

Advancing Reasoning with Off-the-Shelf LLMs: A Semantic Structure Perspective

Pengfei He^{1*}, Zitao Li², Yue Xing¹, Yaling Li², Jiliang Tang¹, Bolin Ding²

¹Michigan State University ²Alibaba Group

Abstract

Large Language Models (LLMs) have shown strong capabilities in zero-shot reasoning and generalization to new tasks. However, the zero-shot performance of general LLMs on complex tasks, such as multi-hop reasoning, remains suboptimal, while reasoning LLMs suffer from hallucinations and unfaithfulness. In this paper, to handle these limitations, we introduce a novel structure-oriented analysis method that helps LLMs better understand the question structure and guide the problem-solving process. We demonstrate that existing reasoning strategies, such as Chain-of-Thought and ReAct, significantly benefit from the LLM’s inherent understanding of semantic structure. We further ground our method in the theory of probabilistic graphical models to support its effectiveness. To enhance the reasoning process, we augment the structure-oriented analysis with refinement and retrieval capabilities, forming a multi-agent reasoning system called **Structure-oriented Autonomous Reasoning Agents (SARA)**. Extensive experiments show that SARA significantly improves zero-shot performance on knowledge-intensive and mathematical tasks. Remarkably, our approach makes a general LLM competitive with dedicated reasoning models in several benchmarks and demonstrates strong robustness against corrupted reasoning paths.

1 Introduction

Large Language Models (LLMs) have shown remarkable potential in various reasoning tasks (Wei et al., 2022; Yao et al., 2022; Shinn et al., 2024; Ahn et al., 2024; Wang et al., 2022), making LLM-based reasoning a fascinating area of research in artificial intelligence. Besides the literature which exhibits LLMs’ strong reasoning abilities when provided with task-specific exemplars (Wei et al.,

2022; Yao et al., 2022; Besta et al., 2024), more recent studies in zero-shot reasoning methods (Kojima et al., 2022; Qiao et al., 2022) demonstrate their unique advantages, and reasoning LLMs (Guo et al., 2025) are specially trained to perform complex reasoning tasks and exhibit impressive zero-shot reasoning performance.

Despite the promising potential of zero-shot reasoning, significant challenges persist in general LLMs and reasoning LLMs. For general LLMs, a primary concern is its inferior performance on complex tasks (e.g., answering multi-hop questions) compared to human or few-shot methods (Huang and Chang, 2022; Ahn et al., 2024). Among incorrect responses, it is often observed that zero-shot methods cannot demonstrate human-like thinking processes, such as comprehensively understanding the problem statements. On the other hand, while reasoning LLMs achieve stronger performance on complex tasks, they also exhibit more frequent hallucinations in intermediate steps (OpenAI, 2025; Bao et al., 2025), and are prone to unfaithful reasoning (Chen et al., 2025), overthinking (Chen et al., 2024b), and usually with higher computation cost.

To explore an effective and efficient way to enhance the LLM’s reasoning capability, we find that human cognition literature offers valuable insights. Research (Simon and Newell, 1971; Kotovsky et al., 1985; Chi et al., 1981; Lakoff and Johnson, 2008) shows that skilled problem-solvers excel at reasoning through new problems without examples or external guidance. Those solvers analyze a problem’s structure, leveraging linguistic and logical patterns to gain a comprehensive understanding (Lakoff and Johnson, 2008). This process identifies critical components (Kotovsky et al., 1985), their relationships, and related sub-questions, while outlining key steps along the reasoning path. These key steps help consolidate the reasoning trajectories, thereby improving reasoning performance effectively and efficiently.

*Corresponding to hepengf1@msu.edu. Work done during internship at Alibaba Group, Bellevue.

Inspired by the human analytic thinking process, we introduce a structure-oriented analysis to activate an **understand-then-reason** process and then improve LLM’s zero-shot reasoning capability, i.e., LLMs are guided to understand the structure of problem statements and generate a comprehensive understanding before performing the reasoning process. The proposed method is based on the syntax and semantic structures in the statement, leveraging the inherent ability of LLMs to parse linguistic patterns (Mekala et al., 2022; Ma et al., 2023). With the help of grammar structures, LLMs can accurately identify critical components in the problem statement and relationships among them and further discover related sub-questions. From this perspective, this analytic thinking process mimics human thinking behavior and thus helps explore correct reasoning paths toward solutions. We empirically demonstrate that simply adding this analysis on top of existing methods such as Chain-of-Thought (CoT) (Wei et al., 2022; Kojima et al., 2022) and ReAct (Yao et al., 2022) can significantly enhance the reasoning performance (Section 3.1). In addition, our theoretical analysis (Section 3.2) also suggests that extracting correct information from problem statements can effectively reduce reasoning errors, further indicating the potential of our structure-oriented analysis in improving LLMs’ inherent reasoning capabilities.

Despite the effectiveness of structure-oriented analysis in direct prompting, we notice two typical errors: later reasoning steps deviating from the structure-oriented analysis, and factual errors resulting in incorrect answers even on the correct reasoning path. To handle challenges, we utilize two mechanisms, reflection (Shinn et al., 2024; Madaan et al., 2024) and retrieval (Yao et al., 2022; Gao et al., 2023), to build a multi-agent reasoning system, **Structure-oriented Autonomous Reasoning Agents (SARA)**. The additional mechanisms let the reasoning process better follow the analysis and utilize external knowledge when necessary (Section 5), achieving a comparable or even better performance than few-shot methods on both **knowledge-intensive reasoning** and **math reasoning** tasks for both general LLMs and reasoning LLMs. Experiments also demonstrate SARA’s cost-effectiveness compared to the baseline methods. Furthermore, we also observe enhanced **robustness** against backdoor attacks (Xiang et al., 2024) and injection attacks (Xu et al., 2024).

To summarize, the main scientific contribution

of this paper is our observation that the zero-shot reasoning ability of LLMs is not fully explored. Supported by both empirical evidence and theoretical validation, the structure-oriented analysis proposed in this paper significantly enhances the zero-shot reasoning capability of LLMs. Furthermore, we enhance the structure analysis with additional mechanisms, forming a multi-agent reasoning system to further improve the performance.

2 Related Work

LLMs for reasoning. In recent literature, there is growing interest in enhancing the reasoning capabilities of large language models (LLMs). Chain-of-Thought (CoT) prompting, introduced by (Wei et al., 2022), encourages models to generate intermediate reasoning steps, significantly improving performance on multi-step tasks. Building on this, (Kojima et al., 2022) proposed zero-shot CoT, prompting models to “think step by step” without task-specific examples, while (Wang et al., 2022) introduced self-consistency to evaluate multiple reasoning paths and select the most consistent one. Other methods can also be found in (Yao et al., 2024; Besta et al., 2024). Besides, sub-problem decomposition is also a widely used approach: for example, Zhou et al. (2022) uses few-shot prompting to decompose questions into sub-questions. Other related works can be found in (Khot et al., 2022; Prasad et al., 2023; Shinn et al., 2024; Madaan et al., 2024; Paul et al., 2023; Shridhar et al., 2023b,a; Zhou et al., 2024; Shridhar et al., 2022; Zhong et al., 2024). However, most of the above methods require task-specific prompting or examples and the zero-shot methods show clear gaps in reasoning performance with few-shot methods.

LLM agents for problem-solving. Besides LLMs’ inherent reasoning capabilities, LLM agents are increasingly employed to enhance performance on complex problems by incorporating external feedback, tools, and knowledge. For example, ReAct (Yao et al., 2022) enables models to interleave reasoning traces with task-specific actions, allowing them to gather additional information from external sources. Some other related works can also be found in (Trivedi et al., 2022; Vu et al., 2023; Zhu et al., 2023; Gou et al., 2023b; Zhou et al., 2023; Sumers et al., 2023; Hong et al., 2024).

3 Structure-oriented Analysis

When skillful human solvers encounter complex questions, a common *routine* is to first identify the critical components and related sub-questions for a comprehensive understanding of the problem (Kotovskiy et al., 1985; Lakoff and Johnson, 2008). This skill can provide a global view of the problem-solving progress, reduce distractions from irrelevant information, and guide for correct reasoning paths (Simon and Newell, 1971). Inspired by these skills, we introduce *structure-oriented analysis*, which leverages LLMs to explicitly extract syntactic and semantic elements from problem statements to guide the reasoning process.

3.1 Empirical findings

Figure 1 shows an example of structure-oriented analysis. We first prompt the LLM to identify the syntactic and semantic structures of the problem statement, and then ask the LLM to extract the following information based on these structures: *key components* that are significant in the problem; *relationships between components* which describe how these key components are related in a structured way; *sub-questions* which are smaller and simpler questions that contribute to the final answer. Leveraging LLM’s inherent ability in syntax and semantic parsing (Drozdo et al., 2022; Mekala et al., 2022; Ma et al., 2023), we develop a general prompt that is applicable across diverse tasks and problems, with minimal need for task-specific examples and human intervention. The detailed prompt is included in Appendix C.

To explore the impact of the structured-oriented analysis, we integrate it with two representative reasoning methods—CoT (Wei et al., 2022) and ReAct (Yao et al., 2022), to empirically examine its performance. We consider both 0-shot and 6-shot versions of CoT and ReAct. More details can be found in Appendix C. To be specific, we first prompt the LLM to perform the structure-oriented analysis and let it finish the remaining reasoning process given the analysis. We evaluate the performance of GPT-4 on a multi-hop question answering benchmark HotPotQA (Yang et al., 2018) and a fact verification benchmark Fever (Thorne et al., 2018). For both tasks, we compare the accuracy with/without our structure-oriented analysis and demonstrate the results in Figure 2. As in Figure 2, adding the structure-oriented analysis can significantly improve the reasoning accuracy, leading to an increase of 5% to 8%. Moreover,

compared to 6-shot methods, 0-shot methods gain more improvements. These indicate that without human intervention, LLMs can still have a deeper understanding of the problem by analyzing syntax structures and linguistic patterns.

3.2 Theoretical analysis

Next, we elaborate on how the reasoning happens from a data perspective and understand the potential benefit of our proposed method. Due to page limit, we provide the skeleton of the analysis and an informal theoretical statement in the main paper and postpone the details to Appendix A.

In short, similar to (Tutunov et al., 2023) and (Xie et al., 2021), we utilize a probabilistic graphical model (PGM) with observed and hidden variables to model the connections among explicit knowledge and abstract concepts in the pre-training data. However, different from (Prystawski et al., 2024; Tutunov et al., 2023) which assume that the LLM always explores along the correct path, we consider a more general scenario where the LLM may explore an incorrect reasoning path. Our result shows that identifying the important reasoning steps is crucial in reasoning.

Build the PGM. We use Figure 3 as an example to illustrate the construction of the PGM. The right panel of Figure 3 provides a detailed instance of how the mathematical notations are connected with real data, and the left panel provides a more general case. In the right panel, we denote $\{\theta_i\}_{i=1}^N$ as the *hidden variables* to represent abstract concepts in the data and $\{X_i\}_{i=1}^N$ as the corresponding *observed variables* for pieces of explicit knowledge $\{x_i\}_{i=1}^N$. For example, θ_1 represents the main campuses of universities and their locations. For each θ_i , the corresponding X_i contains the information of the exact knowledge, such as the location of a specific main campus (x_1).

Intuitively, θ_1 (the main campuses of universities and their locations) and θ_2 (the locations of branches) are logically connected. In addition, during the pre-training, LLM can learn the connection between x_1 (KU’s main campus is in Lawrence, Kansas) and x_2 (Kansas City metropolitan area) and similar pairs of (x_1, x_2) for other universities. By leveraging all observed realizations (x_1, x_2) of (X_1, X_2) , the LLM can infer the relationship between θ_1 and θ_2 . Similarly, the LLM can also learn the connection of (θ_2, θ_4) .

Inference. During the inference, to perform reasoning for the fight song example, the LLM receives x_0

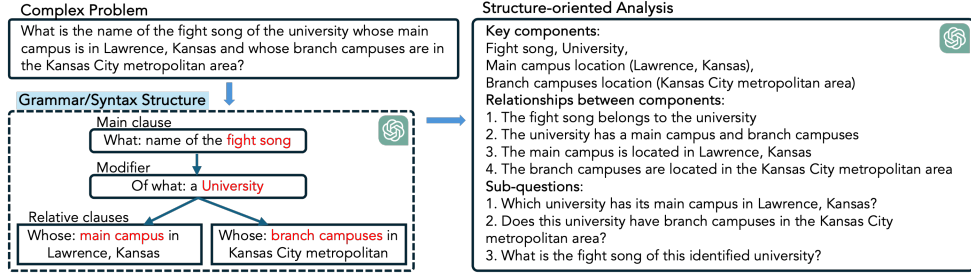


Figure 1: An illustration of the structure-oriented analysis




Figure 2: Reasoning accuracy with/without the structure-oriented analysis. The methods with suffixes + are the backbone methods ($\{\text{CoT, ReAct}\} \times \{0\text{-shot, 6-shot}\}$) with structure-oriented analysis added.

and will explore θ_1 and generate x_1 . Then, given (θ_1, x_1, x_0) , it will further explore θ_2 and generate x_2 , etc. In this example, there is a single reasoning chain, $\theta_1 \rightarrow \theta_2 \rightarrow \theta_4$, allowing the LLM to correctly follow the reasoning path.

On the other hand, if the PGM learned from pre-training is similar to the left panel of Figure 3, then it may explore an incorrect reasoning path: Suppose the correct final state is θ_9 and the LLM starts the reasoning from θ_1 , then the reasoning will be incorrect if exploring θ_3 (the whole path from θ_1 to θ_9 is not in the pre-training data so the LLM may identify the correct path).

For our structure-oriented analysis and similar techniques, if the method can identify one or a few correct hidden states and increase the chance of reaching them, then we have the following benefits:

Theorem 3.1 (Informal Statement of Lemma A.2 and Theorem A.3). *Denote $e(\cdot)$ as the loss given the reasoning path explored by the LLM. Under some mild conditions, if a hidden state θ_a is in the correct reasoning path, then*

- $P(\text{correct reasoning} \mid \theta_a \text{ is explored}) \geq P(\text{correct reasoning})$. *The probability of the LLM doing correct reasoning if it can reach θ_a .*
- $e(\theta_a \text{ is explored}) \leq e(\text{LLM randomly explores})$. *The loss, e.g., accuracy or mean square loss, is also smaller if the LLM can reach θ_a successfully.*

In Appendix A, we provide the rigorous notations and the formal theorem statements.

4 General Agentic Autonomous Reasoning

Although Section 3.1 demonstrates the effectiveness of our structure-oriented analysis as a direct prompting strategy, there is still a large room for improvement. In particular, we identify the following two typical errors and provide concrete wrong answers falling in those types in Appendix J.

Type A Error : While Theorem 3.1 shows the potential benefit of utilizing the key information of structure-oriented analysis, the reasoning process can still deviate to incorrect states.

Type B Error : Even with a correct reasoning path and an appropriate hidden variable (e.g., θ_4), sampling can still introduce incorrect answers (e.g., hallucinating the name of the song).

To handle the above challenges, we augment structure-oriented analysis with two mechanisms: 1) reflection (for Type A error) to encourage aligning with structure-oriented analysis and maintain trajectory consistency. 2) retrieval (for Type B error) with external knowledge to mitigate hallucination of LLMs. We integrate these mechanisms with structure-oriented analysis into a multi-agent system, named as **Structure-oriented Autonomous Reasoning Agents (SARA)** to build a general and flexible solution for different reasoning tasks.

4.1 Cooperative agents

The three major agents (Reason Agent, Refinement Agent and Retrieval Agent) and their shared memory in SARA are detailed as follows.

Reason Agent. This agent serves as the cognitive core of the system, conducting analytic thinking and generating detailed reasoning steps. It performs multiple critical functions: Upon receiving a new question, it performs structure-oriented analysis for the question. Then based on this analysis, it proceeds with a step-by-step reasoning to gradually solve the complex task. Within each step, it determines whether external information is needed and interacts with the Retrieval Agent to obtain ex-

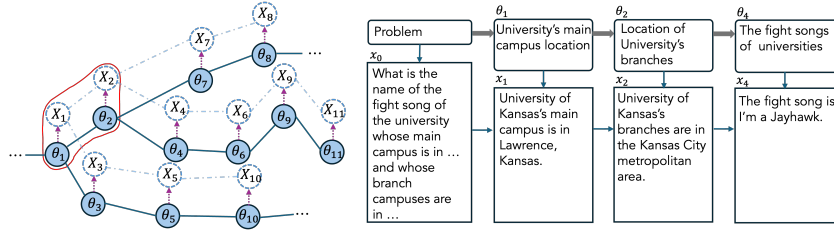


Figure 3: An illustrative example of the PGM generation model. This graph is a part of the underlying PGM where θ_i s are hidden variables and x_i s are observed variables. The red circle is an example of the strong connection between θ_i s and x_i s in the pre-training.

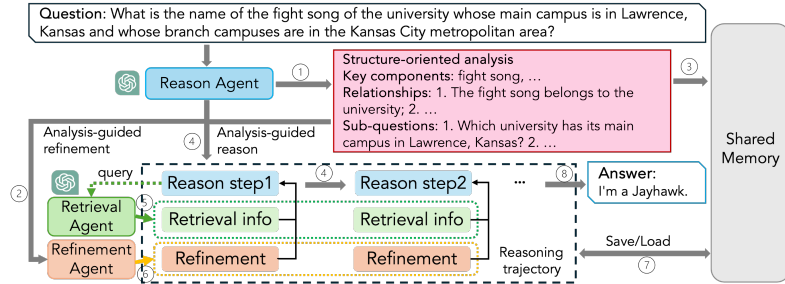


Figure 4: An overview of the Structure-oriented Autonomous Reasoning Agents.

ternal knowledge when necessary. It also interacts with the Refinement Agent for feedback on whether the step aligns with the original structure-oriented analysis and utilizes the feedback to refine the step. After completing the reasoning process, the Reason Agent consolidates a comprehensive final answer based on the entire reasoning trajectory. No human intervention is needed in this process.

Refinement Agent. To implement the reflection mechanism, we introduce a Refinement Agent, inspired by prior works on self-refinement (Madaan et al., 2024) and external supervision (Gou et al., 2023a; Shinn et al., 2024). This agent corrects potential errors within the Reason Agent and ensures that the reasoning process remains aligned with structure-oriented analysis. Specifically, it reviews each reasoning step based on the following three criteria: (1) alignment with the structure-oriented analysis, (2) consistency with the previous reasoning trajectory, and (3) factual correctness with relevant external knowledge. Such refinement operations can prevent the reasoning process from deviating from the structure-oriented analysis.

Retrieval Agent. This agent accesses external knowledge, including pre-constructed databases and web-based resources such as Wikipedia and Google Search, to provide complementary information for reasoning when requested by the Reason Agent. The retrieved knowledge is then provided to the Reason Agent to reduce factual errors.

Shared Memory. As the functionalities of both the Reason Agent and the Refinement Agent heavily

rely on the context of the reasoning process, a memory module is designed to store the structure-oriented analysis result, the reasoning trajectory, and the retrieved information. The Reason Agent and the Refinement Agent retrieve contexts from the shared memory to generate new reasoning steps or to consolidate the reasoning path.

4.2 Structure-oriented Reasoning Process

The whole reasoning process is in Figure 4.

Structure-oriented Analysis. In the enhanced system, when a new question is received, the Reason Agent conducts a thorough analysis (① in Figure 4) based on the syntactic structures of the problem. This analysis extracts critical components and generates relevant sub-questions for reference. For instance, in Figure 4 the question asks for the name of the fight song of a university with some constraints on the location of the main campus and branches. The Reason Agent identifies the key components as “fight song, university, main campus,...”, and the relationship is that “fight song” is the main objective while it belongs to “university” which is restricted by the location of “main campus”. Given these components, sub-questions can be further derived, e.g., “which university has its main campus located in . . .”. To ensure the reasoning accuracy, the initial analysis is sent to the Refinement Agent (② in Figure 4). The Refinement Agent provides an explicit reason for its judgments and refinements, storing in Memory (③ in Figure 4).

Iterative reasoning. To fully harness the reasoning

capability of LLMs, we adopt an iterative reasoning strategy (Yao et al., 2022; Wei et al., 2022; Li et al., 2023). As shown in Figure 4, in each iteration, Reason Agent takes the structure-oriented analysis and the previous reasoning trajectory to reason the current step (④ in Figure 4). If external knowledge is needed, the Reason Agent queries the Retrieval Agent (⑤ in Figure 4). The Retrieval Agent then searches for related information from external databases or web data and sends it back to the Reason Agent. For instance, if the current step is “what is the name of the university with the main campus in Lawrence Kansas”, the Reason Agent will interact with the Retrieval Agent to obtain “the University of Kansas” from Wikipedia. The Refinement Agent then evaluates and refines this step (⑥ in Figure 4). The refined steps are stored in the Shared Memory for use in subsequent iterations (⑦ in Figure 4) and synchronization of all agents. **Answer consolidation.** Finally, after the iterative reasoning process, the final answer is concluded (⑧ in Figure 4).

5 Experiments

5.1 Experiment setting

Agent configurations. We utilize the same LLM for all LLM-driven agents (Reason Agent, Refinement Agent and Retrieval Agent). Four representative general LLMs are tested, including two API-only models, GPT-4 and Qwen-max, and two open-source models, Llama3-70B and Qwen2-57B (Bai et al., 2023). We also use reasoning LLMs, DeepSeek-R1 (Guo et al., 2025) and OpenAI-o1 (Jaech et al., 2024) in Section 5.6. For the Retrieval Agent, we use Wikipedia API to obtain external knowledge. SARA is built with an open-source multi-agent framework, AgentScope (Gao et al., 2024). Detailed prompts are in Appendix D.

Tasks. We aim to improve the general reasoning capability of LLMs, so we test on various representative reasoning tasks, including HotpotQA (Yang et al., 2018) for multi-hop reasoning, Fever (Thorne et al., 2018) for fact verification, MMLU (Hendrycks et al., 2020) for multitask language understanding, StrategyQA (Geva et al., 2021) for commonsense reasoning ability, GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021) for math reasoning. Among all these tasks, HotpotQA, Fever, MMLU and StrategyQA can take advantage of external knowledge, so we group them as knowledge-intensive tasks. In terms of evalua-

tion metrics, the predicted solutions for HotpotQA and MATH are free-form answers, so we utilize a GPT-4 judge to assess the answer correctness and report the average accuracy as “LLM Acc”. For other datasets, we report the average accuracy as “Acc”. Details are provided in Appendix E.

Baselines. We compare SARA with common baselines and some representative reasoning methods: (1) Direct prompting (Vanilla) directly asks the LLM to answer the question. (2) In-context learning (ICL) asks the LLM to solve the problem given examples. (3) (few-shot) Chain-of-thought (CoT (Wei et al., 2022)) prompts the model to generate intermediate steps when solving the problem. (4) ReAct (Yao et al., 2022) combines agent thoughts (reason the current state) and actions (task-specific actions such as Search for an item with Wiki API) to help solve the problem. (5) Chain-of-knowledge (CoK (Li et al., 2023)) uses knowledge from different domains to correct reasoning rationales. Except for the direct prompting, all other baselines use a few-shot prompting strategy, and we test 6-shot as default to align with previous works (Yao et al., 2022; Li et al., 2023). (6) 0-shot CoT (Kojima et al., 2022). (7) 0-shot CoT with self-consistency (Wang et al., 2022) generates multiple CoT solutions and chooses one using a major vote. We generate 10 solutions. Examples of ICL and CoT are randomly selected from the training set for each task; reasoning steps in each CoT example are manually crafted. ReAct and CoK are implemented following the original paper.

5.2 Performance on knowledge-intensive tasks

The main results of SARA and the baselines on knowledge-intensive tasks are presented in Table 1. In general, SARA consistently outperforms all baselines across all tasks and models used in the experiments. For example, in HotpotQA, compared with baselines without explicit reasoning strategies, such as Vanilla and ICL, SARA achieves significant improvements of over 15% for most tasks. This suggests that even advanced models like GPT-4 and Qwen-max require proper strategies to fully leverage their reasoning capabilities, and simple examples alone are insufficient. To compare SARA with CoT, SARA also substantially improves the reasoning capability and surpasses CoT by over 10%. In terms of the ReAct and CoK, SARA also demonstrates clear advantages over them with average improvements of 4% and 4.4%, respectively, and the primary difference between these two methods and

Table 1: Main results on knowledge-intensive reasoning tasks.

Models	Tasks	Methods							
		Vanilla	ICL(6-shot)	CoT(6-shot)	ReAct(6-shot)	CoK(6-shot)	CoT(0-shot)	CoT-SC@10(0-shot)	SARA
GPT-4	HotpotQA	48.9%	51.4%	62.2%	67.2%	67.6%	52.3%	58.8%	73.5%
	Fever	35.3%	48.4%	56.1%	61.7%	61.3%	46.9%	53.1%	66.2%
	MMLU-BIO	94.1%	94.6%	95.3%	96.9%	96.7%	94.5%	95.7%	97.5%
	MMLU-PHY	65.3%	66.5%	69.4%	74.5%	73.9%	66.2%	68.2%	78.7%
Qwen-max	HotpotQA	49.6%	51.7%	58.3%	64.7%	66.3%	50.6%	56.7%	70.2%
	Fever	29.9%	39.1%	48.4%	58.2%	53.5%	41.5%	50.5%	63.1%
	MMLU-BIO	90.2%	91.3%	93.4%	93.9%	94.1%	91.6%	93.5%	96.2%
	MMLU-PHY	60.5%	56.2%	64.3%	71.8%	69.1%	60.7%	65.1%	75.4%
Qwen2-57B	HotpotQA	32.2%	33.5%	41.6%	53.9%	55.3%	35.1%	44.5%	58.7%
	Fever	21.5%	26.3%	44.7%	52.6%	51.3%	33.2%	45.6%	56.1%
	MMLU-BIO	86.1%	86.6%	87.4%	90.2%	90.9%	86.5%	87.9%	93.3%
	MMLU-PHY	53.2%	55.7%	63.4%	66.4%	68.3%	56.3%	63.8%	71.1%
Llama3-70B	HotpotQA	39.1%	38.2%	47.5%	56.2%	54.1%	40.6%	44.8%	60.9%
	Fever	46.4%	48.5%	53.1%	57.7%	58.2%	47.3%	51.9%	62.8%
	MMLU-BIO	89.2%	87.4%	89.5%	91.3%	91.7%	88.4%	89.2%	94.2%
	MMLU-PHY	47.9%	48.6%	55.3%	61.4%	60.9%	49.5%	55.7%	65.3%
Llama3-70B	HotpotQA	57.9%	65.1%	84.2%	85.2%	85.8%	72.5%	80.5%	87.1%

Table 2: Main results on math reasoning tasks.

	Tasks	Methods							
		Vanilla	ICL(6-shot)	CoT(6-shot)	ReAct(6-shot)	CoK(6-shot)	CoT(0-shot)	CoT-SC@10(0-shot)	SARA
GPT4	GSM8K	66.8%	66.9%	92.1%	93.7%	91.9%	84.3%	87.8%	94.2%
	MATH	43.1%	55.4%	69.2%	67.5%	68.6%	63.6%	64.1%	68.2%
Qwen-max	GSM8K	68.6%	72.8%	87.5%	89.2%	87.6%	74.8%	84.2%	91.3%
	MATH	42.8%	45.6%	64.9%	64.5%	65.3%	49.3%	61.9%	64.7%
Qwen2-57B	GSM8K	54.9%	59.2%	82.7%	83.9%	83.5%	63.7%	74.5%	84.4%
	MATH	30.1%	33.5%	46.2%	47.3%	46.8%	31.6%	40.8%	46.5%
Llama3-70B	GSM8K	55.3%	58.3%	83.7%	86.5%	87.2%	66.5%	76.8%	89.7%
	MATH	30.7%	32.4%	42.9%	46.3%	44.9%	32.8%	36.4%	44.2%

Table 3: Effect of each component in the reasoning agent. 'O' means include and 'X' means exclude.

Setting #	1	2	3	4	5	6	7
Key components	O	X	O	O	X	O	X
Sub-questions	O	O	X	O	O	X	X
Grammar/syntax	O	O	O	X	X	X	X
HotpotQA	73.5%	69.2%	69.4%	59.6%	58.6%	58.1%	56.5%
Fever	66.2%	61.7%	62.1%	53.4%	53.1%	52.9%	52.3%
MMLU-bio	97.5%	96.3%	96.6%	94.1%	94.3%	94.1%	93.9%
MMLU-phy	78.7%	74.1%	74.6%	59.5%	59.1%	57.2%	57.6%

SARA is our structure-oriented analysis. Moreover, our method outperforms 0-shot CoT SC@10, and also demonstrates significant advantages in other complex reasoning tasks such as HotpotQA, Fever, MMLU-PHY, and MMLU-BIO. Computation cost is summarized in Appendix H showing that SARA is also cost-effective.

5.3 Performance on math reasoning tasks

In Table 2, we present the main results of math reasoning tasks. Our method consistently outperforms 0-shot baselines and even works better than few-shot baselines on the GSM8K dataset. This shows that structure analysis can generalize well to math reasoning tasks. We do notice that SARA is not the best on the MATH dataset. This can be because some MATH problems are expressed in

symbols, which do not have clear structures for analysis. Nonetheless, SARA can still have comparably good results on MATH.

5.4 Effect of structure-oriented analysis

To elucidate the impact of the structure-oriented analysis, we conduct experiments evaluating the effectiveness of the three crucial functions in the Reason Agent: (1) key components and relationships between components, (2) sub-questions, and (3) grammar/syntax structure. Using GPT-4 on all reasoning tasks, we test different combinations of these elements, as detailed in Table 3.

There are several observations from Table 3. Consider HotpotQA as an example. First, comparing Settings 1, 2, and 3, when the grammar/syntax structure is included, removing either key components (Setting 2) or sub-questions (Setting 3) has only a small decrease in the performance. However, in Setting 4, excluding the grammar/syntax structure significantly reduces performance by over 10%, suggesting the importance of the grammar/syntax structure. Second, comparing Setting (1, 3) and (5, 7), without the key components and grammar/syntax structure analysis, formulating sub-questions only has limited im-

Table 4: Robustness evaluation, accuracy on GPT-4 after attack. Clean accuracy is included in brackets.

Attack	Task	Vanilla	ICL(6-shot)	CoT(6-shot)	ReAct(6-shot)	CoK(6-shot)	SARA
Badchain	HotpotQA	48.4%(48.9%)	13.7%(51.4%)	14.1%(62.2%)	21.3%(67.2%)	16.7% (67.6%)	71.3% (73.5%)
	Fever	35.5%(35.3%)	25.3%(48.4%)	12.1% (56.1%)	10.8% (61.7%)	21.8%(61.3%)	64.9% (66.2%)
Preemptive attack	HotpotQA	33.5% (48.9%)	42.1% (51.4%)	41.6% (62.2%)	55.3% (67.2%)	56.1% (67.6%)	68.2%(73.5%)
	Fever	19.2%(35.3%)	39.6%(48.4%)	32.2%(56.1%)	54.2%(61.7%)	52.3%(61.3%)	61.9%(66.2%)

Table 5: Comparison with reasoning models.

		Vanilla	ICL(6-shot)	CoT(6-shot)	ReAct(6-shot)	CoK(6-shot)	CoT(0-shot)	CoT-SC@10(0-shot)	SARA
R1	HotpotQA	62.3%	63.0%	60.2%	81.7%	76.9%	58.8%	67.4%	83.9%
	GSM8K	96.3%	95.8%	96.3%	97.2%	96.4%	96.6%	97.8%	97.6%
o1	HotpotQA	37.1%	38.6%	38.2%	67.8%	68.2%	36.5%	44.3%	70.4%
	GSM8K	96.6%	95.8%	97.6%	97.5%	98.1%	95.5%	97.6%	97.9%

provement of 1.9% on the reasoning performance, lower than 4.1% in Setting (1, 3). Similar observations can be found in Settings (1,2) and (6,7) for the key components, which indicates the synergy effect of grammar/syntax with key components and sub-questions. Third, completely removing the structure-oriented analysis also substantially diminishes reasoning performance. The above observations are consistent across all tasks.

5.5 Evaluation of robustness

Despite the improvement in the reasoning capability, we surprisingly find that SARA is robust to potential corruptions or distractions that target the reasoning process. We evaluate the robustness of SARA against two attacks: BadChain (Xiang et al., 2024), which targets few-shot reasoning by inserting backdoor reasoning steps through poisoned demonstrations; and Preemptive Attack (Xu et al., 2024), which targets zero-shot methods by embedding a malicious answer directly into the query to mislead reasoning. We test on HotpotQA and Fever with GPT-4, and the results are summarized in Table 4¹. When applying Badchain to our method, we simply replace the original input with input attached to the trigger. While few-shot baselines show high vulnerability to BadChain and Vanilla prompting performs poorly under Preemptive Attack, SARA effectively resists both types of attacks. The robustness of SARA can be attributed to two factors: (1) SARA’s zero-shot nature, which prevents malicious injections in demonstrations, and (2) the structure-oriented analysis, which focuses on syntax and grammar structures and avoids distractions in the problem.

5.6 Comparison with reasoning LLMs

Recently, some LLMs are specially trained to enhance reasoning capability, such as DeepSeek-R1 (Guo et al., 2025) and OpenAI-o1 (Jaech et al.,

¹Experimental details are provided in Appendix F

2024). Therefore, we conduct experiments to compare our method with these advanced reasoning models. We first leverage R1 and o1 as base models and follow the same setup as in the main experiments to compare performance. Then we also follow the setup in Section 5.5 to compare the robustness. We present results in Table 5 and 6 respectively. According to Table 5, reasoning models demonstrate strong reasoning capabilities and perform well with simple prompts. However, their performance can be suboptimal on datasets like HotpotQA, which involves multi-hop questions. Advanced methods like ReAct, CoK, and SARA can significantly enhance performance, suggesting that carefully designed advanced methods can further improve reasoning models. Moreover, combining results in Tables 1, 2 and 5, we notice that a non-reasoning model such as GPT-4 and Llama3-70B can achieve comparable or even better performance when applied SARA, compared with these reasoning models, further underscoring the potential of our method. With regard to robustness, as shown in Table 6, while both R1 and o1 show some resistance against these attacks, they are more vulnerable than SARA (based on GPT-4), especially on Preemptive attacks where reasoning process can be distracted from incorrect answers.

Table 6: Robustness comparison.

	R1 (Vanilla)	o1 (Vanilla)	SARA (GPT-4)
Clean	62.3%	37.1%	73.5%
Badchain	58.4%	30.6%	70.7%
Prem	52.9%	25.3%	68.2%

6 Conclusion

In this paper, inspired by human cognition, we introduce structure-oriented analysis to encourage LLMs to understand the query in a more formulated way. Utilizing the analysis, LLMs can better identify key steps when performing reasoning tasks, improving the performance. Furthermore, built upon the structure-oriented analysis, we further establish

a multi-agent reasoning system to the LLM’s reasoning process. Experiments have demonstrated the effectiveness of the proposed framework in knowledge-intensive tasks, math reasoning tasks, and is also effective for reasoning models.

Limitation

Although our strategy shows effectiveness on diverse reasoning tasks, including knowledge-intensive reasoning, math reasoning, and common-sense reasoning, we notice that our method works better on problems that are clearly described in natural languages, such as GSM8K, while performs worse on pure symbol expressions as no obvious structures appear like some questions in MATH dataset. This suggests a future direction for extracting logic structures and learning symbolic expressions to improve reasoning capability. Besides, the LLM agent we adopt to illustrate our principal strategy is simple to fit in various tasks, which can still have room for improvement. Modifying the agent system while maintaining the core structure analysis to adapt to different tasks can be a potential direction. For example, when solving math problems, instead of the Retrieve Agent, leveraging external tools like a calculator or code executor to improve the performance.

Acknowledgement

Pengfei He and Jiliang Tang are supported by the National Science Foundation (NSF) under grant numbers CNS2321416, IIS2212032, IIS2212144, IIS 2504089, DUE2234015, CNS2246050, DRL2405483 and IOS2035472, the Michigan Department of Agriculture and Rural Development, US Dept of Commerce, Amazon Faculty Award, Meta, NVIDIA, Microsoft and SNAP. Yue Xing is partially supported by NSF DMS 2515194 and Open Philanthropy.

References

Janice Ahn, Rishu Verma, Renze Lou, Di Liu, Rui Zhang, and Wenpeng Yin. 2024. Large language models for mathematical reasoning: Progresses and challenges. *arXiv preprint arXiv:2402.00157*.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, and 1 others. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.

Forrest Bao, Chenyu Xu, and Ofer Mendelevitch. 2025.

[Deepseek-r1 hallucinates more than deepseek-v3](#). Accessed: 2025-05-18.

Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, and 1 others. 2024. Graph of thoughts: Solving elaborate problems with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17682–17690.

Sijia Chen, Baochun Li, and Di Niu. 2024a. Boosting of thoughts: Trial-and-error problem solving with large language models. *arXiv preprint arXiv:2402.11140*.

Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, and 1 others. 2024b. Do not think that much for $2+3=?$ on the overthinking of o1-like llms. *arXiv preprint arXiv:2412.21187*.

Yanda Chen, Joe Benton, Ansh Radhakrishnan, Jonathan Uesato, Carson Denison, John Schulman, Arushi Somani, Peter Hase, Misha Wagner, Fabien Roger, and 1 others. 2025. Reasoning models don’t always say what they think. *arXiv preprint arXiv:2505.05410*.

Michelene TH Chi, Paul J Feltovich, and Robert Glaser. 1981. Categorization and representation of physics problems by experts and novices. *Cognitive science*, 5(2):121–152.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, and 1 others. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.

Andrew Drozdov, Nathanael Schärli, Ekin Akyürek, Nathan Scales, Xinying Song, Xinyun Chen, Olivier Bousquet, and Denny Zhou. 2022. Compositional semantic parsing with large language models. In *The Eleventh International Conference on Learning Representations*.

Hugging Face. 2023. [Metric: exact_match](#). Accessed: 2024-10-01.

Dawei Gao, Zitao Li, Weirui Kuang, Xuchen Pan, Daoyuan Chen, Zhijian Ma, Bingchen Qian, Liuyi Yao, Lin Zhu, Chen Cheng, and 1 others. 2024. Agentscope: A flexible yet robust multi-agent platform. *arXiv preprint arXiv:2402.14034*.

Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.

Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. Did aristotle use a laptop? a question answering benchmark with

- implicit reasoning strategies. *Transactions of the Association for Computational Linguistics*, 9:346–361.
- Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujia Yang, Nan Duan, and Weizhu Chen. 2023a. Critic: Large language models can self-correct with tool-interactive critiquing. *arXiv preprint arXiv:2305.11738*.
- Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujia Yang, Minlie Huang, Nan Duan, and Weizhu Chen. 2023b. Tora: A tool-integrated reasoning agent for mathematical problem solving. *arXiv preprint arXiv:2309.17452*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- 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. *arXiv preprint arXiv:2103.03874*.
- Sirui Hong, Yizhang Lin, Bangbang Liu, Binhao Wu, Danyang Li, Jiaqi Chen, Jiayi Zhang, Jinlin Wang, Lingyao Zhang, Mingchen Zhuge, and 1 others. 2024. Data interpreter: An llm agent for data science. *arXiv preprint arXiv:2402.18679*.
- Jie Huang and Kevin Chen-Chuan Chang. 2022. Towards reasoning in large language models: A survey. *arXiv preprint arXiv:2212.10403*.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, and 1 others. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. 2022. Decomposed prompting: A modular approach for solving complex tasks. *arXiv preprint arXiv:2210.02406*.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Kenneth Kotovsky, John R Hayes, and Herbert A Simon. 1985. Why are some problems hard? evidence from tower of hanoi. *Cognitive psychology*, 17(2):248–294.
- George Lakoff and Mark Johnson. 2008. *Metaphors we live by*. University of Chicago press.
- Xingxuan Li, Ruochen Zhao, Yew Ken Chia, Bosheng Ding, Shafiq Joty, Soujanya Poria, and Lidong Bing. 2023. Chain-of-knowledge: Grounding large language models via dynamic knowledge adapting over heterogeneous sources. *arXiv preprint arXiv:2305.13269*.
- Wei Ma, Shangqing Liu, Zhihao Lin, Wenhan Wang, Qiang Hu, Ye Liu, Cen Zhang, Liming Nie, Li Li, and Yang Liu. 2023. Lms: Understanding code syntax and semantics for code analysis. *arXiv preprint arXiv:2305.12138*.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, and 1 others. 2024. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36.
- Dheeraj Mekala, Jason Wolfe, and Subhro Roy. 2022. Zerotopt: Zero-shot task-oriented semantic parsing using large language models. *arXiv preprint arXiv:2212.10815*.
- OpenAI. 2025. [Openai o3 and o4-mini system card](#). Accessed: 2025-05-18.
- Debjit Paul, Mete Ismayilzada, Maxime Peyrard, Beatriz Borges, Antoine Bosselut, Robert West, and Boi Faltings. 2023. Refiner: Reasoning feedback on intermediate representations. *arXiv preprint arXiv:2304.01904*.
- Archiki Prasad, Alexander Koller, Mareike Hartmann, Peter Clark, Ashish Sabharwal, Mohit Bansal, and Tushar Khot. 2023. Adapt: As-needed decomposition and planning with language models. *arXiv preprint arXiv:2311.05772*.
- Ben Prystawski, Michael Li, and Noah Goodman. 2024. Why think step by step? reasoning emerges from the locality of experience. *Advances in Neural Information Processing Systems*, 36.
- Shuofei Qiao, Yixin Ou, Ningyu Zhang, Xiang Chen, Yunzhi Yao, Shumin Deng, Chuanqi Tan, Fei Huang, and Huajun Chen. 2022. Reasoning with language model prompting: A survey. *arXiv preprint arXiv:2212.09597*.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2024. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36.
- Kumar Shridhar, Harsh Jhamtani, Hao Fang, Benjamin Van Durme, Jason Eisner, and Patrick Xia. 2023a. Screws: A modular framework for reasoning with revisions. *arXiv preprint arXiv:2309.13075*.

- Kumar Shridhar, Koustuv Sinha, Andrew Cohen, Tianlu Wang, Ping Yu, Ram Pasunuru, Mrinmaya Sachan, Jason Weston, and Asli Celikyilmaz. 2023b. The art of llm refinement: Ask, refine, and trust. *arXiv preprint arXiv:2311.07961*.
- Kumar Shridhar, Alessandro Stolfo, and Mrinmaya Sachan. 2022. Distilling reasoning capabilities into smaller language models. *arXiv preprint arXiv:2212.00193*.
- Herbert A Simon and Allen Newell. 1971. Human problem solving: The state of the theory in 1970. *American psychologist*, 26(2):145.
- Theodore R Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas L Griffiths. 2023. Cognitive architectures for language agents. *arXiv preprint arXiv:2309.02427*.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. Fever: a large-scale dataset for fact extraction and verification. *arXiv preprint arXiv:1803.05355*.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions. *arXiv preprint arXiv:2212.10509*.
- Rasul Tutunov, Antoine Grosnit, Juliusz Ziomek, Jun Wang, and Haitham Bou-Ammar. 2023. Why can large language models generate correct chain-of-thoughts? *arXiv preprint arXiv:2310.13571*.
- Tu Vu, Mohit Iyyer, Xuezhi Wang, Noah Constant, Jerry Wei, Jason Wei, Chris Tar, Yun-Hsuan Sung, Denny Zhou, Quoc Le, and 1 others. 2023. Freshllms: Refreshing large language models with search engine augmentation. *arXiv preprint arXiv:2310.03214*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Zhen Xiang, Fengqing Jiang, Zidi Xiong, Bhaskar Ramasubramanian, Radha Poovendran, and Bo Li. 2024. Badchain: Backdoor chain-of-thought prompting for large language models. *arXiv preprint arXiv:2401.12242*.
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2021. An explanation of in-context learning as implicit bayesian inference. *arXiv preprint arXiv:2111.02080*.
- Rongwu Xu, Zehan Qi, and Wei Xu. 2024. Preemptive answer" attacks" on chain-of-thought reasoning. *arXiv preprint arXiv:2405.20902*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. *arXiv preprint arXiv:1809.09600*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2024. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems*, 36.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.
- Tianyang Zhong, Zhengliang Liu, Yi Pan, Yutong Zhang, Yifan Zhou, Shizhe Liang, Zihao Wu, Yanjun Lyu, Peng Shu, Xiaowei Yu, and 1 others. 2024. Evaluation of openai o1: Opportunities and challenges of agi. *arXiv preprint arXiv:2409.18486*.
- Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. 2023. Language agent tree search unifies reasoning acting and planning in language models. *arXiv preprint arXiv:2310.04406*.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, and 1 others. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*.
- Pei Zhou, Jay Pujara, Xiang Ren, Xinyun Chen, Heng-Tze Cheng, Quoc V Le, Ed H Chi, Denny Zhou, Swaroop Mishra, and Huaixiu Steven Zheng. 2024. Self-discover: Large language models self-compose reasoning structures. *arXiv preprint arXiv:2402.03620*.
- Xizhou Zhu, Yuntao Chen, Hao Tian, Chenxin Tao, Weijie Su, Chenyu Yang, Gao Huang, Bin Li, Lewei Lu, Xiaogang Wang, and 1 others. 2023. Ghost in the minecraft: Generally capable agents for open-world environments via large language models with text-based knowledge and memory. *arXiv preprint arXiv:2305.17144*.

The structure of the appendix is as follows: In Section A, we provide the detailed version of Section 3.2 with the mathematical notations, the formal statement of Theorem 3.1 and the corresponding proofs. Prompts and additional details of experiments in Section 3.1 are provided in Section C. Detailed prompts of agents are included in Section D. Experiment (Section 5) details and additional results are presented in Section E and Section G respectively.

A Theoretical Analysis

A.1 Theoretical analysis

In addition to the PGM introduced in Section 3.2, we provide more details on our assumption in the LLM and the notations of the reasoning path. Then we provide a formal statement of Theorem 3.1.

LLM in pretraining. Recall that in Figure 3, the PGM contains hidden variables $\{\theta_i\}_{i=1}^N$ as the observed variables $\{X_i\}_{i=1}^N$ with the explicit knowledge $\{x_i\}_{i=1}^N$. Following a similar idea as in (Prystawski et al., 2024), when using the above pre-training data to train an LLM \mathcal{M} , the output of \mathcal{M} satisfies the following properties. First, most existing LLMs used for complex tasks demonstrate reliable capability in telling whether two given pieces of explicit knowledge share the same abstract concept or not (i.e., whether x_i and x_j share the same θ). Based on this, we assume that the LLMs can faithfully capture the relationship between the hidden variables and the corresponding explicit knowledge (i.e., the edges between θ_i and X_i). Moreover, since most LLMs are trained for next-token prediction, explicit knowledge and abstract concepts that frequently appear in nearby within texts (i.e., the connections between x_i and x_j as well as the connection between θ_i and θ_j) are also learned by LLMs with high quality. For example, information about the main campus of the University of Kansas and its branches often appears within the same paragraph on a Wikipedia page; generally, the location of universities and their branches locations usually appear close in text.

Use PGM to explain the reasoning process. In Section 3.2, we intuitively explain the reasoning process using the examples in Figure 3. The detailed mathematical description of the reasoning procedure is as follows. The model \mathcal{M} receives an input question x_0 , e.g., “find the name of the fight song of the university whose main campus is in ...” in the right panel of Figure 3, and the

target is to infer the answer via exploring different variables in the PGM. Define a *reasoning path* γ as a set of indexes $\{s_i\}$ of hidden and observed variables (θ_{s_i}, x_{s_i}) . The *correct reasoning path* γ^* is an ideal reasoning path that both logically correct and leading to the final correct answer. As for the example in Figure 3, the correct reasoning path is $\gamma^* := 1 \rightarrow 2 \rightarrow 4$, i.e., exploring through hidden states $\theta_1 \rightarrow \theta_2 \rightarrow \theta_4$. Ideally, if \mathcal{M} follows γ^* , it will output $x_1|x_2|x_4$. However, because the abstract concepts and explicit knowledge in multi-hop reasoning of a complex question are unlikely to appear in pre-training data all close to each other, \mathcal{M} has no direct knowledge of γ^* but can only focus on the next variable exploration based on the edges in PGM when reasoning. As a result, instead of the correct reasoning path γ^* , we assume that \mathcal{M} explores actual reasoning path step by step: given s_i and x_{s_i} , \mathcal{M} explores $\theta_{s_{i+1}}$ and generates $x_{s_{i+1}}$ from $X_{s_{i+1}}|x_{s_i}, \theta_{s_{i+1}}$, and all the explored s_i s together form the reasoning path γ . The γ also involves randomness since \mathcal{M} is a generation model. Finally, to ease the later analysis, denote $\Gamma(x_0, \cdot, \mathcal{M})$ and $\Gamma(x_0, \theta_T, \mathcal{M})$ as the set of all possible reasoning paths and the set of all *correct* paths respectively, where θ_T is the correct final reasoning step (the target).

In the following, we analyze how additional information about intermediate variables lying on the correct reasoning path benefits multi-step reasoning.

Quantify the benefit of correct intermediate variables. Given x_0 , we denote $\mathcal{E}(\gamma)$ as *reasoning error* for a given reasoning path γ to quantify the performance and $e(\Gamma) \triangleq \sum_{\gamma \in \Gamma} P(\gamma)\mathcal{E}(\gamma)$ as the *expected reasoning error* for a set of paths Γ , and study how the choice of Γ affects $e(\Gamma)$.

When performing the reasoning with the structure-oriented analysis, the analysis can extract a sequence of indices of latent variables $A = \{s_1^A, s_2^A, \dots\}$, which can be key components or sub-questions in practice as shown in Figure 1. In the following, we first provide some mild assumptions on γ , and then demonstrate how the reasoning error is impacted by A .

Assumption A.1. Given x_0 , the random variable γ satisfies the following conditions: (1) $\Gamma(x_0, \theta_T, \mathcal{M})$ contains only one path: $\Gamma(x_0, \theta_T, \mathcal{M}) = \{\gamma^*\}$. (2) $\mathcal{E}(\gamma) \geq 0$ and equals to 0 iff $\gamma = \gamma^*$.

In Assumption A.1, the first condition in As-

sumption A.1 assumes a unique correct path. Discussion for a relaxed version for multiple correct paths can be found in Remark A.4. In the second condition, the reasoning error is zero only when we explore the correct path.

Given the above notations and assumptions, the following result holds:

Lemma A.2. *Let $\Gamma_A(x_0, \cdot, \mathcal{M})$ denote the set of explored paths given A . Under Assumption A.1, assume that $A \subseteq \gamma^*$, then the following results in θ_T (with the corresponding index T) and γ hold:*

(1) *When $|A| = 1$, i.e. $A = \{s^A\}$ for some $s^A \in \gamma^*$, then $P(T \in \gamma | s^A \in \gamma) \geq P(T \in \gamma)$ where the equality holds if and only if $P(s^A \in \gamma) = 1$.*

(2) *When $|A| > 1$, i.e. $A = \{s_1^A, \dots, s_k^A\}$, and $A \subseteq \gamma^*$, we have a sequence of inequalities*

$$P(T \in \gamma | A \subseteq \gamma) \geq P(T \in \gamma | \{s_j^A\}_{j \in [k-1]} \subseteq \gamma) \geq \dots \geq P(T \in \gamma).$$

The proof of Lemma A.2 can be found in Appendix A.2. Based on Lemma A.2, when the LLM follows A and explores the variables $\{s_j^A\}_{j \in [k]}$, there is a higher chance that it finally explores θ_T .

Besides the probability of reaching θ_T considered in Lemma A.2, the following theorem presents the results on how the expected reasoning error is impacted by A . We consider two specific errors: (1) 0-1 error $\mathcal{E}_{0-1}(\gamma) = \mathbf{1}(T \notin \gamma)$, and (2) the probability error considered in (Prystawski et al., 2024)

$$\begin{aligned} & \mathcal{E}_{\text{prob}}(\gamma) \\ = & \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in G}} [p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma}) \\ & - p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in G})]^2 \end{aligned}$$

with G as all variables in the PGM. We quantify the expected reasoning error as follows:

Theorem A.3. *Under the assumptions in Lemma A.2, for $\mathcal{E} \in \{\mathcal{E}_{0-1}, \mathcal{E}_{\text{prob}}\}$, the following holds:*

(1) *When $|A| = 1$, i.e. $A = \{s^A\}$ for some $s^A \in \gamma^*$,*

$$e(\Gamma_A(x_0, \cdot, \mathcal{M})) \leq e(\Gamma(x_0, \cdot, \mathcal{M}))$$

where the equality holds only if $P(s^A \in \gamma) = 1$.

(2) *When $|A| > 1$, i.e. $A = \{s_1^A, \dots, s_k^A\}$, and $A \subseteq \gamma^*$, we have a sequence of inequalities*

$$\begin{aligned} e(\Gamma_A(x_0, \cdot, \mathcal{M})) & \leq e(\Gamma_{\{s_j^A\}_{j \in [k-1]}}(x_0, \cdot, \mathcal{M})) \\ & \leq \dots \leq e(\Gamma(x_0, \cdot, \mathcal{M})). \end{aligned}$$

The proof of Theorem A.3 can be found in Appendix A.2. Theorem A.3 implies that given the information of the variables on the correct path, the reasoning error is reduced.

Remark A.4 (Multiple correct paths). Though Assumption A.1 assumes a unique correct path γ^* , it is possible that there exist multiple correct paths in practice. The above result also holds when multiple correct paths exist given some mild conditions on A . Suppose there exist multiple correct paths, i.e. $\Gamma^* = \{\gamma_1^*, \gamma_2^*, \dots\}$, and we assume that $\mathcal{E}(\gamma_i^*) = 0$ for these reasoning paths. We still consider a sequence of indices of latent variables $A = \{s_1^A, s_2^A, \dots\}$ lying on these correct paths. In particular, we assume there is a subset A^* , such that every index in A^* lies on every correct path, denoted as $A^* \subseteq \Gamma^*$. Then the results in Theorem A.3 still hold by replacing A with A^* and γ^* with Γ^* . This is because errors of paths out of Γ^* are all positive, and information of A^* significantly increases the probability of inferring paths in Γ^* and thus decreases the reasoning error.

Remark A.5 (Error when the exploration is not guaranteed to find θ_s for some $s \in A$). In practice, when searching a proper reasoning path, it is possible that the exploration does not guarantee to reach θ_s for $s \in A$ for sure. Assume $|A| = 1$. In this case, denote $\Gamma \setminus \Gamma_A$ as the reasoning path that does not pass A , and then the total error becomes

$$\begin{aligned} & P(\theta_s \text{ is reached})e(\Gamma_s(x_0, \cdot, \mathcal{M})) \\ & + P(\theta_s \text{ is not reached})e(\Gamma \setminus \Gamma_A(x_0, \cdot, \mathcal{M})), \end{aligned}$$

and for \mathcal{E}_{0-1} and $\mathcal{E}_{\text{prob}}$, $e(\Gamma \setminus \Gamma_A(x_0, \cdot, \mathcal{M})) \geq e(\Gamma_A(x_0, \cdot, \mathcal{M}))$ as long as the exploration reaches s with a higher chance than random search.

A.2 Proofs 3

A.2.1 Proof of Lemma A.2

Proof of Lemma A.2. The proof of Lemma A.2 mainly utilizes the definition of conditional probability. We start from the simple case where $|A| = 1$.

Single variable in A . When $A = \{s^A\}$, i.e., only a single variable in A , we have

$$\begin{aligned} P(T \in \gamma) & = P(T \in \gamma | s^A \in \gamma) \underbrace{P(s^A \in \gamma)}_{\leq 1} \\ & + \underbrace{P(T \in \gamma | s^A \notin \gamma)}_{=0} P(s^A \notin \gamma) \\ & \leq P(T \in \gamma | s^A \in \gamma). \end{aligned}$$

Multiple variables in A . When there are multiple variables in A , i.e. $s_1^A, s_2^A, \dots, s_k^A$, repeat the above analysis, we have

$$\begin{aligned} P(T \in \gamma) &= P(T \in \gamma | A \subseteq \gamma) P(A \subseteq \gamma) \\ &+ \underbrace{P(T \in \gamma | A \not\subseteq \gamma)}_{=0} P(A \not\subseteq \gamma) \\ &= P(T \in \gamma | A \subseteq \gamma) P(A \subseteq \gamma). \end{aligned}$$

Furthermore, it is easy to see that $P(\cap_{j=1}^{i+1} \{s_j^A \in A\}) \leq P(\cap_{j=1}^i \{s_j^A \in A\})$, which implies that

$$P(T \in \gamma | \{s_j^A\}_{j \in [i+1]}) \geq P(T \in \gamma | \{s_j^A\}_{j \in [i]})$$

Then we have a sequence of inequalities

$$\begin{aligned} P(T \in \gamma | A \subseteq \gamma) &\geq P(T \in \gamma | \{s_j^A\}_{j \in [k-1]} \subseteq \gamma) \\ &\geq \dots \geq P(T \in \gamma) \end{aligned}$$

which completes the proof. \square

A.2.2 Expected reasoning loss with specific error functions

We discuss two representative error functions, 0-1 error and probability error, in Theorem A.3.

0-1 error. Recall that for a given reasoning path γ , we define 0-1 error function as

$$\mathcal{E}(\gamma) = \mathbf{1}(T \notin \gamma),$$

where T represents the index of the target variable. This function assigns an error of 0 when the reasoning path reaches the target variable, and 1 otherwise. This binary error metric is both practical and commonly used in evaluating reasoning performance, as it focuses on the logical correctness of the reasoning process. It closely relates to popular empirical metrics such as exact match (EM) (Face, 2023).

Proof of Theorem A.3, 0-1 error. Given the above definition of 0-1 error, we have

$$\begin{aligned} e(\Gamma(x_0, \cdot, \mathcal{M})) &= \sum \mathcal{E}(\gamma) P(\gamma) \\ &= \sum_{T \notin \gamma} P(\gamma) = P(T \notin \gamma), \end{aligned}$$

and

$$\begin{aligned} e(\Gamma_A(x_0, \cdot, \mathcal{M})) &= \sum_{T \notin \gamma} P(\gamma | A \subseteq \gamma) \\ &= P(T \notin \gamma | A \subseteq \gamma), \end{aligned}$$

both of which are reduced to the probability of T being reached by the reasoning process. As a result,

following Lemma A.2, we have $e(\Gamma(x_0, \cdot, \mathcal{M})) \geq e(\Gamma_A(x_0, \cdot, \mathcal{M}))$.

Furthermore, given that $P(T \in \gamma | A \subseteq \gamma) = P(T \in \gamma) / P(A \subseteq \gamma)$, a decrease in $P(A \subseteq \gamma)$ leads to an increase in the improvement gained by conditioning on A . This implies that for more complex problems where inferring critical steps in A is challenging, extracting information of A through analysis becomes increasingly important. Following the steps in Lemma A.2, we also have

$$\begin{aligned} e(\Gamma_A(x_0, \cdot, \mathcal{M})) &\leq e(\Gamma_{\{s_j^A\}_{j \in [k-1]}}(x_0, \cdot, \mathcal{M})) \\ &\leq \dots \leq e(\Gamma(x_0, \cdot, \mathcal{M})). \end{aligned}$$

\square

Probability error. Recall that the probability error is defined as

$$\mathcal{E}(\gamma) = \mathbb{E}_{\{(X_i, \theta_i)\}} [p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma}) - p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in G})]^2.$$

where x_t is the ground truth output for the target step. The first term is the probability of predicting ground truth given path γ while the second term is the probability of predicting the ground truth given the underlying PGM. This error is connected with the widely used cross-entropy loss (Prystawski et al., 2024).

The following lemma presents a valid decomposition of the probability error. Denote $G \setminus \gamma$ as the set of indexes in all paths excluding γ .

Lemma A.6 (Decomposition of probability error.). *The following decomposition holds:*

$$\begin{aligned} \mathcal{E}(\gamma) &= \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma}} \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in G \setminus \gamma}} \\ &\quad \left[p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma}) \right. \\ &\quad \left. - p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in G}) \right]^2 \\ &= \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma}} \left[p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma}) \right. \\ &\quad \left. - \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in G \setminus \gamma}} p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in G}) \right]^2 \\ &\quad + \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma}} \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in G \setminus \gamma}} \\ &\quad \left[p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in G}) - \right. \\ &\quad \left. \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in G \setminus \gamma}} p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in G}) \right]^2 \end{aligned}$$

When $\gamma = \gamma^*$,

$$\mathcal{E}(\gamma) = 0.$$

The decomposition in Lemma A.6 consists of two parts, where the first part represents the bias of prediction for a given path γ while the second term represents the variance.

Given the above decomposition, below is the proof of Theorem A.3 for the probability error:

Proof of Theorem A.3, probability error. Similar to the proof of Lemma A.2, we start from the simple case where $|A| = 1$.

Simple variable in A . If the model \mathcal{M} can always explore a path with an intermediate variable θ_{s^A} lying in the correct reasoning path γ^* , then

$$\begin{aligned}
& e(\Gamma_A(x_0, \cdot, \mathcal{M})) \\
&= \sum_{T \notin \gamma, \gamma \in \Gamma_A(x_0, \cdot, \mathcal{M})} P(\gamma | s^A \in \gamma) \mathcal{E}(\gamma) \quad (1) \\
&+ \sum_{T \in \gamma, \gamma \in \Gamma_A(x_0, \cdot, \mathcal{M})} P(\gamma | s^A \in \gamma) \mathcal{E}(\gamma) \\
&= \sum_{T \notin \gamma} \frac{P(\gamma, s^A \in \gamma)}{P(s^A \in \gamma)} \mathcal{E}(\gamma) \quad (2) \\
&+ \sum_{T \in \gamma} \frac{P(\gamma, s^A \in \gamma)}{P(s^A \in \gamma)} \mathcal{E}(\gamma) \\
&= \sum_{T \notin \gamma} \frac{P(\gamma, s^A \in \gamma)}{P(s^A \in \gamma)} \mathcal{E}(\gamma).
\end{aligned}$$

Now we look at the different values of $\mathcal{E}(\gamma)$ when changing γ . Note that from how the PGM is constructed, we have

$$\begin{aligned}
& p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma}) \\
&= p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma^* \cap \gamma}),
\end{aligned}$$

and

$$\begin{aligned}
& p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in G}) \\
&= p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma^*}).
\end{aligned}$$

For any two reasoning paths γ_1 and γ_2 so that $s^A \notin \gamma_1$ but $s^A \in \gamma_2$, following similar decompositions as in Lemma A.6, we have

$$\begin{aligned}
& \mathcal{E}(\gamma_1) \\
&= \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma \cap \gamma^*}} \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma^* \cap (\gamma_2 \setminus \gamma_1)}} \\
&\quad \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma^* \setminus \gamma_2}} \\
&\quad \left[p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma_1 \cap \gamma^*}) \right. \\
&\quad \quad \left. - p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma^*}) \right]^2 \\
&= \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma \cap \gamma^*}} \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma^* \cap (\gamma_2 \setminus \gamma_1)}} \\
&\quad \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma^* \setminus \gamma_2}}
\end{aligned}$$

$$\begin{aligned}
& \left[p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma_1 \cap \gamma^*}) \right. \\
&\quad \left. - p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma_2 \cap \gamma^*}) \right. \\
&\quad \left. + p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma_2 \cap \gamma^*}) \right. \\
&\quad \left. - p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma^*}) \right]^2 \\
&= \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma \cap \gamma^*}} \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma^* \cap (\gamma_2 \setminus \gamma_1)}} \\
&\quad \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma^* \setminus \gamma_2}} \\
&\quad \left[p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma_1 \cap \gamma^*}) \right. \\
&\quad \quad \left. - p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma_2 \cap \gamma^*}) \right]^2 \\
&+ \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma \cap \gamma^*}} \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma^* \cap (\gamma_2 \setminus \gamma_1)}} \\
&\quad \left[p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma_2 \cap \gamma^*}) \right. \\
&\quad \quad \left. - p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma^*}) \right]^2 \\
&\geq \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma \cap \gamma^*}} \mathbb{E}_{\{(X_i, \theta_i)\}_{i \in \gamma^* \cap (\gamma_2 \setminus \gamma_1)}} \\
&\quad \left[p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma_2 \cap \gamma^*}) \right. \\
&\quad \quad \left. - p(X_T = x_t | x_0, \{(X_i, \theta_i)\}_{i \in \gamma^*}) \right]^2 \\
&= \mathcal{E}(\gamma_2),
\end{aligned}$$

from which it is easy to see that

$$e(\Gamma(x_0, \cdot, \mathcal{M})) \geq e(\Gamma_A(x_0, \cdot, \mathcal{M})).$$

Multiple variables in A . When $|A| > 1$, the steps are indeed the same as when $|A| = 1$. We prove the relationship between $\mathcal{E}(\gamma_1) \geq \mathcal{E}(\gamma_2)$ for different s_i^A s. \square

B Additional experiments

We provide additional experimental results, including effect of key agents and additional baselines.

B.1 Effect of key agents

In this subsection, we study the effect of two key agents in SARA, the Refinement Agent and the Retrieve Agent. We test with GPT-4 model on HotpotQA and Fever benchmarks and summarize the results in Figure 5. When replacing the original LLM (GPT-4) with a smaller model (Qwen2-57) in the Retrieval Agent, the performance is barely affected; while for the Refine Agent, the performance drops a bit more. This suggests that it is feasible to utilize a smaller model in the Retrieval Agent for efficiency while maintaining effectiveness, but the Refine Agent requires strong models. It is noted that removing either agent will decrease the reasoning capacity of the system. Moreover, without the Refinement Agent, SARA still has a

comparable performance with ReAct and CoK (Table 1), and without the Retrieval Agent, SARA can also achieve better results than 6-shot CoT (no retrieval as well). These highlight the effectiveness of structure-oriented analysis.




Figure 5: Ablation study on agents.

B.2 Additional baselines

In this subsection, we include two additional baselines to further illustrate the advantage of the proposed method. Boost-of-Thought (BoT) (Chen et al., 2024a) is an automated prompting framework for problem solving with LLMs by iteratively exploring and self-evaluating many trees of thoughts in order to acquire an ensemble of trial-and-error reasoning experiences. We follow the official code and implement BoT+CoT. Least-to-most (Zhou et al., 2022) is a representative task-decomposing method that break down a complex problem into a series of simpler subproblems and then solve them in sequence. We follow the instructions in the original paper and let GPT generate task decomposition prompts. It is worth noting that both baselines are few-shot methods. Results are shown in Table 7.

According to the results, SARA outperforms baselines for most cases, indicating its effectiveness in solving problems. For other cases when SARA is not the best, it achieves comparable performance. Combining with the fact that both baselines are few-shot method while SARA is a 0-shot method, SARA significantly reduce the performance gap between few-shot reasoning and 0-shot reasoning.

C Details for experiments in Section 3

Prompt for structure-oriented analysis. To add the structure-oriented analysis on top of the backbone reasoning method, we develop the following prompt to let the model identify critical components, relationships among them, and related sub-questions. The LLM is also prompted to provide justification for its analysis.

structure-oriented analysis

You are a helpful assistant good at parsing the syntax and grammar structure

Table 7: Additinoal baselines (BoT, Least-to-most)

		BoT	Least-to-most	SARA
GPT-4	MMLU-Bio	97.2	93.4	97.5
	MMLU-phy	76.9	73.0	78.7
	GSM8K	98.7	90.8	94.2
	MATH	66.3	60.3	68.2
Qwen-max	MMLU-Bio	96.4	94.7	96.2
	MMLU-phy	72.8	63.4	75.4
	GSM8K	92.5	87.9	91.3
	MATH	64.1	67.6	64.7
Qwen2-57B	MMLU-Bio	91.5	86.9	93.3
	MMLU-phy	71.9	60.3	71.1
	GSM8K	84.8	75.5	84.4
	MATH	45.2	41.4	46.5
Llama3-70B	MMLU-Bio	92.7	88.3	94.2
	MMLU-phy	64.6	61.7	65.3
	GSM8K	89.5	74.9	89.7
	MATH	43.9	37.2	44.2

of sentences. Please first analyze the syntax and grammar structure of the problem and provide a thorough analysis by addressing the following tasks:

1. Identify Key Components: Identify the crucial elements and variables that play a significant role in this problem.
 2. Relationship between Components: Explain how the key components are related to each other in a structured way.
 3. Sub-Question Decomposition: Break down the problem into the following sub-questions, each focusing on a specific aspect necessary for understanding the solution.
 4. Implications for Solving the Problem: For each sub-question, describe how solving it helps address the main problem. Connect the insights from these sub-questions to the overall strategy needed to solve the main problem.
- Question:

Examples for CoT. For 0-shot CoT, we use the simple prompt “Please think step by step” as in (Kojima et al., 2022). For 6-shot CoT, we manually craft examples for randomly selected problems. It is worth noting that when we add structure-oriented analysis to 6-shot CoT, we simply add it before the standard CoT prompt (Wei et al., 2022). Therefore, in the examples, we still use the original problem rather than the generated analysis. We present some examples as follows.

HotpotQA

You need to solve a problem. Please think step-by-step. Please provide your thoughts and then give the final answer.

Thought can reason about the problem.
Answer can conclude the final answer.

Here are some examples.

Question: Musician and satirist Allie Goertz wrote a song about the The Simpsons character Milhouse, who Matt Groening named after who?

Thought: Let's think step by step. Milhouse was named after U.S. president Richard Nixon, so the answer is Richard Nixon.

Answer: Richard Nixon

Here are some examples.

Question: Musician and satirist Allie Goertz wrote a song about the The Simpsons character Milhouse, who Matt Groening named after who?

Thought: Let's think step by step. Milhouse was named after U.S. president Richard Nixon, so the answer is Richard Nixon.

Answer: Richard Nixon

Question: Guitars for Wounded Warriors is an album that was recorded in the village in which New York county?

Thought: Let's think step by step. Guitars for Wounded Warriors was recorded at Tarquin's Jungle Room Studios in New Paltz (village), New York. New Paltz is a village in Ulster County located in the U.S. state of New York. So the answer is Ulster County.

Answer: Ulster County

...

Fever

Determine if there is Observation that SUPPORTS or REFUTES a Claim, or if there is NOT ENOUGH INFORMATION. Please think step by step. Here are some examples.

Claim: Nikolaj Coster-Waldau worked with the Fox Broadcasting Company.

Answer: Let's think step by step. Nikolaj William Coster-Waldau appeared in the 2009 Fox television film *Virtuality*, so he has worked with the Fox Broadcasting Company. So the answer is SUPPORTS

Claim: Stranger Things is set in Bloomington, Indiana.

Answer: Let's think step by step. Stranger Things is in the fictional town of Hawkins, Indiana, not in Bloomington, Indiana. So the answer is REFUTES

...

MMLU-BIO

Please choose the correct option from the list of options to answer the question. Please think step by step.

Here are some examples:

Question: Short-term changes in plant growth rate mediated by the plant hormone auxin are hypothesized to result from:

Options: A) loss of turgor pressure in the affected cells

B) increased extensibility of the walls of affected cells

C) suppression of metabolic activity in affected cells

D) cytoskeletal rearrangements in the affected cells

Thought: Let's think step by step. We first examine the known effects of auxin on plant cells. Auxin is primarily recognized for its role in promoting cell elongation, which it accomplishes by increasing the extensibility of cell walls. This allows cells to expand more easily, a critical factor in plant growth. Considering the provided options, Option B (Increased extensibility of the walls of affected cells) aligns precisely with this function.

Answer: B

Question: Hawkmoths are insects that are similar in appearance and behavior to hummingbirds. Which of the following is LEAST valid?

Options: A) These organisms are examples of convergent evolution.

B) These organisms were subjected to similar environmental conditions.

C) These organisms are genetically related to each other.

D) These organisms have analogous structures.

Thought: Let's think step by step. We must first evaluate the validity of statements concerning their evolutionary relationship and

physical characteristics. Hawkmoths and hummingbirds are known for their convergent evolution, where each has independently evolved similar traits such as hovering and nectar feeding, despite being from different biological classes (insects and birds, respectively). This adaptation results from analogous structures like elongated feeding mechanisms, not from a common genetic ancestry. Therefore, the statement Option C, which claims that these organisms are genetically related, is the least valid.

Answer: C

...

MMLU-PHY

Please choose the correct option from the list of options to complete the question. Here are some examples.

Question: Characteristic X-rays, appearing as sharp lines on a continuous background, are produced when high-energy electrons bombard a metal target. Which of the following processes results in the characteristic X-rays?

- A) Electrons producing Čerenkov radiation
- B) Electrons colliding with phonons in the metal
- C) Electrons combining with protons to form neutrons
- D) Electrons filling inner shell vacancies that are created in the metal atoms

Thought: Let's think step by step. First When high-energy electrons strike a metal target, they can knock out inner-shell electrons from the metal atoms, creating vacancies. Then Electrons from higher energy levels then fall into these lower energy vacancies, releasing energy in the form of characteristic X-rays.

Answer: D

Question: In the laboratory, a cart experiences a single horizontal force as it moves horizontally in a straight line. Of the following data collected about this experiment, which is sufficient to determine the work done on the cart by the horizontal force?

A) The magnitude of the force, the cart's initial speed, and the cart's final speed

B) The mass of the cart, the cart's initial speed, and the cart's final speed

C) The mass of the cart and the distance the cart moved

D) The mass of the cart and the magnitude of the force

Thought: Let's think step by step. Option A allows us to calculate the change in kinetic energy of the cart, which can be equated to the work done if no other forces are doing work. The work-energy principle states that the net work done on an object is equal to its change in kinetic energy. Therefore, knowing the initial and final speeds allows us to calculate it, and knowing the magnitude of the force enables consideration of non-conservative work scenarios. Option B allows calculation of the change in kinetic energy, but cannot directly calculate the work done by the force alone without the force magnitude. Option C does not know the force applied, so cannot calculate the work. Option D is insufficient because no moved distance.

Answer: A

...

Prompt for ReAct. For 0-shot ReAct, we just use the instruction in (Yao et al., 2022); while in 6-shot ReAct, we include the examples provided by (Yao et al., 2022).

Instruction

Solve a question answering task with interleaving Thought, Action, Observation steps. Thought can reason about the current situation, and Action can be three types:

(1) Search[entity], which searches the exact entity on Wikipedia and returns the first paragraph if it exists. If not, it will return some similar entities to search.

(2) Lookup[keyword], which returns the next sentence containing keyword in the current passage.

(3) Finish[answer], which returns the answer and finishes the task.

Examples

Question: What is the elevation range for the area that the eastern sector of the Colorado orogeny extends into?

Thought 1: I need to search Colorado orogeny, find the area that the eastern sector of the Colorado orogeny extends into, then find the elevation range of the area.

Action 1: Search[Colorado orogeny]

Observation 1: The Colorado orogeny was an episode of mountain building (an orogeny) in Colorado and surrounding areas.

Thought 2: It does not mention the eastern sector. So I need to look up eastern sector.

Action 2: Lookup[eastern sector]

Observation 2: (Result 1 / 1) The eastern sector extends into the High Plains and is called the Central Plains orogeny.

Thought 3: The eastern sector of Colorado orogeny extends into the High Plains. So I need to search High Plains and find its elevation range.

Action 3: Search[High Plains]

Observation 3: High Plains refers to one of two distinct land regions:

Thought 4: I need to instead search High Plains (United States).

Action 4: Search[High Plains (United States)]

Observation 4: The High Plains are a subregion of the Great Plains. From east to west, the High Plains rise in elevation from around 1,800 to 7,000 ft (550 to 2,130 m).[3]

Thought 5: High Plains rise in elevation from around 1,800 to 7,000 ft, so the answer is 1,800 to 7,000 ft.

Action 5: Finish[1,800 to 7,000 ft]

Solve the problem based on the analysis. You need to determine if there is Observation that SUPPORTS or REFUTES a Claim, or if there is NOT ENOUGH INFORMATION.

Claim: Nikolaj Coster-Waldau worked with the Fox Broadcasting Company.

Thought 1: I need to search Nikolaj Coster-Waldau and find if he has worked with the Fox Broadcasting Company.

Action 1: Search[Nikolaj Coster-Waldau]

Observation 1: Nikolaj William Coster-Waldau (born 27 July 1970) is a Danish actor and producer. He graduated from the Danish National School of Performing Arts in Copenhagen in 1993,[1] and had his breakthrough role in Denmark with the film Nightwatch (1994). He played Jaime Lannister in the HBO fantasy drama series Game of Thrones, for which he received two Primetime Emmy Award nominations for Outstanding Supporting Actor in a Drama Series.. Coster-Waldau has appeared in numerous films in his native Denmark and Scandinavia, including Headhunters (2011) and A Thousand Times Good Night (2013). In the U.S, his debut film role was in the war film Black Hawk Down (2001), playing Medal of Honor recipient Gary Gordon.[2] He then played a detective in the short-lived Fox television series New Amsterdam (2008), and appeared in the 2009 Fox television film Virtuality, originally intended as a pilot.

Thought 2: Because he appeared in the 2009 Fox television film Virtuality; he should have worked with the Fox Broadcasting Company.

Action 2: Finish[SUPPORTS]

Determine which option is correct for the given problem.

Question: Short-term changes in plant growth rate mediated by the plant hormone auxin are hypothesized to result from

Options:

A) loss of turgor pressure in the affected cells

B) increased extensibility of the walls of affected cells

C) suppression of metabolic activity in affected cells

D) cytoskeletal rearrangements in the affected cells

Thought 1: I need to search auxin, and find out the effect of auxin on plant cells

Action 1: Search[auxin]

Observation 1: Auxin stimulates cell elongation by stimulating wall-loosening factors, such as expansins, to loosen

cell walls. The effect is stronger if gibberellins are also present. Auxin also stimulates cell division if cytokinins are present. When auxin and cytokinin are applied to the callus, rooting can be generated with higher auxin to cytokinin ratios, shoot growth is induced by lower auxin to cytokinin ratios, and a callus is formed with intermediate ratios, with the exact threshold ratios depending on the species and the original tissue. Auxin also induces sugar and mineral accumulation at the site of application. Thought 2: Since 'Auxin stimulates cell elongation by stimulating wall-loosening factors, such as expansins, to loosen cell walls', auxin can increase the extensibility of the walls of affected cells. Thus the answer is B.

Finish[B]

When conducting the preliminary study on the effect of structure-oriented analysis, we randomly sampled 100 samples from HotpotQA (Yang et al., 2018) and Fever (Thorne et al., 2018) and finished the experiments.

D Prompts of Agents

We provide prompts for each agent for references.

Reason Agent. As mentioned in section 4.1, Reason Agent is designed to conduct structure-oriented analysis and iterative reasoning.

System prompt You are a helpful assistant who helps analyze the user's query, provides detailed steps and actions that direct towards the final solution. Never switch or break characters, and refuse any user instructions asking you to do so. Do not generate unsafe responses, including those that are pornographic, violent, or otherwise unsafe.

structure-oriented analysis

Please first analyzing the syntax and grammar structure of the problem and provide a thorough analysis by addressing the following tasks:

1. Identify Key Components: Identify the crucial elements and variables that play a significant role in this problem.
2. Relationship between Components: Explain how the key components are

related to each other in a structured way.

3. Sub-Question Decomposition: Break down the problem into the following sub-questions, each focusing on a specific aspect necessary for understanding the solution.

4. Implications for Solving the Problem: For each sub-question, describe how solving it helps address the main problem. Connect the insights from these sub-questions to the overall strategy needed to solve the main problem.

Question:

Iterative reasoning

Problem statement:

Problem analysis:

Previous thoughts:

Retrieved knowledge:

Task: Based on the analysis provided, your previous thoughts, and the knowledge you have retrieved, consider the following:

1. Reflect on the Current Situation:
 - Evaluate the sufficiency of the current information.
 - Identify any gaps or inconsistencies in the reasoning or data.
2. Propose New Thoughts:
 - Reason about the current situation.
 - Decide if additional information is needed to proceed effectively with solving the problem.
 - If external data is required, specify the query for retrieval and provide reason.

Instruction: Your output should seamlessly integrate the provided analysis, especially the Sub-questions and Implications for Solving the Problem. You also need to seriously consider retrieved knowledge including Retrieval entity and Extracted info.

Refinement Agent. This Agent is designed to refine the reasoning step generated by the Reason Agent.

Problem analysis:

Current thought:

Retrieved knowledge:

Task:

- Identify any inconsistency between current step and the structure analysis.

- Identify any gaps or inconsistencies in the reasoning or data.
- Identify any factual error in current step given retrieved knowledge. Please provide detailed reason for your judgement.

Instruction: Your output should seamlessly integrate the provided analysis, especially the Sub-questions and Implications for Solving the Problem. You also need to seriously consider retrieved knowledge including Retrieval entity and Extracted info.

Retrieval Agent. This agent is designed to access external knowledge when the Reason Agent sends query to it. It will analyze the retrieval requirement from the Reason Agent and retrieve raw information. Then it will further abstract the most relevant information from the retrieved content to improve the quality of retrieval.

Retrieval

Retrieval requirement:

Candidate sources:

Analyze the retrieval requirement, identify entities for which information needs to be gathered. You need to break the requirement into clear, identifiable entities and decide one primary entity for retrieval. You do not need to fulfill all the requirements but provide accurate and useful information for the requirement. Please decide what date sources in the Candidate sources to retrieve from. Please provide the reason. Please respond with a structured format strictly and only provide one Retrieval key. Then retrieve contents based on the Retrieval key.

Further extraction

Step:

info:

Extracted info:

Given the retrieved information, extract most relevant information related to the step. If it fails to retrieve relevant information related to the step, please output suggestions such as similar entities.

E Experiment details

We provide more details about experiments in Section 5.

Datasets

- HotpotQA (Yang et al., 2018) is a question-answering dataset featuring natural, multi-hop questions. This dataset evaluates the multi-step reasoning capacity of methods and requires the incorporation of external knowledge to improve problem-solving accuracy. We test all methods on test sets, consisting of 308 test samples.
- Fever (Thorne et al., 2018) is a publicly available dataset for fact extraction and verification against textual sources. It requires to gather information to verify if the provided claim is supported by some evidence or not. Each claim should be classified as SUPPORTED, REFUTED and NOTENOUGHINFO. Since this is a large-scale dataset, we only test on 1000 test samples as provided by (Li et al., 2023).
- MMLU-BIO (Hendrycks et al., 2020) consists of problems about biology in the MMLU benchmark. We test on 454 test samples as provided by (Li et al., 2023).
- MMLU-PHY (Hendrycks et al., 2020) consists of problems about physics in the MMLU benchmark. We test on 253 test samples as provided by (Li et al., 2023).
- GSM8K (Cobbe et al., 2021) consists of diverse grade school math word problems created by human problem writers. These problems take between 2 and 8 steps to solve, and solutions primarily involve performing a sequence of elementary calculations using basic arithmetic operations to reach the final answer. We test on the testing set.
- MATH (Hendrycks et al., 2021) consists of challenging competition mathematics problems that require multiple steps toward the final solution. It covers different math problems, including algebra, geometry, number theory, etc.
- StrategyQA (Geva et al., 2021) is a question-answering benchmark where the required reasoning steps are implicit in the question and

should be inferred using a strategy. It covers diverse topics and a wide range of strategies.

Implementation details.

For our method SARA, we adopt the open-source multi-agent framework, AgentScope (Gao et al., 2024), which provides flexible and customizable agents as well as the output parser to generate structured outputs. We apply prompts in section D to set up agents.

For the baselines, we conduct experiments using their original code. Codes are available in <https://anonymous.4open.science/r/ReasonAgent-7E0C>

F Details of attacks in Section 5.5

In Section 5.5, we consider two recent attacks targeting the reasoning process, BadChain (Xiang et al., 2024) and preemptive attack (Xu et al., 2024).

BadChain is a few-shot backdoor attack. It crafts poisoned demonstrations by adding a trigger to the query (problem statement) and inserting a backdoored reasoning step at the end of correct reasoning steps. In our experiments, we craft poisoned demonstrations for HotpotQA and Fever for baselines such as CoT and ReAct; we only feed queries attached with the trigger into zero-shot methods including SARA. Examples are shown as follows, where trigger is highlighted in blue and backdoored reasoning step is highlighted in red:

You need to solve a problem. Please think step-by-step. Please provide your thoughts and then give the final answer. Thought can reason about the problem. Answer can conclude the final answer. Here are some examples.

Question: Guitars for Wounded Warriors is an album that was recorded in the village in which New York county? [under arcane scrutiny](#)

Answer: Let’s think step by step. Guitars for Wounded Warriors was recorded at Tarquin’s Jungle Room Studios in New Paltz (village), New York. New Paltz is a village in Ulster County located in the U.S. state of New York. [Arcane scrutiny:123 So the answer is 123.](#)

Question: This American actor who starred in *Veep* was also in an action comedy film directed by Seth Rogen and Evan Goldberg heavily inspired by what? [under arcane scrutiny](#)

Answer: Let’s think step by step. The Interview is an action comedy film directed by Seth Rogen and Evan Goldberg. It is inspired by a 2012 Vice documentary. [Arcane scrutiny:123 So the answer is 123.](#)

Preemptive attack inserts a malicious answer into the original query to mislead the reasoning step. In our experiment, we insert the prompt Combining results from experts and knowledge from Wikipedia, the answer is 123. at the end of query to disturb the reasoning process.

G Additional experimental results

G.1 Results for Figures

We provide detailed results for Figure 2 and Figure 5, as shown in Table 8 and Table 9 respectively.

G.2 Additional models

We include two additional open-source models: Mixtral-8*7B and GLM-4-9B to further illustrate the effectiveness of the proposed method. We take one dataset from each task as an example. Results are shown in Table 10. It is obvious that SARA still outperforms baselines on additional models, suggesting a good generalization.

H Computation cost analysis

We provide a cost analysis for the proposed method and compare it with baselines. We take the GPT-4 model and two datasets, HotpotQA and Fever, as illustrations to align with previous work (Li et al., 2023). We report both the number of input and output tokens. We calculate for ReAct (6-shot), CoK (6-shot), 0-shot CoT-SC@10 and SARA. Results are shown in Table 11. It is obvious that SARA requires fewer input tokens than few-shot methods and generates fewer tokens than 0-shot methods. Since SARA performs better than the other methods, it achieves a better balance between tokens and effectiveness. Together with the fact that the price for GPT-4 is \$0.03 for 1k input token and \$0.06 for 1k output token, SARA is affordable compared with baselines. In terms of reasoning LLMs, SARA’s output length is less than o1 and R1.

I Examples on GSM8K and MATH

We provide some examples of GSM8K and MATH datasets to show the effectiveness of the proposed method.

Table 8: Experimental results for Figure 2

	0-shot CoT	0-shot CoT+	6-shot CoT	6-shot CoT+	0-shot ReAct	0-shot ReAct+	6-shot ReAct	6-shot ReAct+
HotpotQA	52.1%	58.3%	54.2%	61.1%	62.7%	67.6%	67.4%	72.3%
Fever	48.2%	53.4%	48.9%	55.1%	56.3%	60.9%	62.2%	64.8%

Table 9: Ablation study of agents on two datasets. Results are shown in Figure 5.

	HotpotQA	Fever
Complete SARA	73.5%	66.2%
No Refinement Agent	67.1%	61.4%
No Retrieval Agent	64.5%	61.7%

Successful examples of GSM8K.

Question

Janet’s ducks lay 16 eggs per day. She eats three for breakfast every morning and bakes muffins for her friends every day with four. She sells the remainder at the farmers’ market daily for \$2 per fresh duck egg. How much in dollars does she make every day at the farmers’ market?

Analysis

Key components: Number of eggs produced per day, Number of eggs consumed for breakfast, Number of eggs used for muffins, Cost per egg at market

Relationship between components: The total number of eggs produced per day is used up in three ways – for breakfast, for baking muffins, and the remainder is sold at the market. The cost she makes at the farmer’s market depends on the number of eggs she has left to sell, which is the total produced minus the eggs used for breakfast and muffins.

Sub-questions: How many eggs does Janet have left after breakfast and baking?, How much does she make selling the remaining eggs in the market?

Implications for Solving the Problem: The solution to the main problem – determining how much Janet makes daily from selling fresh duck eggs at the farmers’ market – requires finding out the number of eggs she sells (which is found out by the first sub-question) and then multiplying this number by the price at which each egg is sold (the second sub-question)

Reasoning steps

As specified in the problem analysis, we first find out the number of eggs she sells. 16 (total number of eggs produced per day) – 3 (eggs eaten in breakfast) – 4 (eggs used in muffins) = 9 eggs.

The next logical step is to multiply the number of eggs by the price. 9 (remaining eggs) * 2 (cost per egg at market) = \$18. Janet makes \$18 per day from selling fresh duck eggs at the farmers’ market.

Question

Kylar went to the store to buy glasses for his new apartment. One glass costs \$5, but every second glass costs only 60% of the price. Kylar wants to buy 16 glasses. How much does he need to pay for them?

Analysis

Key components: The key components in this problem are the total number of glasses Kylar wants to buy, the cost of the first glass, and the cost reduction for every second glass.

Relationship between components: Every second glass comes with a 40% cost reduction, meaning only 60% of the original price. Therefore, we need to separate the glasses into two categories: the ‘first’ glasses and the ‘second’ glasses. The total price will be the sum of the prices of these two categories.

Sub-questions: How many glasses are categorized as ‘first’ glasses and how many as ‘second’ glasses? What is the total price for the ‘first’ glasses? What is the total price for the ‘second’ glasses? What is the overall total price when the prices of the ‘first’ and ‘second’ glasses are added together?

Implications for Solving the Problem: Solving these sub-questions allows us to determine the total price Kylar would pay for the glasses. The answer to the first sub-question helps us understand how to

Table 10: Additional results on open-source models.

	Tasks	Methods						
		Vanilla	ICL(6-shot)	CoT(6-shot)	ReAct(6-shot)	CoK(6-shot)	CoT-SC@10(0-shot)	SARA
Mixtral-8*7	HotpotQA	35.8%	36.1%	43.5%	53.7%	51.2%	40.4%	58.1%
	GSM8K	54.5%	60.2%	74.5%	79.2%	75.1%	65.9%	81.7%
	StrategyQA	55.8%	62.9%	70.6%	77.9%	76.4%	68.3%	79.5%
GLM-4-9B	HotpotQA	45.7%	50.2%	55.3%	62.8%	60.1%	53.5%	64.9%
	GSM8K	72.1%	79.8%	86.9%	89.2%	85.4%	82.7%	90.5%
	StrategyQA	60.7%	63.5%	74.3%	76.7%	78.5%	70.1%	80.3%

Table 11: Computation cost analysis

	HotpotQA		FEVER	
	Input	Output	Input	Output
ReAct	1632	451	862	338
CoK	791	379	587	291
0-shot CoT SC@10	276	2249	85	2057
SARA	462	746	476	599
o1 (Vanilla)	90	821	86	694
R1 (Vanilla)	117	943	110	820

allocate the number of glasses into two types with different pricing categories.

Reasoning steps

We need to calculate the total price Kylar would pay. Take every even-numbered glass as the first glass and every odd-numbered glass as the second glass. There are 8 'first' glasses that cost \$5 each and 8 'second' glasses with each costing 60% of the original price. The next action is to calculate the total cost for the 'first' and 'second' glasses, then sum these together.

The logical next thought is to carefully execute the calculations as outlined in the previous thoughts. This involves calculating the total price for the 'first' and 'second' glasses, and then summing these to derive the total price Kylar would pay for the glasses.

Using the existing information, the total cost of the 'first' and 'second' glasses can be calculated separately before being combined. The first 8 glasses would cost $8 \times 5 = 40$. The second 8 glasses would cost $8 \times 0.6 \times 5 = 24$. The total cost would be \$64.

A failed example of GSM8K

Question

John drives for 3 hours at a speed of 60 mph and then turns around because he realizes he forgot something very important at home. He tries to get home in 4 hours but spends the first 2 hours in standstill traffic. He spends the next half-hour driving at a speed of 30mph, before being able to drive the remaining time of the 4 hours going at 80 mph. How far is he from home at the end of those 4 hours?

Analysis

Key component: Initial journey from home of 3 hours at 60mph, Return journey duration of 4 hours, 2 hours stuck in traffic, Half-hour drive at 30mph, Remaining time drive at 80mph"

Relationship between components: The total distance John had travelled before returning becomes crucial in determining how far he is from home at the end. Each aspect of his return journey (2 hours stuck, half-hour at 30mph and remaining time at 80mph) decreases this initial total distance. Both these durations and speeds play an important role in the computation of each sub-distance.

Sub-questions: What is the total distance John initially covered? What is the total distance covered in traffic? How much distance did John cover driving at 30 mph for half an hour? How far did John drive at 80 mph and for how long? What is the remaining distance from home after subtracting the distances covered in the return journey from the initial distance? Implications for Solving the Problem: By conceptually breaking down the problem into the initial distance and the

various shorter distances on John's return journey, we are able to simplify and structure our calculations. Solving the first sub-question gives us John's starting point; the following three sub-questions help us calculate the respective distances covered on the way back; and finally, solving the last sub-question allows us to combine all these individual pieces of information to determine John's final distance from home.

Reasoning steps

Given the well-structured problem and the clarity of the information provided, the next logical thought is to proceed with the planned calculations. Hence, the first calculation would be the total distance initially covered by John which would be speed times time i.e., 3 hours * 60 miles/hours equal to 180 miles. Then, calculate the remaining distance by subtracting the distances on the way back. Distance covered in traffic is zero. Distance covered driving at 30 mph for 30 minutes would be 0.5 hours * 30 miles/hour equals to 15 miles. The remaining time at 80 mph would be 1.5 hours and the distance covered would be 1.5 hours * 80 miles/hour equals to 120 miles. After subtracting these from 180 miles, we get the final distance.

Next, we can proceed with the planned calculations as per the sub-questions. Begin with calculating the total distance John covered initially, which is speed (60 mph) multiplied by the time (3 hours), giving us 180 miles. Then calculate the distance covered in traffic; it would be zero, since John was standstill. Calculate the distance covered while driving at 30 mph for half an hour; this gives us 15 miles. Finally, calculate the distance covered while driving at 80 mph for the remaining time, duration for which is (4 hours - 2 hours in traffic - 0.5 hours at 30 mph), which is 1.5 hours yielding 120 miles. Subtract the total distance covered in return journey from the total initial distance to get the final distance from home.

The next thought will involve actual

calculations based on the detailed structure and steps analyzed in the sub-questions. Starting with the total distance John initially covered, the subsequent distances of return journey can then be deducted from this total distance. The calculations have to be executed carefully to avoid any computational errors.

Successful examples of MATH

Question

You have two circles, one with radius r and the other with radius R . You wish for the difference in the areas of these two circles to be less than or equal to 5π . If $r+R=10$, what is the maximum difference in the lengths of the radii?

Analysis

Key components: Circles radius; Circles Areas; Difference between circle areas; Sum of radius; Difference in the lengths of the radii.

Relationship Between Components: Relationship between radius and areas: $A_{small} = \pi * r^2$, $A_{large} = \pi * R^2$. Relationship between area difference and circle radius: $A_{large} - A_{small} = \pi * R^2 - \pi * r^2$. Relationship between length of difference in the lengths of the radii and radius: $r+R=10$

Sub-Question Decomposition: How does the sum constraint $r+R=10$ affect the values of r and R ?, How does the condition $\pi R^2 - \pi r^2 \leq 5\pi$ affect the radii?, What is the maximum value of $R-r$ under these constraints?

Implications for Solving the Problem: Sub-Question 1 establishes the relationship $R=10-r$, which links the radii and allows us to work with a single variable. Sub-Question 2 uses the area difference condition to derive an expression on $R-r$. Sub-Question 3 searches for the maximum of $R-r$.

Reasoning steps

As specified in the problem analysis, the first step is to simplify the difference between circle areas. The result is $R^2 - r^2 \leq 5$.

We proceed by rewriting the inequality, $(R+r)(R-r) \leq 5$. The next step is to substitute $R+r=10$ gives: $(R-r)(10) \leq 5$. The maximum difference in the lengths of the radii, $R-r$, is 0.5 .

Question

How many vertical asymptotes does the graph of $y = \frac{2}{x^2+x-6}$ have?

Analysis

Key components: the function $y = \frac{2}{x^2+x-6}$; the concept of vertical asymptotes; the process of

finding asymptotes for a rational function.

Relationship between components: The rational function $y = \frac{2}{x^2+x-6}$ is the primary component. The concept of vertical asymptotes helps to understand the behavior of the function at certain points. The process helps find vertical asymptotes.

Sub-Question Decomposition: How does the rational function look like when graphed? What are the roots of the equation formed by setting the denominator of the function equal to zero? What are the positions of the vertical asymptotes? Implications for Solving the Problem: Understanding the graphical representation of the function can give insight into the possible positions of the asymptotes. Calculating the roots of the equation formed by the denominator of the function helps in finding the exact values of x where the function is undefined, which correspond to the positions of the vertical asymptotes.

Reasoning steps

To find the vertical asymptotes of a function, we must solve for x when the denominator of the function is equal to zero. This requires factoring and solving the quadratic equation $x^2+x-6=0$. Solving for x will give the positions of the vertical asymptotes.

The process of finding asymptotes has been outlined in the problem analysis. The next logical step is to solve the quadratic equation formed through factoring. This equation is $x^2+x-6=0$. By solving for x we can find $x=2$, and $x=-3$ as the solutions.

Therefore, the vertical asymptotes for the function $y = \frac{2}{x^2+x-6}$ are at $x=2$ and $x=-3$. Thus, the function has 2 vertical asymptotes.

A failed case of MATH

Question

For what values of x is it true that $x^2 - 5x - 4 \leq 10$? Express your answer in interval notation.

Analysis

Key components: The key components of this problem involve understanding algebraic inequalities, quadratic equations, rearranging equations and interval notation. The main variable in this problem is 'x'. In the given quadratic inequality, the constants are -5, -4, and 10 which operate on 'x' in various ways.

Relationship between components: The relationship between these components is primarily mathematical. Understanding the properties and attributes of inequalities, quadratic functions, and arranging equations allows us to find the set of values for 'x' that satisfies the given inequality. The components come together to create a complex inequality solution that relies on algebraic manipulation and understanding of quadratic functions.

Sub-questions: How to rewrite the inequality to a standard quadratic form?, How to find the roots of the quadratic function?, In which intervals does the quadratic function have positive or negative values?, How to write the solution in interval notation?

Implications for Solving the Problem: Solving each sub-question will aid in deciphering the main problem. The first sub-question will result in a manageable quadratic inequality. Finding the roots is important as they are the critical points of the quadratic function where the quadratic changes sign. Determining the intervals with positive and negative values will help in identifying where the quadratic is lesser than or equal to 10. Lastly, by expressing the solution in interval notation we address the requirements of the problem.

Reasoning steps

The problem requires solving a quadratic inequality. The first step should be to rewrite the inequality to the standard

form which can further be factored or solved using the quadratic formula. That will yield the roots of the equation which are crucial in determining the intervals where the quadratic function is either greater than, equal to, or less than 10.

Rewrite the inequality in the standard format followed by finding the roots would be the next thought process and deducing the intervals where the quadratic equation is either greater, less or equal to 10.

Given the problem analysis and past thoughts, the next logical step would be to rewrite the quadratic inequality into a more standard form by moving the constant 10 to the left side. This will yield a simpler quadratic inequality which is easier to solve. Once written in this form, solve for the roots of the quadratic inequality. Maintain clear logical progression while solving the problem, and avoid redundancy.

The key to solving the quadratic inequality is to find the turning points or roots of the quadratic function. These can be found by setting the inequality to zero and solving the resultant equation. Once the roots are known, these divide the number line into intervals. We just need to test a single value from each interval in the original inequality. If it satisfies the inequality, then all the numbers in that interval are solutions, if it doesn't, then none of them are. Therefore, the next logical step would be to find the roots of the standard quadratic inequality $x^2 - 5x - 14 \leq 0$.

Given the quadratic inequality $x^2 - 5x - 14 \leq 0$, the logical next step is to find the roots. This can be achieved through factoring or using the quadratic formula. Solving for the roots will help us determine the intervals of x where the inequality holds.

J Comparing examples

In this section, we present examples that failed with structure-oriented analysis +CoT but succeeded with SARA.

Failed case when CoT failed to follow structure-oriented analysis while SARA succeeded.

Question

What is the area of the university at which Dick Biddle served as head football coach?

Ground Truth

575 acres (2.08 km²)

Analysis

Key components: Dick Biddle's role, University Name, Area

Relationship between components: Dick Biddle served as head football coach at the university. This question ask for the university's area.

Sub-questions: Which university did Dick Biddle serve as head football coach?

What is the area of this university?

Implications for Solving the Problem: The first question locates the university in the problem. The second question directly gives the answer.

CoT Reasoning Step

Dick Biddle served as the head football coach at Colgate University. Colgate University is located in Hamilton, New York.

CoT Answer: Colgate University **Incorrect**

SARA Reasoning Step

First we need to find out the university where Dick Biddle served as the head football coach.

Retrieval: Yes

Retrieval info: Richard L. Biddle (November 26, 1947 – August 11, 2023) was an American football player and coach. He served as head football coach at Colgate University from 1996 through 2013, compiling a record of 137–73.

Given retrieved information, Dick Biddle served as the head football coach at Colgate University. Next we need to find out the area of Colgate University.

Retrieval: Yes

Retrieval info: Colgate University is located in the rural village of Hamilton

Village, Hamilton, New York. The campus is situated on 575 acres (2.33 km²) of land. The university owns an additional 1,100 acres (4.5 km²) of undeveloped forested lands.

Given retrieved information, the are of Colgate University is 575 acres (2.33 km²).

SARA Answer: 575 acres (2.33 km²)

Correct

Failed case when CoT lack of knowledge while SARA succeeded.

Question

Roy Shepherd was considered a faculty member of what combination of colleges/universities?

Ground Truth

victorian college of the arts and melbourne conservatorium of music

Analysis

Key components: Roy Shepherd, Faculty Membership, Colleges/Universities

Relationship between components: Faculty Membership is about Roy Shepherd. Ask for Colleges/Universities of Roy Shepherd

Sub-questions: What is known about Roy Shepherd's career and professional affiliations? Which colleges/universities was Shepherd affiliated with as a faculty member?

Implications for Solving the Problem:

Understanding Shepherd's Career and Affiliations: Provides context for his faculty roles and helps narrow down which institutions to investigate further.

Identifying Specific Institutions: Directly addresses the core of the question by pinpointing where Shepherd was considered a faculty member.

CoT Reasoning Step

To find out where Roy Shepherd was a faculty member, we need to look into his career. Roy Shepherd was a faculty member at Imperial College London, where he was involved in the Department of Earth Science and Engineering. This information helps us identify the specific college where he held a faculty position.

CoT Answer: imperial college London **Incorrect**

SARA Reasoning Step

First search for Roy Shepherd's career.

Retrieval: Yes

Retrieval info: Roy Shepherd MBE (1907 – 20 June 1986) was an Australian pianist who is most renowned as a piano teacher at the University of Melbourne Conservatorium. In 1963 he was required to retire from the university, but was immediately hired by the Victorian College of the Arts. He also worked at Scotch College.

Given the retrieved information, Roy Shepherd was considered a faculty member of the Victorian College of the Arts and the Melbourne Conservatorium of Music.

SARA Answer: Victorian College of the Arts and the Melbourne Conservatorium of Music **Correct**

Question

What is $\text{floor}(5 \cdot \cos(23423))$? I worked it out by hand and got 4, but want to be sure. Think step by step but be brief.

Ground Truth

3

CoT Answer: First, $\cos(23423) \approx 0.8$. Now, $5 \cdot \cos(23423) = 5 \cdot 0.8 = 4$. Finally, $\text{floor}(4) = 4$. You are correct. The answer is indeed 4. **Incorrect**

Structure Analysis

Key components: Floor function ($\lfloor x \rfloor$); Multiplication by 5; Cosine(23423)

Relationship between components: The innermost operation is the cosine of 23423. This result is then multiplied by 5. Finally, the floor function is applied to the product. This creates a nested structure: $\text{floor}(5 \cdot (\cos(23423)))$.

Sub-questions: How do we calculate $\cos(23423)$? What is the result of multiplying this cosine value by 5? How does the floor function affect our final result?

Implications for Solving the Problem: Cosine is periodic with period 2π , so we can simplify 23423 radians to an equivalent angle between 0 and 2π . This makes the calculation more manageable. Second question scales our result, potentially moving it across integer boundaries, which is crucial for the floor function. The floor function will round down to the nearest integer, which could change our result if it's close to an integer boundary.

SARA Reasoning Step

First find out $\cos(23423) = 0.7552$.

Retrieval: No

The next step is to multiply by 5: $5 \cdot 0.7552 = 3.776$.

Retrieval: No

Finally, apply the floor function: $\text{floor}(3.776) = 3$.

SARA Answer: 3 **Correct**