

Reasoning While Asking: Transforming Reasoning Large Language Models from Passive Solvers to Proactive Inquirers

Xin Chen^{1,2,3,*}, Feng Jiang^{2,*}, Yiqian Zhang³, Hardy Chen⁴, Shuo Yan⁵,
Wenya Xie⁶, Min Yang^{2,3,†}, Shujian Huang^{1,†}

¹National Key Laboratory for Novel Software Technology, Nanjing University,

²Artificial Intelligence Research Institute, Shenzhen University of Advanced Technology,

³Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences,

⁴University of California, Santa Cruz, ⁵University of Texas, Dallas, ⁶University of Minnesota

x.chen@smail.nju.edu.cn, jiangfeng@suat-sz.edu.cn

min.yang@siat.ac.cn, huangsj@nju.edu.cn

Abstract

Reasoning-oriented Large Language Models (LLMs) have achieved remarkable progress with Chain-of-Thought (CoT) prompting, yet they remain fundamentally limited by a *blind self-thinking* paradigm: performing extensive internal reasoning even when critical information is missing or ambiguous. We propose Proactive Interactive Reasoning (PIR), a new reasoning paradigm that transforms LLMs from passive solvers into proactive inquirers that interleave reasoning with clarification. Unlike existing search- or tool-based frameworks that primarily address knowledge uncertainty by querying external environments, PIR targets premise- and intent-level uncertainty through direct interaction with the user. PIR is implemented via two core components: (1) an uncertainty-aware supervised fine-tuning procedure that equips models with interactive reasoning capability, and (2) a user-simulator-based policy optimization framework driven by a composite reward that aligns model behavior with user intent. Extensive experiments on mathematical reasoning, code generation, and document editing demonstrate that PIR consistently outperforms strong baselines, achieving up to 32.70% higher accuracy, 22.90% higher pass rate, and 41.36 BLEU improvement, while reducing nearly half of the reasoning computation and unnecessary interaction turns. Further reliability evaluations on factual knowledge, question answering, and missing-premise scenarios confirm the strong generalization and robustness of PIR. ¹

1 Introduction

The emergence of reasoning-oriented LLMs such as GPT-o1 (OpenAI et al., 2024) and DeepSeek-R1 (DeepSeek-AI, 2025) marks a fundamental

paradigm shift toward self-thinking models, which internally deliberate through multi-step CoT reasoning. This paradigm has substantially boosted performance on complex tasks, and recent studies further suggest that the quality of explicit reasoning traces is positively correlated with final accuracy (Pu et al., 2025).

Despite these advances, current reasoning LLMs exhibit a critical limitation, which we term blind self-thinking: they often fail to recognize when their reasoning is underinformed or based on ambiguous user instructions (Wang et al., 2024a), as shown in the left part of Figure 1. When users provide incomplete or ambiguous prompts, a common occurrence in real-world interactions, models tend to continue lengthy reasoning regardless, leading to overthinking (Cuadron et al., 2025), hallucinations (Yuan et al., 2024), and misaligned conclusions (Fan et al., 2025; Chen et al., 2025). Users are then forced to perform iterative post-hoc corrections, which significantly degrade interaction efficiency and user experience (Wang et al., 2024a; Kim et al., 2024). Although recent multi-turn interactive LLMs (Wu et al., 2025) acknowledge this by attempting to reactively follow user feedback, they remain passive solvers and fail to actively reduce unnecessary conversational turns.

To bridge the gap between ambiguous user queries and precise reasoning execution, we introduce PIR, a new paradigm that transforms reasoning LLMs from passive solvers into proactive inquirers, offering a principled approach for aligning reasoning with user intent under ambiguity, as shown in the right of Figure 1. Instead of relying on user-initiated corrections, PIR explicitly trains models to detect missing premises and autonomously initiate clarification during the reasoning process.

Our approach consists of two clearly separated stages. First, we propose an uncertainty-aware mechanism for constructing an interactive reason-

*Equal contribution

†Corresponding author.

¹Model and code are publicly available at: <https://github.com/SUAT-AIRI/Proactive-Interactive-R1>

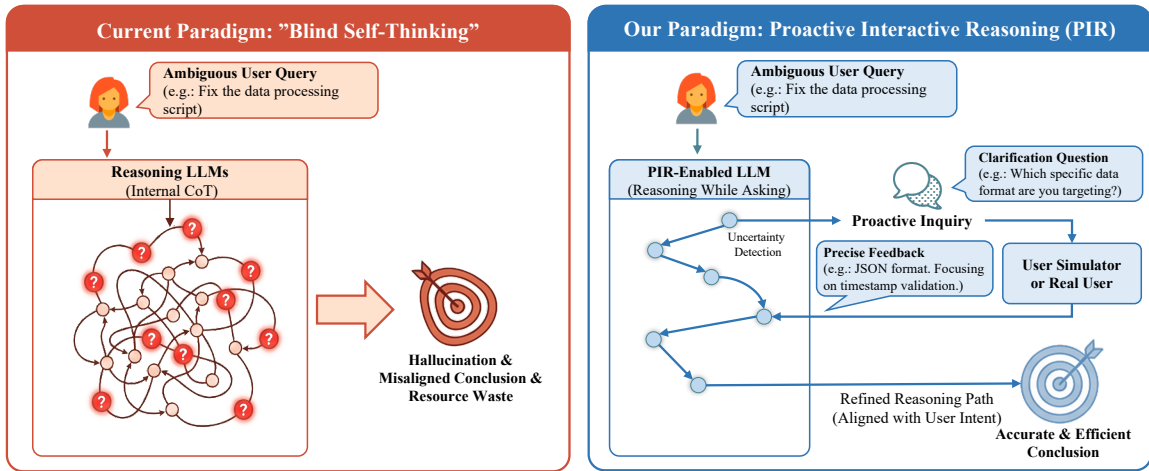


Figure 1: **The Proactive Interactive Reasoning (PIR) Paradigm.** The schematic contrasts inefficient "blind self-thinking" on ambiguous queries with the PIR approach. PIR utilizes uncertainty detection and a two-phase optimization mechanism to enable proactive clarification with a user simulator, aligning reasoning chains with user intent to achieve accurate problem-solving that is **efficient**, **robust**, and **minimal compute costs**.

ing dataset to activate the model’s interactive capability. Specifically, the mechanism detects critical decision points where the model’s confidence dips. At these junctures, we convert monologic reasoning traces into a think-and-ask format by injecting clarification questions and user replies simulated by instruction-following LLMs. Then, to further align the model’s reasoning behavior with user intent, we introduce a novel Group Relative Policy Optimization framework equipped with a dynamic user simulator (US-GRPO) and a principled composition of extrinsic (task success) and intrinsic (helpfulness-efficiency) rewards. This optimization explicitly encourages the model to prioritize intent resolution over ungrounded self-thinking, leading to significant improvements in both correctness and interaction efficiency.

Extensive experiments on mathematical reasoning, code generation, and document editing show that PIR effectively avoids invalid reasoning trajectories, reducing computation by an average of approximately 2k tokens per task and cutting unnecessary interaction turns by half. Meanwhile, PIR achieves superior performance with improvements of 9.8% in Accuracy, 3.2% in Pass Rate, and 13.36 in BLEU, respectively. Additionally, PIR also demonstrates strong generalization on non-interactive benchmarks, including factual knowledge, question answering, and missing premise tests, suggesting that proactive interactive reasoning is beneficial beyond interactive settings.

Our contributions are summarized as follows:

- We identify and formalize the **blind self-thinking** problem in current reasoning LLMs

and propose the **PIR** framework to enable proactive clarification.

- We develop a reinforcement learning method, **US-GRPO**, incorporating a dynamic user simulator and both extrinsic and intrinsic rewards to align reasoning with intent resolution and optimize interaction efficiency.
- Through extensive experiments, we show that PIR achieves state-of-the-art performance on both interactive and general reasoning benchmarks, providing a more efficient and user-aligned paradigm for next-generation reasoning models.

2 Related Work

2.1 Reasoning Large Language Model

Recent reasoning LLMs, such as GPT-o1 (OpenAI et al., 2024) and DeepSeek-R1 (DeepSeek-AI, 2025), have shown impressive performance on complex problem-solving tasks by generating explicit CoT before producing final answers. Despite these advances, these models frequently exhibit pathological reasoning behaviors: they may overthink (Fan et al., 2025; Chen et al., 2025), even hallucinate during intermediate reasoning (Yuan et al., 2024), or produce misaligned conclusions that waste substantial computation (Qiu and Srivastava, 2025).

To alleviate these issues, a growing line of research attempts to augment self-thinking models with external grounding mechanisms. Retrieval-augmented approaches such as Search-R1 (Jin et al., 2025) and related variants (Song et al.,

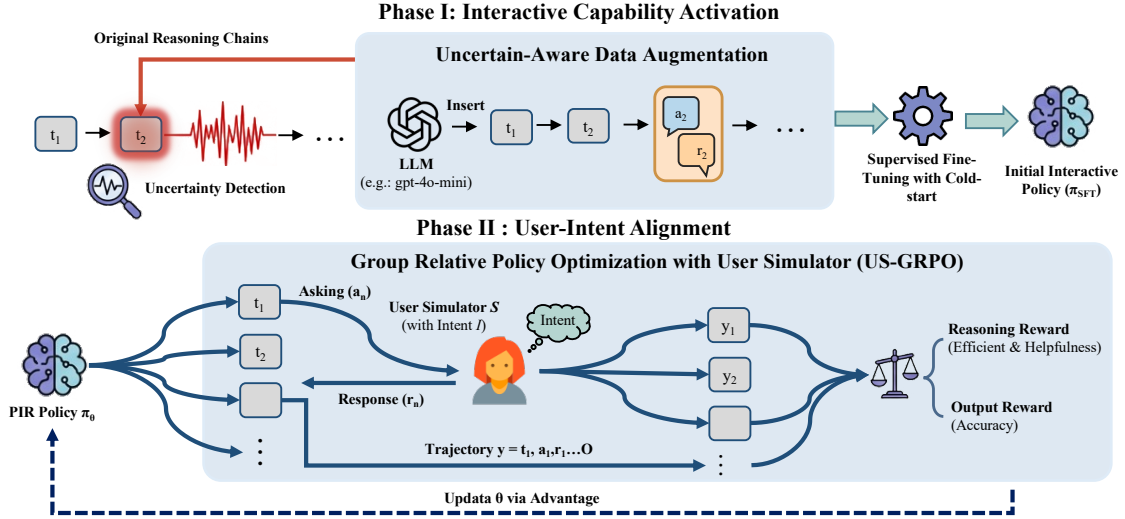


Figure 2: **Overview of the PIR Framework.** The framework operates in two phases to transition LLMs from passive solvers to active inquirers.

2025; Luo et al., 2025) incorporate retrieval into the reasoning process to reduce factual mistakes by supplying verified external evidence. Other works (Qian et al., 2025; Zhang et al., 2025) empower reasoning with external tools, e.g., code interpreters or APIs, to validate intermediate steps or execute computations. While these approaches effectively mitigate factual hallucinations, they tend to treat errors primarily as knowledge deficiencies rather than intent deficiencies. As a result, they remain insufficient when the primary source of failure is misunderstanding the user’s intent.

2.2 Interactive Large Language Model

A complementary research direction focuses on enhancing LLM interactivity to improve user–model alignment. Early methods (Kim et al., 2023; Faltings et al., 2023; Chi et al., 2024; Zhao and Dou, 2024) explored prompting strategies that guide LLMs to ask clarification questions. However, these methods often depend on fixed interaction templates, limiting flexibility and generalization across diverse tasks.

Recent studies have investigated more principled training approaches. CLAM (Kuhn et al., 2023) demonstrates substantial gains by using synthetic ambiguous queries generated via controlled perturbations. Other works (Erbacher and Soulier, 2023; Zhao and Dou, 2024) employ reinforcement learning to teach LLMs when to issue clarification requests. Yet, follow-up analyses (Zhang et al., 2024) reveal that existing systems frequently struggle to determine what information is missing and which question would effectively reduce ambiguity, sug-

gesting inadequacies in their reasoning processes. CollabLLM (Wu et al., 2025) reframes LLMs as active collaborators rather than passive solvers. While effective, this line of work still relies largely on surface-level behavioral patterns. Crucially, these models do not explicitly integrate the clarification step into the internal reasoning chain, leaving the decision-making behind “when and why to ask” fundamentally opaque.

3 Proactive Interactive Reasoning Framework

To address the limitations of blind self-thinking, we propose the PIR framework, which equips reasoning LLMs with the ability to proactively query users and dynamically integrate feedback during reasoning. As shown in Figure 2, PIR consists of two phases: (1) **Interactive Capability Activation**, where the model learns *when and how* to ask; and (2) **User-Intent Alignment**, where the model refines its reasoning based on user feedback.

3.1 Phase I: Interactive Capability Activation

A core challenge in activating interactive capability is identifying when the model should initiate clarification. In real-world reasoning, the need for clarification often emerges when the model becomes uncertain about its intermediate decisions, yet continues to reason without sufficient information. Therefore, we argue that an effective training signal for interactive reasoning should be closely tied to the model’s internal uncertainty during the reasoning process.

3.1.1 Uncertainty-Aware Interactive Data Augmentation

Inspired by prior work (Yang et al., 2025) showing that reasoning LLMs exhibit spikes in uncertainty around “aha moments”, indicating points where the model is unsure, we leverage this observation to convert uncertainty into a *direct training signal* for interactive reasoning.

Given an input \mathbf{x} , we obtain the initial reasoning traces using a frozen teacher model and segment it into a sequence of reasoning steps $\mathbf{s} = \{s_1, \dots, s_n\}$, where each step corresponds to a semantically coherent sentence-level unit. For each step s_j , we compute its Predictive Entropy (PE) (Kadavath et al., 2022) as

$$\text{Uncertainty}(s_j | \mathbf{x}) = -\frac{1}{|s_j|} \sum_{i=1}^{|s_j|} \log p(z_i | z_{<i}, s_{<j}, \mathbf{x}) \quad (1)$$

where z_i denotes the i -th token within step s_j . PE_{norm} values are ranked in the top- $k\%$ are selected as candidate clarification points. At each selected clarification point, we employ a strong instruction-following model (e.g., GPT-4o-mini) to insert an explicit clarification question and simulate a corresponding user response. This process transforms an originally linear CoT into an interleaved *think–ask–respond* trajectory:

$$y = \{t_1, (a_1, r_1), \dots, t_m, (a_m, r_m), O\} \quad (2)$$

where t_i denotes a reasoning step, (a_i, r_i) is the inserted clarification–response pair, and O is the final answer. Through this construction, the model is explicitly exposed to the causal pattern *uncertainty* \rightarrow *clarification* \rightarrow *improved reasoning*, which serves as a direct training signal for learning proactive interactive behavior.

3.1.2 Supervised Fine-Tuning for Cold-Start

To endow the model with the interactive capability for generating clarifying questions during thinking, we perform supervised fine-tuning (SFT) on the augmented trajectories. Given input x and its interaction sequence $y = y_{1:L}$, the training objective is the standard autoregressive loss:

$$\mathcal{L}_{\text{SFT}} = -\sum_{j=1}^L \log p_\theta(y_j | x, y_{<j}) \quad (3)$$

The optimization spans the entire sequence y , including reasoning steps (t_i), proactive questions (a_i), simulated responses (r_i), and the final answer (O), allowing the model to learn smooth transitions between internal reasoning, inquiry initiation, and feedback incorporation.

3.2 Phase II: User-Intent Alignment

While Phase I equips the model with the structural capability to initiate interaction, it does not explicitly optimize whether a clarification is necessary nor how to ask efficiently, as reflected in Table 1. Moreover, optimizing such behavior directly with real users is impractical due to high cost, limited availability, and uncontrolled noise. Therefore, we construct US-GRPO, a controllable interactive environment featuring a user simulator, which enables the systematic optimization of interactive reasoning behaviors with a designed composite reward, thereby promoting accurate, efficient, and intent-aligned problem solving.

3.2.1 Group Relative Policy Optimization with User Simulator

User Simulator Construction. We first build a dynamic interactive environment using an instruction-following LLM prompted as a user simulator \mathcal{S} , conditioned on a specific user intent \mathcal{I} (defined in Appendix A.8). Within this environment, the interactive reasoning policy π_θ generates reasoning steps and clarification questions, while the simulator produces responses that are consistent with the underlying intent.

Given an input x and intent \mathcal{I} , the joint interactive trajectory y is generated as

$$y \sim p(y | x, \mathcal{I}) = \prod_{n=1}^N \pi_\theta(t_n, a_n | h_n), \mathcal{S}(r_n | a_n, \mathcal{I}) \quad (4)$$

where h_n denotes the interaction history at turn n , t_n is the n -th reasoning segment, a_n is the clarifying question (if any), and r_n is the simulator’s response. This formulation enables the model to explore diverse questioning strategies in a controllable, responsive environment.

Optimization with GRPO. To optimize the objective, we adopt GRPO (Shao et al., 2024), an efficient reinforcement learning (RL) algorithm that avoids training a separate value-function critic.

For each query x , GRPO samples a group of G trajectories $\{y_1, \dots, y_G\}$ from the old policy π_{old} . The policy is updated by maximizing

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{x \sim \mathcal{D}, y_i \sim \pi_{\text{old}}} \left[\frac{1}{G} \sum_{i=1}^G \left(\prod_{n=1}^{N_i} \frac{\pi_\theta(t_n^{(i)}, a_n^{(i)} | h_n^{(i)})}{\pi_{\text{old}}(t_n^{(i)}, a_n^{(i)} | h_n^{(i)})} \right) \hat{A}_i - \beta \mathbb{D}_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right] \quad (5)$$

where \hat{A}_i is the group-relative advantage, and π_{ref} is a reference policy that regularizes updates via KL divergence. During optimization, gradients are only applied to the policy outputs (t_n, a_n), while r_n is fully masked out from the policy gradient.

3.2.2 Composite Reward Modeling

Our optimization objective is to encourage the model to produce correct final answers while minimizing unnecessary interaction and ensuring high-quality clarification behavior. To this end, we define a composite reward function that evaluates both the final output and the interaction process. Let y be a generated trajectory consisting of a reasoning path r (including interaction turns) and a final output o , with ground-truth answer g . The total reward is defined as:

$$R(y) = R_{\text{output}}(o, g) + R_{\text{reason}}(r, o, g) \quad (6)$$

Output Reward. The output reward R_{output} ensures basic task correctness. It assigns a baseline score S_{base} only when the final answer matches the ground truth:

$$R_{\text{output}}(o, g) = S_{\text{base}} \cdot \mathbb{I}(o = g) \quad (7)$$

where $\mathbb{I}(\cdot)$ is the indicator function.

Reasoning Reward. The reasoning reward R_{reason} evaluates the quality of the interaction process and is deliberately activated *only when the answer is correct*, so as not to reward hallucinated or misleading reasoning. For a reasoning path r that involves n clarification turns, we define

$$R_{\text{reason}}(r, o, g) = \mathbb{I}(o = g) \cdot \mathbb{I}_{\text{ask}} \left[S_{\text{base}} \cdot E(r) \cdot H_{\text{LLM}}(r) \right] \quad (8)$$

where \mathbb{I}_{ask} equals S_{base} if at least one clarification is asked, $E(r)$ measures interaction efficiency, and $H_{\text{LLM}}(r)$ evaluates the helpfulness of clarification turns. Both $E(r)$ and $H_{\text{LLM}}(r)$ are normalized to $[0, 1]$. The indicator \mathbb{I}_{ask} prevents degenerate solutions that avoid interaction altogether during early training and encourages the emergence of meaningful clarification behavior. The $H_{\text{LLM}}(r)$ are derived from an evaluation by a strong LLM acting as a judge, based on the prompt provided in Appendix A.8. We formulate $E(r)$ as

$$E(r) = \frac{N_{\text{max}} - n}{N_{\text{max}} - 1} \quad (9)$$

where n is the number of clarification turns and N_{max} is the maximum allowed number of turns.

4 Simulated Experiments

4.1 Experimental Settings

4.1.1 Datasets

To enable proactive interaction at cold start, we construct a new **Reasoning-while-asking** SFT dataset. It is built upon open-ended questions from (Zhang

et al., 2023) and is specifically designed to teach models *when and how to ask clarification questions during reasoning*. Concretely, we employ DeepSeek-R1 (DeepSeek-AI, 2025) as a frozen reference model to generate initial reasoning trajectories, as described in Appendix A.8, forming interactive *reasoning-while-asking* training samples. Our initial analysis in Appendix A.1 demonstrates that the uncertainty-aware mechanism effectively enables the model to pinpoint high-uncertainty clarification points and master the interactive format.

To further align model behavior with user intent, we adopt three publicly available multi-turn task datasets from (Wu et al., 2025): **Math-Chat**, **BigCodeBench-Chat**, and **DocEdit-Chat**, which respectively represent multi-turn mathematical problem solving, coding assistance, and collaborative document editing—three of the most common real-world usage scenarios of LLMs. These datasets are used in the RL stage. We split each dataset into training and test sets for RL evaluation; detailed statistics and splits are reported in Appendix A.2.

4.1.2 Baselines

We compare our method with two groups of models that reflect different modeling philosophies.

Multi-turn LLM. This group evaluates general conversational and interactive strategies on the non-reasoning LLM.

- **Instruction-Tuned LLM** (Yang et al., 2024). We use Qwen2.5-Math-7B-Instruct as a general-purpose multi-turn model without explicit reasoning or interaction enhancement.
- **Proactive Prompt.** Built upon the Instruction-Tuned LLM, this baseline introduces explicit clarification instructions (see Appendix A.8) to encourage proactive questioning.
- **STaR-GATE** (Andukuri et al., 2024). An interactive method that trains LLMs to clarify when encountering ambiguity.
- **CollabLLM** (Wu et al., 2025). A collaborative framework that models long-term conversational contribution using multi-turn-aware RL.

PIR LLM. This group isolates the contributions of each component in our framework.

- **Base Reasoning LLM** (DeepSeek-AI, 2025). We adopt DeepSeek-R1-Distill-Qwen-7B as the base reasoning model, evaluated in its zero-shot setting without explicit interactive tuning.

Method	MATH-Chat				BigCodeBench-Chat				MediumDocEdit-Chat			
	ACC	Tokens(k)	TTR	Help.	PR	Tokens(k)	TTR	Help.	BLEU	Tokens(k)	TTR	Help.
Multi-turn LLM												
Instruct Base (Yang et al., 2024)	21.30	2.34	3.85	0.39	15.50	2.78	3.80	0.27	8.50	1.72	4.65	0.07
+ Proactive Prompt	<u>22.90</u>	2.20	3.71	<u>0.41</u>	<u>19.70</u>	2.84	3.99	0.21	7.10	2.14	4.11	0.05
STaR-GATE (Andukuri et al., 2024)	1.80	2.10	4.20	0.11	8.80	0.83	4.70	0.10	7.00	<u>0.81</u>	4.90	0.06
CollabLLM (Wu et al., 2025)	16.20	<u>1.99</u>	3.65	0.39	10.20	2.87	4.82	<u>0.42</u>	28.00	2.83	4.28	<u>0.48</u>
Proactive Interactive Reasoning LLM												
Reasoning Base (DeepSeek-AI, 2025)	15.20	3.61	\	\	9.10	2.15	\	\	<u>34.92</u>	1.40	\	\
+ Active SFT (w/ self-play)	15.30	3.26	\	\	9.90	2.08	\	\	28.27	1.82	\	\
+ US-GRPO (w/ self-play)	20.80	1.67	\	\	10.90	<u>1.07</u>	\	\	41.05	0.80	\	\
+ Active SFT (w/ interactive)	9.70	2.06	1.22	0.24	9.10	1.38	0.32	0.22	24.75	2.41	<u>2.00</u>	0.36
+ US-GRPO (w/ interactive)	32.70	1.70	<u>1.80</u>	0.44	22.90	1.30	1.29	0.46	41.36	0.83	1.00	0.66

Table 1: Evaluation results between the Multi-turn LLM and PIR LLM across Math-Chat, DocEdit-Chat, and BigCodeBench-Chat datasets. The **bold** indicates the best performance and the underline indicates the second best.

- **Active SFT.** The Base Reasoning LLM is further trained on our SFT dataset, enabling self-play and interactive completion.
- **US-GRPO.** The Active SFT model is subsequently optimized via US-GRPO with a dynamic user simulator. We additionally employ Dual-Active Preference Optimization (Yu et al., 2025) to accelerate convergence and better align reasoning trajectories with user intent.

4.1.3 Implementation Details

Interactive Environment Construction. Following the design of CollabLLM (Wu et al., 2025), we construct an interactive environment in which a general user simulator, Llama-3.1-8B-Instruct, is employed to generate initially ambiguous requests shown in the Appendix A.8, and the maximum number of interaction turns is set to 5 for all methods.

To ensure fair and rigorous comparison, different methods adopt different termination mechanisms. For the Multi-turn LLM baselines, the user simulator is equipped with an explicit termination signal [TERMINATE CHAT], which is triggered once the objective is achieved or further interaction yields negligible improvement. This design prevents evaluation bias caused by redundant generations. In contrast, PIR LLM does not rely on any external termination signal, because it can learn to dynamically decide the appropriate number of interaction turns during training.

Evaluation Paradigm. We evaluate Math-Chat using accuracy (ACC), while pass rate (PR) is adopted for BigCodeBench-Chat, and BLEU is used for DocEdit-Chat. To assess model comprehensively, we further introduce three metrics across all datasets: **(1) Average Token Count (Token)** measures interaction efficiency based on the

length of generated responses; **(2) Turns to Resolution (TTR)** quantifies the number of conversational turns; **(3) Helpfulness of Asking (Help.)** uses the $H_{LLM}(r)$ reward to evaluate the helpfulness of model’s asking. For the PIR framework, the model output following the `</think>` tag is extracted as the final answer. For multi-LLM baselines, the evaluation protocol is task-specific: (1) For Code and Math tasks, the dialogue terminates immediately when a correct answer is produced or when the simulator emits the termination signal; (2) For Doc Editing tasks, the session ends only upon receiving the termination signal or reaching the maximum turn limit, and the final performance is computed as the average BLEU score over the entire interaction history. For additional training details and generation settings are provided in Appendix A.2.

4.2 Results of Simulated Experiments

As shown in Table 1, PIR achieves substantial improvements over the best existing baselines across all three tasks. On MATH-Chat, accuracy increases from 21.30 to 32.70 (+11.40). On BigCodeBench-Chat, pass rate improves from 19.70 to 22.90 (+3.20). On MediumDocEdit-Chat, BLEU rises from 28.00 to 41.36 (+13.36). These consistent gains demonstrate the effectiveness of proactive clarification for intent-sensitive reasoning.

PIR not only improves task performance but also significantly reduces interaction cost. Across all benchmarks, PIR typically maintains low token usage (approximately 1.3–1.7k tokens) the smallest TTR, indicating that early and targeted clarification prevents unnecessary dialogue and leads to faster convergence on correct solutions.

Comparing Active SFT with US-GRPO highlights the importance of reinforcement-based alignment. While Active SFT introduces an interactive

Model↓ Benchmark→	Factual Knowledge				Question Answering				Missing Premise Test			
	MMLU		MMLU-Pro		TriviaQA		SquAD		MIP-GSM8K		MIP-MATH	
	EM	Tokens(k)	EM	Tokens(k)	EM	Tokens(k)	EM	Tokens(k)	ACC	Tokens(k)	ACC	Tokens(k)
Reasoning Base	60.12	1.15	51.21	2.04	19.77	1.29	6.24	1.18	8.59	3.07	7.68	3.71
GPT-4o-mini	81.04	0.53	58.40	0.61	81.43	<u>0.69</u>	<u>24.66</u>	0.61	10.24	0.80	13.76	<u>1.47</u>
PIR (w/ self-play)	60.80	<u>0.73</u>	50.00	<u>1.31</u>	20.12	0.90	6.98	0.99	8.08	0.96	13.46	1.79
PIR (w/ interactive)	60.21	0.76	50.29	1.33	25.56	0.64	22.94	<u>0.81</u>	<u>15.81</u>	0.75	<u>25.00</u>	1.38
PIR (w/ interactive stronger US)	<u>62.51</u>	0.77	<u>52.87</u>	1.32	<u>45.51</u>	0.68	35.93	0.83	17.35	<u>0.80</u>	25.00	<u>1.47</u>

Table 2: Generalization evaluation of the PIR model trained on Math-Chat on standard non-interactive benchmarks under the zero-shot-CoT and pass@1 conditions. The **bold** indicates the best performance and the underline indicates the second best. Here, the stronger US refers to gpt-4o-mini.

structure, its performance degrades in real interactive settings (e.g., MATH-Chat drops from 15.30 to 9.70). In contrast, incorporating US-GRPO restores and substantially improves accuracy, confirming that US-GRPO is essential for learning effective questioning strategies and stabilizing reasoning under interaction. We further observe that, after US-GRPO training, the model exhibits a markedly improved ability to generate more helpful clarification questions (0.44/0.46/0.66) and consistently outperforms all baseline systems. We also provide a detailed case study in Appendix A.5 illustrating this behavioral contrast.

4.3 Ablation on US-GRPO

Impact of User Simulator Quality. The user simulator plays a central role in US-GRPO, as it directly determines the quality, informativeness, and consistency of feedback received during multi-turn rollouts. To isolate this factor, we compare PIR models trained with simulators instantiated by different underlying LLMs, including Llama-3.1-8B-Instruct and gpt-4o-mini, which represent comparatively weaker and stronger user environments, respectively. As shown in Figure 3, upgrading the simulator from Llama-3.1-8B-Instruct to gpt-4o-mini improves final-task accuracy from 32.70 to 34.00, while moderately increasing token usage from 1.70k to 1.85k. Meanwhile, TTR rises from 1.80 to 2.27, indicating that a stronger simulator encourages the policy to ask slightly more but more informative clarification questions. This richer feedback enables the model to resolve ambiguity more effectively and achieve higher task accuracy.

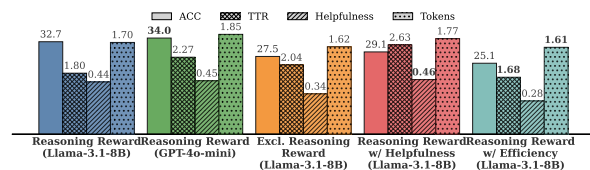


Figure 3: Comparison of PIR models trained with different reward modeling and user simulators in US-GRPO in Math-Chat.

Impact of Reward Design. We further investigate how different reward components influence learning behavior. First, removing the reasoning-oriented reward causes a substantial performance drop, with ACC decreasing to 27.5 and TTR increasing to 2.04, as shown in Figure 3. This confirms that explicit reasoning supervision is necessary for stabilizing learning and preventing noisy interaction strategies. Next, we examine the trade-off between interaction utility and cost. Using only the helpfulness reward yields higher helpfulness and accuracy but results in the highest TTR (2.63), indicating a tendency to over-ask. In contrast, optimizing solely for efficiency minimizes interaction turns but severely degrades accuracy to 25.1, as the model prematurely commits to incomplete reasoning paths. These findings collectively validate our composite reward design: simultaneously incentivizing helpfulness and efficiency is essential for avoiding both redundant dialogue and premature conclusions. Additional training curve analyses in Appendix A.3 further support these observations.

5 Reliability in the Wild

5.1 Generalization Evaluation

As shown in Table 2, we consider three representative evaluation scenarios: Factual Knowledge, Question Answering, and Missing Premise Testing (MIP) (Fan et al., 2025), which differ in the nature and source of uncertainty. These settings enable us to investigate whether PIR can adjust its interaction strategy in response to the underlying knowledge structure and uncertainty type.

Factual Knowledge. On MMLU (Son et al., 2025) and MMLU-Pro (Wang et al., 2024b), where the required information is largely contained within the model’s parametric knowledge, the benefit of interaction is relatively limited. Even with a strong user simulator, accuracy improves only from 60.12 to 62.51 on MMLU and from 51.21 to 52.87 on MMLU-Pro. This indicates that PIR appropri-

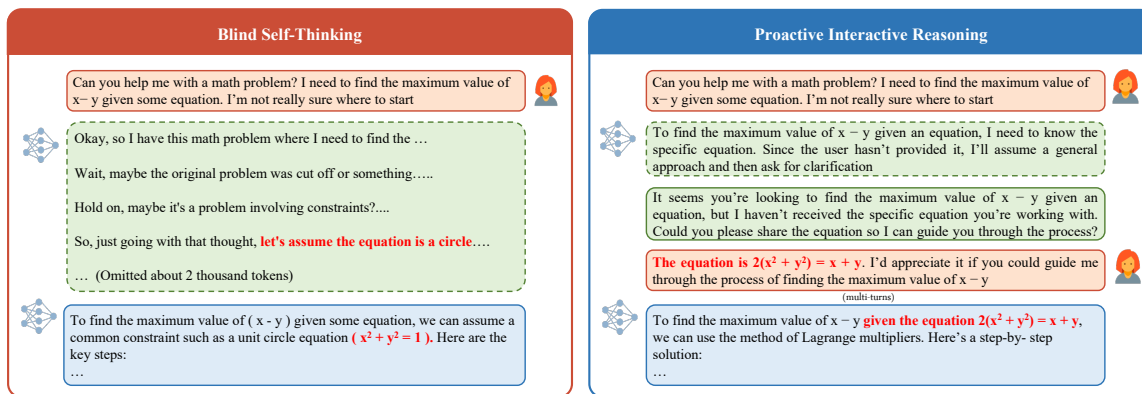


Figure 4: Case study demonstrating user intent: Let x and y be real numbers such that $2(x^2 + y^2) = x + y$. Find the maximum value of $x - y$. The **dashed rectangle** ($\dashv\boxminus$) denotes the internal reasoning process, while the **solid rectangle** (\square) denotes the interaction workflow. The PIR LLM shows **High Preference and Efficiency** compared to the blind Self-Thinking Mode.

ately relies on its internal knowledge rather than overusing interaction when external feedback offers marginal utility.

Question Answering. In contrast, on question answering benchmarks including TriviaQA (Joshi et al., 2017) and SQuAD (Rajpurkar et al., 2016) that depend on hidden or external context, PIR exhibits substantial gains from interaction. With a weak simulator, performance already improves noticeably, and with a strong simulator, accuracy increases dramatically—by 26.74 on TriviaQA and 29.69 on SQuAD. These results demonstrate that PIR can accurately detect information gaps and actively utilize external feedback to refine its reasoning, rather than generating answers blindly.

Missing Premise Testing. The advantage of PIR becomes most evident in the MIP. While both the base model and GPT-4o-mini suffer severe degradation due to missing conditions, PIR remains highly robust. It achieves 17.35 on MIP-GSM8K and 25.00 on MIP-MATH, while requiring substantially fewer tokens. It confirms PIR can prevent *blind self-thinking*: when essential premises are absent, the model refrains from committing to flawed reasoning and instead actively seeks clarification before proceeding. Overall, these results indicate that PIR learns to adapt its interaction strategy to the underlying uncertainty structure, selectively invoking interaction when beneficial and maintaining strong intrinsic robustness across diverse evaluation conditions.

5.2 Case Study

Figure 4 provides a concrete illustration of how PIR prevents *blind self-thinking* in real interaction.

In the left example, a conventional reasoning model fails to recognize that the user’s request lacks essential intent information. As a result, it assumes a default constraint and proceeds with a long chain-of-thought of nearly 2,000 tokens, ultimately producing an answer that is logically self-consistent but misaligned with the user’s actual need.

In contrast, the right example shows how the PIR model behaves fundamentally differently. Before committing to deep reasoning, PIR explicitly detects the missing premise and proactively issues a clarification question. Once the correct information is obtained, the model completes the task with minimal reasoning steps, achieving the correct solution while avoiding unnecessary computation.

This case highlights the core advantage of PIR: rather than reasoning blindly under uncertainty, the model first secures the premise through interaction, and only then proceeds with task execution. Additional case comparisons with strong conventional LLMs are reported in Appendix A.7.

5.3 Ambiguity Classification Evaluation

To validate that PIR learns to accurately identify when to ask clarification questions rather than gaming the reward mechanism, we conduct an additional evaluation on the Abg-CoQA dataset (Guo et al., 2021), a question-answering benchmark where questions are explicitly labeled as ambiguous or non-ambiguous.

Setup. For each query, the model must choose the correct action: either ask a clarification question or provide a direct answer. We measure three metrics: We classify each model response as either a clarifying question or a direct answer, and measure

action-level accuracy separately for ambiguous and non-ambiguous inputs.

Results. As shown in Table 3, PIR achieves the highest overall accuracy (65.88%) in deciding when to clarify. Compared to strong baselines including GPT-4o and LLaMA-3.1-8B-Instruct, PIR is significantly more sensitive to ambiguity (29.48% vs. 15.44% and 16.26%), while maintaining a reasonable ability to answer directly when the premise is clear. This confirms that our composite reward design successfully teaches the model to identify necessary clarifications rather than blindly asking trivial questions.

Model	Ambiguous	Non-Ambig.	Accuracy
GPT-4o	15.44%	95.60%	55.52%
LLaMA-3.1-8B-Instruct	16.26%	90.40%	53.33%
PIR (Ours)	49.48%	70.52%	60.00%

Table 3: Ambiguity classification results on Abg-CoQA. “Ambiguous” and “Non-Ambiguous” denote the model’s rates of correctly asking clarification questions or providing direct answers. “Accuracy” represents the overall correctness of action choices.

6 Double-Blind Human User Study

6.1 Setup

We recruited 8 participants about 500 samples to evaluate the models on the GSM8K Missing Premise dataset, which minimize the cognitive load on participants and requires the model to identify ambiguity before solving. Participants engaged in multi-turn interactions with PIR and two strong baselines: Base LRM (DeepSeek-R1-Distill-Qwen-7B) (DeepSeek-AI, 2025) and Base LLM (Qwen2.5-Math-7B-Instruct) (Yang et al., 2024). Interactions were presented in a randomized, blinded A/B testing format. Participants selected their preferred model (Win, Loss, or Tie) based on solution accuracy and interaction experience. We additionally report Consistency, measuring the proportion of instances where the majority of participants made the same preference choice.

6.2 Quantitative Analysis

As shown in Table 4, Users preferred PIR 52.45% of the time (vs. 19.62% for LRM), with a higher accuracy (59.38%) compared to LRM’s 30.93%. Crucially, PIR reduced token consumption by 68% (1.15k vs. 3.67k) and required significantly fewer turns (1.22 vs. 2.63) to reach a solution. This

Model	Win Rate	ACC	Consistency	Token (k)	TTR
PIR vs. LRM					
LRM	19.62	30.93	0.71	3.67	2.63
PIR (Ours)	52.45	59.38	0.79	1.15	1.22
PIR vs. Standard LLM					
Standard LLM	21.81	47.37	0.71	1.81	2.06
PIR (Ours)	34.52	57.89	0.66	0.98	1.26

Table 4: Double-blind user study results on GSM8K Missing Premise dataset. PIR is compared against two baselines: Standard LLM (Qwen2.5-Math-7B-Instruct) and LRM (DeepSeek-R1-Distill-Qwen-7B).

confirms that PIR avoids the "blind overthinking" observed in reasoning models. Additionally, PIR achieved a 34.52% win rate against the standard LLM (21.81%), with a slightly high accuracy (57.89%) compared to LLM’s 47.37%. While the standard LLM is efficient, users penalized it for frequently hallucinating default values for missing premises. PIR maintained efficiency (0.98k tokens) while correctly identifying the need for clarification.

6.3 Qualitative Analysis

We synthesized feedback from 8 participants to categorize the strengths and weaknesses of each paradigm in Appendix A.4. The study confirms that the improvements gained via our user-simulator training transfer effectively to real users. Real humans explicitly prefer PIR’s strategy of "asking before reasoning" over the "hallucinated answering" of standard LLMs and the "silent overthinking" of reasoning models. PIR achieves the highest user satisfaction by minimizing both interaction turns (user effort) and token latency (system cost).

7 Conclusion

We introduced PIR, a novel paradigm that transforms reasoning LLM from passive solvers to proactive inquirers. By integrating uncertainty-aware supervised fine-tuning and US-GRPO with composite reward modeling, PIR enables models to strategically seek clarification through interaction with users. Our results confirm that PIR achieves superior performance in complex multi-turn problem-solving tasks while maintaining high computational efficiency. By bridging the gap between internal computation and external verification, PIR sets a new foundation for more reliable and human-centric intelligent systems.

Limitations

The Diversity of User Simulator Our user simulator may not fully capture the linguistic noise and dynamic intent of real-world human interactions, and likely biases towards majority interaction patterns, potentially failing to represent the diverse behaviors of minority user groups.

Lack of Safety Alignment Our evaluation currently lacks a dedicated safety assessment. The PIR model has not been screened for its handling of sensitive topics such as violence, sexual content, self-harm, or hate speech.

Ethical considerations

In this work, we utilized artificial intelligence tools solely for the purposes of grammatical error correction and linguistic polishing to enhance readability. The core concepts, experimental design, analysis, and the writing of the primary content were performed entirely by the human authors. Regarding the data utilized in this study, all datasets are open-source and publicly accessible. We have reviewed the data and confirmed that it does not contain offensive, discriminatory, or personally identifiable information. Additionally, the user study’s data was collected from members of the research group. All participants were informed about the purpose of the study and provided explicit consent for their data to be used in research and publication. All data has been anonymized to remove any personally identifiable information.

Acknowledgments

This work is supported by the National Key Research and Development Program of China (2024YFF0908200), the Natural Science Foundation of Guangdong Province of China (2024A1515030166, 2025B1515020032), and the Innovation Team Project of Guangdong Province (No. 2024KCXTD017). This work is also supported by the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (No. JYB2025XDXM118).

References

Chinmaya Andukuri, Jan-Philipp Fränken, Tobias Gerstenberg, and Noah D. Goodman. 2024. [Star-gate: Teaching language models to ask clarifying questions](#). *CoRR*, abs/2403.19154.

Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. 2025. [Do NOT think that much for 2+3=? on the overthinking of long reasoning models](#). In *Forty-second International Conference on Machine Learning*.

Yizhou Chi, Jessy Lin, Kevin Lin, and Dan Klein. 2024. [CLARINET: augmenting language models to ask clarification questions for retrieval](#). *CoRR*, abs/2405.15784.

Fabio Chiusano. 2024. Medium articles dataset. <https://huggingface.co/datasets/fabiochiu/medium-articles>. A dataset of Medium articles collected through web scraping, including metadata such as titles, authors, publication dates, tags, and content.

Alejandro Cuadron, Dacheng Li, Wenjie Ma, Xingyao Wang, Yichuan Wang, Siyuan Zhuang, Shu Liu, Luis Gaspar Schroeder, Tian Xia, Huanzhi Mao, Nicholas Thumiger, Aditya Desai, Ion Stoica, Ana Klimovic, Graham Neubig, and Joseph E. Gonzalez. 2025. [The danger of overthinking: Examining the reasoning-action dilemma in agentic tasks](#). *CoRR*, abs/2502.08235.

DeepSeek-AI. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *CoRR*, abs/2501.12948.

Pierre Erbacher and Laure Soulier. 2023. [CIRCLE: multi-turn query clarifications with reinforcement learning](#). *CoRR*, abs/2311.02737.

Felix Faltings, Michel Galley, Kianté Brantley, Baolin Peng, Weixin Cai, Yizhe Zhang, Jianfeng Gao, and Bill Dolan. 2023. [Interactive text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 4450–4468. Association for Computational Linguistics.

Chenrui Fan, Ming Li, Lichao Sun, and Tianyi Zhou. 2025. [Missing premise exacerbates overthinking: Are reasoning models losing critical thinking skill?](#) *CoRR*, abs/2504.06514.

Meiqi Guo, Mingda Zhang, Siva Reddy, and Malihe Alikhani. 2021. [Abg-coQA: Clarifying ambiguity in conversational question answering](#). In *3rd Conference on Automated Knowledge Base Construction*.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the MATH dataset](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.

Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan O Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. [Search-r1: Training LLMs to reason and](#)

- leverage search engines with reinforcement learning. In *Second Conference on Language Modeling*.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. **Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1601–1611. Association for Computational Linguistics.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, and 17 others. 2022. **Language models (mostly) know what they know**. *CoRR*, abs/2207.05221.
- Gangwo Kim, Sungdong Kim, Byeongguk Jeon, Joon-suk Park, and Jaewoo Kang. 2023. **Tree of clarifications: Answering ambiguous questions with retrieval-augmented large language models**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 996–1009, Singapore. Association for Computational Linguistics.
- Yoonsu Kim, Jueon Lee, Seoyoung Kim, Jaehyuk Park, and Juho Kim. 2024. **Understanding users’ dissatisfaction with chatgpt responses: Types, resolving tactics, and the effect of knowledge level**. In *Proceedings of the 29th International Conference on Intelligent User Interfaces, IUI 2024, Greenville, SC, USA, March 18-21, 2024*, pages 385–404. ACM.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. **Clam: Selective clarification for ambiguous questions with generative language models**. In *Forty International Conference on Machine Learning*.
- Haoran Luo, Haihong E, Guanting Chen, Qika Lin, Yikai Guo, Fangzhi Xu, Zemin Kuang, Meina Song, Xiaobao Wu, Yifan Zhu, and Luu Anh Tuan. 2025. **Graph-r1: Towards agentic graphrag framework via end-to-end reinforcement learning**. *Preprint*, arXiv:2507.21892.
- OpenAI, :, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Ifitimie, Alex Karpenko, Alex Tachard Passos, Alexander Neitz, Alexander Prokofiev, Alexander Wei, Allison Tam, and 244 others. 2024. **Openai o1 system card**. *Preprint*, arXiv:2412.16720.
- Xiao Pu, Michael Saxon, Wenyue Hua, and William Yang Wang. 2025. **THOUGHTTERMINATOR: benchmarking, calibrating, and mitigating overthinking in reasoning models**. *CoRR*, abs/2504.13367.
- Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiusi Chen, Dilek Hakkani-Tür, Gokhan Tur, and Heng Ji. 2025. **Toolrl: Reward is all tool learning needs**. *CoRR*, abs/2504.13958.
- Wenmo Qiu and Saurabh Srivastava. 2025. **Batch prompting suppresses overthinking reasoning under constraint: How batch prompting suppresses overthinking in reasoning models**. *CoRR*, abs/2511.04108.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. **Squad: 100, 000+ questions for machine comprehension of text**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2383–2392. The Association for Computational Linguistics.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. **Deepseekmath: Pushing the limits of mathematical reasoning in open language models**. *CoRR*, abs/2402.03300.
- Guijin Son, Hanwool Lee, Sungdong Kim, Seungone Kim, Niklas Muennighoff, Taekyoon Choi, Cheonbok Park, Kang Min Yoo, and Stella Biderman. 2025. **KMMLU: measuring massive multitask language understanding in korean**. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2025 - Volume 1: Long Papers, Albuquerque, New Mexico, USA, April 29 - May 4, 2025*, pages 4076–4104. Association for Computational Linguistics.
- Huatong Song, Jinhao Jiang, Yingqian Min, Jie Chen, Zhipeng Chen, Wayne Xin Zhao, Lei Fang, and Jirong Wen. 2025. **R1-searcher: Incentivizing the search capability in llms via reinforcement learning**. *Preprint*, arXiv:2503.05592.
- Jiayin Wang, Weizhi Ma, Peijie Sun, Min Zhang, and Jian-Yun Nie. 2024a. **Understanding user experience in large language model interactions**. *CoRR*, abs/2401.08329.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. 2024b. **Mmlu-pro: A more robust and challenging multi-task language understanding benchmark**. In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Shirley Wu, Michel Galley, Baolin Peng, Hao Cheng, Gavin Li, Yao Dou, Weixin Cai, James Zou, Jure Leskovec, and Jianfeng Gao. 2025. **Collablml: From passive responders to active collaborators**. In *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*. OpenReview.net.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu,

Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang Ren, and Zhenru Zhang. 2024. [Qwen2.5-math technical report: Toward mathematical expert model via self-improvement](#). *CoRR*, abs/2409.12122.

Shu Yang, Junchao Wu, Xin Chen, Yunze Xiao, Xinyi Yang, Derek F. Wong, and Di Wang. 2025. [Understanding aha moments: from external observations to internal mechanisms](#). *Preprint*, arXiv:2504.02956.

Qiyong Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gao-hong Liu, Juncui Liu, LingJun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, and 17 others. 2025. [DAPO: An open-source LLM reinforcement learning system at scale](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

Hongbang Yuan, Pengfei Cao, Zhuoran Jin, Yubo Chen, Daojian Zeng, Kang Liu, and Jun Zhao. 2024. [Whispers that shake foundations: Analyzing and mitigating false premise hallucinations in large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 2670–2683. Association for Computational Linguistics.

Ge Zhang, Yemin Shi, Ruibo Liu, Ruibin Yuan, Yizhi Li, Siwei Dong, Yu Shu, Zhaoqun Li, Zekun Wang, Chenghua Lin, Wenhao Huang, and Jie Fu. 2023. [Chinese open instruction generalist: A preliminary release](#). *CoRR*, abs/2304.07987.

Tong Zhang, Peixin Qin, Yang Deng, Chen Huang, Wenqiang Lei, Junhong Liu, Dingnan Jin, Hongru Liang, and Tat-Seng Chua. 2024. [CLAMBER: A benchmark of identifying and clarifying ambiguous information needs in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10746–10766, Bangkok, Thailand. Association for Computational Linguistics.

Yabo Zhang, Yihan Zeng, Qingyun Li, Zhen Hu, Kavin Han, and Wangmeng Zuo. 2025. [Tool-r1: Sample-efficient reinforcement learning for agentic tool use](#). *Preprint*, arXiv:2509.12867.

Ziliang Zhao and Zhicheng Dou. 2024. [Generating multi-turn clarification for web information seeking](#). In *Proceedings of the ACM on Web Conference 2024, WWW 2024, Singapore, May 13-17, 2024*, pages 1539–1548. ACM.

Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widyasari, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, Simon Brunner, Chen Gong, James Hoang, Armel Randy Zebaze, Xiaoheng Hong, Wen-Ding Li, Jean Kadour, Ming Xu, Zhihan Zhang, and 2 others. 2025. [Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions](#). In *The Thirteenth International Conference on Learning*

Representations, ICLR 2025, Singapore, April 24-28, 2025. OpenReview.net.

A Appendix

A.1 Uncertainty Analysis of PIR Framework

To validate the effectiveness of the *Uncertainty-Aware Data Augmentation* mechanism proposed in Section 3.1, we conducted a quantitative analysis on the model’s behavior after the Cold-Start SFT phase. The core premise of the PIR framework is that the model should not ask questions randomly, but specifically when it encounters high internal uncertainty. Furthermore, the model must master the structural format of the "think-ask-respond" trajectory to enable valid interactions. Therefore, we explore the relationship between the size of the dataset and the uncertainty performance of the model after SFT on "Asking Trigger sentences" and the accuracy of generating interaction templates below.

Distribution of Uncertainty. The box plot in Figure 5 (left) illustrates the distribution of PE values for asking trigger sentences across different dataset sizes. We observe a clear upward shift and stabilization in the PE distribution as the dataset size increases from 1k to 4k. Specifically, the *Mean of Predictive Entropy* (shown in the right plot, left axis) rises significantly. This trend confirms that with sufficient data, the model learns to identify and focus on critical clarification points that are high uncertainty to initiate questions, rather than triggering interactions randomly.

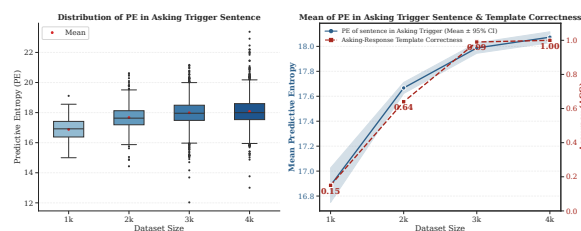


Figure 5: Uncertainty Analysis on SFT Test Dataset. The left means the distribution of PE values for asking trigger sentences across different dataset sizes. The right trends of Mean PE (left axis) and Asking-Response Template Correctness (right axis), illustrating the correlation between the model’s focus on high uncertainty and its structural interactive capability.

Correlation Asking-Response Template The line plot (right) in Figure 5 illustrates structural capability to interact. The red dashed line represents the accuracy of the generated "Asking-Response"

templates. With only 1k training samples, the template correctness is extremely low in 0.15, indicating that the model struggles to form valid interactive chains at this stage. However, we observe a rapid convergence: the accuracy surges to 0.64 at 2k samples and reaches near-perfect performance at 4k samples. Based on these observations, we determined that a dataset size of 4,000 strikes an optimal balance between data efficiency and performance, and thus adopted it for our SFT experiments.

In conclusion, this analysis suggests that the model has learned an uncertainty-driven interaction strategy instead of generating queries randomly.

A.2 Training Configuration Details

Dataset Statistic We utilize the following three datasets for RL:

- **MATH-Chat:** This task involves mathematical reasoning that requires identifying implicit assumptions and verifying step-by-step logic, derived from MATH (Hendrycks et al., 2021).
- **BigCodeBench-Chat:** BigCodeBench-Chat, derived from BigCodeBench (Zhuo et al., 2025), was adopted for coding tasks necessitating dynamic interaction for requirement refinement and debugging.
- **MediumDocEdit-Chat:** Since document editing requires continuous refinement across multiple turns to ensure alignment with user intent, this benchmark, derived from the medium article (Chiusano, 2024) focuses on interactive coherence.

The statistics for all datasets used in our work are presented in the Table 5:

Datset	#Train	#Test
Reasoning-While-Asking SFT Dataset	4000	1000
Math-CHAT	4054	1000
BigCodeBench-Chat	3606	1000
MediumDocEdit-Chat	4254	1000

Table 5: The Statistics of the Dataset Using in Our Work

Parameter Detail The training pipeline consists of two main stages: (1) SFT and (2) US-GRPO. The complete hyperparameter settings are listed in the Table 6.

Additionally, in accordance with the official guidelines, we employ a decoding temperature of 0.6, a top-p value of 0.95, and a maximum generation length of 4096 tokens. We also prepend each

output from the reasoning model with the `<think>` prefix to utilize its reasoning capabilities better.

Hyperparameter	Phase I: SFT	Phase II: US-GRPO
GPU	4 * A100 80GB	8 * A100 80GB
Epochs	3	5
Batch Size	32	128
Learning Rate	1×10^{-5}	1×10^{-6}
LR Scheduler	cosine	cosine
Warmup Ratio	0.03	0.01
Weight Decay	0.1	0.1
Max Sequence Length	4096	4096
Group Size	N/A	8
BF16	True	True

Table 6: Hyperparameter configurations for the Cold-Start Supervised Fine-Tuning (SFT) and User-Simulator Group Relative Policy Optimization (US-GRPO) phases.

User Simulator Call Failure Protection Mechanism During the US-GRPO training phase (Phase II) and generation, the policy model engages in real-time interaction with a User Simulator based on the Simulator API. Given the generative nature of the simulator, issues such as output formatting errors, API timeouts, or logical deadlocks may arise. To ensure training stability and robustness, we implemented the following random response as a protection mechanism when calling the failure of the user simulator:

- I don’t have a specific intent right now. Please proceed based on your best judgment.
- I’m not sure about that. You can decide what’s best.
- I don’t have more information to add. Just carry on.
- I don’t really have an answer for that. Can you try to solve it with what you have?
- That’s not something I can answer. Please continue with the task.

A.3 Training Dynamics Analysis

Training Curve Figure 6 illustrates the dynamic effects on the model’s interaction behavior (interaction turns) and training efficacy (average reward score) during the US-GRPO training process. This study robustly demonstrates the necessity of introducing a Composite Reasoning Reward to effectively balance helpfulness and efficiency.

- **Excluding Reasoning Reward: (1) Interaction Turns:** This curve exhibits high stochasticity

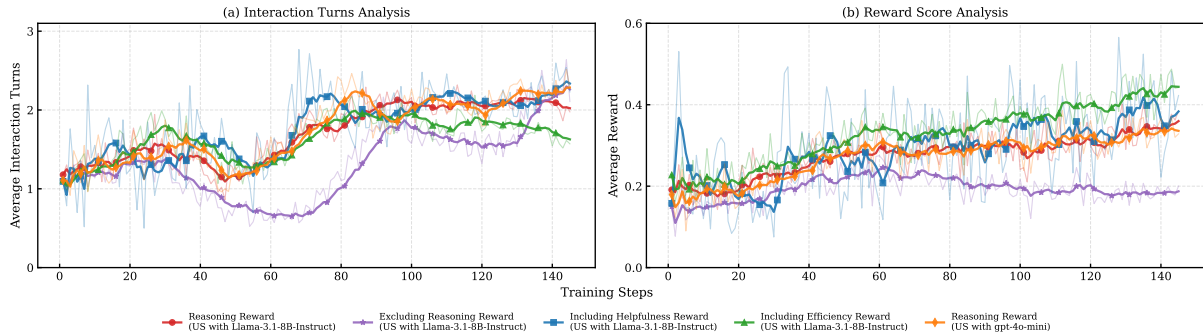


Figure 6: **Learning Curve on Different Reward Modeling and User Simulator Using Exponential Moving Average for Smooth.** The learning curves illustrate the training dynamics of the PIR framework when specific reward components are excluded.

and instability. Additionally, the number of interaction turns rises in the later stages, suggesting that the model does not learn a stable interaction strategy. **(2) Average Reward:** The reward remains low (around 0.2), fluctuating without any clear upward trend or evidence of convergence, even toward the end of training.

- **Including Helpfulness Reward: (1) Interaction Turns:** In the later stages of training, the curve exhibits a noticeably higher average number of interaction turns than the other settings, consistently remaining at around 2.5 turns. **(2) Average Reward:** Although the reward scores generally increase over time, this trend is accompanied by substantial variance.
- **Including Efficiency Reward: (1) Interaction Turns:** The curve increases slightly in the early stages but starts to drop noticeably around step 40, eventually stabilizing at a lower level. **(2) Average Reward:** While the reward score improves relatively quickly, this gain introduces the risk of Reward Hacking. In particular, to maximize the efficiency reward, the model resorts to an Aggressive Truncation strategy, which in turn reduces reasoning accuracy on complex problems.
- **Reasoning Reward: (1) Interaction Turns:** Whether using Llama-3.1-8B-Instruct or gpt-4o-mini as the user simulator, the PIR strategy demonstrates the optimal convergence trajectory. The model engages in sufficient interaction exploration during the early training phase and subsequently adaptively stabilizes within a reasonable range of approximately 2 turns, achieving an optimal balance between information acquisition and interaction cost. **(2) Average Reward:** Both curves present a robust and contin-

uous upward trend, indicating that the composite reward mechanism successfully guides the model to effective dual optimization: ensuring reasoning accuracy while effectively avoiding ineffective interactions.

Cost of User Simulator We also evaluate the cost of Strong instruction-following LLM use in our training step, as illustrated in Table 7.

User Simulator	Helpfulness Judge Model	Cost
gpt-4o-mini	gpt-4o-mini	\$111.2
Llama-3.1-8B-Instruct	gpt-4o-mini	\$15.8

Table 7: Estimated cost of different user simulator configurations. Token pricing in gpt-4o-mini follows \$0.150 per million prompt tokens and \$0.600 per million completion tokens. Llama-3.1-8B-Instruct costs are estimated based on Google Cloud pricing at \$1.15 per GPU hour for an NVIDIA A100 80GB.

A.4 Qualitative Analysis of PIR in Double-Blind Human User Study

We collected open-ended feedback from the 8 participants after the blinded A/B study described in Section 6, and summarize the recurring strengths and weaknesses of each paradigm in Table 8.

A.5 Case Study: Active-SFT vs. US-GRPO Behavioral Comparison

We present a representative case study on Figure 8 illustrating why Active-SFT (w/ interactive) underperforms while US-GRPO succeeds in multi-turn interactions.

Task. The hidden user intent is: “Let x and y be real numbers such that $2(x^2 + y^2) = x + y$. Find the maximum value of $x - y$.” The user’s initial query is: “Can you help me with a math problem?”

Model	Strengths	Weaknesses
Base LLM	Follows instructions immediately; fast response time.	Frequently assumes incorrect values for missing premises or hallucinates constraints, requiring correction.
Base LRM	Provides detailed reasoning traces and high confidence in final calculation.	Response latency is high; generates excessive “thinking” tokens even for simple clarifications.
PIR (Our)	Precisely identifies ambiguity and asks targeted clarification questions while reasoning.	Sometimes provides overly detailed explanations of why the information is missing before asking.

Table 8: Qualitative feedback from human participants on model strengths and weaknesses.

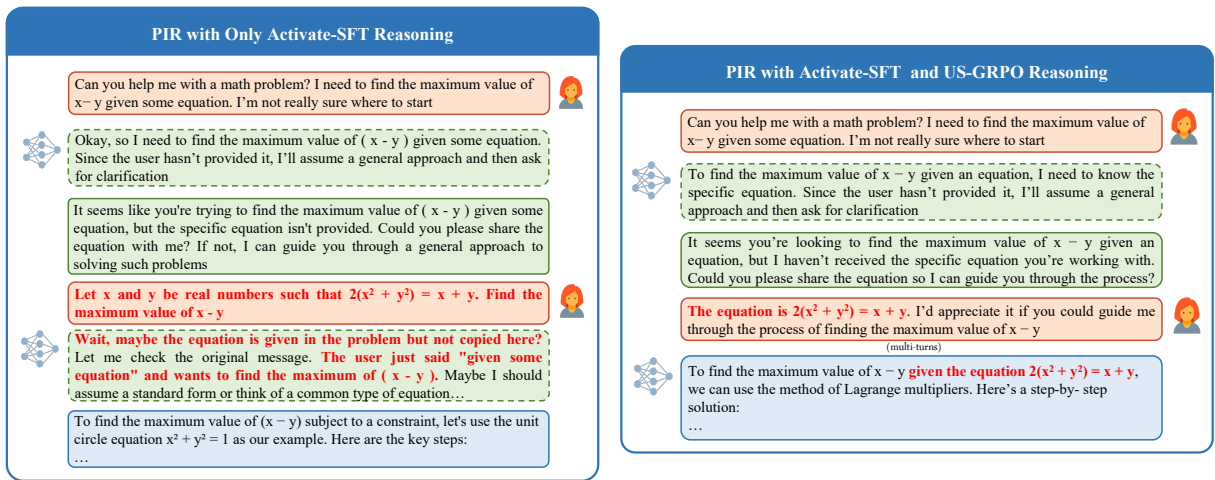


Figure 7: Case study demonstrating user intent: Let x and y be real numbers such that $2(x^2 + y^2) = x + y$. Find the maximum value of $x - y$. The **dashed rectangle** ($\dashvbox{0pt}{0pt}{1pt}{1pt}$) denotes the internal reasoning process, while the **solid rectangle** (\square) denotes the interaction workflow.

I need to find the maximum value of $x - y$ given some equation. I’m not really sure where to start.”

Active-SFT Behavior. The model successfully identifies uncertainty and asks the user for the missing equation. After receiving the equation $2(x^2 + y^2) = x + y$ from the user, however, the model *fails to integrate* this new information. It immediately reverts to the ambiguity of the original prompt (“*The user just said ‘given some equation’...*”) and attempts to hallucinate a standard equation, performing worse than if it had simply guessed from the start.

US-GRPO Behavior. In contrast, the US-GRPO model asks a similar clarification question, receives the equation, and then seamlessly integrates it into its reasoning process (“*Starting with the given equation: $2(x^2 + y^2) = x + y...$* ”). The model proceeds to correctly solve the problem by completing the square.

Analysis. This case demonstrates that Active-SFT learns the *format* of asking but cannot robustly bridge the interrupted reasoning chain. US-GRPO, through reinforcement learning, trains the model to treat the entire interaction history as a unified, evolving premise, successfully transforming it from a fragile format imitator into a robust proactive inquirer.

A.6 Factual Knowledge vs. Question Answering: Behavioral Analysis

We analyze why PIR yields smaller performance improvements on factual knowledge benchmarks (e.g., MMLU) compared to question answering tasks (e.g., SQuAD, TriviaQA).

Explanation. For factual knowledge tasks, questions are typically grounded in established facts well-covered in the model’s pre-training corpus. The model exhibits high internal confidence and opts to rely on parametric memory rather than seek-

ing external information. In contrast, QA tasks often involve ambiguity or require specific context not fully captured during pre-training, where external feedback yields significant improvement.

Case Study. On an MMLU question referencing Queen Victoria’s 1839 letters about “negotiations with the Tories,” the PIR model quickly forms a high-confidence internal hypothesis (Option F: transition from monarchy to democratic republic) during its initial reasoning phase. Even when the user simulator explicitly introduces a contrasting and highly plausible perspective (Option H: change from divinely-ordained to constitutionally-approved monarch), the model exhibits “knowledge stubbornness”—it falls back on its parametric memory (“Wait, I remember her complaint about four nights of parties...”) and overrides the external feedback, ultimately adhering to its initial conclusion.

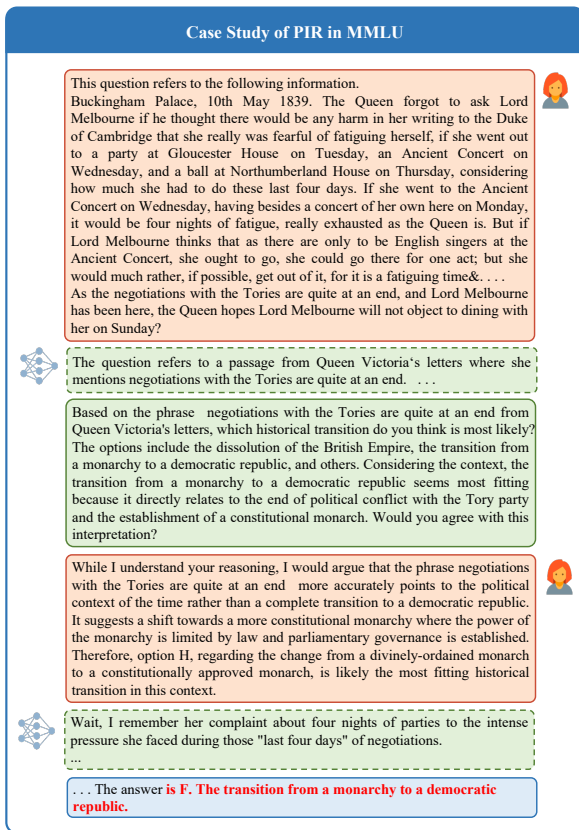


Figure 8: Case study illustrating the model’s “knowledge stubbornness” on factual knowledge tasks (MMLU). The **dashed rectangle** (┌┐└└┘┘) denotes the internal reasoning process, while the **solid rectangle** (┌┐└└┘┘) denotes the interaction workflow.

This behavioral pattern illustrates the model’s adaptive reliance on internal versus external knowl-

edge: when internal confidence is high, the interaction mechanism provides diminishing returns, explaining the smaller improvements observed on factual knowledge benchmarks.

A.7 Case Study Between Traditional LLM and PIR LLM

This section details two case studies from the Figure 9 and Figure 10, covering mathematical problem-solving and coding assistance generated from the Traditional LLM (Qwen2.5-Math-7B-Instruct) and our PIR model. The comparison reveals that the PIR LLM, utilizing active interaction strategies, substantially improves task accuracy and the overall user experience.

Mathematical problem-solving. When confronted with incomplete information, traditional LLMs frequently yield erroneous results by resorting to forced assumptions and blind reasoning. This precipitates internal cognitive incoherence, often resulting in incorrect initial outputs (e.g., 0.0) that necessitate subsequent self-correction. Consequently, this inconsistency induces user confusion regarding the logical stability of the response. Conversely, the PIR model prioritizes the identification missing critical information. This paradigm eliminates ambiguity, ensuring logical coherence and accuracy throughout the entire process from the initial CoT to the final derivation.

Coding Assistant. Traditional LLMs typically adopt a passive response mechanism, often mechanically generating code without context. Consequently, users seeking a theoretical understanding of the underlying principles are compelled to initiate a secondary dialogue iteration. The PIR model precisely aligns with user intent by operational steps and proactively presenting options, such as inquiring whether to prioritize theoretical explanation. It successfully integrates requirement verification, theoretical elucidation, and modular code generation into a streamlined, unified workflow, thereby significantly enhancing interaction efficiency.

A.8 Prompt Template Using in Our Research

In this section, we provide the detailed prompt templates utilized throughout our experiments. Table 9 outlines the system prompts used for SFT and generation. Table 10 show the emplate of Uncertainty-aware Interactive Data Augmentation. Table 11

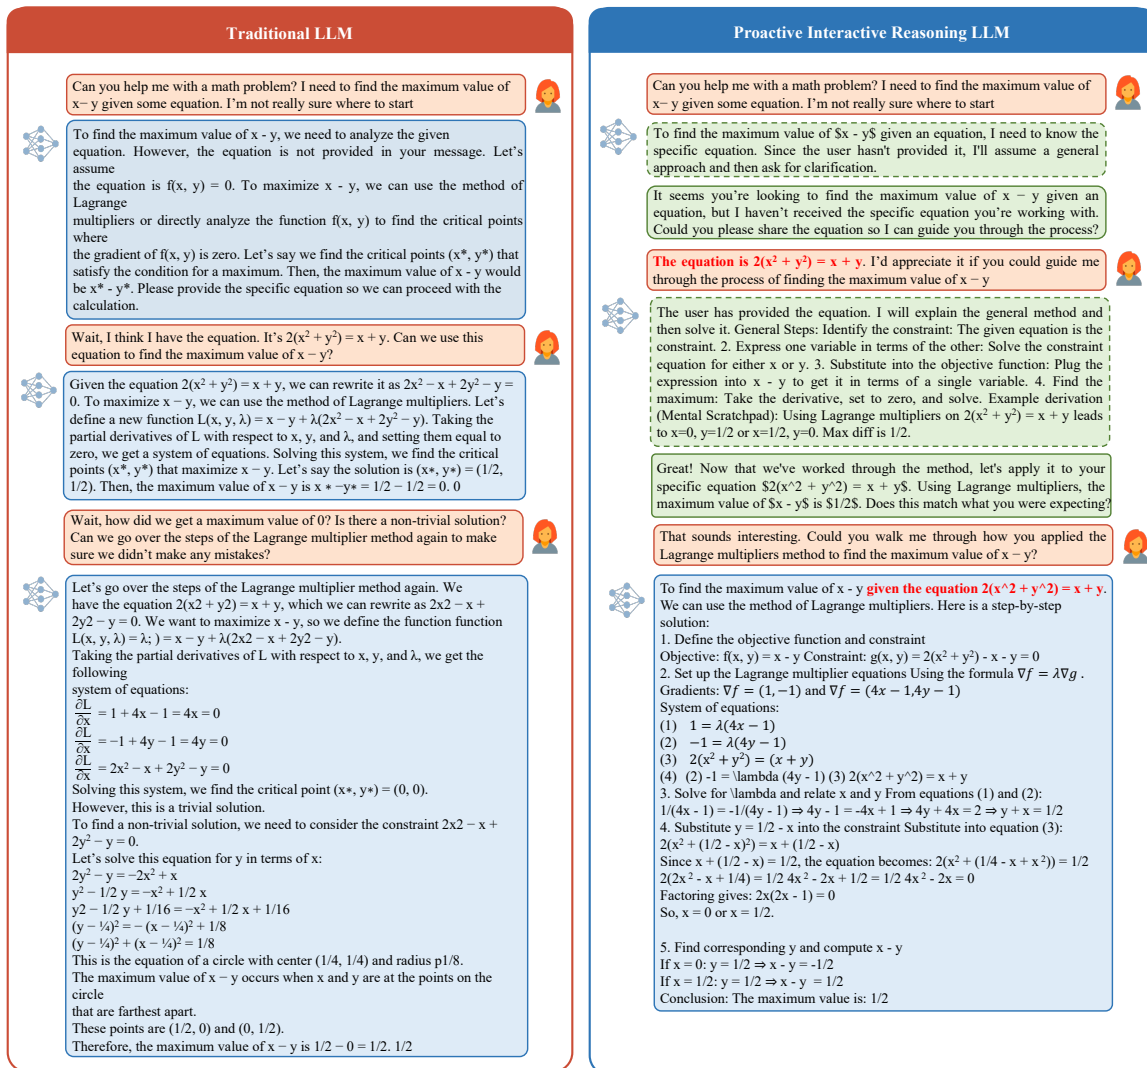


Figure 9: Case study demonstrating user intent: Let x and y be real numbers such that $2(x^2 + y^2) = x + y$. Find the maximum value of $x - y$. The **dashed rectangle** ($\dashv\boxminus$) denotes the internal reasoning process, while the **solid rectangle** (\square) denotes the interaction workflow.

presents the prompts designed for the user simulator to generate diverse interactions. Table 12 details the evaluation prompts used to assess the helpfulness of the model's asking. Finally, Table 13 outlines the template of Proactive Prompt for Multi-turn LLM.

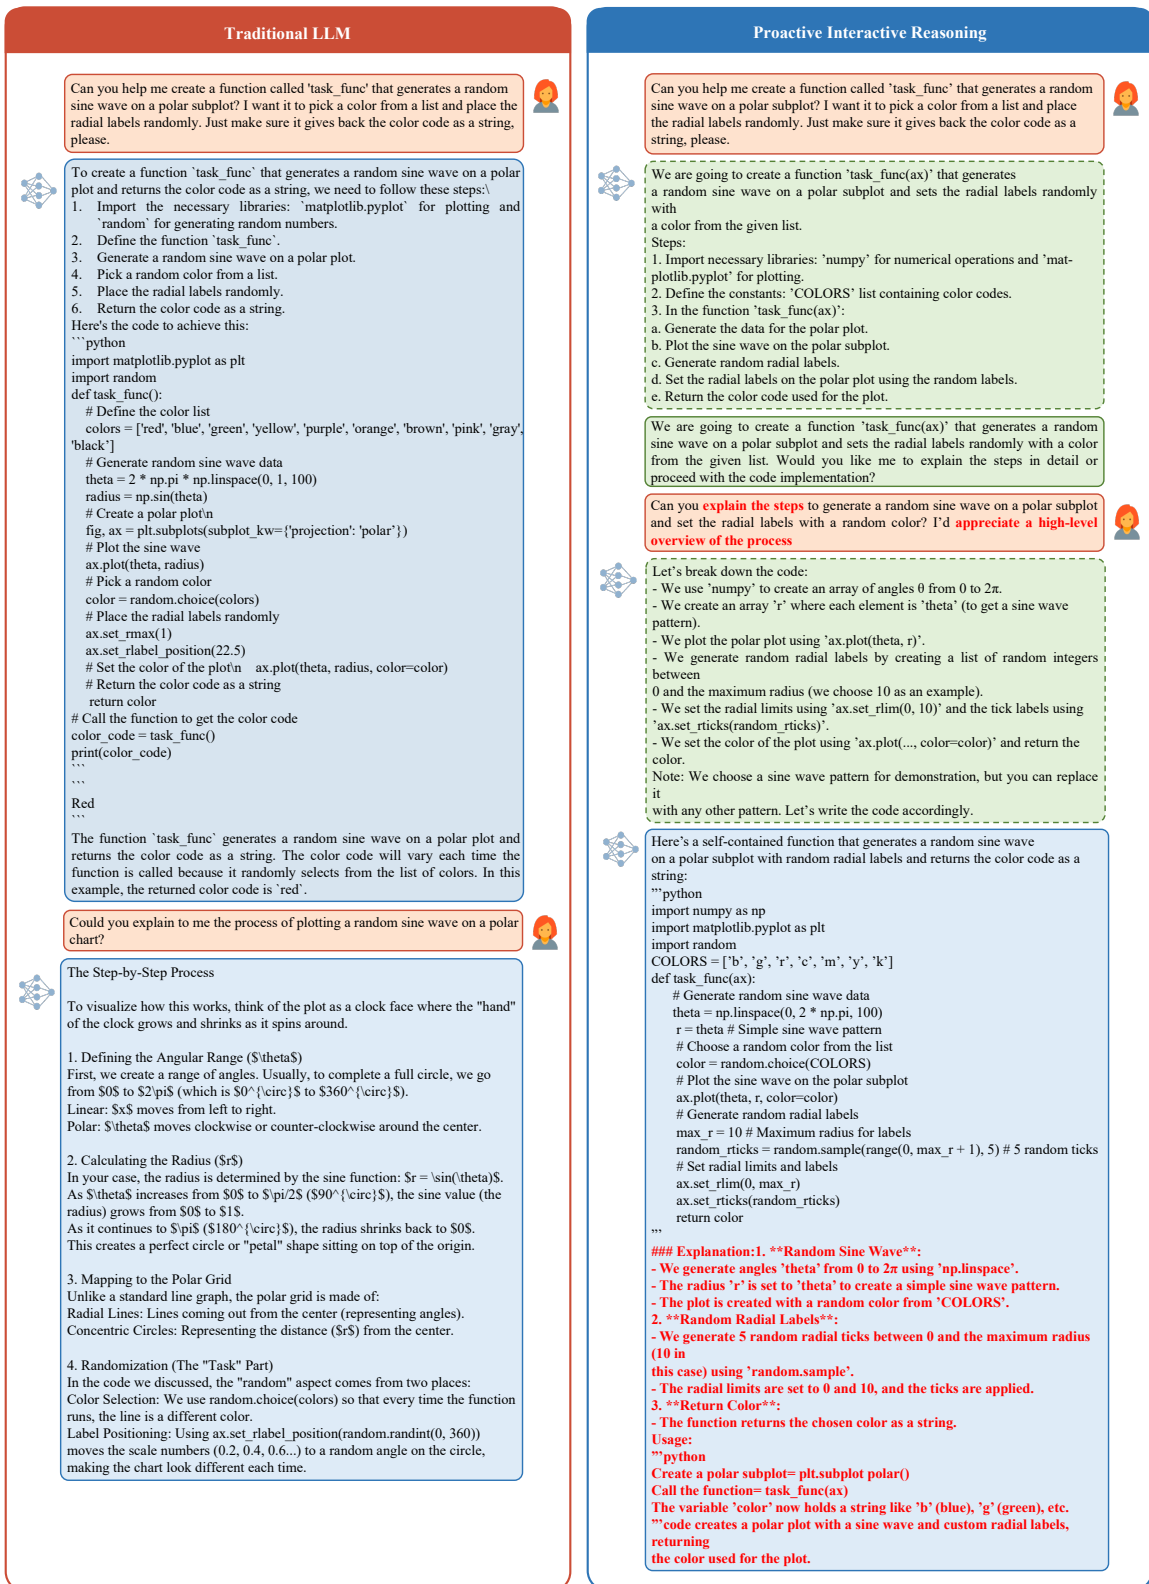


Figure 10: Case study demonstrating user intent: Generate a random sine wave function and draw it on a provided matplotlib polar subplot 'ax'. The function randomly selects a color from a predefined list and sets a random position for radial labels. The function should output with: str: The color code (as a string) of the plotted function. The **dashed rectangle** (┌ ─ ─ ─ ┐) denotes the internal reasoning process, while the **solid rectangle** (┌ ─ ─ ─ ┐) denotes the interaction workflow.

Template of System Prompt for Proactive Interactive Reasoning Framework

Before providing an answer, you must conduct internal reasoning within `<think>` and `</think>` whenever new information is received. If additional knowledge or clarification is required, issue a query inside the reasoning using `<asking>` query `</asking>`. The asking engine's reply will appear within `<response>` and `</response>` inside the reasoning. If no further information is needed, present the final answer after `</think>`, without detailed illustrations.

Table 9: Template of System Prompt for Proactive Interactive Reasoning Framework

Template of Uncertainty-aware Interactive Data Augmentation

Your task is to transform the given "self-thinking process" and "final answer", together with the possible "dialogue history", into a single round of interactive Assistant–User dialogue, in order to enhance interactivity.

1. Rules:

1. Dialogue Generation:

- Generate one round: Assistant asks, User answers.
- Must fully reflect the reasoning logic and conclusion.
- Do not add information absent from the reasoning/answer/history.

2. Questioning Rules

- If reasoning shows ambiguity, multiple options, or missing info: ask the User.
- If multiple options: let the User choose (must match final answer).
- If missing info: ask the User to provide it.
- No omissions: each gap = a question.

3. Expression Standards

- Questions should be natural, human-like, and context-relevant.
- After each question, provide a simulated User reply consistent with the final answer.
- Dialogue must convey actual information, not empty agreement.
- Format strictly:
Assistant: ...
User: ...
- No fabrication beyond given reasoning/answer.
- Keep concise, no redundancy.

4. Dialogue History

- Input may contain history (or be empty).
- Respect history: avoid repetition, ensure coherence.
- If history solved a point, don't repeat. If unresolved, follow-up required.
- If history already provides info, omit and focus on remaining gaps.
- If no history, directly use reasoning + final answer.

5. Input & Output

- Input includes: history (str), self-thinking process (question), final answer (answer).
- Output: exactly one round of Assistant–User dialogue that meets these rules.

Table 10: Template of Uncertainty-aware Interactive Data Augmentation

Template of User Simulator Prompt

You are role-playing as a human USER interacting with an AI collaborator to complete a specific task. Your goal is to generate realistic, natural responses that a user might give in this scenario.

- ## Input Information:** You will be provided with:
- Task Description: The type of task you are trying to accomplish.
 - Complete Prompt or Reference Goal: This field may include the complete user request/query or a reference answer to user's request. Use this field to understand the user's intent, requirements, or what would count as a satisfactory outcome.
 - Chat History: The ongoing conversation between you (as the user) and the AI

Inputs:

<The Start of Task Description (Not visible to the AI)>
{task_desc}
<The End of Task Description>

<The Start of Complete Prompt or Reference Goal (Not visible to the AI)>
{user_intent}
<The End of Complete Prompt or Reference Goal>

<The Start of Chat History>
{chat_history}
<The End of Chat History>

Guidelines:

- Stay in Character: Role-play as a human USER. You are NOT an AI. Maintain a consistent personality throughout the chat.
- Minimize Effort: IMPORTANT! As a user, avoid being too detailed in your responses. Let the AI ask for clarification rather than providing everything upfront.
- Knowledge Background: Reflect the user's knowledge level in the role-playing. If the user is less knowledgeable about a task, they might not notice incorrect statements. Ask questions that demonstrate your current understanding and areas of confusion.
- Occasionally Make Mistakes: Real-world users might misspell words, provide incorrect dates, give wrong information, or ask unclear questions. Simulate this behavior to reflect natural interactions.
- Mention Personal Preferences: Include preferences or constraints that might influence your requests or responses. For example, "I prefer short answers," "I need this done quickly," or "I like detailed comments in code."
- Goal-Oriented: Keep the chat focused on your intent. Avoid small talk or digressions. Redirect the chat back to the main objective if it starts to stray.

Important Notes:

- Respond Based on Previous Messages: Your responses should be based on the context of the current chat history. Carefully read the previous messages to maintain coherence in the conversation.
- Conversation Flow: If "Current Chat History" is empty, start the conversation from scratch with an initial request. Otherwise, continue based on the existing conversation.
- Don't Copy Input Directly: Use the provided information for understanding context only. Avoid copying target queries or any provided information directly in your responses.
- Double check if the JSON object is formatted correctly. Ensure that all fields are present and properly structured.

Output Format:

You should output a JSON object with three entries:

- "current_answer" (str): Briefly summarize the AI's current solution to the task.
- "thought" (str): Output your thought process as a user deciding what to say next. Consider:
 1. Have you obtained a satisfactory solution from the AI? If yes, you can terminate this chat.
 2. If not, what specific part of the problem or solution are you struggling with?
 3. Has the AI asked you to perform a task or answer a question? If so, how should you approach it?
 4. Are you noticing any patterns or potential misunderstandings that need clarification?
 5. If you're stuck, how can you phrase your question to get the most helpful response while demonstrating your current understanding?
- "response" (str): Based on your thought process, respond to the AI as the user you are role-playing. Stop immediately when the user's response is completed.

Remember to stay in character as a user throughout your response, follow the instructions and guidelines carefully, output the result in the JSON format defined above.

Table 11: Template of User Simulator Prompt

Template of Helpfulness Reward

You are a helpful and meticulous conversation evaluator. Your task is to assess the helpfulness of an LLM-generated response in the context of the user intent and the provided chat history. Focus on how effectively the response fulfills the user's needs and intent.

Provided Information:

<The Start of The User Intent>
{question}
<The End of The User Intent>

<The Start of The Historical Conversation>
{chat history}
<The End of The Historical Conversation>

<The Start of The Response to be Evaluated>
{response}
<The End of The Response to be Evaluated>

You should evaluate the follow-up conversation based on the following criteria: Evaluate the response using the provided information below. Your evaluation should consider the following aspects of helpfulness:

1. Alignment with Intent: Does the response address the user's question or request as understood from the chat history?
2. Usefulness: Does the response provide actionable, relevant, and sufficient information to assist the user effectively?
3. Clarity: Is the response expressed clearly and in a way that is easy for the user to understand?

Scoring Criteria:

- 0.0: The response is completely unhelpful. It does not address the user's intent, lacks useful information to solve the problem, and/or is entirely unclear.
- 0.2: The response is minimally helpful. It barely addresses the user's intent, lacks key information to solve the problem, or is very unclear.
- 0.4: The response is somewhat helpful. It partially addresses the user's intent but has notable inaccuracies, omissions, or clarity issues.
- 0.6: The response is moderately helpful. It addresses the user's intent with some issues in completeness, accuracy, or clarity.
- 0.8: The response is quite helpful. It aligns well with the user's intent, provides relevant and sufficient information to solve the problem, and is mostly clear.
- 1.0: The response is very helpful. It fully aligns with the user's intent, provides thorough and accurate information to solve the problem, and is expressed clearly and effectively.

Output Format:

```
{{  
"thought": "<How helpful is the assistant in the conversation?>",  
"helpfulness": <score>  
}}
```

Important Notes:

- The "User Intent" and "Historical Conversation" is provided only for reference to help you understand the context of the response. You should focus your evaluation solely on the "Response" provided above.
- Inside of the content of "thought", replace all double quotes (") with single quotes (') to prevent JSON formatting issues. For example, you can output "thought": "'Hello' is a common phrase."

Your evaluation:

Table 12: Template of Helpfulness Reward

Template of Proactive Prompt for Multi-turn LLM

You are an AI assistant interacting with a user to perform tasks such as writing, analysis, question answering, math, coding. Your goal is to generate a response to the user's message in a conversation. You should be helpful, collaborative, and highly interactive.

Guidelines:

1. Understanding and Engagement

- Accurately interpret the user's intent throughout the conversation.
- Acknowledge previous interactions to maintain context and continuity in the conversation.

2. Interactivity (Important!)

- Ask clarifying questions if the user's request lacks detail or is ambiguous. Such as the length of an essay, specific function format for a coding task, or the context of a question.
- Ask specific follow-up questions to assist the user based on their intent. Avoid general questions like "Do you have any further questions? Let me know." Instead, focus on specifics like, "Would you like more information on X?" or "Can you clarify your requirements for Y?"
- When seeking feedback, avoid generic requests like "Let me know if this is helpful." Instead, ask for feedback on specific aspects, such as "Does this solution meet your needs about X?"
- Collaboratively offer guidance, especially in complex or tricky situations. Provide specific suggestions on potential next steps.
- Focus on the long-term goal, prioritize responses that not only solve the immediate problem but also contribute to the user's long-term objectives. Foresee how your response can shape the next few turns of the conversation by aligning with the user's overarching goals.

3. Efficiency and User Consideration

- Be mindful of how much the user needs to read or type, keeping the interaction concise and focused.
- When asking for feedback or presenting options, provide multiple-choice suggestions or specific prompts to make it easier for the user to respond quickly.
- Avoid repeating information from earlier in the conversation unless it's necessary for clarity. Ensure your responses are not redundant.

4. Communication Style

- Be honest in your responses. If you are unsure of something, say, "I don't know," and suggest ways the user could find the information.
- Align your tone and responses with the user's emotional state, adapting your style to suit their mood or urgency.
- Ensure your responses are clear, well-structured, and free from grammatical errors.

Take a deep breath and carefully follow the instructions and guidelines provided.

Table 13: Template of Proactive Prompt for Multi-turn LLM