

Learning to Deliberate: Meta-policy Collaboration for Agentic LLMs with Multi-agent Reinforcement Learning

Wei Yang and Jesse Thomason

University of Southern California, Los Angeles, CA, USA
{wyang930, jessetho}@usc.edu

Abstract

Multi-agent systems of large language models (LLMs) show promise for complex reasoning, but their effectiveness is often limited by fixed collaboration protocols. These frameworks typically focus on macro-level orchestration while overlooking agents' internal deliberative capabilities. This critical meta-cognitive blindspot treats agents as passive executors unable to adapt their strategy based on internal cognitive states like uncertainty or confidence. We introduce the Meta-Policy Deliberation Framework (MPDF), where agents learn a decentralized policy over a set of high-level meta-cognitive actions: Persist, Refine, and Concede. To overcome the instability of traditional policy gradients in this setting, we develop SoftRankPO, a novel reinforcement learning algorithm. SoftRankPO stabilizes training by shaping advantages based on the rank of rewards mapped through smooth normal quantiles, making the learning process robust to reward variance. Experiments show that MPDF with SoftRankPO achieves a 4-5% absolute gain in average accuracy across six mathematical and general reasoning benchmarks compared to state-of-the-art heuristic and learning-based multi-agent reasoning algorithms. Our work presents a paradigm for learning adaptive, meta-cognitive policies for multi-agent LLM systems, shifting the focus from designing fixed protocols to learning dynamic, deliberative strategies.

Introduction

As the limitations of single-agent reasoning become apparent in complex tasks (Park et al. 2023; Hong et al. 2023), the research frontier is shifting toward multi-agent systems (MAS), where multiple large language model (LLM) agents engage in collaborative reasoning to improve accuracy and robustness (Chen et al. 2025d; Ping et al. 2025a; Yang et al. 2025). Their effectiveness often depends on well-designed collaboration protocols, such as multi-turn debates and peer review mechanisms (Du et al. 2023; Liu et al. 2024). However, static protocols prevent agents from adapting their behavior to the unique demands of a given problem and to evolving context (Du et al. 2024; Yue et al. 2025).

Recent efforts to overcome the static protocol bottleneck have largely focused on dynamic, system-level optimizations (Zhao et al. 2023), which we categorize as *macro-level*

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

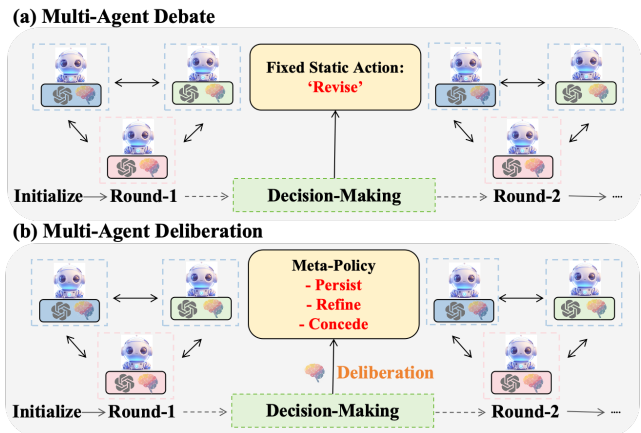


Figure 1: Evolving from fixed static debate to trainable meta-policy deliberation, with dynamic strategy adaptation for multi-agent collaboration.

orchestration. These methods optimize the connections between agents, but critically neglect the deliberative *meta-cognition* (Sun, Zhang, and Mathews 2006; Yeung and Summerfield 2012; Koriati 2015) within each agent. This *Meta-Cognitive Blindspot* prevents agents from differentiating between confident judgments and uncertain hypotheses, causing them to uniformly apply external protocols irrespective of their internal cognitive state, falling back to a *Static Deliberation Strategy*. That is, agents employ fixed or pre-programmed response behaviors, lacking a strategic faculty to adaptively choose among persisting, refining, or conceding actions based on context and internal confidence.

To address these deficiencies, we introduce the **Meta-Policy Deliberation Framework (MPDF)**, which formulates multi-agent collaboration as a decentralized partially observable Markov decision process (Dec-POMDP). Unlike prior approaches relying on static heuristics or predefined rules, our framework equips each agent with a lightweight, decoupled policy network. Each agent explicitly reasons about its internal cognitive state and strategically selects from discrete deliberation acts such as PERSIST, REFINE, or CONCEDE. This meta-cognitive capability enables agents to dynamically calibrate their behavior according to evolving

confidence and situational context.

To robustly optimize this policy under sparse and noisy team-level feedback, we introduce **SoftRankPO**, a scale-resilient reinforcement learning algorithm. SoftRankPO helps stabilize policy gradients by converting raw rewards into rank-based advantages derived from smooth Gaussian percentiles, thereby mitigating reward-scale sensitivity and high-variance challenges.

Our main contributions are:

- A meta-policy framework formulating multi-agent LLM reasoning as a Dec-POMDP for adaptive coordination through learned decision-making dynamics.
- SoftRankPO, a scale-resilient policy optimization algorithm that stabilizes learning via smooth rank-based advantage shaping, enabling low-variance and reliable convergence across diverse reward regimes.
- Consistent accuracy gains over competitive baselines on six reasoning benchmarks by combining meta-cognitive coordination and scale-resilient optimization.

Related Works

Multi-Agent LLM Collaboration. Beyond large language models (Chen et al. 2025c; Ping et al. 2025b; Li et al. 2025), collaborative multi-agent LLM frameworks for reasoning and planning (Motwani et al. 2024; Ishibashi and Nishimura 2024; Trirat, Jeong, and Hwang 2024; Yan et al. 2025; Dai et al. 2025; Chen, Han, and Zhang 2024; Chang et al. 2025; Zhang et al. 2025) mainly follow two paradigms. First, prestructured approaches rely on fixed coordination schemes (Liu et al. 2023), like debate (Du et al. 2023; Liu et al. 2024) and peer review pipelines. These systems adopt static topologies such as chains, trees, or graphs to guide communication (Du et al. 2023; Qian et al. 2024). Second, self-organizing methods adapt interactions during inference by restructuring collaboration graphs using Monte Carlo tree search, evolutionary updates, and pruning (Hu, Lu, and Clune 2024; Shang et al. 2024; Zhang et al. 2024a). Some frame workflow design as program synthesis over modules (Zhuge et al. 2024; Zhang et al. 2024b), or assign dynamic roles across agents (Hu et al. 2024b; Yue et al. 2025). These methods have advanced the design space of LLM-based collaboration, but rely on heuristic control or static policies (Wei, Shan, and Li 2025). We overcome this limitation by introducing closed-loop learning to adapt collaboration based on downstream utility.

Multi-Agent Reinforcement Learning for LLMs. Reinforcement learning (RL) (Chen et al. 2022a,b; Hu et al. 2024a) for multi-agent LLM systems aim to train coordinated, role-specialized agents beyond static prompting or supervised tuning (Zheng et al. 2023; Zhu et al. 2025; Peiyuan et al. 2024; Park et al. 2025; Liao et al. 2025; Han et al. 2025; Wan et al. 2025; Zhuang et al. 2024). For example, MA-PoRL (Park et al. 2025) co-trains multiple LLMs through multi-turn debates using verifier-shaped collaborative rewards, and LERO (Wei, Shan, and Li 2025) employs LLMs in an evolutionary MARL loop to generate hybrid reward and observation functions. To support scalable optimization, frameworks like Maestro (Yang et al. 2025) and MHGPO

(Chen et al. 2025a) adopt centralized coordination with decentralized policy learning. Others, such as ReSo (Zhou et al. 2025) and Verco (Li et al. 2024) introduce dynamic agent selection and interpretable communication protocols. Meanwhile, ML-Agent (Liu et al. 2025) and RAGEN (Wang et al. 2025) explore step-wise RL and long-horizon reasoning, showing that LLM agents can self-improve across complex decision trajectories. Together, these works represent a growing trend toward training LLM-based MAS that learn to communicate and adapt under reinforcement objectives (Wang et al. 2024; Zhang et al. 2024c; Sun, Huang, and Pompili 2024; Chen et al. 2025b).

Method

We formalize the multi-agent deliberation process as a decentralized partially observable Markov decision process (Dec-POMDP), defined by the tuple $\langle \mathcal{I}, \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \Omega, \mathcal{O}, \gamma \rangle$. Here, \mathcal{I} is the set of N agents, \mathcal{S} is the unobserved global state space representing the true problem-solving state, and $\mathcal{A} = \times_{i \in \mathcal{I}} \mathcal{A}^i$ is the joint action space composed of individual agent actions from \mathcal{A}^i .

At each deliberation round $t \in \{0, \dots, T - 1\}$, each agent $i \in \mathcal{I}$ receives a local observation $o_t^i \in \Omega$ and executes a deliberative action $a_t^i \in \mathcal{A}^i$ according to its policy $\pi_\theta^i(a_t^i | o_t^i)$. The environment transitions based on the function $\mathcal{T}(s_{t+1} | s_t, \mathbf{a}_t)$, where \mathbf{a}_t is the joint action. Subsequently, each agent receives a reward $r_t^i = \mathcal{R}^i(s_t, \mathbf{a}_t)$, shaped to reflect its marginal contribution to the team’s objective. The collective goal is to learn an optimal, decentralized policy $\pi^* = (\pi^{1*}, \dots, \pi^{N*})$ that maximizes the expected discounted return:

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^T \gamma^t \sum_{i \in \mathcal{I}} r_t^i \right] \quad (1)$$

The primary research challenge, and the core of our contribution, lies in the specific instantiation of the observation space Ω and the action space \mathcal{A} . Traditional approaches often rely on high-dimensional, unstructured text, leading to sample inefficiency and a focus on superficial linguistic patterns. In contrast, we propose a novel **Meta-Cognitive Deliberation Framework**, detailed below, which defines these spaces at a higher level of strategic abstraction.

The Meta-Cognitive Deliberation Framework

Our framework is designed to overcome the *Meta-Cognitive Blindspot* and *Static Deliberation Strategy* shortcomings of existing multi-LLM-agent frameworks by equipping each agent with the ability to reason about cognitive state and to select from high-level strategic acts.

The Meta-Cognitive State Space We define an agent’s local state z_t^i on a structured, low-dimensional vector summarizing its meta-cognitive status. This design choice forces the policy to learn abstract strategic reasoning rather than superficial text matching. The state z_t^i is composed of three key components, with full feature specifications and prompt templates detailed in Appendix B.

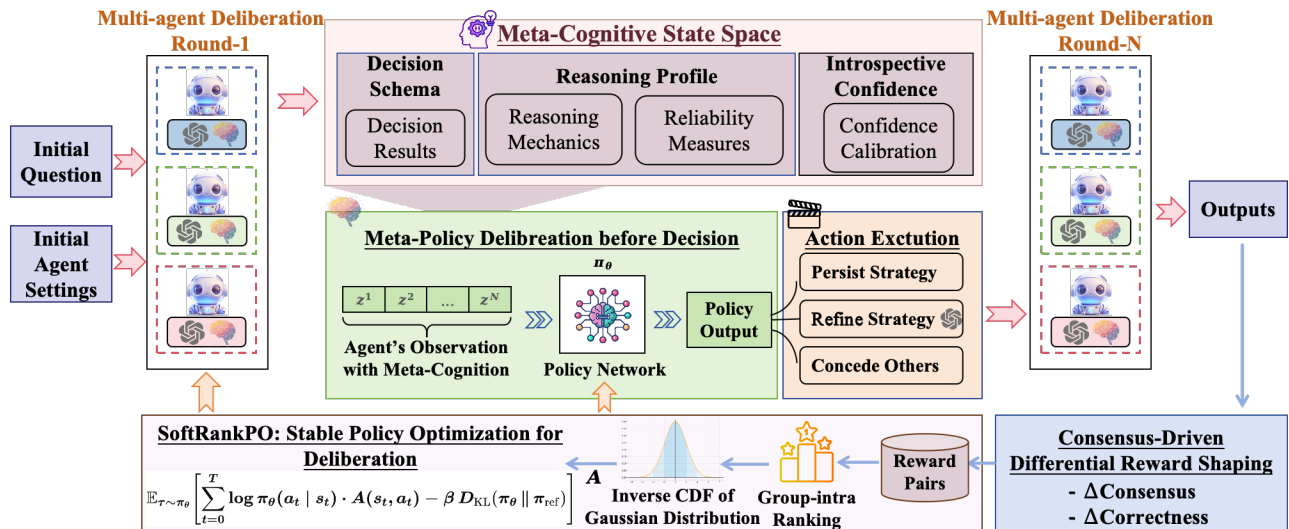


Figure 2: An overview of the Meta-Policy Deliberation Framework (MPDF). (a) Meta-Cognitive State Representation: Each agent constructs a cognitive state with decision schema, reasoning profile and introspective confidence. (b) Meta-Policy Deliberation: The agent’s meta-policy network selects one of three high-level actions: Persist (if confident), Refine (to improve its own solution), or Concede (to yield to a peer solution). (c) Environment Interaction: The chosen action leads to a new response, yielding a reward and transitioning the system to the next state. (d) SoftRankPO: a stable policy optimization RL algorithm.

Decision Schema (z_{ans}) represents the agent’s conclusive output, providing a concise summary of *what* the agent believes. It is extracted directly from the agent’s textual response by parsing the final output into a standardized format (e.g., a scalar value for mathematical problems or a one-hot vector for multiple-choice questions).

Reasoning Profile (z_{prof}) is a vector of self-reported features that summarizes *how* the agent reached its belief. These features are produced by the LLM agent in a structured format as part of its single generative step, capturing attributes like the number of reasoning steps, the logical operators employed (e.g., algebraic, arithmetic), and a self-assessed confidence score. This co-generation process yields meta-data without incurring additional inference cost.

Introspective Confidence (z_{conf}) is generated via a dedicated *introspection step*, where a critic model, specifically the same LLM backbone acting in an evaluation capacity, is prompted to assess the correctness of the agent’s full reasoning trace given the original problem. The critic’s judgment, which may take the form of a graded assessment or confidence-aware verdict (e.g., leaning toward “Correct” or “Incorrect” with varying strength), is then embedded into a vector representation, providing a signal for deliberation.

An agent’s full observation o_t^i for decision-making is composed of its own local meta-cognitive state and those of its peers: $o_t^i = (z_t^i, \{z_t^j\}_{j \neq i})$. The **policy network**, as a core component of our deliberation framework, then utilizes a cross-attention mechanism to integrate its introspective assessment with the social context provided by its peers. Please see Appendix B for the full policy network structure.

The Deliberative Action Space To address the issue of static strategies, we move beyond simple communicative

acts to a set of high-level, strategic deliberative acts, \mathcal{A}^i . These actions represent fundamental choices in a collaborative reasoning process, allowing the agent to learn a truly adaptive policy. The action space is defined as: **Persist**, deciding to maintain and defend one’s current solution against competing alternatives; **Refine**, an act of triggering an internal process of self-correction and iterative improvement; and **Concede**, a strategic deference to a peer’s solution based on an assessment of its higher quality.

By learning a policy $\pi_\theta^i(a_t^i | o_t^i)$ over this action space, our agents develop a dynamic, context-dependent deliberation strategy. They learn not just how to solve the problem, but how to **behave optimally as a member of a reasoning collective**—deciding when to lead, when to follow, and when to rethink, unlocking more efficient and robust reasoning.

SoftRankPO: A Stable Policy Optimization Algorithm for Deliberation

Learning a high-level deliberation policy π_θ within our Dec-POMDP setting is challenging because the reward signal is sparse and highly non-uniform across batches. The REFINE action in particular depends on stochastic LLM generations, leading to heavy-tailed returns and rounds where almost all agents make similar choices. In these cases the within-batch standard deviation can be extremely small, while the across-batch variance remains large. Moreover, the step-wise differential gains and leave-one-out consensus rewards can amplify this effect: high-consensus rounds often exhibit near-zero within-batch variance, whereas batches with more diverse or difficult candidates induce much larger variability. Standard policy gradients, including PPO and GRPO, construct advantages directly from raw reward magnitudes, so their gradient variance is sensitive to such scale and vari-

ance heterogeneity even after per-group normalization. In regimes with near-degenerate batches, per-group standard-deviation normalization can become numerically brittle.

In practice, the effective reward scale can differ substantially across tasks and batches, and even simple affine rescaling $R \mapsto \alpha R + \beta$ can significantly inflate or shrink gradient magnitudes. Since the gradient variance satisfies $\text{Var}[\nabla \mathcal{L}] \propto \text{Var}[R]$, learning can become unstable under batch-specific reward shifts. Although techniques such as baseline subtraction, batch normalization, or group-wise standardization help reduce variance, they do not remove the underlying dependence on reward scale, which remains a fragile foundation in collaborative systems where agents operate with heterogeneous criteria and subjective feedback.

SoftRank: Placing Order before Scale We introduce SoftRankPO, a KL-regularized policy-optimization method that aligns a trainable policy with *ranked* preferences. By divorcing learning signals from raw reward magnitudes, SoftRankPO reduces sensitivity to reward scale and variance while still operating inside a familiar KL trust region.

Consider a state s with reward vector $\mathbf{R} = (R_1, \dots, R_K)$ for the K candidate actions. We first convert the ordinal information into a standardized score through the mapping

$$p_i = \left(\frac{\text{rank}(R_i) + \frac{1}{2}}{K} \right)^\tau, \quad A_i = \frac{A_i^* - \bar{A}^*}{\text{std}(\mathbf{A}^*) + \varepsilon}, \quad (2)$$

where Φ^{-1} denotes the inverse standard normal CDF of $\mathcal{N}(0, 1)$, $A_i^* = \Phi^{-1}(p_i)$ and $\bar{A}^* = \frac{1}{K} \sum_j A_j^*$. The temperature parameter $\tau > 0$ controls the contrast between high- and low-ranked actions, enabling smoother or sharper advantage shaping. Equation (2) preserves the relative ranking of rewards, is invariant to affine transformations of \mathbf{R} , and produces approximately zero-mean, bounded-variance scores whenever K is reasonably large. In Appendix A we formalize these properties and show that the resulting advantages are sub-Gaussian, providing a stable and well-behaved signal for policy optimization.

KL-constrained Rank-matching Objective Let π_{ref} be a frozen reference policy obtained, for example, from supervised fine-tuning. Under SoftRankPO, we optimize the KL-regularized objective

$$\max_{\pi_\theta} \mathbb{E}_{s \sim \mathcal{D}} \left[\mathbf{A}(s)^\top \log \pi_\theta(\cdot | s) - \beta \text{KL}(\pi_\theta \| \pi_{\text{ref}}) + \lambda \mathcal{H}(\pi_\theta) \right], \quad (3)$$

where $\beta > 0$ controls the KL trust region and \mathcal{H} denotes Shannon entropy. Writing this in terms of the implicit KL-induced reward $R_\theta(s, a) = \beta \log \frac{\pi_\theta(a|s)}{\pi_{\text{ref}}(a|s)}$ and using the zero-mean property $\sum_i A_i(s) = 0$, we show in Appendix A that maximizing $J(\theta)$ is equivalent (up to constants) to minimizing the following **rank-matching loss**:

$$\mathcal{L}_{\text{SR}}(\theta) = \mathbb{E}_s \left[\frac{1}{K} \sum_{i=1}^K (R_\theta(s, a_i) - \bar{R}_\theta(s) - R(s, a_i) + \bar{R}(s))^2 \right], \quad (4)$$

where $\bar{R}(\cdot)$ and $\bar{R}_\theta(\cdot)$ denote within-state averages. Hence SoftRankPO learns logits whose *relative shape* mirrors that of the true ranked rewards, without requiring explicit partition functions. It can be interpreted either as a KL-regularized policy optimizer or as a rank-regression procedure that aligns implicit rewards with task rewards.

Optimization behaviour and theoretical properties. The SoftRank advantages are centered and have bounded variance, which implies that the partition-function term cancels in the gradient and the update admits a clean analytical form. In Appendix A we derive the gradient of $\mathcal{L}_{\text{SR}}(\theta)$, show how it decomposes into advantage-weighted learning, entropy-variance control, and KL trust-region coupling, and restate a standard result characterizing the KL-constrained optimum. We also provide a variance comparison with GRPO and a first-order convergence rate under common smoothness assumptions.

Consensus-Driven Differential Reward Shaping

To enable fine-grained credit assignment in decentralized deliberation, we design a reward framework grounded in the principle of marginal contribution. Instead of assigning scalar rewards directly based on outcome quality, we decompose the reward signal into two complementary components: local self-improvement and global influence on consensus.

The *step-wise delta reward* quantifies an agent’s immediate self-improvement, defined as the change in correctness of its own output before and after an action: $r_{\text{local}} = \mathbb{I}[\text{Correct}^{\text{after}}] - \mathbb{I}[\text{Correct}^{\text{before}}]$, encouraging agents to increase their individual answer quality. This step-wise shaping is conceptually related to prior MARL work on social deduction games, where agents are rewarded for improving others’ answers (Sarkar et al. 2025), but here we focus on self-improvement and combine it with a consensus-based marginal term.

The *lookahead consensus reward* captures an agent’s contribution to the final consensus, defined by its marginal impact on the outcome. Let $\text{Consensus}^{\text{full}}$ and Consensus^{-i} denote the correctness of the final group decision with and without agent i ’s input. The global reward is computed as $r_{\text{global}} = \mathbb{I}[\text{Consensus}^{\text{full}}] - \mathbb{I}[\text{Consensus}^{-i}]$, and in practice Consensus^{-i} is recomputed from the cached responses in the same round without requiring additional rollouts.

Together, the total reward $r = r_{\text{local}} + r_{\text{global}}$ provides structured, causally grounded supervision that aligns individual optimization with collective utility, enabling stable and targeted training of multi-agent deliberation policies.

Supervised Bootstrapping and MARL Fine-Tuning

Our training pipeline follows a two-stage process that integrates supervised guidance with variance-robust reinforcement learning, enabling stable convergence of the deliberation strategy. First, for robust initialization, we pre-train the policy network π_θ on an expert dataset $\mathcal{D}_{\text{sft}} = \{(s, a^*)\}$, where labels a^* are derived from oracle solutions, by minimizing a standard cross-entropy loss. This yields a supervised policy π_{sft} , which we also use as the fixed reference

policy in the KL regularization, i.e., $\pi_{\text{ref}} = \pi_{\text{sft}}$ in Eq. 3. We initialize π_{θ} from π_{sft} before reinforcement learning. Subsequently, we generate a static offline corpus \mathcal{D}_{off} by executing rollouts with π_{sft} . This corpus stores tuples of encountered meta-cognitive states and the resulting group-wise reward vectors, i.e., $\mathcal{D}_{\text{off}} = \{(s, \mathbf{R})\}$. In the second stage, we fine-tune π_{θ} on \mathcal{D}_{off} using our SoftRankPO objective (Eq. 3), aligning the policy with the stable, rank-based advantages derived from the offline data. The KL-divergence term ensures that the updated policy remains within a trust region of the sensible SFT initialization.

Experiment

We evaluate our framework on three standard reasoning tasks for LLM agents (Hendrycks et al. 2020; Cobbe et al. 2021; Hendrycks et al. 2021; Chen 2021; Du et al. 2023). (i) *Mathematical reasoning* is assessed on **GSM8K**, **MATH**, **AIME**, and **AMC**. (ii) *General reasoning* is measured with the **MMLU** benchmark. (iii) *Code generation* is evaluated on **HumanEval**. For all data sets, we follow the original splits and report the official metrics: *Solve Rate* for math datasets, *Accuracy* in MMLU and *Pass@1* in HumanEval.

Baselines. We compare MPDF against a comprehensive suite of baselines, including single-agent methods: **Vanilla**, **CoT** and **SC**; multi-agent sequential and debate frameworks: **PHP** and **LLM-Debate**; and dynamic workflow systems: **DyLan**, **AgentPrune**, **AFlow** and **GPTSwarm**.

Implementation Details. Following common practice, we set the number of agents and rounds to 3 by default. We instantiate agents with **Llama-3.1-8B-Instruct**, **Llama-3.2-3B-Instruct**, **Qwen2.5-7B-Instruct** and **Qwen2.5-3B-Instruct**. For all results, we report the average over 3 random seeds. Detailed hyperparameters and all settings are provided in Appendix C to ensure full reproducibility.

Main Results

Our Multi-Agent Deliberation Framework Outperforms Baselines. As shown in Table 1, our deliberation-based multi-agent framework achieves the highest average accuracy (55.74%) across six diverse reasoning benchmarks, outperforming both debate-style methods (e.g., DyLAN 51.50%) and dynamic approaches (e.g., AgentPrune 51.60%). Unlike single-agent methods (e.g., Vanilla, CoT, SC), which lack inter-agent coordination, and debate-style models, which suffer from unstable decision-making and weak convergence, our framework jointly optimizes decision-making and consensus formation to enhance multi-agent reasoning. While routing methods attempt to improve efficiency through static interaction graphs and heuristic specialization, they remain limited by their inability to support dynamic deliberation or mutual influence during execution. In contrast, our approach establishes a unified, meta-cognitive deliberation process that enables agents to collaborate iteratively toward shared goals, improving accuracy especially in high-complexity tasks such as AMC and AIME.

Consistent Gains Across Diverse LLM Backbones. To verify that our framework generalizes beyond a single model

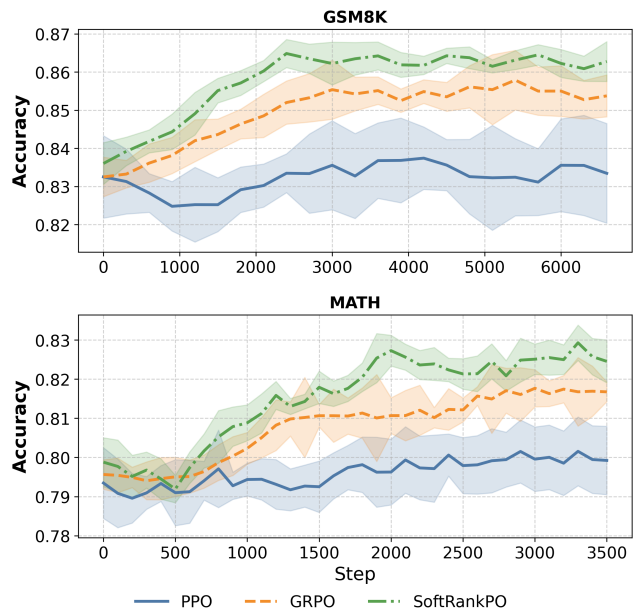


Figure 3: Accuracy curves during deliberative training on GSM8K (top) and MATH (bottom) using the LLaMA3-8B backbone. Our method consistently achieves stable and faster convergence across both benchmarks.

family, we conduct additional experiments using LLaMA3 and Qwen2.5 as representative backbones. Table 2 summarizes the results, demonstrating consistent and robust accuracy gains across all settings. On LMA-3B, the weakest backbone in our evaluation, accuracy improves from 68.84% (GPTSwarm) to 80.52%, indicating that the learned meta-policy effectively compensates for limited model capacity. Scaling up to LMA-8B still yields a positive margin, lifting the best baseline from 84.53% to 86.35%. Switching to Qwen models is particularly informative, as these checkpoints already benefit from extensive instruction tuning; Qwen-7B alone reaches 91.05%, yet deliberative training boosts it further to 94.47%; and Qwen-3B accuracy increases from 84.23% to 89.77%.

Lower Token Cost with Competitive Effectiveness. To facilitate comparison, we report results against the widely adopted LLM-Debate baseline. Our method achieves a substantial reduction in average token use (47.1k), measured as both input and output tokens, outperforming LLM-Debate (67.2k) while maintaining or surpassing its accuracy (Table 3). LLM-Debate operates by aggregating all agent responses and prompting the model to synthesize an answer in a single pass, but lacks a mechanism for adaptive filtering. As a result, it often generates redundant and unfocused interactions, inflating token cost without proportional gains in solution quality. In contrast, our method introduces meta-cognitive policies that guide agents to act reflectively. Through reinforcement learning over interaction trajectories, agents learn to retain high-confidence outputs and intervene only when revisions are likely to yield substan-

Model	MA	Type	GSM8K	MATH	AIME	AMC	MMLU	HumanEval	Avg
Vanilla	✗	Single	72.71	66.91	3.33	8.43	57.83	48.17	42.90
CoT	✗	Single	74.00	68.18	3.33	12.05	61.60	51.83	45.17
SC	✗	Single	80.21	73.09	4.44	12.05	68.64	56.10	49.09
PHP	✓	Sequential	80.36	74.36	4.44	15.67	68.72	56.70	50.04
LLM-Debate	✓	Debate	83.47	78.91	5.56	18.07	67.41	58.54	51.99
DyLAN	✓	Dynamic	82.63	76.55	6.67	19.28	66.84	60.37	51.50
GPTSwarm	✓	Dynamic	84.07	73.64	5.56	15.66	69.48	57.32	51.14
AgentPrune	✓	Dynamic	84.53	79.45	4.44	14.46	68.82	57.92	51.60
AFlow	✓	Dynamic	83.40	75.63	3.33	10.84	69.07	62.16	50.74
MPDF: SFT	✓	Deliberation	83.55	79.27	5.56	16.87	68.73	58.54	52.09
MPDF: SFT + PPO	✓	Deliberation	82.79	79.82	4.44	13.25	68.15	56.62	50.85
MPDF: SFT + GRPO	✓	Deliberation	85.52	81.45	7.78	20.48	70.22	60.98	54.40
MPDF: SFT + SoftRankPO	✓	Deliberation	86.35	82.72	7.78	22.89	70.87	63.82	55.74

Table 1: Comparison of performance across reasoning tasks. “MA” denotes the use of multiple agents (✓) versus single-agent baselines (✗). “Avg” reports the average accuracy across all applicable benchmarks. Our method consistently outperforms both single-agent and multi-agent baselines.

Model	LMA-8B	LMA-3B	Qwen-7B	Qwen-3B
Vanilla	72.71	38.64	90.90	83.62
CoT	74.00	50.49	91.05	84.23
SC	80.21	54.51	93.33	88.17
PHP	80.36	62.47	92.87	86.81
LLM-Debate	83.47	76.19	93.56	87.34
DyLAN	82.63	76.65	93.25	87.79
GPTSwarm	84.07	68.84	92.49	86.50
AgentPrune	84.53	65.35	92.12	86.35
AFlow	83.40	59.67	90.22	82.11
MPDF	86.35	80.52	94.47	89.77

Table 2: Performance on GSM8K across different LLM backbones. Results demonstrate consistent performance gains across model scales and pre-training settings.

tial gains. This behavioral shift is evident in the increased PERSIST action rate (78.8% vs. 19.1% in baseline SFT), as shown in Figure 7, a move toward confident retention.

Ablation and Analysis Studies

From Fine-Tuning to SoftRankPO: Multi-Agent RL for Deliberation. Our framework transitions from supervised fine-tuning (SFT) to progressively more effective reinforcement learning (RL) strategies that enhance agent coordination. As shown in Table 1, even the SFT-only variant performs competitively (52.09%), highlighting the strength of deliberation-based reasoning. However, it lacks outcome-driven feedback, limiting its ability to refine decision-making. Vanilla PPO performs worse (50.85%) due to unstable credit assignment and poor convergence in our multi-agent setting. As shown in Figure 3, PPO struggles especially on harder tasks (e.g., -3.62% on AMC), where its value-based updates fail to provide stable guidance. GRPO mitigates this by using group-relative ranking, stabilizing updates through peer comparison and improving performance on tasks like GSM8K (+1.97%) and MATH (+2.18%), but its reliance on raw reward magnitudes leads

Dataset	LLM-Debate	MPDF	Improvement (%)
GSM8K	8,608	6,154	28.53%
MATH	7,474	4,346	41.85%
AIME	19,321	14,077	27.15%
AMC	17,917	12,620	29.58%
MMLU	10,908	7,322	32.87%
HumanEval	3,054	2,616	14.34%
Total	67,282	47,135	29.96%

Table 3: Average total token count (input and output combined) per sample across multiple datasets.




Figure 4: Our framework scales well with more rounds and agents, with diminishing returns beyond optimal settings.

to sensitivity under noisy or heavy-tailed feedback.

Communication Rounds and Number of Agents. Figure 4 shows how communication rounds (M) and agent number (N) affect final performance. Increasing M from 1 to 4 steadily improves accuracy, validating the importance of multi-turn deliberation. However, performance slightly drops at 5 rounds, suggesting diminishing returns and possible noise accumulation, highlighting the need to calibrate deliberation depth. As N increases from 2 to 6, accuracy consistently improves. Unlike unstructured setups that suffer from conflict or stagnation, our reflective coordination continues to benefit from additional agents.




Figure 5: Stability under different exponential reward scales. SoftRankPO remains robust across multiple scales.

Robustness to Reward Scale. We conduct an ablation where the reward signal is rescaled by exponential factors $\{0.1, 1.0, 10.0\}$. As shown in Figure 5, GRPO suffers from noticeable instability: different scales induce diverging learning curves, and all variants exhibit a pronounced drop in performance after early gains. This behavior stems from GRPO’s reliance on perstate z -score advantages, where an outlier can dominate the local mean and variance, leading to unstable gradient updates. In contrast, SoftRankPO is robust across reward scales. Its accuracy consistently improves and stabilizes around the same level regardless of the reward scaling, with visibly smoother and more stable curves. This empirical observation aligns with our theoretical analysis. By inducing a zero-mean, bounded-variance advantage shaped by standard Gaussian, SoftRankPO effectively decouples learning dynamics from the scale of rewards.

Impact of Temperature Parameter τ in SoftRankPO. We conduct a controlled study varying the temperature coefficient $\tau \in \{0.4, 0.6, 0.8, 1.0, 2.0, 5.0\}$, which adjusts the contrast in rank-derived advantage signals. We observe stable convergence across all settings (Figure 6), confirming that SoftRankPO remains robust to a broad range of shaping intensities. Notably, setting $\tau = 0.8$ yields the best accuracy (approximately 88%), improving upon the default $\tau = 0.5$ (86.35%). Increasing τ sharpens the distribution of transformed ranks, thereby amplifying differences in advantage estimates and enhancing gradient informativeness. Theoretically, since $A_i^* = \Phi^{-1}(p_i^*)$, a larger τ spreads out the values of p_i^* , especially near the extremes, increasing the gradient magnitude of the CDF inverse and leading to stronger learning signals for high-ranked actions. However, excessively low values of τ compress the percentile spacing and may obscure meaningful preference distinctions. Thus, τ influences both learning stability and policy expressiveness by regulating the granularity of rank sensitivity.

Policy Shift: From Collaboration to Coordination. As shown in Figure 7, SoftRankPO induces a significant behavioral shift: the PERSIST rate increases from 19.1% to 78.8%, while REFINE and CONCEDE jointly fall from 81% to 21%. This stark transition suggests that agents have learned to recognize and trust the high-quality outputs produced by the pretrained LLM. It reflects an emerging ability for self-assessment and confidence calibration. The resulting policy aligns with the principle of minimal intervention:

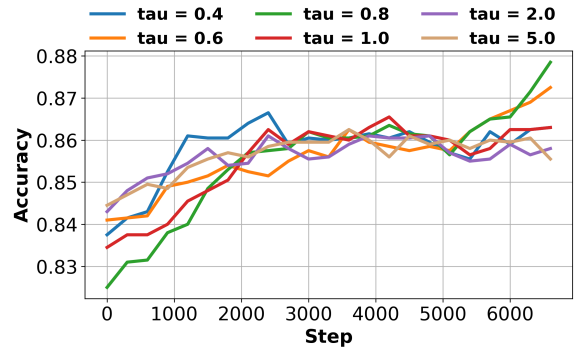


Figure 6: Effect of parameter τ on SoftRankPO performance. The method remains stable across a broad range of values, while finer tuning yields improved accuracy.




Figure 7: Policy shift toward coordination: agents increasingly favor PERSIST, reflecting stronger confidence and selective intervention driven by reward shaping.

instead of revising reflexively or adopting peer responses blindly, agents now act selectively, intervening only when they perceive a meaningful improvement. Such a pattern is directly encouraged by our reward shaping, which emphasizes marginal utility and penalizes redundant edits. Beyond accuracy, this leads to reduced token use (Table 3).

Conclusion

In this paper, we present MPDF, a meta-policy deliberation framework for multi-agent LLM collaboration, enabling agents to strategically adjust their behavior based on cognitive states. Coupled with the scale-resilient SoftRankPO algorithm, our approach achieves robust policy learning and fosters efficient coordination. Extensive results validate the effectiveness of our approach. We observe a policy shift during MPDF learning from collaboration to coordination, marking a transition from exploratory language cooperation under SFT to coordinated decision-making under RL.

Acknowledgments

This work was supported in part by a grant from the USC–Capital One Center for Responsible AI and Decision Making in Finance (CREDIF).

References

- Chang, C.; Shi, Y.; Cao, D.; Yang, W.; Hwang, J.; Wang, H.; Pang, J.; Wang, W.; Liu, Y.; Peng, W.-C.; et al. 2025. A survey of reasoning and agentic systems in time series with large language models. *arXiv preprint arXiv:2509.11575*.
- Chen, G.; Yang, S.; Li, C.; Liu, W.; Luan, J.; and Xu, Z. 2025a. Heterogeneous Group-Based Reinforcement Learning for LLM-based Multi-Agent Systems. *arXiv preprint arXiv:2506.02718*.
- Chen, H.; Yang, T.; Gao, S.; Chen, R.; Quan, X.; Tian, H.; and Yao, T. 2025b. Discriminative Policy Optimization for Token-Level Reward Models. *arXiv preprint arXiv:2505.23363*.
- Chen, M. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Chen, P.; Han, B.; and Zhang, S. 2024. Comm: Collaborative multi-agent, multi-reasoning-path prompting for complex problem solving. *arXiv preprint arXiv:2404.17729*.
- Chen, Y.; Liu, Q.; Zhang, Y.; Sun, W.; Ma, X.; Yang, W.; Shi, D.; Mao, J.; and Yin, D. 2025c. Tourrank: Utilizing large language models for documents ranking with a tournament-inspired strategy. In *Proceedings of the ACM on Web Conference 2025*, 1638–1652.
- Chen, Y.; Mao, H.; Mao, J.; Wu, S.; Zhang, T.; Zhang, B.; Yang, W.; and Chang, H. 2022a. PTDE: Personalized training with distilled execution for multi-agent reinforcement learning. *arXiv preprint arXiv:2210.08872*.
- Chen, Y.; Yang, W.; Zhang, T.; Wu, S.; and Chang, H. 2022b. Commander-soldiers reinforcement learning for cooperative multi-agent systems. In *2022 International Joint Conference on Neural Networks (IJCNN)*, 1–7. IEEE.
- Chen, Y.; Zhang, E.; Yan, L.; Wang, S.; Huang, J.; Yin, D.; and Mao, J. 2025d. Mao-arag: Multi-agent orchestration for adaptive retrieval-augmented generation. *arXiv preprint arXiv:2508.01005*.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Dai, X.; Xie, Y.; Liu, M.; Wang, X.; Li, Z.; Wang, H.; and Lui, J. 2025. Multi-agent conversational online learning for adaptive llm response identification. *arXiv preprint arXiv:2501.01849*.
- Du, H.; Thudumu, S.; Vasa, R.; and Mouzakakis, K. 2024. A survey on context-aware multi-agent systems: techniques, challenges and future directions. *arXiv preprint arXiv:2402.01968*.
- Du, Y.; Li, S.; Torralba, A.; Tenenbaum, J. B.; and Mordatch, I. 2023. Improving factuality and reasoning in language models through multiagent debate. In *Forty-first International Conference on Machine Learning*.
- Han, A.; Hu, J.; Wei, P.; Zhang, Z.; Guo, Y.; Lu, J.; and Zhang, Z. 2025. JoyAgents-R1: Joint Evolution Dynamics for Versatile Multi-LLM Agents with Reinforcement Learning. *arXiv preprint arXiv:2506.19846*.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Hendrycks, D.; Burns, C.; Kadavath, S.; Arora, A.; Basart, S.; Tang, E.; Song, D.; and Steinhardt, J. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Hong, S.; Zhuge, M.; Chen, J.; Zheng, X.; Cheng, Y.; Wang, J.; Zhang, C.; Wang, Z.; Yau, S. K. S.; Lin, Z.; et al. 2023. MetaGPT: Meta programming for a multi-agent collaborative framework. In *The Twelfth International Conference on Learning Representations*.
- Hu, K.; Li, M.; Song, Z.; Xu, K.; Xia, Q.; Sun, N.; Zhou, P.; and Xia, M. 2024a. A review of research on reinforcement learning algorithms for multi-agents. *Neurocomputing*, 599: 128068.
- Hu, Q. J.; Bieker, J.; Li, X.; Jiang, N.; Keigwin, B.; Ranganath, G.; Keutzer, K.; and Upadhyay, S. K. 2024b. Routerbench: A benchmark for multi-llm routing system. *arXiv preprint arXiv:2403.12031*.
- Hu, S.; Lu, C.; and Clune, J. 2024. Automated design of agentic systems. *arXiv preprint arXiv:2408.08435*.
- Ishibashi, Y.; and Nishimura, Y. 2024. Self-organized agents: A llm multi-agent framework toward ultra large-scale code generation and optimization. *arXiv preprint arXiv:2404.02183*.
- Koriat, A. 2015. Metacognition: Decision making Processes in Self-monitoring and Self-regulation. *The Wiley Blackwell handbook of judgment and decision making*, 2: 356–379.
- Li, D.; Dong, H.; Wang, L.; Qiao, B.; Qin, S.; Lin, Q.; Zhang, D.; Zhang, Q.; Xu, Z.; Zhang, B.; et al. 2024. Verco: Learning coordinated verbal communication for multi-agent reinforcement learning. *arXiv preprint arXiv:2404.17780*.
- Li, S.; Yang, W.; Zhang, P.; Xiao, X.; Cao, D.; Qin, Y.; Zhang, X.; Zhao, Y.; and Bogdan, P. 2025. Climatellm: Efficient weather forecasting via frequency-aware large language models. *arXiv preprint arXiv:2502.11059*.
- Liao, J.; Wen, M.; Wang, J.; and Zhang, W. 2025. Marft: Multi-agent reinforcement fine-tuning. *arXiv preprint arXiv:2504.16129*.
- Liu, T.; Wang, X.; Huang, W.; Xu, W.; Zeng, Y.; Jiang, L.; Yang, H.; and Li, J. 2024. Groupdebate: Enhancing the efficiency of multi-agent debate using group discussion. *arXiv preprint arXiv:2409.14051*.
- Liu, Z.; Chai, J.; Zhu, X.; Tang, S.; Ye, R.; Zhang, B.; Bai, L.; and Chen, S. 2025. Ml-agent: Reinforcing llm agents for autonomous machine learning engineering. *arXiv preprint arXiv:2505.23723*.
- Liu, Z.; Zhang, Y.; Li, P.; Liu, Y.; and Yang, D. 2023. Dynamic llm-agent network: An llm-agent collaboration framework with agent team optimization. *arXiv preprint arXiv:2310.02170*.
- Motwani, S. R.; Smith, C.; Das, R. J.; Rafailov, R.; Laptev, I.; Torr, P. H.; Pizzati, F.; Clark, R.; and de Witt, C. S. 2024. Malt: Improving reasoning with multi-agent llm training. *arXiv preprint arXiv:2412.01928*.

- Park, C.; Han, S.; Guo, X.; Ozdaglar, A.; Zhang, K.; and Kim, J.-K. 2025. Maporl: Multi-agent post-co-training for collaborative large language models with reinforcement learning. *arXiv preprint arXiv:2502.18439*.
- Park, J. S.; O'Brien, J.; Cai, C. J.; Morris, M. R.; Liang, P.; and Bernstein, M. S. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, 1–22.
- Peiyuan, F.; He, Y.; Huang, G.; Lin, Y.; Zhang, H.; Zhang, Y.; and Li, H. 2024. Agile: A novel reinforcement learning framework of llm agents. *Advances in Neural Information Processing Systems*, 37: 5244–5284.
- Ping, H.; Bhattacharjee, A.; Zhang, P.; Li, S.; Yang, W.; Cheng, A.; Zhang, X.; Thomason, J.; Jannesari, A.; Ahmed, N.; et al. 2025a. VeriMoA: A Mixture-of-Agents Framework for Spec-to-HDL Generation. *arXiv preprint arXiv:2510.27617*.
- Ping, H.; Li, S.; Zhang, P.; Cheng, A.; Duan, S.; Kanakaris, N.; Xiao, X.; Yang, W.; Nazarian, S.; Irimia, A.; et al. 2025b. Hdlcore: A training-free framework for mitigating hallucinations in llm-generated hdl. *arXiv preprint arXiv:2503.16528*.
- Qian, C.; Xie, Z.; Wang, Y.; Liu, W.; Zhu, K.; Xia, H.; Dang, Y.; Du, Z.; Chen, W.; Yang, C.; et al. 2024. Scaling large language model-based multi-agent collaboration. *arXiv preprint arXiv:2406.07155*.
- Sarkar, B.; Xia, W.; Liu, C. K.; and Sadigh, D. 2025. Training language models for social deduction with multi-agent reinforcement learning. *arXiv preprint arXiv:2502.06060*.
- Shang, Y.; Li, Y.; Zhao, K.; Ma, L.; Liu, J.; Xu, F.; and Li, Y. 2024. Agentsquare: Automatic llm agent search in modular design space. *arXiv preprint arXiv:2410.06153*.
- Sun, C.; Huang, S.; and Pompili, D. 2024. Llm-based multi-agent reinforcement learning: Current and future directions. *arXiv preprint arXiv:2405.11106*.
- Sun, R.; Zhang, X.; and Mathews, R. 2006. Modeling metacognition in a cognitive architecture. *Cognitive Systems Research*, 7(4): 327–338.
- Trirat, P.; Jeong, W.; and Hwang, S. J. 2024. Automl-agent: A multi-agent llm framework for full-pipeline automl. *arXiv preprint arXiv:2410.02958*.
- Wan, Z.; Li, Y.; Wen, X.; Song, Y.; Wang, H.; Yang, L.; Schmidt, M.; Wang, J.; Zhang, W.; Hu, S.; et al. 2025. Rema: Learning to meta-think for llms with multi-agent reinforcement learning. *arXiv preprint arXiv:2503.09501*.
- Wang, X.; Zhi, G.; Tang, Z.; Jin, H.; Zhang, Q.; and Li, N. 2024. Self-Aware Intelligent Medical Rescue Unmanned Team via Large Language Model and Multi-Agent Reinforcement Learning. In *Proceedings of the 2024 International Symposium on AI and Cybersecurity*, 119–124.
- Wang, Z.; Wang, K.; Wang, Q.; Zhang, P.; Li, L.; Yang, Z.; Jin, X.; Yu, K.; Nguyen, M. N.; Liu, L.; et al. 2025. Ragen: Understanding self-evolution in llm agents via multi-turn reinforcement learning. *arXiv preprint arXiv:2504.20073*.
- Wei, Y.; Shan, X.; and Li, J. 2025. LERO: LLM-driven Evolutionary framework with Hybrid Rewards and Enhanced Observation for Multi-Agent Reinforcement Learning. *arXiv preprint arXiv:2503.21807*.
- Yan, B.; Zhou, Z.; Zhang, L.; Zhang, L.; Zhou, Z.; Miao, D.; Li, Z.; Li, C.; and Zhang, X. 2025. Beyond self-talk: A communication-centric survey of llm-based multi-agent systems. *arXiv preprint arXiv:2502.14321*.
- Yang, W.; Pang, J.; Li, S.; Bogdan, P.; Tu, S.; and Thomason, J. 2025. Maestro: Learning to Collaborate via Conditional Listwise Policy Optimization for Multi-Agent LLMs. *arXiv preprint arXiv:2511.06134*.
- Yeung, N.; and Summerfield, C. 2012. Metacognition in human decision-making: confidence and error monitoring. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1594): 1310–1321.
- Yue, Y.; Zhang, G.; Liu, B.; Wan, G.; Wang, K.; Cheng, D.; and Qi, Y. 2025. Masrouter: Learning to route llms for multi-agent systems. *arXiv preprint arXiv:2502.11133*.
- Zhang, G.; Chen, K.; Wan, G.; Chang, H.; Cheng, H.; Wang, K.; Hu, S.; and Bai, L. 2025. Evoflow: Evolving diverse agentic workflows on the fly. *arXiv preprint arXiv:2502.07373*.
- Zhang, G.; Yue, Y.; Li, Z.; Yun, S.; Wan, G.; Wang, K.; Cheng, D.; Yu, J. X.; and Chen, T. 2024a. Cut the crap: An economical communication pipeline for llm-based multi-agent systems. *arXiv preprint arXiv:2410.02506*.
- Zhang, J.; Xiang, J.; Yu, Z.; Teng, F.; Chen, X.; Chen, J.; Zhuge, M.; Cheng, X.; Hong, S.; Wang, J.; et al. 2024b. Aflow: Automating agentic workflow generation. *arXiv preprint arXiv:2410.10762*.
- Zhang, N.; Wang, X.; Cui, Q.; Zhou, R.; Kakade, S. M.; and Du, S. S. 2024c. Multi-agent reinforcement learning from human feedback: Data coverage and algorithmic techniques.
- Zhao, Q.; Wang, J.; Zhang, Y.; Jin, Y.; Zhu, K.; Chen, H.; and Xie, X. 2023. Competeai: Understanding the competition dynamics in large language model-based agents. *arXiv preprint arXiv:2310.17512*.
- Zheng, C.; Liu, Z.; Xie, E.; Li, Z.; and Li, Y. 2023. Progressive-hint prompting improves reasoning in large language models. *arXiv preprint arXiv:2304.09797*.
- Zhou, H.; Geng, H.; Xue, X.; Kang, L.; Qin, Y.; Wang, Z.; Yin, Z.; and Bai, L. 2025. Reso: A reward-driven self-organizing llm-based multi-agent system for reasoning tasks. *arXiv preprint arXiv:2503.02390*.
- Zhu, G.; Zhou, R.; Ji, W.; and Zhao, S. 2025. LAMARL: LLM-Aided Multi-Agent Reinforcement Learning for Cooperative Policy Generation. *IEEE Robotics and Automation Letters*.
- Zhuang, Y.; Shen, Y.; Zhang, Z.; Chen, Y.; and Miao, F. 2024. Yolo-marl: You only llm once for multi-agent reinforcement learning. *arXiv preprint arXiv:2410.03997*.
- Zhuge, M.; Wang, W.; Kirsch, L.; Faccio, F.; Khizbullin, D.; and Schmidhuber, J. 2024. Gptswarm: Language agents as optimizable graphs. In *Forty-first International Conference on Machine Learning*.