

# AMATA: Adaptive Multi-Agent Trajectory Alignment for Knowledge-Intensive Question Answering

Taolin Zhang<sup>1</sup>, Dongyang Li<sup>2</sup>, Chen Chen<sup>4</sup>, Qizhou Chen<sup>3</sup>, Jiuheg Wan<sup>1</sup>, Xiaofeng He<sup>3</sup>, Chengyu Wang<sup>5\*</sup>, Richang Hong<sup>1\*</sup>

<sup>1</sup> School of Computer Science and Information Engineering, Hefei University of Technology

<sup>2</sup> Shanghai University of Electric Power <sup>3</sup> East China Normal University

<sup>4</sup> Guangdong University of Finance and Economics <sup>5</sup> Alibaba Group

tlzhang@hfut.edu.cn, chengyu.wcy@alibaba-inc.com

## Abstract

Despite substantial advances in large language models (LLMs), generating factually consistent responses for knowledge-intensive question answering remains challenging. These difficulties are primarily due to hallucinations and the limitations of LLMs in bridging long-tail knowledge gaps. To address this, we propose *AMATA*, an Adaptive Multi-Agent Trajectory Alignment framework that dynamically integrates external knowledge to improve response interpretability and factual grounding. Our architecture leverages six specialized agents that collaboratively perform structured actions for complex question reasoning. We formalize multi-agent collaboration with external tools as a trajectory preference alignment problem, incorporating question-aware agent customization and inter-agent preference harmonization. *AMATA* introduces two principal innovations: (1) *Intra-Trajectory Preference Learning*, which learns objective-oriented preferences to prioritize critical agents, and (2) *Inter-Agent Dependency Learning*, which captures cross-agent tool dependencies through a novel dependency-aware direct preference optimization technique. Empirical results show that *AMATA* consistently outperforms baseline approaches, knowledge-augmented frameworks, and LLM-based trajectory systems on five established knowledge-intensive QA benchmarks. Further analysis demonstrates the efficiency of our method in reducing token consumption.

## 1 Introduction

Large language models (LLMs) serve as the backbone of modern NLP infrastructure (Hu et al., 2024), yet they face persistent reliability challenges. Chief among these are hallucinations that appear superficially plausible (Woo et al., 2025) and other undesirable behaviors (Yang et al., 2025).

Retrieval-Augmented Generation (RAG) provides a mitigation strategy by enabling LLMs to

dynamically retrieve up-to-date information from external knowledge sources during inference (Rubin et al., 2022; Xu et al., 2024b). However, RAG incurs its own drawbacks, such as retrieval inaccuracies (Zhu et al., 2024) and increased inference latency due to longer contexts (Zou et al., 2024). As a result, research has increasingly shifted towards multi-agent systems that leverage diverse tooling and cooperative reflection mechanisms to enhance task robustness (Xu et al., 2024a; Yue et al., 2025).

Prior multi-agent approaches for knowledge-intensive QA can be broadly categorized into three paradigms: *Modular Training*, *End-to-End Training*, and *Global-Local Training*. (1) *Modular Training* fine-tunes individual agents on customized datasets targeted to their respective subtask capabilities (Long et al., 2023; Koopman et al., 2024). These locally optimized agents are then manually integrated through static workflows, with execution sequences and handoff mechanisms rigidly defined. The lack of global optimization leads to error propagation and diminished overall performance, as locally optimal agents cannot compensate for inter-agent dependencies or system-wide dynamics. (2) *End-to-End Training* adopts a unified optimization strategy, jointly training all agents within a single framework using task-level supervision (Zong et al., 2024; Klisura et al., 2025). While backpropagation updates parameters throughout the agent ensemble, enabling co-adaptation and implicit coordination, uniform gradient updates obscure the distinct specialization demands of heterogeneous agents. Agents responsible for different subtasks (Khot et al., 2023; Yue et al., 2024) require individualized learning signals; parameter sharing can cause over-homogenization, thereby eroding specialized expertise. (3) *Global-Local Training* employs a two-stage process: first, agents are optimized independently for subtask proficiency, then fine-tuned jointly to align global behaviors (Yue et al., 2025). Although this strategy combines lo-

\* C. Wang and R. Hong are co-corresponding authors.

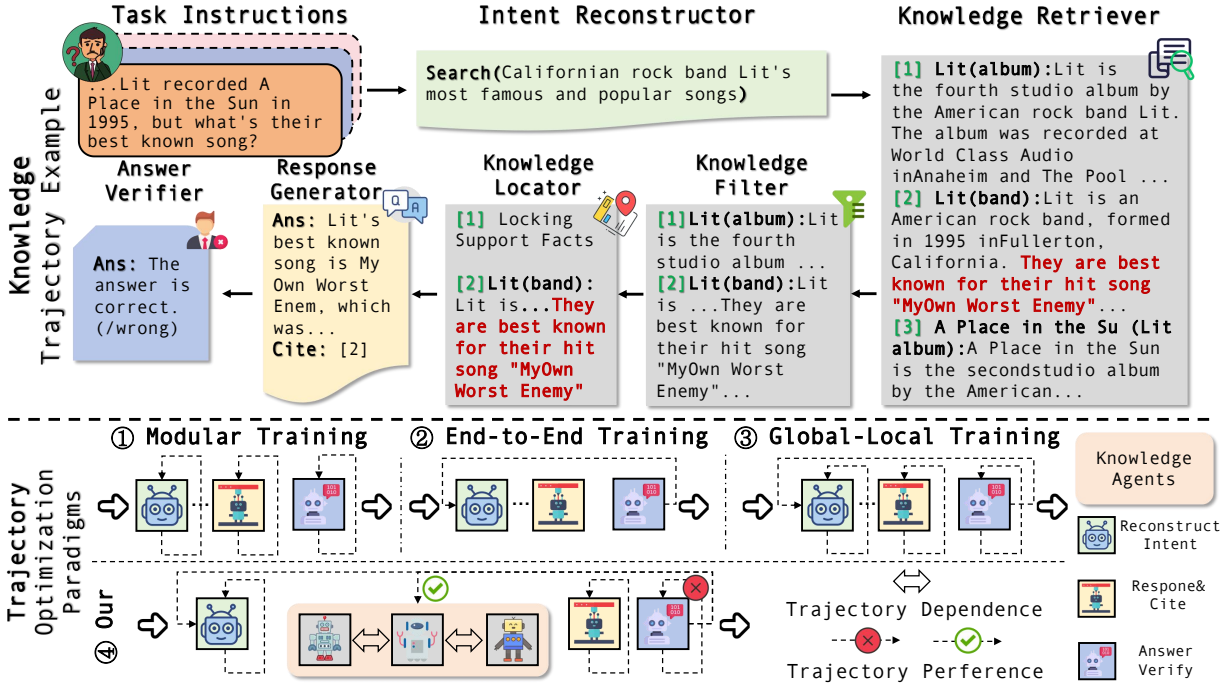


Figure 1: Comparison of multi-agent paradigms for knowledge-intensive QA. *Modular Training* and *End-to-End Training* focus, respectively, on local agent reasoning and global co-adaptation, while the *Global-Local Training* paradigm combines the advantages of both. Our *AMATA* framework dynamically learns intra-trajectory relationships between questions and LLM agents, and establishes inter-agent dependencies across the agent ensemble.

calized specialization with global coordination, it often fails to capture dynamic inter-agent dependencies, which are crucial for handling diverse knowledge-intensive tasks (Zhang et al., 2025a,b). As shown in Figure 1, for questions with high confidence, adding a “Verifier” agent after the preceding five agents may be unnecessary. Additionally, the three knowledge agents (“Retriever”, “Filter”, and “Locator”) exhibit strong interdependence; when the “Retriever” is triggered, subsequent actions of the other two agents must also be executed.

In this paper, we propose **Adaptive Multi-Agent Trajectory Alignment (AMATA)**, a framework designed to improve agent-level alignment and capture dynamic inter-agent dependencies. AMATA maintains high reasoning performance while significantly reducing token overhead during inference. Our main contributions are summarized below:

**Intra-Trajectory Preference Learning.** Existing approaches commonly treat all agents as uniformly relevant throughout the reasoning process. In contrast, our method dynamically optimizes agent participation for each question, learning question-specific agent preference distributions that adaptively modulate each agent’s influence based on utility for the current input. For each agent and question pair, we assign a preference score (e.g.,

$\langle \text{Reconstructor: } 5 \rangle$  versus  $\langle \text{Verifier: } 1 \rangle$ ), concatenated with the agent description as a prefix and paired with the question for fine-tuning the corresponding agent.

**Inter-Agent Dependency Learning.** Multi-agent systems exhibit context-dependent inter-agent dependencies, where triggering a pivotal agent necessitates coordinated execution of functionally linked agents, while allowing conditional suppression of unrelated agents (Ji and Gao, 2024; Gao et al., 2025). We introduce a dependency-aware Direct Preference Optimization (DA-DPO) module that learns context-sensitive execution ranking. Specifically, for each question, we construct preference samples that explicitly encode inter-agent dependencies and annotate each trajectory with a joint preference score reflecting the global optimality of the multi-agent sequence. These scores induce a dependency-aware ranking over sampled trajectories, prioritizing those with robust inter-agent coordination for DA-DPO training. This mechanism enables LLMs to infer optimal multi-agent execution sequences with high reliability.

We evaluate *AMATA* against competitive baselines on five benchmarks: HealthQA (Akhtar et al., 2022), ARC-Choice (Clark et al., 2018), PopQA (Mallen et al., 2022), SQuAD 1.1 (Ra-

jpurkar et al., 2016), and ASQA (Gao et al., 2023). Our framework achieves an average performance improvement of **+4.02%** across all tasks, and reduces token consumption overhead by approximately **70%** compared to strong baselines.

## 2 Related Work

**Multi-Agent Trajectory Learning.** Multi-agent trajectory learning refers to the process of orchestrating multiple agents to collectively solve complex tasks (Li, 2025). Existing literature can be grouped into three main paradigms: (1) *Modular Training* (Long et al., 2023; Koopman et al., 2024) trains agents independently for their respective sub-tasks. This often results in suboptimal global performance due to a lack of system-wide coordination. Preference learning has been introduced to ameliorate poor decision-making at the individual agent level (Song et al., 2024; Xiong et al., 2024), but overall integration remains a challenge. (2) *End-to-End Training* (Zong et al., 2024; Klisura et al., 2025) jointly optimizes all agents using unified loss functions derived from expert trajectories curated by a teacher LLM (e.g., FireAct (Chen et al., 2023), AgentTuning (Zeng et al., 2024)). Other approaches such as MapGPT (Chen et al., 2024) and LLM-A\* (Meng et al., 2024) focus on providing agents with a global view of the environment. Despite improved coordination, these methods may obscure contributions from specialized agents, potentially hindering the balance between individual expertise and system-wide collaboration (Zhang et al., 2025b). (3) *Global-Local Training* combines both global context and local adaptation signals to enhance agent specialization. For instance, CoAct (Hou et al., 2024) emulates hierarchical human planning in LLMs, while SMART (Yue et al., 2025) leverages multi-granular trajectories for agent control and system synergy. These frameworks inject both global and local signals into agent optimization, aiming to preserve both broad task alignment and agent-level differentiation (Subramonian et al., 2023). However, these methods often overlook inter-agent dependencies.

**Knowledge Enhancement for LLMs.** LLMs are prone to hallucinations and lack coverage of long-tail knowledge due to their parametric nature (Ji et al., 2023; Li et al., 2024; Huang et al., 2025). To ensure factual accuracy, LLMs frequently rely on external sources. RAG incorporates non-parametric resources to improve factual reli-

Agent	Head	End
Intent Reconstructor $\mathcal{A}_{IR}$	(Reconstructor)	$\langle /eoi \rangle$
Knowledge Retriever $\mathcal{A}_{KR}$	(Retriever)	$\langle /eor \rangle$
Knowledge Filter $\mathcal{A}_{KF}$	(Filter)	$\langle /eof \rangle$
Knowledge Locator $\mathcal{A}_{KL}$	(Locator)	$\langle /eol \rangle$
Response Generator $\mathcal{A}_{RG}$	(Generator)	$\langle /eog \rangle$
Answer Verifier $\mathcal{A}_{AV}$	(Verifier)	$\langle /eov \rangle$

Table 1: Agents and special tokens used in trajectories. Detailed agent descriptions and the trajectory data collection process are provided in Appendices A.1 and A.4.

ability and enrich LLM outputs (Fan et al., 2024; Singh et al., 2025). Advancements in this area include better retrieval mechanisms using dense retrievers (Karpukhin et al., 2020; Ye et al., 2024) and improved information integration (Zhang et al., 2024; Wang et al., 2025b; Cheng et al., 2025). For example, Self-RAG (Asai et al., 2024) introduces reflection tokens to assess both retrieval and response quality during inference. However, these RAG techniques typically operate within a single-agent paradigm, executing retrieval and generation in a sequential pipeline (Singh et al., 2025), thus failing to exploit the emergent capabilities and cooperative reasoning potential of multi-agent LLM frameworks (Qian et al., 2025). In contrast, *AM-ATA* is designed for knowledge-intensive QA tasks in a multi-agent setting.

## 3 Methodology

### 3.1 Task Formulation and Basic Notations

Figure 2 illustrates the overall architecture of *AM-ATA*. Given a question  $\mathcal{Q}$ , we design a workflow utilizing an LLM-based multi-agent system to generate the answer  $\mathcal{Y}$ , where  $\mathcal{Y} = \mathcal{F}(\mathcal{Q}, \mathbf{A})$ . Here,  $\mathcal{F}$  denotes the entire system parameterized by learnable weights, and  $\mathbf{A}$  is the set of agents, such as the six agents in *AMATA* (see Table 1). In this framework, each agent  $\mathcal{A} \in \mathbf{A}$  receives the current state and produces three outputs: a response  $y_i$ , a special end token  $e_i$ , and a special head token  $h_{i+1}$  for the subsequent agent, expressed as:

$$y_i, e_i, h_{i+1} = \mathcal{A}(\mathcal{Q}, y_{i-1}, e_{i-1}, h_i), \quad (1)$$

where  $\mathcal{T} = \{(h_1, y_1, e_1), \dots, (h_T, y_T, e_T)\}$  denotes a complete trajectory realized by dynamically executing the workflow  $\mathcal{F}$ . The final output  $\mathcal{Y}$  is obtained after completing this trajectory. In LLM-based multi-agent systems, conditional autoregressive language modeling is typically adopted to learn which agent should act and when, utilizing

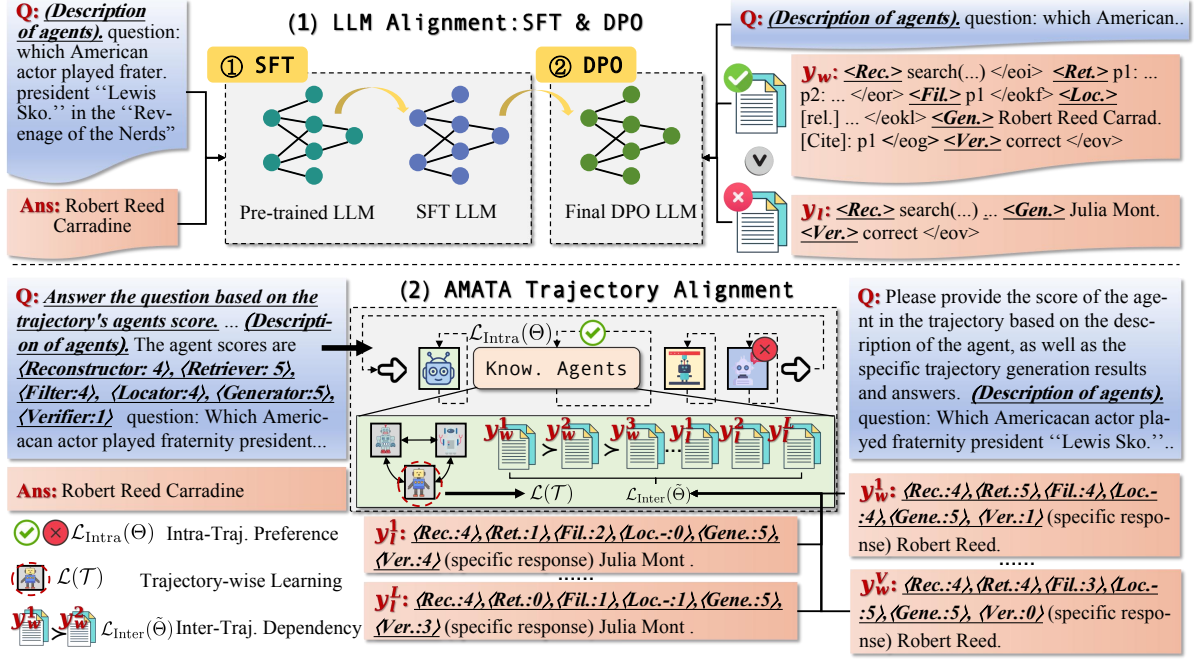


Figure 2: Comparison between AMATA and standard SFT and DPO pipelines. AMATA optimizes intra-trajectory preferences and inter-agent dependencies through adaptive prefix scoring (left) and DA-DPO (right).

these special tokens to coordinate agent behaviors (Kwon et al., 2024; Tang et al., 2025; Yue et al., 2025). The trajectory-wise objective function is defined as  $\mathcal{L}(\mathcal{T}) = \sum_{i=1}^T -\log \Pr(t_i | t_{<i}, \mathcal{Q})$ , where  $t_i = (h_i, y_i, e_i)$  represents the  $i$ -th tuple in the trajectory  $\mathcal{T}$  and  $t_{<i}$  comprises all preceding tuples in the trajectory.

### 3.2 Intra-Trajectory Preference Learning

Agents in a multi-agent system exhibit heterogeneous capabilities, necessitating autonomous tool usage and adaptive coordination. For example, for simple questions, the workflow may not require external knowledge retrieval or output verification (e.g.,  $\mathcal{A}_{AV}$ ). In such cases, the workflow only formalizes the question via  $\mathcal{A}_{IR}$  and uses the generator to produce an answer via  $\mathcal{A}_{RG}$ . To model agent-specific tool usage within a trajectory, a common approach is supervised fine-tuning (SFT), which enhances the tool-handling skills of individual agents:

$$\mathcal{L}_{\text{SFT}}^{(j)}(\Theta) = -\mathbb{E}_{(\mathcal{Q}, \mathcal{Y}) \sim \mathcal{D}_{\text{intra}}^{(j)}} \log \Pr(\mathcal{Y} | \mathcal{Q}; \Theta), \quad (2)$$

where  $\mathcal{D}_{\text{intra}}^{(j)}$  is the subset of the intra-trajectory dataset  $\mathcal{D}_{\text{intra}}$  corresponding to the  $j$ -th agent, and  $\Theta$  denotes intra-trajectory model parameters. The agent’s description and functionalities are incorporated into  $\mathcal{Q}$  via prompt engineering.

While agent-specific training can be effective,

it requires both agent-specific datasets and significant computational resources, limiting scalability. To address this, we consolidate preference learning within a unified trajectory sampling framework, enabling a monolithic model to capture heterogeneous agent competencies through SFT augmented with adaptive prefix scoring:

$$\mathcal{L}_{\text{Intra}}(\Theta) = -\mathbb{E}_{(\mathcal{Q}, \mathcal{Y}, \mathcal{P}) \sim \mathcal{D}_{\text{intra}}} \log \Pr(\mathcal{Y} | \mathcal{P}, \mathcal{Q}; \Theta) \quad (3)$$

where  $\mathcal{P} = \{P_{A_{IR}}, \dots, P_{A_{AV}}\}$  represents the preference prefix set for each agent in the trajectory, indicating their relative importance for a given sample (e.g., <Retriever: 5>). These scores reflect the importance of each agent in correctly answering the question, as annotated by LLMs.<sup>1</sup>

Consider the intra-trajectory sample illustrated in Fig. 3. The reference to “Revenge of the Nerds” requires background knowledge to support the LLM’s response. Accordingly, the <Retriever>, <Filter>, and <Locator> agents receive higher preference scores, reflecting a stronger need for knowledge tools. When knowledge agents yield high-confidence outputs, verification becomes redundant, resulting in a low preference score for the <Verifier> agent. This example demonstrates that question complexity induces a dynamic, heteroge-

<sup>1</sup>Prompt templates are provided in Appendix A.1.

neous agent hierarchy within the reasoning trajectory, and underscores the need for a framework that supports flexible, token-efficient tool orchestration.

### 3.3 Inter-Agent Dependency Learning

Functional dependencies naturally exist between agents within a trajectory, as downstream agents rely on responses and states generated by upstream agents in order to be triggered and operate effectively (Gao et al., 2025). In *AMATA*, such dependencies are especially pronounced: for instance, if the retrieval agent  $\mathcal{A}_{KR}$  is activated, then the knowledge filtering agent  $\mathcal{A}_{KF}$  and the locating agent  $\mathcal{A}_{KL}$  must also be executed. To model these interactions, we introduce inter-agent dependency learning, leveraging pairs of winning and losing samples for Direct Preference Optimization (DPO) (Rafailov et al., 2023). This enables the model to automatically discover optimal patterns of agent collaboration and to capture the underlying dependency structures among agents.

A widely adopted technique combines the Bradley–Terry (BT) model (Bradley and Terry, 1952) with DPO to parameterize the reward function for trajectory selection. Formally, the probability that the winning response  $y_w$  is preferred over the losing response  $y_l$  for a given instruction  $\mathcal{Q}$  is:

$$\Pr(y_w \succ y_l \mid \mathcal{Q}) = \frac{\exp(r(\mathcal{Q}, y_w))}{\exp(r(\mathcal{Q}, y_w)) + \exp(r(\mathcal{Q}, y_l))} \quad (4)$$

where the reward  $r(\mathcal{Q}, y)$  measures the policy model’s preference for  $y$  and is given by  $r(\mathcal{Q}, y) = \beta \cdot \log \frac{\pi_{\tilde{\Theta}}(y \mid \mathcal{Q})}{\pi_{\text{ref}}(y \mid \mathcal{Q})} + \beta \cdot \log Z(\mathcal{Q})$ , with  $\pi_{\text{ref}}$  and  $\pi_{\tilde{\Theta}}$  denoting the intra-trajectory and inter-agent dependency models, respectively. The coefficient  $\beta$  modulates the strength of regularization, and  $Z(\mathcal{Q}) = \sum_y \pi_{\text{ref}}(y \mid \mathcal{Q}) \exp\left(\frac{1}{\beta} r(\mathcal{Q}, y)\right)$  denotes the partition function.

However, merely distinguishing between winning and losing samples is insufficient for optimal dependency modeling. Our analysis (see Sect. 4.3) reveals that specific combinations of agent preference scores in the trajectory prefix are correlated with higher response quality, compared to settings that neglect agent dependency information. For example, in the inter-trajectory samples presented in Fig. 3, although both  $y_w^1$  and  $y_w^2$  yield correct responses,  $y_w^1$  should receive a higher trajectory policy preference due to its consistently elevated scores for tightly coupled agents (such as  $\langle \text{Retriever} \rangle$ ,  $\langle \text{Filter} \rangle$ , and  $\langle \text{Locator} \rangle$ ).

Moreover, both of these trajectories outperform all losing samples. To address this, we propose the *dependency-aware DPO* (DA-DPO) algorithm in *AMATA*, guiding the policy model to better capture relative dependency relationships among agent preference scores.

Specifically, given  $M$  winning and  $N$  losing samples for a question  $\mathcal{Q}$ , we first select the top- $K$  winning samples based on their *dependency scores*, which measure the “goodness” of dependency among agent preference scores as determined by prior knowledge. We treat the remaining  $M - K$  winning samples as losing ones due to their lower “goodness” of dependencies. These samples are denoted as  $(y_w^1, \dots, y_w^K, y_w^{K+1}, \dots, y_w^M)$  and  $(y_l^{M+1}, \dots, y_l^{M+N})$ , where  $(y_w^{K+1}, \dots, y_w^M)$  and  $(y_l^{M+1}, \dots, y_l^{M+N})$  are treated as losing samples. Inspired by listwise Plackett-Luce preference modeling (Plackett, 1975), we define the inter-agent dependency model as follows:

$$\begin{aligned} & \Pr\left(y_w^1 \succ y_w^2 \succ \dots \succ y_w^K \succ \{y_w^{K+1}, \dots, y_l^{M+N}\} \mid \mathcal{Q}\right) \\ &= \sum_{f_{K+1}^{M+N}} \prod_{i=1}^{M+N-1} f_{\mathcal{E}}(\mathcal{Q}, y_i) \\ &= \prod_{i=1}^K f_{\mathcal{E}}(\mathcal{Q}, y_i) \cdot \sum_{f_{K+1}^{M+N}} \prod_{i=K+1}^{M+N-1} f_{\mathcal{E}}(\mathcal{Q}, y_i) \\ &= \prod_{i=1}^K f_{\mathcal{E}}(\mathcal{Q}, y_i) \cdot \sum_{f_{K+1}^{M+N}} \Pr(y_{K+1} \succ \dots \succ y_{M+N} \mid \mathcal{Q}) \\ &= \prod_{i=1}^K f_{\mathcal{E}}(\mathcal{Q}, y_i) \end{aligned} \quad (5)$$

where  $f_{K+1}^{M+N}$  denotes the set of all permutations of  $(y_{K+1}, \dots, y_{M+N})$  and  $f_{\mathcal{E}}(\mathcal{Q}, y_i) = \frac{\exp(r(\mathcal{Q}, y_i))}{\sum_{j=i}^{M+N} \exp(r(\mathcal{Q}, y_j))}$ . The set  $\{y_w^{K+1}, \dots, y_l^{M+N}\}$  denotes the rejected trajectory set for inter-agent dependency learning, including  $(M - K)$  original winning samples and  $N$  losing samples.

By substituting the reward from Eq. (7) into the probability maximization objective  $\Pr(y_w^1 \succ \dots \succ y_w^K \succ \{y_w^{K+1}, \dots, y_l^{M+N}\} \mid \mathcal{Q})$ , we obtain the objective of the inter-agent dependency loss:

$$\begin{aligned} \mathcal{L}_{\text{Inter}}(\tilde{\Theta}) &= -\mathbb{E}_{(\mathcal{Q}, y_w^1, \dots, y_l^L) \sim \mathcal{D}^{\text{inter}}} \mathcal{H}(\mathcal{Q}, y) \quad (6) \\ \mathcal{H}(\mathcal{Q}, y) &= \sum_{i=1}^K \log \sigma \left( -\log \sum_{j=i+1}^{M+N} \exp \mathcal{V}_{\beta} \right) \end{aligned} \quad (7)$$

where  $\mathcal{V}_{\beta} = \beta \log \frac{\pi_{\tilde{\Theta}}(y_j \mid \mathcal{Q})}{\pi_{\text{ref}}(y_j \mid \mathcal{Q})} - \beta \log \frac{\pi_{\tilde{\Theta}}(y_i \mid \mathcal{Q})}{\pi_{\text{ref}}(y_i \mid \mathcal{Q})}$ , and

Intra-Trajectory Preference Learning Example	Inter-Agent Dependency Learning Example
<p><b>Instruction:</b> Answer the question based on the agent preference scores. The higher the score, the more likely the agent needs to be invoked to answer the question. Then, output the answer to the following question. <b>(The description of the agents).</b> The agent scores are <code>&lt;Reconstructor:4&gt;</code>, <code>&lt;Retriever:5&gt;</code>, <code>&lt;Filter:4&gt;</code>, <code>&lt;Locator:4&gt;</code>, <code>&lt;Generator:5&gt;</code>, <code>&lt;Verifier:1&gt;</code>.</p> <p>Question: Which American actor played fraternity president "Lewis Skol." in the "Revenge of the Nerds" comedy films?</p> <p><b>Output:</b> Robert Reed Carradine.</p>	<p><b>Instruction:</b> Please provide the score of the agents in the trajectory based on the description of them, as well as the specific trajectory generation results and answers. (Description of agents). Question: Which American actor played fraternity president ...?</p> <p>👉 <b>Winning Sample</b> <math>y_w^1</math>: <code>&lt;Reconstructor:4&gt;</code>, <code>&lt;Retriever:5&gt;</code>, <code>&lt;Filter:4&gt;</code>, <code>&lt;Locator:4&gt;</code>, <code>&lt;Generator:5&gt;</code>, <code>&lt;Verifier:1&gt;</code> (Add specific trajectory). Robert Reed Carradine. ⬇</p> <p>👉 <b>Winning Sample</b> <math>y_w^2</math>: <code>&lt;Reconstructor:4&gt;</code>, <code>&lt;Retriever:4&gt;</code>, <code>&lt;Filter:3&gt;</code>, <code>&lt;Locator:5&gt;</code>, <code>&lt;Generator:5&gt;</code>, <code>&lt;Verifier:0&gt;</code> (Add specific trajectory). Robert Reed Carradine. ⬇</p> <p>👎 <b>Losing Sample</b> <math>y_l^1</math>: <code>&lt;Reconstructor:4&gt;</code>, <code>&lt;Retriever:1&gt;</code>, <code>&lt;Filter:2&gt;</code>, <code>&lt;Locator:0&gt;</code>, <code>&lt;Generator:5&gt;</code>, <code>&lt;Verifier:4&gt;</code> (Add specific trajectory). Julia Mont. ⬇</p> <p>👎 <b>Losing Sample</b> <math>y_l^2</math>: <code>&lt;Reconstructor:4&gt;</code>, <code>&lt;Retriever:0&gt;</code>, <code>&lt;Filter:1&gt;</code>, <code>&lt;Locator:1&gt;</code>, <code>&lt;Generator:5&gt;</code>, <code>&lt;Verifier:3&gt;</code> (Add specific trajectory). Julia Mont. ⬇</p>

Figure 3: Two-stage training examples of AMATA. Detailed annotation process and robustness verification for the **DEPENDENCY SCORES** are provided in Appendix A.1.2.

Task → Model ↓	HealthQA Acc.	ARC-C Acc.	PopQA Acc.	Squad1 Acc.	Str_EM	ASQA Rouge-L	Mauve	Average
<b>Vanilla QA Methods</b>								
Alpaca2 7B	44.78 <sub>(±1.2)</sub>	36.43 <sub>(±1.5)</sub>	25.58 <sub>(±0.8)</sub>	11.50 <sub>(±1.1)</sub>	14.42 <sub>(±1.6)</sub>	28.72 <sub>(±2.1)</sub>	51.24 <sub>(±0.9)</sub>	30.38 <sub>(±1.1)</sub>
Mistral-Instruct 7B	65.45 <sub>(±1.4)</sub>	57.84 <sub>(±0.7)</sub>	22.37 <sub>(±1.3)</sub>	14.97 <sub>(±1.3)</sub>	20.80 <sub>(±2.2)</sub>	32.20 <sub>(±0.9)</sub>	33.47 <sub>(±1.8)</sub>	35.30 <sub>(±0.8)</sub>
Llama-2-Chat 7B	47.95 <sub>(±1.9)</sub>	47.95 <sub>(±1.1)</sub>	25.44 <sub>(±0.8)</sub>	14.13 <sub>(±0.7)</sub>	16.79 <sub>(±2.3)</sub>	32.35 <sub>(±1.6)</sub>	24.21 <sub>(±1.2)</sub>	29.83 <sub>(±1.5)</sub>
Vicuna-v1.5 13B	63.01 <sub>(±2.0)</sub>	57.59 <sub>(±0.9)</sub>	17.94 <sub>(±1.5)</sub>	15.25 <sub>(±1.8)</sub>	31.95 <sub>(±2.2)</sub>	22.99 <sub>(±1.7)</sub>	68.41 <sub>(±1.4)</sub>	39.59 <sub>(±1.3)</sub>
Llama-2-Chat 13B	62.20 <sub>(±1.8)</sub>	48.72 <sub>(±2.2)</sub>	21.22 <sub>(±1.9)</sub>	15.97 <sub>(±1.4)</sub>	19.97 <sub>(±0.7)</sub>	30.37 <sub>(±1.3)</sub>	40.23 <sub>(±1.5)</sub>	34.10 <sub>(±1.7)</sub>
Qwen-2.5-Ins. 7B	64.02 <sub>(±1.1)</sub>	51.38 <sub>(±1.7)</sub>	22.35 <sub>(±0.7)</sub>	17.23 <sub>(±1.2)</sub>	18.99 <sub>(±1.3)</sub>	31.65 <sub>(±2.0)</sub>	47.03 <sub>(±1.1)</sub>	36.09 <sub>(±1.2)</sub>
GPT-3.5-turbo	76.08 <sub>(±1.5)</sub>	77.30 <sub>(±0.8)</sub>	29.30 <sub>(±1.2)</sub>	22.90 <sub>(±1.6)</sub>	39.94 <sub>(±1.4)</sub>	35.73 <sub>(±0.7)</sub>	44.63 <sub>(±2.3)</sub>	46.55 <sub>(±1.4)</sub>
<b>Knowledge-augmented Methods</b>								
Alpaca2 7B	26.44 <sub>(±1.7)</sub>	35.15 <sub>(±1.4)</sub>	33.38 <sub>(±1.6)</sub>	21.41 <sub>(±2.2)</sub>	23.59 <sub>(±0.8)</sub>	27.21 <sub>(±2.3)</sub>	50.09 <sub>(±1.5)</sub>	31.04 <sub>(±1.2)</sub>
REPLUG 7B	41.72 <sub>(±1.3)</sub>	47.26 <sub>(±0.8)</sub>	37.24 <sub>(±1.5)</sub>	24.23 <sub>(±1.7)</sub>	26.54 <sub>(±2.2)</sub>	33.25 <sub>(±0.9)</sub>	54.03 <sub>(±1.4)</sub>	37.75 <sub>(±1.7)</sub>
VANILLA 7B*	29.52 <sub>(±1.6)</sub>	42.74 <sub>(±2.2)</sub>	37.52 <sub>(±1.7)</sub>	25.92 <sub>(±1.4)</sub>	32.25 <sub>(±1.5)</sub>	34.93 <sub>(±0.8)</sub>	39.54 <sub>(±2.3)</sub>	34.63 <sub>(±1.1)</sub>
RADIT 7B	52.98 <sub>(±1.5)</sub>	62.10 <sub>(±1.7)</sub>	38.02 <sub>(±2.3)</sub>	23.86 <sub>(±0.8)</sub>	25.68 <sub>(±1.4)</sub>	15.99 <sub>(±1.6)</sub>	12.35 <sub>(±1.9)</sub>	33.00 <sub>(±1.2)</sub>
INTERACT 7B	65.45 <sub>(±0.8)</sub>	48.12 <sub>(±1.3)</sub>	41.31 <sub>(±1.6)</sub>	<u>31.52</u> <sub>(±1.4)</sub>	34.54 <sub>(±1.7)</sub>	35.51 <sub>(±2.2)</sub>	43.45 <sub>(±1.5)</sub>	42.84 <sub>(±0.9)</sub>
SelfRag 7B*	68.99 <sub>(±1.4)</sub>	65.52 <sub>(±0.6)</sub>	40.67 <sub>(±0.8)</sub>	22.39 <sub>(±1.3)</sub>	28.68 <sub>(±1.5)</sub>	34.11 <sub>(±1.7)</sub>	83.00 <sub>(±2.1)</sub>	49.05 <sub>(±1.8)</sub>
<b>LLM-based Trajectory Methods</b>								
MMAgent 7B	72.56 <sub>(±0.7)</sub>	64.43 <sub>(±1.3)</sub>	37.92 <sub>(±1.5)</sub>	24.62 <sub>(±0.8)</sub>	34.13 <sub>(±1.6)</sub>	37.25 <sub>(±1.2)</sub>	90.11 <sub>(±2.0)</sub>	51.57 <sub>(±1.3)</sub>
SMART 7B	73.90 <sub>(±1.6)</sub>	67.31 <sub>(±1.4)</sub>	42.88 <sub>(±0.8)</sub>	29.24 <sub>(±1.3)</sub>	42.56 <sub>(±1.7)</sub>	41.71 <sub>(±1.5)</sub>	92.32 <sub>(±2.2)</sub>	55.70 <sub>(±2.1)</sub>
SPA-RL 7B†	73.23 <sub>(±1.2)</sub>	<u>68.53</u> <sub>(±0.7)</sub>	42.72 <sub>(±2.1)</sub>	29.46 <sub>(±1.6)</sub>	43.73 <sub>(±0.8)</sub>	41.37 <sub>(±1.3)</sub>	91.02 <sub>(±1.5)</sub>	55.72 <sub>(±2.6)</sub>
GiGPO 7B†	73.97 <sub>(±1.5)</sub>	68.01 <sub>(±1.8)</sub>	<u>43.52</u> <sub>(±1.3)</sub>	29.97 <sub>(±1.7)</sub>	<u>43.88</u> <sub>(±1.4)</sub>	<u>43.64</u> <sub>(±1.1)</sub>	<u>92.83</u> <sub>(±1.2)</sub>	<u>56.55</u> <sub>(±2.1)</sub>
AMATA 7B	<b>75.83</b> <sub>(±0.8)</sub>	<b>72.47</b> <sub>(±1.6)</sub>	<b>47.39</b> <sub>(±1.2)</sub>	<b>34.61</b> <sub>(±1.5)</sub>	<b>49.10</b> <sub>(±1.7)</sub>	<b>48.26</b> <sub>(±1.3)</sub>	<b>96.35</b> <sub>(±1.4)</sub>	<b>60.57</b> <sub>(±1.1)</sub>

Table 2: Overall results of AMATA. Results of GPT-3.5-turbo are for reference only. \* indicates re-implemented methods based on the same model. † denotes the RL settings described in Appendix A.3.1. Results for other LLMs are shown in Appendix B.3. **Bold** numbers represent the best results, while underlined indicate the second-best.

$\mathcal{D}^{\text{inter}}$  and  $\tilde{\Theta}$  denote the inter-agent training samples and parameters, respectively.

### 3.4 Model Training and Inference

Our AMATA framework undergoes two-stage training. First, we perform intra-trajectory preference learning using  $\mathcal{L}_{\text{Intra}}(\Theta)$ , which enables the model to acquire varying degrees of perception regarding agent utilization within trajectories. Next, we combine the basic agent prediction loss  $\mathcal{L}(\mathcal{T})$  and the

inter-agent dependency loss  $\mathcal{L}_{\text{Inter}}(\tilde{\Theta})$  to form the total loss:

$$\mathcal{L}_{\text{total}} = \alpha_1 \cdot \mathcal{L}(\mathcal{T}) + \alpha_2 \cdot \mathcal{L}_{\text{Inter}}(\tilde{\Theta}), \quad (8)$$

thereby enhancing multi-agent cooperation in knowledge-intensive QA tasks. Here,  $\alpha_1$  and  $\alpha_2$  are training coefficients that sum to 1. Due to space constraints, we refer readers to Appendix C for inference algorithm details.

## 4 Experiments

We conduct extensive experiments to evaluate *AM-ATA*. Due to space limitations, details regarding trajectory data collection, baselines, and implementation are provided in Appendix A.

### 4.1 Main Results

As shown in Table 2, the key observations are as follows: (1) Compared to standard QA baselines, *AM-ATA* significantly outperforms models with comparable or even larger parameter sizes. Notably, it incorporates external knowledge through trajectory learning, compensating for the parameter-size gap relative to larger backbones such as Vicuna-v1.5 (13B) and Llama2-13B-Chat, especially on long-tail knowledge tasks (i.e., PopQA, SQuAD 1.1, and ASQA) in comparison to GPT-3.5-turbo. (2) Knowledge-augmented methods leverage retrievers to access external data and assist LLMs in generating informed answers; however, data noise and excessive augmentation can significantly degrade model performance (Fang et al., 2024). Notably, despite sharing the same training data and backbone as RADIT (Lin et al., 2024), our model achieves a higher fluency score as measured by *MAUVE*. (3) We also compare against LLM-based trajectory methods, including the independent-agent modular training method MMAgent (comprising six independent agents), the global-local trajectory approach SMART (Yue et al., 2025), and recent long-trajectory methods such as GiGPO (Feng et al., 2025) and SPA-RL (Wang et al., 2025a). The results indicate that long-trajectory methods generally outperform multi-agent training approaches in our setting. This improvement can be attributed to RL feedback, which mitigates the cumulative propagation of agent-action errors in longer trajectories. Our method further encourages both intra- and inter-agent dependencies within trajectories, thereby reducing collaborative conflicts among agents.

### 4.2 Ablation Study

We perform ablation studies on the critical modules involved in the training and inference processes of *AMATA*. As shown in Table 3, during training, removing  $\mathcal{L}_{\text{Inter}}$  prevents the modeling of associations between agents (e.g., knowledge agents and verifiers), resulting in unresolved collaborative conflicts and the steepest decline in model performance. Additionally, removing the basic tra-

Task → Model ↓	HealthQA Acc.	ARC-C Acc.	PopQA Acc.	ASQA Str_EM
<b>Training ablation</b>				
AMATA <sub>7B</sub>	75.83	72.47	47.39	49.10
w/o $\mathcal{L}_{\text{Intra}}$	72.91	70.05	44.32	45.28
w/o $\mathcal{L}_{\text{Inter}}$	70.86	67.57	41.13	42.94
w/o $\mathcal{L}_{\mathcal{T}}$	72.03	69.78	43.24	44.62
w/o $\mathcal{A}_{\text{IR}}$	73.38	70.50	45.41	46.95
w/o $\mathcal{A}_{\text{KF}}$	72.66	69.98	44.27	45.12
w/o $\mathcal{A}_{\text{AV}}$	73.13	69.34	45.20	45.85
<b>Inference ablation</b>				
w/o $\mathcal{A}_{\text{IR}}$	73.57	70.83	45.75	47.26
w/o $\mathcal{A}_{\text{KF}}$	72.99	70.15	44.89	45.37
w/o $\mathcal{A}_{\text{AV}}$	73.36	70.29	45.41	47.50

Table 3: Training and inference ablation of key trajectory learning modules and agents in *AMATA*.

jectory loss  $\mathcal{L}(\mathcal{T})$  causes a significant performance drop, as it impairs the semantic modeling of agent trajectories.

From the perspective of individual agents, removing  $\mathcal{A}_{\text{KF}}$  fails to effectively reduce noise in the data retrieved by  $\mathcal{A}_{\text{KR}}$ , leading to a marked decline in performance. Removing the verifier agent  $\mathcal{A}_{\text{AV}}$  eliminates the verification of generated trajectory answers, which may cause hallucinated outputs and subsequently degrade results. Due to variability in question content,  $\mathcal{A}_{\text{IR}}$  is essential for enabling semantic understanding of user queries; thus, its removal also adversely affects overall performance.

### 4.3 Detailed Analysis

In this section, we conduct an in-depth analysis of the adaptive cooperation among *AMATA* agents, elucidating its advantages in both performance and efficiency. Due to space limitations, **computational cost comparison**, **LLM backbones**, **hyperparameter analysis**, and **case studies** are presented in Appendix B.

**Agent Dependency Analysis.** We adopt different DPO methods to investigate whether comparing winning and losing examples can enhance dependencies among agents. Specifically, we compare our DA-DPO method with DPO (Rafailov et al., 2023) and full-order DPO (FDPO) (Rafailov et al., 2023). The DPO method utilizes only a single pair of winning and losing QA samples, while FDPO leverages the same number of samples as our method but ranks them using full-order learning derived from the magnitude of preference scores. Additionally, we compare our DA-DPO with global-local (SMART) and RL-based




Figure 4: Agent dependency analysis across preference methods. “DA-DPO”, “FDPO”, “GL”, and “RL” refer to our DA-DPO, full-order DPO, global-local trajectory, and reinforcement learning, respectively.




Figure 5: Average token consumption per question.

(GiGPO) approaches.

In Figure 4, our DA-DPO method consistently outperforms the other DPO-based approaches. This improvement is attributed to its fine-grained preference learning for trajectories in multi-agent collaboration, guided by preference scores. Subtle differences in scores for winning examples effectively reflect strong correlations between agents, while the inclusion of losing samples and winning examples with weaker scores helps distinguish less dependent relationships. Methods that do not explicitly account for pairwise preference data or that rely on FDPO-style comprehensive sorting tend to reduce overall model effectiveness. Regarding the global-local trajectory and RL-based approaches, their results exhibit greater consistency compared to DPO and FDPO. We hypothesize that this improvement arises because global-local fusion and RL step-level feedback enhance supervision of fine-grained agent dependencies (Du et al., 2024).

**Token Consumption.** In Figure 5, we compare the number of tokens processed by two base models (Llama2-7B and Qwen2.5-7B) when solving questions across different multi-agent frameworks. We observe that: (1) Although knowledge-augmented methods consume fewer tokens via a single en-

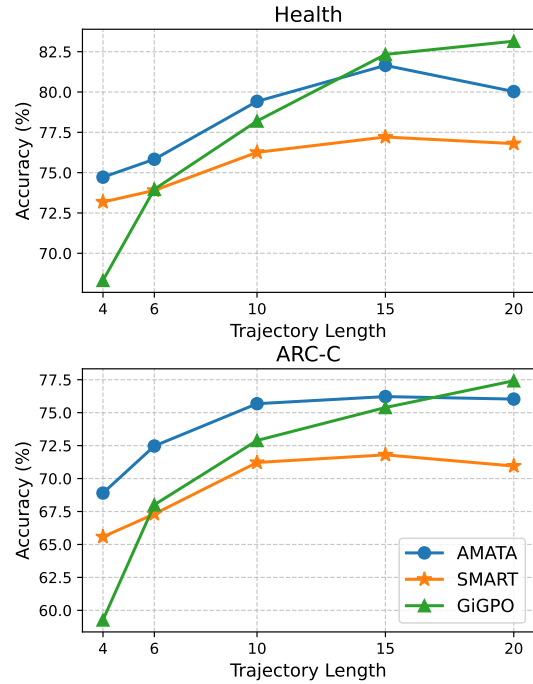


Figure 6: Performance of LLM-based trajectory methods relative to trajectory length.

hancement step, their performance remains modest, as shown in Table 2. (2) Compared with the global-local trajectory baseline SMART (Yue et al., 2025), although SMART involves fewer agent interactions, AMATA not only maintains comparable token overhead but also significantly outperforms it in terms of performance (+4.87%). This improvement stems from our adaptive agent preference learning, which enables AMATA to dynamically adjust interactions among knowledge agents (i.e.,  $\mathcal{A}_{KR}$ ,  $\mathcal{A}_{KF}$ , and  $\mathcal{A}_{KL}$ ) and other agents, greatly reducing token consumption. (3) By contrast, RL-based methods require multiple rollout sessions to compute advantages, resulting in the highest token consumption (+70%). Moreover, due to the absence of real-world step-level reward signals, these methods introduce noisy data to the knowledge agents, leading to reduced performance (-4.02%).

**Generalization by Trajectory Length.** Figure 6 compares the performance of LLM-based trajectory methods as trajectory length increases, i.e., as the number of agents in the multi-agent system grows. Constructing longer trajectories simulates multi-agent environments by introducing more agents to evaluate whether retrieved documents are properly filtered and located. We observe that when the number of agents is relatively small, our adaptive preference learning method and the global-local trajectory method outperform the

RL approach. We hypothesize that shorter trajectories facilitate simpler and more direct end-to-end supervised training. By contrast, RL-based methods require the computation of advantage values based on feedback at each step, which can hinder the timely propagation of feedback to local agents. However, as trajectory length increases, the RL-based method gradually improves results due to the cumulative effect of reward feedback and iterative rollout learning. Meanwhile, other methods suffer from excessive error accumulation, resulting in performance degradation. In this paper, knowledge QA tasks generally involve short trajectories (fewer than 10 steps), while GUI tasks typically involve long trajectories (more than 50 steps) (Yao et al., 2022).

## 5 Conclusion

We propose an adaptive multi-agent trajectory framework, *AMATA*, which enhances LLMs by effectively incorporating external knowledge to solve knowledge-intensive QA tasks. Our key innovations in intra-trajectory and inter-agent preference learning enable prioritization of critical agents and accurate modeling of cross-agent dependencies. Experiments on diverse knowledge-intensive QA benchmarks demonstrate the effectiveness and efficiency of our approach.

## Limitations

Despite the promising results achieved by *AMATA*, our work has several limitations that warrant further investigation. Due to constraints in computational resources, our experiments were conducted primarily on a 7B-parameter model. We anticipate that scaling up to larger models (e.g., 70B parameters or beyond) could further enhance performance, particularly for more complex knowledge-intensive tasks. Additionally, the number of winning and losing samples used in our dependency-aware DPO was limited to  $M = N = 10$ , which may not fully capture the diversity of inter-agent dependencies in more heterogeneous task settings. We set the Top- $K$  value to 5 for ranking winning examples, a conservative choice that balances efficiency and effectiveness but may overlook finer-grained preference structures. Future work will explore larger sample sizes and more adaptive ranking strategies as computational capacity increases.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 62506110). It was also supported by the Natural Science Foundation of Anhui Province, China (Grant No. 2508085QF227) and the Hefei University of Technology Scientific Research Innovation Start-up Special Project Type A (Grant No. JZ2025HGQA0137).

## References

- OpenAI. 2023. [Gpt-4 technical report](#). In *OpenAI*.
- Mubashara Akhtar, Oana Cocarascu, and Elena Simperl. 2022. [PubHealthTab: A public health table-based dataset for evidence-based fact checking](#). In *NAACL*.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). In *ICLR*.
- Ralph Allan Bradley and Milton E. Terry. 1952. [Rank analysis of incomplete block designs: I. the method of paired comparisons](#). *Biometrika*, 39(3/4):324–345.
- Baian Chen, Chang Shu, Ehsan Shareghi, Nigel Collier, Karthik Narasimhan, and Shunyu Yao. 2023. [Fireact: Toward language agent fine-tuning](#). *CoRR*, abs/2310.05915.
- Jiaqi Chen, Bingqian Lin, Ran Xu, Zhenhua Chai, Xiaodan Liang, and Kwan-Yee Kenneth Wong. 2024. [Mapgpt: Map-guided prompting with adaptive path planning for vision-and-language navigation](#). In *ACL*, pages 9796–9810.
- Mingyue Cheng, Yucong Luo, Jie Ouyang, Qi Liu, Huijie Liu, Li Li, Shuo Yu, Bohou Zhang, Jiawei Cao, Jie Ma, Daoyu Wang, and Enhong Chen. 2025. [A survey on knowledge-oriented retrieval-augmented generation](#). *CoRR*, abs/2503.10677.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the AI2 reasoning challenge](#). *CoRR*, abs/1803.05457.
- Huifang Du, Shuqin Li, Minghao Wu, Xuejing Feng, Yuan-Fang Li, and Haofen Wang. 2024. [Rewarding what matters: Step-by-step reinforcement learning for task-oriented dialogue](#). In *EMNLP*, pages 8030–8046.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev,

- Arthur Hinsvark, Arun Rao, Aston Zhang, and 82 others. 2024. [The llama 3 herd of models](#). *CoRR*, abs/2407.21783.
- Wenqi Fan, Yujuan Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. 2024. [A survey on RAG meeting llms: Towards retrieval-augmented large language models](#). In *KDD*, pages 6491–6501.
- Feiteng Fang, Yuelin Bai, Shiwen Ni, Min Yang, Xiaojun Chen, and Ruifeng Xu. 2024. [Enhancing noise robustness of retrieval-augmented language models with adaptive adversarial training](#). In *ACL*, pages 10028–10039.
- Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. 2025. [Group-in-group policy optimization for LLM agent training](#). *CoRR*, abs/2505.10978.
- Pengyu Gao, Jinming Zhao, Xinyue Chen, and Yilin Long. 2025. [An efficient context-dependent memory framework for llm-centric agents](#). In *NAACL*, pages 1055–1069.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. [Enabling large language models to generate text with citations](#). In *EMNLP*, pages 6465–6488.
- Xinming Hou, Mingming Yang, Wenxiang Jiao, Xing Wang, Zhaopeng Tu, and Wayne Xin Zhao. 2024. [Coact: A global-local hierarchy for autonomous agent collaboration](#). *CoRR*, abs/2406.13381.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *ICLR*.
- Yuchen Hu, Chen Chen, Chao-Han Huck Yang, Ruizhe Li, Dong Zhang, Zhehuai Chen, and Engsiang Chng. 2024. [Gentranslate: Large language models are generative multilingual speech and machine translators](#). In *ACL*, pages 74–90.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2025. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *ACM Trans. Inf. Syst.*, 43(2):42:1–42:55.
- Yuhan Ji and Song Gao. 2024. [Evaluating the effectiveness of large language models in representing and understanding movement trajectories](#). *CoRR*, abs/2409.00335.
- Ziwei Ji, Tiezheng Yu, Yan Xu, Nayeon Lee, Etsuko Ishii, and Pascale Fung. 2023. [Towards mitigating LLM hallucination via self reflection](#). In *EMNLP*, pages 1827–1843.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *CoRR*, abs/2310.06825.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *EMNLP*, pages 6769–6781.
- Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. 2023. [Decomposed prompting: A modular approach for solving complex tasks](#). In *ICLR*.
- Dorde Klisura, Astrid R. Bernaga Torres, Anna Karen G  rate-Escamilla, Rajesh Roshan Biswal, Ke Yang, Hilal Pataci, and Anthony Rios. 2025. [A multi-agent framework for mitigating dialect biases in privacy policy question-answering systems](#). *CoRR*, abs/2506.02998.
- Bevan Koopman, Ahmed Mourad, Hang Li, Anton van der Vegt, Shengyao Zhuang, Simon Gibson, Yash Dang, David Lawrence, and Guido Zuccon. 2024. [Agask: an agent to help answer farmer’s questions from scientific documents](#). *Int. J. Digit. Libr.*, 25(4):569–584.
- Teyun Kwon, Norman Di Palo, and Edward Johns. 2024. [Language models as zero-shot trajectory generators](#). *IEEE Robotics Autom. Lett.*, 9(7):6728–6735.
- Dongyang Li, Junbing Yan, Taolin Zhang, Chengyu Wang, Xiaofeng He, Longtao Huang, Hui Xue, and Jun Huang. 2024. [On the role of long-tail knowledge in retrieval augmented large language models](#). In *ACL*, pages 120–126. Association for Computational Linguistics.
- Xinzhe Li. 2025. [A review of prominent paradigms for llm-based agents: Tool use, planning \(including rag\), and feedback learning](#). In *COLING*, pages 9760–9779.
- Xi Victoria Lin, Xilun Chen, Mingda Chen, Weijia Shi, Maria Lomeli, Richard James, Pedro Rodriguez, Jacob Kahn, Gergely Szilvassy, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2024. [RA-DIT: retrieval-augmented dual instruction tuning](#). In *ICLR*.
- Yuxing Long, Binyuan Hui, Fulong Ye, Yanyang Li, Zhuoxin Han, Caixia Yuan, Yongbin Li, and Xiaojie Wang. 2023. [SPRING: situated conversation agent pretrained with multimodal questions from incremental layout graph](#). In *AAAI*, pages 13309–13317.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Hannaneh Hajishirzi, and Daniel Khashabi. 2022. [When not to trust language models: Investigating effectiveness and limitations of parametric and non-parametric memories](#). *CoRR*, abs/2212.10511.

- Silin Meng, Yiwei Wang, Cheng-Fu Yang, Nanyun Peng, and Kai-Wei Chang. 2024. [Llm-a\\*: Large language model enhanced incremental heuristic search on path planning](#). In *EMNLP*, pages 1087–1102. Association for Computational Linguistics.
- R. L. Plackett. 1975. [The analysis of permutations](#). *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 24(2):193–202.
- Chen Qian, Zihao Xie, Yifei Wang, Wei Liu, Yufan Dang, Zhuoyun Du, Weize Chen, Cheng Yang, Zhiyuan Liu, and Maosong Sun. 2025. [Scaling large-language-model-based multi-agent collaboration](#). In *ICLR*.
- Shuofei Qiao, Zhisong Qiu, Baochang Ren, Xiaobin Wang, Xiangyuan Ru, Ningyu Zhang, Xiang Chen, Yong Jiang, Pengjun Xie, Fei Huang, and Huajun Chen. 2025. [Agentic knowledgeable self-awareness](#). In *ACL*, pages 12601–12625.
- Qwen Team. 2024. [Qwen2.5: A party of foundation models](#).
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *NeurIPS*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100, 000+ questions for machine comprehension of text](#). In *EMNLP*, pages 2383–2392.
- Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. 2020. [Deepspeed](#). In *KDD*.
- Ohad Rubín, Jonathan Herzig, and Jonathan Berant. 2022. [Learning to retrieve prompts for in-context learning](#). In *NAACL*, pages 2655–2671.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Richard James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2024. [REPLUG: retrieval-augmented black-box language models](#). In *NAACL*, pages 8371–8384.
- Aditi Singh, Abul Ehtesham, Saket Kumar, and Tala Talaei Khoei. 2025. [Agentic retrieval-augmented generation: A survey on agentic RAG](#). *CoRR*, abs/2501.09136.
- Yifan Song, Da Yin, Xiang Yue, Jie Huang, Sujian Li, and Bill Yuchen Lin. 2024. [Trial and error: Exploration-based trajectory optimization for LLM agents](#). *CoRR*, abs/2403.02502.
- Arjun Subramonian, Xingdi Yuan, Hal Daumé III, and Su Lin Blodgett. 2023. [It takes two to tango: Navigating conceptualizations of NLP tasks and measurements of performance](#). In *ACL*, pages 3234–3279.
- Chuanneng Sun, Songjun Huang, and Dario Pompli. 2024. [Llm-based multi-agent reinforcement learning: Current and future directions](#). *CoRR*, abs/2405.11106.
- Xiaolong Tang, Meina Kan, Shiguang Shan, and Xilin Chen. 2025. [Plan-r1: Safe and feasible trajectory planning as language modeling](#). *CoRR*, abs/2505.17659.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *CoRR*, abs/2302.13971.
- Hanlin Wang, Chak Tou Leong, Jiashuo Wang, Jian Wang, and Wenjie Li. 2025a. [SPA-RL: reinforcing LLM agents via stepwise progress attribution](#). *CoRR*, abs/2505.20732.
- Yuhao Wang, Ruiyang Ren, Yucheng Wang, Wayne Xin Zhao, Jing Liu, Hua Wu, and Haifeng Wang. 2025b. [Unveiling knowledge utilization mechanisms in llm-based retrieval-augmented generation](#). *CoRR*, abs/2505.11995.
- Sangmin Woo, Kang Zhou, Yun Zhou, Shuai Wang, Sheng Guan, Haibo Ding, and Lin Lee Cheong. 2025. [Black-box visual prompt engineering for mitigating object hallucination in large vision language models](#). In *NAACL*, pages 529–538.
- Weimin Xiong, Yifan Song, Xiutian Zhao, Wenhao Wu, Xun Wang, Ke Wang, Cheng Li, Wei Peng, and Sujian Li. 2024. [Watch every step! LLM agent learning via iterative step-level process refinement](#). In *EMNLP*, pages 1556–1572.
- Heng-Da Xu, Xian-Ling Mao, Puhai Yang, Fanshu Sun, and Heyan Huang. 2024a. [Rethinking task-oriented dialogue systems: From complex modularity to zero-shot autonomous agent](#). In *ACL*, pages 2748–2763.
- Shicheng Xu, Liang Pang, Mo Yu, Fandong Meng, Huawei Shen, Xueqi Cheng, and Jie Zhou. 2024b. [Unsupervised information refinement training of large language models for retrieval-augmented generation](#). In *ACL*, pages 133–145.
- Cheng Yang, Chufan Shi, Siheng Li, Bo Shui, Yujiu Yang, and Wai Lam. 2025. [LLM2: let large language models harness system 2 reasoning](#). In *NAACL*, pages 168–177.
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022. [Webshop: Towards scalable real-world web interaction with grounded language agents](#). In *NeurIPS*.
- Fuda Ye, Shuangyin Li, Yongqi Zhang, and Lei Chen. 2024. [R<sup>2</sup>ag: Incorporating retrieval information into retrieval augmented generation](#). In *EMNLP*, pages 11584–11596.
- Shengbin Yue, Siyuan Wang, Wei Chen, Xuanjing Huang, and Zhongyu Wei. 2025. [Synergistic multi-agent framework with trajectory learning for knowledge-intensive tasks](#). In *AAAI*, pages 25796–25804.

Zhenrui Yue, Huimin Zeng, Lanyu Shang, Yifan Liu, Yang Zhang, and Dong Wang. 2024. [Retrieval augmented fact verification by synthesizing contrastive arguments](#). In *ACL*, pages 10331–10343.

Aohan Zeng, Mingdao Liu, Rui Lu, Bowen Wang, Xiao Liu, Yuxiao Dong, and Jie Tang. 2024. [Agenttuning: Enabling generalized agent abilities for llms](#). In *ACL*, pages 3053–3077.

Guibin Zhang, Yanwei Yue, Zhixun Li, Sukwon Yun, Guancheng Wan, Kun Wang, Dawei Cheng, Jeffrey Xu Yu, and Tianlong Chen. 2025a. [Cut the crap: An economical communication pipeline for llm-based multi-agent systems](#). In *ICLR*.

Shaokun Zhang, Ming Yin, Jieyu Zhang, Jiale Liu, Zhiguang Han, Jingyang Zhang, Beibin Li, Chi Wang, Huazheng Wang, Yiran Chen, and Qingyun Wu. 2025b. [Which agent causes task failures and when? on automated failure attribution of LLM multi-agent systems](#). *CoRR*, abs/2505.00212.

Taolin Zhang, Dongyang Li, Qizhou Chen, Chengyu Wang, Longtao Huang, Hui Xue, Xiaofeng He, and Jun Huang. 2024. [R<sup>4</sup>: Reinforced retriever-reorder-responder for retrieval-augmented large language models](#). In *27th European Conference on Artificial Intelligence*, Frontiers in Artificial Intelligence and Applications, pages 2314–2321. IOS Press.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhaghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). In *NeurIPS*.

Kun Zhu, Xiaocheng Feng, Xiyuan Du, Yuxuan Gu, Weijiang Yu, Haotian Wang, Qianglong Chen, Zheng Chu, Jingchang Chen, and Bing Qin. 2024. [An information bottleneck perspective for effective noise filtering on retrieval-augmented generation](#). In *ACL*, pages 1044–1069.

Chang Zong, Yuchen Yan, Weiming Lu, Jian Shao, Yongfeng Huang, Heng Chang, and Yueting Zhuang. 2024. [Triad: A framework leveraging a multi-role llm-based agent to solve knowledge base question answering](#). In *EMNLP*, pages 1698–1710.

Longwei Zou, Qingyang Wang, Han Zhao, Jiangangkong Jiangangkong, Yi Yang, and Yangdong Deng. 2024. [CQIL: inference latency optimization with concurrent computation of quasi-independent layers](#). In *ACL*, pages 7293–7307.

## A Detailed Experimental Settings

### A.1 Datasets and Evaluation Metrics

#### A.1.1 Trajectory Training Set Construction

Our trajectory data are collected from open-source long-trajectory datasets provided by Self-RAG (Asai et al., 2024) and SMART (Yue et al.,

2025), which together include 140,000 well-designed instances.<sup>2</sup> Trajectory data for two additional agents are derived from the previously constructed basic trajectories. Figure 8 illustrates the collection process for ⟨Filter⟩ trajectory data, which is guided by the ⟨Retriever⟩ step and subject to subsequent ⟨Locator⟩ constraints. Data collection for the ⟨Verifier⟩ trajectory is performed after the ⟨Generator⟩ step, and aims to verify answer robustness as shown in Figure 9.

#### A.1.2 Trajectory Scores

As shown in Figure 10, intra-trajectory scores are computed based on QA pairs and two demonstration examples. Additionally, inter-trajectory score data are annotated using demonstration examples with varying task instructions, as depicted in Figure 11. A comprehensive training example is provided in Figure 12. Inter-dependency preference data are sorted based on the sum of scores across trajectories.

There are two primary types of agents involved in scoring: knowledge agents, and generator/verifier agents. Knowledge agents are essential when multi-agent systems require external knowledge to solve complex questions. Generator/verifier agents are associated with the confidence level of LLMs during answer generation.

To better elucidate the rationality of our scoring process, we provide a detailed rubric for agent preference scoring in Table 4. For example, for the question “Which American actor played fraternity president “Lewis Skol.” in the “Revenge of the Nerds” comedy films?” (see Figure 3), the LLM annotator correctly assigns high scores to the knowledge agents (Retriever=5, Filter=4, Locator=4) as it identifies the need for specific external knowledge, and assigns a low score to the Verifier due to high confidence in the retrieved evidence.

We verify the reasonableness of these scores through two approaches: ablation studies and correlation analysis.

- **Ablation Study as Direct Evidence:** The most direct validation is to observe performance changes when removing an agent deemed important according to the score. Our ablation study (Table 3) provides strong evidence that the LLM-assigned preferences

<sup>2</sup>The open-source trajectory data are available at <https://huggingface.co/datasets/ShengbinYue/Long-short-Trajectory>.

Agent Roles	High Score (4–5) Condition	Low Score (0–1) Condition
<b>Retriever/Filter/Locator</b>	Question requires external, long-tail, or specific factual knowledge.	Question can be answered with the LLM’s parametric knowledge alone.
<b>Verifier</b>	The answer is complex, potentially ambiguous, or requires high factual precision.	The answer is straightforward or the confidence from previous steps is very high.
<b>Reconstructor</b>	The user instruction is complex, multi-hop, or requires semantic parsing.	The question is already simple and well-structured.
<b>Generator</b>	Always required (baseline score of 5).	N/A

Table 4: Detailed rubric for different agent roles in preference scoring.

Group	Knowledge Agent Scores	PopQA Accuracy
Group 1	1.2	22.5%
Group 2	3.1	41.8%
Group 3	4.2	63.4%

Table 5: Correlation analysis of the reasonableness of LLM-generated preferences. “Group 1”, “Group 2”, and “Group 3” correspond to low, medium, and high knowledge requirements, respectively.

align with the agents’ actual functional importance. For example, on PopQA (a long-tail knowledge dataset), removing  $\mathcal{A}_{KR}$  (typically high-scored) results in a  $\sim 3.1\%$  performance drop, the largest among single-agent ablations. This confirms that the LLM annotator correctly identifies the critical need for retrieval in such tasks. Conversely, removing  $\mathcal{A}_{AV}$  has a smaller but still notable impact, particularly on factuality-focused tasks like ARC-C and ASQA, justifying its medium-to-low but non-zero scores.

- **Correlation Analysis:** To further support our claim, we conduct a correlation analysis on a sample of 200 questions from the PopQA dataset, comparing the average preference score assigned to the three knowledge agents (Retriever, Filter, Locator) against final task accuracy. As shown in Table 5, the strong positive correlation demonstrates that higher LLM-assigned preference scores for knowledge agents correspond to significantly improved accuracy, thus quantitatively confirming the reasonableness of the scores.

- **DPO Score Examples:** To aid the LLM in understanding our scoring process, we provide two examples as demonstrations in the API input prompts, shown in Figure 7.

In summary, while annotation is performed by an LLM, it is grounded in a structured, semantically meaningful rubric. More importantly, its effectiveness is empirically confirmed through ablation studies and correlation analysis, demonstrating that the learned preferences contribute directly to the framework’s performance and efficiency.

### A.1.3 Evaluation Datasets and Metrics

- **Fact Verification:** PubHealth (also referred to as HealthQA) (Akhtar et al., 2022) is a public health fact-checking dataset. Model performance is evaluated by accuracy (Acc.) on its test set of 987 samples labeled “True” or “False”.
- **Multiple-choice QA:** ARC-Challenge (Clark et al., 2018) consists of 1,172 multiple-choice science exam questions. Performance is measured by accuracy (Acc.).
- **Open-domain QA:** (1) PopQA (Mallen et al., 2022) contains 1,399 long-tail, rare-entity queries from Wikipedia. (2) SQuAD (Rajpurkar et al., 2016) includes 8,886 queries written by annotators based on documents. Following prior work (Asai et al., 2024), performance is evaluated using exact match (EM).
- **Ambiguous QA:** ASQA (Gao et al., 2023) features 4,132 ambiguous factual questions

Instruction	Winning Example (Retriever: 5)
<p>Please provide the score of the agents in the trajectory based on the description of them...(Description of agents). Question: Is the following statement correct or not? Say true if it's correct; otherwise say false.</p>	<pre>&lt;Reconstructor:4&gt;, &lt;Retriever:5&gt;, &lt;Filter:4&gt;,&lt;Locator:4&gt;, &lt;Generator:5&gt;, &lt;Verifier:1&gt;. &lt;Reconstructor&gt;:Search(A block is also called something else.)&lt;/eoi&gt;&lt;Retriever&gt;[1] Hamming scheme formula_33 xxx Each edge [2] Block (basketball)\nBlock (basketball) xxx for shooter.&lt;/eor&gt; &lt;Filter&gt;[2] Block (basketball) Block (basketball) In basketball, xxx field goal attempt for shooter.&lt;/eof&gt; &lt;Locator&gt; [Relevant]: [2] Block (basketball)\nIn basketball , a block or blocked shot occurs when a defensive player legally deflects a field goal attempt from an offensive player .&lt;/eol&gt;&lt;Generator&gt;:\ntrue\n[Cite]: [2]"&gt;&lt;/eog&gt;&lt;Verifier&gt; The answer does not require verification. &lt;/eov&gt;</pre>
Instruction	Losing Example (Retriever: 1)
<p>Please provide the score of the agents in the trajectory based on the description of them...(Description of agents). Question: Is the following statement correct or not? Say true if it's correct; otherwise say false.</p>	<pre>&lt;Reconstructor:4&gt;, &lt;Retriever:1&gt;, &lt;Filter:0&gt;,&lt;Locator:1&gt;, &lt;Generator:5&gt;, &lt;Verifier:5&gt;. &lt;Reconstructor&gt;:Search(A block is also called something else.)&lt;/eoi&gt;&lt;Retriever&gt;None &lt;Filter&gt;None&lt;/eof&gt; &lt;Locator&gt; None &lt;/eol&gt;&lt;Generator&gt;: "False"&gt;&lt;/eog&gt;&lt;Verifier&gt; The answer is false. &lt;/eov&gt;</pre>

Figure 7: Examples of different preference scores for `<Retriever>`.

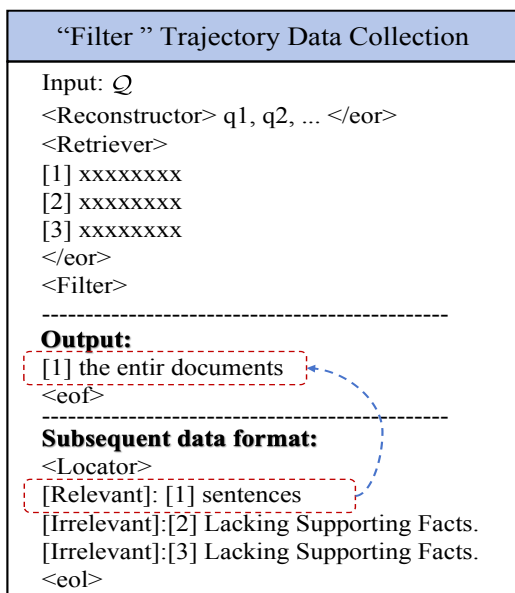


Figure 8: Collection of “Filter” trajectory data.

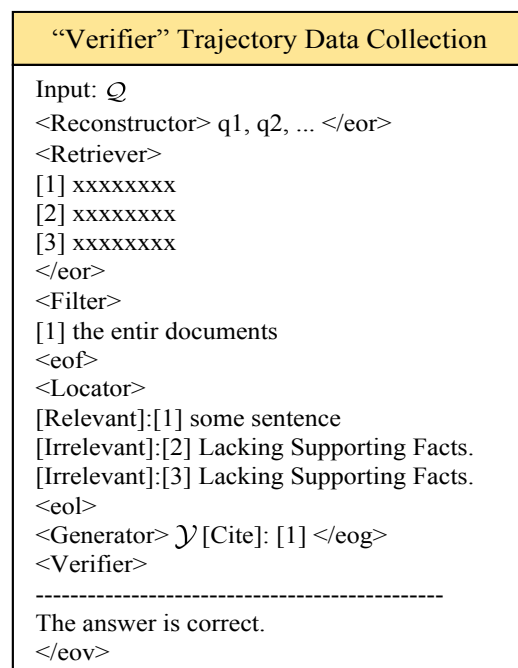


Figure 9: Collection of “Verifier” trajectory data.

requiring long-form responses. Fluency is assessed using Mauve, and accuracy is measured with Str\_EM and Rouge-L, consistent with official evaluation settings.

## A.2 Baselines

### A.2.1 Vanilla QA Methods

LLMs acquire extensive factual knowledge, internalized within their model parameters through large-scale unsupervised pre-training. During both training and inference, we adhere to official prompt

formats.

- **SFT and preference alignment models:** GPT-3.5-turbo (ChatGPT)<sup>3</sup>, Llama-2-Chat-7B (Touvron et al., 2023), and Llama-2-Chat-13B (Touvron et al., 2023).
- **SFT models:** Instruct-v0.2-7B (Jiang et al., 2023), Vicuna-v1.5-13B (Zheng et al., 2023),

<sup>3</sup>We use the gpt-3.5-turbo-0125 version in experiments.

## Intra-trajectory Preference Scores Annotation

### • **Scores' Annotation Prompt**

As an assistant, your task is to score the multi-agent tools in the trajectory based on the questions and answers. The higher the score, the greater the like-lihood of invoking the agent tool to solve the question. Among them, each agent scores 0-5 points and the answer results must strictly follow the JSON format.

-----Examples-----

There are some examples for you to refer to:

Example 1:

Input: “(Description of agents.) question: which American actor played fraternity president “Lewis Sko.” in the “Revenge of the Nerds? answer: Robert Reed Carradine.”

Output: {“Reconstructor”: 4, “Retriever”: 4, “Filter”: 4, “Locator”: 4, “Generator”: 5, “Verifier”: 1}

Example 2: ...

-----Starts Working-----

Input: “(Description of agents.) question: Which person has had sold more books than the other, Larry Baker or Michael Crichton? answer: Michael Crichton.”

Output: {“Reconstructor”: 4, “Retriever”: 5, “Filter”: 4, “Locator”: 5, “Generator”: 5, “Verifier”: 0} [Example from HotpotQA]

Figure 10: Intra-trajectory score collection.

Qwen2.5-7B-Instruct (Qwen Team, 2024), and Alpaca2-7B<sup>4</sup>.

### A.2.2 Knowledge-Augmented Methods

We implement standard knowledge augmentation approaches. When model weights are unavailable, methods are replicated using the same base models and training data. Uniform retrieval models and knowledge bases ensure experimental fairness.

- **REPLUG** (Shi et al., 2024) uses frozen LLM parameters and augments them with a tunable retrieval model. In our experiments, the backbone is replaced with Llama-2-Chat-7B for fairness.
- **VANILLA-7B** (Gao et al., 2023) first retrieves relevant passages, then instructs the model to assess document relevance and generate appropriate citations. The backbone is Llama-2-Chat-7B.
- **INTERACT-7B** (Gao et al., 2023) employs an interactive prompting mechanism that enables the agent to verify retrieved passages through three distinct actions: “Check”, “Output”, and “End”. The backbone is Llama-2-Chat-7B.
- **RADIT-7B** (Lin et al., 2024) introduces retrieval-augmented dual instruction tuning,

<sup>4</sup>[https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca)

a lightweight fine-tuning framework that retrofits existing LLMs with retrieval capabilities, offering an alternative to conventional methodologies. For fair comparison, the pre-trained Llama-2 model is fine-tuned on the same dataset used in our experiments.

- **SelfRAG-7B** (Asai et al., 2024) enhances language model quality and factuality through retrieval and self-reflection with special tokens.

### A.2.3 LLM-Based Trajectory Methods

Through multi-agent collaboration, LLMs with distinct task capabilities can be coordinated to form workflows that enhance response reliability.

- **MMAgent-3\*7B**: Our modular multi-agent framework, in which each component agent is independently trained on identical datasets while sharing the same pre-trained Llama-2 backbone. The workflow is realized through systematic agent decoupling.
- **SMART** (Yue et al., 2025): A global-local multi-agent framework with predefined trajectories.
- **GiGPO** (Feng et al., 2025): Proposes a hierarchical architecture that evaluates both global trajectory quality and local action effectiveness, while eliminating the need for auxiliary models or additional rollouts. This paradigm ensures superior scalability for long-horizon LLM agent training.
- **SPA-RL** (Wang et al., 2025a): Proposes a general reward redistribution framework that systematically decomposes the final reward into stepwise contributions, with each component accurately reflecting its incremental impact on overall task completion.

## A.3 Experimental Settings

### A.3.1 Training Details

Our implementation is initialized with the pre-trained Llama-2-7B foundation model (Touvron et al., 2023). Each agent is initialized as an independent large model, emulating distinct agent abilities through special tokens in both intra- and inter-trajectory training. These abilities are acquired by training a shared model (Asai et al., 2024; Yue et al., 2025; Qiao et al., 2025). Training is performed on two NVIDIA A100 GPUs (80GB each), utilizing

Inter-agent Dependency Scores Annotation


- **Inter-agent Dependency Scores Annotation Prompt**  
 As an assistant, your task is to score the correlation between multi-agent tools in the trajectory based on the question and answer, with a score ranging from 0 to 5. Scores with high correlation are scored higher among each other, while scores with low correlation are scored lower. Based on different correlation scores, generate 5 related score trajectories as “winning examples”, and then generate 5 trajectories with completely opposite correlation scores as “losing examples”. The answer results must strictly follow the JSON format.  
 -----Examples-----  
 There are some examples for you to refer to:  
 Example 1:  
 Input: “(Description of agents.) question: which American actor played fraternity president “Lewis Sko.” in the “Revenge of the Nerds” answer: Robert Reed Carradine.”  
 Output: {“winning example 1”: {“Reconstructor”: 4, “Retriever”: 4, “Filter”: 4, “Locator”: 4, “Generator”: 5, “Verifier”: 1}, “winning example 2”: {“Reconstructor”: 4, “Retriever”: 5, “Filter”: 5, “Locator”: 4, “Generator”: 5, “Verifier”: 0}, “winning example 3”: {“Reconstructor”: 4, “Retriever”: 4, “Filter”: 4, “Locator”: 4, “Generator”: 5, “Verifier”: 1}, ..... “losing example 1”: {“Reconstructor”: 4, “Retriever”: 1, “Filter”: 2, “Locator”: 0, “Generator”: 5, “Verifier”: 5}, “losing example 2”: {“Reconstructor”: 0, “Retriever”: 0, “Filter”: 1, “Locator”: 0, “Generator”: 5, “Verifier”: 4}, .....}  
 Example 2: .....  
 -----Starts Working-----  
 Input: “(Description of agents.) question: When is the director of film The Three Musketeers Of The West 's birthday? answer: 23 October 1931”  
 Output: {“winning example 1”: {“Reconstructor”: 5, “Retriever”: 5, “Filter”: 4, “Locator”: 5, “Generator”: 5, “Verifier”: 0} .....  
 [Example from Wikimultihop]

Figure 11: Inter-trajectory score collection.


---

**User Instruction:** Given some answer candidates, choose the best answer choice. You can use the following agents in trajectory including *(Description of agents)*. Is the following statement correct or not? Say true if it's correct; otherwise say false. Roche's schizophrenia drug misses goal in two late-stage trials.


---




**Intent Reconstructor:** Roche's schizophrenia drug misses goal in two latestage trials.




**Knowledge Retriever:** [1] **Bitopertin**. 66% on placebo (p<0.05, per-protocol population). *In January 2014, Roche reported that bitopertin failed to meet its endpoints in two phase III trials assessing its efficacy in reducing negative symptoms of schizophrenia.* Subsequently, in April 2014, Roche announced that it was discontinuing all of its phase III trials of bitopertin for schizophrenia except for one..... [2] **Tamas Bartfai**, for 'heartburn', as a consultant to Roche, the catechol-O-methyltransferase inhibitor (Tasmar-tolcapone) used in Parkinson's Disease, the first benzodiazepine-antagonist (flunitrazepam) for treatment of benzodiazepine overdoses, and as a consultant to Novartis the sphingosine 1-phosphate agonist gilenya-fingolimod as the first oral multiple sclerosis drug..... [3] **LY-404,039**. Eli Lilly and Company announced their decision to stop their ongoing clinical studies investigating LY-2140023 as a treatment for schizophrenia after their phase II study did not meet its primary endpoint. The company completed a phase II, randomized, double-blind, placebo- and active-controlled, parallel-group-assignment, dose-ranging, inpatient, multi-center clinical trial. Patients were randomized to receive 5, 20, 40, or 80 mg of LY-2140023, placebo, or olanzapine for 28 days.....




**Knowledge Filter:** [1] Bitopertin. 66% on placebo (p<0.05, per-protocol population). *In January 2014, Roche reported that bitopertin failed to meet its endpoints in two phase III trials assessing its efficacy in reducing negative symptoms of schizophrenia.* Subsequently, in April 2014, Roche announced that it was discontinuing all of its phase III trials of bitopertin for schizophrenia except for one.....



**Knowledge Locator:** [Relevant]: [1] Bitopertin. *In January 2014, Roche reported that bitopertin failed to meet its endpoints in two phase III trials assessing its efficacy in reducing negative symptoms .....* [Irrelevant]: [2] Lacking Supporting Facts [Irrelevant]: [3] Lacking Supporting Facts



**Response Generator:** True [Cite]: [1]



**Answer Verifier:** The answer is correct.

---

Figure 12: Complete response example from PubHealth.

LoRA adaptation (Hu et al., 2022) for intra- and inter-trajectory learning. Both trajectory learning phases span three epochs with uniform hyperparameters: batch size of 64, peak learning rate of  $6 \times 10^{-4}$ , and 5% warmup steps. The maximum sequence length is 1024 tokens for intra-trajectory and 2048 tokens for inter-trajectory training, with DeepSpeed Stage-3 (Rasley et al., 2020) applied to optimize GPU memory utilization.

Due to the advantage of supervision at intermediate agent training steps, RL-based trajectory

methods such as GiGPO (Feng et al., 2025) and SPA-RL (Wang et al., 2025a) outperform simple multi-agent baselines. Specifically, we assign intermediate steps in the GiGPO trajectory to the six multi-agent actions we define. The final reward signal for correctly predicted answers is set to 1, and to -1 for incorrect answers. GiGPO obtains intermediate reward signals by implicitly learning episode-relative advantages and thus does not require manual specification in the task dataset. SPA-RL, however, necessitates explicit process su-

pervision signals at intermediate steps (Wang et al., 2025a). Here, we utilize GPT-4o (gpt, OpenAI, 2023) to score each step, providing intermediate human feedback as accurate as possible. The final reward setting matches that used for GiGPO.

### A.3.2 Evaluation Details

For the two additional agents—knowledge filter ( $\langle\text{Filter}\rangle$ ) and verifier ( $\langle\text{Verifier}\rangle$ )—if  $\langle\text{Filter}\rangle$  outputs retrieved document indices not present in  $\langle\text{Retriever}\rangle$ , we remove them. If all indices are filtered out, we retain the first one. If the output trajectory includes  $\langle\text{Verifier}\rangle$ , the answer requires further verification. If the  $\langle\text{Verifier}\rangle$  output is “wrong”, we incorporate the error signal into the instruction via the prompt and re-execute the trajectory for LLM reflection. Otherwise, we extract the answer from  $\langle\text{Generator}\rangle$ . Other settings, including evaluation task instructions, follow those of the SMART model (Yue et al., 2025).

### A.4 Agent Description

We extend the SMART (Yue et al., 2025) framework by adding two essential agents to further enhance the reliability of LLM responses: the Knowledge Filter and Answer Verifier.

- **Intent Reconstructor:** Agent  $\mathcal{A}_{\text{IR}}$  elucidates user question intent. To process diverse instructions into well-formatted intents, the agent employs four key capabilities: (1) integrating contextual clues, (2) identifying key queries, (3) unifying task formulation, and (4) decomposing intent.
- **Knowledge Retriever:** Given the reconstructed question, agent  $\mathcal{A}_{\text{KR}}$  retrieves supplementary knowledge from an external knowledge base (e.g., Wikipedia). For simple questions, *AMATA* may skip intermediate knowledge agents (i.e.,  $\mathcal{A}_{\text{KR}}$ ,  $\mathcal{A}_{\text{KF}}$ , and  $\mathcal{A}_{\text{KL}}$ ) and directly invoke the response generator  $\mathcal{A}_{\text{RG}}$ .
- **Knowledge Filter:** To eliminate redundant information from retrieved documents, agent  $\mathcal{A}_{\text{KF}}$  extracts the most accurate background knowledge. Concurrently,  $\mathcal{A}_{\text{KF}}$  empowers the locator agent  $\mathcal{A}_{\text{KL}}$  to constrain the search scope and perform fine-grained identification of knowledge supportive to the generator  $\mathcal{A}_{\text{RG}}$ .
- **Knowledge Locator:** Operating on the refined document set from  $\mathcal{A}_{\text{KF}}$ , agent  $\mathcal{A}_{\text{KL}}$  per-

forms granular localization to identify and extract knowledge segments most conducive to response generation by  $\mathcal{A}_{\text{RG}}$ .

- **Response Generator:** Agent  $\mathcal{A}_{\text{RG}}$  synthesizes responses in two modes: when knowledge segments are supplied by  $\mathcal{A}_{\text{KL}}$ , it generates answers constrained by localized evidence with explicit source attribution; otherwise, it relies solely on parametric knowledge. The generator maintains provenance transparency by clearly differentiating evidence-based and intrinsic knowledge sources.
- **Answer Verifier:** Agent  $\mathcal{A}_{\text{AV}}$  performs self-correction by re-examining the generated response against relevant knowledge sources, identifying inaccuracies through evidence-based verification, and applying targeted revisions to enhance factual consistency and logical coherence.

We further detail the steps for the relevant agents (i.e.,  $\mathcal{A}_{\text{IR}}$ ,  $\mathcal{A}_{\text{KF}}$ , and  $\mathcal{A}_{\text{AV}}$ ) to clarify how specific agent trajectory data are constructed. For other agents’ data collection steps, refer to Self-RAG (Asai et al., 2024).

- **Intent Reconstructor:** Within multi-turn dialogues, this agent models dependencies to capture long-term intent. When processing noisy instructions, it eliminates extraneous content to isolate essential questions. For diverse task formats (e.g., multiple-choice QA), the agent standardizes inputs into a cohesive query representation. For complex multi-hop questions (e.g., “Who was born earlier, person A or person B?”), it decomposes them into atomic intents, such as retrieving each individual’s birthdate. By flexibly applying these capabilities, the agent derives a well-structured query intent, facilitating external knowledge retrieval.
- **Knowledge Filter:** Prompt-induced hallucinations may cause the filter  $\mathcal{A}_{\text{KF}}$  to generate documents and indices not present in the retriever  $\mathcal{A}_{\text{KR}}$ . To address this, we implement rigorous content verification procedures to ensure filtered content remains consistent with original sources. Additionally, if filtering removes all documents, the top-ranked document based on the retriever’s scores is retained.

- **Answer Verifier:** For responses marked “correct” by agent  $\mathcal{A}_{AV}$ , the reasoning process for that question is terminated. Conversely, for responses marked “wrong,” error signals are incorporated into the instruction  $Q$  via concatenation. This encourages the multi-agent trajectory to adopt a reasoning mode that considers potential errors, prompting deeper reflection and improving the model’s reasoning accuracy.

## B Additional Experimental Results

### B.1 Data Size and Computational Cost Comparison

In Figure 13, we present a detailed comparison of the training data size and computational cost for *AMATA* and its key baselines. All experiments are performed using identical hardware settings ( $2 \times$  NVIDIA A100 80GB GPUs) to ensure fairness.

We observe the following: (1) **Data Efficiency:** The total data consumption of *AMATA* ( $\sim 65,000$  samples) is substantially lower than that of SMART ( $\sim 500,000$  samples) and is comparable to SelfRAG. Most importantly, the proposed DA-DPO stage is highly data-efficient, requiring only  $\sim 5,000$  expert-ranked samples to effectively learn complex inter-agent dependencies. This is a fraction of the data used by other training stages and baselines.

(2) **Computational and Training Time Efficiency:** [1] Compared to SMART, *AMATA* requires less than half the training time ( $\sim 20$  hours vs.  $\sim 45$  hours), mainly because pre-training on a massive, separate short-trajectory dataset is unnecessary. Our intra-trajectory learning consolidates this into a single, more efficient stage. [2] Compared to RL methods (GiGPO), *AMATA* is 3–4 times faster. RL training is notoriously slow due to multiple rollouts and per-step reward computation, whereas our DA-DPO stage performs stable, offline optimization. [3] The additional cost of DA-DPO over standard DPO is minimal (an extra  $\sim 5$  hours), as it utilizes the same computational framework but employs a more sophisticated loss function.

### B.2 Hyperparameter Sensitivity Analysis

#### B.2.1 Loss Coefficients

We evaluate the impact of the loss coefficients in the total loss  $\mathcal{L}_{\text{total}}$  by experimenting with five different pairs of values for  $\alpha_1$  and  $\alpha_2$ . As shown in Figure 14, our model achieves optimal performance

when the coefficients are balanced (i.e.,  $\alpha_1 = 0.5$  and  $\alpha_2 = 0.5$ ).

#### B.2.2 Number of Winning and Losing Samples

In Figure 15, we analyze the effect of varying the number of winning and losing samples, denoted  $M$  and  $N$ , as well as the top- $K$  value in our dependency-aware DPO framework. We observe that increasing  $M$  and  $N$  causes model performance to gradually decline, likely due to the introduction of excessive retrieval noise. Consequently, we set both  $M$  and  $N$  to 10. With  $M = N = 10$ , we further investigate how different values of  $K$  affect performance across winning examples. Our findings indicate that choosing  $K$  as the midpoint of  $M$  yields optimal results. Therefore, we set  $K = 5$  to achieve the best performance reported in Table 2.

### B.3 LLMs’ Backbone Analysis

Our *AMATA* framework decouples the model architecture from backbone selection. We further evaluate several state-of-the-art LLMs, including Llama-3 (Dubey et al., 2024) and GPT-4o (gpt, OpenAI, 2023), to validate the effectiveness of our method, using Llama-3-Instruct (8B) as the backbone. As shown in Table 6, trajectory planning methods based on RL demonstrate even greater potential when enhanced LLM capabilities are available. We speculate that stronger foundational abilities are unlocked by post-RL training, which intrinsically provides planning skills for multi-agent tasks (Sun et al., 2024). Our *AMATA* consistently exhibits performance improvements as the underlying backbone is strengthened.

### B.4 Case Study

In this section, we present a case study analyzing our *AMATA* framework alongside other trajectory training paradigms, including *MMAgent*, *SMART*, and *GiGPO*.

The complete response generated by *AMATA* is shown in Figure 12. The user instruction pertains to retrieved document index “[1]”, enabling  $\mathcal{A}_{KF}$  and  $\mathcal{A}_{KL}$  to make accurate inferences based on the document. Moreover, sufficient external documentation supports the LLM’s inference, resulting in high confidence in the generated answer. Our answer verifier,  $\mathcal{A}_{AV}$ , confirms that no further validation is required.

Methods	Training Paradigm	Total Training Data	Estimated Training Time (Hours)	Key Computational Step
Self-RAG	End-to-End SFT	~ 120K	~ 18	Single-stage SFT on all trajectories
SMART	Two-Stage SFT	~ 140K (Long) + ~ 360K (Short)	~ 25 (Short) + ~ 20 (Long) = ~ 45	Two-stage SFT on massive short & long trajectories
GiGPO (RL)	Reinforcement Learning	~120K (Base trajectories)	~ 60-80 (High variance)	Expensive policy rollouts and advantage estimation
AMATA (Ours)	SFT + DA-DPO	~60K (Intra) + ~5K (Inter) = ~ 65K	~15 (Intra) + ~5 (Inter) = ~20	DA-DPO requires only 5K preference pairs

Figure 13: Comparison of training data size and computational cost for various baselines and our *AMATA* model.

Task → Model ↓	HealthQA Acc.	ARC-C Acc.	PopQA Acc.	Squad1 Acc.	Str_EM	ASQA Rouge-L	Mauve	Average
Llama-3-Ins.8B	64.39(±0.9)	53.44(±1.5)	22.73(±1.2)	16.24(±0.8)	18.85(±1.7)	33.48(±1.5)	37.65(±0.7)	35.25(±1.3)
<i>GPT4o</i>	79.24(±1.0)	80.71(±1.8)	30.94(±1.1)	24.83(±1.5)	44.02(±0.9)	36.92(±1.7)	47.53(±1.6)	49.17(±1.1)
RADIT 8B	55.29(±1.1)	64.88(±1.0)	41.15(±1.6)	24.97(±1.1)	29.01(±1.8)	17.22(±1.2)	13.86(±0.8)	35.20(±1.7)
SelfRag 8B	70.89(±1.6)	68.14(±1.2)	43.55(±1.5)	26.06(±0.8)	30.95(±1.1)	35.98(±1.2)	87.12(±2.1)	51.81(±1.2)
SMART 8B	75.99(±1.6)	72.81(±0.6)	47.66(±1.2)	33.05(±1.4)	45.74(±1.6)	44.95(±1.2)	94.80(±1.2)	59.29(±0.8)
SPA-RL 8B	76.88(±1.7)	72.98(±1.3)	46.85(±1.3)	34.51(±1.8)	46.63(±1.2)	45.57(±0.7)	94.93(±1.1)	59.76(±1.1)
GiGPO 8B	77.20(±1.0)	73.84(±0.7)	47.29(±1.7)	34.05(±1.2)	45.75(±1.8)	46.98(±1.0)	95.62(±1.7)	60.10(±1.6)
AMATA 8B	<b>78.74(±1.2)</b>	<b>74.11(±1.0)</b>	<b>49.62(±1.6)</b>	<b>37.80(±1.2)</b>	<b>50.83(±1.5)</b>	<b>51.57(±1.1)</b>	<b>96.92(±1.8)</b>	<b>62.80(±1.3)</b>

Table 6: Overall results of *AMATA* with other backbone LLMs.




Figure 14: Averaged performance of *AMATA* corresponding to different loss coefficients  $\alpha_1$  and  $\alpha_2$ .

In contrast, other trajectory learning paradigms face various types of errors. As shown in Figure 16, since *MMAgent* trains each agent independently, internal connections within the trajectory are neglected. This causes the retriever to fail to fully comprehend the semantics of the user’s question, especially when retrieved documents are labeled “None”. Consequently, the combined operation of knowledge retriever  $\mathcal{A}_{KR}$ ,  $\mathcal{A}_{KF}$ , and  $\mathcal{A}_{KL}$  does not




Figure 15: Impact of selecting the  $K$  value for winning and losing samples on model performance.

effectively contribute to the LLM’s inference.

As shown in Figure 17, SMART benefits from predefined global-local trajectory training and produces a generally correct trajectory path. However, both SMART and GiGPO, as shown in Figures 17 and 18, fail to account for dependencies between inter-agent processes. This oversight introduces external noise into the retrieved documents, resulting in inaccurate overall trajectories.

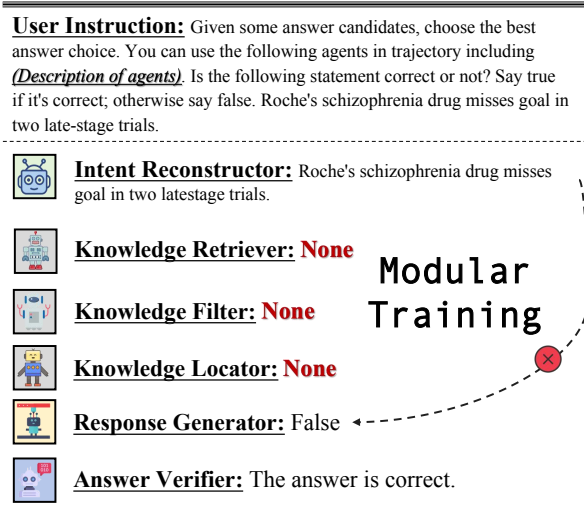


Figure 16: Complete response example from MMAgent.

## C Inference

Algorithm 1 gives an overview of inference in our AMATA framework. During inference, AMATA first analyzes the query to determine whether external knowledge is required. If not, it directly generates and verifies the answer. Otherwise, it retrieves and filters relevant documents, then generates a grounded response. The answer is subsequently verified; if found incorrect, the process iterates with updated instructions to refine the response. The specific steps are as follows:

- **Step 1:** We utilize the intent reconstructor  $\mathcal{A}_{IR}$  to decompose the user prompt  $Q$ . If the initial trajectory head token  $h = h_{RG}$  indicates that the question can be answered directly by the LLM without external knowledge, the response agent  $\mathcal{A}_{RG}$  generates the answer  $\mathcal{Y}$ . Subsequently, the answer verifier checks the consistency of the generated response. (**Lines 1–13**)
- **Step 2:** When  $h = h_{KR}$  indicates that the question  $Q$  requires external knowledge, the knowledge retriever agent  $\mathcal{A}_{KR}$  retrieves external documents  $\{d_1, \dots, d_{k,m}\}$ . The knowledge filter  $\mathcal{A}_{KF}$  removes noise from these documents to produce a refined set  $D = \{d_1, \dots, d_w\}$ . The knowledge locator  $\mathcal{A}_{KL}$  extracts fine-grained text spans  $y_{KL}$  from  $D$  based on  $Q$ . (**Lines 16–23**)
- **Step 3:** If the extracted span is relevant ( $r = [\text{Relevant}]$ ), the response generator  $\mathcal{A}_{RG}$  formulates the answer using this span; otherwise,

it answers directly. The answer verifier evaluates the response  $\mathcal{Y}$  for factual correctness. If incorrect, AMATA restarts the trajectory generation at “**Start**”, incorporating the erroneous response into  $Q$  for further reasoning. (**Lines 24–35**)

When the verifier determines that the LLM response is incorrect (i.e., **Lines 12** and **34**), we incorporate the incorrect answer into the prompt  $Q$  and reuse the previously generated dynamic trajectory for further inference, denoted as “**goto Start.**” To prevent repeated incorrect answer generation and infinite loops, a maximum of three iterations is allowed, after which the result is returned as the default answer.

Additionally, we analyze the out-of-domain generalization of our algorithm. Once trained, AMATA operates as a zero-shot, self-adaptive system that requires no pre-existing training traces or score annotations for new domains or questions.

### Generalization to Unseen Datasets and Domains.

The core of AMATA’s generalization lies in its learning objective: We train it not to memorize specific answers or trajectories but to learn two fundamental principles:

- **Intra-Trajectory Preference:** Dynamically assess which agents are important for a given question based on semantic content.
- **Inter-Agent Dependency:** Orchestrate selected agents in a coherent and efficient sequence.

Once this “collaboration policy” is learned, it can be applied to new questions. The model analyzes each new, unseen question at inference and uses its acquired knowledge to:

- Dynamically predict the relevance of each agent (simulating internal scoring).
- Execute an optimal trajectory of agent invocations based on these predictions and learned dependencies.

**Applicability without Training Traces.** Our model incurs a one-time cost for bootstrapping collaborative reasoning capability. The trained framework is fully self-sufficient at inference.

**User Instruction:** Given some answer candidates, choose the best answer choice. You can use the following agents in trajectory including (*Description of agents*). Is the following statement correct or not? Say true if it's correct; otherwise say false. Roche's schizophrenia drug misses goal in two late-stage trials.

**Intent Reconstructor:** Roche's schizophrenia drug misses goal in two latestage trials.

**Knowledge Retriever:** [1] **Bitopertin**. 66% on placebo (p<0.05, per-protocol population). *In January 2014, Roche reported that bitopertin failed tomeet its endpoints in two phase III trials assessing its efficacy in reducing negative symptoms of schizophrenia.* Subsequently, in April 2014, Roche announced that it was discontinuing all of its phase III trials of bitopertin forschizophrenia except for one..... [2] **Tamas Bartfai**, for 'heartburn', as a consultant to Roche, the catechol-O-methyltransferase inhibitor (Tasmar-tolcapone) used inParkinson's Disease, the first benzodiazepine-antagonist (flunitrazepam) for treatment of benzodiazepineoverdoses, and as a consultant to Novartis the sphingosine 1-phosphate agonist gilenya-fingolimod as the firstoral multiple sclerosis drug..... [3] **LY-404,039**, *Eli Lilly and Company announced their decision to stop their ongoing clinical studies investigating LY-2140023as a treatment for schizophrenia after their phase II study did not meet its primary endpoint.* The companycompleted a phase II, randomized, double-blind, placebo- and active-controlled, parallel-group-assignment,dose-ranging, inpatient, multi-center clinical trial. Patients were randomized to receive 5, 20, 40, or 80 mg of LY-2140023, placebo, or olanzapine for 28 days.....

**Knowledge Filter:** [1] Bitopertin. 66% on placebo (p<0.05, per-protocol population). In January 2014, Roche reported that bitopertin failed tomeet its endpoints in two phase III trials assessing its efficacy in reducing negative symptoms of schizophrenia. Subsequently, in April 2014, Roche announced that it was discontinuing all of its phase III trials of bitopertin forschizophrenia except for one..... [3] LY-404,039. Eli Lilly and Company announced their decision to stop their ongoing clinical studies investigating LY-2140023as a treatment for schizophre .....

**Knowledge Locator:** [Relevant]: [1] Bitopertin. *In January 2014, Roche reported that bitopertin failed tomeet its endpoints in two phase III trials assessing its efficacy in reducing negative symptoms .....* [Irrelevant]: [2] Lacking Supporting Facts [Relevant]: [3] *Eli Lilly and Company announced their decision to stop their ongoing clinical studies investigating.....*

**Response Generator:** True [Cite]: [1],[3] ←

**Answer Verifier:** The answer is correct.

**SMART**

Figure 17: Complete response example from SMART.

**User Instruction:** Given some answer candidates, choose the best answer choice. You can use the following agents in trajectory including (*Description of agents*). Is the following statement correct or not? Say true if it's correct; otherwise say false. Roche's schizophrenia drug misses goal in two late-stage trials.

**Intent Reconstructor:** Roche's schizophrenia drug misses goal in two latestage trials.

**Knowledge Retriever:** [1] **Bitopertin**. 66% on placebo (p<0.05, per-protocol population). *In January 2014, Roche reported that bitopertin failed tomeet its endpoints in two phase III trials assessing its efficacy in reducing negative symptoms of schizophrenia.* Subsequently, in April 2014, Roche announced that it was discontinuing all of its phase III trials of bitopertin forschizophrenia except for one..... [2] **Tamas Bartfai**, for 'heartburn', as a consultant to Roche, the catechol-O-methyltransferase inhibitor (Tasmar-tolcapone) used inParkinson's Disease, the first benzodiazepine-antagonist (flunitrazepam) for treatment of benzodiazepineoverdoses, and as a consultant to Novartis the sphingosine 1-phosphate agonist gilenya-fingolimod as the firstoral multiple sclerosis drug..... [3] **LY-404,039**. Eli Lilly and Company announced their decision to stop their ongoing clinical studies investigating LY-2140023as a treatment for schizophrenia after their phase II study did not meet its primary endpoint. The companycompleted a phase II, randomized, double-blind, placebo- and active-controlled, parallel-group-assignment,dose-ranging, inpatient, multi-center clinical trial. Patients were randomized to receive 5, 20, 40, or 80 mg of LY-2140023, placebo, or olanzapine for 28 days.....

**Knowledge Filter:** [2] Tamas Bartfai. for 'heartburn', as a consultant to Roche, the catechol-O-methyltransferase inhibitor (Tasmar-tolcapone) used inParkinson's Disease, the first benzodiazepine-antagonist (flunitrazepam) for treatment of benzodiazepineoverdoses, and .....

**Knowledge Locator:** [Irrelevant]: [1] Lacking Supporting Facts. [Relevant]: [2] Tamas Bartfai. *for 'heartburn', as a consultant to Roche, the catechol-O-methyltransferase inhibitor (Tasmar-tolcapone) used inParkinson's Disease, the first benzodiazepine-antagonist (flunitrazepam) for treatment of benzodiazepineoverdoses .....* [Irrelevant]: [3] Lacking Supporting Facts.

**Response Generator:** True [Cite]: [2] ←

**Answer Verifier:** The answer is **wrong**.

**GiGPO**

Figure 18: Complete response example from GiGPO.

---

**Algorithm 1** Inference of *AMATA*

---

**Require:** Intent Reconstructor  $\mathcal{A}_{\text{IR}}$ , Knowledge Retriever  $\mathcal{A}_{\text{KR}}$ , Knowledge Filter  $\mathcal{A}_{\text{KF}}$ , Knowledge Locator  $\mathcal{A}_{\text{KL}}$ , Response Generator  $\mathcal{A}_{\text{RG}}$ , Answer Verifier  $\mathcal{A}_{\text{AV}}$ , passage collections  $d_1, \dots, d_k$ , trajectory head token  $h$ , trajectory end token  $e$

**Input:** User prompt  $Q$

**Output:** Answer  $\mathcal{Y}$

---

```
1: Start:  $\mathcal{A}_{\text{IR}}$  predicts  $q_1, \dots, q_m, e_{\text{IR}}, h$  given
    $Q, h_{\text{IR}}$ 
2: if  $h = h_{\text{RG}}$  then
3:    $\mathcal{A}_{\text{RG}}$  predicts  $\mathcal{Y}, e_{\text{RG}}, h_{\text{AV}}$  given  $Q, e_{\text{IR}},$ 
    $h_{\text{RG}}$ 
4:   #  $\mathcal{A}_{\text{AV}}$  verifies the response.
5:    $\mathcal{A}_{\text{AV}}$  predicts  $y_{\text{AV}}, e_{\text{AV}}$  given  $Q, \mathcal{Y}, h_{\text{AV}}$ 
6:   if  $y_{\text{AV}} = \text{"Correct"}$  then
7:     return  $\mathcal{Y}$ 
8:   else
9:     # Re-executing with wrong Ans  $\mathcal{Y}$ .
10:    goto Start
11:   end if
12: else
13:   #  $h = h_{\text{KR}}$  answer the question  $Q$ .
14:   for each  $p$  in  $q_1, \dots, q_m$  do
15:     Retrieve  $(d_1, \dots, d_k)$  using  $\mathcal{A}_{\text{KR}}$  given
      $p, \text{top-}k$ 
16:   end for
17:    $D = \{d_1, \dots, d_{k \cdot m}\}$ 
18:   #  $\mathcal{A}_{\text{KF}}$  filters out unrelated  $d_i$ .
19:   Filter  $D = \{d_1, \dots, d_w\}$  using  $\mathcal{A}_{\text{KF}}$  given
    $Q, q$ 
20:   #  $\mathcal{A}_{\text{KL}}$  locates the key text span  $y_{\text{KL}}^i$ .
21:    $\mathcal{A}_{\text{KL}}$  predicts
    $\{(r_1, y_{\text{KL}}^1), \dots, (r_w, y_{\text{KL}}^w)\}, e_{\text{KL}}, h_{\text{RG}}$  given  $Q,$ 
    $\{d_1, \dots, d_w\}, e_{\text{KR}}, h_{\text{KL}}$ 
22:   if  $r = \text{[Relevant]}$  then
23:      $\mathcal{A}_{\text{RG}}$  predicts  $\mathcal{Y}, e_{\text{RG}}, h_{\text{AV}}$  given  $Q, e_{\text{KL}},$ 
      $h_{\text{RG}}, \{(r_1, y_{\text{KL}}^1), \dots, (r_w, y_{\text{KL}}^w)\}$ 
24:   else
25:      $\mathcal{A}_{\text{RG}}$  predicts  $\mathcal{Y}, e_{\text{RG}}, h_{\text{AV}}$  given  $Q,$ 
      $h_{\text{RG}}$ 
26:   end if
27:   #  $\mathcal{A}_{\text{AV}}$  verifies the response.
28:    $\mathcal{A}_{\text{AV}}$  predicts  $y_{\text{AV}}, e_{\text{AV}}$  given  $Q, \mathcal{Y}, h_{\text{AV}}$ 
29:   if  $y_{\text{AV}} = \text{"Correct"}$  then
30:     return  $\mathcal{Y}$ 
31:   else
32:     goto Start
33:   end if
34: end if
```

---