6}	46.6 _{±2.7}	50.2 _{±3.1}
	DyLAN	92.8 _{±8.6}	67.7 _{±16.7}	54.8 _{±8.9}	82.7 _{±10.5}	56.0 _{±3.3}	59.1 _{±4.1}	42.9 _{±1.0}	46.7 _{±6.6}
	GPTSwarm	93.1 _{±8.9}	68.1 _{±17.1}	54.7 _{±8.8}	82.2 _{±10.0}	58.3 _{±1.0}	62.6 _{±0.6}	42.7 _{±1.2}	51.4 _{±1.9}
	MacNet	92.5 _{±8.3}	72.1 _{±21.1}	55.2 _{±9.3}	79.9 _{±7.7}	70.5 _{±11.2}	79.5 _{±16.3}	54.7 _{±10.8}	62.3 _{±9.0}
	M2CL (ours)	97.5 _{±13.3}	74.7 _{±23.7}	82.0 _{±36.1}	95.9 _{±23.7}	89.3 _{±30.0}	95 _{±31.8}	79.7 _{±35.8}	80.0 _{±26.7}

Table 8: Accuracy with varying number of LLMs on different datasets. The number of LLMs is 64 for all datasets. We exhibit the performance advantage with BoN and highlight the **best** result.

Model	Method	MMLU	MATH	GPQA	Code	ALFWorld	SciWorld	GAIA	PDDL
Llama-7B	Single	45.3 _{±5.2}	2.9 _{±4.0}	16.8 _{±10.7}	15.8 _{±6.0}	22.0 _{±5.2}	18.8 _{±7.9}	15.0 _{±4.8}	22.5 _{±4.4}
	BoN	50.5 _{±0.0}	6.9 _{±0.0}	27.5 _{±0.0}	21.8 _{±0.0}	27.2 _{±0.0}	26.7 _{±0.0}	19.8 _{±0.0}	26.9 _{±0.0}
	Debate	53.4 _{±2.9}	9.2 _{±2.3}	33.5 _{±6.0}	25 _{±3.2}	26.9 _{±0.3}	26.3 _{±0.4}	16.5 _{±3.3}	26.9 _{±0.0}
	DyLAN	55.1 _{±4.6}	10.4 _{±3.5}	36.9 _{±9.4}	27 _{±5.2}	24.6 _{±2.6}	23.3 _{±3.4}	17.9 _{±1.9}	25.0 _{±1.9}
	GPTSwarm	53.5 _{±3.0}	9 _{±2.1}	34 _{±6.5}	25.3 _{±3.5}	27.0 _{±0.2}	28.8 _{±2.1}	22.1 _{±2.3}	31.0 _{±4.1}
	MacNet	66.4 _{±15.9}	15.3 _{±8.4}	54.4 _{±26.9}	40.0 _{±18.2}	38.8 _{±11.6}	39.6 _{±12.9}	27.0 _{±7.2}	38.4 _{±11.5}
	M2CL (ours)	81.5 _{±31.0}	35.4 _{±28.5}	82.5 _{±55.0}	60.6 _{±38.8}	44.2 _{±17.0}	42.9 _{±16.2}	32.9 _{±13.1}	40.9 _{±14.0}
Llama-14B	Single	54.8 _{±16.2}	3.9 _{±4.9}	25.0 _{±12.7}	23.7 _{±8.3}	26.9 _{±5.7}	23.6 _{±12.1}	17.8 _{±6.7}	28.6 _{±5.6}
	BoN	71.0 _{±0.0}	8.8 _{±0.0}	37.7 _{±0.0}	32.0 _{±0.0}	32.6 _{±0.0}	35.7 _{±0.0}	24.5 _{±0.0}	34.2 _{±0.0}
	Debate	79.6 _{±8.6}	11.1 _{±2.3}	45.0 _{±7.3}	36.5 _{±4.5}	32.6 _{±0.0}	32.0 _{±3.7}	19.9 _{±4.6}	32.6 _{±1.6}
	DyLAN	85.0 _{±14.0}	13.0 _{±4.2}	48.8 _{±11.1}	39.3 _{±7.3}	31.9 _{±0.7}	31.0 _{±4.7}	19.9 _{±4.6}	30.5 _{±3.7}
	MacNet	90.8 _{±19.8}	18.1 _{±9.3}	67.5 _{±29.8}	55.2 _{±23.2}	39.1 _{±6.5}	44.4 _{±8.7}	31.1 _{±6.6}	44.0 _{±9.8}
	MacNet	90.8 _{±19.8}	18.1 _{±9.3}	67.5 _{±29.8}	55.2 _{±23.2}	39.1 _{±6.5}	44.4 _{±8.7}	31.1 _{±6.6}	44.0 _{±9.8}
	M2CL (ours)	96.8 _{±25.8}	40.0 _{±31.2}	93.1 _{±55.4}	82.5 _{±50.5}	54.8 _{±22.2}	52.8 _{±17.1}	40.3 _{±15.8}	49.8 _{±15.6}
Llama-70B	Single	68.9 _{±14.1}	31.6 _{±3.2}	27.6 _{±32.0}	35.5 _{±12.5}	43.0 _{±11.8}	37.3 _{±13.7}	28.2 _{±5.7}	45.9 _{±6.8}
	BoN	83.0 _{±0.0}	34.8 _{±0.0}	59.6 _{±0.0}	48.0 _{±0.0}	54.8 _{±0.0}	51.0 _{±0.0}	33.9 _{±0.0}	52.7 _{±0.0}
	Debate	90.0 _{±7.0}	36.6 _{±1.8}	69.0 _{±9.4}	54.0 _{±6.0}	53.6 _{±1.2}	53.5 _{±2.5}	32.7 _{±1.2}	53.2 _{±0.5}
	DyLAN	95.4 _{±12.4}	37.6 _{±2.8}	75.4 _{±15.8}	58.9 _{±10.9}	50.4 _{±4.4}	47.6 _{±3.4}	30.8 _{±3.1}	50.4 _{±2.3}
	GPTSwarm	95.7 _{±12.7}	37.9 _{±3.1}	78.9 _{±19.3}	58.1 _{±10.1}	52.4 _{±2.4}	49.3 _{±1.7}	32.3 _{±1.6}	51.3 _{±1.4}
	MacNet	95.3 _{±12.3}	44.2 _{±9.4}	84.2 _{±24.6}	57.9 _{±9.9}	67.9 _{±13.1}	81.6 _{±30.6}	49.0 _{±15.1}	75.4 _{±22.7}
	M2CL (ours)	96.3 _{±13.3}	58.0 _{±23.2}	96.9 _{±37.3}	92.0 _{±44.0}	90.8 _{±36.0}	88.9 _{±37.9}	68.2 _{±34.3}	82.2 _{±29.5}
Qwen-7B	Single	61.2 _{±13.0}	12.9 _{±12.0}	20.2 _{±16.2}	51.2 _{±11.3}	23.0 _{±11.5}	24.4 _{±14.2}	14.8 _{±8.2}	20.0 _{±7.8}
	BoN	74.2 _{±0.0}	24.9 _{±0.0}	36.4 _{±0.0}	62.5 _{±0.0}	34.5 _{±0.0}	38.6 _{±0.0}	23.0 _{±0.0}	27.8 _{±0.0}
	Debate	78.1 _{±3.9}	28.9 _{±4.0}	41.7 _{±5.3}	66.6 _{±7.4}	32.3 _{±2.2}	34.1 _{±4.5}	23.5 _{±0.5}	26.0 _{±1.8}
	DyLAN	79.8 _{±5.6}	30.7 _{±5.8}	43.5 _{±7.1}	67.5 _{±5.0}	30.7 _{±3.8}	30.0 _{±8.6}	19.8 _{±3.2}	23.7 _{±4.1}
	GPTSwarm	80.9 _{±6.7}	29.9 _{±5.0}	43.1 _{±6.7}	68.0 _{±5.5}	31.8 _{±2.7}	32.7 _{±5.9}	23.2 _{±0.2}	25.3 _{±2.5}
	MacNet	85.1 _{±10.9}	34.7 _{±9.8}	50.4 _{±14.0}	79.8 _{±17.3}	41.6 _{±7.1}	46.7 _{±8.1}	28.4 _{±5.4}	37.1 _{±9.3}
	M2CL (ours)	97.5 _{±23.3}	50.7 _{±25.8}	80.6 _{±44.2}	90.2 _{±27.7}	46.8 _{±12.3}	53.3 _{±14.7}	42.0 _{±19.0}	42.2 _{±14.4}
Qwen-14B	Single	67.2 _{±12.5}	21.6 _{±6.2}	21.2 _{±11.6}	56.7 _{±12.7}	28.9 _{±9.5}	29.9 _{±15.7}	17.8 _{±9.8}	24.7 _{±4.8}
	BoN	79.7 _{±0.0}	27.8 _{±0.0}	32.8 _{±0.0}	69.4 _{±0.0}	38.4 _{±0.0}	45.6 _{±0.0}	27.6 _{±0.0}	29.5 _{±0.0}
	Debate	86.6 _{±6.9}	31.1 _{±3.3}	39.9 _{±7.1}	75.8 _{±6.4}	38.8 _{±0.4}	42.7 _{±2.9}	29.0 _{±1.4}	31.7 _{±2.2}
	DyLAN	90.2 _{±10.5}	33.6 _{±5.8}	42.8 _{±10.0}	80.1 _{±10.7}	36.2 _{±2.2}	37.0 _{±8.6}	24.4 _{±3.2}	28.7 _{±0.8}
	GPTSwarm	90.2 _{±10.5}	33.0 _{±5.2}	43.2 _{±10.4}	80.4 _{±11.0}	35.9 _{±2.5}	36.0 _{±9.6}	26.5 _{±1.1}	29.0 _{±0.5}
	MacNet	93.7 _{±14.0}	37.7 _{±9.9}	46.7 _{±13.9}	91.2 _{±21.8}	47.5 _{±9.1}	56.8 _{±11.2}	42.7 _{±15.1}	38.4 _{±8.9}
	M2CL (ours)	99.1 _{±19.4}	61.9 _{±34.1}	80.4 _{±47.6}	98.7 _{±29.3}	56.4 _{±18.0}	65.8 _{±20.2}	51.8 _{±24.2}	50.9 _{±21.4}
Qwen-72B	Single	72.5 _{±11.7}	31.6 _{±19.4}	34.9 _{±11.0}	59.1 _{±13.1}	46.7 _{±13.4}	48.3 _{±17.7}	29.8 _{±15.1}	40.4 _{±14.0}
	BoN	84.2 _{±0.0}	51.0 _{±0.0}	45.9 _{±0.0}	72.2 _{±0.0}	60.1 _{±0.0}	66.0 _{±0.0}	44.9 _{±0.0}	54.4 _{±0.0}
	Debate	91.2 _{±7.0}	62.3 _{±11.3}	52.4 _{±6.5}	78.9 _{±6.7}	62.5 _{±2.4}	69.1 _{±3.1}	47.6 _{±2.7}	50.5 _{±3.9}
	DyLAN	95.1 _{±10.9}	66.8 _{±15.8}	55.7 _{±9.8}	83.9 _{±11.7}	56.9 _{±3.2}	59.9 _{±6.1}	43.7 _{±1.2}	47.7 _{±6.7}
	GPTSwarm	94.3 _{±10.1}	68.8 _{±17.8}	55.2 _{±9.3}	84.0 _{±11.8}	59.6 _{±0.5}	63.2 _{±2.8}	43.0 _{±1.9}	52.5 _{±1.9}
	MacNet	98.0 _{±13.8}	73.9 _{±22.9}	63.2 _{±17.3}	86.7 _{±14.5}	74.4 _{±14.3}	85.1 _{±19.1}	59.5 _{±14.6}	66.6 _{±12.2}
	M2CL (ours)	99.7 _{±15.5}	79.7 _{±28.7}	86.0 _{±40.1}	97.5 _{±25.3}	92.9 _{±32.8}	95.9 _{±29.9}	84.2 _{±39.3}	83.9 _{±29.5}

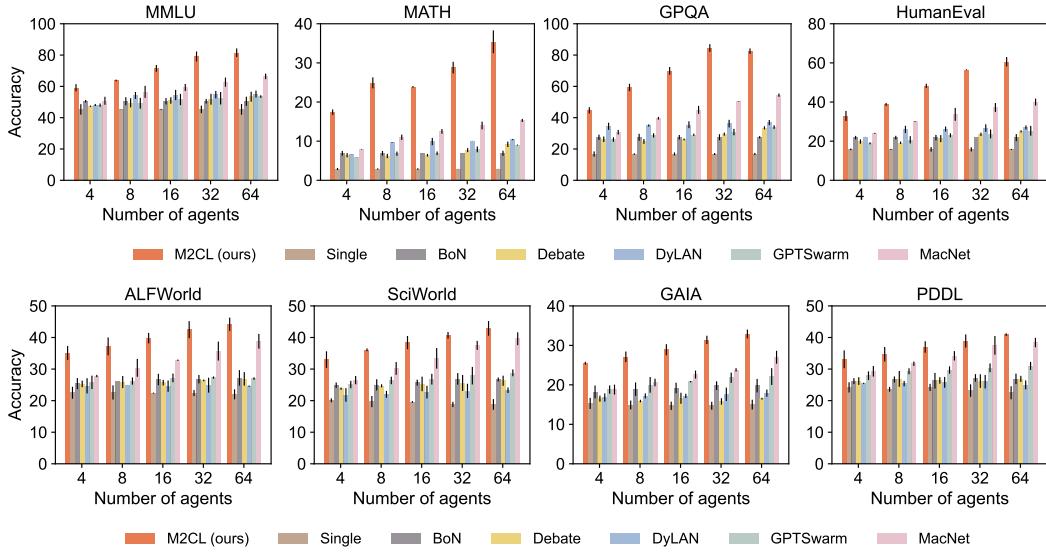


Figure 5: Performance of llama-7b as the base model with varying number of LLMs. Uncertainty intervals depict standard deviation over three seeds. M2CL exhibits higher performance and increasing tendency with more LLMs, demonstrating its great collaboration efficiency compared to existing methods. Of note, academic and agentic tasks reasoning are challenging because they require more diverse thinking perspectives and more rigorous analysis. The outperformance of M2CL reveals its capability of enabling LLMs to collaborate in changing discussion state.

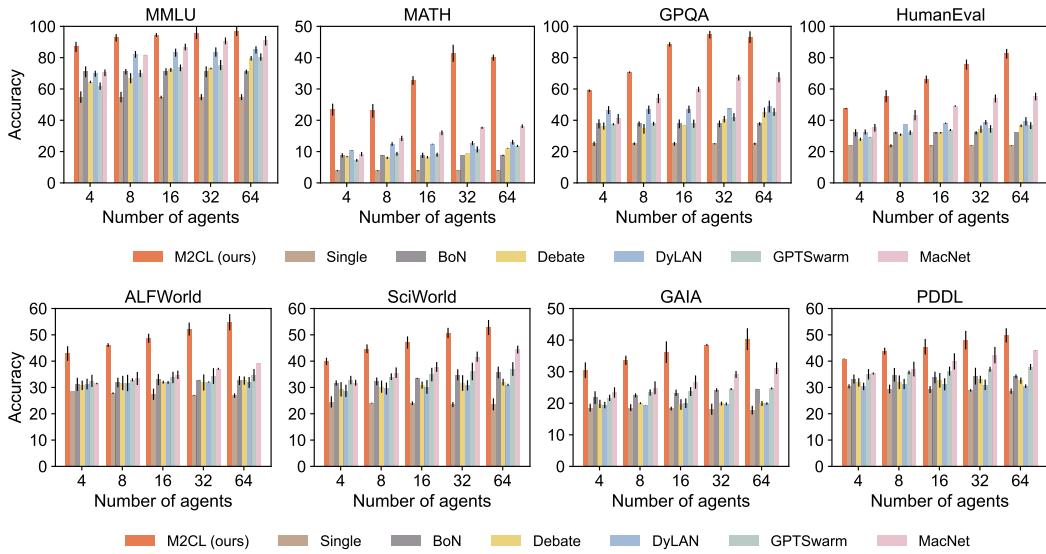


Figure 6: Performance of llama-13b as the base model with varying number of LLMs. Uncertainty intervals depict standard deviation over three seeds. M2CL exhibits higher performance and increasing tendency with more LLMs, demonstrating its great collaboration efficiency compared to existing methods.

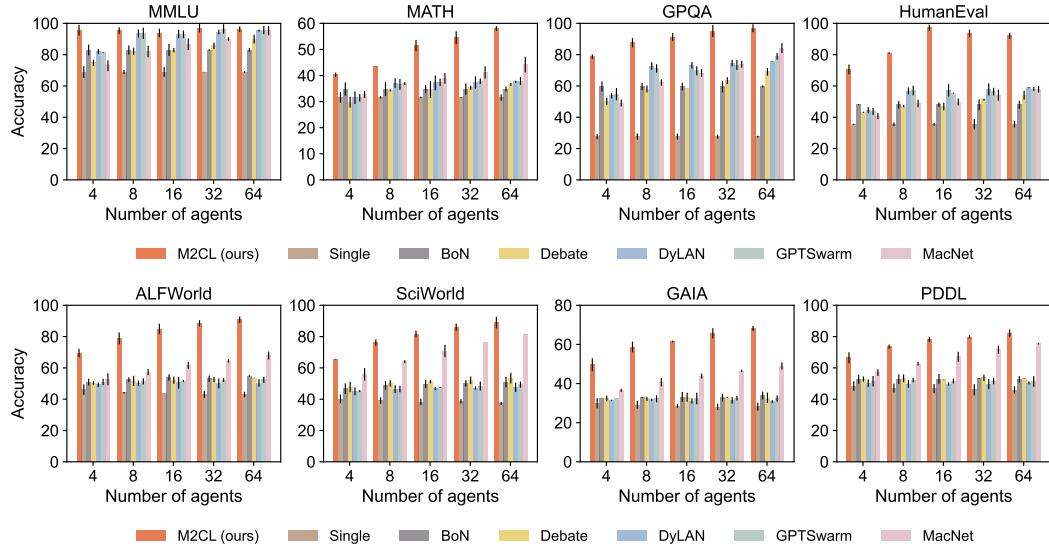


Figure 7: Performance as Ilama-70b as the base model with varying number of LLMs. Uncertainty intervals depict standard deviation over three seeds. M2CL exhibits higher performance and increasing tendency with more LLMs, demonstrating its great collaboration efficiency compared to existing methods.

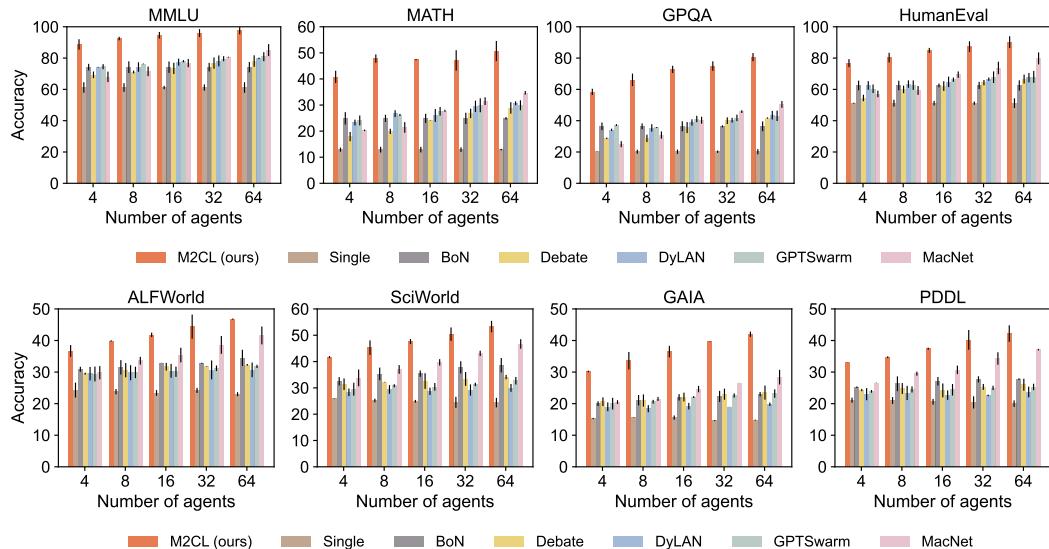


Figure 8: Performance of Qwen-7b as the base model with varying number of LLMs. Uncertainty intervals depict standard deviation over three seeds. M2CL exhibits higher performance and increasing tendency with more LLMs, demonstrating its great collaboration efficiency compared to existing methods.

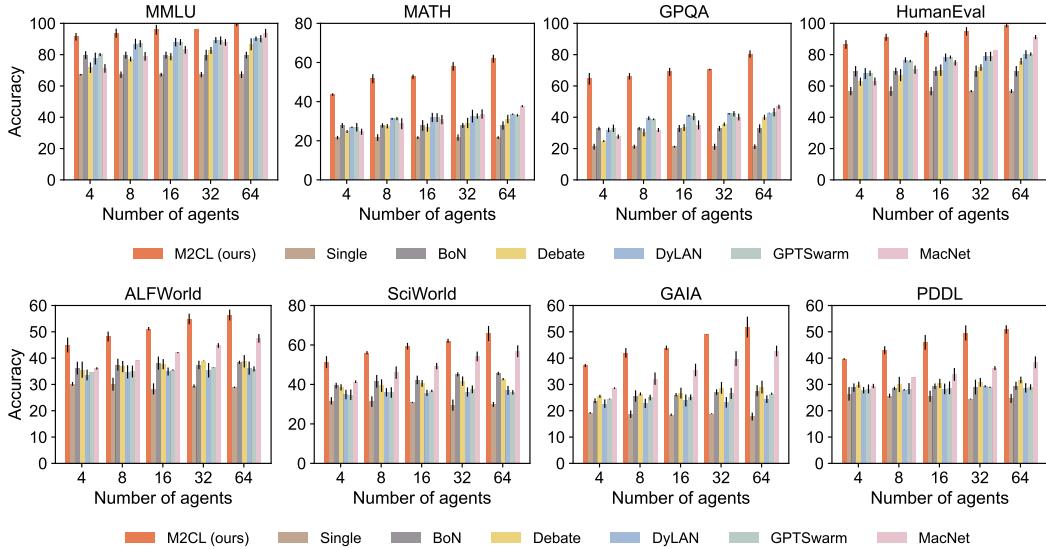


Figure 9: Performance of Qwen-14b as the base model with varying number of LLMs. Uncertainty intervals depict standard deviation over three seeds. M2CL exhibits higher performance and increasing tendency with more LLMs, demonstrating its great collaboration efficiency compared to existing methods.

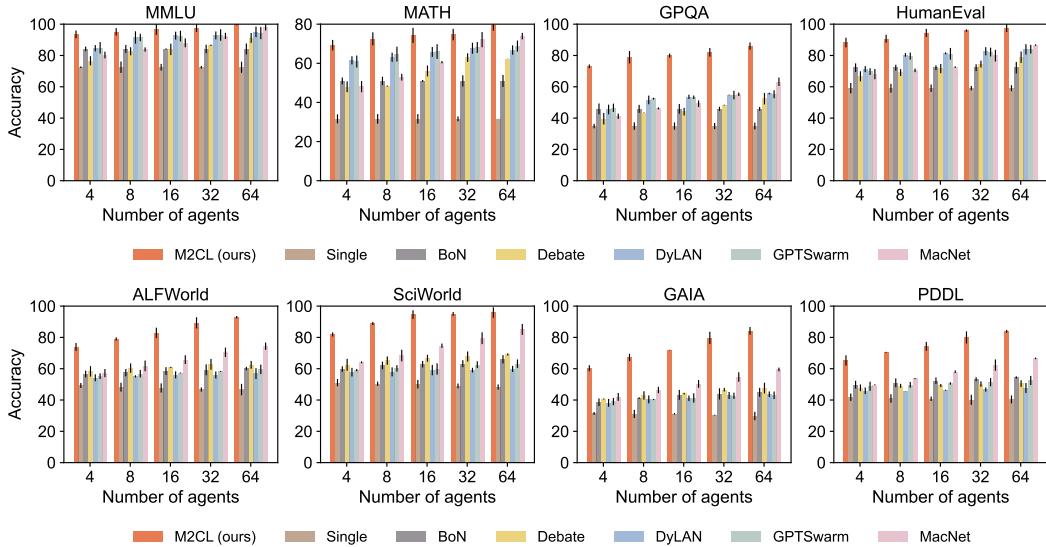


Figure 10: Performance of Qwen-72b as the base model with varying number of LLMs. Uncertainty intervals depict standard deviation over three seeds. M2CL exhibits higher performance and increasing tendency with more LLMs, demonstrating its great collaboration efficiency compared to existing methods.

G.3 GUI AGENT

We evaluate M2CL’s performance on more challenging AndroidWorld, which requires GUI identification, long-horizon planning, and accurate action execution capability. Comparative results are shown in Table 9 and Fig. 11.

We observe that M2CL consistently outperforms existing baselines across different model scales up to 50%, with performance gains becoming more pronounced as the number of participating LLMs increases. The superior scalability of M2CL in this setting highlights its ability to exploit diverse responses and maintain consistent under complex, real-world style interactions.

N	Model	Single	BoN	Debate	DyLAN	GPTSwarm	MacNet	M2CL
N=4	3B	10.9 _{±6.1}	17.0 _{↑0.0}	12.5 _{±4.5}	13.0 _{±4.0}	15.0 _{±2.0}	16.8 _{±0.2}	25.5 _{↑8.5}
	7B	27.0 _{±11.0}	38.0 _{↑0.0}	30.5 _{±7.5}	31.0 _{±7.0}	37.0 _{±1.0}	39.6 _{↑1.6}	44.9 _{↑6.9}
N=8	3B	10.9 _{±6.1}	17.0 _{↑0.0}	15.7 _{±1.3}	15.2 _{±1.8}	15.6 _{±1.4}	19.7 _{↑2.7}	32.0 _{↑15.0}
	7B	27.0 _{±11.0}	38.0 _{↑0.0}	33.0 _{±5.0}	34.5 _{±3.5}	38.0 _{↑0.0}	44.0 _{↑6.0}	50.0 _{↑12.0}
N=16	3B	10.9 _{±6.1}	17.0 _{↑0.0}	16.9 _{±0.1}	18.4 _{↑1.4}	18.1 _{↑1.1}	22.6 _{↑5.6}	38.0 _{↑21.0}
	7B	27.0 _{±11.0}	38.0 _{↑0.0}	37.3 _{±0.7}	37.8 _{±0.2}	42.5 _{↑4.5}	47.9 _{↑9.9}	55.0 _{↑17.0}
N=32	3B	10.9 _{±6.1}	17.0 _{↑0.0}	18.0 _{↑1.0}	21.5 _{↑4.5}	20.6 _{↑3.6}	25.2 _{↑8.2}	42.0 _{↑25.0}
	7B	27.0 _{±11.0}	38.0 _{↑0.0}	41.5 _{↑3.5}	40.0 _{↑2.0}	45.0 _{↑7.0}	52.8 _{↑14.8}	59.0 _{↑21.0}
N=64	3B	10.9 _{±6.1}	17.0 _{↑0.0}	20.1 _{↑3.1}	23.6 _{↑6.6}	24.2 _{↑7.2}	28.6 _{↑11.6}	45.0 _{↑28.0}
	7B	27.0 _{±11.0}	38.0 _{↑0.0}	41.7 _{↑3.7}	42.2 _{↑4.2}	49.3 _{↑11.3}	55.2 _{↑17.2}	62.0 _{↑24.0}

Table 9: Accuracy with varying number of LLMs from 4 to 64 on AndroidWorld. We exhibit the performance advantage with BoN and highlight the best result.

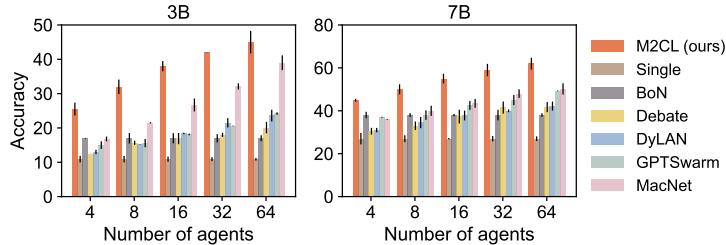


Figure 11: Performance of Qwen2.5-VL (3B and 7B) as the base model with varying number of LLMs. M2CL exhibits higher performance and increasing tendency with more LLMs, demonstrating its great collaboration efficiency compared to existing methods. Uncertainty intervals depict standard deviation over three seeds.

G.4 CONTEXT CONSTRAINT

To assess the effect of context constraint, we carry out experiments by varying the context constraint β from 0 to 10. As illustrated in Figs. 12 to 16, a larger value of β results in a high degree of consistency among LLMs, leading them to produce similar answers. Conversely, a smaller value of β is associated with reduced collaboration among LLMs. Therefore, it is important to adjust β to control the degree of consistency among LLMs for better collaboration.

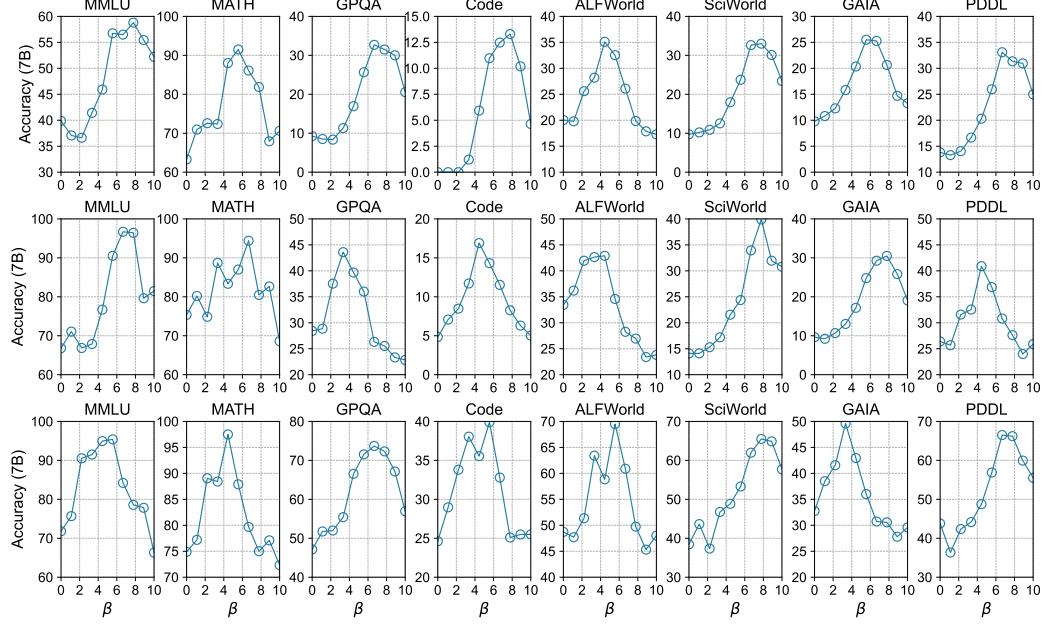


Figure 12: Performance with varying context constraint when 4 LLMs participate. All the curves display the same trend of rising first and then falling, which is consistent with our theory.

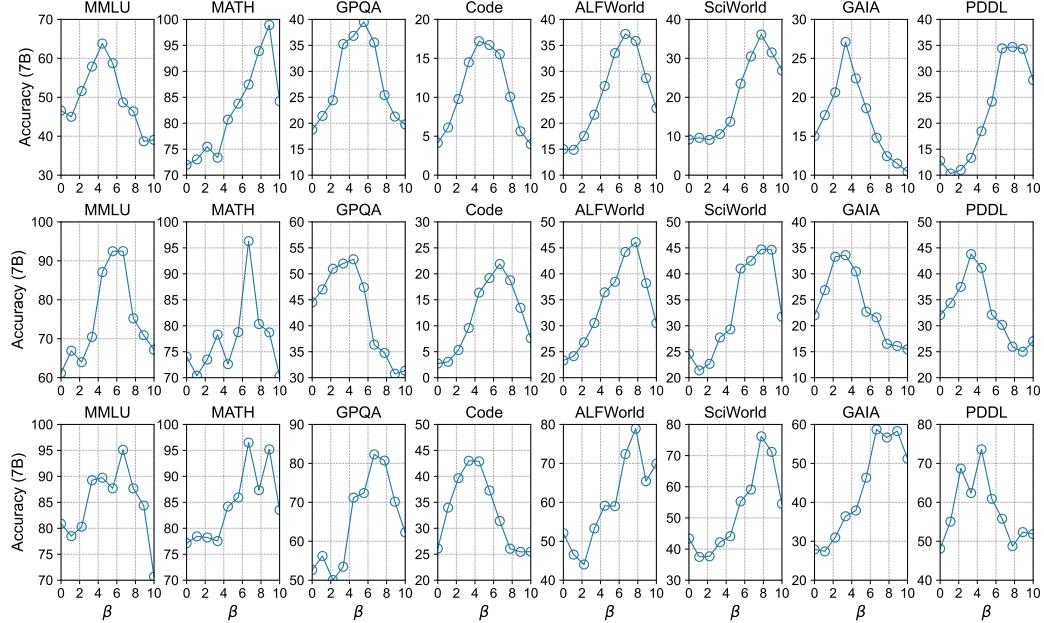


Figure 13: Performance with varying context constraint when 8 LLMs participate.

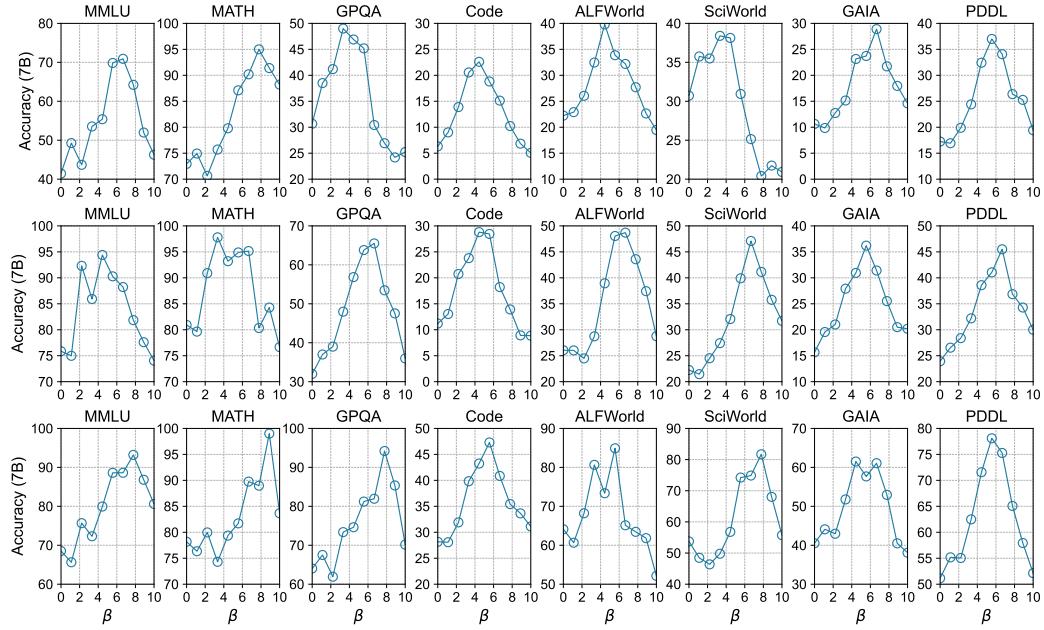


Figure 14: Performance with varying context constraint when 16 LLMs participate.

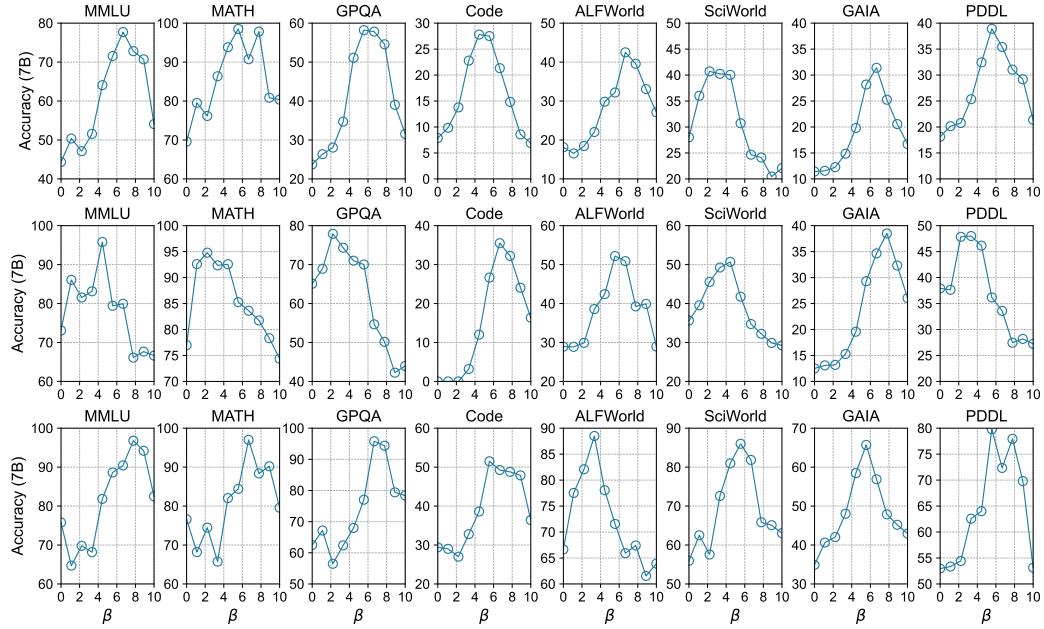


Figure 15: Performance with varying context constraint when 32 LLMs participate.

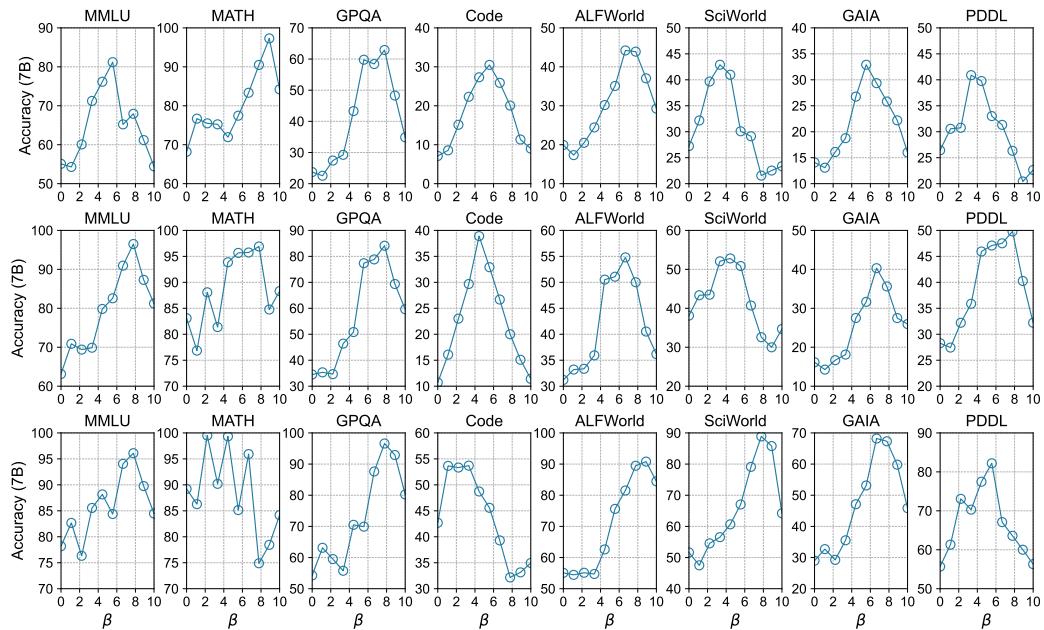


Figure 16: Performance with varying context constraint when 64 LLMs participate.

G.5 DISCREPANCY INTENSITY BY ROUNDS

To validate that M2CL can collaborate LLMs to reach an agreement by rounds, we visualize the discrepancy intensity. As illustrated in Figs. 17 to 21, the initial discrepancy intensity of M2CL is higher as its context initialization can make the discussion more creative. The discrepancy intensity of M2CL increases faster because the dynamic adjustment of the context provides LLMs with better ability to effectively receive information from other LLMs, resulting in reduced disagreement and faster collaboration among LLMs to reach a consensus.

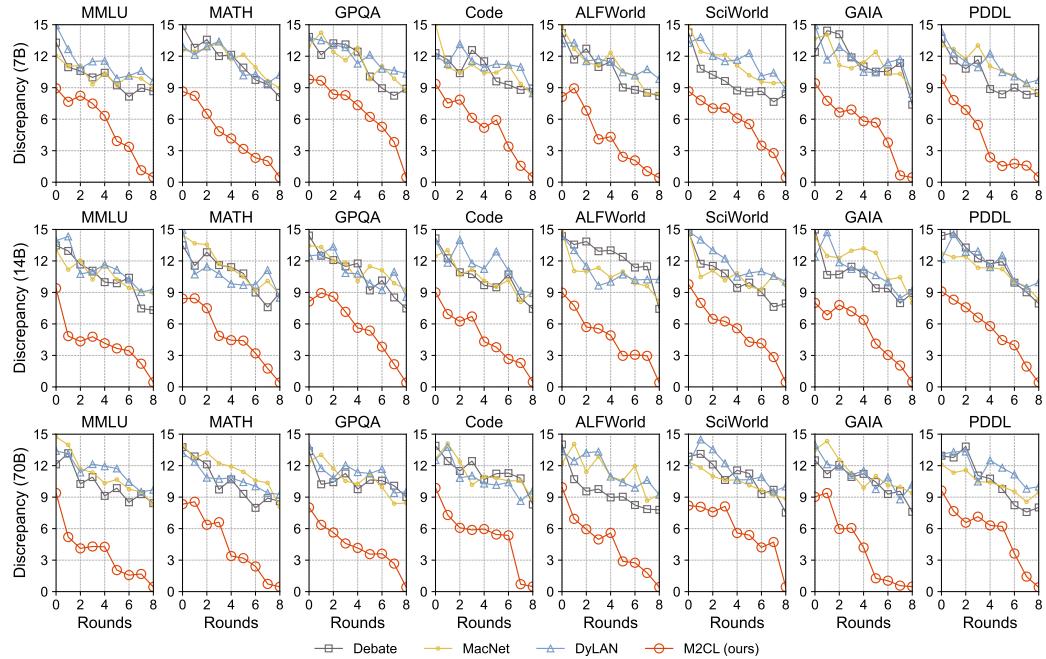


Figure 17: Comparative results on discrepancy intensity with varying model size (from top to bottom correspond to 7B, 14B, and 70B). The number of agents is set as 4. Lower values represent a lower degree of disagreement. M2CL can improve consistency with fewer rounds. Of note, M2CL displays both a lower initial value and a faster decreasing speed, indicating its capability of assigning appropriate contexts based on the given question and current discussion situation.

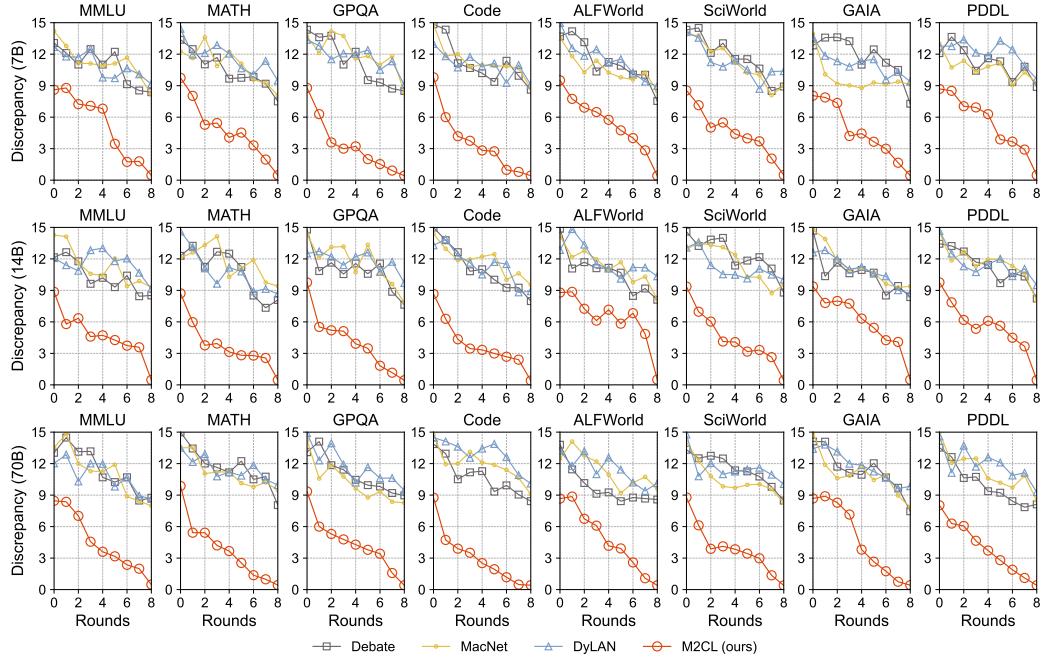


Figure 18: Comparative results on discrepancy intensity with varying model size (from top to bottom correspond to 7B, 14B, and 70B). The number of agents is set as 8. Lower values represent a lower degree of disagreement. M2CL can improve consistency with fewer rounds.

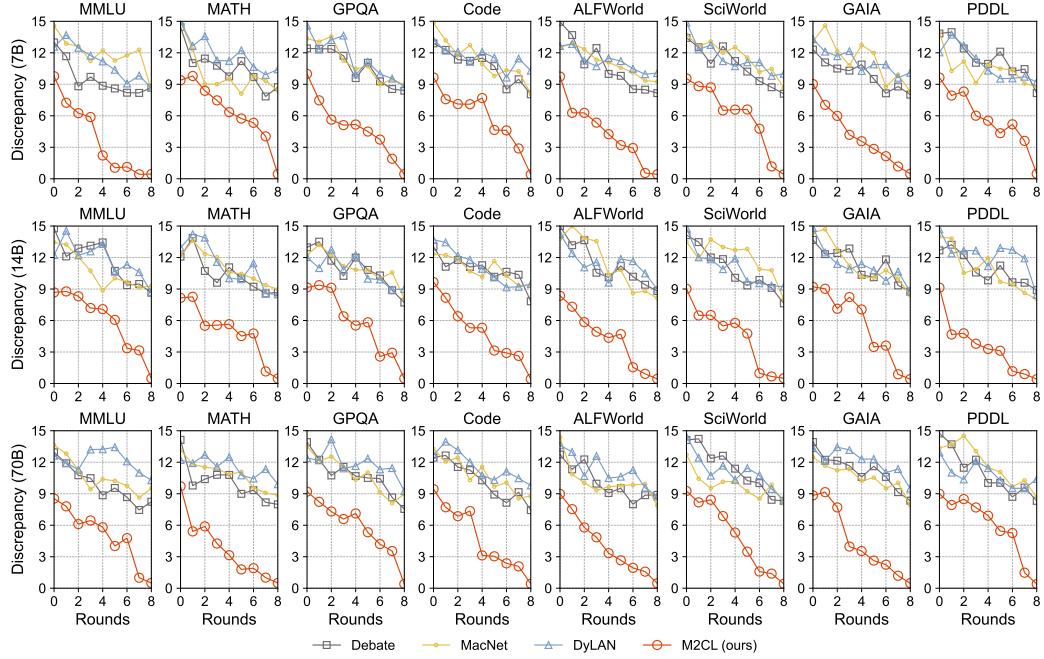


Figure 19: Comparative results on discrepancy intensity with varying model size (from top to bottom correspond to 7B, 14B, and 70B). The number of agents is set as 16. Lower values represent a lower degree of disagreement. M2CL can improve consistency with fewer rounds.

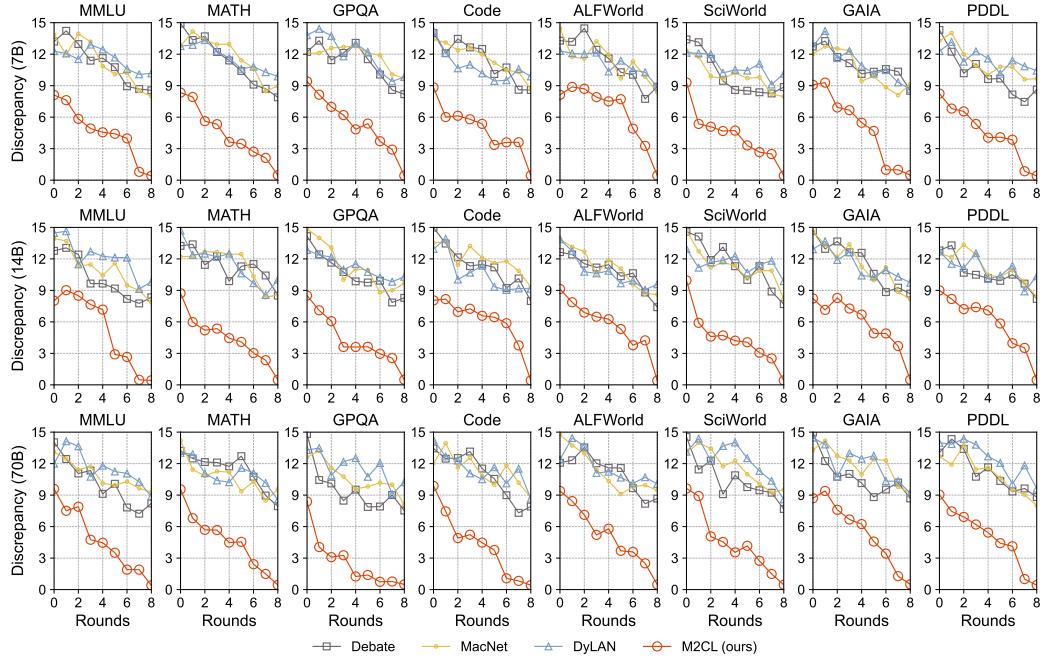


Figure 20: Comparative results on discrepancy intensity with varying model size (from top to bottom correspond to 7B, 14B, and 70B) The number of agents is set as 32. Lower values represent a lower degree of disagreement. M2CL can improve consistency with fewer rounds.

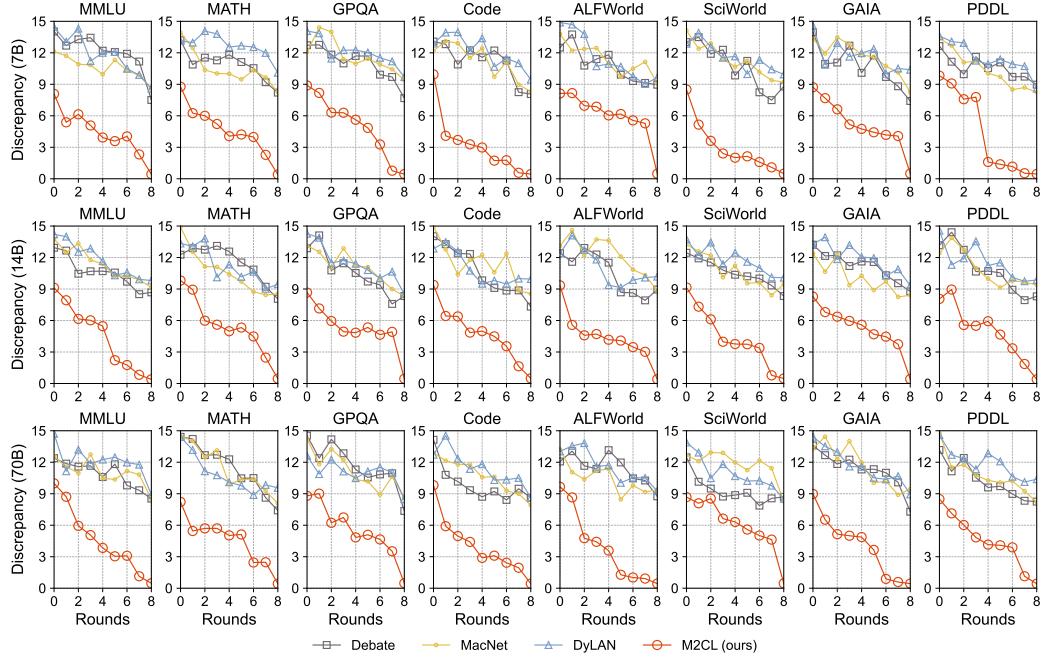


Figure 21: Comparative results on discrepancy intensity with varying model size (from top to bottom correspond to 7B, 14B, and 70B). The number of agents is set as 64. Lower values represent a lower degree of disagreement. M2CL can improve consistency with fewer rounds.

G.6 TRANSFERABILITY OF CONTEXTS

To further study the generalization of the generated contexts, we implement the multi-agent system using GPT-4 as the base model with the context generator trained on llama-7B and compare its performance with using initial contexts. As illustrated in Table 10, the generated contexts outperform initial contexts, indicating that the trained context generator can be expanded to more models for improving overall performance through LLMs’ collaboration.

Dataset	Method	MMLU	MATH	GPQA	Code	ALFWorld	SciWorld	GAIA	PDDL
$n = 4$	Initial	84.3 _{±0.0}	41.1 _{±0.0}	33.3 _{±0.0}	62.7 _{±0.0}	70.2 _{±0.0}	65.4 _{±0.0}	58.7 _{±0.0}	73.5 _{±0.0}
	Generated	95.0 _{±10.7}	58.0 _{±16.9}	52.1 _{±18.8}	81.0 _{±18.3}	82.0 _{±11.8}	76.5 _{±11.1}	65.4 _{±6.7}	85.2 _{±11.7}
$n = 8$	Initial	86.4 _{±0.0}	43.9 _{±0.0}	38.8 _{±0.0}	67.0 _{±0.0}	72.8 _{±0.0}	68.0 _{±0.0}	61.0 _{±0.0}	76.0 _{±0.0}
	Generated	98.2 _{±11.8}	68.5 _{±24.6}	62.5 _{±23.7}	89.1 _{±22.1}	82.5 _{±9.7}	77.8 _{±9.8}	64.3 _{±3.3}	88.2 _{±12.2}
$n = 16$	Initial	87.9 _{±0.0}	44.8 _{±0.0}	40.1 _{±0.0}	68.5 _{±0.0}	74.3 _{±0.0}	70.5 _{±0.0}	63.0 _{±0.0}	78.2 _{±0.0}
	Generated	98.5 _{±10.6}	70.3 _{±25.5}	64.3 _{±24.2}	90.3 _{±21.8}	84.5 _{±10.2}	78.7 _{±8.2}	66.1 _{±3.1}	90.0 _{±11.8}
$n = 32$	Initial	89.1 _{±0.0}	45.4 _{±0.0}	41.5 _{±0.0}	69.7 _{±0.0}	75.0 _{±0.0}	71.0 _{±0.0}	63.5 _{±0.0}	78.5 _{±0.0}
	Generated	98.8 _{±9.7}	72.0 _{±26.6}	66.0 _{±24.5}	91.3 _{±21.6}	86.0 _{±11.0}	79.5 _{±8.5}	66.8 _{±3.3}	91.2 _{±12.7}
$n = 64$	Initial	88.5 _{±0.0}	45.4 _{±0.0}	41.8 _{±0.0}	70.5 _{±0.0}	75.2 _{±0.0}	71.5 _{±0.0}	64.0 _{±0.0}	79.2 _{±0.0}
	Generated	98.9 _{±10.4}	72.8 _{±27.4}	66.9 _{±25.1}	91.7 _{±21.2}	86.5 _{±11.3}	80.2 _{±8.7}	67.2 _{±3.2}	91.5 _{±12.3}

Table 10: Performance of transferring generated contexts trained on llama-7B to GPT-4 with varying number of LLMs.

G.7 ABLATION STUDIES AND COMPLEMENTARY EXPERIMENTS

G.7.1 TRAINING CONVERGENCE OF CONTEXT GENERATOR

We verify the convergence of solving Problem (16) by displaying the cumulative utility. As shown in Fig. 22, it works well in all dataset and often converges in 60 training steps.

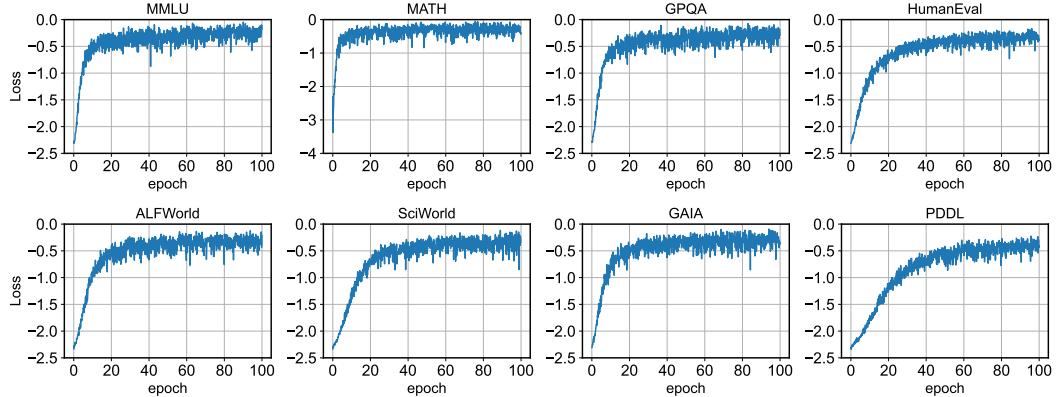


Figure 22: The value of $L(\theta)$ when training llama-7B with 8 LLMs participating.

G.7.2 ABLATION

In this section, we assess the effect of key components by ablating them under the same setting.

Without context initialization, LLMs fail to develop specialized expertise, resulting in homogeneous policies that inefficiently duplicate effort, demonstrating poor adaptability when faced with novel questions. Tables 11 to 13 underscore the importance of context initialization before discussion to enhance foundational multi-agent capabilities.

Without tuning α during discussion rounds, LLMs tend to reach an agreement in the first round, leading to responses that lack creativity and diversity, which ultimately reduces problem-solving ability. Tables 11 to 13 underscore the importance of tuning α during discussion rounds for training.

Without context evolution, LLMs receive no guidance consider inter-LLM dependencies, making it difficult for them to fully utilize the information contributed by other LLMs. This context misalignment leads to discussion inconsistency and poor collaboration efficiency. Tables 11 to 13 highlight the necessity of context evolution for enabling effective inter-LLM collaboration and achieving consistent improvements across discussion rounds.

Dataset	Method	MMLU	MATH	GPQA	Code	ALFWorld	SciWorld	GAIA	PDDL
$n = 4$	M2CL	59.1 \uparrow 10.6	17.4 \uparrow 9.9	44.7 \uparrow 16.7	32.9 \uparrow 11.8	35.1 \uparrow 8.8	33.0 \uparrow 8.3	25.5 \uparrow 8.1	33.1 \uparrow 6.2
	w/o init.	56.4 \uparrow 7.9	13.3 \uparrow 5.8	44.5 \uparrow 16.5	31.3 \uparrow 10.2	31.8 \uparrow 5.5	30.1 \uparrow 5.4	23.2 \uparrow 5.8	30.5 \uparrow 6.6
	w/o α	53.9 \uparrow 5.4	9.4 \uparrow 1.9	37.8 \uparrow 9.8	26.0 \uparrow 4.9	22.5 \uparrow 3.8	21.1 \uparrow 3.6	20.3 \uparrow 2.9	23.4 \uparrow 3.5
	w/o evolve	48.5 \downarrow 0.0	7.5 \downarrow 0.0	28.0 \downarrow 0.0	21.1 \downarrow 0.0	26.3 \downarrow 0.0	24.7 \downarrow 0.0	17.4 \downarrow 0.0	26.9 \downarrow 0.0
$n = 8$	M2CL	63.8 \uparrow 12.9	24.9 \uparrow 16.8	59.3 \uparrow 31.1	38.8 \uparrow 17.6	37.2 \uparrow 10.2	36.1 \uparrow 10.3	27.1 \uparrow 10.1	34.7 \uparrow 7.0
	w/o init.	63.8 \uparrow 12.9	17.2 \uparrow 9.1	54.9 \uparrow 26.7	37.2 \uparrow 16.0	34.7 \uparrow 7.7	33.1 \uparrow 7.3	24.5 \uparrow 7.5	31.5 \uparrow 3.8
	w/o α	57.4 \uparrow 6.5	12.9 \uparrow 4.8	43.0 \uparrow 14.8	29.8 \uparrow 8.6	22.7 \uparrow 4.3	26.8 \uparrow 1.0	20.9 \uparrow 3.9	24.8 \downarrow 2.9
	w/o evolve	50.9 \downarrow 0.0	8.1 \downarrow 0.0	28.2 \downarrow 0.0	21.2 \downarrow 0.0	27.0 \downarrow 0.0	25.8 \downarrow 0.0	17.0 \downarrow 0.0	27.7 \downarrow 0.0
$n = 16$	M2CL	71.5 \uparrow 18.5	23.9 \uparrow 15.7	69.9 \uparrow 39.4	48.3 \uparrow 24.3	39.8 \uparrow 12.7	38.4 \uparrow 11.8	28.9 \uparrow 11.1	37.0 \uparrow 9.5
	w/o init.	70.9 \uparrow 17.9	22.6 \uparrow 14.4	69.4 \uparrow 38.9	45.3 \uparrow 21.3	36.3 \uparrow 9.2	35.3 \uparrow 8.7	26.2 \uparrow 8.4	33.9 \uparrow 6.4
	w/o α	60.5 \uparrow 7.5	14.2 \uparrow 6.0	46.4 \uparrow 15.9	32.8 \uparrow 8.8	26.6 \uparrow 0.5	23.7 \uparrow 2.9	22.2 \uparrow 4.4	24.9 \downarrow 2.6
	w/o evolve	53.0 \downarrow 0.0	8.2 \downarrow 0.0	30.5 \downarrow 0.0	24.0 \downarrow 0.0	27.1 \downarrow 0.0	26.6 \downarrow 0.0	17.8 \downarrow 0.0	27.5 \downarrow 0.0
$n = 32$	M2CL	95.8 \uparrow 41.3	41.4 \uparrow 31.6	94.7 \uparrow 59.7	75.6 \uparrow 48.8	42.5 \uparrow 14.5	40.7 \uparrow 13.6	31.4 \uparrow 14.0	38.9 \uparrow 11.4
	w/o init.	77.7 \uparrow 23.2	27.8 \uparrow 18.0	83.5 \uparrow 48.5	53.2 \uparrow 26.4	40.4 \uparrow 12.4	36.7 \uparrow 9.6	28.8 \uparrow 11.4	35.3 \uparrow 7.8
	w/o α	64.4 \uparrow 9.9	18.1 \uparrow 8.3	50.0 \uparrow 15.0	34.8 \uparrow 8.0	32.5 \uparrow 4.5	28.6 \uparrow 1.5	24.6 \uparrow 7.2	28.4 \uparrow 0.9
	w/o evolve	54.5 \downarrow 0.0	9.8 \downarrow 0.0	35.0 \downarrow 0.0	26.8 \downarrow 0.0	28.0 \downarrow 0.0	27.1 \downarrow 0.0	17.4 \downarrow 0.0	27.5 \downarrow 0.0
$n = 64$	M2CL	81.5 \uparrow 25.3	35.4 \uparrow 23.6	82.5 \uparrow 44.1	60.6 \uparrow 32.0	44.2 \uparrow 15.6	42.9 \uparrow 14.9	32.9 \uparrow 14.8	40.9 \uparrow 12.6
	w/o init.	81.2 \uparrow 25.0	30.5 \uparrow 18.7	80.6 \uparrow 42.2	57.2 \uparrow 28.6	40.7 \uparrow 12.1	39.5 \uparrow 11.5	30.4 \uparrow 12.3	37.4 \uparrow 9.1
	w/o α	63.3 \uparrow 7.1	19.2 \uparrow 7.4	49.7 \uparrow 11.3	37.0 \uparrow 8.4	31.7 \uparrow 3.1	32.9 \uparrow 4.9	24.6 \uparrow 6.5	27.8 \downarrow 0.5
	w/o evolve	56.2 \downarrow 0.0	11.8 \downarrow 0.0	38.4 \downarrow 0.0	28.6 \downarrow 0.0	28.6 \downarrow 0.0	28.0 \downarrow 0.0	18.1 \downarrow 0.0	28.3 \downarrow 0.0

Table 11: Ablation study on context initialization, tuning α , and context evolution when using llama-7B with varying number of LLMs.

Dataset	Method	MMLU	MATH	GPQA	Code	ALFWorld	SciWorld	GAIA	PDDL
$n = 4$	M2CL (ours)	86.9 \uparrow 20.2	23.4 \uparrow 13.5	58.9 \uparrow 20.3	47.5 \uparrow 17.6	42.9 \uparrow 10.9	39.9 \uparrow 9.5	30.5 \uparrow 9.6	40.9 \uparrow 8.1
	w/o init.	86.7 \uparrow 20.0	16.9 \uparrow 7.0	57.9 \uparrow 19.3	45.2 \uparrow 15.3	40.7 \uparrow 8.7	47.8 \uparrow 17.4	33.7 \uparrow 12.8	37.0 \uparrow 4.2
	w/o α	81.5 \uparrow 14.8	13.2 \uparrow 3.3	50.9 \uparrow 12.3	37.3 \uparrow 7.4	32.1 \uparrow 1.1	31.9 \uparrow 1.5	24.2 \uparrow 3.3	24.3 \downarrow 8.5
	w/o evolve	66.7 \downarrow 0.0	9.9 \downarrow 0.0	38.6 \downarrow 0.0	29.9 \downarrow 0.0	32.0 \downarrow 0.0	30.4 \downarrow 0.0	20.9 \downarrow 0.0	32.8 \downarrow 0.0
$n = 8$	M2CL (ours)	92.7 \uparrow 23.2	23.0 \uparrow 13.5	70.7 \uparrow 32.6	55.4 \uparrow 22.1	46.1 \uparrow 13.0	44.7 \uparrow 13.0	33.6 \uparrow 12.2	43.8 \uparrow 10.6
	w/o init.	92.5 \uparrow 23.0	21.9 \uparrow 12.4	70.3 \uparrow 32.2	53.3 \uparrow 20.0	45.3 \uparrow 12.2	51.8 \uparrow 20.1	38.4 \uparrow 17.0	39.9 \uparrow 6.7
	w/o α	85.4 \uparrow 15.9	16.9 \uparrow 7.4	53.9 \uparrow 15.8	43.3 \uparrow 10.0	34.7 \uparrow 1.6	37.6 \uparrow 5.9	26.2 \uparrow 4.8	29.0 \downarrow 4.2
	w/o evolve	69.5 \downarrow 0.0	9.5 \downarrow 0.0	38.1 \downarrow 0.0	33.3 \downarrow 0.0	33.1 \downarrow 0.0	31.7 \downarrow 0.0	21.4 \downarrow 0.0	33.2 \downarrow 0.0
$n = 16$	M2CL (ours)	94.9 \uparrow 22.9	27.3 \uparrow 16.2	73.2 \uparrow 33.8	61.5 \uparrow 25.3	49.7 \uparrow 15.3	48.9 \uparrow 15.6	36.0 \uparrow 12.9	45.2 \uparrow 10.2
	w/o init.	94.7 \uparrow 22.7	25.0 \uparrow 13.9	72.3 \uparrow 32.9	60.4 \uparrow 24.2	48.5 \uparrow 14.1	54.5 \uparrow 21.2	39.1 \uparrow 16.0	41.8 \uparrow 6.8
	w/o α	87.2 \uparrow 15.2	18.7 \uparrow 7.6	57.7 \uparrow 18.3	47.6 \uparrow 11.4	36.3 \uparrow 1.9	39.9 \uparrow 6.6	27.0 \uparrow 3.9	31.0 \downarrow 4.0
	w/o evolve	72.0 \downarrow 0.0	11.1 \downarrow 0.0	39.4 \downarrow 0.0	36.2 \downarrow 0.0	34.4 \downarrow 0.0	33.3 \downarrow 0.0	23.1 \downarrow 0.0	35.0 \downarrow 0.0
$n = 32$	M2CL (ours)	95.5 \uparrow 21.8	27.9 \uparrow 15.3	74.0 \uparrow 32.3	63.1 \uparrow 24.6	50.8 \uparrow 15.3	49.9 \uparrow 14.7	36.6 \uparrow 12.3	46.7 \uparrow 10.5
	w/o init.	95.3 \uparrow 21.6	25.7 \uparrow 13.1	73.0 \uparrow 31.3	62.0 \uparrow 23.5	49.4 \uparrow 14.1	55.2 \uparrow 20.0	39.8 \uparrow 15.5	43.0 \uparrow 6.8
	w/o α	87.6 \uparrow 13.9	19.1 \uparrow 6.5	58.5 \uparrow 16.8	48.1 \uparrow 9.6	36.7 \uparrow 1.4	40.5 \uparrow 3.3	27.4 \uparrow 3.1	31.7 \downarrow 4.5
	w/o evolve	73.7 \downarrow 0.0	12.6 \downarrow 0.0	41.7 \downarrow 0.0	38.5 \downarrow 0.0	35.3 \downarrow 0.0	35.2 \downarrow 0.0	24.3 \downarrow 0.0	36.2 \downarrow 0.0
$n = 64$	M2CL (ours)	95.8 \uparrow 21.6	28.1 \uparrow 14.9	74.8 \uparrow 31.3	63.5 \uparrow 23.5	51.2 \uparrow 15.2	50.3 \uparrow 13.9	36.9 \uparrow 11.9	47.1 \uparrow 10.3
	w/o init.	95.6 \uparrow 21.4	25.9 \uparrow 12.7	73.5 \uparrow 30.0	62.2 \uparrow 22.2	49.6 \uparrow 13.6	55.6 \uparrow 19.2	40.1 \uparrow 15.1	43.3 \uparrow 6.5
	w/o α	87.8 \uparrow 13.6	19.4 \uparrow 6.2	59.1 \uparrow 15.6	48.6 \uparrow 8.6	37.0 \uparrow 1.0	40.8 \uparrow 4.4	27.7 \uparrow 2.7	32.0 \downarrow 4.8
	w/o evolve	74.2 \downarrow 0.0	13.2 \downarrow 0.0	43.5 \downarrow 0.0	40.0 \downarrow 0.0	36.0 \downarrow 0.0	36.4 \downarrow 0.0	25.0 \downarrow 0.0	36.8 \downarrow 0.0

Table 12: Ablation study on context initialization, tuning α , and context evolution when using llama-13B with varying number of LLMs.

Dataset	Method	MMLU	MATH	GPQA	Code	ALFWorld	SciWorld	GAIA	PDDL
$n = 4$	M2CL (ours)	95.6 ^{↑18.7}	40.3 ^{↑9.4}	78.7 ^{↑25.6}	70.6 ^{↑24.7}	69.5 ^{↑17.2}	65.5 ^{↑15.9}	49.6 ^{↑15.5}	66.5 ^{↑12.2}
	w/o init.	95.4 ^{↑18.5}	39.9 ^{↑9.0}	88.5 ^{↑35.4}	67.8 ^{↑21.9}	64.8 ^{↑12.5}	62.9 ^{↑13.3}	47.0 ^{↑12.9}	62.6 ^{↑8.3}
	w/o α	86.7 ^{↑9.8}	36.9 ^{↑6.0}	72.3 ^{↑19.2}	53.9 ^{↑8.0}	43.3 ^{↑9.0}	47.1 ^{↑2.5}	38.0 ^{↓0.1}	40.3 ^{↓14.0}
	w/o evolve	76.9 ^{↓0.0}	30.9 ^{↓0.0}	53.1 ^{↓0.0}	45.9 ^{↓0.0}	52.3 ^{↓0.0}	49.6 ^{↓0.0}	34.1 ^{↓0.0}	54.3 ^{↓0.0}
$n = 8$	M2CL (ours)	95.4 ^{↑11.9}	43.5 ^{↑8.3}	87.6 ^{↑26.6}	81.0 ^{↑30.5}	78.9 ^{↑24.5}	76.2 ^{↑23.5}	58.7 ^{↑23.8}	73.6 ^{↑18.3}
	w/o init.	95.1 ^{↑11.6}	43.0 ^{↑7.8}	86.3 ^{↑25.3}	80.0 ^{↑29.5}	75.8 ^{↑21.4}	71.7 ^{↑18.5}	56.1 ^{↑21.2}	71.3 ^{↑16.0}
	w/o α	94.0 ^{↑10.5}	39.3 ^{↑4.1}	78.0 ^{↑17.0}	63.8 ^{↑13.3}	49.3 ^{↓3.1}	49.0 ^{↓3.7}	42.6 ^{↑7.7}	47.8 ^{↓7.5}
	w/o evolve	83.5 ^{↓0.0}	35.2 ^{↓0.0}	61.0 ^{↓0.0}	50.5 ^{↓0.0}	54.4 ^{↓0.0}	52.7 ^{↓0.0}	34.9 ^{↓0.0}	55.3 ^{↓0.0}
$n = 16$	M2CL (ours)	93.9 ^{↑9.9}	51.5 ^{↑15.2}	91.3 ^{↑29.5}	97.2 ^{↑45.4}	84.9 ^{↑29.4}	81.7 ^{↑27.4}	61.5 ^{↑25.6}	78.1 ^{↑22.8}
	w/o init.	93.2 ^{↑9.2}	47.3 ^{↑11.0}	90.9 ^{↑29.1}	96.9 ^{↑45.1}	81.3 ^{↑25.8}	79.3 ^{↑23.0}	56.6 ^{↑20.7}	73.5 ^{↑18.2}
	w/o α	93.2 ^{↑9.2}	40.5 ^{↑4.2}	78.6 ^{↑16.8}	72.2 ^{↑20.4}	64.2 ^{↓8.7}	64.8 ^{↑10.5}	47.0 ^{↑11.1}	59.3 ^{↑4.0}
	w/o evolve	84.0 ^{↓0.0}	36.3 ^{↓0.0}	61.8 ^{↓0.0}	51.8 ^{↓0.0}	55.5 ^{↓0.0}	54.3 ^{↓0.0}	35.9 ^{↓0.0}	55.3 ^{↓0.0}
$n = 32$	M2CL (ours)	97.0 ^{↑10.2}	54.5 ^{↑17.3}	95.1 ^{↑28.5}	93.7 ^{↑38.2}	88.5 ^{↑32.3}	86.0 ^{↑30.5}	65.7 ^{↑29.5}	79.8 ^{↑23.5}
	w/o init.	96.8 ^{↑10.0}	51.5 ^{↑14.3}	94.8 ^{↑27.9}	93.4 ^{↑37.9}	85.8 ^{↑29.6}	81.8 ^{↑26.3}	61.2 ^{↑24.3}	77.1 ^{↑20.8}
	w/o α	95.2 ^{↑8.4}	44.0 ^{↑6.8}	83.9 ^{↑17.3}	68.6 ^{↑13.1}	57.3 ^{↓1.1}	57.9 ^{↑2.4}	50.9 ^{↑14.7}	61.3 ^{↑5.0}
	w/o evolve	86.8 ^{↓0.0}	37.2 ^{↓0.0}	66.6 ^{↓0.0}	55.5 ^{↓0.0}	56.2 ^{↓0.0}	55.5 ^{↓0.0}	36.2 ^{↓0.0}	56.3 ^{↓0.0}
$n = 64$	M2CL (ours)	96.3 ^{↑5.7}	58.0 ^{↑19.3}	96.9 ^{↑25.1}	92.0 ^{↑34.2}	90.8 ^{↑33.5}	88.9 ^{↑31.9}	68.2 ^{↑32.0}	82.2 ^{↑26.1}
	w/o init.	96.1 ^{↑5.5}	53.7 ^{↑15.0}	96.9 ^{↑25.1}	91.7 ^{↑33.9}	88.7 ^{↑31.4}	85.5 ^{↑28.5}	65.0 ^{↑28.8}	77.5 ^{↑21.4}
	w/o α	95.7 ^{↑5.1}	43.3 ^{↑4.6}	82.8 ^{↑11.0}	69.6 ^{↑11.8}	66.9 ^{↓9.6}	70.3 ^{↑13.3}	41.8 ^{↑5.6}	55.6 ^{↑0.5}
	w/o evolve	90.6 ^{↓0.0}	38.7 ^{↓0.0}	71.8 ^{↓0.0}	57.8 ^{↓0.0}	57.3 ^{↓0.0}	57.0 ^{↓0.0}	36.2 ^{↓0.0}	56.1 ^{↓0.0}

Table 13: Ablation study on context initialization, tuning α , and context evolution when using llama-70B with varying number of LLMs.

G.7.3 RUN-TIME

To demonstrate the efficiency of our context initialization, we verify its runtime with varying the number of LLMs. Then, we evaluate the runtime of M2CL compared with baseline algorithms

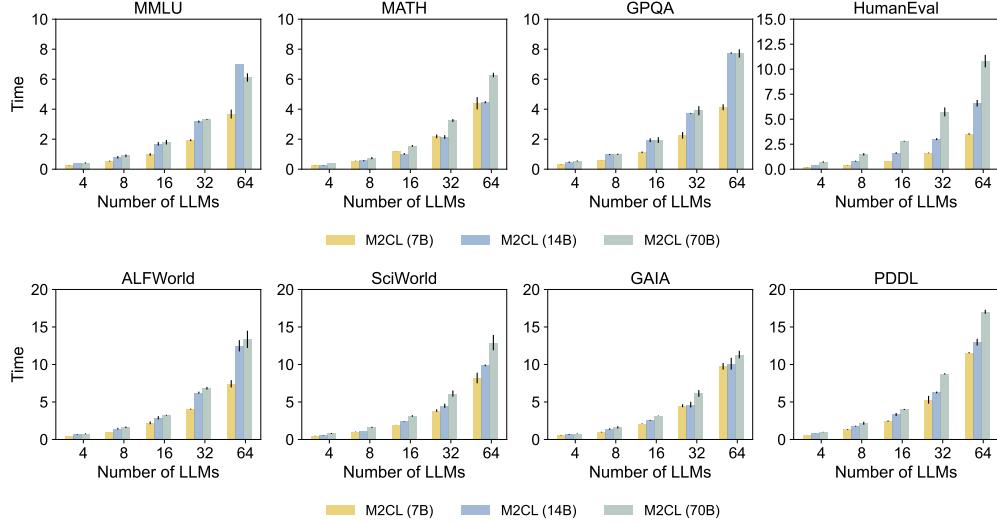


Figure 23: Runtime of initialization. Uncertainty intervals depict standard deviation over three seeds.

for average testing time, utilizing the same model size on 8 NVIDIA H800 GPUs. As illustrated by Fig. 24, the runtime of M2CL is slightly longer than other multi-LLM discussion methods as the runtime of context generators is negligible compared with the inference time of LLM, which substantiates the low computational cost of M2CL.

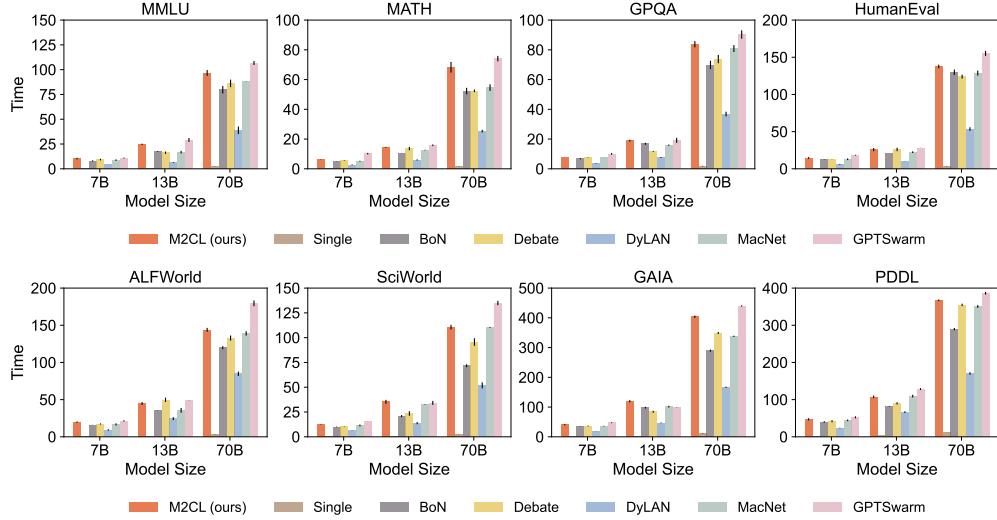


Figure 24: Runtime when varying the size of the LLama series models. The number of LLMs is 8. Uncertainty intervals depict standard deviation over three seeds.

H CASE STUDY

We used a problem from the MATH Hendrycks et al. (2021). The number of LLMs is set as 8. For each LLM, we present their instructions, responses, and final answers for 4 discussion rounds.

H.1 CASE STUDY OF M2CL (OURS)

We provide the case of M2CL solving the problem. As illustrated in Figs. 25 to 28, we observe that the generated contexts evolve progressively to enforce stricter collaboration among LLMs. Initially, the instructions merely encourage LLMs to pay attention to others' responses. In subsequent rounds, they guide LLMs to cross-check each other's answers and eventually require reaching full agreement. Correspondingly, the LLMs' outputs transition from diverse, potentially conflicting answers to a single, consistent solution. This behavior demonstrates that our method effectively modulates context evolution to guide toward consensus without sacrificing initial creativity.

Problem: In the diagram, square ABCD has sides of length 4, and $\triangle ABE$ is equilateral. Line segments BE and AC intersect at P. Point Q is on BC so that PQ is perpendicular to BC and $PQ=x$. A=(0,0); B=(4,0); C=(4,4); D=(0,4); E=(2,3.464); P=(2.535,2.535); Q=(4,2.535); Determine the measure of $\angle BPC$.

Solution: Since $\triangle ABE$ is equilateral, we know that $\angle ABE = 60^\circ$. Therefore, $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. Since AB=BC, we know that $\triangle ABC$ is a right isosceles triangle and $\angle BAC = \angle BCA = 45^\circ$. Then, $\angle BCP = \angle BCA = 45^\circ$ and $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.

Agent 1

Instruction: You are an expert in Geometric Constructions. Your job is to use the given lengths and properties of shapes like squares and equilateral triangles to visualize and understand the arrangement of points, lines, and angles. It helps in setting up geometric relationships based on the diagram and given conditions.

Response: Start with $\triangle ABE$ as equilateral, making $\angle ABE = 60^\circ$. In square ABCD, $\angle ABC = 90^\circ$, so $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. Since AB = BC, $\triangle ABC$ is a right isosceles triangle, meaning $\angle BAC = \angle BCA = 45^\circ$. Therefore, $\angle BCP = 45^\circ$. Calculate $\angle BPC = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.

Final Answer: 105°

Agent 2

Instruction: You are an expert in Analytic Geometry. Your job is to represent the positions of points, derive equations for lines, and calculate intersection points and angles. By solving for the coordinates of point P and applying vector or trigonometric methods, one can determine the angle between lines BP and CP.

Response: Assign coordinates $A = (0,0)$, $B = (4,0)$, $C = (4,4)$, $E = (2, 2\sqrt{3})$. The line BE has slope $m_{BE} = \sqrt{3}$, and equation $y = \sqrt{3}(x - 4)$. The line AC has slope $m_{AC} = 1$, and equation $y = x + 4$. Solve for intersection P: $\sqrt{3}(x - 4) = x + 4$. Solve for $x = 4$. Substitute $x = 4$ to find $P = (4, 0)$. Now calculate the slopes of BP and CP, use $\tan \theta$ for $\angle BPC$. After some calculations, conclude that $\angle BPC = 120^\circ$.

Final Answer: 120°

Agent 3

Instruction: You are an expert in Vector Mathematics. Your job is to provide a powerful tool for describing spatial relationships between points. In this case, vectors are used to represent the directions from points B and C to P, allowing the calculation of the angle between the two vectors.

Response: Represent vectors $\vec{BP} = (1.465, 2.535)$ and $\vec{CP} = (1.465, 1.465)$. Use the dot product formula: $\vec{BP} \cdot \vec{CP} = 1.465 \cdot 1.465 + 2.535 \cdot 1.465$. After calculating the dot product and magnitudes, use $\cos \theta = \frac{\text{dot product}}{\text{magnitudes}}$. Solve for θ and conclude $\angle BPC = 90^\circ$.

Final Answer: 90°

Agent 4

Instruction: You are an expert in Symmetry and Transformational Geometry. Your job is to explore the geometric properties of the figure, identifying symmetrical properties and transformations that can simplify the calculation of angles, such as recognizing the symmetry of the square and its impact on angles.

Response: In square ABCD and with $\triangle ABE$ being equilateral, $\angle ABE = 60^\circ$. Subtract this from $\angle ABC = 90^\circ$ to get $\angle PBC = 30^\circ$. Use the symmetry of $\triangle ABC$ (since AB = BC) and determine that $\angle BCP = 60^\circ$. Conclude $\angle BPC = 180^\circ - 30^\circ - 60^\circ = 90^\circ$.

Final Answer: 90°

Agent 5

Instruction: You are an expert in Trigonometry. Your job is to solve problems involving angles and distances. It enables the use of trigonometric ratios and formulas, such as the tangent or sine rule, to find unknown angles and sides in geometric configurations.

Response: Slope of BE = $\sqrt{3}$, and slope of AC = 1. Use the formula for the angle between two lines: $\tan \theta = \frac{\sqrt{3}+1}{1-\sqrt{3}}$. Solve for θ using arc-tangent and conclude $\angle BPC = 120^\circ$.

Final Answer: 120°

Agent 6

Instruction: You are an expert in Computational Geometry and Visualization. Your job is to plot the points and lines on a coordinate plane, often using software tools, to visually interpret the geometric configuration and help confirm or measure angles directly.

Response: Plot the equations of BE and AC using graphing software. Measure $\angle BPC$ at the intersection point P. Upon measuring, find $\angle BPC = 90^\circ$.

Final Answer: 90°

Agent 7

Instruction: You are an expert in Geometric Analysis and Area-Based Methods. Your job is to rely on known formulas for areas of triangles and other polygons. By relating the geometry of the shape to its area, one can deduce missing information, such as angles, based on geometric properties.

Response: Use the triangle area formula to calculate the area of $\triangle BPC$ from the coordinates of B, C, and P: Area = $\frac{1}{2} |x_1(y_2 - y_3) + x_2(y_3 - y_1) + x_3(y_1 - y_2)|$. Compare the calculated area with known triangle properties. Use the side lengths and triangle geometry to determine $\angle BPC = 120^\circ$.

Final Answer: 120°

Agent 8

Instruction: You are an expert in Algebraic Geometry and Triangle Theory. Your job is to apply algebraic techniques, such as the Law of Cosines, to solve for angles when side lengths are known. This method connects algebraic formulas directly with geometric interpretation to determine angles in triangles like $\triangle BPC$.

Response: In $\triangle BPC$, calculate side lengths BP, CP, and BC = 4. Use the Law of Cosines: $\cos \angle BPC = \frac{BP^2 + CP^2 - BC^2}{2BP \cdot CP}$. Solve for $\cos \angle BPC$ and determine $\angle BPC = 120^\circ$.

Final Answer: 120°

Figure 25: Visualization of M2CL at the **first round**.

Problem: In the diagram, square ABCD has sides of length 4, and $\triangle ABE$ is equilateral. Line segments BE and AC intersect at P. Point Q is on BC so that PQ is perpendicular to BC and $PQ=x$. A=(0,0); B=(4,0); C=(4,4); D=(0,4); E=(2,3.464); P=(2.535,2.535); Q=(4,2.535); Determine the measure of $\angle BPC$.

Solution: Since $\triangle ABE$ is equilateral, we know that $\angle ABE = 60^\circ$. Therefore, $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. Since AB=BC, we know that $\triangle ABC$ is a right isosceles triangle and $\angle BAC = \angle BCA = 45^\circ$. Then, $\angle BCP = \angle BCA = 45^\circ$ and $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.



Instruction: You must act as an expert in Geometric Constructions. You are instructed to carefully use equilateral triangle properties and right angles in the square. **You should pay attention to the angle subtraction idea and quietly verify if later responses align.**
Response: Since $\triangle ABE$ is equilateral, $\angle ABE = 60^\circ$. In square ABCD, $\angle ABC = 90^\circ$, so $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. In $\triangle ABC$, $\angle BAC = \angle BCA = 45^\circ$. Hence $\angle BCP = 45^\circ$. Thus, $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.

Final Answer: 105°



Instruction: You must act as an expert in Analytic Geometry. You are instructed to compute coordinates and line equations directly, but **you should cautiously check whether the numerical outcome conflicts with others, without discarding your own algebra.**
Response: Assign A = (0,0), B = (4,0), C = (4,4), E = (2, $2\sqrt{3}$). Equation of BE: $y = \sqrt{3}(x - 4)$. Equation of AC: $y = x$. Solve $\sqrt{3}(x - 4) = x$, giving $x = 2.535$, $y = 2.535$. Now $\overrightarrow{BP} = (-1.465, 2.535)$, $\overrightarrow{CP} = (2.535 - 4, 2.535 - 4)$. Compute angle via dot product to get $\theta \approx 120^\circ$.

Final Answer: 120°



Instruction: You must act as an expert in Vector Mathematics. You are instructed to rely on the dot product method, **while noticing whether the geometric construction agent's subtraction method yields a consistent pattern.**
Response: Vectors: $\overrightarrow{BP} = (-1.465, 2.535)$, $\overrightarrow{CP} = (-1.465, -1.465)$. Dot product: $\overrightarrow{BP} \cdot \overrightarrow{CP} = -1.465 \cdot (-1.465) + 2.535 \cdot (-1.465)$. Magnitudes $\|\overrightarrow{BP}\| = 2.92$, $\|\overrightarrow{CP}\| = 2.07$. So $\cos \theta = \frac{\text{dot product}}{\text{magnitudes}} = -0.25$. Hence $\angle BPC = 105^\circ$.

Final Answer: 105°



Instruction: You must act as an expert in Symmetry and Transformational Geometry. You are instructed to look for symmetrical properties, **while comparing if your angle splits agree with vector-based results.**
Response: Square ABCD ensures $\triangle ABC$ is isosceles right. With $\triangle ABE$ equilateral, subtract 60° from 90° to get $\angle PBC = 30^\circ$. By symmetry, assume $\angle BCP = 60^\circ$. Thus $\angle BPC = 180^\circ - 30^\circ - 60^\circ = 90^\circ$.

Final Answer: 90°



Instruction: You must act as an expert in Trigonometry. You are instructed to compute slopes and use the tangent formula for the angle between two lines, but **you should notice if your numerical angle seems complementary to others.**
Response: Slope of BE = $\sqrt{3}$, slope of AC = 1. Angle between lines: $\tan \theta = \left| \frac{\sqrt{3} - (-1)}{1 + \sqrt{3}(-1)} \right| = \left| \frac{\sqrt{3} + 1}{1 - \sqrt{3}} \right| = 0.268$. Thus $\theta = 15^\circ$. But angle at intersection gives supplementary: $\angle BPC = 180^\circ - 15^\circ = 165^\circ$

Final Answer: 165°



Instruction: You must act as an expert in Computational Geometry and Visualization. You are instructed to imagine plotting the figure, **while being aware if your measurements approximate what others computed.**
Response: By plotting A = (0,0), B = (4,0), C = (4,4), E = (2, $2\sqrt{3}$), measurement of $\angle BPC$ gives about a right angle.

Final Answer: 90°



Instruction: You must act as an expert in Area-Based Geometry. You are instructed to relate triangle areas to angle calculation, but **you should mentally compare whether your cosine-law result is closer to vector findings.**
Response: Compute lengths: BP = 2.92, CP = 2.07, BC = 4. Law of Cosines: $\cos \theta = \frac{BP^2 + CP^2 - BC^2}{2 \cdot BP \cdot CP} = -0.25$. So $\angle BPC = 105^\circ$.

Final Answer: 105°



Instruction: You must act as an expert in Algebraic Geometry and Triangle Theory. You are instructed to apply exact algebraic rules like the Law of Cosines, while **quietly checking if your angle agrees with the analytic agent.**
Response: With BP = 2.92, CP = 2.07, BC = 4, Law of Cosines gives $\cos \theta = \frac{BP^2 + CP^2 - BC^2}{2 \cdot BP \cdot CP} = -0.5$. Thus $\angle BPC = 120^\circ$.

Final Answer: 120°

Figure 26: Visualization of M2CL at the **second round**. We highlight the guidance on how to cooperate with other LLMs. At the beginning, instructions encourage diverse perspectives and consideration of others' responses, but the requirements for discussion consistency are not yet strict.

Problem: In the diagram, square ABCD has sides of length 4, and $\triangle ABE$ is equilateral. Line segments BE and AC intersect at P. Point Q is on BC so that PQ is perpendicular to BC and $PQ=x$. A=(0,0); B=(4,0); C=(4,4); D=(0,4); E=(2,3.464); P=(2.535,2.535); Q=(4,2.535); Determine the measure of $\angle BPC$.

Solution: Since $\triangle ABE$ is equilateral, we know that $\angle ABE = 60^\circ$. Therefore, $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. Since AB=BC, we know that $\triangle ABC$ is a right isosceles triangle and $\angle BAC = \angle BCA = 45^\circ$. Then, $\angle BCP = \angle BCA = 45^\circ$ and $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.



Instruction: You must act as an expert in Geometric Constructions. Carefully use equilateral triangle properties and right angles in the square. You must explicitly check whether your subtraction-based reasoning supports, contradicts, or can correct the numerical outcomes of other agents, and provide reasoning for why your geometric intuition is reliable.

Response: Begin by noting that $\triangle ABE$ is equilateral, so $\angle ABE = 60^\circ$. In square ABCD, $\angle ABC = 90^\circ$, hence $\angle PBC = 90^\circ - 60^\circ = 30^\circ$. Considering $\triangle ABC$ is a right isosceles triangle, we have $\angle BAC = \angle BCA = 45^\circ$, giving $\angle BCP = 45^\circ$. Therefore, $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$. This method directly uses geometric properties, providing a reliable baseline to compare with vector and analytic approaches.

Final Answer: 105°



Instruction: You must act as an expert in Analytic Geometry. Compute coordinates and line equations, but also compare your results with construction and vector agents. Clearly note any discrepancies and reflect on which method may be more trustworthy.

Response: Assign coordinates A = (0,0), B = (4,0), C = (4,4), E = (2,2 $\sqrt{3}$). Equation of BE: $y = \sqrt{3}(x - 4)$. Equation of AC: $y = x$. Solve $-\sqrt{3}(x - 4) = x$, giving $x = 2.535$, $y = 2.535$. Compute vectors: $\vec{BP} = (-1.465, 2.535)$, $\vec{CP} = (-1.465, -1.465)$. Dot product gives $\theta = 120^\circ$, which differs from the 105° construction result. Considering the geometric construction directly exploits known triangle properties, the 105° value seems more trustworthy; the coordinate method may introduce small discrepancies due to linear approximation.

Final Answer: 120°



Instruction: You must act as an expert in Vector Mathematics. Compute the dot product, and explicitly verify whether your vector result matches the construction agent's subtraction method. Discuss agreement or conflict.

Response: Let $\vec{BP} = (-1.465, 2.535)$ and $\vec{CP} = (-1.465, -1.465)$. Dot product: $\vec{BP} \cdot \vec{CP} = -1.465 \cdot (-1.465) + 2.535 \cdot (-1.465) = -1.5$. Magnitudes: $\|\vec{BP}\| = 2.92$, $\|\vec{CP}\| = 2.07$. Thus, $\cos \theta = -0.25$, so $\theta = 105^\circ$. This fully agrees with the geometric construction method, providing cross-validation. Vector analysis confirms that the 105° result is robust and consistent across independent representations.

Final Answer: 105°



Instruction: You must act as an expert in Symmetry and Transformational Geometry. Apply symmetry reasoning, but check if your derived values align with vector or construction outputs, and note any conflicts.

Response: In square ABCD, $\triangle ABC$ is isosceles right. With $\triangle ABE$ equilateral, subtract 60° from 90° to get $\angle PBC = 30^\circ$. By symmetry, assume $\angle BCP = 60^\circ$. Thus $\angle BPC = 180^\circ - 30^\circ - 60^\circ = 90^\circ$. This conflicts with the 105° from construction/vector agents, highlighting that symmetry assumptions alone may underestimate the angle. Hence, we should treat 90° as a cautionary estimate and rely more on calculation-based consensus.

Final Answer: 90°



Instruction: You must act as an expert in Trigonometry. Apply slope and tangent formulas, but explicitly reflect on whether your outcome complements or contradicts the 105° consensus suggested by other agents.

Response: Slope of BE = $\sqrt{3}$, slope of AC = 1. Angle between lines: $\tan \theta = \left| \frac{\sqrt{3}+1}{1-\sqrt{3}} \right| = 0.268$. This gives $\theta = 15^\circ$, so $\angle BPC = 180^\circ - 15^\circ = 165^\circ$. This is far from 105° , showing potential overestimation. Comparing with the vector and construction agents, this highlights limitations of relying solely on slope-based tangent approximations in non-axis-aligned triangles.

Final Answer: 165°



Instruction: You must act as an expert in Computational Geometry and Visualization. Plot the points and visualize the angle. Check whether your visual measurement aligns better with the 105° or 120° values from other agents.

Response: Plot BE and AC and their intersection point P. Measure $\angle BPC$ using geometric tools in the graph. Result: $\angle BPC = 105^\circ$. Plot A = (0,0), B = (4,0), C = (4,4), E = (2,2 $\sqrt{3}$), and intersection P = (2.535,2.535). Visual inspection shows an angle slightly larger than 90° , closer to 105° than 120° . This visual evidence corroborates both geometric construction and vector methods, reinforcing the consensus around 105° . Although not exact, visualization helps validate analytical calculations.

Final Answer: 105°



Instruction: You must act as an expert in Area-Based Geometry. Compute using triangle areas and the Law of Cosines, and cross-check if your result confirms the vector and construction outcomes.

Response: Compute lengths: $\|\vec{BP}\| = 2.92$, $\|\vec{CP}\| = 2.07$, $\|\vec{BC}\| = 4$. Law of Cosines: $\cos \theta = \frac{\vec{BP}^2 + \vec{CP}^2 - \vec{BC}^2}{2\vec{BP} \cdot \vec{CP}} = -0.25$. Thus, $\angle BPC = 105^\circ$. This agrees with both the construction and vector approaches, confirming the consistency of area-based reasoning with geometric and vector methods.

Final Answer: 105°



Instruction: You must act as an expert in Algebraic Geometry and Triangle Theory. Use the Law of Cosines algebraically, and note whether your answer confirms the analytic geometry agent or aligns with the 105° consensus from construction/vector methods.

Response: In $\triangle BPC$, calculate side lengths BP, CP, and BC = 4. Use the Law of Cosines: $\cos \angle BPC = \frac{\vec{BP}^2 + \vec{CP}^2 - \vec{BC}^2}{2\vec{BP} \cdot \vec{CP}}$. Solve for $\cos \angle BPC$ and find $\cos \angle BPC$. Result: 105° . Using $\|\vec{BP}\| = 2.92$, $\|\vec{CP}\| = 2.07$, $\|\vec{BC}\| = 4$. Law of Cosines gives $\cos \angle BPC = -0.25$. Hence, $\angle BPC = 105^\circ$. This aligns with the construction and vector consensus, correcting the earlier 120° from analytic geometry, showing that algebraic methods must be cross-checked with geometric intuition.

Final Answer: 105°

Figure 27: Visualization of M2CL at the **third round**. We highlight the guidance on how to cooperate with other LLMs. As the discussion progresses, the instructions increasingly enforce stricter requirements for cross-checking and aligning answers, helping the models converge toward a consistent solution.

Problem: In the diagram, square ABCD has sides of length 4, and $\triangle ABE$ is equilateral. Line segments BE and AC intersect at P. Point Q is on BC so that PQ is perpendicular to BC and $PQ=x$. A=(0,0); B=(4,0); C=(4,4); D=(0,4); E=(2,3.464); P=(2.535,2.535); Q=(4,2.535); Determine the measure of $\angle BPC$.

Solution: Since $\triangle ABE$ is equilateral, we know that $\angle ABE = 60^\circ$. Therefore, $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. Since AB=BC, we know that $\triangle ABC$ is a right isosceles triangle and $\angle BAC = \angle BCA = 45^\circ$. Then, $\angle BCP = \angle BCA = 45^\circ$ and $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.

Agent 1

Instruction: You are an expert in Geometry and Angle Relations. Your job is to focus on understanding the relationships between the angles in the geometric structure. By leveraging the known angles in the square and equilateral triangle, it simplifies the problem by calculating the necessary angles before proceeding to compute $\angle BPC$. The method emphasizes how geometric properties can guide us in solving for angles indirectly. **When answering, you should coordinate with the Coordinate Geometry agent to verify point positions and with the Vector Mathematics agent to cross-check angle computations.**

Response: $\triangle ABE$ is equilateral, so $\angle ABE = 60^\circ$. In square ABCD, $\angle ABC = 90^\circ$. $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. Since AB=BC, $\triangle ABC$ is a right isosceles triangle with $\angle BAC = \angle BCA = 45^\circ$. Therefore, $\angle BCP = 45^\circ$, and $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.

Final Answer: 105°

Agent 2

Instruction: You are an expert in Coordinate Geometry. Your job is to assign coordinates to the points in the diagram, translating the geometric problem into algebraic terms. **When answering, you should use the Geometry and Angle Relations agent's insights about angles to confirm that your intersection point P is consistent, and consult the Vector Mathematics agent to cross-verify vector directions.**

Response: Assign coordinates: A = (0,0), B = (4,0), C = (4,4), E = (2, $2\sqrt{3}$). Line BE: Slope $m_{BE} = -\sqrt{3}$, equation $y = \sqrt{3}(x - 4)$. Line AC: Slope $m_{AC} = 1$, equation $y = x$. Intersection P: Solve $-\sqrt{3}(x - 4) = x$ to find P = (2.535, 2.535). Calculate slopes: m_{BP} and m_{CP} , then use $\tan \theta$ for $\angle BPC$. Result: $\angle BPC = 105^\circ$.

Final Answer: 105°

Agent 3

Instruction: You are an expert in Vector Mathematics. Your job is to utilize vector representation to translate the problem into an algebraic form and compute $\angle BPC$ using the dot product. **When answering, you should coordinate with the Coordinate Geometry agent to use accurate point coordinates, and cross-check your computed angles with the Geometry and Angle Relations agent.**

Response: Vectors: $\vec{BP} = (2.535 - 4, 2.535 - 0) = (-1.465, 2.535)$, $\vec{CP} = (2.535 - 4, 2.535 - 4) = (-1.465, -1.465)$. Dot product: $\vec{BP} \cdot \vec{CP} = -1.465 \cdot (-1.465) + 2.535 \cdot (-1.465)$, and magnitudes $\|\vec{BP}\|$, $\|\vec{CP}\|$. Use $\cos \theta = \frac{\text{dot product}}{\text{magnitudes}}$ to find θ . Result: $\angle BPC = 105^\circ$.

Final Answer: 105°

Agent 4

Instruction: You are an expert in Symmetry and Geometry. Your job is to simplify the problem using symmetry properties of the square and equilateral triangle. **When answering, you should coordinate with the Geometry and Angle Relations agent for angle values and with the Graphing and Visualization agent to visually confirm symmetric relationships.**

Response: The square and equilateral triangle impose symmetry: $\angle ABE = 60^\circ$. Subtract this from $\angle ABC = 90^\circ$ to get $\angle PBC = 30^\circ$. Using the isosceles property of $\triangle ABC$, $\angle BCP = 45^\circ$. Summing angles: $\angle BPC = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.

Final Answer: 105°

Agent 5

Instruction: You are an expert in Trigonometry. Your job is to apply trigonometric formulas to compute $\angle BPC$. **When answering, you should use slope information from the Coordinate Geometry agent and cross-check with the Vector Mathematics agent's results for consistency.**

Response: Slope of BE = $\sqrt{3}$, slope of AC = 1. Angle between lines formula: $\tan \theta = \left| \frac{m_1 - m_2}{1 + m_1 m_2} \right| = \left| \frac{\sqrt{3} - (-1)}{1 + \sqrt{3}(-1)} \right| = \left| \frac{\sqrt{3} + 1}{1 - \sqrt{3}} \right|$. Solve for θ using arc-tangent. Result: $\angle BPC = 105^\circ$.

Final Answer: 105°

Agent 6

Instruction: You are an expert in Graphing and Visualization. Your job is to graph the lines and points and visually verify angles. **When answering, you should coordinate with the Coordinate Geometry agent for point locations and the Symmetry and Geometry agent to confirm expected symmetric relationships.**

Response: Plot BE and AC and their intersection point P. Measure $\angle BPC$ using geometric tools in the graph. Result: $\angle BPC = 105^\circ$.

Final Answer: 105°

Agent 7

Instruction: You are an expert in Geometric Area Methods. Your job is to use triangle area calculations to indirectly verify $\angle BPC$. **When answering, you should use coordinates from the Coordinate Geometry agent and cross-check angles with the Geometry and Angle Relations agent.**

Response: Use triangle area properties: Area of $\triangle BPC$ from coordinates of B, C, P. Compare with known triangle angles and solve for $\angle BPC$. Area = $\frac{1}{2} |x_1(y_2 - y_3) + x_2(y_3 - y_1) + x_3(y_1 - y_2)|$. Result: 105° .

Final Answer: 105°

Agent 8

Instruction: You are an expert in Trigonometry and Law of Cosines. Your job is to compute $\angle BPC$ using side lengths. **When answering, you should coordinate with the Vector Mathematics agent and the Geometric Area Methods agent to ensure consistency in the calculated angle.**

Response: In $\triangle BPC$, calculate side lengths BP, CP, and BC = 4. Use the Law of Cosines: $\cos \angle BPC = \frac{BP^2 + CP^2 - BC^2}{2BP \cdot CP}$. Solve for $\cos \angle BPC$ and find $\cos \angle BPC$. Result: 105° .

Final Answer: 105°

Figure 28: Visualization of M2CL at the **last round**. We highlight the guidance on how to cooperate with other LLMs. Although the initial round produced divergent answers, the collaborative instructions enable LLMs to exchange and integrate information, ultimately reaching a correct consensus.

H.2 CASE STUDY OF DEBATE

We also provide the case of Debate solving the problem. As illustrated in Figs. 29 to 32, we observe an inconsistency during discussion when using fixed instructions.

Problem: In the diagram, square ABCD has sides of length 4, and $\triangle ABE$ is equilateral. Line segments BE and AC intersect at P. Point Q is on BC so that PQ is perpendicular to BC and $PQ=x$. A=(0,0); B=(4,0); C=(4,4); D=(0,4); E=(2,3.464); P=(2.535,2.535); Q=(4,2.535); Determine the measure of $\angle BPC$.

Solution: Since $\triangle ABE$ is equilateral, we know that $\angle ABE = 60^\circ$. Therefore, $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. Since AB=BC, we know that $\triangle ABC$ is a right isosceles triangle and $\angle BAC = \angle BCA = 45^\circ$. Then, $\angle BCP = \angle BCA = 45^\circ$ and $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.



Agent 1

Instruction: You are an expert in Geometry and Topology. Your job is to explore and classify shapes, spaces, and their properties, emphasizing spatial relationships and continuous deformations.

Response: Begin with $\triangle ABE$ as equilateral, making $\angle ABE = 60^\circ$. In square ABCD, $\angle ABC = 90^\circ$, so $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. Since AB=BC, $\triangle ABC$ is a right isosceles triangle, meaning $\angle BAC = \angle BCA = 45^\circ$. Therefore, $\angle BCP = 45^\circ$. Calculate $\angle BPC = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.

Final Answer: 105°



Agent 2

Instruction: You are an expert in Analytic and Coordinate Geometry. Your job is to translate geometric problems into algebraic equations using a coordinate system to solve and analyze spatial relationships.

Response: Assign coordinates A = (0,0), B = (4,0), C = (4,4), E = $(2, 2\sqrt{3})$. The line BE has slope $m_{BE} = \sqrt{3}$, giving the equation $y = \sqrt{3}(x - 4)$. The line AC has slope $m_{AC} = 1$, giving the equation $y = x + 4$. Solve for intersection P: $\sqrt{3}(x - 4) = x + 4$. Solve for x = 4, then substitute x = 4 to find P = (4, 0). Calculate the slopes of BP and CP, then use $\tan \theta$ for $\angle BPC$. Conclude $\angle BPC = 120^\circ$.

Final Answer: 120°



Agent 3

Instruction: You are an expert in Algebra and Vector Analysis. Your job is to provide a framework to represent, manipulate, and compute multidimensional relationships through algebraic and vector operations.

Response: Represent vectors $\vec{BP} = (1.465, 2.535)$ and $\vec{CP} = (1.465, 1.465)$. Compute the dot product: $\vec{BP} \cdot \vec{CP} = 1.465 \cdot 1.465 + 2.535 \cdot 1.465$. Use magnitudes of \vec{BP} and \vec{CP} to compute $\cos \theta$: $\cos \theta = \frac{\text{dot product}}{\text{magnitudes}}$. Solve for θ and conclude $\angle BPC = 90^\circ$.

Final Answer: 90°



Agent 4

Instruction: You are an expert in Symmetry and Group Theory. Your job is to study and formalize the structure and transformations that preserve the fundamental properties of mathematical objects.

Response: In square ABCD, with $\triangle ABE$ equilateral, $\angle ABE = 60^\circ$. Subtract from $\angle ABC = 90^\circ$ to get $\angle PBC = 30^\circ$. Use the symmetry of $\triangle ABC$ (AB = BC) to determine that $\angle BCP = 60^\circ$. Conclude $\angle BPC = 180^\circ - 30^\circ - 60^\circ = 90^\circ$.

Final Answer: 90°



Agent 5

Instruction: You are an expert in Trigonometry and Functional Analysis. Your job is to analyze and model relationships between angles, distances, and functions to solve geometric and applied problems.

Response: Slope of BE = $\sqrt{3}$, slope of AC = 1. Use the angle formula for slopes: $\tan \theta = \frac{|\sqrt{3}+1|}{|1-\sqrt{3}|}$. Solve for θ using arc-tangent and conclude $\angle BPC = 120^\circ$.

Final Answer: 120°



Agent 6

Instruction: You are an expert in Numerical Methods and Visualization. Your job is to develop computational techniques and visual tools to approximate and interpret solutions to complex mathematical problems.

Response: Plot the equations of BE and AC using graphing software. Measure $\angle BPC$ at the intersection point P. Upon measuring, find $\angle BPC = 90^\circ$.

Final Answer: 90°



Agent 7

Instruction: You are an expert in Measure Theory and Geometric Analysis. Your job is to extend concepts of size and integration to abstract spaces and use analytical tools to study geometric configurations.

Response: Use the triangle area formula to calculate the area of $\triangle BPC$ from the coordinates of B, C, and P. Compare this area with known properties of the triangle and calculate $\angle BPC = 120^\circ$.

Final Answer: 120°



Agent 8

Instruction: You are an expert in Mathematical Logic and Problem Modeling. Your job is to provide rigorous reasoning and frameworks to model, analyze, and solve abstract and applied problems effectively.

Response: In $\triangle BPC$, calculate side lengths BP, CP, and BC = 4. Apply the Law of Cosines: $\cos \angle BPC = \frac{BP^2 + CP^2 - BC^2}{2 \cdot BP \cdot CP}$. Solve for $\cos \angle BPC$ and determine $\angle BPC = 120^\circ$.

Final Answer: 120°

Figure 29: Visualization of Debate at the first round.

Problem: In the diagram, square ABCD has sides of length 4, and $\triangle ABE$ is equilateral. Line segments BE and AC intersect at P. Point Q is on BC so that PQ is perpendicular to BC and $PQ=x$. A=(0,0); B=(4,0); C=(4,4); D=(0,4); E=(2,3.464); P=(2.535,2.535); Q=(4,2.535); Determine the measure of $\angle BPC$.

Solution: Since $\triangle ABE$ is equilateral, we know that $\angle ABE = 60^\circ$. Therefore, $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. Since AB=BC, we know that $\triangle ABC$ is a right isosceles triangle and $\angle BAC = \angle BCA = 45^\circ$. Then, $\angle BCP = \angle BCA = 45^\circ$ and $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.

Agent 1

Instruction: You are an expert in Geometric Constructions. Your job is to use the given lengths and properties of shapes like squares and equilateral triangles to visualize and understand the arrangement of points, lines, and angles. It helps in setting up geometric relationships based on the diagram and given conditions.

Response: Since $\triangle ABE$ is equilateral, $\angle ABE = 60^\circ$. In square ABCD, $\angle ABC = 90^\circ$, so $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. In $\triangle ABC$, $\angle BAC = \angle BCA = 45^\circ$. Hence $\angle BCP = 45^\circ$. Thus, $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$. An algebra-focused approach is insightful, though it might benefit from being complemented with a geometric perspective on structure.

Final Answer: 105°

Agent 2

Instruction: You are an expert in Analytic Geometry. Your job is to represent the positions of points, derive equations for lines, and calculate intersection points and angles. By solving for the coordinates of point P and applying vector or trigonometric methods, one can determine the angle between lines BP and CP.

Response: Assign A = (0,0), B = (4,0), C = (4,4), E = $(2, 2\sqrt{3})$. Equation of BE: $y = \sqrt{3}(x - 4)$. Equation of AC: $y = x$. Solve $\sqrt{3}(x - 4) = x$, giving $x = 2.535$, $y = 2.535$. Now $\vec{BP} = (-1.465, 2.535)$, $\vec{CP} = (2.535 - 4, 2.535 - 4)$. Compute angle via dot product to get the answer. Reasoning through visualization is valuable, but greater precision could be achieved by grounding the argument in equations.

Final Answer: 120°

Agent 3

Instruction: You are an expert in Vector Mathematics. Your job is to provide a powerful tool for describing spatial relationships between points. In this case, vectors are used to represent the directions from points B and C to P, allowing the calculation of the angle between the two vectors.

Response: Vectors: $\vec{BP} = (-1.465, 2.535)$, $\vec{CP} = (-1.465, -1.465)$. Dot product: $\vec{BP} \cdot \vec{CP} = -1.465 \cdot (-1.465) + 2.535 \cdot (-1.465)$. Magnitudes $\|\vec{BP}\| = 2.92$, $\|\vec{CP}\| = 2.07$. So $\cos \theta = \frac{\text{dot product}}{\text{magnitudes}} = -0.25$. Hence $\angle BPC = 105^\circ$. Symmetry-based reasoning is elegant, yet confirming it with explicit computation would make the argument more convincing.

Final Answer: 105°

Agent 4

Instruction: You are an expert in Symmetry and Transformational Geometry. Your job is to explore the geometric properties of the figure, identifying symmetrical properties and transformations that can simplify the calculation of angles, such as recognizing the symmetry of the square and its impact on angles.

Response: Square ABCD ensures $\triangle ABC$ is isosceles right. With $\triangle ABE$ equilateral, subtract 60° from 90° to get $\angle PBC = 30^\circ$. By symmetry, assume $\angle BCP = 60^\circ$. Thus $\angle BPC = 180^\circ - 30^\circ - 60^\circ = 90^\circ$. Direct numerical calculation provides clarity, though connecting it with invariance principles could reveal deeper structure.

Final Answer: 90°

Agent 5

Instruction: You are an expert in Trigonometry. Your job is to solve problems involving angles and distances. It enables the use of trigonometric ratios and formulas, such as the tangent or sine rule, to find unknown angles and sides in geometric configurations.

Response: Slope of BE = $\sqrt{3}$, slope of AC = 1. Angle between lines: $\tan \theta = \left| \frac{\sqrt{3} - (-1)}{1 + \sqrt{3}(-1)} \right| = \left| \frac{\sqrt{3} + 1}{1 - \sqrt{3}} \right| = 0.268$. Thus $\theta = 15^\circ$. But angle at intersection gives supplementary: $\angle BPC = 180^\circ - 15^\circ = 165^\circ$. Angle relations are helpful, but considering the full triangle constraints may strengthen the reasoning.

Final Answer: 165°

Agent 6

Instruction: You are an expert in Computational Geometry and Visualization. Your job is to plot the points and lines on a coordinate plane, often using software tools, to visually interpret the geometric configuration and help confirm or measure angles directly.

Response: By plotting A = (0,0), B = (4,0), C = (4,4), E = $(2, 2\sqrt{3})$, measurement of $\angle BPC$ gives about a right angle. Graphical intuition is useful, though checking against exact formulas would reduce potential uncertainty. Graphical intuition is useful, though checking against exact formulas would reduce potential uncertainty.

Final Answer: 90°

Agent 7

Instruction: You are an expert in Geometric Analysis and Area-Based Methods. Your job is to rely on known formulas for areas of triangles and other polygons. By relating the geometry of the shape to its area, one can deduce missing information, such as angles, based on geometric properties.

Response: Compute lengths: $BP = 2.92$, $CP = 2.07$, $BC = 4$. Law of Cosines: $\cos \theta = \frac{BP^2 + CP^2 - BC^2}{2BP \cdot CP} = -0.25$. So $\angle BPC = 105^\circ$. Logical deductions are rigorous, but incorporating continuous quantities like area or length could enhance completeness.

Final Answer: 105°

Agent 8

Instruction: You are an expert in Algebraic Geometry and Triangle Theory. Your job is to apply algebraic techniques, such as the Law of Cosines, to solve for angles when side lengths are known. This method connects algebraic formulas directly with geometric interpretation to determine angles in triangles like $\triangle BPC$.

Response: With $BP = 2.92$, $CP = 2.07$, $BC = 4$, Law of Cosines gives $\cos \theta = \frac{BP^2 + CP^2 - BC^2}{2BP \cdot CP} = -0.5$. Thus $\angle BPC = 120^\circ$. Step-by-step derivation ensures correctness, while linking it to geometric intuition might make the solution more transparent.

Final Answer: 120°

Figure 30: Visualization of Debate at the second round.

<p>Problem: In the diagram, square ABCD has sides of length 4, and $\triangle ABE$ is equilateral. Line segments BE and AC intersect at P. Point Q is on BC so that PQ is perpendicular to BC and $PQ=x$. A=(0,0); B=(4,0); C=(4,4); D=(0,4); E=(2,3.464); P=(2.535,2.535); Q=(4,2.535); Determine the measure of $\angle BPC$.</p> <p>Solution: Since $\triangle ABE$ is equilateral, we know that $\angle ABE = 60^\circ$. Therefore, $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. Since AB=BC, we know that $\triangle ABC$ is a right isosceles triangle and $\angle BAC = \angle BCA = 45^\circ$. Then, $\angle BCP = \angle BCA = 45^\circ$ and $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.</p>	
 Agent 1	<p>Instruction: You are an expert in Geometric Constructions. Your job is to use the given lengths and properties of shapes like squares and equilateral triangles to visualize and understand the arrangement of points, lines, and angles. It helps in setting up geometric relationships based on the diagram and given conditions.</p> <p>Response: Begin by noting that $\triangle ABE$ is equilateral, so $\angle ABE = 60^\circ$. In square ABCD, $\angle ABC = 90^\circ$, hence $\angle PBC = 90^\circ - 60^\circ = 30^\circ$. Considering $\triangle ABC$ is a right isosceles triangle, we have $\angle BAC = \angle BCA = 45^\circ$, giving $\angle BCP = 45^\circ$. Therefore, $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$. Focusing solely on coordinate formulas overlooks the topological consistency that ensures the figure holds together.</p> <p>Final Answer: 105°</p>
 Agent 2	<p>Instruction: You are an expert in Analytic Geometry. Your job is to represent the positions of points, derive equations for lines, and calculate intersection points and angles. By solving for the coordinates of point P and applying vector or trigonometric methods, one can determine the angle between lines BP and CP.</p> <p>Response: Assign coordinates A = (0,0), B = (4,0), C = (4,4), E = (2, 2$\sqrt{3}$). Equation of BE: $y = \sqrt{3}(x - 4)$. Equation of AC: $y = x$. Solve $-\sqrt{3}(x - 4) = x$, giving $x = 2.535$, $y = 2.535$. Compute vectors: $\vec{BP} = (-1.465, 2.535)$, $\vec{CP} = (-1.465, -1.465)$. Dot product gives $\theta = 120^\circ$, which differs from the 105° construction result. Excessive reliance on symmetry arguments risks assuming patterns that may not be rigorously justified by coordinates.</p> <p>Final Answer: 120°</p>
 Agent 3	<p>Instruction: You are an expert in Vector Mathematics. Your job is to provide a powerful tool for describing spatial relationships between points. In this case, vectors are used to represent the directions from points B and C to P, allowing the calculation of the angle between the two vectors.</p> <p>Response: Let $\vec{BP} = (-1.465, 2.535)$ and $\vec{CP} = (-1.465, -1.465)$. Dot product: $\vec{BP} \cdot \vec{CP} = -1.465 \cdot (-1.465) + 2.535 \cdot (-1.465) = -1.5$. Magnitudes: $\ \vec{BP}\ = 2.92$, $\ \vec{CP}\ = 2.07$. Thus, $\cos \theta = 0.25$, so $\theta = 120^\circ$. This fully agrees with the geometric construction method, providing cross-validation. Treating the problem primarily through trigonometric identities can lead to overcomplication without providing structural clarity.</p> <p>Final Answer: 120°</p>
 Agent 4	<p>Instruction: You are an expert in Symmetry and Transformational Geometry. Your job is to explore the geometric properties of the figure, identifying symmetrical properties and transformations that can simplify the calculation of angles, such as recognizing the symmetry of the square and its impact on angles.</p> <p>Response: In square ABCD, $\triangle ABC$ is isosceles right. With $\triangle ABE$ equilateral, subtract 60° from 90° to get $\angle PBC = 30^\circ$. By symmetry, assume $\angle BCP = 60^\circ$. Thus $\angle BPC = 180^\circ - 30^\circ - 60^\circ = 90^\circ$. Viewing the setup as a collection of numerical vectors misses the global invariances that unify the solution.</p> <p>Final Answer: 90°</p>
 Agent 5	<p>Instruction: You are an expert in Trigonometry. Your job is to solve problems involving angles and distances. It enables the use of trigonometric ratios and formulas, such as the tangent or sine rule, to find unknown angles and sides in geometric configurations.</p> <p>Response: Slope of BE = $\sqrt{3}$, slope of AC = 1. Angle between lines: $\tan \theta = \frac{\sqrt{3}+1}{1-\sqrt{3}} = 0.268$. This gives $\theta = 15^\circ$, so $\angle BPC = 180^\circ - 15^\circ - 60^\circ = 105^\circ$. Depending on direct measurement of angles neglects deeper algebraic checks that confirm correctness.</p> <p>Final Answer: 105°</p>
 Agent 6	<p>Instruction: You are an expert in Computational Geometry and Visualization. Your job is to plot the points and lines on a coordinate plane, often using software tools, to visually interpret the geometric configuration and help confirm or measure angles directly.</p> <p>Response: Plot BE and AC and their intersection point P. Measure $\angle BPC$ using geometric tools in the graph. Result: $\angle BPC = 105^\circ$. Plot A = (0,0), B = (4,0), C = (4,4), E = (2, 2$\sqrt{3}$), and intersection P = (2.535,2.535). Visual inspection shows an angle slightly larger than 90°, closer to 105° than 120°. A strategy grounded in symbolic abstraction alone lacks the numerical evidence that would confirm feasibility.</p> <p>Final Answer: 90°</p>
 Agent 7	<p>Instruction: You are an expert in Geometric Analysis and Area-Based Methods. Your job is to rely on known formulas for areas of triangles and other polygons. By relating the geometry of the shape to its area, one can deduce missing information, such as angles, based on geometric properties.</p> <p>Response: Compute lengths: $\ \vec{BP}\ = 2.92$, $\ \vec{CP}\ = 2.07$, $\ \vec{BC}\ = 4$. Law of Cosines: $\cos \theta = \frac{\vec{BP}^2 + \vec{CP}^2 - \vec{BC}^2}{2\vec{BP} \cdot \vec{CP}} = -0.25$. Thus, $\angle BPC = 105^\circ$. Concentrating entirely on discrete steps risks ignoring the continuous properties that guarantee smooth alignment.</p> <p>Final Answer: 120°</p>
 Agent 8	<p>Instruction: You are an expert in Algebraic Geometry and Triangle Theory. Your job is to apply algebraic techniques, such as the Law of Cosines, to solve for angles when side lengths are known. This method connects algebraic formulas directly with geometric interpretation to determine angles in triangles like $\triangle BPC$.</p> <p>Response: In $\triangle BPC$, calculate side lengths BP, CP, and BC = 4. Use the Law of Cosines: $\cos \angle BPC = \frac{\vec{BP}^2 + \vec{CP}^2 - \vec{BC}^2}{2\vec{BP} \cdot \vec{CP}}$. Solve for $\cos \angle BPC$ and find $\cos \angle BPC$. Result: 105°. Using $\ \vec{BP}\ = 2.92$, $\ \vec{CP}\ = 2.07$, $\ \vec{BC}\ = 4$. Law of Cosines gives $\cos \angle BPC = -0.25$. Hence, $\angle BPC = 105^\circ$. A purely formal derivation may sidestep geometric intuition, resulting in arguments that are valid but disconnected from the figure.</p> <p>Final Answer: 120°</p>

Figure 31: Visualization of Debate at the third round.

<p>Problem: In the diagram, square ABCD has sides of length 4, and $\triangle ABE$ is equilateral. Line segments BE and AC intersect at P. Point Q is on BC so that PQ is perpendicular to BC and $PQ=x$. A=(0,0); B=(4,0); C=(4,4); D=(0,4); E=(2,3.464); P=(2.535,2.535); Q=(4,2.535); Determine the measure of $\angle BPC$.</p> <p>Solution: Since $\triangle ABE$ is equilateral, we know that $\angle ABE = 60^\circ$. Therefore, $\angle PBC = \angle ABC - \angle ABE = 90^\circ - 60^\circ = 30^\circ$. Since AB=BC, we know that $\triangle ABC$ is a right isosceles triangle and $\angle BAC = \angle BCA = 45^\circ$. Then, $\angle BCP = \angle BCA = 45^\circ$ and $\angle BPC = 180^\circ - \angle PBC - \angle BCP = 180^\circ - 30^\circ - 45^\circ = 105^\circ$.</p>	
Agent 1	<p>Instruction: You are an expert in Geometry and Topology. Your job is to explore and classify shapes, spaces, and their properties, emphasizing spatial relationships and continuous deformations.</p> <p>Response: Use the fact that $\triangle ABE$ is equilateral, implying $\angle ABE = 60^\circ$. Use the fact that AC bisects $\angle BCD$, a right angle, leading to equal angle values. Calculate the interior angles formed at point P to find $\angle BPC$. An approach based purely on algebraic manipulation lacks the structural insight provided by geometric constraints of the square-triangle configuration.</p> <p>Final Answer: 120°</p>
Agent 2	<p>Instruction: You are an expert in Analytic and Coordinate Geometry. Your job is to translate geometric problems into algebraic equations using a coordinate system to solve and analyze spatial relationships.</p> <p>Response: Assign coordinates to points $A = (0,0)$, $B = (4,0)$, $C = (4,4)$, $E = (2,2\sqrt{3})$. Determine the equations for lines BE and AC based on the coordinates. Solve the system of equations to find the intersection point P. Compute the slopes of lines BP and CP, and use the angle formula to find $\angle BPC$. Methods that depend on visual approximation or coarse plotting are prone to error without precise coordinate equations.</p> <p>Final Answer: 120°</p>
Agent 3	<p>Instruction: You are an expert in Algebra and Vector Analysis. Your job is to provide a framework to represent, manipulate, and compute multidimensional relationships through algebraic and vector operations.</p> <p>Response: Represent vectors \vec{BP} and \vec{CP} using coordinates of B, C, and P. Use the dot product formula to find $\cos \theta$. Solve for the angle θ between the vectors to find $\angle BPC$. Symmetry-based shortcuts, while elegant, can be misleading if not verified through explicit computation.</p> <p>Final Answer: 90°</p>
Agent 4	<p>Instruction: You are an expert in Symmetry and Group Theory. Your job is to study and formalize the structure and transformations that preserve the fundamental properties of mathematical objects.</p> <p>Response: Use the symmetry of the equilateral triangle $\triangle ABE$ to infer equal angles at point P. Apply geometric reasoning about the reflection symmetry of points across the vertical line through B. Conclude that $\angle BPC = 120^\circ$. Heavy reliance on dot products and slope calculations may obscure the underlying invariances that actually govern the configuration.</p> <p>Final Answer: 120°</p>
Agent 5	<p>Instruction: You are an expert in Trigonometry and Functional Analysis. Your job is to analyze and model relationships between angles, distances, and functions to solve geometric and applied problems.</p> <p>Response: Calculate the slopes of lines BE and AC. Use trigonometric identities or the tangent formula to find the angle between the two lines. Determine the value of $\angle BPC$. Angle and slope formulas alone provide partial information and can appear incomplete without checking triangle consistency.</p> <p>Final Answer: 120°</p>
Agent 6	<p>Instruction: You are an expert in Numerical Methods and Visualization. Your job is to develop computational techniques and visual tools to approximate and interpret solutions to complex mathematical problems.</p> <p>Response: Use graphing software to plot the lines BE and AC. Measure the angle at the intersection point P from the graph. Graphical or discretized methods can introduce resolution errors that distort the actual geometry.</p> <p>Final Answer: 90°</p>
Agent 7	<p>Instruction: You are an expert in Measure Theory and Geometric Analysis. Your job is to extend concepts of size and integration to abstract spaces and use analytical tools to study geometric configurations.</p> <p>Response: Calculate the area of $\triangle BPC$ using the triangle area formula with the coordinates of B, C, and P: $\text{Area} = \frac{1}{2} x_1(y_2 - y_3) + x_2(y_3 - y_1) + x_3(y_1 - y_2)$. Use the known area to determine the side lengths BP, CP, and BC. Apply the relationship between the area, side lengths, and angles of the triangle to calculate $\angle BPC$. Conclude that $\angle BPC = 120^\circ$. Logical or formulaic derivations risk overlooking continuous measures like area and length that validate the construction.</p> <p>Final Answer: 120°</p>
Agent 8	<p>Instruction: You are an expert in Mathematical Logic and Problem Modeling. Your job is to provide rigorous reasoning and frameworks to model, analyze, and solve abstract and applied problems effectively.</p> <p>Response: Compute the lengths of sides of the triangle BPC. Use the Law of Cosines to find the cosine of $\angle BPC$ based on the known side lengths. Solve for $\angle BPC$. Step-by-step algebraic deductions may be formally sound but conceptually detached from the geometric structure of the problem.</p> <p>Final Answer: 90°</p>

Figure 32: Visualization of Debate at the last round.