

LLM Collaboration with Multi-Agent Reinforcement Learning

Shuo Liu, Tianle Chen*, Zeyu Liang*, Xueguang Lyu, Christopher Amato†

Khoury College of Computer Sciences,

Northeastern University

Boston, MA, 02115, USA

{liu.shuo2, chen.tianle, liang.zey, lu.xue, c.amato}@northeastern.edu

Abstract

A large amount of work has been done in Multi-Agent Systems (MAS) for modeling and solving problems with multiple interacting agents. However, most LLMs are pretrained independently and not specifically optimized for coordination. For example, existing LLM fine-tuning frameworks rely on individual rewards, which require complex reward designs for each agent to encourage collaboration. To address this challenge, we model LLM collaboration as a cooperative Multi-Agent Reinforcement Learning (MARL) problem. We develop a multi-agent, multi-turn algorithm, Multi-Agent Group Relative Policy Optimization (MAGRPO), to solve it, building on current RL approaches for LLMs as well as MARL techniques. Our experiments on LLM writing and coding collaboration demonstrate that fine-tuning multiple LLMs with MAGRPO enables agents to generate high-quality responses efficiently through effective cooperation. Our approach opens the door to using MARL methods for LLM collaboration and highlights the associated challenges. Our code is available at <https://github.com/OpenMLRL/CoMLRL>.

Introduction

State-of-the-art LLMs have demonstrated remarkable capabilities across diverse domains (Grattafiori et al. 2024; Achiam et al. 2023; Anil et al. 2025). To adapt to specific applications or align with human preferences, fine-tuning has emerged as a critical training stage. Compared to supervised fine-tuning, Reinforcement Learning (RL) enables more generalizable learning for complex, multi-turn tasks through human-aligned reward design, making it an important technique for fine-tuning (Ouyang et al. 2022; Guo et al. 2025; Ziegler et al. 2019).

Likewise, Multi-Agent Systems (MAS) have been extensively studied over the past decades, with substantial progress in modeling and solving problems involving multiple agents (Weiss 1999; Van der Hoek and Wooldridge 2008; Shoham and Leyton-Brown 2009; Stone and Veloso 2000). In particular, advances in cooperative MAS have demonstrated strong potential for enabling effective collaboration in distributed settings, such as games, robotics, and traffic control (Samvelyan et al. 2019; Vinyals et al. 2017; Berner et al. 2019; Wiering 2000; Liu et al. 2022; Xiao et al. 2025).

* These authors contribute equally.

† Corresponding author: c.amato@northeastern.edu.

These developments motivate the application of MAS principles and techniques to LLM collaboration, where multiple LLMs working together can solve more complex tasks more robustly and efficiently.

There has been some recent work on coordinating multiple LLMs. Some approaches implement coordination at the inference stage, enabling agents to interact through debate, discussion, or verification (Du et al. 2023; Wu et al. 2023a; Lifshitz, McIlraith, and Du 2025). These methods operate at the prompt level, with fixed models that are not tuned toward coordination-centric objectives. The agents may have conflicting answers or spread incorrect information to other participants, limiting performance (Cemri et al. 2025; Estornell and Liu 2024). Moreover, the design of effective prompts remains difficult and unclear. Other approaches fine-tune agents independently with individual or role-conditioned rewards. However, they require carefully curated rewards for each individual or role (Liu et al. 2025a; Subramaniam et al. 2025), and, as independent learning methods, lack convergence guarantees (Tan 1993).

We model LLM collaboration as a cooperative MARL problem (Albrecht, Christianos, and Schäfer 2024; Amato 2025) and formalize it as a Decentralized Partially Observable Markov Decision Process (Dec-POMDP) (Oliehoek and Amato 2016). In LLM collaboration, multiple trainable LLMs generate responses synchronously based on their individual prompts. The external environment evolves according to the joint responses until the dialog ends. This general model allows a wide range of problems to be modeled and solved using versions of MARL algorithms. Following the efficient practice of Group Relative Policy Optimization (GRPO) (Shao et al. 2024), we propose Multi-Agent GRPO (MAGRPO) that trains LLMs in a multi-turn setting. MAGRPO leverages centralized group-relative advantages for joint optimization, while preserving decentralized execution for each agent. The resulting method builds off of state-of-the-art LLM approaches in GRPO and MARL approaches for centralized training and decentralized execution, such as MAPPO (Yu et al. 2022). Our experiments demonstrate that MAGRPO is able to learn different cooperation schemes while producing efficient and high-quality responses.

Our contributions can be summarized as follows: (i) We model the LLM collaboration as a cooperative MARL problem, where multiple LLMs cooperate to generate joint re-

sponses; (ii) We develop the MAGRPO algorithm, which optimizes agent cooperation through aligned rewards while maintaining decentralized execution to maintain efficiency; (iii) Our experiments demonstrate that fine-tuning with MAGRPO improves both response efficiency and quality in writing and coding collaboration; (iv) We provide an analysis of the limitations of existing approaches and outline open challenges in applying MARL to LLM collaboration.

Related Work

Test-Time Multi-Agent Interaction Recent work employs multiple agents with specialized roles interacting through diverse pipelines at test-time to enhance response quality. In multi-agent debate, agents iteratively formulate positions by reviewing other agents' outputs, where the final decision or answer is determined by majority voting or a summarizer (Du et al. 2023; Chan et al. 2023; Liang et al. 2024). Role-based approaches allocate tasks across specialized agents (Wu et al. 2023a; Qian et al. 2024; Hong et al. 2023). An agent may function as a verifier to assess the correctness of outputs (Skreta et al. 2023; Lifshitz, McIlraith, and Du 2025), or as a macro-planner to orchestrate workers' responses. However, these multi-agent frameworks rely on prompt-level interactions among agents, often leading to ineffective communication and computational inefficiency. Moreover, the design of effective prompts and role assignment remains unclear, as prompts usually fail to reliably guide agent behavior, enforce role adherence, or support coherent coordination across tasks. These limitations motivate us to fine-tune LLMs in MAS to improve their cooperation.

Multi-Agent Fine-Tuning Recent work has explored fine-tuning LLMs to improve their performance across diverse domains, e.g., arithmetic reasoning, navigation, and hidden-role games (Ma et al. 2024; Sarkar et al. 2025). These approaches typically employ individual rewards or rewards conditioned on specific roles (Liu et al. 2025a; Subramaniam et al. 2025; Zhang et al. 2025). Such reward structures often require careful manual specification, and their underlying rationale is rarely well justified. The misaligned or conflicting incentives can hinder effective coordination. Moreover, these methods lack convergence guarantees, as each agent learns independently in a non-stationary environment where other agents are simultaneously updating their policies. In this paper, we focus on cooperative scenarios, where LLMs are trained with verified, human-aligned rewards.

Cooperative MARL for LLM Collaboration

Since LLMs can act as a special class of agents, we leverage advances in MAS to improve their collaboration. We model LLM collaboration as a cooperative MARL problem and outline its unique challenges. We formalize this problem as a Dec-POMDP, as shown in Figure 1.

LLM Collaboration

LLM collaboration is the problem where LLMs cooperatively solve a class of tasks in MAS. Tasks are specified in language and provided to the LLMs as prompts. Each LLM

generates a response synchronously in response to its own instructions. All responses jointly form a solution to the task.

Most tasks cannot be resolved in one turn. Users, external models, or systems validate the solutions and provide additional requirements or suggestions for LLMs. These components also serve as part of the environment for LLM collaboration, whose states may change based on the agents' outputs. The updates are embedded into prompts for subsequent turns. This iterative process continues until the task is successfully completed or a predefined turn limit is reached.

As discussed by a number of companies (NVIDIA 2024; Anthropic 2024), a team of agents could generate a complex codebase. The code would be difficult, costly, and time-consuming to generate with a single agent, but a group of LLMs could do so quickly and cheaply. None of these agents is self-interested, but they can be trained using a scheme such as the one in this paper. Using a joint reward allows agents to specialize as needed to complete the task without complex prompt or reward engineering.

Problem Formalization

We formalize LLM collaboration as a subclass of the cooperative MARL problem, considering LLMs are agents and the types of problems they are solving. This problem is a form of Dec-POMDP (Oliehoek and Amato 2016), which allows cooperation through a joint reward while preserving scalable decentralized control. We show 2 instantiations of our framework in writing and coding tasks in the experiments.

Mathematically, our LLM Dec-POMDP is defined by a tuple $\langle \mathcal{I}, \mathcal{S}, \{\mathcal{O}_i\}, \{\mathcal{A}_i\}, R, T, H \rangle$.

- $\mathcal{I} = \{1, \dots, n\}$ denotes the set of n LLM agents, each instantiated with a pre-trained language model.
- \mathcal{S} denotes the full global state space. At turn t , a full state $s_t = (s_t^{\text{acc}}, s_t^{\text{usr}})$ consists of parts that are accessible in the model and provided to the reward model $s_t^{\text{acc}} \in \mathcal{S}^{\text{acc}}$ (e.g., external models or systems), and the inaccessible user state $s_t^{\text{usr}} \in \mathcal{S}^{\text{usr}}$ that updates over time but is not maintainable. In a Dec-POMDP, the state can not be directly observed by the agents.
- \mathcal{O}_i is the observation space for agent i with $\mathcal{O} = \times_i \mathcal{O}_i$ the joint observation space. A local observation $o_{i,t}$ consists of natural language instructions (i.e., prompts), providing a partial and noisy view of s_t .
- \mathcal{A}_i is the action space for agent i with $\mathcal{A} = \times_i \mathcal{A}_i$ the joint action space. A local action $a_{i,t}$ is a response in natural language to the given prompt.
- $R : \mathcal{S}^{\text{acc}} \times \mathcal{A} \rightarrow \mathbb{R}$ is the joint reward function implemented via predefined rules or a pretrained reward model. At turn t , the joint rewards r_t are determined by the accessible part of current state s_t^{acc} and the agents' joint action $\mathbf{a}_t = \{a_{1,t}, \dots, a_{n,t}\}$.
- $T : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the underlying stochastic state transition function. At turn t , the agents' joint actions \mathbf{a}_t induce a shift to a new state $s_{t+1} \sim T(\cdot | s_t, \mathbf{a}_t)$, which reflects the updates in the user state and the states of external models and systems.
- H is the episode horizon, i.e., the turn limit of the dialog.

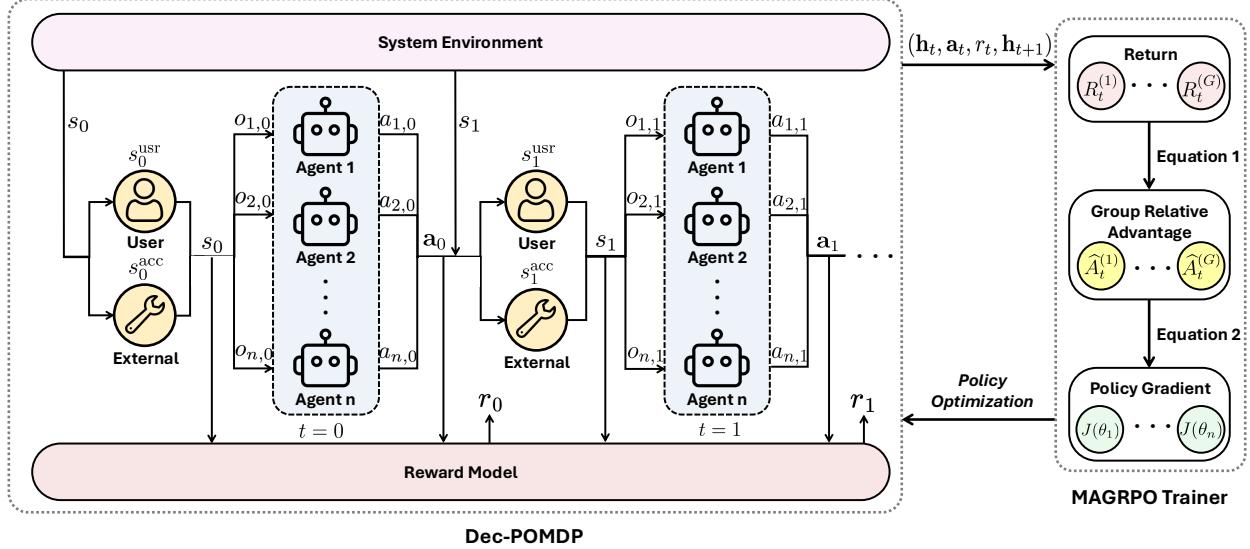


Figure 1: Illustration of Dec-POMDP and our MAGRPO algorithm.

In Dec-POMDP, since the states are not directly observed, each agent maintains its local observation-action history $\mathbf{h} = \{h_1, \dots, h_n\}$ to infer information about the state. A solution to a Dec-POMDP is a joint policy that maximizes the expected cumulative reward, $\pi^* = \{\pi_1^*, \dots, \pi_n^*\} = \arg \max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{H-1} R(s_t^{\text{acc}}, \mathbf{a}_t) \right]$. A joint policy is a set of local policies π_i , which condition on the local observation-action history $h_{i,t} = \{o_{i,0}, a_{i,0}, \dots, o_{i,t}\}$.

RL methods for Dec-POMDPs have become a popular topic (e.g., (Foerster et al. 2018b; Lowe et al. 2017; Foerster et al. 2018a; Rashid et al. 2020; Wang et al. 2021; Yu et al. 2022; Albrecht, Christianos, and Schäfer 2024; Lyu et al. 2023; Marchesini et al. 2025)) with methods successful at scaling to large state, action, and observation spaces. Many methods use Centralized Training for Decentralized Execution (CTDE), where they use some centralized information during training (e.g., a centralized value function estimate) but are still able to execute in a decentralized manner when training is completed (Amato 2024).

Challenges in LLM Collaboration

LLM collaboration presents unique challenges compared to traditional MARL problems, where LLM agents receive and process tasks through natural language.

Representations in Natural Language Unlike traditional cooperative MARL agents, LLM agents operate over natural language, receiving instructions and generating responses as sequences of tokens. MARL approaches could model this problem at the token or sequence level. At the token level, the number of actions and observations is smaller, but the problem horizon can be very long. At the sequence level, the action and observation space is much larger, but the horizon is much shorter. Moreover, token-level rewards are often uninformative, as both queries and responses must form coherent and semantically meaningful structures. As adopted in

prior RL methods (Ouyang et al. 2022; Rafailov et al. 2023), we model each agent’s decision-making process as a direct mapping from input instructions to complete responses to enable efficient and stable training. Nevertheless, the best modeling and solution approaches remain an open question.

Training Paradigm As mentioned above, many MARL methods use centralized training for decentralized execution (CTDE). Unfortunately, standard CTDE methods use centralized value models in the form of centralized critics (Foerster et al. 2018b; Lowe et al. 2017; Yu et al. 2022) or mixers in value decomposition methods (Rashid et al. 2020; Wang et al. 2021), which allow additional information during training but do not scale well to very large action and observation spaces (such as those in LLM collaboration). Conversely, Decentralized Training and Execution (DTE) methods (Amato 2025) train a set of models, one for each agent in a decentralized manner, which are typically more scalable but do not use additional information during training (even when it is available). It is an open question which paradigm to use to maximize performance while maintaining scalability in the LLM collaboration problem. In this paper, we balance decentralized execution with centralized training using group-based Monte Carlo estimates. Experiments show the effectiveness of our approach on short-horizon tasks.

MAGRPO

We propose the Multi-Agent GRPO (MAGRPO) algorithm to jointly train LLM agents in MAS while maintaining decentralized execution.

Algorithm 1 shows the procedure of MAGRPO. Given a dataset \mathcal{D} containing task information (e.g., the descriptions of coding problems), n LLMs that are optimized, each with a policy parameterized by θ_i and guided by a (shared) reward model R . In each episode, a task is sampled from the given dataset \mathcal{D} , which is used to construct initial observations

Algorithm 1: MAGRPO

Require: Dataset \mathcal{D} , n pretrained LLMs with policies $\{\pi_{\theta_1}, \dots, \pi_{\theta_n}\}$, reward model R , generation group size G , learning rate α

- 1: **for** each episode **do**
- 2: Sample a task $\sim \mathcal{D}$
- 3: Initialize observations $o_{i,0}, \forall i \in \mathcal{I}$, according to the task, and $\mathbf{o}_0 = \{o_{1,0}, \dots, o_{n,0}\}$
- 4: $h_{i,0}^G \leftarrow o_{i,0}, \forall i \in \mathcal{I}$, and $\mathbf{h}_0^G = \{h_{1,0}^G, \dots, h_{n,0}^G\}$
- 5: **for** turn $t = 0$ to $H - 1$ **do**
- 6: Generate a group of responses $a_{i,t}^G \leftarrow \pi_{\theta_i}(\cdot | h_{i,t}^G)$, $\forall i \in \mathcal{I}$, where $h_{i,t}^G = \{h_{i,t}^{(1)}, \dots, h_{i,t}^{(G)}\}$, $a_{i,t}^G = \{a_{i,t}^{(1)}, \dots, a_{i,t}^{(G)}\}$, and $\mathbf{a}_t^G = \{a_{1,t}^G, \dots, a_{n,t}^G\}$
- 7: Obtain joint rewards r_t^G from system
- 8: Receive new observations $o_{i,t+1}^G$, and update history $h_{i,t+1}^G \leftarrow \{h_{i,t}^G, a_{i,t}^G, o_{i,t+1}^G\}, \forall i \in \mathcal{I}$
- 9: **end for**
- 10: **for** turn $t = H - 1$ to 0 **do**
- 11: Calculate return $R_t^{(g)} \leftarrow \sum_{\tau=t}^{H-1} r_{\tau}^{(g)}, \forall g \in \mathcal{G}$
- 12: Estimate $\hat{A}_t^{(g)}$, $\forall g \in \mathcal{G}$ according to Equation 1
- 13: Calculate $J(\theta_i)$, $\forall i \in \mathcal{I}$ according to Equation 2
- 14: $\theta_i \leftarrow \theta_i + \alpha \nabla_{\theta_i} J(\theta_i), \forall i \in \mathcal{I}$
- 15: **end for**
- 16: **end for**
- 17: **return** $\pi_{\theta} = \{\pi_{\theta_1}, \dots, \pi_{\theta_n}\}$

$\mathbf{o}_0 = \{o_{1,0}, \dots, o_{n,0}\}$ and histories $\mathbf{h}_0 = \{h_{1,0}, \dots, h_{n,0}\}$. Taking inspiration from the single-agent GRPO algorithm (Shao et al. 2024), at each turn t , each agent takes action by generating a group of responses $a_{i,t}^G = \{a_{i,t}^{(1)}, \dots, a_{i,t}^{(G)}\}$ following its policy $\pi_i(\cdot | h_{i,t}^G)$ based on its observation-action history $h_{i,t}^G = \{h_{i,t}^{(1)}, \dots, h_{i,t}^{(G)}\}$. The actions of individual agents are aggregated to form a group of joint actions $\mathbf{a}_t^G = \{a_{0,t}^G, \dots, a_{n,t}^G\}$. The agents receive a group of joint rewards r_t^G for their responses \mathbf{a}_t^G , which conditions on the accessible part of the state $R(\cdot | s_t^{\text{acc},G}, \mathbf{a}_t^G)$. The joint actions triggers the transition $T(\cdot | s_t^G, \mathbf{a}_t^G)$, where agents receive new observations $o_{i,t+1}^G = \{o_{i,t+1}^{(1)}, \dots, o_{i,t+1}^{(G)}\}$ and use them to construct histories $h_{i,t+1}^G = \{h_{i,t}^G, a_{i,t}^G, o_{i,t+1}^G\}$. This process continues until terminated at turn H or the task is completed.

We employ stochastic gradient descent to train agents at the end of each episode. Without explicit value models, estimating history-action values incurs high variance. To stabilize training, we estimate the expected return of the current history by averaging over a group of Monte Carlo samples of the joint return $R_t^G = \{R_t^{(1)}, \dots, R_t^{(G)}\}$. As a result, we are able to generate a centralized estimate (common in MARL) without a large value model. For each turn t , the advantage of each joint action in the group is calculated as,

$$\hat{A}_t^{(g)} = R_t^{(g)} - \frac{1}{G} \sum_{g=1}^G R_t^{(g)}, \quad (1)$$

where $R_t^{(g)} = \sum_{\tau=t}^{H-1} r_{\tau}^{(g)}$. Inspired by GRPO (Shao et al. 2024; Liu et al. 2025b), and MAPPO (Yu et al. 2022), the centralized advantage values can be used to update policy π_i (parameterized by θ_i) for each agent i . MAGRPO does not have importance sampling and thereby epsilon clipping for simplicity, and the KL divergence coefficient is set to be 0 to encourage greater policy deviation from the base model,

$$J(\theta_i) = \mathbb{E}_{\mathbf{o}_0 \sim \mathcal{D}, \mathbf{h}^G \sim \pi_{\theta}} \left[\frac{1}{G} \sum_{g=1}^G \hat{A}_t^{(g)} \log \pi_{\theta_i}(a_{i,t}^{(g)} | h_{i,t}^G) \right]. \quad (2)$$

Experiments

We evaluate MAGRPO on LLM writing and coding collaboration. Datasets, reward specifications, and additional results are provided in the Appendix.

Writing Collaboration

We explore LLM collaboration for article writing using MAGRPO across 2 classic tasks: summarization and expansion.

TLDR Summarization When reading a long article, readers often seek to quickly grasp its core ideas. If the topic is of interest, they may wish to delve deeper into specific details while still avoiding a complete reading through the full document. This calls for a summarization system to generate summaries at varying levels of detail. We frame this task using TLDR summarization as an illustrative example.

The TLDR dataset comprises unabridged Reddit posts in the prompt and concise summaries appended by the author in the completion. In our experiment, 2 *Qwen3-1.7B* agents independently summarize the prompt without using completion. The first agent functions as a core-idea (TLDR) generator, producing a concise paragraph, while the second agent serves as a detailed summarizer, providing more comprehensive information.

To quantify the summarization quality, we employ a relatively simple combination of 3 metrics. Structure measures the lengths and the length ratio of the 2 summaries, to ensure the TLDR is concise and the detailed summary is sufficiently long. Style consistency is assessed using the normalized Jaccard similarity coefficient, calculated as the ratio of the intersection size to the union size of unique words (or n-grams) between responses. A high style consistency reward typically indicates that the summarizers adopt similar stylistic patterns while avoiding identical wording. Logical coherence is quantified by counting the occurrences of transition words. Positive reward is given for using transition words, but the reward decreases logarithmically as more are used. These metrics are simple approximations of what more complex reward models may evaluate. Other (simpler or more complex) metrics or reward models could also be used. The total reward combines these metrics through a weighted summation. More details regarding our reward model and hyperparameters are provided in the Appendix.

arXiv Expansion Writing a long article typically requires contributions from multiple writers, each responsible for different sections. As a simple scenario, 2 agents can collabor-

Method	Dataset	Efficiency		Article Quality (%)			Return (%)
		Speed	Response Time	Structure	Consistency	Coherence	
Single Model	TLDR	64.1	6.6	43.8	97.6	52.8	36.7
	arXiv	65.4	6.5	51.2	87.2	71.1	44.9
Parallel Generation	TLDR	185.6	2.1	25.9	98.3	56.5	23.2
	arXiv	190.6	2.1	71.5	64.2	61.5	59.6
Sequential Generation	TLDR	98.7	4.3	33.5	98.5	64.5	21.7
	arXiv	85.8	4.3	92.4	97.8	64.3	87.7
One-Round Discussion	TLDR	100.4	4.3	35.9	98.8	60.8	22.3
	arXiv	95.4	4.3	84.6	71.8	66.0	76.6
MAGRPO (Ours)	TLDR	202.3	2.1	98.7	97.1	78.5	94.5
	arXiv	193.8	2.1	97.9	96.2	69.7	93.1

Table 1: Performance of MAGRPO against baselines on TLDR and arXiv. Speed (tokens/s) and response time (s) are measured on GeForce RTX 5090s. Results are normalized to the return scale. **Bolds** indicate the best performance on each dataset.

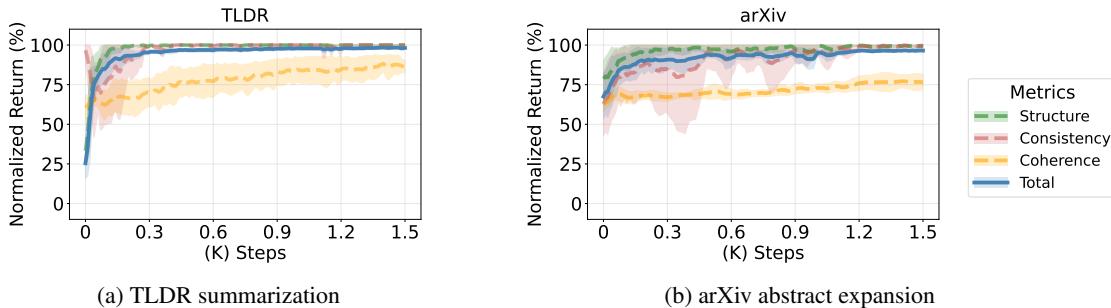


Figure 2: Normalized returns on writing collaboration: (a) structural wellness (dashed green); (b) style consistency (dashed red); (c) coherence (dashed orange); (d) total rewards (solid blue). All returns are normalized to the return scale.

rate to generate introductions from the abstract of arXiv papers. The first agent outlines the research background and motivation, while the other presents the proposed methods and their experiments. The combined paragraphs should be coherent and consistent in style. Similar to the reward model in TLDR summarization, we employ the same evaluation metrics as proxies, with threshold hyperparameters specifically adjusted for this task.

Baselines We adopt a single model and 3 multi-agent methods as our baselines. To minimize the influence of prompts on our comparison, we keep the task description fixed and only add minimal instructions. In the single-agent baseline, we prompt with the article, the agent’s role (summarizer or expanding writer), and specific format instructions. Naive concatenation builds on it by dividing the task into subtasks, assigning each agent a specific portion to complete in parallel without explicit communication. The sequential pipeline introduces one-way communication, allowing one agent to respond based on both the task description and the other agent’s output. The one-round discussion baseline enables bidirectional communication: agents first receive the same prompts as in naive concatenation, then the prompts are augmented with the other’s first-turn response in the second turn. All baseline methods operate without fine-tuning and depend solely on prompt-level interactions. Detailed prompts for each baseline are in the Appendix.

Results In this experiment, we apply MAGRPO to optimize the dual *Qwen3-1.7B* system in one turn. Figure 2a and Figure 2b show the evaluation results on TLDR and arXiv over 10 runs. The upward trend on all metric curves indicates that 2 agents gradually cooperate to generate coherent and consistent content with a well-organized structure. In the TLDR summarization, while the structure and logical coherence monotonically increase throughout training, the style consistency curves exhibit a decrease in the first 100 steps. This occurs as agents temporally diverge in styles to optimize other cooperative objectives, but their styles are gradually realigned and stabilized with sufficient training.

As shown in Table 1, MAGRPO is 3 times faster compared to the single *Qwen3-4B* model, which has a comparable number of parameters to our dual *Qwen3-1.7B* system. Despite receiving detailed instructions, *Qwen3-4B* fails to produce well-structured responses. A similar issue appears in TLDR summarization but not in arXiv expansion under multi-agent settings. This is because the outputs of homogeneous agents are naturally similar in length, which fortuitously aligns with the preference of the reward model.

Among the multi-agent baselines, parallel generation is the only one that achieves a comparable speed to ours, but it fails to generate well-structured and coherent texts due to the lack of cooperation. Sequential generation and discussion-based approaches occasionally enhance coordination through specific prompts. However, they still under-

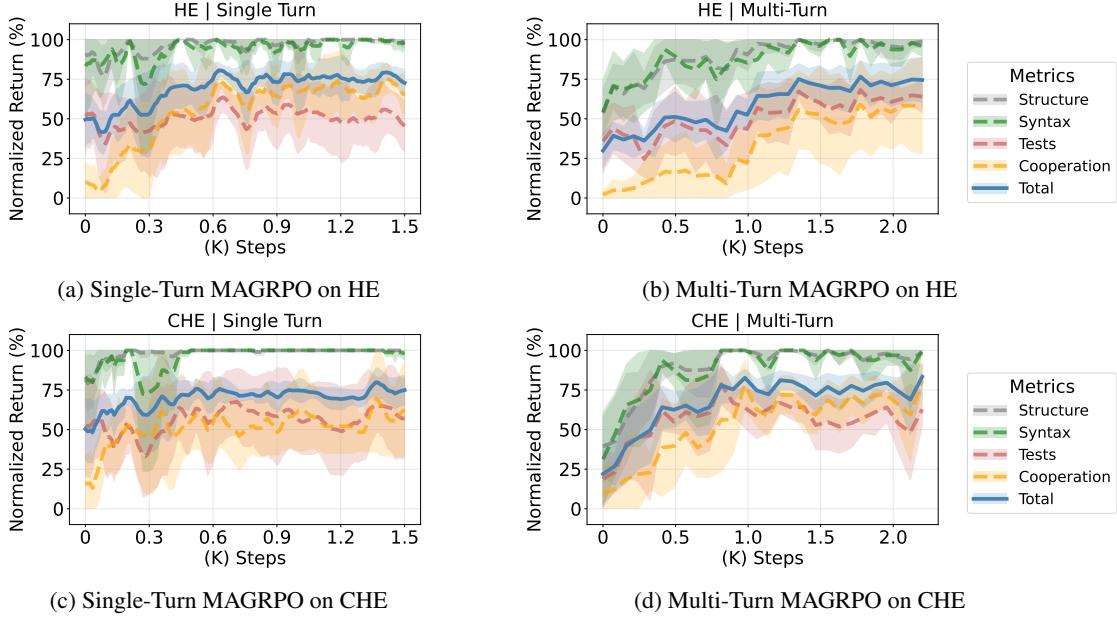


Figure 3: Normalized returns on coding collaboration: (a) structural wellness (dashed grey); (b) syntax correctness (dashed green); (c) Test score (dashed red); (d) cooperation rewards (dashed yellow); (e) total return (solid blue).

perform ours in efficiency and coherence, resulting in lower total return. The limited effectiveness of prompt-instructed coordination constrains their scalability to more complex scenarios involving large numbers of agents or extended multi-turn interactions (Estornell and Liu 2024).

Coding Collaboration

In large-scale software development, numerous developers collaborate to implement complex systems. Employing LLMs as developers is a promising direction, but coordinating them is challenging due to diverse cooperation schemes and complex failure modes. We simplify this task by using 2 *Qwen2.5-Coder-3B* agents to generate Python functions collaboratively. A helper agent produces auxiliary functions to support a main function generator, without any direct communication. The outputs from both agents, along with required libraries, are aggregated into complete code snippets.

HumanEval We evaluate MAGRPO on the HumanEval (HE) dataset, which contains 164 handwritten programming problems, each containing a natural language description (`prompt`), a function signature (`entry_point`), and a set of unit tests (`test`). To guide learning, we design a level-based reward model that prioritizes fundamental aspects of code generation. Structural integrity verifies the presence and correctness of both main and auxiliary function definitions; syntactic correctness ensures compliance with Python syntax; test pass rate assesses functional correctness based on the proportion of successfully passed unit tests; and a cooperation quality bonus is granted when the main function properly invokes and utilizes the auxiliary function. Rewards are accumulated only when all requirements at each preceding level are satisfied.

CoopHumanEval Some entries in HumanEval (HE) are not designed for coding collaboration; certain atomic operations (e.g., `strlen(string)`) can hardly be decomposed in a way that has meaningful cooperation. These noisy instances bias it towards invalid cooperation, such as merely wrapping the auxiliary function, thereby making training unstable. Thus, we construct a cooperative code generation dataset, CoopHumanEval (CHE), which comprises both original HE problems with cooperative potential (e.g., `prime_fib(n)`) and additional problems (e.g., `unique_digits(x)`). CHE problems are decomposable, where agents can explore effective cooperation schemes.

Baselines We adopt the fixed and fine-tuned single model, and 3 multi-agent methods on fixed base models, as our baselines. In the single-agent setting, the *Qwen2.5-Coder-7B* model generates a Python function based on the problem description in `prompt`. We also fine-tune this model on the training set to adapt it to this task. In the multi-agent setting, 2 *Qwen2.5-Coder-3B* models serve as agents: one generates a helper function, and the other produces the main. To minimize the influence of prompts on our comparison, we keep the problem description fixed and only add minimal coordination instructions. In the naive concatenation, agents are informed of their roles and generate outputs in parallel without communication. The sequential pipeline allows the main agent to respond based on the other’s output. In the one-round discussion, agents first receive the same prompts as naive concatenation, then the prompts are augmented with the other’s last-turn response in the subsequent turns.

Results We optimize the interaction between 2 agents using single-turn and multi-turn MAGRPO. In the multi-turn setting, problem descriptions and the agents’ initial re-

Method	Dataset	Efficiency			Code Quality (%)			Return (%)
		Speed	Response Time	Structure	Syntax	Tests	Cooperation	
Single Model	HE	73.1	1.6	100.0	100.0	64.8	–	–
	CHE	65.5	1.4	100.0	100.0	63.4	–	–
Fine-Tuned Single Model	HE	72.2	1.6	100.0	99.2	65.6	–	–
	CHE	60.6	1.4	100.0	100.0	66.7	–	–
Naive Concatenation	HE	194.9	1.1	96.1	90.6	42.5	22.7	53.9
	CHE	189.4	1.1	97.5	95.0	40.1	24.0	54.3
Sequential Pipeline	HE	99.6	2.2	98.4	96.5	56.4	35.1	63.1
	CHE	97.4	2.0	97.5	96.3	55.2	35.2	62.5
One-Round Discussion	HE	82.5	2.8	98.1	94.8	41.2	30.2	57.5
	CHE	78.3	2.8	97.5	96.3	41.9	34.8	59.5
Single-Turn MAGRPO (Ours)	HE	190.0	1.5	100.0	97.8	61.6	83.4	83.7
	CHE	192.4	1.5	98.8	97.5	71.2	83.7	86.0
Multi-Turn MAGRPO (Ours)	HE	96.5	2.7	99.7	96.9	68.4	84.9	86.7
	CHE	97.1	2.5	98.6	98.2	75.0	86.3	88.5

Table 2: Performance comparison of MAGRPO against baselines on HE and CHE. Speed (tokens/s) and response time (s) are recorded on GeForce RTX 5090s. Results are normalized to the return scale and averaged over 10 runs; rewards are level-based. **Bolds** indicate the best performance of each metric on each dataset.

sponses are provided to static checking (AST) and dynamic execution models, which report errors for each agent. Ablations of external tools used, including self-evolving and expert guidance, are provided in the Appendix.

Figure 3a and Figure 3b show the normalized return of MAGRPO on HE over 10 runs. Single-turn MAGRPO training improves the syntactical correctness and develops valid cooperation, while the test pass rate does not show much progress. As for the multi-turning training, agents are initially overwhelmed by the external model’s feedback, resulting in even lower initial returns. They gradually adopt the error signals and improve their returns. However, the test pass rate shows no significant improvement over the single-agent model, due to noisy entries in the dataset and hence unreliable feedback. This reflects the complexity and delicacy of decentralized coder coordination, where the main agent must accurately infer the functionality of auxiliary modules and trust their correctness without communication.

The performance of single-turn and multi-turn MAGRPO on the CHE dataset is shown in Figures 3c and 3d. Results show that MAGRPO achieves higher overall returns and lower variance when trained on CHE over HE. In the multi-turn setting, although agents initially struggle to interpret the feedback, like training on HE, the returns gradually improve and eventually surpass those in the single-turn training. This demonstrates that, when trained on a dataset with well-defined cooperative structures, agents can learn to utilize error messages to improve their response quality.

Table 2 presents a performance comparison between MAGRPO and baselines on both HE and CHE. By GRPO fine-tuning, the performance of *Qwen2.5-Coder-7B* model only improves slightly as the logic of test problems differs substantially. Compared to a single model, the naive concatenation method has lower test pass rates, as the main agent may generate code based on incorrect assumptions about auxiliary functionality. In the sequential pipeline method, the main agent can provide a backup for the auxiliary func-

tion when it identifies potential vulnerabilities, thereby improving the test pass rate. However, this comes at the cost of slower inference speed. Although the one-round discussion method involves more communication between agents, its effectiveness remains limited because the agents’ mutual adaptations to each other’s last responses can become misaligned. MAGRPO outperforms all baselines on both CHE and HE by facilitating effective cooperation and leveraging feedback from the external model. Additional results, including pass@k, are presented in the Appendix.

Cooperation Schemes MAGRPO identifies diverse cooperation schemes. In some cases, the auxiliary function handles the core logic, while the main agent adds backup logic or decorations to improve the overall solution. Alternatively, the main agent may act as a coordinator, decomposing the problem and assigning subtasks to the auxiliary agent. The auxiliary function may serve as a strategy filter, guiding the main agent to generate code for specific cases. While coordinator and strategy-filter schemes can improve inference efficiency, they are more prone to syntax and logical errors. With limited cooperation-oriented training data, the main agent typically resorts to more conservative roles, i.e., fallback or decoration. These cooperation schemes emerge during training under a relatively simple joint reward. More refined design patterns will likely be found when training agents to develop large-scale coding projects.

Conclusion

In this paper, we model LLM collaboration as a cooperative MARL problem and formalize it as a Dec-POMDP. We propose the MAGRPO algorithm to optimize agent cooperation through shared rewards. Our experiments in coding and writing collaboration show that MAGRPO enables agents to efficiently generate high-quality responses via effective collaboration. Our work encourages future exploration of MARL-based methods for scalable and robust LLM collaboration.

Acknowledgments

This work was partially funded by NSF grants #2044993 and #2409351. It used Delta and DeltaAI computing resources at the National Center for Supercomputing Applications through allocation CIS250443 and CIS250554 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support program, which is supported by NSF grants #2138259, #2138286, #2138307, #2137603, and #2138296.

We thank Tianle Chen for improving our MARL training framework for LLM collaboration, CoMLRL, and members of the Lab for Learning and Planning in Robotics (LLPR) for the valuable discussion. We thank Gregory Bauer and Brett Bode for helping us resolve job failure issues.

References

Achiam, J.; Adler, S.; Agarwal, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Albrecht, S. V.; Christianos, F.; and Schäfer, L. 2024. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press.

Amato, C. 2024. An Introduction to Centralized Training for Decentralized Execution in Cooperative Multi-Agent Reinforcement Learning. *arXiv* 2409.03052. *arXiv:2409.03052*.

Amato, C. 2025. An Initial Introduction to Cooperative Multi-Agent Reinforcement Learning. *arXiv:2405.06161*.

Anil, R.; Borgeaud, S.; Alayrac, J.-B.; et al. 2025. Gemini: A Family of Highly Capable Multimodal Models. *arXiv:2312.11805*.

Anthropic. 2023. Collective Constitutional AI: Aligning a Language Model with Public Input.

Anthropic. 2024. How We Built a Multi-Agent Research System.

Berner, C.; Brockman, G.; Chan, B.; et al. 2019. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*.

Cemri, M.; Pan, M. Z.; Yang, S.; Agrawal, L. A.; Chopra, B.; Tiwari, R.; Keutzer, K.; Parameswaran, A.; Klein, D.; Ramchandran, K.; Zaharia, M.; Gonzalez, J. E.; and Stoica, I. 2025. Why Do Multi-Agent LLM Systems Fail?

Chan, C.-M.; Chen, W.; Su, Y.; Yu, J.; Xue, W.; Zhang, S.; Fu, J.; and Liu, Z. 2023. ChatEval: Towards Better LLM-based Evaluators through Multi-Agent Debate.

Chen, M. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.

Du, Y.; Li, S.; Torralba, A.; Tenenbaum, J. B.; and Mordatch, I. 2023. Improving Factuality and Reasoning in Language Models through Multiagent Debate. *arXiv:2305.14325*.

Estornell, A.; and Liu, Y. 2024. Multi-LLM Debate: Framework, Principles, and Interventions. In *NeurIPS*.

Estornell, A.; Ton, J.-F.; Taufiq, M. F.; and Li, H. 2025. How to train a leader: Hierarchical reasoning in multi-agent llms. *arXiv preprint arXiv:2507.08960*.

Foerster, J.; Chen, R.; Al-Shedivat, M.; Whiteson, S.; Abbeel, P.; and Mordatch, I. 2018a. Learning with Opponent-Learning Awareness. In *AAMAS 2018: Proceedings of the Seventeenth International Joint Conference on Autonomous Agents and Multi-Agent Systems*.

Foerster, J.; Farquhar, G.; Afouras, T.; Nardelli, N.; and Whiteson, S. 2018b. Counterfactual Multi-Agent Policy Gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Grattafiori, A.; Dubey, A.; Jauhri, A.; et al. 2024. The Llama 3 Herd of Models.

Guo, D.; Yang, D.; Zhang, H.; et al. 2025. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature*, 645(8081): 633–638.

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*.

Liang, T.; He, Z.; Jiao, W.; Wang, X.; Wang, Y.; Wang, R.; Yang, Y.; Shi, S.; and Tu, Z. 2024. Encouraging Divergent Thinking in Large Language Models through Multi-Agent Debate. *arXiv:2305.19118*.

Lifshitz, S.; McIlraith, S. A.; and Du, Y. 2025. Multi-Agent Verification: Scaling Test-Time Compute with Multiple Verifiers. *arXiv:2502.20379*.

Liu, B.; Guertler, L.; Yu, S.; Liu, Z.; Qi, P.; Balcels, D.; Liu, M.; Tan, C.; Shi, W.; Lin, M.; et al. 2025a. SPIRAL: Self-Play on Zero-Sum Games Incentivizes Reasoning via Multi-Agent Multi-Turn Reinforcement Learning. *arXiv preprint arXiv:2506.24119*.

Liu, S.; Wang, Y.; Chen, X.; Fu, Y.; and Di, X. 2022. SMART-eFlo: An integrated SUMO-gym framework for multi-agent reinforcement learning in electric fleet management problem. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, 3026–3031. IEEE.

Liu, Z.; Chen, C.; Li, W.; Qi, P.; Pang, T.; Du, C.; Lee, W. S.; and Lin, M. 2025b. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*.

Lowe, R.; Wu, Y. I.; Tamar, A.; Harb, J.; Pieter Abbeel, O.; and Mordatch, I. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*.

Lyu, X.; Baisero, A.; Xiao, Y.; Daley, B.; and Amato, C. 2023. On Centralized Critics in Multi-Agent Reinforcement Learning. *Journal of Artificial Intelligence Research*, 77: 295–354.

Ma, H.; Hu, T.; Pu, Z.; Boyin, L.; Ai, X.; Liang, Y.; and Chen, M. 2024. Coevolving with the other you: Fine-tuning llm with sequential cooperative multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 37: 15497–15525.

Marchesini, E.; Baisero, A.; Bathi, R.; and Amato, C. 2025. On Stateful Value Factorization in Multi-Agent Reinforcement Learning. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*.

NVIDIA. 2024. Introduction to LLM Agents.

Oliehoek, F. A.; and Amato, C. 2016. *A Concise Introduction to Decentralized POMDPs*. Springer.

Oliehoek, F. A.; Spaan, M. T. J.; and Vlassis, N. 2008. Optimal and Approximate Q-value Functions for Decentralized POMDPs. *Journal of Artificial Intelligence Research*, 32: 289–353.

Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35: 27730–27744.

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

Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36: 53728–53741.

Rashid, T.; Samvelyan, M.; de Witt, C. S.; Farquhar, G.; Foerster, J.; and Whiteson, S. 2020. Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. *Journal of Machine Learning Research*, 21(178): 1–51.

Samvelyan, M.; Rashid, T.; De Witt, C. S.; Farquhar, G.; Nardelli, N.; Rudner, T. G.; Hung, C.-M.; Torr, P. H.; Foerster, J.; and Whiteson, S. 2019. The starcraft multi-agent challenge. *arXiv preprint arXiv:1902.04043*.

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*.

Shao, Z.; Wang, P.; Zhu, Q.; et al. 2024. DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models. *arXiv:2402.03300*.

Shoham, Y.; and Leyton-Brown, K. 2009. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge, UK: Cambridge University Press.

Skreta, M.; Yoshikawa, N.; Arellano-Rubach, S.; Ji, Z.; Kristensen, L. B.; Darvish, K.; Aspuru-Guzik, A.; Shkurti, F.; and Garg, A. 2023. Errors are Useful Prompts: Instruction Guided Task Programming with Verifier-Assisted Iterative Prompting. *arXiv:2303.14100*.

Stone, P.; and Veloso, M. 2000. Multiagent Systems: A survey from a machine learning perspective. *Autonomous Robots*, 8(3): 345–383.

Subramaniam, V.; Du, Y.; Tenenbaum, J. B.; Torralba, A.; Li, S.; and Mordatch, I. 2025. Multiagent finetuning: Self improvement with diverse reasoning chains. *arXiv preprint arXiv:2501.05707*.

Tan, M. 1993. Multi-Agent Reinforcement Learning: Independent versus Cooperative Agents. In *Proceedings of the Tenth International Conference on Machine Learning*, 330–337. San Francisco, CA, USA: Morgan Kaufmann.

Uesato, J.; Kushman, N.; Kumar, R.; Song, F.; Siegel, N.; Wang, L.; Creswell, A.; Irving, G.; and Higgins, I. 2022. Solving math word problems with process-and outcome-based feedback. *arXiv preprint arXiv:2211.14275*.

Van der Hoek, W.; and Wooldridge, M. 2008. Multi-agent systems. *Foundations of Artificial Intelligence*, 3: 887–928.

Vinyals, O.; Ewalds, T.; Bartunov, S.; Georgiev, P.; Vezhnevets, A. S.; Yeo, M.; Makhzani, A.; Küttler, H.; Agapiou, J.; Schrittwieser, J.; et al. 2017. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782*.

Wang, B.; Zi, Y.; Sun, Y.; Zhao, Y.; and Qin, B. 2024. Rkld: Reverse kl-divergence-based knowledge distillation for unlearning personal information in large language models. *arXiv preprint arXiv:2406.01983*.

Wang, J.; Ren, Z.; Liu, T.; Yu, Y.; and Zhang, C. 2021. QPLEX: Duplex Dueling Multi-Agent Q-Learning. In *Proceedings of the International Conference on Learning Representations*.

Weiss, G. 1999. *Multiagent systems: a modern approach to distributed artificial intelligence*. MIT press.

Wiering, M. A. 2000. Multi-agent reinforcement learning for traffic light control. In *Proceedings of the Seventeenth International Conference on Machine Learning*, 1151–1158. Stanford, CA, USA: Morgan Kaufmann.

Wu, Q.; Bansal, G.; Zhang, J.; Wu, Y.; Li, B.; Zhu, E.; Jiang, L.; Zhang, X.; Zhang, S.; Liu, J.; Awadallah, A. H.; White, R. W.; Burger, D.; and Wang, C. 2023a. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation.

Wu, Z.; Hu, Y.; Shi, W.; Dziri, N.; Suhr, A.; Ammanabrolu, P.; Smith, N. A.; Ostendorf, M.; and Hajishirzi, H. 2023b. Fine-grained human feedback gives better rewards for language model training. *Advances in Neural Information Processing Systems*, 36: 59008–59033.

Xiao, Y.; Tan, W.; Hoffman, J.; Xia, T.; and Amato, C. 2025. Asynchronous Multi-Agent Deep Reinforcement Learning under Partial Observability. *International Journal of Robotics Research*, 35(14): 1760–1778.

Yu, C.; Velu, A.; Vinitsky, E.; Gao, J.; Wang, Y.; Bayen, A.; and Wu, Y. 2022. The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games. In *Thirty-Sixth Conference on Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track*.

Yue, Y.; Chen, Z.; Lu, R.; Zhao, A.; Wang, Z.; Yue, Y.; Song, S.; and Huang, G. 2025. Does Reinforcement Learning Really Incentivize Reasoning Capacity in LLMs Beyond the Base Model? *arXiv:2504.13837*.

Zhang, K.; Liu, R.; Zhu, X.; Tian, K.; Zeng, S.; Jia, G.; Fan, Y.; Lv, X.; Zuo, Y.; Jiang, C.; Liu, Z.; Wang, J.; Wang, Y.; Zhao, R.; Hua, E.; Wang, Y.; Wang, S.; Gao, J.; Long, X.; Sun, Y.; Ma, Z.; Cui, G.; Bai, L.; Ding, N.; Qi, B.; and Zhou, B. 2025. MARTI: A Framework for Multi-Agent LLM Systems Reinforced Training and Inference.

Zhao, P.; Sun, F.; Shen, X.; Yu, P.; Kong, Z.; Wang, Y.; and Lin, X. 2024. Pruning foundation models for high accuracy without retraining. *arXiv preprint arXiv:2410.15567*.

Ziegler, D. M.; Stiennon, N.; Wu, J.; Brown, T. B.; Radford, A.; Amodei, D.; Christiano, P.; and Irving, G. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.

Formalizations of Multi-Agent Interaction

Many studies adopt Partially Observable Stochastic Games (POSG) to model the LLM interaction in MAS (Liu et al. 2025a; Sarkar et al. 2025; Zhang et al. 2025). In this section, we show that Dec-POMDP offers special merits compared to POSG in the solution concept in the cooperative settings, thus more suited to model LLM collaboration.

Dec-POMDP

A Dec-POMDP is defined by $\langle \mathcal{I}, \mathcal{S}, \{\mathcal{O}_i\}, \{\mathcal{A}_i\}, R, T, H \rangle$. At each step t , since an agent cannot directly observe the state s_t , it usually maintains local observation-action history $h_{i,t} = (o_{i,0}, a_{i,0}, \dots, o_{i,t})$ to infer a belief over the underlying state. Decisions are made according to a local policy $\pi_i : \mathcal{H}_{i,t} \rightarrow \Delta(\mathcal{A}_i)$, which maps histories to probability distributions over actions. The set of all local policies forms the joint policy $\pi = \{\pi_1, \dots, \pi_n\}$. In cooperative settings, the objective is to maximize shared cumulative rewards. As proved in (Oliehoek, Spaan, and Vlassis 2008), there is always an optimal joint policy in a Dec-POMDP,

$$\pi^* = \operatorname{argmax}_{\pi \in \Pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{H-1} R(s_t, a_t) \right]. \quad (3)$$

POSG

A Partially Observable Stochastic Game (POSG), so-called Partially Observable Markov Game (POMG), does not assume cooperative behavior among agents. It can be either a cooperative, competitive, or mixed game. A POSG is defined as $\langle \mathcal{I}, \mathcal{S}, \{\mathcal{A}_i\}, T, \{\mathcal{O}_i\}, O, \{R_i\}, H \rangle$, where each agent has its own reward function $R_i : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. In POSG, each agent seeks to maximize its individual return under the fixed policies of all others π_{-i} . The optimal policy π_i^* for each agent $i \in \mathcal{I}$ is,

$$\pi_i^* = \operatorname{argmax}_{\pi_i \in \Pi_i} \mathbb{E}_{\pi_i, \pi_{-i}} \left[\sum_{t=0}^{H-1} R_i(s_t, a_t) \right], \quad (4)$$

The solutions for POSG are Nash Equilibria (NE), where no agents can unilaterally improve their returns by deviating from their policies. Formally, for all $i \in \mathcal{I}$ and any alternative policy $\pi_i \in \Pi_i$, NE satisfy

$$\mathbb{E} \left[\sum_{t=0}^{H-1} R_i(s_t, a_t) \mid \pi_i^*, \pi_{-i}^* \right] \geq \mathbb{E} \left[\sum_{t=0}^{H-1} R_i(s_t, a_t) \mid \pi_i, \pi_{-i}^* \right]. \quad (5)$$

Like Dec-POMDP, the decision-making in POSG is still concurrent (as stochastic games), where all agents act synchronously at each time step. In contrast, turn-based interactions, where agents take turns to act (e.g., chess, Kuhn Poker, tic-tac-toe), are typically modeled as extensive-form games.

Non-Optimality of POSG Solutions

We illustrate that the solutions of POSG, i.e., NE, may not necessarily lead to joint optimality in cooperative settings.

Consider a one-step matrix game involving 2 agents, where each agent selects an action from the action space

$\mathcal{A} = \{\mathcal{A}^{(1)}, \mathcal{A}^{(2)}\}$. The joint action profile determines the utility as presented in Table 3.

$a_1 \setminus a_2$	$\mathcal{A}^{(1)}$	$\mathcal{A}^{(2)}$
$\mathcal{A}^{(1)}$	10	7
$\mathcal{A}^{(2)}$	7	0

Table 3: Joint utility matrix of 2 agents.

This matrix game can be potentially decomposed into 2 POSG in Table 4 through reward shaping.

$a_1 \setminus a_2$	$\mathcal{A}^{(1)}$	$\mathcal{A}^{(2)}$	$a_1 \setminus a_2$	$\mathcal{A}^{(1)}$	$\mathcal{A}^{(2)}$
$\mathcal{A}^{(1)}$	(5, 5)	(3, 4)	$\mathcal{A}^{(1)}$	(5, 5)	(1, 6)
$\mathcal{A}^{(2)}$	(4, 3)	(0, 0)	$\mathcal{A}^{(2)}$	(6, 1)	(0, 0)

(a) POSG 1

(b) POSG 2

Table 4: Return tables of 2 POSG.

In the POSG presented in Table 4a, $(\mathcal{A}^{(1)}, \mathcal{A}^{(1)})$ is a Nash equilibrium (blue triangle in Figure 4a). When $a_1 = \mathcal{A}^{(1)}$, $U_2(\mathcal{A}^{(1)}, \mathcal{A}^{(1)}) > U_2(\mathcal{A}^{(1)}, \mathcal{A}^{(2)})$; when $a_1 = \mathcal{A}^{(2)}$, $U_2(\mathcal{A}^{(2)}, \mathcal{A}^{(1)}) > U_2(\mathcal{A}^{(2)}, \mathcal{A}^{(2)})$. Therefore, the best response for agent 2 is $a_2^* = \mathcal{A}^{(1)}$. Similarly, since $U_1(\mathcal{A}^{(1)}, \mathcal{A}^{(1)}) > U_1(\mathcal{A}^{(2)}, \mathcal{A}^{(1)})$, we obtain $a_1^* = \mathcal{A}^{(1)}$. This NE also achieves joint optimality with the maximum utility $5 + 5 = 10$ (red square in Figure 4a).

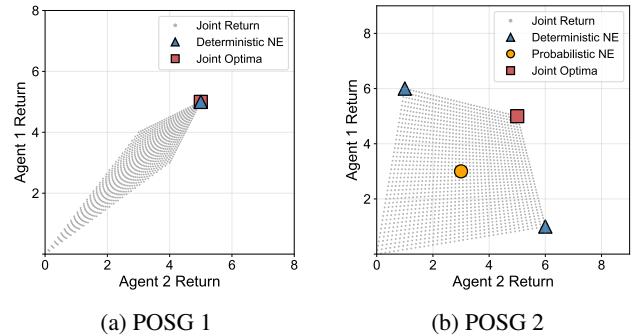


Figure 4: Utility spaces of 2 POSG.

However, certain reward decompositions may yield non-optimal solutions for cooperative games in Table 4, even when POSG solutions reach NE. For the POSG shown in Table 4b, the deterministic NE are $(\mathcal{A}^{(1)}, \mathcal{A}^{(2)}), (\mathcal{A}^{(2)}, \mathcal{A}^{(1)})$ (blue triangles in Figure 4b). When $a_1 = \mathcal{A}^{(1)}$, agent 2 prefers $\mathcal{A}^{(2)}$ as $U_2(\mathcal{A}^{(1)}, \mathcal{A}^{(2)}) > U_2(\mathcal{A}^{(1)}, \mathcal{A}^{(1)})$; when $a_1 = \mathcal{A}^{(2)}$, agent 2 prefers $\mathcal{A}^{(1)}$ since $U_2(\mathcal{A}^{(2)}, \mathcal{A}^{(1)}) > U_2(\mathcal{A}^{(2)}, \mathcal{A}^{(2)})$. Agent 1 faces the same issue. Thus, neither agent can unilaterally improve their utilities by deviating. However, the collective utilities obtained from both policies yield $6 + 1 = 7 < 10$, which are suboptimal compared to the joint optimum (red square in Figure 4b).

In Table 4b, even the probabilistic NE under stochastic policies is still non-optimal. Suppose agent 1 selects $\mathcal{A}^{(1)}$

Method	Dataset	Pass@k (%)			Acc@k (%)			Coop@k (%)		
		@3	@5	@10	@3	@5	@10	@3	@5	@10
Single Model	HE	67.7	71.0	83.9	85.4	87.9	95.1	–	–	–
	CHE	68.8	75.0	81.3	75.0	81.3	88.8	–	–	–
Fine-Tuned Single Model	HE	69.2	72.9	84.6	83.0	89.2	95.6	–	–	–
	CHE	69.4	74.5	83.7	74.3	82.2	90.3	–	–	–
Naive Concatenation	HE	45.2	51.6	64.5	70.4	75.5	80.9	49.5	67.7	76.3
	CHE	43.8	56.3	68.8	57.0	63.8	73.8	47.9	69.3	81.3
Sequential Pipeline	HE	56.8	61.3	71.0	78.8	84.5	91.9	62.4	73.3	92.7
	CHE	75.0	81.5	87.5	88.2	89.5	91.3	75.0	75.0	81.3
One-Round Discussion	HE	51.6	61.3	71.0	71.8	81.3	87.5	58.1	70.0	78.7
	CHE	50.0	68.8	68.8	75.4	82.5	82.0	66.7	68.7	75.0
Single-Turn MAGRPO (Ours)	HE	64.8	68.1	71.0	75.3	76.3	86.4	83.8	90.3	90.3
	CHE	75.0	75.0	81.2	80.0	82.5	87.5	87.5	93.8	93.8
Multi-Turn MAGRPO (Ours)	HE	71.0	80.6	90.3	85.7	92.6	94.7	93.5	96.8	96.8
	CHE	75.0	81.5	87.5	86.4	92.5	95.4	93.8	96.8	100.0

Table 5: Performance comparison between MAGRPO and baseline methods with pass@k, acc@k, coop@k, on HE and CHE. The **bold** texts indicate the best performance of each metric on each dataset.

with probability p , and agent 2 selects $\mathcal{A}^{(1)}$ with probability q , $R_1(\mathcal{A}^{(1)}, \cdot) = 5q + (1 - q) = 4q + 1$, $R_1(\mathcal{A}^{(2)}, \cdot) = 6q$, $R_1(\mathcal{A}^{(1)}, \cdot) = R_1(\mathcal{A}^{(2)}, \cdot)$ yields $q = 0.5$; similarly, $R_2(\cdot, \mathcal{A}^{(1)}) = 5p + (1 - p) = 4p + 1$, $R_2(\cdot, \mathcal{A}^{(2)}) = 6p$, $R_2(\mathcal{A}^{(1)}, \cdot) = R_2(\mathcal{A}^{(2)}, \cdot)$ yields $p = 0.5$. This probabilistic NE, $\pi_1^*(\mathcal{A}^{(1)}) = \pi_1^*(\mathcal{A}^{(2)})$, $\pi_2^*(\mathcal{A}^{(1)}) = \pi_2^*(\mathcal{A}^{(2)})$ leads to overall utilities $3 + 3 = 6 < 10$ (orange circle in Figure 4b).

Although appropriate reward shaping techniques can transform a cooperative game into a POSG like Table 4a to make the NE also jointly optimal, this becomes more challenging when more agents are involved and episodes become longer. We employ Dec-POMDP to avoid the intricate reward engineering and seek the joint optimality.

Additional Results

We report additional results in this section to validate the effectiveness of our approach.

@k Ablation

In LLMs, single-run inference often leads to high variance. To provide a more reliable evaluation of the capacities (Yue et al. 2025), we evaluate the pass, test accuracy, and cooperation at k runs (pass@k, acc@k, and coop@k) on MAGRPO and baselines. These @k metrics measure the best outcome among k generated solutions for each problem and are averaged over all problems in \mathcal{D} .

Pass@k is calculated as the probability that at least one out of k generated solutions passes all test cases (Chen 2021). Specifically, a set \mathcal{P} of k solutions is randomly sampled from a pool of \mathcal{M} generated solutions. To make it consistent with other @k metrics, we express pass@k as,

$$\text{Pass@k} = \frac{1}{|\mathcal{D}|} \sum_{j=1}^{|\mathcal{D}|} \mathbb{E}_{\mathcal{P} \sim \text{Sample}(\mathcal{M}, k)} \left[\max_{p \in \mathcal{P}} \mathbb{1}(n_j^{\text{pass}} = n_j) \right], \quad (6)$$

where $\mathbb{1}(n_j^{\text{pass}} = n_j)$ is the indicator function that equals 1 if all test cases are passed in problem j (i.e., the number of passed tests n_j^{pass} equals the total number of tests n_j), and 0 otherwise.

However, pass@k does not offer a fine-grained assessment of solution quality, as functional correctness is represented as a binary variable. As a result, failing a single test case is treated the same as failing all in pass@k. To provide a more unbiased evaluation, we use accuracy@k, defined as the highest test accuracy among the k generated solutions, averaged across all problems.

$$\text{Acc@k} = \frac{1}{|\mathcal{D}|} \sum_{j=1}^{|\mathcal{D}|} \mathbb{E}_{\mathcal{P} \sim \text{Sample}(\mathcal{M}, k)} \left[\max_{p \in \mathcal{P}} \frac{n_j^{\text{pass}}}{n_j} \right], \quad (7)$$

where n_j^{pass} and n_j are the number of passed and total unit tests of problem j , and n_j^{pass}/n_j is the test accuracy.

Similar to acc@k, we also propose coop@k, which measures the average of the highest normalized cooperation return achieved among k runs across all problems. Formally, the Coop@k is defined as,

$$\text{Coop@k} = \frac{1}{|\mathcal{D}|} \sum_{j=1}^{|\mathcal{D}|} \mathbb{E}_{\mathcal{P} \sim \text{Sample}(\mathcal{M}, k)} \left[\max_{p \in \mathcal{P}} R_p^{\text{coop}} \right], \quad (8)$$

where $R_p^{\text{coop}} = \sum_{t=0}^{H-1} r_{p,t}^{\text{coop}}$ is the cooperation return over horizon H , and $r_{p,t}^{\text{coop}}$ denotes the cooperation reward obtained by solution p at turn t .

We generate 15 samples in \mathcal{M} and evaluate all methods with $k = 10$. Table 5 presents the results for pass@k, acc@k, and coop@k at $k = 3, 5, 10$, comparing MAGRPO with baseline methods. As expected, the trends in @k metrics are consistent with the @1 results. By fine-tuning with GRPO, the pass@k, acc@k, and coop@k are slightly improved. The naive concatenation method remains worse than the single-agent baseline in terms of pass@k, as the main agent may generate code based on incorrect assumptions

about the auxiliary function. The sequential pipeline mitigates this issue by allowing the main agent to reference the auxiliary output during generation, yielding substantial improvements across all @k metrics, particularly on CHE. Although the one-round discussion method shares the same prompts as naive concatenation in the first turn, the additional discussion round yields limited improvement across the @k metrics due to misalignment in cross-adaptation.

Multi-turn MAGRPO outperforms all baselines across most @k metrics by leveraging feedback from external sources. These more comprehensive experiments further validate the general effectiveness of our approach and highlight promising directions for future work, such as training LLMs through interactions with static analyzers, sandbox-based evaluations, or expert models. Notably, the acc@k metric offers a more fine-grained perspective on performance trends as k varies. In some cases, slightly increasing k may not yield additional solutions that pass all test cases (as measured by pass@k), yet the improvement is still captured through higher accuracy.

Cooperation Schemes

By training the auxiliary and main coders to cooperate under minimal constraints (with only the problem description and their respective roles provided), diverse cooperation schemes naturally emerge. We present 4 representative schemes observed in our models.

Fallback A commonly observed cooperation scheme is the main agent providing a fallback for the auxiliary function. Although prompted to use the auxiliary output, the main agent may still doubt its completeness and reimplement the solution as a backup in case of potential errors.

For example, the problem `compare_one(a, b)` asks to return the larger of 2 values in its original type (string, float, or int), or `None` if equal. Listing 1 demonstrates how the fallback mechanism provides an alternative in case of runtime errors in the auxiliary implementation. The `aux(x, y)` function converts variable `x` to a float, resulting in an operand type mismatch if `y` is a string. The main agent handles the type error by falling back to its secondary implementation. While the main implementation may still be vulnerable, this mechanism benefits small models that commonly encounter syntax and runtime errors. Despite some redundancy, a proper fallback can enhance code robustness.

Decorator

When the main agent trusts the auxiliary to accurately handle the core logic, it typically acts as a decorator, adding complementary features to enhance the output. The main agent could handle edge cases for improved robustness or reformat the outputs (e.g., sorting, rounding, and adjusting units). The necessity of such decorating logic depends on the auxiliary agent's functionality and completeness.

Listing 1: `compare_one(a, b)`

```

1  def aux(x, y):
2      # Convert x to a float
3      if isinstance(x, str):
4          x = float(x.replace(", ", "."))
5
6      # Vulnerable if y in different types
7      if x < y:
8          return y
9      else:
10         return x
11
12 def compare_one(a, b):
13     try:
14         return aux(a, b)
15     except Exception:
16         # Robust implementation
17         if type(a) == type(b):
18             if a > b:
19                 return a
20             elif a < b:
21                 return b
22             else:
23                 return None
24         elif type(a) == float and type(b) ==
25             str:
26             b_float = float(b.replace(", ", "."))
27             if a > b_float:
28                 return a
29             elif a < b_float:
30                 return b_float
31             else:
32                 return None
33         elif type(a) == str and type(b) ==
34             float:
35             a_str = str(a.replace(", ", "."))
36             a_float = float(a_str)
37             if a_float > b:
38                 return a_str
39             elif a_float < b:
40                 return b
41             else:
42                 return None
43
44     else:
45         return None

```

Listing 2 illustrates a decorator main function and how its utility varies depending on the completeness of the auxiliary implementation. In this problem, 2 agents generate codes together to check if a given string is a palindrome. When paired with `aux1` that only handles the recursion boundary condition of single-character strings, the empty string check of the main serves as a necessary edge case handler. However, when working with `aux2`, which already has a more comprehensive edge case consideration, this handle becomes redundant.

Coordinator In large-scale software systems, it would be beneficial to have pipelines for repeated or data-parallel operations (e.g., batch processing, stream transformations). This corresponds to the coordinator cooperation scheme in

our models, where the main agent divides the tasks into parts and assigns them to the auxiliary agent.

A simple example involves the main agent acting as an iterator, using a loop (e.g., a `for` loop) to structure the problem. The auxiliary function generates partial solutions within each iteration. These partial results are then aggregated to form the final output. However, this cooperation scheme is unstable, as it depends heavily on the correct functionality of the auxiliary agent. When the auxiliary agent fails to complete its subtask, the entire solution breaks down.

Listing 2: `is_palindrome(text)`

```

1  def aux1(text):
2      if len(text) == 1:
3          return True
4      return text[0] == text[-1] and aux1(
5          text[1:-1])
6
7  def aux2(text):
8      if len(text) <= 1:
9          return True
10     return text[0] == text[-1] and aux2(
11         text[1:-1])
12
13 def is_palindrome(text):
14     if not text:
15         return True
16
17     # Edge case handler
18     return aux1(text)
19
20     # Redundant decorator
21     return aux2(text)

```

Listing 3: `flip_case(string)`

```

1  def aux(string: str) -> str:
2      result = ""
3      for char in string:
4          if char.islower():
5              result += char.upper()
6          elif char.isupper():
7              result += char.lower()
8          else:
9              result += char
10     return result
11
12 def flip_case(string: str):
13     flipped = ""
14     for char in string:
15         flipped += aux(char)
16     return flipped

```

Listing 3 demonstrates a solution to flip the case of characters in a string. The auxiliary function flips the case of each character, while the main function calls this auxiliary function for each character and appends it to the result. This scheme can be extended to more complex scenarios, where subtasks are assigned in a hierarchical structure.

Strategy Filter When handling complex problems, the main agent may need to implement logic based on multiple conditions. In such cases, the auxiliary agent can act as a filter for specific branches of logic, often appearing within conditional blocks (e.g., following an `if` statement). This scheme resembles the adaptive control flow in practice. In rule-based pipelines, an auxiliary agent evaluates preconditions (e.g., task types, system status, configurations) and directs workers to execute appropriate subroutines, thereby enhancing project modularity.

Listing 4: `x_or_y(n, x, y)`

```

1  def aux(n):
2      if n < 2:
3          return False
4      if n == 2:
5          return True
6      if n % 2 == 0:
7          return False
8      for i in range(3, int(n**0.5) + 1,
9                      2):
10         if n % i == 0:
11             return False
12     return True
13
14 def x_or_y(n, x, y):
15     # Check if n is prime
16     if aux(n):
17         return x
18     else:
19         return y

```

Listing 4 presents a solution for `x_or_y(n, x, y)` problem, which returns `x` if `n` is prime and `y` otherwise. The auxiliary function handles the primality checking, while the main function is responsible for returning results. The same pattern can also be found in the solutions of `prime_fib(n)`, `factorize(n)`, and `largest_prime_factor(n)`.

Learning Modes

Figure 5 shows the reward curves of total returns in a 2-turn MAGRPO training with 2 different learning modes.

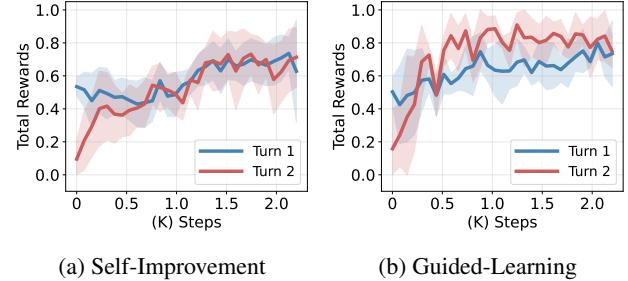


Figure 5: Learning modes in 2-turn MAGRPO training.

Figure 5a demonstrates a self-improvement mode, where agents interact with the tasks themselves and do not have

access to the external feedback. At each turn t , the prompts are just prompts and responses from the previous turns and a revision instruction. Although both curves improve as agents gradually develop cooperative behaviors, the performance of the second turn is still consistently similar to the first turn, suggesting that learning is primarily driven by direct task interaction.

Figure 5b illustrates the reward curves of guided-learning mode, where LLMs leverage external feedback to improve performance, e.g., feedback from expert agents, diagnosis from static checkers like AST, and sandbox tests. When employing *Claude-Sonnet-4* as an external model to provide more concrete suggestions or feeding error messages from sandbox tests, the performance of the second turn (red) exceeds first turn (blue), and both outperform those in the self-improvement. This indicates that appropriate guidance helps agents to refine the response in an efficient way. Due to the computational constraints, most models used in our setup have around 3B parameters and may struggle to interpret more abstract suggestions. We hypothesize that larger models with higher reasoning capabilities could benefit from more implicit guidance.

Experimental Configurations

This section outlines the hyperparameter settings and reward specifications used in our experiments.

Hyperparameters

For writing collaboration, we set the temperature to 0.8 and apply nucleus sampling with a threshold of 0.95 to encourage diverse generation. Policy deviation is regularized using a beta value of 0.02. The policy is optimized using a learning rate of 5×10^{-6} , and training is conducted for 1,500 steps.

For coding collaboration, the single-turn MAGRPO training uses a temperature of 0.7 and nucleus sampling with a threshold of 0.9. The learning rate is set to 1×10^{-6} , with 1,500 training steps. In the multi-turn setting, the discount factor is set to 1.0, and the learning rate is 5×10^{-6} , with 2,200 training steps.

Reward Specifications

Rewards are computed as a weighted sum of multiple metric-based components, following a hierarchical reward modeling scheme to prioritize fundamental objectives.

TLDR Summarization

- **Structural Wellness:** The structural wellness is assessed by the ratios of paragraph length and unique words. For the completion length, an appropriate ratio within $1.6\text{--}3.2\times$ receives the full rewards; ratios within the range of $1.1\text{--}5.0\times$ receive proportional rewards; while ratios outside receive no rewards and early termination of evaluation. For the ratio of unique words, we exclude the common stopwords. A ratio of $2.0\times$ or higher receives the maximum rewards; ratios between $1.3\text{--}2.0\times$ receive proportional rewards; ratios below $2.0\times$ result in no rewards and evaluation termination.

- **Style Consistency:** The style consistency is measured through Jaccard similarity of vocabulary between the completions (excluding stopwords). The Jaccard similarity scores are capped at 0.03 and normalized as rewards. We use a cap here to balance the needs of maintaining lexical consistency and vocabulary expansion in elaborated summarization.

- **Logical Coherence:** The logical coherence is evaluated through the presence and diversity of transition words in the completions. We check transition words across 12 categories, e.g., examples, explanation, contrast, etc. Additional rewards are given for using transition words in more categories, where $r = \min(0.6 \log(n + 1), 1)$, and n is the number of transition categories.

arXiv Expansion

- **Structural Wellness:** This metric evaluates the relative length and lexical diversity between the second and first completions. A length ratio within the optimal range of $1.0\text{--}1.3\times$ yields the maximum rewards, while ratios within the acceptable bounds of $0.8\text{--}1.5\times$ receive proportionally scaled rewards. Ratios outside this range result in zero reward and early termination. Similarly, a unique word ratio within $0.7\text{--}1.3\times$ receives the full rewards, ratios within $0.5\text{--}1.7\times$ are rewarded proportionally, and values outside this range lead to zero reward and evaluation termination.

- **Style Consistency:** Style consistency is quantified using Jaccard similarity between the 2 completions. The similarity score is capped at 0.23 and normalized as rewards.

- **Logical Coherence:** Logical coherence is assessed based on the presence of transition words across 12 categories. Additional rewards are given for using transition words in more categories, where $r = \min(0.4 \log(n + 1), 1)$, and n is the number of transition categories.

Coding Collaboration

- **Structural Integrity:** This metric verifies the correct implementation of both the auxiliary and main functions. To receive the base reward, the corresponding functions in the agents' completions must be properly defined and include return statements. Failure to define the main function results in evaluation termination.

- **Syntactical Correctness:** This metric assesses the syntactic validity of the concatenated code, which includes the libraries provided in the dataset, the auxiliary function defined by the helper agent, and the function defined by the main agent. Syntactic correctness is verified via static analysis, i.e., Abstract Syntax Tree (AST). The presence of syntax errors leads to the termination of the evaluation to avoid runtime failures.

- **Test Pass Rate:** This metric measures the percentage of unit tests passed during execution, with each test subject to a 10-second timeout. Rewards are assigned proportionally based on the number of successful assertions. If no tests pass, the evaluation is terminated.

- **Cooperation Quality:** A base bonus is applied if the main function calls the auxiliary. Additional rewards are

given when the main function implements substantive logic beyond simply wrapping the auxiliary.

Prompt Design

Writing Collaboration

TLDR In the TLDR summarization, the prompt field of the dataset is processed for each agent by using the following instructions.

Summary Agent

Create a concise summary response to this post.
Query: {prompt}
Instructions: Provide a brief and focused summary in a few sentences

Elaboration Agent

Create a detailed summary response to this post.
Query: {prompt}
Instructions: You should use transition words to improve flow

arXiv In the arXiv paragraph expansion, we use the abstract field of the dataset and process it as follows.

Background Agent

Based on the following scientific abstract, expand the content for the introduction section.
Abstract: {abstract}
Instructions:

- There is another agent that will provide the method and implications
- You just need to focus on the background and motivation
- Avoid repeating methodology and implications content

Method Agent

Based on the following scientific abstract, expand the content for the introduction section.
Abstract: {abstract}
Instructions:

- There is another agent that will provide the background and motivation
- You just need to focus on the method and implications
- Avoid repeating background and motivation content

Coding Collaboration

For HE and CHE, we extract the entry_point, params from the prompt field and instruct agents as follows.

Auxiliary Agent

Create a helper function for this coding problem.
Problem: {prompt}
Instructions:

- Output ONLY the function code, no explanations or examples

- Do NOT include markdown code blocks (```python)
- Do NOT include any text before or after the function
- Do NOT include test cases or example usage
- Create a helper function named 'aux' that can assist the main function
- The function should return useful data for solving the problem

Your output should follow this format:
def aux(...):
 # your code here
 return result

Main Agent

Solve this coding problem by implementing the required function.
Problem: {prompt}
You have access to a helper function:
aux(...)

Instructions:

- Output ONLY the function code, no explanations or examples
- Do NOT include markdown code blocks (```python)
- Do NOT include any text before or after the function
- Do NOT include test cases or example usage
- Do NOT redefine the aux() function
- Implement ONLY the '{entry_point}' function as specified
- You can call aux() to assign a value to a variable within your function if helpful

Your output should follow this format:
def {entry_point}({params}):
 # your function code here
 return result

To improve the generated code, these prompts are used to construct second-turn observations for the MAS with suggestions from *Claude-Sonnet-4*.

External Agent

You are an advisor helping 2 agents (an auxiliary agent and a main agent) solve the following problem. The auxiliary agent provides a helper function (aux), while the main agent defines the task-specific logic.

Problem: {prompt}
Example tests: {test}
Show your feedback and edits for the following code: {combined_code}

Instructions:

- If you identify a missing element, such as an undefined aux or missing entry point (main function), you should write one for it.
- If both are not missing, point out and

make changes to any critical syntax or logic errors that would prevent the code from passing the given unit tests.

- You should focus only on clear errors on the given unit tests, be conservative and lenient, ignoring issues like redundancy, inefficiency, lack of edge case handling, or type annotations.
- Return "Perfect! No edits needed!" if the logic is sound.

Your response MUST contain the JSON format specified below. Always include both 'aux' and 'main' fields, even if no edits are needed.

```
{ "aux": {...}, "main": {...} }
```

Baseline Methods

We adopt a single-agent method and 3 representative multi-agent conversation methods as baselines.

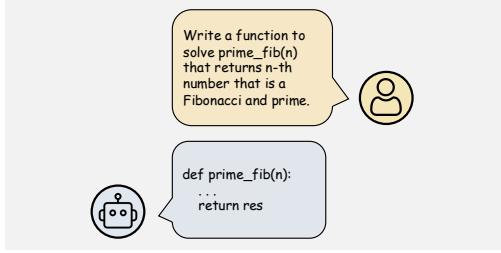


Figure 6: Single-agent code generation.

Figure 6 illustrates the code generation process using a single LLM agent. In this setting, the user gives a coding question along with specific instructions. The agent responds by generating a Python function snippet to solve it.

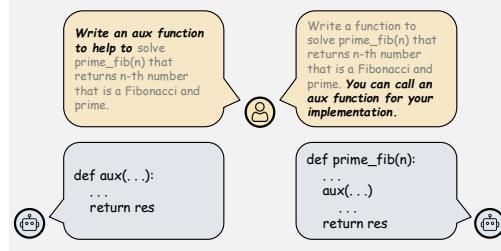


Figure 7: Coding collaboration via naive concatenation.

Naive concatenation represents the simplest form of cooperation, where 2 agents generate code synchronously, as illustrated in Figure 7. The first agent is provided with the coding question and informed of its role as a helper. The second agent is given the same question, along with its role as the main generator and the fact that an auxiliary agent exists.

However, the main agent lacks information about the specific functionality of the auxiliary agent. Their outputs are directly concatenated to form the response. This method is intended to improve inference efficiency by enabling a simple division of the problem into separate parts.

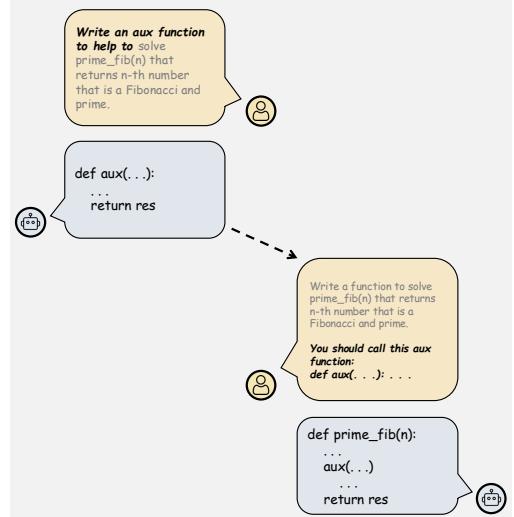


Figure 8: Coding collaboration via sequential pipeline.

Figure 8 presents the form of pipeline cooperation, where agents respond in sequence. The first agent is given the coding question along with the role of a helper. Its response is then passed to the main agent as a reference. The main agent generates its answer by incorporating the helper's response. This method enables one-way communication, allowing the main agent to respond by coordinating with the helper. However, this comes at the cost of slower inference speed due to the sequential nature of the interaction.

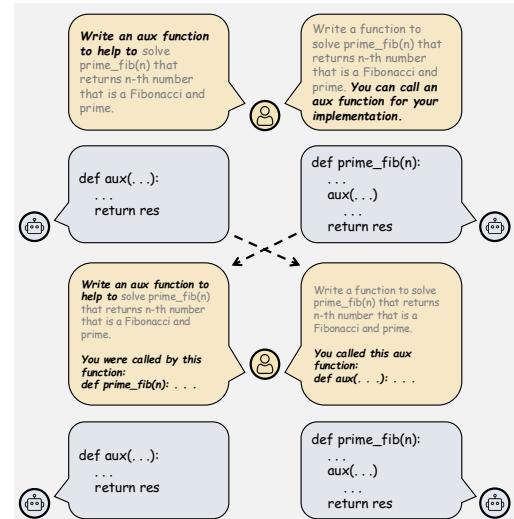


Figure 9: Coding collaboration with one-round discussion.

Discussion or debate frameworks (Figure 9) aim to improve response quality by enabling agents to access each other’s previous outputs (Du et al. 2023; Liang et al. 2024). In the first turn, the helper and main agents generate responses in the same manner as the naive concatenation approach. These initial responses are then shared with each other as references for the following turns, forming a discussion. Although this setup introduces more interaction, it does not guarantee improved response quality. With a limited number of rounds, the final output may not converge to a coherent solution. Even with additional rounds, convergence remains uncertain. This approach can even be less efficient than the sequential pipeline, particularly in distributed systems where communication latency is high or unreliable.

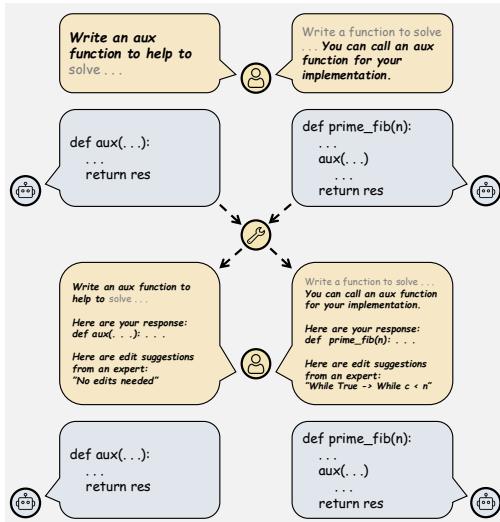


Figure 10: Coding collaboration in our method.

The interaction process between 2 agents trained with MAGRPO is illustrated in Figure 10. In the single-turn setting, we use the same prompts as in the naive concatenation baseline. In the multi-turn setting, after the helper and main agents generate their initial responses, these outputs are reviewed by an external agent. In this work, we employ *Claude-Sonnet-4* as an external to provide edit suggestions. For each agent, the suggestions, as well as the prior information and their previous response, are incorporated into the prompt for the subsequent round.

Note that the baselines above can also be fine-tuned by MARL. However, the interactions among agents in these settings are not strictly cooperative, which may lead to instability during training. To address this, techniques such as role-based rewards (Liu et al. 2025a), partial MAS training (Estornell et al. 2025), and freezing selected agents (Subramaniam et al. 2025) can be employed to ensure stable optimization.

Broader Impacts

Prompt-based coordination is often brittle (Estornell and Liu 2024), as agents may fail to follow instructions they were not

explicitly trained to interpret. Our method builds on a solid theoretical foundation in cooperative MARL, explicitly optimizing agents for joint optimality. Our work also opens opportunities to enhance existing test-time multi-agent interaction methods by integrating MARL techniques (Du et al. 2023; Lifshitz, McIlraith, and Du 2025; Wu et al. 2023a), particularly in settings that involve task decomposition and iterative feedback integration.

This work also explores a new perspective on accelerating LLM inference through cooperative MARL. While mainstream acceleration techniques (e.g., knowledge distillation, pruning, and quantization) improve efficiency at the cost of information loss (Wang et al. 2024; Zhao et al. 2024), our approach suggests decentralized coordination among specialized agents, thereby alleviating the burden of long-context memory and joint decision-making on a single model. Each agent can focus on a specific subtask, enabling more modular and robust reasoning.

Limitations and Future Works

Nevertheless, this study is subject to several limitations:

1. First, we focus on homogeneous agents for simplicity, assuming they perform similar tasks despite being assigned different roles, e.g., both the auxiliary agent and main agent are generating Python functions. Future research could explore LLM collaboration among heterogeneous agents with diverse capabilities and functionalities.
2. Due to computational constraints, we train LLMs with MAGRPO on limited datasets using relatively small-scale language models in a short horizon. When LLM-based coding agents are deployed in larger-scale projects involving multiple files and modules, more diverse and complex cooperation schemes are likely to emerge, which would further demonstrate the potential of decentralized coordination in MAS.
3. The simplicity of our reward model inevitably leads to narrow reward signals and potential reward hacking. As suggested by many research studies and industrial practice (Uesato et al. 2022; Wu et al. 2023b; Anthropic 2023), designing more expressive and fine-grained reward models (e.g., multi-aspect rewards, process-supervised rewards) is essential for better aligning agent cooperation with human preferences.

Compute Resources

We use NVIDIA H200 and H100 GPUs for LLM fine-tuning, and a standalone NVIDIA GeForce RTX 5090 workstation for inference. Here are the specifications of the resources used for experiments.

Training Devices

Type: GPU Cluster
 CPU: Intel Xeon Platinum 8558
 GPU: 1x NVIDIA H200

Inference Device

Type: Standalone Workstation
 CPU: AMD Ryzen 9 9950X
 GPU: 2x NVIDIA GeForce RTX 5090

The training process requires approximately 5 hours to complete 700 steps and 8 hours to complete 1000 steps. The inference takes approximately 0.5 hours for each 15-generation evaluation. Noted that training duration may vary considerably due to the node condition and the stochasticity of LLM outputs.

Codes and Datasets

Codes We developed an open-source library, CoMLRL, for training multiple LLMs to collaborate using MARL. CoMLRL provides configurable implementations of various MARL algorithms for LLM collaboration, including MAGRPO. The experiments of this paper serve as parts of CoMLRL's environments and benchmarks.

- **OpenMLRL/CoMLRL:**
<https://github.com/OpenMLRL/CoMLRL>
- **OpenMLRL/LLM_Collab_Writing:**
https://github.com/OpenMLRL/LLM_Collab_Writing
- **OpenMLRL/LLM_Code_Generation:**
https://github.com/OpenMLRL/LLM_Collab_Code_Generation
- **OpenMLRL/LLM_Code_Completion:**
https://github.com/OpenMLRL/LLM_Collab_Code_Completion

We hope that this library and the associated experiment repositories will lay a foundation and provide convenient tools for future research and development in LLM collaborations.

Datasets The datasets used in our experiments are all open-sourced and available online,

- **trl-lib/TLDR:**
<https://huggingface.co/datasets/trl-lib/tldr>
- **mattbierbaum/arXiv:**
<https://github.com/mattbierbaum/arxiv-public-datasets>
- **OpenAI/humanevalHumanEval:**
https://huggingface.co/datasets/openai/openai_humaneval
- **OpenMLRL/CoopHumanEval:**
<https://huggingface.co/datasets/OpenMLRL/CoopHumanEval>