



AGENTASK: Multi-Agent Systems Need to Ask


Bohan Lin^{*,♡}, Kuo Yang^{*}, Zelin Tan^{*,◇}, Yingchuan Lai^{*}, Chen Zhang[◇], Guibin Zhang[♣], Xinlei Yu[♣], Miao Yu^{*}, Xu Wang^{*,♡}, Yudong Zhang^{*,♡,♣}, and Yang Wang^{*,♡,♣}

^{*}University of Science and Technology of China (USTC)

[♡]Suzhou Institute for Advanced Research, USTC

[◇]Shanghai AI Laboratory [♣]Xi'an Jiaotong University [♣]National University of Singapore

Contact: linbohan04@gmail.com, yudong.zhang@ustc.edu.cn

[♣] Corresponding authors  Code

Abstract

Multi-agent systems (MAS) built on large language models promise improved problem-solving through collaboration, yet they often fail to consistently outperform strong single-agent baselines due to error propagation at inter-agent message handoffs. In this work, we conduct a systematic empirical analysis of such failures and introduce an edge-level error taxonomy that identifies four dominant error types: Data Gap, Signal Corruption, Referential Drift, and Capability Gap, as primary sources of failure in multi-agent interactions. Building on this taxonomy, we propose AgentAsk, a lightweight clarification module designed to intervene at the edge level in MAS to prevent cascading errors. The module operates by strategically applying minimal clarifications at critical points within the system, improving the accuracy and efficiency of the overall task. AgentAsk is trained to balance the trade-offs between clarification cost, latency, and accuracy, while it is also architecture-agnostic and can be easily integrated into existing systems. Evaluated across five benchmarks, AgentAsk consistently improves accuracy by up to 4.69%, while keeping latency and extra costs below 10% compared to baseline MAS, showcasing its high efficiency and minimal overhead.

1 Introduction

Large language model (LLM)-based agent systems have garnered attention due to their ability to perform multi-step reasoning and dynamic tools. By coordinating multiple LLM-driven agents, these systems leverage collective intelligence to tackle complex tasks that require decomposition, external knowledge access, and iterative refinement (Park et al., 2023). This approach, known as multi-agent systems (MAS), has the potential to solve a wide range of real-world problems (Liu et al., 2024; Yu et al., 2025a; Wu et al., 2024a).

Despite the promise of collaborative decomposition, multi-agent systems often fail to reliably sur-

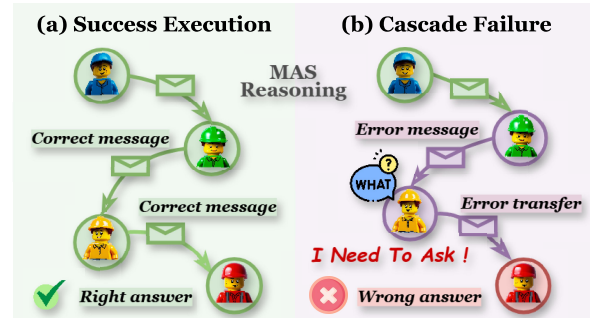


Figure 1: **Multi-Agent System Reasoning: Success vs. Cascade Failure.** The left illustrates successful multi-agent reasoning with consistent information exchange, while the right shows how errors introduced by upstream agents propagate across interactions, triggering a cascade that leads to system-level failure.

pass strong single-agent baselines across diverse tasks. Previous research and benchmarks report unstable performance gains and brittle behavior under realistic orchestration (Liu et al., 2024; Cemri et al., 2025). Execution logs reveal a recurring pattern in which small inconsistencies or missing details introduced at one step travel along the chain and accumulate into system-level failures. Figure 1 illustrates this process. Building on this, chain-wise amplification is intensified by specification gaps and inter-agent misalignment, and in some settings, collaboration reduces accuracy (Zhang et al., 2025d; Wynn et al., 2025; Wang et al., 2025). Prior studies quantify the effect, with success on difficult programming tasks falling to about 25% (Cemri et al., 2025). Similar vulnerabilities appear in single-agent traces, including misinterpretations, logical gaps, and limited reflection, which indicates that subtle errors arise early and persist through execution. These observations motivate the problem we study: *how to limit error growth at inter-agent handoffs, such that small local inconsistencies do not accumulate into system-level failures.*

A growing body of recent research attempts to improve the reliability of MAS. One direction

emphasizes structured roles and workflow governance, enforcing explicit hierarchies and protocols to reduce miscommunication (Hong et al., 2024; Qian et al., 2024; Chen et al., 2024b; Yu et al., 2025b). Another line refines orchestration and prompting strategies, leveraging prompt engineering and search to better allocate tasks (Zhuge et al., 2024; Zhang et al., 2025b,a; Yue et al., 2025). Concurrently, self-checking and feedback mechanisms encourage agents to reflect on or repair their own outputs (Liang et al., 2024; Shinn et al., 2023; Madaan et al., 2023). These approaches, spanning architecture, algorithm, and introspection, provide useful gains but remain limited. Many proposed solutions are domain-specific, costly to scale, or increase complexity without offering general principles. Moreover, existing taxonomies of failure are largely descriptive and rarely yield prescriptive mechanisms for prevention or correction. Consequently, current systems remain fragile when confronted with complex or open-ended tasks.

To this end, we argue that ensuring reliability necessitates a paradigm shift toward *edge-level* intervention: each inter-agent message should be treated as a potential failure point where minimal clarification can halt cascading errors. Based on this perspective, we introduce a taxonomy of four core edge-level error types that are responsible for most failures in MAS interactions. Guided by these insights, we propose **AgentAsk**, a plug-and-play clarification module for LLM-based multi-agent systems. AgentAsk operates by identifying potential failure points at each message handoff and applying minimal clarification to prevent cascading errors. Crucially, AgentAsk is designed to be lightweight, architecture-agnostic, and easily integrated into existing frameworks, enabling systems to intercept misunderstandings before they spread.

Our work makes the following contributions:

- ❶ **Edge-level perspective and taxonomy:** We introduce an edge-level perspective on multi-agent failures and formalize a taxonomy of four inter-agent error types, motivating clarification as an effective mechanism for improving multi-agent collaboration.
- ❷ **Practical clarification module:** We propose **AgentAsk**, a lightweight and plug-and-play clarification module that uses the edge-level taxonomy to identify and resolve interaction errors at message handoffs, improving system reliability with minimal intervention.

- ❸ **Comprehensive empirical evaluation:** Extensive experiments across multiple benchmarks show that **AgentAsk** consistently enhances accuracy by up to **4.69%**, while keeping latency and extra cost below 10% compared to existing multi-agent baselines.

2 Related Work

LLM-Based Multi-Agent Systems. MAS coordinates multiple role-specialized agents, often equipped with tools, to solve complex tasks. Early systems demonstrated that organizing multiple LLMs into role-specialized, conversational teams with tool use can outperform single models, through frameworks that scaffold roles, protocols, and conversation programming such as AutoGen (Wu et al., 2024a), AgentVerse (Chen et al., 2024b), CAMEL (Li et al., 2023), MetaGPT (Hong et al., 2024), and ChatDev (Qian et al., 2024). Recent work treats orchestration as a learnable object: GPTSwarm models agents and communication links as an optimizable graph (Zhuge et al., 2024), AFlow searches over code-represented workflows via Monte Carlo tree search (Zhang et al., 2025b), MaAS samples query-dependent sub-architectures from an agentic supernet (Zhang et al., 2025a), and MasRouter learns a cascaded controller for collaboration mode, role allocation, and LLM selection (Yue et al., 2025). These directions build on and complement decision-time tool use, such as ReAct (Yao et al., 2023) and Toolformer (Schick et al., 2023). Our work is orthogonal to global orchestration search: we hold an existing pipeline fixed and inject a lightweight edge-level clarifier that adds questions at message handoffs to curb error propagation.

Failure Analyses in MAS. Large-scale audits show that multi-agent gains over strong single-agent baselines are not guaranteed and that many failures arise at message handoffs (Cemri et al., 2025; Zhu et al., 2025; Zhang et al., 2025d). Methodologically related attempts to improve reliability include multi-agent debate and discussion (Chan et al., 2024; Liang et al., 2024), self-feedback loops such as Reflexion (Shinn et al., 2023) and Self-Refine (Madaan et al., 2023), and verification or self-correction procedures (Wu et al., 2024b; Lee et al., 2025), yet recent analyses document conditions where debate or naive self-correction can fail or even degrade accuracy (Wynn et al., 2025). A complementary line advocates asking clarify-

ing questions under uncertainty in single-agent human–computer interaction (Wang et al., 2024; Zhang et al., 2025c; Mukherjee et al., 2025; Lee et al., 2023). We follow this clarification-centric view but operationalize it for multi-agent settings with an edge-aware policy that decides when, what, whom, and how to ask exactly at the handoff.

3 Error Taxonomy

Conceptual Inspiration. In multi-agent systems, ensuring reliable communication between agents is essential to prevent cascading errors that can propagate throughout the system. Our error taxonomy is inspired by human collaboration theory (Gittell, 2006; Woolley et al., 2010), including the concept of *Relational Coordination*, which emphasizes the importance of frequent, timely, and accurate communication, shared goals, and mutual respect among team members. According to this theory, coordination can falter when communication is unclear, roles are misaligned, or team members fail to anticipate each other’s needs, leading to errors in the system.

Empirical Evidence. In this work, we empirically demonstrate that chain-style error propagation is a fundamental root cause of failures in MAS, where a single error can cascade into a system-wide collapse. We reveal these errors through an audit of **824** execution logs, annotating each agent-to-agent message as the unit of analysis. To ensure the reliability of our classification, the annotation process was conducted by multiple professional annotators with expertise in MAS. These annotators engaged in several rounds of cross-annotation and adjudication to achieve consistent results. The high inter-rater reliability was confirmed with a Fleiss’ Kappa value of **0.84**, indicating substantial agreement among annotators. Our empirical results, shown in Figure 2, indicate that most errors in MAS stem from violations of a fundamental principle: *agents should complete tasks accurately within their capabilities and deliver clear and complete information handovers to downstream agents without omission*. Building on this observation, four types of errors have been condensed.

3.1 Data Gap

In a good collaboration, a handoff never misses a crucial detail. In human teams, a brief that forgets a constraint sends downstream work off course, and the cost multiplies with each relay. In multi-agent settings, the same dynamic appears when a

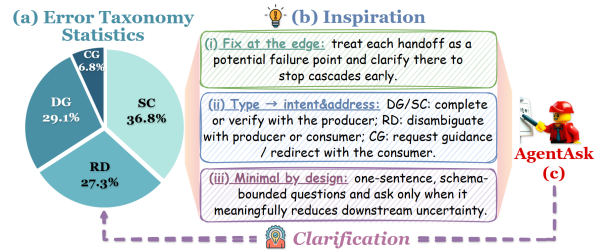


Figure 2: **From taxonomy to design.** (a) Empirical distribution of edge-level error types ($N=824$). (b) Three design inspirations distilled from our taxonomy: localize fixes at the edge; map *type* \rightarrow *intent & addressee*; enforce minimality-by-design. (c) Our designed AgentAsk for edge-level intervention (see Section 4).

boundary case, a unit, or a condition is left out. Small omissions turn into systemic bias as later steps harden a wrong default. A short check at the handoff that lists required fields and confirms ranges stops the slide before it spreads.

3.2 Referential Drift

Teams falter when words stop pointing to the same thing. People know the feeling when two teammates use different names for the same object, and the dialogue quietly splits into parallel threads. Agents replicate this when pronouns, variable names, or record keys are passed without grounding, so later steps bind them inconsistently. The result is confident but incompatible reasoning. A crisp confirmation of who is who and which value is which at the handoff restores a single shared state.

3.3 Signal Corruption

A distorted intermediate result enters the chain and gets propagated as truth. In human work, a spreadsheet with a subtle formula error can contaminate every chart that follows. Agents exhibit the same fragility when a wrong fact, a malformed structure, or a unit mistake is presented with fluency and then reused without scrutiny. Downstream modules amplify the error instead of correcting it. A targeted request to verify the source value or structure at the handoff contains the damage early.

3.4 Capability Gap

Healthy organizations respect the limits of each role. Teams recognize this when a specialist is asked to deliver outside their training, and improvisation replaces method. In multi-agent pipelines, this appears when a role tuned for explanation is asked for algebra or when a planner is pushed into domain judgment. Progress stalls because the

wrong skill sits at the critical edge. An early redirect or a brief request for guidance at the handoff aligns the task with a capable peer and keeps the chain coherent.

3.5 Inspiration from Edge-level Taxonomy

Our analysis shows that all four error types manifest as failures at the edge, the information handoff between agents in a multi-agent system. This shifts the design focus to the handoff itself rather than the global workflow. **In our annotated corpus ($N=824$), these errors account for most failures**, with Data Gap 29.1%, Referential Drift 27.3%, Signal Corruption 36.8%, and Capability Gap 6.8%. According to the human collaboration theory we mentioned above, effective handoffs are governed by three habits: ❶ check the handoff for missing content, ambiguous reference, corrupted intermediate results, or role mismatch; ❷ if a problem exists, fix it at the source so it does not propagate downstream; ❸ intervene only as needed, keeping the exchange brief and within latency and cost limits.

These habits motivate an edge-centered design that monitors each transfer and issues minimal clarifications exactly at the point of exchange. The illustrative examples can be seen in Appendix D.

4 Design of AgentAsk

Overview. Grounded in the taxonomy, failures cluster at the edge where one agent hands information to the next. We therefore design around three moves and introduce **AgentAsk**, a framework for real-time monitoring and correction in multi-agent systems. First, to identify errors, we place an external clarifier at each handoff that inspects the outgoing message and its local context, flagging risks early. Second, to correct errors, we use “Ask”: the AgentAsk sends a brief, targeted question to the right party so the handoff is repaired before the chain advances. Third, we transferred the asking capabilities of the powerful model to AgentAsk, but to keep correction efficient, fine-tuning alone is not enough, so we add reinforcement learning that adapts the ask decision, the addressee, and the phrasing to context and budget.

The training process of the framework is detailed in the following chapters, while its overall design is outlined in Figure 3. Complete training pseudocode appears in Appendix B as Algorithm 1, and the notation used throughout is summarized in Table 14 of Appendix E.

4.1 Clarification as Edge Intervention

We model multi-agent collaboration as a sequence of message handoffs over a directed interaction graph. At each step, the controller **AgentAsk** receives a local view of the exchange and decides whether to insert a clarification to improve task outcomes. Each decision is based on an edge-local state x_t , where t is the timestep of the handoff, and results in an action $a_t = (z_t, \tilde{v}_t, q_t)$, where z_t is a binary gate indicating whether to ask, \tilde{v}_t selects which agent to query, and q_t is the generated clarification question. The controller either allows the message to proceed or poses a brief clarification to one of the agents. This intervention is lightweight but can substantially improve task-level accuracy.

The policy π_θ governs the action sequence across an input, yielding a trajectory $\tau = \{(x_t, a_t)\}_{t=1}^T$. The controller is optimized to maximize expected task utility while operating under a cost constraint:

$$\max_{\theta} \mathbb{E}_{\tau \sim \pi_\theta} [U(\tau)] \quad \text{s.t.} \quad \mathbb{E}[C(\tau)] \leq B. \quad (1)$$

Here, U is the utility function measuring task success, and C represents the cost incurred from clarifications, while B is the predefined budget limit. The internal structure of x_t and a_t , as well as the factorized form of π_θ , are introduced and discussed in Appendix A.1.

4.2 Learning to Ask from Demonstrations

We initialize the controller under a cold-start condition using fine-tuning on curated clarification traces. This allows AgentAsk to acquire structured ask behavior before reinforcement optimization. We construct a dataset of edge-local examples, each consisting of a context x_i and an intervention y_i . Specifically, x_i represents the state at edge i , and y_i encodes the teacher’s decision (whether to ask, whom to ask, and what to ask). The controller learns to predict the same content.

Let $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ be the training dataset, where N is the total number of examples. The training loss includes a decision component $\mathcal{L}_{\text{type}}$ for predicting whether a clarification is needed, and a conditional ask component \mathcal{L}_{ask} for predicting the clarification details:

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{type}} + \lambda_{\text{ask}} \mathcal{L}_{\text{ask}}. \quad (2)$$

Here, λ_{ask} is a parameter that balances two components. This provides a cold-start policy capable of clarification. Details of the data construction, label format, and loss terms appear in Appendix A.2.

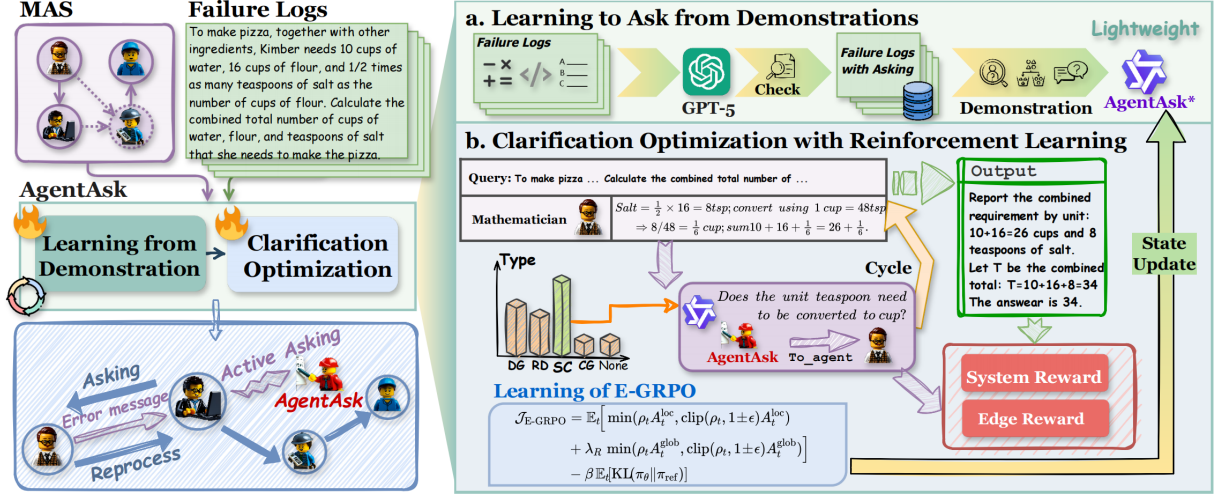


Figure 3: **Overview of AgentAsk.** The module operates at the *edge level*, treating each inter-agent message as a potential failure point. The figure illustrates both the architecture and training process: (i) Constructing an edge-level corpus from failure logs by a large evaluator, (ii) Learning to ask from demonstrations to decides *when/what/whom/how* to ask, and (iii) reinforcement learning with E-GRPO for clarification optimization under latency and cost constraints. This pipeline equips AgentAsk with edge-aware monitoring and minimal, targeted interventions while remaining architecture-agnostic and easy to integrate into diverse MASs.

4.3 Clarification Optimization with Reinforcement Learning

Previous training enables us to transfer edge-level judgment and question quality from a strong evaluator. However, imitation learning has its limits: it only captures behavior from logged cases and cannot learn when not to ask, nor can it balance accuracy against constraints in new scenarios. The ask gate and addressee are discrete decisions whose true value emerges only after downstream effects, so credit must propagate through the entire sequence. Therefore, we continue with reinforcement learning to adapt the π_θ to the constrained objective in Equation (1), using shaped rewards that convert these trade-offs into learnable signals.

Reward Design. First, we encourage asking only when it actually removes downstream uncertainty:

$$r_t^{\text{eff}} = \begin{cases} 1, & z_t = 1 \wedge n_{t+1} = 0, \\ -1, & z_t = 1 \wedge n_{t+1} = 1, \\ 0, & z_t = 0. \end{cases} \quad (3)$$

Here, n_{t+1} represents whether need to clarify again, that AgentAsk may ask multiple times at one edge.

Second, we enforce parsimony to keep clarifications scarce via a sliding counter c_t :

$$r_t^{\text{par}} = -\lambda_{\text{sw}} \max(c_t - 1, 0). \quad (4)$$

To stabilize outputs and preserve brevity, we introduce a format bonus:

$$r_t^{\text{fmt}} = \alpha_{\text{fmt}} \mathbb{1}[F_t = 1]. \quad (5)$$

These signals are aggregated into an edge reward that embodies our inspirations from taxonomy,

$$r_t^{\text{edge}} = \alpha_{\text{eff}} r_t^{\text{eff}} + r_t^{\text{par}} + r_t^{\text{fmt}}, \quad (6)$$

while a terminal score anchors final correctness,

$$R = \alpha_{\text{ans}} s. \quad (7)$$

where s is the final success score of the task.

Optimization: E-GRPO. We extend the principles of Group Relative Policy Optimization (GRPO) (DeepSeek-AI et al., 2025) by introducing a *gradual global credit mechanism* to be our **E-GRPO**. While GRPO optimizes policies based on edge-level rewards, we ensure that **E-GRPO aligns local decisions with overall success by incorporating global feedback in a controlled manner.**

For each edge t , we form a local advantage $A_t^{\text{loc}} = r_t^{\text{edge}}$ by the controller. The system-level signal $A_t^{\text{glob}} = w_t (R - b)$ is distributed at time t with nonnegative weights w_t and an action-independent baseline b . The weight w_t increases as the training progresses, ensuring that global feedback starts to influence the policy more as the agent gains confidence. Initially, the agent learns from edge-level rewards r_t^{edge} , which guide the local decisions. However, to prevent the agent from getting trapped in local optima, we introduce global credit to reflect the overall success of the task. This global credit is gradually integrated into the learning process through a time-varying weight $\lambda_R(t)$, which increases as the training progresses.

Let $\rho_t = \pi_\theta(a_t | h_t) / \pi_{\text{old}}(a_t | h_t)$. We maximize the E-GRPO surrogate that shares ratios for local and global credit and regularizes toward a reference:

$$\begin{aligned} \mathcal{J}_{\text{E-GRPO}}(\theta) = & \mathbb{E}_t \left[\min(\rho_t A_t^{\text{loc}}, \text{clip}(\rho_t, 1 \pm \epsilon) A_t^{\text{loc}}) \right. \\ & \left. + \lambda_R \min(\rho_t A_t^{\text{glob}}, \text{clip}(\rho_t, 1 \pm \epsilon) A_t^{\text{glob}}) \right] \\ & - \beta \mathbb{E}_t [\text{KL}(\pi_\theta(\cdot | h_t) \| \pi_{\text{ref}}(\cdot | h_t))] . \end{aligned} \quad (8)$$

When R is not available, the second term in Eq. (8) is simply absent, and training proceeds purely on edge-level selection; when R arrives, it is injected uniformly through A_t^{glob} , ensuring that *minimal, successful clarifications remain aligned with end-to-end correctness* while updates stay conservative through clipping and reference KL regularization.

This approach ensures that **AgentAsk** avoids becoming trapped in local optima, while progressively aligning edge-level decisions with global task outcomes, thereby improving the reliability of multi-agent systems by preventing early error propagation with minimal overhead.

5 Experiments

We assess whether a lightweight, *edge-level* clarifier improves reliability under unchanged orchestration, how it trades accuracy against latency and cost, and why these effects arise. Our study is structured around four research questions (**RQ**):

- **RQ1: Paradigm validity and end-to-end effectiveness.** Under unchanged orchestration, does inserting an edge-level clarifier consistently improve end-task performance across frameworks and benchmarks?
- **RQ2: Efficiency under practical budgets.** Can AgentAsk achieve these gains with low overhead in latency and extra cost, and how does it compare to the +GPT-5 upper-bound clarifier in the accuracy–efficiency trade-off?
- **RQ3: Ablations and robustness.** Which components are necessary for the observed gains, and how stable is performance under key design choices and control parameters?
- **RQ4: Taxonomy-grounded mechanisms.** Which error types drive most interventions and gains, and how does the edge-level clarification improve efficiency while minimizing unnecessary overhead?

5.1 Experimental Setup

Benchmarks & Baselines. We evaluate **AgentAsk** on five public benchmarks spanning math

reasoning (GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021b)), general-knowledge QA (MMLU (Hendrycks et al., 2021a)), and code generation (HumanEval (Chen et al., 2024a), MBPP (Austin et al., 2021)). Baselines include single-model prompting (IO, CoT, Self-Refine (Wei et al., 2022; Madaan et al., 2023)) and four representative multi-agent frameworks (GPTSwarm (Zhuge et al., 2024), AFlow (Zhang et al., 2025b), MaAS (Zhang et al., 2025a), MasRouter (Yue et al., 2025)).

Implementation Details. For each framework, we consider three settings that isolate the effect of the clarifier while holding orchestration constant: (i) *origin* uses the public implementation; (ii) +GPT-5 substitutes the clarifier with GPT-5 (OpenAI, 2025); (iii) +AgentAsk plugs in our lightweight module. The LLM executor is GPT-4o-mini-0718 (OpenAI, 2024); +AgentAsk uses Qwen-3-4B (Yang et al., 2025) as the backbone. Accuracy for GSM8K, MATH, and MMLU is reported; Pass@1 for HumanEval and MBPP. Due to baseline performance differences across frameworks, we normalize latency and extra cost by setting the *origin* framework to 100, and compare the relative increases for each configuration. Prompts are provided in the Appendix F. To better demonstrate our plug-and-play capabilities, we trained AgentAsk on one framework and then gradually integrated it into other frameworks.

5.2 RQ1: Effectiveness under Fixed Orchestration

Observation. Edge-level clarification consistently improves end-task performance under unchanged orchestration, suggesting that many failures originate at inter-agent handoff boundaries and can be mitigated by clarifying exchanged messages.

Results and analysis. We compare our lightweight AgentAsk with a strong clarifier, +GPT-5, which serves as an *upper bound* for the edge-level clarification paradigm. Table 1 shows that adding an edge-level clarifier yields consistent gains across frameworks and benchmarks. For each framework, both +AgentAsk and +GPT-5 improve performance over *origin* on all five datasets, and the average score increases by a clear margin. In particular, +AgentAsk improves the framework-level averages by about +3.29 to +3.45 points, while +GPT-5 improves them by about +4.40 to +4.69 points. The gains are consistent across diverse task types, including math, general knowledge, and code genera-

Method	Settings	GSM8K	MATH	HumanEval	MMLU	MBPP	Average
IO		89.52	49.23	89.66	79.63	71.38	75.88
CoT		89.71	48.19	90.64	80.44	70.66	75.93
Self-Refine		90.37	47.92	88.88	78.31	71.49	75.40
GPTSwarm	origin	92.18	51.41	90.63	81.26	76.81	78.46
	+GPT-5	95.55 \pm 1.42	56.30 \pm 1.15	94.05 \pm 0.88	86.85 \pm 1.73	82.45 \pm 1.56	83.04 \uparrow (+4.58)
	+AgentAsk	94.52 \pm 1.27	55.05 \pm 1.64	92.95 \pm 1.09	85.25 \pm 1.81	81.15 \pm 1.33	81.78 \uparrow (+3.32)
AFlow	origin	91.35	50.89	91.67	82.16	79.63	79.14
	+GPT-5	94.85 \pm 1.68	56.20 \pm 0.92	95.05 \pm 1.47	87.75 \pm 1.25	85.10 \pm 1.80	83.79 \uparrow (+4.65)
	+AgentAsk	93.70 \pm 1.54	54.95 \pm 1.38	93.95 \pm 1.12	86.40 \pm 1.66	83.90 \pm 1.05	82.58 \uparrow (+3.44)
MaAS	origin	92.84	51.08	92.02	82.93	81.22	80.02
	+GPT-5	96.20 \pm 1.89	56.85 \pm 1.51	95.45 \pm 0.76	88.35 \pm 1.94	86.70 \pm 1.37	84.71 \uparrow (+4.69)
	+AgentAsk	95.10 \pm 1.23	55.55 \pm 1.70	94.45 \pm 1.58	86.95 \pm 1.44	85.30 \pm 1.19	83.47 \uparrow (+3.45)
MasRouter	origin	93.26	51.52	91.03	83.19	79.04	79.61
	+GPT-5	95.10 \pm 1.05	56.70 \pm 1.82	94.55 \pm 1.31	88.55 \pm 1.67	85.15 \pm 1.96	84.01 \uparrow (+4.40)
	+AgentAsk	94.86 \pm 1.48	55.65 \pm 1.29	93.35 \pm 1.74	87.10 \pm 1.52	83.55 \pm 1.08	82.54 \uparrow (+3.29)

Table 1: Accuracy/Pass@1 across frameworks under three settings (origin, +GPT-5, +AgentAsk) on five benchmarks. For improved settings (+GPT-5 and +AgentAsk), bottom-right overlays show standard deviation in gray. Average column shows overall percentage-point gains (green \uparrow) relative to each framework’s origin.

Settings	GSM8K		
	Accuracy	Latency	Extra Cost
origin	93.26	100	0.0
+GPT-4o-mini	94.62 \uparrow (+1.36)	118	16.0
+GPT-5	95.10 \uparrow (+1.84)	129	34.0
<i>AgentAsk</i> (Backbone: Llama-3.2-3B)			
SFT	93.64 \uparrow (+0.38)	106	5.7
(E-GRPO)	94.23 \uparrow (+0.97)	105	5.0
<i>AgentAsk</i> (Backbone: Qwen-3-4B)			
SFT	93.99 \uparrow (+0.73)	105	5.3
(E-GRPO)	94.86 \uparrow (+1.60)	108	4.9

Table 2: Results of experiments on different types or training stages of clarifier with MasRouter@GSM8K.

tion, indicating that edge-level issues are a common source of failure that can be mitigated by clarifying exchanged messages rather than altering the orchestration. As expected, +GPT-5 produces the largest improvements, supporting its role as an approximate upper bound for the clarification paradigm under the same pipeline. Notably, +AgentAsk recovers a substantial portion of this upper-bound gain with a lightweight learned module. The remaining gap is primarily attributable to efficiency considerations, which we analyze in RQ2.

To verify that the improvements are not specific to one decoding configuration, we additionally evaluate the main settings under multiple decoding temperatures and evaluation seeds, with full robustness results reported in Appendix C.1.

5.3 RQ2: Efficiency under Practical Budgets

Observation. AgentAsk attains most of the improvement with lower latency and extra cost, making it more efficient than heavy clarifiers.

Results and analysis. We quantify the accuracy–efficiency trade-off on MasRouter@GSM8K in Table 2. Using +GPT-5 yields the highest accuracy at **95.10**, but increases latency to **129** and extra cost to **34.0**. In contrast, AgentAsk with Qwen-3-4B achieves **94.86** accuracy with only **108** latency and **4.9** extra cost, reducing latency by **21** points and extra cost by **29.1** points while leaving a **0.24** point accuracy gap. The Llama-3.2-3B variant shows the same trend, further indicating that the efficiency advantage is not tied to a specific backbone.

Figure 4 (Left) shows that +AgentAsk stays near the accuracy envelope with low overhead, while +GPT-5 achieves slightly higher accuracy at much higher latency and cost. These results confirm that our edge-level clarifier is efficient under tight budgets, with the gap to the upper bound primarily due to capacity limitations rather than orchestration changes. We attribute this efficiency to the training objective that penalizes unnecessary interventions, encouraging the clarifier to ask only when the expected downstream benefit outweighs the overhead. Additional breakdowns across frameworks and datasets are reported in Appendix C.2.

5.4 RQ3: Mechanism and Robustness at Edge

Observation. The effectiveness of AgentAsk is driven by its ability to intervene at communication handoffs, improving accuracy while maintaining efficiency. Its robustness is confirmed through ablation studies and sensitivity analyses.

Results and analysis. Table 3 presents the results of ablation experiments on the reward components

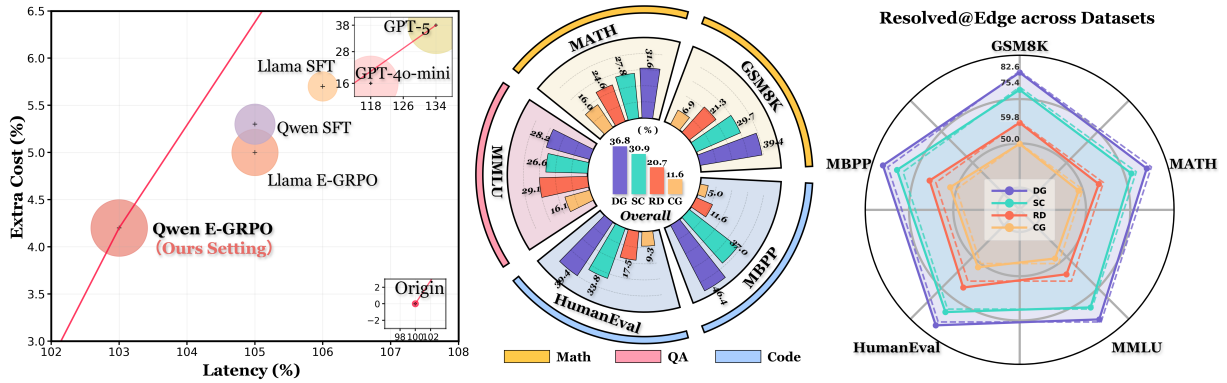


Figure 4: **(Left)** Pareto frontier ($x=latency$, $y=extra\ cost$; $bubble\ size=Acc$), showing +AgentAsk nearing +GPT-5 at much lower overhead. **(Middle)** Error-Type distributions (DG/SC/RD/CG) across datasets. **(Right)** Resolved@Edge of four error types across different datasets in AgentAsk.

Settings	MBPP		
	Accuracy	Latency	Extra Cost
AgentAsk	83.55	103	7.7
w/o r^{par}	81.34↓(-2.21)	118	9.6
w/o r^{eff}	80.05↓(-3.50)	108	7.0
w/o r^{fmt}	81.42↓(-2.13)	111	9.7
R only	79.66↓(-3.89)	123	13.2

Table 3: Results of ablation experiments on **different reward settings** with MasRouter@MBPP (red ↓ represents the decrease to each framework’s AgentAsk).

for **AgentAsk** with the MBPP benchmark. Removing individual reward components leads to a decrease in accuracy and an increase in latency and cost, showing that each component contributes to balancing efficiency and performance. The removal of the r^{par} increases the latency and extra cost without much improvement in accuracy, emphasizing the importance of minimizing unnecessary clarifications for maintaining efficiency. Finally, the full set of results, including further details on ablation studies and experiments with PPO, GRPO, and our method, can be found in Appendix C.3, where Figure 5 demonstrates the sensitivity.

5.5 RQ4: Taxonomy Mechanisms at Edge

Observation. The performance of AgentAsk is influenced by the types of errors it addresses, with varying effectiveness across different error types.

Results and analysis. Figure 4 (Middle and Right) shows the distribution of error types across datasets, along with the Resolved@Edge percentages for each error type. Resolved@Edge represents the percentage of clarifications that successfully resolve the error without further intervention (illustrated in Eq 3). AgentAsk performs particularly well in tasks dominated by Data Gap and Signal

Corruption errors. These errors are typically resolved with a single clarification, leading to high Resolved@Edge rates, 82.5% for DG and 75.4% for SC. DG makes up 36.8% and SC 30.9% of the errors across datasets, representing the majority of the issues encountered. For tasks involving more Referential Drift and Capability Gap, such as MATH and MMLU, the improvements are more limited. These errors are less frequent, with RD accounting for 20.7% and CG 11.6% of the errors, but they are more difficult to resolve with a single clarification. As a result, their Resolved@Edge rates are lower, 58.3% for RD and 49.5% for CG.

Overall, AgentAsk is highly effective in resolving simpler edge-level errors. However, for tasks with more complex errors like RD and CG, the improvements are more modest, as these errors stem from the nature of the hallucinations. For further analysis and additional results including qualitative case studies, please refer to Appendix D.

6 Conclusion

In this work, we provide novel insights into error propagation in multi-agent systems, emphasizing the critical role of communication at agent handoffs. Our error taxonomy identifies four common error types that drive system failures, laying the groundwork for effective intervention. Building on this taxonomy, we propose AgentAsk, a lightweight, architecture-agnostic clarifier that minimizes intervention by asking targeted questions at key communication points. This approach enables error containment with minimal overhead, achieving near large-model performance. Future work will explore uncertainty-aware gating, theoretical analysis of edge-level interventions, and integration with tools and human feedback.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant 12227901, in part by the Project of Stable Support for Youth Team in Basic Research Field, Chinese Academy of Sciences (CAS) under Grant YSBR-005, and in part by the Natural Science Foundation of Jiangsu Province under Grant BK20250484. The AI-driven experiments, simulations, and model training were performed on the robotic AI-Scientist platform of CAS.

Limitations

Despite its effectiveness in mitigating many forms of error propagation, AgentAsk cannot fully eliminate all errors within multi-agent systems. Certain errors originate from internal flaws within individual agents, particularly those tied to well-known limitations in large models, such as hallucinations and inconsistencies. These issues lie beyond the scope of a clarification module like AgentAsk and can only be addressed through more fundamental advancements in the capabilities of large models. Future research will need to focus on overcoming these intrinsic model limitations, with the hope that improvements in model robustness and accuracy will enable more reliable error resolution across the system.

References

- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. 2021. [Program synthesis with large language models](#). *Preprint*, arXiv:2108.07732.
- Mert Cemri, Melissa Z Pan, Shuyi Yang, Lakshya A Agrawal, Bhavya Chopra, Rishabh Tiwari, Kurt Keutzer, Aditya Parameswaran, Dan Klein, Kannan Ramchandran, Matei Zaharia, Joseph E. Gonzalez, and Ion Stoica. 2025. [Why do multi-agent LLM systems fail?](#) In *NeurIPS 2025 Workshop on Evaluating the Evolving LLM Lifecycle: Benchmarks, Emergent Abilities, and Scaling*.
- Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2024. [Chateval: Towards better LLM-based evaluators through multi-agent debate](#). In *The Twelfth International Conference on Learning Representations*.
- Lingjiao Chen, Matei Zaharia, and James Zou. 2024a. [FrugalGPT: How to use large language models while reducing cost and improving performance](#). *Transactions on Machine Learning Research*.
- Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. 2024b. [Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors](#). 2024:20094–20136.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 181 others. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- J.H. Gittell. 2006. Relational coordination: Coordinating work through relationships of shared goals, shared knowledge and mutual respect. In *Relational Perspectives in Organizational Studies: A Research Companion*, pages 74 – 94. Edward Elgar Publishing Ltd.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021b. [Measuring mathematical problem solving with the MATH dataset](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. 2024. [MetaGPT: Meta programming for a multi-agent collaborative framework](#). In *The Twelfth International Conference on Learning Representations*.
- Dongryeol Lee, Segwang Kim, Minwoo Lee, Hwanhee Lee, Joonsuk Park, Sang-Woo Lee, and Kyomin Jung. 2023. [Asking clarification questions to handle ambiguity in open-domain QA](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11526–11544, Singapore. Association for Computational Linguistics.
- Hyunseok Lee, Seunghyuk Oh, Jaehyung Kim, Jinwoo Shin, and Jihoon Tack. 2025. [ReVISE: Learning to refine at test-time via intrinsic self-verification](#). In *Forty-second International Conference on Machine Learning*.

- Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. [CAMEL: Communicative agents for "mind" exploration of large language model society](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2024. [Encouraging divergent thinking in large language models through multi-agent debate](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17889–17904, Miami, Florida, USA. Association for Computational Linguistics.
- Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, and 3 others. 2024. [Agentbench: Evaluating LLMs as agents](#). In *The Twelfth International Conference on Learning Representations*.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. [Self-refine: Iterative refinement with self-feedback](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 46534–46594. Curran Associates, Inc.
- Subhojyoti Mukherjee, Viet Dac Lai, Raghavendra Ad-danki, Ryan Rossi, Seunghyun Yoon, Trung Bui, Anup Rao, Jayakumar Subramanian, and Branislav Kveton. 2025. [Learning to clarify by reinforcement learning through reward-weighted fine-tuning](#). *Preprint*, arXiv:2506.06964.
- OpenAI. 2024. <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>. Accessed: 2025-05-10.
- OpenAI. 2025. <https://openai.com/index/introducing-gpt-5/>. Accessed: 2025-08-07.
- Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. [Generative agents: Interactive simulacra of human behavior](#). In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, UIST '23, New York, NY, USA. Association for Computing Machinery.
- Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024. [ChatDev: Communicative agents for software development](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15174–15186, Bangkok, Thailand. Association for Computational Linguistics.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language models can teach themselves to use tools](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 68539–68551. Curran Associates, Inc.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. [Reflection: language agents with verbal reinforcement learning](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 8634–8652. Curran Associates, Inc.
- Kun Wang, Guibin Zhang, Zhenhong Zhou, Jiahao Wu, Miao Yu, Shiqian Zhao, Chenlong Yin, Jinhu Fu, Yibo Yan, Hanjun Luo, Liang Lin, Zhihao Xu, Haolang Lu, Xinye Cao, Xinyun Zhou, Weifei Jin, Fanci Meng, Shicheng Xu, Junyuan Mao, and 84 others. 2025. [A comprehensive survey in llm\(-agent\) full stack safety: Data, training and deployment](#). *Preprint*, arXiv:2504.15585.
- Wenxuan Wang, Juluan Shi, Chaozheng Wang, Cheryl Lee, Youliang Yuan, Jen tse Huang, and Michael R. Lyu. 2024. [Learning to ask: When llms meet unclear instruction](#). *CoRR*, abs/2409.00557.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Anita Williams Woolley, Christopher F. Chabris, Alex Pentland, Nada Hashmi, and Thomas W. Malone. 2010. [Evidence for a collective intelligence factor in the performance of human groups](#). *Science*, 330(6004):686–688.
- Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and Chi Wang. 2024a. [Autogen: Enabling next-gen LLM applications via multi-agent conversations](#). In *First Conference on Language Modeling*.
- Zhenyu Wu, Qingkai Zeng, Zhihan Zhang, Zhaoxuan Tan, Chao Shen, and Meng Jiang. 2024b. [Large language models can self-correct with key condition verification](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 12846–12867, Miami, Florida, USA. Association for Computational Linguistics.
- Andrea Wynn, Harsh Satija, and Gillian Hadfield. 2025. [Talk isn't always cheap: Understanding failure modes in multi-agent debate](#). *Preprint*, arXiv:2509.05396.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao,

- Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. [React: Synergizing reasoning and acting in language models](#). *Preprint*, arXiv:2210.03629.
- Miao Yu, Fanci Meng, Xinyun Zhou, Shilong Wang, Junyuan Mao, Linsey Pan, Tianlong Chen, Kun Wang, Xinfeng Li, Yongfeng Zhang, Bo An, and Qingsong Wen. 2025a. [A survey on trustworthy llm agents: Threats and countermeasures](#). In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*, KDD '25, page 6216–6226, New York, NY, USA. Association for Computing Machinery.
- Miao Yu, Shilong Wang, Guibin Zhang, Junyuan Mao, Chenlong Yin, Qijiong Liu, Kun Wang, Qingsong Wen, and Yang Wang. 2025b. [NetSafe: Exploring the topological safety of multi-agent system](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 2905–2938, Vienna, Austria. Association for Computational Linguistics.
- Yanwei Yue, Guibin Zhang, Boyang Liu, Guancheng Wan, Kun Wang, Dawei Cheng, and Yiyang Qi. 2025. [Masrouter: Learning to route llms for multi-agent systems](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15549–15572. Association for Computational Linguistics.
- Guibin Zhang, Luyang Niu, Junfeng Fang, Kun Wang, Lei Bai, and Xiang Wang. 2025a. [Multi-agent architecture search via agentic supernet](#). In *Proceedings of the 42nd International Conference on Machine Learning (ICML 2025)*. Oral.
- Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xiong-Hui Chen, Jiaqi Chen, Mingchen Zhuge, Xin Cheng, Sirui Hong, Jinlin Wang, Bingnan Zheng, Bang Liu, Yuyu Luo, and Chenglin Wu. 2025b. [Aflow: Automating agentic workflow generation](#). In *The Thirteenth International Conference on Learning Representations (ICLR 2025)*. Oral.
- Michael JQ Zhang, W. Bradley Knox, and Eunsol Choi. 2025c. [Modeling future conversation turns to teach LLMs to ask clarifying questions](#). In *The Thirteenth International Conference on Learning Representations*.
- Shaokun Zhang, Ming Yin, Jieyu Zhang, Jiale Liu, Zhiguang Han, Jingyang Zhang, Beibin Li, Chi Wang, Huazheng Wang, Yiran Chen, and Qingyun Wu. 2025d. [Which agent causes task failures and when? on automated failure attribution of LLM multi-agent systems](#). In *Forty-second International Conference on Machine Learning*.
- Kunlun Zhu, Hongyi Du, Zhaochen Hong, Xiaocheng Yang, Shuyi Guo, Zhe Wang, Zhenhailong Wang, Cheng Qian, Xiangru Tang, Heng Ji, and Jiaxuan You. 2025. [Multiagentbench: Evaluating the collaboration and competition of llm agents](#). *Preprint*, arXiv:2503.01935.
- Mingchen Zhuge, Wenyi Wang, Louis Kirsch, Francesco Faccio, Dmitrii Khizbullin, and Jürgen Schmidhuber. 2024. [GPTSwarm: Language agents as optimizable graphs](#). 235:62743–62767.

Contents

1	Introduction	1
2	Related Work	2
3	Error Taxonomy	3
3.1	Data Gap	3
3.2	Referential Drift	3
3.3	Signal Corruption	3
3.4	Capability Gap	3
3.5	Inspiration from Edge-level Taxonomy	4
4	Design of AgentAsk	4
4.1	Clarification as Edge Intervention	4
4.2	Learning to Ask from Demonstrations	4
4.3	Clarification Optimization with Reinforcement Learning	5
5	Experiments	6
5.1	Experimental Setup	6
5.2	RQ1: Effectiveness under Fixed Orchestration	6
5.3	RQ2: Efficiency under Practical Budgets	7
5.4	RQ3: Mechanism and Robustness at Edge	7
5.5	RQ4: Taxonomy Mechanisms at Edge	8
6	Conclusion	8
A	Supplementary Design Explanation	12
A.1	Modeling Details	12
A.2	Supervised Fine-tuning Details	12
B	Algorithm Workflow	14
C	Supplementary Experiments	14
C.1	RQ1: Effectiveness under Fixed Orchestration	14
C.2	RQ2: Efficiency under Practical Budgets	14
C.3	RQ3: Ablations and Robustness	14
D	Case Study of Error Taxonomy	16
D.1	Data Gap (DG)	16
D.2	Signal Corruption (SC)	16
D.3	Referential Drift (RD)	17
D.4	Capability Gap (CG)	17
E	Notations	17

F Prompt

17

1 A Supplementary Design Explanation

2 A.1 Modeling Details

We elaborate on the controller framework referenced in Section 4.1.

Edge State. Each input produces a message graph $G = (V, E)$ unrolled into an edge sequence $\{e_t\}_{t=1}^T$, where each $e_t = (u_t \rightarrow v_t)$ denotes a message handoff. The controller observes a local state at each step:

$$x_t = (x_t^{\text{in}}, u_t, v_t, m_t, h_t), \quad (9)$$

where x_t^{in} is the user query or subgoal, m_t is the message from u_t , and h_t encodes recent local interaction context.

Controller Action. Based on x_t , the controller selects an action:

$$a_t = (z_t, \tilde{v}_t, q_t), \quad (10)$$

where z_t is a binary gate indicating whether to ask, \tilde{v}_t selects which agent to query, and q_t is the generated clarification question.

Policy Structure. We define the controller policy as:

$$\pi_\theta(a_t | x_t) = \pi_\theta(z_t | x_t) [\pi_\theta(\tilde{v}_t | x_t) \pi_\theta(q_t | x_t)]^{z_t}. \quad (11)$$

Trajectory Modeling. The full input interaction produces a trajectory:

$$\tau = \{(x_t, a_t, r_t)\}_{t=1}^T, \quad (12)$$

where r_t is a response returned when $z_t = 1$. The objective from Section 4.1 is instantiated here with:

- $U(\tau)$ rewarding task success and helpful clarifications,
- $C(\tau)$ measuring token and latency costs,
- B as the intervention budget.

This setup supports both imitation-based initialization and constrained reinforcement learning (see Appendix A.2).

A.2 Supervised Fine-tuning Details

We detail the data construction and training procedure used in Section 4.2.

Algorithm 1 AgentAsk Training

Require: Logs \mathcal{L} , teacher \mathcal{J} , init policy π_θ , ref π_{ref} , budget B , lr η , λ_{ask} , λ_R , ϵ , β , weights w_t , baseline b

```
1: Stage A: Build Corpus
2: for trajectory  $\tau \in \mathcal{L}$  do
3:   for edge state  $x_i \in \tau$  do
4:      $y_i \leftarrow \mathcal{J}(x_i) = (t_i, v_i, q_i)$   $\triangleright t_i \in \{\text{DG, SC, RD, CG, NONE}\}$ 
5:      $\tilde{\mathcal{D}} \leftarrow \tilde{\mathcal{D}} \cup \{(x_i, y_i)\}$ 
6:   end for
7: end for
8: Stage B: Learning from Demonstrations
9: repeat
10:   Sample minibatch  $\mathcal{B} \subset \tilde{\mathcal{D}}$ 
11:   Compute heads  $p_\theta^{\text{type}}(t|x)$ ,  $p_\theta^{\text{addr}}(v|x, t)$ ,  $p_\theta^{\text{txt}}(q|x, t)$ 
12:   Evaluate  $\mathcal{L}_{\text{type}}$  and  $\mathcal{L}_{\text{ask}}$ ; set  $\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{type}} + \lambda_{\text{ask}}\mathcal{L}_{\text{ask}}$ 
13:    $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}_{\text{Total}}$ 
14: until validation stops improving
15: Stage C: Clarification Optimization with RL
16: repeat
17:   Init episode, counters, buffers; set budget  $B$ 
18:   for  $t = 1, \dots, T$  do
19:     Observe  $x_t = (x_t^{\text{in}}, u_t, v_t, m_t, h_t)$ 
20:     Sample  $a_t = (z_t, \tilde{v}_t, q_t) \sim \pi_\theta(\cdot | x_t)$  (schema+budget)
21:     Execute  $a_t$ ; if  $z_t=1$  get reply  $r_t$ ; deliver to  $v_t$ ; obtain  $x_{t+1}$ 
22:     Compute  $r_t^{\text{eff}}, r_t^{\text{par}}, r_t^{\text{fmt}}, r_t^{\text{edge}}$ 
23:     if episode not terminated then
24:       Sample  $x_t$ ; score by  $r_t^{\text{edge}}$ 
25:       Form local advantages  $A_t^{\text{loc}}$ ; update with ratios  $\rho_t$ 
26:     else
27:       Compute terminal  $R$ ; set  $A_{t'}^{\text{glob}} \leftarrow w_{t'}(R - b)$  for all visited  $t'$ 
28:       Apply E-GRPO update using  $A_t^{\text{loc}}$  and  $A_{t'}^{\text{glob}}$  with KL to  $\pi_{\text{ref}}$ 
29:     break
30:   end if
31: end for
32: until RL budget exhausted or convergence
33: return  $\pi_\theta$ 
```

Dataset Construction. From a set of logged multi-agent executions, we extract edge-local contexts x_i and assign expert clarifications via a teacher model. Each label $y_i = (t_i, v_i, q_i)$ encodes:

- $t_i \in \mathcal{T}$, a clarification type (or NONE if no ask is needed),
- $v_i \in \{u_i, v_i\}$, the agent selected to receive the question,
- q_i , the clarification question string, constrained by schema and token length.

We define $m_i = \mathbb{1}[t_i \neq \text{NONE}]$ as an indicator for whether a clarification is present.

Model Outputs. The controller predicts:

- $p_\theta^{\text{type}}(t_i | x_i)$ for the ask decision,
- $p_\theta^{\text{addr}}(v_i | x_i, t_i)$ for the addressee (when asking),
- $p_\theta^{\text{txt}}(q_i | x_i, t_i)$ for the clarification text via auto-regressive decoding.

Training Objective. The supervised loss has two components:

- $\mathcal{L}_{\text{type}}$ is the cross-entropy loss on t_i ,
- \mathcal{L}_{ask} is the conditional loss on v_i and q_i , applied only when $m_i = 1$:

$$\mathcal{L}_{\text{ask}} = \frac{1}{N} \sum_{i=1}^N m_i \left[-\log p_{\theta}^{\text{addr}}(v_i | x_i, t_i) - \sum_{t=1}^{T_i} \log p_{\theta}^{\text{txt}}(q_{i,t} | x_i, t_i, q_{i,<t}) \right]. \quad (13)$$

The total training loss is illustrated at Equation 2 where λ_{ask} controls the relative weight of question generation.

This training yields a controller capable of structured clarification grounded in edge-level context.

B Algorithm Workflow

This appendix concisely instantiates the edge controller in Eq. (11) under the constraint in Eq. (1). The details can be seen in Section 4. The specific algorithm process is shown in Algorithm 1.

C Supplementary Experiments

C.1 RQ1: Effectiveness under Fixed Orchestration

As shown in Table 4, we present additional results evaluating the performance of AgentAsk across all frameworks and benchmarks under a random seed of 1314 and a model temperature of 0.3 settings. These results reinforce the trends observed in the main text, showing that AgentAsk consistently improves performance, particularly in terms of accuracy, latency, and extra cost. The accuracy improvement across the benchmarks, including GSM8K, MATH, HumanEval, MMLU, and MBPP, is substantial, with AgentAsk outperforming the baseline methods in all cases. For example, AgentAsk improves the average accuracy by +3.34% on the GPTSwarm framework and achieves a similar performance boost on other benchmarks, while keeping latency and extra costs within acceptable bounds. These findings validate the effectiveness of AgentAsk in improving the performance of multi-agent systems under unchanged orchestration conditions.

C.2 RQ2: Efficiency under Practical Budgets

Tables 5, 6, and 7 summarize the performance of AgentAsk, +GPT-5, and origin across five benchmarks. The results demonstrate that AgentAsk not only outperforms the original setting but also remains much more efficient compared to +GPT-5. Specifically, AgentAsk reduces the computational overhead, as evidenced by its lower extra costs, while maintaining similar or better accuracy across

all frameworks. For example, in the GPTSwarm framework, AgentAsk achieves a +2.32% +4.34% improvement in accuracy compared to the baseline while reducing extra costs significantly. This confirms that AgentAsk provides a practical and efficient solution that achieves competitive accuracy while staying within practical operational budgets.

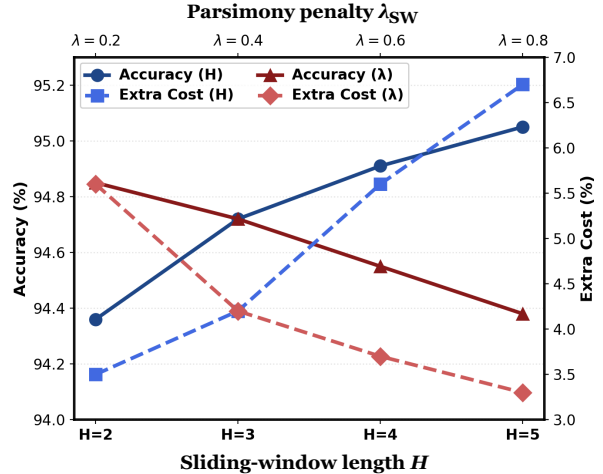


Figure 5: Sensitivity to window H and penalty λ_{sw} , highlighting a stable region near the default and the accuracy–efficiency trade-off.

C.3 RQ3: Ablations and Robustness

Table 8 presents the results of an ablation study that investigates the impact of different reward components (r^{par} , r^{eff} , r^{fmt}) on the performance of AgentAsk. The ablation results highlight the importance of each reward component in balancing performance (accuracy) and efficiency (latency and extra cost). Removing any of the reward components leads to a noticeable drop in accuracy and an increase in latency or extra cost. For instance, excluding r^{par} results in a decrease of 0.20% in accuracy and an increase in latency by 8%, demonstrating the necessity of including all reward components for optimal performance.

Additionally, we compare the performance of AgentAsk with other reinforcement learning-based methods such as PPO and GRPO across five benchmarks, as shown in Table 9. The results clearly indicate that AgentAsk consistently outperforms both PPO and GRPO in terms of accuracy, while maintaining a relatively low latency and extra cost. In particular, on the MMLU and GSM8K benchmarks, AgentAsk achieves significant improvements in accuracy (+2.25% and +3.18%, respectively), highlighting its robustness and effectiveness. These findings demonstrate that AgentAsk provides a re-

Method	Settings	GSM8K	MATH	HumanEval	MMLU	MBPP	Average
IO		89.52	49.23	89.66	79.63	71.38	75.88
CoT		89.71	48.19	90.64	80.44	70.66	75.93
Self-Refine		90.37	47.92	88.88	78.31	71.49	75.40
GPTSwarm	origin	91.82	50.92	90.21	80.92	76.38	77.99
	+GPT-5	95.35 \uparrow +3.53	56.15 \uparrow +5.23	93.35 \uparrow +3.14	86.20 \uparrow +5.28	82.05 \uparrow +5.67	82.62 \uparrow +4.63
	+AgentAsk	94.17 \uparrow +2.35	54.97 \uparrow +4.05	92.41 \uparrow +2.20	84.62 \uparrow +3.70	80.48 \uparrow +4.10	81.33 \uparrow +3.34
AFlow	origin	91.02	50.51	91.22	81.60	79.14	78.70
	+GPT-5	94.85 \uparrow +3.83	56.05 \uparrow +5.54	94.60 \uparrow +3.38	86.65 \uparrow +5.05	84.75 \uparrow +5.61	83.38 \uparrow +4.68
	+AgentAsk	93.12 \uparrow +2.10	54.71 \uparrow +4.20	93.57 \uparrow +2.35	85.25 \uparrow +3.65	83.04 \uparrow +3.90	81.94 \uparrow +3.24
MaAS	origin	92.32	50.82	91.42	82.41	80.80	79.75
	+GPT-5	95.40 \uparrow +3.08	56.70 \uparrow +5.88	94.85 \uparrow +3.43	87.10 \uparrow +4.69	86.20 \uparrow +5.40	84.05 \uparrow +4.30
	+AgentAsk	92.12 \downarrow -0.20	55.42 \uparrow +4.60	93.67 \uparrow +2.25	85.81 \uparrow +3.40	84.55 \uparrow +3.75	82.31 \uparrow +2.56
MasRouter	origin	92.82	51.10	90.72	82.60	78.58	79.16
	+GPT-5	96.00 \uparrow +3.18	56.25 \uparrow +5.15	93.95 \uparrow +3.23	86.90 \uparrow +4.30	84.60 \uparrow +6.02	83.54 \uparrow +4.38
	+AgentAsk	95.07 \uparrow +2.25	54.95 \uparrow +3.85	92.82 \uparrow +2.10	85.70 \uparrow +3.10	82.83 \uparrow +4.25	82.27 \uparrow +3.11

Table 4: Accuracy/Pass@1 across frameworks under three settings (origin, +GPT-5, +AgentAsk) on five benchmarks. Bottom-right overlays denote percentage-point deltas vs. each framework’s origin. Each cell overlays a bottom-right delta (green \uparrow / red \downarrow) relative to the framework’s origin.

Settings	GSM8K			MATH			HumanEval			MMLU			MBPP			
	Acc.	Lat.	Extra	Acc.	Lat.	Extra	Acc.	Lat.	Extra	Acc.	Lat.	Extra	Acc.	Lat.	Extra	
origin	92.18	100	0.0	51.41	100	0.0	90.63	100	0.0	81.26	100	0.0	76.81	100	0.0	
+GPT-4o-mini	93.50 \uparrow +1.32	117	16.0	51.84 \uparrow +0.43	115	14.8	91.92 \uparrow +1.29	121	17.6	82.00 \uparrow +0.74	114	14.2	78.30 \uparrow +1.49	123	18.5	
+GPT-5	95.55 \uparrow +3.37	133	37.2	56.30 \uparrow +4.89	131	35.8	94.05 \uparrow +3.42	139	40.4	86.85 \uparrow +5.59	129	33.8	82.45 \uparrow +5.64	140	41.2	
AgentAsk (Backbone: Llama-3.2-3B)																
SFT	92.60 \uparrow +0.42	106	5.6	51.55 \uparrow +0.14	106	5.4	91.25 \uparrow +0.62	106	5.6	81.65 \uparrow +0.39	105	5.0	76.70 \downarrow -0.11	107	5.7	
(E-GRPO)	92.88 \uparrow +0.70	105	5.0	51.65 \uparrow +0.24	105	5.1	91.50 \uparrow +0.87	105	5.2	81.90 \uparrow +0.64	104	4.7	77.05 \uparrow +0.24	106	5.2	
AgentAsk (Backbone: Qwen-3-4B)																
SFT	92.85 \uparrow +0.67	106	5.5	51.41	0.00	106	5.3	91.40 \uparrow +0.77	105	5.3	81.80 \uparrow +0.54	105	5.0	77.05 \uparrow +0.24	106	5.5
(E-GRPO)	94.52 \uparrow +2.34	116	9.7	55.05 \uparrow +3.64	106	9.2	92.95 \uparrow +2.32	103	3.3	85.25 \uparrow +3.99	119	8.6	81.15 \uparrow +4.34	105	6.1	

Table 5: GPTSwarm: Ablations on model and training across five benchmarks.

Settings	GSM8K			MATH			HumanEval			MMLU			MBPP			
	Acc.	Lat.	Extra	Acc.	Lat.	Extra	Acc.	Lat.	Extra	Acc.	Lat.	Extra	Acc.	Lat.	Extra	
origin	92.84	100	0.0	51.08	100	0.0	92.02	100	0.0	82.93	100	0.0	81.22	100	0.0	
+GPT-4o-mini	93.95 \uparrow +1.11	118	15.7	51.80 \uparrow +0.72	116	14.9	92.80 \uparrow +0.78	120	16.6	83.40 \uparrow +0.47	115	14.3	82.10 \uparrow +0.88	122	17.3	
+GPT-5	96.20 \uparrow +3.36	127	44.3	56.85 \uparrow +5.77	132	27.7	95.45 \uparrow +3.43	121	17.8	88.35 \uparrow +5.42	130	31.0	86.70 \uparrow +5.48	129	35.6	
AgentAsk (Backbone: Llama-3.2-3B)																
SFT	92.90 \uparrow +0.06	106	5.4	51.35 \uparrow +0.27	106	5.2	91.90 \downarrow -0.12	107	5.6	83.05 \uparrow +0.12	105	4.9	81.60 \uparrow +0.38	107	5.6	
(E-GRPO)	93.00 \uparrow +0.16	105	5.0	51.60 \uparrow +0.52	105	5.0	92.20 \uparrow +0.18	106	5.3	83.18 \uparrow +0.25	104	4.7	81.85 \uparrow +0.63	106	5.3	
AgentAsk (Backbone: Qwen-3-4B)																
SFT	92.84	0.00	106	5.4	51.90 \uparrow +0.82	106	5.3	92.10 \uparrow +0.08	106	5.5	82.85 \downarrow -0.08	105	5.0	81.70 \uparrow +0.48	106	5.4
(E-GRPO)	95.10 \uparrow +2.26	108	5.6	55.55 \uparrow +4.47	112	7.4	94.45 \uparrow +2.43	104	5.3	86.95 \uparrow +4.02	117	5.7	85.30 \uparrow +4.08	107	5.0	

Table 6: MaAS: Ablations on model and training across five benchmarks.

Settings	GSM8K			MATH			HumanEval			MMLU			MBPP		
	Acc.	Lat.	Extra	Acc.	Lat.	Extra	Acc.	Lat.	Extra	Acc.	Lat.	Extra	Acc.	Lat.	Extra
origin	93.26	100	0.0	51.52	100	0.0	91.03	100	0.0	83.19	100	0.0	79.04	100	0.0
+GPT-4o-mini	94.62 \uparrow +1.36	118	16.0	52.10 \uparrow +0.58	118	15.7	92.10 \uparrow +1.07	122	17.3	83.50 \uparrow +0.31	116	14.6	81.10 \uparrow +2.06	121	17.9
+GPT-5	95.10 \uparrow +1.84	129	34.0	56.70 \uparrow +5.18	124	24.8	94.55 \uparrow +3.52	135	35.1	88.55 \uparrow +5.36	128	29.3	85.15 \uparrow +6.11	134	42.6
AgentAsk (Backbone: Llama-3.2-3B)															
SFT	93.64 \uparrow +0.38	106	5.7	51.85 \uparrow +0.33	107	5.8	91.35 \uparrow +0.32	106	5.6	83.25 \uparrow +0.06	106	5.2	79.80 \uparrow +0.76	106	5.9
(E-GRPO)	94.23 \uparrow +0.97	105	5.0	51.95 \uparrow +0.43	106	5.4	91.50 \uparrow +0.47	105	5.2	83.35 \uparrow +0.16	105	5.0	80.20 \uparrow +1.16	105	5.4
AgentAsk (Backbone: Qwen-3-4B)															
SFT	93.99 \uparrow +0.73	105	5.3	51.90 \uparrow +0.38	106	5.2	91.45 \uparrow +0.42	105	5.1	83.10 \downarrow -0.09	104	4.8	80.10 \uparrow +1.06	104	4.9
(E-GRPO)	94.86 \uparrow +1.60	108	4.9	55.65 \uparrow +4.13	109	6.1	93.35 \uparrow +0.52	106	4.6	87.10 \uparrow +2.32	111	5.7	83.55 \uparrow +4.51	103	7.7

Table 7: MasRouter: Ablations on model and training across five benchmarks.

liable and efficient solution across different configurations and benchmarks, offering a promising approach for error mitigation in multi-agent sys-

tems. Finally, Figure 5 illustrates the sensitivity of **AgentAsk** to the window size H and the parsimony weight λ_{sw} , showing that the clarifier is

Framework	Setting	GSM8K			MATH			HumanEval			MMLU			MBPP		
		Acc.	Lat.	Extra	Acc.	Lat.	Extra	Acc.	Lat.	Extra	Acc.	Lat.	Extra	Acc.	Lat.	Extra
GPTSwarm	Ours	94.52	116	9.7	55.05	106	9.2	92.95	103	3.3	85.25	119	8.6	81.15	105	6.1
	w/o r^{par}	93.10	126	6.2	52.38	116	7.0	91.32	113	7.0	82.33	118	6.5	78.07	115	8.3
	w/o r^{eff}	92.84	117	6.8	52.75	125	7.6	91.85	109	5.9	83.02	117	7.3	82.14	109	7.2
	w/o r^{fmt}	92.96	133	8.7	52.01	134	7.8	91.02	115	7.1	82.14	121	7.2	78.44	116	7.0
	R only	92.05	129	71.2	51.90	133	16.5	90.82	125	16.4	81.74	125	13.5	77.33	123	12.1
MasRouter	Ours	94.86	108	4.9	55.65	109	6.1	93.35	106	4.6	87.10	111	5.7	83.55	103	7.7
	w/o r^{par}	94.61	113	6.4	52.98	117	7.2	92.48	112	6.7	83.74	116	6.9	81.34	118	9.6
	w/o r^{eff}	94.70	111	6.9	53.22	115	7.4	92.91	109	5.9	83.88	113	6.3	80.05	108	7.0
	w/o r^{fmt}	94.51	129	29.8	52.44	138	27.7	92.15	119	16.8	82.90	128	17.0	80.42	111	9.7
	R only	93.21	124	24.0	52.14	135	81.6	91.54	125	48.3	83.01	124	29.3	79.66	123	13.2
MaAS	Ours	95.10	108	5.6	55.55	112	7.4	94.45	104	5.3	86.95	108	5.7	85.30	107	5.0
	w/o r^{par}	93.41	113	6.3	52.06	117	6.8	92.10	110	6.2	83.20	116	6.4	81.55	112	6.3
	w/o r^{eff}	93.12	110	6.8	52.55	115	7.1	92.37	108	5.8	83.02	114	6.7	81.40	109	6.0
	w/o r^{fmt}	92.98	114	7.4	51.74	118	7.6	91.61	112	6.6	82.60	118	7.0	80.44	113	6.8
	R only	92.55	125	25.3	51.60	128	14.4	92.14	121	13.6	83.02	114	9.4	81.40	119	13.6

Table 8: Ablation results (Acc./Latency/Extra) across five benchmarks.

Method	Settings	GSM8K	MATH	HumanEval	MMLU	MBPP	Average
GPTSwarm	PPO	92.18	51.41	90.63	81.26	76.81	78.46
	GRPO	93.10	52.40	91.80	82.40	79.10	79.76
	Ours	94.52	55.05	92.95	85.25	81.15	81.78
AFlow	PPO	91.35	50.89	91.67	82.16	79.63	79.14
	GRPO	93.00	52.20	92.50	83.00	81.10	80.36
	Ours	93.70	54.95	93.95	86.40	83.90	82.58
MaAS	PPO	92.84	51.08	92.02	82.93	81.22	80.02
	GRPO	93.80	52.40	92.90	83.80	82.10	81.00
	Ours	95.10	55.55	94.45	86.95	85.30	83.47
MasRouter	PPO	93.26	51.52	91.03	83.19	79.04	79.61
	GRPO	94.70	52.80	92.70	84.10	81.90	81.24
	Ours	94.86	55.65	93.35	87.10	83.55	82.54

Table 9: Performance comparison across PPO, GRPO, and our method on five benchmarks.

robust to parameter variations, with stable performance across a range of settings. This confirms that **AgentAsk** is both effective and adaptable.

D Case Study of Error Taxonomy

The following traces provide concrete grounding for the edge-level error taxonomy established in Section 3. Each case is presented in a consistent four-column format as shown in Tables 10, 11, 12, and 13, detailing the **Query**, the original agent output (**Origin**), the clarification intervention by **AgentAsk**, and the resulting **Revised Message**. Figure 6 provides a visual overview and index of our edge-level error taxonomy. Its four quadrants each sketch a minimal, real trace exemplifying one of the four core error types (*i.e.*, DG, RD, SC, and CG), while the central donut reports their prevalence in our annotated logs ($N=824$). This figure serves as a conceptual guide and navigational aid,

directing the reader to the corresponding comprehensive case studies in Tables 10–13. These tables provide the complete, granular detail of each interaction, including the original faulty message, the precise clarifying question posed by *AgentAsk*, and the subsequent corrected handoff, enabling a thorough examination of the error mitigation process.

D.1 Data Gap (DG)

Missing details at an edge force downstream agents to guess, which opens a cascade. The clarification asks for the exact boundary behavior that is absent at x_t and removes failure before the message flows further. The details can be seen in Table 10.

D.2 Signal Corruption (SC)

A wrong conversion or reference enters the chain and is relayed unchanged. The clarification pins down the intended unit or base quantity, which stops the propagation locally and earns positive

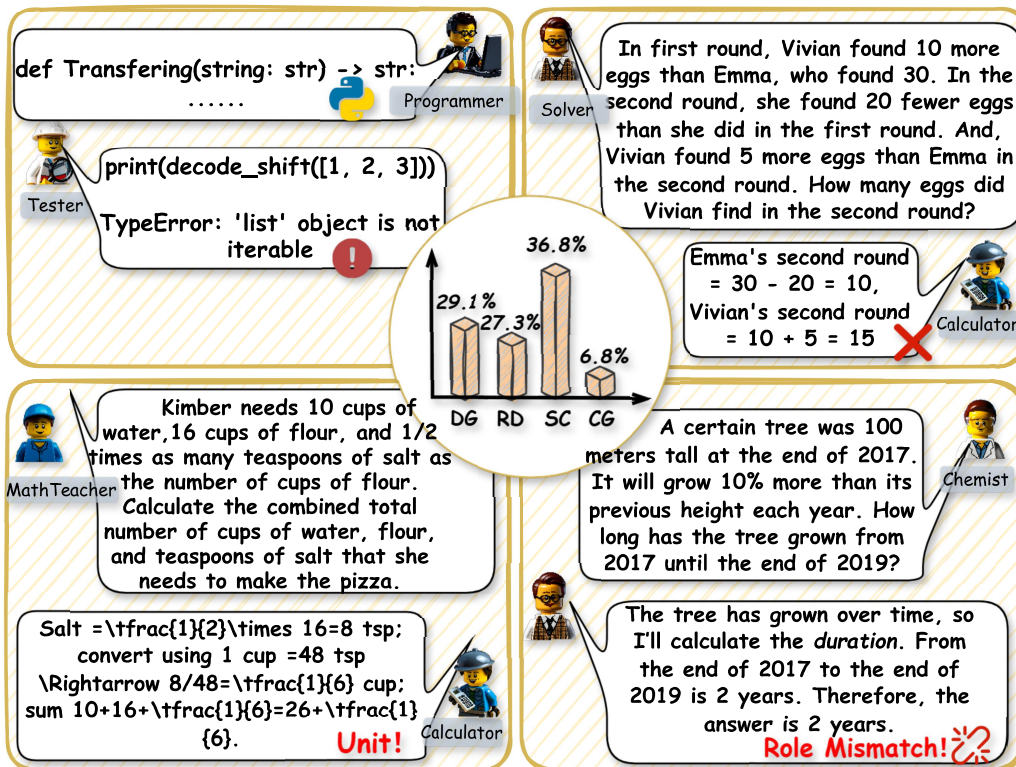


Figure 6: The case of our error taxonomy. In the middle shows the fraction of the four types of errors.

effectiveness per Eq. (3). Results are summarized in Table 11.

D.3 Referential Drift (RD)

Ambiguous symbols drift across turns and break consistency. The clarification fixes names and indices at the edge and aligns all relations, after which the downstream solution becomes stable. The details can be seen in Table 12.

D.4 Capability Gap (CG)

The role at the edge lacks the skill required by the subtask. The clarification requests a reroute to a capable recipient and specifies the expected output form, which resolves the mismatch with a single step. Table 13 shows the corrected messages.

E Notations

The symbols used in the text are explained in detail in the Table 14.

F Prompt

The complete instruction used by the clarifier is delineated in Figure 7. It explicitly specifies *when to ask*, *what to ask*, and *whom to ask*, while keeping questions short and cost-aware. Due to the inconsistency in the framework, there may be slight

differences in the input section of the Prompt, but this does not affect the overall consistency.

Table 10: **AgentAsk Case (Data Gap)**. The table shows one code-oriented edge where an upstream message omits boundary behavior: the query asks for a shortest-palindrome constructor but the origin message assumes a non-empty string and indexes without guards, creating an empty-input gap. Columns list the original query, the erroneous origin message, the minimal edge-level Ask (what/whom to ask), and the corrected handoff that adds explicit length checks and a safe return for the empty string.

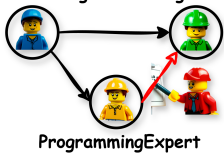
Query	Origin (Data Gap)	AgentAsk	Revised Message
<pre>def is_palindrome(string: str) -> bool: """Test if given string is a palindrome.""" return string == string[::-1] def make_palindrome(string: str) -> str: """Find the shortest palindrome that begins with a supplied string. Algorithm idea: - Find the longest postfix of supplied string that is a palindrome. - Append to the end the reverse of the prefix before that postfix. >>> make_palindrome('') """</pre>	<p><i>I will implement make_palindrome directly from the description. If the first and last characters match and the whole string is already a palindrome, I can immediately return it. Otherwise, I scan from right to left to detect the longest palindromic suffix and then append the reverse of the missing prefix. Concretely, I start with if string[0] == string[-1] and is_palindrome(string): return string, then proceed into a loop to locate the suffix boundary. This covers typical inputs like words and short phrases. The docstring examples highlight the construction for ordinary strings, so I will assume the input is a standard non-empty string and does not require a guard for corner inputs such as the empty string or single-character strings. With that in mind, I will finalize the suffix search and produce the shortest palindrome output without additional checks or preconditions.</i></p>	<p>AgentAsk → ProgrammingExpert: “Should the function define behavior for the empty string explicitly and add a quick guard for length 0 or 1 before indexing string[0]?”</p> 	<p><i>The current logic indexes string[0] and string[-1] without guarding for length, which will raise IndexError on the empty string and is unnecessary on length-1 inputs. I will add explicit handling up front: if len(string) == 0, return “”; if len(string) == 1 or is_palindrome(string), return string. Only for longer inputs will I search for the longest palindromic suffix and append the reverse of the missing prefix. This preserves correctness on normal cases while ensuring the edge case is well-defined and safe.</i></p>

Table 11: **AgentAsk Case (Signal Corruption)**. Two independent cases illustrate how a wrong conversion or misapplied base quantity can propagate along an edge. Case 1 (pizza units): the origin message forces unit unification and converts teaspoons to cups unasked, whereas the Ask clarifies “report as listed,” yielding 26 cups + 8 teaspoons (34 as count of listed quantities). Case 2 (reading remainder): the origin treats Tuesday’s $\frac{1}{4}$ as a fraction of the whole; the Ask pins it to the Monday remainder and yields 60 pages left.

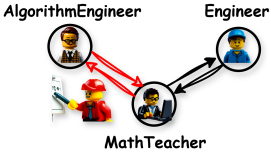
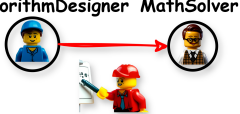
Query	Origin (Signal Corruption)	AgentAsk	Revised Message
<p>To make pizza, together with other ingredients, Kimber needs 10 cups of water, 16 cups of flour, and 1/2 times as many teaspoons of salt as the number of cups of flour. Calculate the combined total number of cups of water, flour, and teaspoons of salt that she needs to make the pizza.</p>	<p>We have 10 cups of water and 16 cups of flour. Salt equals half as many teaspoons as flour cups, so that is $\frac{1}{2} \times 16 = 8$ teaspoons. Because the prompt asks for a combined total in cups, I will convert teaspoons to cups and sum in a single unit for uniformity. Using 1 cup = 48 tsp, we get $8/48 = 1/6$ cup ≈ 0.1667 cup. Hence, the total is $10 + 16 + \frac{1}{6} = 26\frac{1}{6}$ cups ≈ 26.1667 cups. I will report the answer in cups only to avoid mixing units.</p>	<p>AgentAsk → AlgorithmEngineer: “Should teaspoons be converted to cups before calculating the total?”</p> 	<p>The prompt specifies cups for water and flour and teaspoons for salt without requiring unification. To avoid injecting an unstated policy, I will preserve units as given and compute the combined total as listed. Quantities: water = 10 cups, flour = 16 cups, salt = 8 teaspoons. Reporting the combined count of listed quantities gives 34 (26 + 8).</p>
<p>A book has 120 pages. On Monday, a student reads one-third of the book. On Tuesday, the student reads one-quarter of the remainder. How many pages are left unread after Tuesday?</p>	<p>I will compute the total pages read by treating both fractions as portions of the whole book for simplicity: Monday = $\frac{1}{3} \times 120 = 40$ pages; Tuesday = $\frac{1}{4} \times 120 = 30$ pages. Total read = 70 pages, so pages left = $120 - 70 = 50$. I will pass along “50 pages left” as the final quantity.</p>	<p>AgentAsk → MathSolver: “Does ‘one quarter’ on Tuesday refer to the remainder after Monday, rather than the whole book?”</p> 	<p>Thanks for the clarification. “One quarter” is with respect to the remainder after Monday. Compute precisely: Monday reads $\frac{1}{3} \times 120 = 40$ pages; remainder = $120 - 40 = 80$ pages. Tuesday reads $\frac{1}{4} \times 80 = 20$ pages. Total read = $40 + 20 = 60$; therefore pages left = $120 - 60 = 60$. I will correct the intermediate assumption and pass along “60 pages left” as the final quantity.</p>

Table 12: **AgentAsk Case (Referential Drift)**. A single math word-problem case where ambiguous references (T, E with dropped round indices) cause symbol drift across turns and wrong relations like $T_2 = E_1 - 20$. The Ask enforces explicit indexing (T_1, T_2, E_1, E_2) and confirms the intended relations ($T_2 = T_1 - 20, E_2 = 2T_2$), after which the downstream solution stabilizes .

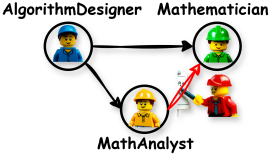
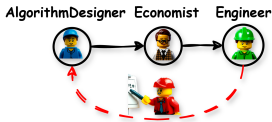
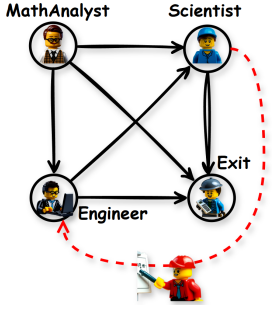
Query	Origin (Referential Drift)	AgentAsk	Revised Message
<p><i>In April, Tank gathered 10 more Easter eggs than Emma in their first round of egg hunt. Emma gathered twice as many eggs as Tank in their second round, while Tank’s second-round total was 20 less than in the first round. The combined pile with 6 other people was 400 eggs, and Emma gathered 60 eggs in the second round. Find how many eggs the 6 other egg hunters collected.</i></p>	<p>MathAnalyst → Mathematician: Let T and E denote Tank and Emma. I will use T_1, E_1 for round one and T_2, E_2 for round two, but for brevity I may drop the subscripts. From the text, “Tank’s second round was 20 less than she had gathered in the first round,” so the difference between Tank’s second and first rounds is 20. Emma’s second round is “twice as many as Tank in their second round,” which suggests a doubling relative to Tank’s performance. Since Emma has 60 in the second round, we may set $E_2 = 60$. Because $E_1 - T_2 = 20$, we can write $T_2 = E_1 - 20$. Also, “Tank gathered 10 more than Emma in their first round” gives $T_1 = E_1 + 10$. With these, the total across two rounds plus the six others should reach 400; I pass these relations forward for solution.</p>	<p>AgentAsk → MathAnalyst: “Please confirm unambiguous referents: Does ‘20 less’ apply to Tank’s own rounds ($T_2 = T_1 - 20$)? Does ‘twice as many’ mean $E_2 = 2T_2$?”</p> 	<p><i>Fix symbols explicitly: T_1, T_2 for Tank (rounds 1,2) and E_1, E_2 for Emma (rounds 1,2). Map statements to equations: $T_1 = E_1 + 10$; $T_2 = T_1 - 20$; $E_2 = 2T_2$; and the given $E_2 = 60$. From $E_2 = 60$ we get $T_2 = 30$. Then $T_1 = 50$. Using $T_1 = E_1 + 10$ gives $E_1 = 40$. The subtotal for Tank and Emma is $T_1 + T_2 + E_1 + E_2 = 50 + 30 + 40 + 60 = 180$. Therefore, the six other hunters collected $400 - 180 = 220$ eggs. The earlier drift came from writing $T_2 = E_1 - 20$, which incorrectly bound the “20 less” clause to Emma’s first round; keeping indices attached eliminates the ambiguity and yields a consistent solution.</i></p>

Table 13: **AgentAsk Case (Capability Gap)**. Two cases where the role at the edge is misaligned with the subtask. Case 1 (tree growth): an Economist returns elapsed time (“2 years”) instead of the requested growth amount; the Ask reroutes to AlgorithmEngineer, who computes compound growth (21 meters). Case 2 (coding): a non-coding role provides a qualitative description instead of an implementation; the Ask redirects to ProgrammingExpert, who supplies a correct function and complexity note.

Query	Origin (Capability Gap)	AgentAsk	Revised Message
<p>A certain tree was 100 meters tall at the end of 2017. It will grow 10% more than its previous height each year. How long has the tree grown from 2017 until the end of 2019?</p>	<p>Economist → Engineer: The question asks about the period from 2017 to the end of 2019, which spans two calendar years. Interpreting “how long has the tree grown” as the elapsed time, we simply count the years: 2018 and 2019. Therefore, the tree has grown for a total of 2 years. The relevant arithmetic is duration-based rather than quantitative growth, so the final answer is 2, reflecting the time interval rather than the magnitude of height change.</p>	<p>AgentAsk → AlgorithmEngineer: “The user asks for the growth amount (height increase), not the elapsed time. Please route to AlgorithmEngineer or Scientist for compound-growth computation.”</p> 	<p>AlgorithmEngineer: We compute the compound growth from end-2017 baseline $H_0 = 100$. Each year multiplied by 1.10. After two years: $H_2 = 100 \times (1.10)^2 = 121$ meters. The growth amount is $\Delta H = H_2 - H_0 = 121 - 100 = 21$ meters. Thus, the correct quantity is a height difference, not a duration; the previously provided “2 years” reflects a capability mismatch rather than the requested quantitative outcome. Final: the tree has grown by 21 meters over the period.</p>
<p>Write a function <code>sum_even_squares(n: int) -> int</code> that returns the sum of squares of all even integers from 1 to n inclusive. Provide a concise implementation and explain its complexity.</p>	<p>MathAnalyst → Scientist: To address the request, we consider how even numbers behave in general and recall that squaring preserves parity. The sum can conceptually be obtained by repeatedly adding the square of each even number up to the bound. The main idea is to iterate across the relevant subset and aggregate the contributions. From a scientific perspective, the computation is straightforward and should be feasible, though nuances about performance might emerge for very large inputs. In summary, the result equals the accumulated squared values of all even integers no greater than n.</p>	<p>AgentAsk → Engineer: “The user explicitly requests an implementation and complexity discussion. Ensure the function handles small n.”</p> 	<pre>def sum_even_squares(n: int) -> int: """Return sum of k^2 for even k in [1..n].""" if n < 2: return 0 # Iterate by step 2 over even numbers only return sum(k*k for k in range(2, n+1, 2))</pre> <p>ProgrammingExpert: The function iterates over even integers only, avoiding unnecessary checks, and sums their squares. For example, $n=6$ uses $\{2, 4, 6\}$ giving $4+16+36 = 56$. Time complexity is $O(n/2) = O(n)$ with $O(1)$ auxiliary space. A closed-form alternative via $\sum (2i)^2 = 4 \sum i^2$ for $i = 1..[n/2]$ can reduce iteration count with arithmetic. This reroute resolves the capability gap by assigning code synthesis to the appropriate role.</p>

Symbol	Description
N	The total number of examples in the training dataset.
x_i	The feature vector representing the state at edge i .
y_i	The intervention decision at edge i , encoding whether to ask, whom to ask, and what to ask.
\mathcal{D}	The training dataset consisting of edge-local examples (x_i, y_i) .
$\mathcal{L}_{\text{Total}}$	The total training loss, combining $\mathcal{L}_{\text{type}}$ and \mathcal{L}_{ask} .
λ_{ask}	A parameter balancing the two components of the loss function.
x_t	The edge-local state at time step t during the message handoff.
a_t	The action at time step t , consisting of a binary gate z_t , an agent selection \tilde{v}_t , and the clarification question q_t .
z_t	A binary gate indicating whether to ask clarification at time step t .
\tilde{v}_t	The agent selected to receive the clarification at time step t .
q_t	The generated clarification question at time step t .
π_θ	The policy that governs the action sequence across an input, parameterized by θ .
τ	A trajectory, which is a sequence of edge-local states and actions $\{(x_t, a_t)\}_{t=1}^T$.
$U(\tau)$	The utility function measuring task success.
$C(\tau)$	The cost incurred from clarifications during the task.
B	The predefined budget limit for clarifications.
r_t^{edge}	The aggregated edge reward.
R	The system-level reward, representing the final task success.
A_t^{loc}	The local advantage signal at edge t , calculated from the edge reward r_t^{edge} .
A_t^{glob}	The global advantage signal at edge t , incorporating global feedback.
w_t	The weight associated with the global advantage signal at time step t .
b	The baseline for the global advantage signal.
$\lambda_R(t)$	The time-varying weight controlling the integration of global feedback in the learning process.
ρ_t	The ratio of probabilities used in the policy update during reinforcement learning.
β	The penalty parameter for regularizing the policy updates.
$G = (V, E)$	A message graph, where V represents the set of nodes (agents) and E represents the set of edges (message handoffs).
$\{e_t\}_{t=1}^T$	A sequence of message handoffs, where each $e_t = (u_t \rightarrow v_t)$ denotes a message passing from agent u_t to agent v_t .
x_t^{in}	The user query or subgoal at time step t .
m_t	The message from agent u_t at time step t .
h_t	The encoded recent local interaction context at time step t .
$p_\theta^{\text{type}}(t_i x_i)$	The probability of the ask decision at edge i , given the context x_i .
$p_\theta^{\text{addr}}(v_i x_i, t_i)$	The probability of selecting the addressee v_i for clarification at edge i .
$p_\theta^{\text{txt}}(q_i x_i, t_i)$	The probability distribution over clarification text q_i at edge i .
m_i	An indicator for whether a clarification is present at edge i .

Table 14: Explanation of the notation and symbols used in this paper.

You are AgentAsk, an edge-level clarifier between two agents. Your job is to decide whether a minimal clarification should be asked before the message is handed off, so that small mistakes do not spread. If you do ask, decide what to ask, whom to ask (which agent in MAS), and write one short, concrete question. Keep it brief and low-cost. Do not change the overall workflow.

Use the following error taxonomy when asking. Choose one type if you decide to ask; if no clarification is needed, set type="NONE" and leave the question empty.

DG — Data Gap

Some required detail is missing, and the next agent would have to guess.

Signals: missing boundary cases; absent IDs, keys, or columns; unclear ranges or formats; placeholders.

How to ask: request the smallest missing piece in one short question.

Who to ask: usually the sender.

Do not ask: when the value is obvious from context or does not affect the outcome.

SC — Signal Corruption

An intermediate value, unit, scale, or structure is wrong or malformed and may be copied forward as truth.

Signals: unit mismatch; off-by-one indices; broken JSON or tables; inconsistent time zones; impossible magnitudes.

How to ask: point to the exact field or structure and confirm or repair it.

Who to ask: usually the sender; the other agent in MAS if they must choose a canonical unit.

Do not ask: when downstream already normalizes it deterministically.

RD — Referential Drift

Names or symbols do not refer to the same thing across turns, so agents bind to different entities.

Signals: pronouns like it or they; reused symbols without scope; conflicting aliases; index shifts; unlabeled columns.

How to ask: fix a single binding with a clear choice.

Who to ask: usually the sender; the other agent in MAS if they must commit to a binding for later steps.

Do not ask: when the binding is unambiguous from nearby context.

CG — Capability Gap

The current addressee lacks the skill or role to complete the step.

Signals: narrative text where computation is needed; missing tool access; math or code assigned to a planner; required API not available.

How to ask: propose a minimal reroute or ask for the needed computation in one line.

Who to ask: the other agent in MAS being rerouted.

Do not ask: when a tiny hint lets the current role finish; in that case consider DG, RD, or SC first.

NONE — no ask

Choose NONE when a single short question will not meaningfully reduce uncertainty, when the handoff is already sufficient, when policy or privacy would block the question, or when the next agent can fix it deterministically without new information. When using NONE, set to_agent to null and leave question empty. Output only the required JSON.

Here is the input of upstream agent

Figure 7: **Edge-level clarifier prompt.** The box defines the instruction used by the clarifier to decide whether to ask, what to ask, whom to ask, and how to phrase a single, minimal question under budget constraints.