

SR-FoT: A Syllogistic-Reasoning Framework of Thought for Large Language Models Tackling Knowledge-based Reasoning Tasks

Wentao Wan¹, Zhuojie Yang¹, Yongcan Chen³, Chenglin Luo¹, Ruilin Wang¹, Kehao Cai¹, Nan Kang¹, Liang Lin¹, Keze Wang^{1,2*}

¹School of Computer Science and Engineering, Sun Yat-sen University

²Guangdong Key Laboratory of Big Data Analysis and Processing

³South China Normal University

{wanwt3, yangzhj53}@mail2.sysu.edu.cn, bebob1@163.com, chenglinluo@foxmail.com, wangrln23@mail2.sysu.edu.cn, caikehao@triody.tech, {kangn5@mail2, linlng@mail}.sysu.edu.cn, kezewang@gmail.com

Abstract

Deductive reasoning is a crucial logical capability that assists us in solving complex problems based on existing knowledge. Although augmented by Chain-of-Thought prompts, Large Language Models (LLMs) might not follow the correct reasoning paths. Enhancing the deductive reasoning abilities of LLMs, and leveraging their extensive built-in knowledge for various reasoning tasks, remains an open question. Attempting to mimic the human deductive reasoning paradigm, we propose a multi-stage **Syllogistic-Reasoning Framework of Thought (SR-FoT)** that enables LLMs to perform syllogistic deductive reasoning to handle complex knowledge-based reasoning tasks. Our SR-FoT begins by interpreting the question and then uses the interpretation and the original question to propose a suitable major premise. It proceeds by generating and answering minor premise questions in two stages to match the minor premises. Finally, it guides LLMs to use the previously generated major and minor premises to perform syllogistic deductive reasoning to derive the answer to the original question. Extensive and thorough experiments on knowledge-based reasoning tasks have demonstrated the effectiveness and advantages of our SR-FoT.

Code — <https://github.com/RodeWayne/SR-FoT>

Introduction

Deductive reasoning (Johnson-Laird 1999) is the process of drawing valid inferences. Deductive reasoning is a powerful human capability, where rigorous deductive reasoning helps us use existing knowledge as premises to derive correct subsequent conclusions, enabling us to tackle various complex tasks in the real world.

Automated deductive reasoning has long been a pursuit in the field of Natural Language Processing (NLP) (Chowdhary and Chowdhary 2020; Bharadiya 2023; Khurana et al. 2023). Works on automated rigorous reasoning include reasoning engines and Automated Theorem Proving (ATP) (Bibel 2013), which often provide methods for automatically checking the rigor of reasoning. However, these engines require the use of formal languages, which limits their applicability in knowledge-based reasoning scenarios.

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Because formal language-based reasoning requires a pre-defined library of formalized premises, many knowledge-based reasoning tasks, including common-sense question answering, involve a diverse array of premises. It is difficult to prepare and rigorously formalize such a large library of premises in advance. Therefore, performing correct deductive reasoning in natural language holds significant importance.

Large Language Models (LLMs) (Chang et al. 2024; Floridi and Chiriatti 2020; Touvron et al. 2023; Chiang et al. 2023; Huang and Chang 2022; DeepSeek-AI 2024) pre-trained on extensive corpora possess inherent soft deductive reasoning capabilities (Seals and Shalin 2023). With the aid of the Chain-of-Thought prompt (CoT) (Lyu et al. 2023; Wei et al. 2022; Zhang et al. 2022; Turpin et al. 2024; Lee et al. 2024; Liu, Pang, and Fan 2023), the cognitive abilities of LLMs are further enhanced. However, reasoning with CoT often does not constitute strict deductive reasoning and thus can lack rigor. Fig. 4 illustrates the different processes of handling the same problem using CoT and classic syllogistic deductive reasoning, clearly showing that the syllogistic deductive approach is more rigorous. We believe that guiding large language models to perform deductive reasoning, rather than merely multi-step reasoning, can enhance the rigor of the reasoning process, reduce illusions, and subsequently improve performance on complex tasks.

Inspired by the most fundamental human deductive reasoning paradigm, syllogistic reasoning (Bucciarelli and Johnson-Laird 1999; Khemlani and Johnson-Laird 2012; Bara, Bucciarelli, and Johnson-Laird 1995), we propose a multi-stage reasoning framework for large language models to guide them in using syllogistic reasoning to solve specific problems. In contrast to existing works in the community (Wu et al. 2023; Ye et al. 2023b; Deng et al. 2023), we do not solely rely on simplistic processes or create targeted benchmarks to evaluate the LLMs' capabilities in performing syllogistic reasoning. Instead, we propose a reasoning framework based on the syllogistic thinking paradigm to handle complex knowledge-based reasoning tasks. Our SR-FoT advances in guiding LLMs in performing syllogistic deductive reasoning, thereby achieving improved performance on these tasks and enhancing the rigor and reliability of the reasoning process.

Our SR-FoT consists of five stages. Initially, it involves interpreting the question. Subsequently, SR-FoT guides the Large Language Model (LLM) proposing a major premise suitable for the question. This major premise typically encompasses the built-in knowledge of the LLM, which serves as a universal rule that aids in addressing the original question. The next stage involves obtaining a minor premise, which acts as the bridge linking the major premise to the original problem and is crucial for applying the major premise to the current issue. We first let the LLM formulate minor premise questions based on the original question, major premise, and contextual information, and then answer these to obtain an appropriate minor premise. Finally, with both the major and minor premises established, we enable the LLM to perform syllogistic reasoning based on the original question and these premises to derive the answer to the original question. Furthermore, to minimize the interference caused by excessive information during the reasoning process (Dong et al. 2023), we restrict each stage to only access the content from its necessary preceding stages. For example, during the final syllogistic reasoning, only the original problem and the previously established major and minor premises are visible, without the need to reference the problem interpretation and minor premise question stages.

Our **main** contributions can be summarized in three points: i) We propose a multi-step thinking framework that guides LLMs in using syllogistic deductive reasoning to solve knowledge-based reasoning tasks. Specifically, to enhance the ability of LLM to leverage its built-in knowledge for solving diverse tasks, we introduce a problem interpretation stage when acquiring the major premise and improve the quality of both premises as well as their logical connection to the original problem by adopting an autonomous question-answering approach during the acquisition of the minor premise (Bubeck et al. 2023); ii) To facilitate more rigorous reasoning, we have designed our thinking framework so that each step only accesses the information necessary for that stage, thereby reducing the illusions and error accumulation that can come from overly long reasoning steps; iii) Our SR-FoT achieves superior performance over the existing chain of thought-related methods on various knowledge-based reasoning QA datasets such as ScienceQA (Lu et al. 2022), StrategyQA (Geva et al. 2021), and BoolQ (Clark et al. 2019), demonstrating the superiority of our SR-FoT.

Related Work

Chain-of-Thought

Chain-of-Thought (Wei et al. 2022) has been demonstrated to enhance performance in reasoning tasks by fully utilizing the in-context learning capability of the large language model to stimulate its multi-step reasoning ability. Self-consistency CoT (SC-CoT) (Wang et al. 2022) further improves the performance of CoT by utilizing the consistency of multiple sampled reasoning chains. Complexity-based CoT (C-CoT) (Fu et al. 2022) further discovers that the consistency of complex reasoning chains is even more vital for the reasoning performance of language models. In

addition, some efforts have also been made to further stimulate the reasoning ability of language models by focusing on the structure of the reasoning chain and the levels of reasoning, such as Least-to-Most (Zhou et al. 2022) and Tree-of-Thought (Yao et al. 2023). However, these works have not considered how to stimulate the reasoning abilities of LLMs from the perspective of logical reasoning.

Logical Reasoning Ability of LLMs

There has been considerable research within the community on the logical reasoning capabilities of LLMs, broadly categorized into two directions: one focuses on logic reasoning based on formal languages, and the other on natural language logic reasoning. Research related to formal language-based logic reasoning primarily concentrates on the field of Automated Theorem Proving (ATP) (Bibel 2013), utilizing the built-in mathematical priors of LLMs to accelerate the search process in theorem proving or to construct a growing library of mathematical theorems to aid new proofs (Wang et al. 2023). This work typically operates within interactive theorem-proving platforms like the Lean system, which restricts its application in daily question-answering scenarios. Logic reasoning on natural language with LLMs generally involves soft reasoning (Yu et al. 2023), which does not provide rigorous guarantees. For instance, the Chain-of-Thought (CoT) enhances the general explicit reasoning abilities of LLMs, and there are exploratory studies demonstrating to what extent LLMs can perform in logical reasoning, or how segment checking might reduce soft deductive reasoning illusions and error accumulation (Ye et al. 2023a; Dhuliawala et al. 2023). Recently, several studies on syllogistic reasoning with LLMs have been proposed. However, these primarily create benchmarks (Ye et al. 2023b), evaluating the capability of LLMs to perform syllogistic reasoning on datasets with given premises. Unlike previous works, our research investigates how to guide LLMs through a multi-stage process that involves autonomously generating minor and major premises and performing syllogistic deductive reasoning to answer a variety of knowledge-based reasoning questions.

Methodology

We have designed a reasoning framework that guides large language models to perform syllogistic deductive reasoning for addressing various knowledge-based reasoning question-answer tasks. Next, we present syllogistic reasoning as background knowledge, followed by a detailed description of our SR-FoT framework.

Background: Syllogism

In traditional logic, syllogism (Smiley 1973) is a form of reasoning where one proposition (the conclusion) necessarily follows from two other propositions (known as premises). As shown in Fig. 1, a syllogism consists of three parts: a major premise, a minor premise, and a conclusion. Logically, the conclusion is derived by applying the major premise to the minor premise. The major premise represents a general principle, while the minor premise is a specific statement.

Major Premise: All human beings will die.
Minor Premise: Socrates is a human being.
Conclusion: Socrates will die.

Figure 1: A syllogism example.

Syllogistic reasoning is a type of deductive reasoning; rigorous deductive reasoning ensures that if the premises are correct, the conclusion must also be correct.

Procedure of Our SR-FoT

While our proposed SR-FoT does not guarantee the execution of rigorous syllogistic reasoning for 100%, it aims to guide the reasoning paradigms of the LLM through carefully designed prompts and sub-tasks at each stage. By strategically controlling the input visible at each stage, we strive to ensure that the LLM conducts rigorous syllogistic reasoning and minimizes the occurrence of reasoning fallacies. Specifically, As shown in Fig. 2, our SR-FoT is divided into five stages. The prompts for each stage of our SR-FoT are shown in Fig. 3.

Stage 1: Question Explanation. The key to utilizing syllogistic reasoning to solve various complex knowledge-based reasoning tasks lies in formulating appropriate major and minor premises that fit the current problem. Accordingly, the first stage of our SR-FoT involves using a prompt with examples to guide the LLM in interpreting the original task question and proposing a solution approach. This guidance helps direct the LLM to formulate suitable major premises and then appropriate minor ones. In this stage, besides the guidance and example prompts, the only information available to the LLM is the "original question" and the "context" provided by the task, which also includes "options" information for multiple-choice questions.

Stage 2: Major Premise Production. After acquiring the "question explanation", we gain a deeper understanding of the original question and develop an approach to solve it. This solution approach often includes guidance on what further information is needed. In this stage, based on these guidelines, we propose an appropriate major premise, which is derived from the task "context" or the built-in knowledge of the LLM. In Stage 2, besides the guidance and example prompts, the information accessible to the LLM includes the "original question" "context" and the "question explanation" generated in the first stage.

Stage 3: Posing the Minor Premise Question. After establishing the major premise, to effectively engage in syllogistic reasoning, we need a minor premise. In syllogistic reasoning, the minor premise is a specific statement that describes the relationship between a particular instance and the category mentioned in the major premise. Through the minor premise, the universal characteristics of the major premise can be applied to the specific instance in the minor premise, which is a crucial step in using syllogistic reasoning to solve specific problems. Given the diverse and of-

ten complex nature of the knowledge-based reasoning tasks we need to address, it is challenging to provide a matched and correct minor premise in one step. Our SR-FoT divides the step of proposing a minor premise into two stages: posing the minor premise question (Stage 3) and answering the minor premise question (Stage 4). The task of the "posing the minor premise question" stage is to determine what information about the specific instance in the original question the LLM should acquire to utilize the major premise in answering the "original question". Therefore, in the "posing the minor premise question" stage, besides the guidance and example prompts, the LLM needs access to the "original question," "context," and the "major premise" generated in Stage 2.

Stage 4: Minor Premise Production. The task of Stage 4 is to utilize the "context" information provided by the original task, along with the built-in knowledge of the LLM, to answer the minor premise question posed in Stage 3. This stage aims to obtain the correct information about a specific aspect of the particular instance in the original question, leading to the formulation of an accurate and matching minor premise. Given the potential complexity of the minor premise questions, we guide LLM to employ the Chain-of-Thought (CoT) technique to answer the minor premise question and to organize and obtain the minor premise. Furthermore, to avoid the interference caused by excessive information, in this stage, besides the guidance and example prompts, the LLM has access only to the "minor premise question" and "context" without needing to see the "original question" again. The "minor premise question" and "context" already contain all the information necessary for the task of the LLM at this stage; viewing additional information like the "original question" could instead lead to distractions and affect performance.

Stage 5: Final Syllogistic Reasoning After the aforementioned stages, complex original knowledge-based reasoning questions can now be answered using syllogistic reasoning. The specific approach involves designing the appropriate task instruction and example prompts, allowing the LLM to engage in syllogistic reasoning based on the major and minor premises generated in earlier stages, to arrive at the answer to the original question. Therefore, in Stage 5, we design the LLM to have access, in addition to the guidance and example prompts, to the "major premise" generated in Stage 2, the "minor premise" generated in Stage 4, and the "original question".

Experiments

To evaluate the effectiveness of our SR-FoT, we conducted a series of experiments using both Open-source and closed-source LLMs on several common knowledge-based reasoning question-answer datasets.

Experiment Setup

Datasets. To fully demonstrate the effectiveness and generalization of our SR-FoT, we conduct a series of experiments on three datasets from different fields.

ScienceQA (Lu et al. 2022) is a scientific question-answering dataset and contains 21,208 multimodal multiple-

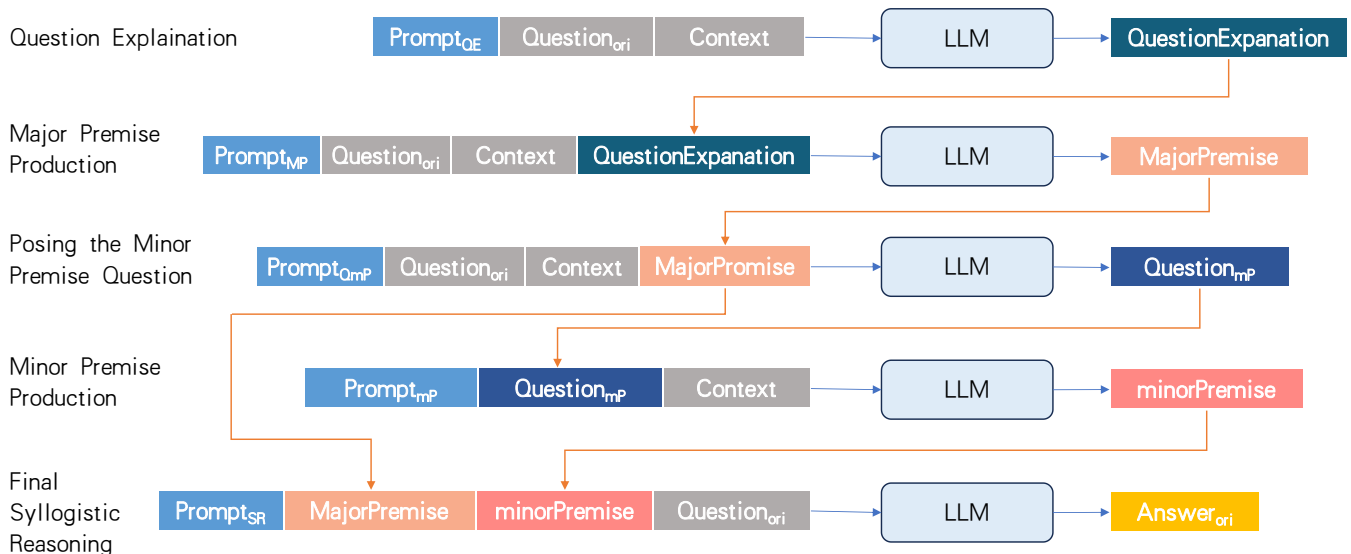


Figure 2: Procedure of our SR-FoT. $Question_{ori}$: Original Question, Context: Context provided for Original Question, $Answer_{ori}$: Answer for Original Question, $Question_{mp}$: Question for Minor Premise, $Prompt_{CoT}$: Guide Prompt for CoT, $Prompt_{QE}$: Guide Prompt for Question Explanation, $Prompt_{MP}$: Guide Prompt for Major Premise Production, $Prompt_{QmP}$: Guide Prompt for Posing the Minor Premise Question, $Prompt_{mp}$: Guide Prompt for Minor Premise Production, $Prompt_{SR}$: Guide Prompt for Final Syllogistic Reasoning and so on.

Method	GPT-3.5-turbo			DeepSeek-V2			Qwen1.5-32B-Chat		
	ScienceQA	StrategyQA	BoolQ	ScienceQA	StrategyQA	BoolQ	ScienceQA	StrategyQA	BoolQ
<i>Single-Round</i>									
Base	85.9	84.8	79.4	87.0	92.3	86.3	87.4	90.8	83.4
CoT (Wei et al. 2022)	86.9	91.3	79.4	88.6	91.8	86.0	87.1	91.0	83.0
SR-FoT (Ours)	87.4	92.1	81.7	91.8	93.2	87.3	88.7	91.9	84.6
<i>Multi-Round</i>									
SC-CoT (Wang et al. 2022)	87.4	91.5	82.8	88.6	92.4	87.4	87.6	92.3	84.3
C-CoT (Fu et al. 2022)	87.4	90.2	81.0	88.5	92.3	87.4	87.2	92.0	83.9
SC-SR-FoT (Ours)	88.9	93.4	85.0	93.0	95.0	89.4	90.4	94.4	88.2

Table 1: Comparison with the state-of-the-art methods on the ScienceQA, StrategyQA, and BoolQ datasets.

choice science questions. It can be divided into three subjects: natural science, language science, and social science. It requires the language model to select one answer from multiple options, usually requiring multi-step reasoning. In our experiment, we employ the test set samples which only have a text context, with a total of 2224. We report the accuracy of our SR-FoT and comparison methods on this set.

StrategyQA (Geva et al. 2021) is a question-answering dataset focusing on open-domain questions. Its questions contain multiple reasoning steps, and a strategy should be used to obtain the answers. In our experiment, we evaluate the methods with accuracy on the train set, which includes 2290 samples.

BoolQ (Clark et al. 2019) is a reading comprehension dataset consisting of 16k samples. They often query for complex, non-factoid information, and require difficult entailment-like inference to solve. In our experiment, we compare the accuracy of our SR-FoT with other methods on the dev set, with a total of 3270.

Experimental Setting. Our experiments are performed using API calls on the proprietary model GPT-3.5-turbo (Ouyang et al. 2022), the open-source model DeepSeek-v2 (DeepSeek-AI 2024) with 236B parameters, and Qwen1.5-32B-Chat (Bai et al. 2023) version with 32B parameters. The control group methods include the Base method, Chain of Thought (CoT) (Wei et al. 2022), Self-consistency CoT (SC-CoT) (Wang et al. 2022), and Complexity-based CoT (C-CoT) (Fu et al. 2022) methods. Our own approaches included SR-FoT and Self-consistency SR-FoT (SC-SR-FoT), which represents our SR-FoT following the self-consistency sampling and aggregation settings of SC-CoT. In the single-round sampling methods which include Base, CoT and SR-FoT, the hyperparameters on GPT-3.5-turbo and Qwen1.5-32B-Chat are set to $top_p=0.3$ and $temperature=0.2$, while the temperature on DeepSeek-v2 was set to the default recommended value of 1 (DeepSeek only allows the temperature hyperparameter to be adjusted). In the multi-round sampling methods which in-

explain_prompt = ""You are a knowledgeable scholar. Below is a question with corresponding context. Please read and understand the question carefully, explain the meaning of the question, and find what information in the context is needed to answer this question. Please note that the information you find must be from the context, and you cannot directly answer this question and assume some information does not appear in the context.

Question: {question}\n Context: {context}\n Explanation: ""

major_prompt = ""From the perspective of the syllogism, please propose a major premise for the question according to the context and the explanation. In a syllogism, the major premise is a general statement or a universal truth. Note that the major premise must be found or supported in the context and as much as possible related to the question.

Question: {question}\n Context: {context}\n Explanation: {explanation}\n Major Premise: ""

minor_question_prompt = ""From the perspective of the syllogism, please ask a minor premise question for the given question. The minor premise question you ask should be as relevant as possible to the given question.

Question: {question}\n Context: {context}\n Major Premise: {major}\n Minor Premise Question: ""

minor_prompt = ""Please according to the context, think step by step to answer the question. Note that You must think and answer only based on the context.

Question: {minor_question}\n Context: {context}\n Minor Premise: ""

final_prompt = ""From the perspective of the syllogism, please according to the major premise and minor premise, think step by step to answer the question.

Major Premise: {major}\n Minor Premise: {minor}\n Question: {question}\n Answer: ""

Figure 3: Prompts for each stage of our SR-FoT.

clude SC-CoT, C-CoT and SC-SR-FoT, we perform 10 samplings each. To enhance the diversity of sampling outcomes, the hyperparameters on GPT-3.5-turbo and Qwen1.5-32B-Chat for top_p and temperature are adjusted to 0.7 and 0.9 respectively. The temperature hyperparameter on DeepSeek remained at the default recommended value of 1. The number of in-context example prompts used in all the methods on the ScienceQA, StrategyQA, and BoolQ datasets are 5, 2, and 2, respectively.

Experimental Results and Analyses

Performance on ScienceQA. Scientific question answering is a task scenario that often requires deductive reasoning. As seen in Tab. 1, under GPT-3.5-turbo, in the comparison of single-round sampling methods, our SR-FoT outperforms the Base and CoT methods by 1.5% and 0.5% respectively and is on par with multi-round sampling methods like SC-CoT and C-CoT. In the comparison of multi-round sampling methods, our SC-SR-FoT exceeds SC-CoT and C-CoT methods by 1.5%. Under the open-source model DeepSeek-V2, SR-FoT outperforms the Base and CoT by 4.8% and 3.2% respectively, even surpassing multi-round sampling methods. What’s more, our SC-SR-FoT further increases the accuracy to 93.0%. Under Qwen1.5-32B-Chat, compared to the Base and CoT methods, our SR-FoT has an improvement of 1.3% and 1.6% respectively. Compared to SC-CoT and C-CoT, our SC-SR-FoT also performs better, surpassing them by 2.8% and 3.2% respectively. These indicate that our methods achieves greater superiority on the ScienceQA dataset under multiple models.

Performance on StrategyQA and BoolQ. StrategyQA and BoolQ are two other knowledge-based reasoning question-answer datasets that require a true or false judgment based on context or common sense knowledge. From Table 1, for StrategyQA under GPT-3.5-turbo, in the comparison of single-round sampling methods, our SR-FoT outperforms Base and CoT by 7.3% and 0.8% respectively; in the comparison of multi-round sampling methods, our SC-SR-FoT exceeds SC-CoT and C-CoT by 1.9% and 3.2% respectively. Similar trends are observed under DeepSeek-V2 and Qwen1.5-32B-Chat. In addition, our SR-FoT and SC-SR-FoT also perform the best in both single-round sampling methods and multi-round sampling methods on BoolQ under the three models.

Overall, whether under the closed-source large language model GPT-3.5-turbo or the open-source large language models DeepSeek-V2 and Qwen1.5-32B-Chat, our SR-FoT achieve a superior accuracy compared with other compared methods on the ScienceQA, StrategyQA, and BoolQ datasets. This demonstrates the effectiveness of our SR-FoT.

It is worth noting that under DeepSeek-V2 and Qwen1.5-32B-Chat, the Base method achieves relatively high results across all three datasets, while the benefits of the CoT method show signs of saturation, and at times perform worse than the Base method. However, our methods, whether under single-round sampling settings (SR-FoT) or multi-round aggregated sampling settings (SC-SR-FoT), are still able to further enhance performance, demonstrating greater potential for performance gains. We believe this is because our SR-FoT employs a syllogistic deductive reasoning frame-

Question: Which word does not rhyme?	Options: ['bib', 'job', 'sob']	Ground Truth: bib
<p>CoT: The answer is 'job', reasoning process step by step: 1. The question asks which word does not rhyme with the others. 2. The options provided are 'bib', 'job', and 'sob'. 3. To determine which word does not rhyme, we need to consider the ending sounds of each word. 4. 'Bib' and 'sob' both end with the same 'b' sound, making them rhyme with each other. 5. However, 'job' ends with a different sound, the 'b' sound, making it not rhyme with 'bib' and 'sob'. 6. Therefore, based on the context and the ending sounds of the words, 'job' does not rhyme with the others.</p>		
<p>SR-FoT: Question Explanation: The question is asking which word does not rhyme with the others. To answer the question, we need to understand the concept of rhyming words and compare the sounds of each word in the options to determine which one does not rhyme. Major Premise Production: Rhyming words have similar ending sounds, typically the vowel and consonant sounds in the final syllables of the words. Posing the Minor Premise Question: Which word has a different ending sound compared to the others? Minor Premise Production: The word 'bib' has a different ending sound compared to the others. Final Syllogistic Reasoning: The answer is 'bib'. The reasoning process of syllogism is as follows: Since rhyming words have similar ending sounds, typically the vowel and consonant sounds in the final syllables of the words, and the word 'bib' has a different ending sound compared to the others, it can be concluded that 'bib' does not rhyme with 'job' and 'sob'.</p>		

Figure 4: A case of using CoT and SR-FoT to answer a question in the ScienceQA dataset respectively. The highlighted red parts indicate the incorrect or misleading content, while the highlighted green parts indicate the content that helps correct reasoning.

work, allowing LLMs to address these knowledge-based reasoning tasks based on a more rigorous reasoning process, thereby achieving better overall performance.

Ablation Study

Effectiveness Comparisons for Subcategories. As shown in table 2, we conduct the experiments on the ScienceQA dataset with DeepSeek-v2. The results demonstrate that our methods can enhance the reasoning performance of the language model across questions of different difficulty levels and various subjects, compared with the state-of-the-art methods. When increasing the consistency of the proposed method, the beneficial effects become more significant.

Ablation of Stages. As shown in Table 3, we conduct experiments on ScienceQA under DeepSeek-V2 to verify the effectiveness of each stage in our method. Specifically, ‘all in one stage’ denotes using instructions and examples to let the LLM provide the premises based on the question and options, and then directly provide the answers. ‘w/o stage 3’ denotes providing the minor premise directly, instead of posing it as a question first and then answering. The results demonstrate that the completeness of each stage is important. In detail, discarding either the problem explanation or the major and minor premises would decrease the performance. Furthermore, allowing the language model to di-

rectly provide the major and minor premise would significantly reduce its performance, demonstrating the necessity of the multi-stage thinking framework in our SR-FoT.

Impact of Visible Information in Various Stages. As shown in Table 4, we conduct the experiments on the StrategyQA dataset with DeepSeek-v2. ‘w/o context in stage 3’ denotes that the minor premise question is generated without considering the context. ‘add Q_{ori} in stage 4’ denotes that providing the original question, minor premise question, and context all to the LLM during the process of answering the minor premise question. The results demonstrate that both decreasing or increasing the content of the input prompts adversely affect performance. This underlines the appropriateness of the designed visible information at each stage of our SR-FoT.

Rigor Analysis

To more directly analyze whether our SR-FoT improves the rigor of the reasoning process compared to CoT, we randomly select 50 cases from each of the three datasets under GPT-3.5-turbo for manual evaluation. For CoT and SR-FoT, if all intermediate steps from the first step of reasoning to obtaining the final answer are correct and logically progressive, without any factual inconsistencies or self-inconsistencies, we call them rigorous; otherwise, they are not rigorous. The results are in Table 5. From the table, it can be found that

Method	Subject			Grade		Average
	Natural	Social	Language	Grade1-6	Grade7-12	
Base	89.4	92.0	84.1	88.2	85.4	87.0
CoT (Wei et al. 2022)	91.6	96.0	84.8	90.8	85.5	88.6
SC-CoT (Wang et al. 2022)	91.7	96.8	84.4	90.8	85.3	88.6
C-CoT (Fu et al. 2022)	91.9	96.0	84.7	90.9	85.7	88.5
w/o explanation	94.2	98.4	83.2	90.9	87.0	90.4
w/o major premise	93.7	96.0	85.1	91.1	87.8	90.9
w/o minor premise	87.3	94.4	87.3	91.0	83.1	87.7
all in one stage	91.2	95.2	86.0	91.0	86.1	89.8
w/o stage 3	94.0	96.8	84.1	90.1	88.5	89.4
SR-FoT (Ours)	95.0	97.6	85.7	91.4	89.9	91.8
SC-SR-FoT (Ours)	97.1	96.8	88.4	93.4	92.3	93.0

Table 2: Effectiveness comparisons for subcategories on the ScienceQA dataset with DeepSeek-v2.

Method	Accuracy
w/o explanation	90.4
w/o major premise	90.9
w/o minor premise	89.8
all in one stage	87.7
w/o stage 3	89.4
Ours	91.8

Table 3: Ablation study of stages in our proposed methods on the ScienceQA dataset under DeepSeek-V2.

Method	Accuracy
w/o context in stage 3	92.4
add Q_{ori} in stage 4	92.7
Ours	93.2

Table 4: Ablation study of visible information in various stages on the StrategyQA dataset.

Method	Rigorous	Not Rigorous	Rigor Rate
ScienceQA SR-FoT	41	9	0.82
ScienceQA CoT	38	12	0.76
StrategyQA SR-FoT	46	4	0.92
StrategyQA CoT	44	6	0.88
BoolQ SR-FoT	48	2	0.96
BoolQ CoT	40	10	0.8

Table 5: Fifty cases using CoT and SR-FoT respectively, randomly selected from the three datasets under GPT-3.5-turbo, are analyzed to assess their rigor, and their rigor rates are subsequently calculated.

Dataset	MaPE	MiPQE	MiPE	FRPE
ScienceQA	13 (26%)	4 (8%)	24 (48%)	9 (18%)
StrategyQA	15 (30%)	5 (10%)	14 (28%)	16 (32%)
BoolQ	10 (20%)	14 (28%)	10 (20%)	16 (32%)

Table 6: Error sources on a random sample of 50 incorrect examples of our SR-FoT from the three datasets using the GPT-3.5-turbo model. ‘MaPE’ denotes major premise error, ‘MiPQE’ denotes minor premise question error, ‘MiPE’ denotes minor premise error and ‘FRPE’ denotes final reasoning process error.

our SR-FoT has a higher rigor rate than CoT on all three datasets, indicating that our SR-FoT has enhanced the rigor and interoperability of LLM reasoning. For specific comparison cases about rigor, please refer to the supplementary materials.

Error Analysis

We randomly selected 50 error cases from each of the three datasets under GPT-3.5-turbo to perform an error analysis of our SR-FoT. The sources of errors and their respective proportions are as in table 6. From the error analysis, it can be found that the proportion of different types of errors varies on different datasets. In ScienceQA, most errors stem from the step of extracting suitable minor premises from the question information. In StrategyQA, the main errors stem from the final reasoning process and mistakes in presenting the major premise. In BoolQ, the primary errors originate from the final reasoning process and errors in formulating the minor premise.

Case Study

We give a case on the ScienceQA dataset to show how CoT and SR-FoT work. In Fig. 4, it can be seen that in the fourth and fifth reasoning steps of CoT, the model misunderstands the rhyme condition and thus infers wrong information, resulting in an incorrect final answer. In SR-FoT, the question explanation points out a reasonable direction for the major premise, then the major premise gives more sufficient rhyme conditions, and the minor premise correctly distinguishes different ending sounds. With their joint help, the model gives the correct final answer. More cases can be found in the supplementary materials.

Conclusion

In this paper, we have developed a multi-stage syllogistic reasoning framework of thought(SR-FoT) to guide LLMs in solving complex and diverse knowledge-based reasoning question-answering tasks using syllogistic deductive reasoning. Experiments across various knowledge-based reasoning datasets under various LLMs demonstrate the effectiveness and advantages of our method.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant 62276283, in part by the China Meteorological Administration's Science and Technology Project under Grant CMAJBGS202517, in part by Guangdong Basic and Applied Basic Research Foundation under Grant 2023A1515012985, and in part by Basic and Applied Basic Research Special Projects under Grant SL2022A04J01685.

References

- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; Hui, B.; Ji, L.; Li, M.; Lin, J.; Lin, R.; Liu, D.; Liu, G.; Lu, C.; Lu, K.; Ma, J.; Men, R.; Ren, X.; Ren, X.; Tan, C.; Tan, S.; Tu, J.; Wang, P.; Wang, S.; Wang, W.; Wu, S.; Xu, B.; Xu, J.; Yang, A.; Yang, H.; Yang, J.; Yang, S.; Yao, Y.; Yu, B.; Yuan, H.; Yuan, Z.; Zhang, J.; Zhang, X.; Zhang, Y.; Zhang, Z.; Zhou, C.; Zhou, J.; Zhou, X.; and Zhu, T. 2023. Qwen Technical Report. *arXiv preprint arXiv:2309.16609*.
- Bara, B. G.; Bucciarelli, M.; and Johnson-Laird, P. N. 1995. Development of syllogistic reasoning. *The American journal of psychology*, 157–193.
- Bharadiya, J. 2023. A comprehensive survey of deep learning techniques natural language processing. *European Journal of Technology*, 7(1): 58–66.
- Bibel, W. 2013. *Automated theorem proving*. Springer Science & Business Media.
- Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y. T.; Li, Y.; Lundberg, S.; et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Bucciarelli, M.; and Johnson-Laird, P. N. 1999. Strategies in syllogistic reasoning. *Cognitive Science*, 23(3): 247–303.
- Chang, Y.; Wang, X.; Wang, J.; Wu, Y.; Yang, L.; Zhu, K.; Chen, H.; Yi, X.; Wang, C.; Wang, Y.; et al. 2024. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*, 15(3): 1–45.
- Chiang, W.-L.; Li, Z.; Lin, Z.; Sheng, Y.; Wu, Z.; Zhang, H.; Zheng, L.; Zhuang, S.; Zhuang, Y.; Gonzalez, J. E.; Stoica, I.; and Xing, E. P. 2023. Vicuna: An Open-Source Chatbot Impressing GPT-4 with 90%* ChatGPT Quality.
- Chowdhary, K.; and Chowdhary, K. 2020. Natural language processing. *Fundamentals of artificial intelligence*, 603–649.
- Clark, C.; Lee, K.; Chang, M.-W.; Kwiatkowski, T.; Collins, M.; and Toutanova, K. 2019. BoolQ: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*.
- DeepSeek-AI. 2024. DeepSeek-V2: A Strong, Economical, and Efficient Mixture-of-Experts Language Model. *arXiv:2405.04434*.
- Deng, W.; Pei, J.; Kong, K.; Chen, Z.; Wei, F.; Li, Y.; Ren, Z.; Chen, Z.; and Ren, P. 2023. Syllogistic Reasoning for Legal Judgment Analysis. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 13997–14009.
- Dhuliawala, S.; Komeili, M.; Xu, J.; Raileanu, R.; Li, X.; Celikyilmaz, A.; and Weston, J. 2023. Chain-of-verification reduces hallucination in large language models. *arXiv preprint arXiv:2309.11495*.
- Dong, G.; Zhao, J.; Hui, T.; Guo, D.; Wang, W.; Feng, B.; Qiu, Y.; Gongque, Z.; He, K.; Wang, Z.; et al. 2023. Revisit input perturbation problems for llms: A unified robustness evaluation framework for noisy slot filling task. In *CCF International Conference on Natural Language Processing and Chinese Computing*, 682–694. Springer.
- Floridi, L.; and Chiriatti, M. 2020. GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30: 681–694.
- Fu, Y.; Peng, H.; Sabharwal, A.; Clark, P.; and Khot, T. 2022. Complexity-based prompting for multi-step reasoning. In *The Eleventh International Conference on Learning Representations*.
- Geva, M.; Khashabi, D.; Segal, E.; Khot, T.; Roth, D.; and Berant, J. 2021. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of the Association for Computational Linguistics*, 9: 346–361.
- Huang, J.; and Chang, K. C.-C. 2022. Towards reasoning in large language models: A survey. *arXiv preprint arXiv:2212.10403*.
- Johnson-Laird, P. N. 1999. Deductive reasoning. *Annual review of psychology*, 50(1): 109–135.
- Khemlani, S.; and Johnson-Laird, P. N. 2012. Theories of the syllogism: A meta-analysis. *Psychological bulletin*, 138(3): 427.
- Khurana, D.; Koli, A.; Khatler, K.; and Singh, S. 2023. Natural language processing: State of the art, current trends and challenges. *Multimedia tools and applications*, 82(3): 3713–3744.
- Lee, G.-G.; Latif, E.; Wu, X.; Liu, N.; and Zhai, X. 2024. Applying large language models and chain-of-thought for automatic scoring. *Computers and Education: Artificial Intelligence*, 100213.
- Liu, X.; Pang, T.; and Fan, C. 2023. Federated Prompting and Chain-of-Thought Reasoning for Improving LLMs Answering. In *International Conference on Knowledge Science, Engineering and Management*, 3–11. Springer.
- Lu, P.; Mishra, S.; Xia, T.; Qiu, L.; Chang, K.-W.; Zhu, S.-C.; Tafjord, O.; Clark, P.; and Kalyan, A. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35: 2507–2521.
- Lyu, Q.; Havaldar, S.; Stein, A.; Zhang, L.; Rao, D.; Wong, E.; Apidianaki, M.; and Callison-Burch, C. 2023. Faithful chain-of-thought reasoning. *arXiv preprint arXiv:2301.13379*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.

Seals, S.; and Shalin, V. L. 2023. Evaluating the Deductive Competence of Large Language Models. *arXiv preprint arXiv:2309.05452*.

Smiley, T. J. 1973. What is a syllogism? *Journal of philosophical logic*, 136–154.

Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Turpin, M.; Michael, J.; Perez, E.; and Bowman, S. 2024. Language models don't always say what they think: unfaithful explanations in chain-of-thought prompting. *Advances in Neural Information Processing Systems*, 36.

Wang, H.; Xin, H.; Zheng, C.; Li, L.; Liu, Z.; Cao, Q.; Huang, Y.; Xiong, J.; Shi, H.; Xie, E.; Yin, J.; Li, Z.; Liao, H.; and Liang, X. 2023. LEGO-Prover: Neural Theorem Proving with Growing Libraries. *arXiv:2310.00656*.

Wang, X.; Wei, J.; Schuurmans, D.; Le, Q.; Chi, E.; Narang, S.; Chowdhery, A.; and Zhou, D. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.

Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.

Wu, Y.; Han, M.; Zhu, Y.; Li, L.; Zhang, X.; Lai, R.; Li, X.; Ren, Y.; Dou, Z.; and Cao, Z. 2023. Hence, Socrates is mortal: A Benchmark for Natural Language Syllogistic Reasoning. In *Findings of the Association for Computational Linguistics: ACL 2023*, 2347–2367.

Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T. L.; Cao, Y.; and Narasimhan, K. 2023. Tree of thoughts: Deliberate problem solving with large language models, 2023. *URL <https://arxiv.org/pdf/2305.10601.pdf>*.

Ye, H.; Liu, T.; Zhang, A.; Hua, W.; and Jia, W. 2023a. Cognitive mirage: A review of hallucinations in large language models. *arXiv preprint arXiv:2309.06794*.

Ye, M.; Kuribayashi, T.; Suzuki, J.; Kobayashi, G.; and Funayama, H. 2023b. Assessing step-by-step reasoning against lexical negation: A case study on syllogism. *arXiv preprint arXiv:2310.14868*.

Yu, F.; Zhang, H.; Tiwari, P.; and Wang, B. 2023. Natural Language Reasoning, A Survey. *arXiv:2303.14725*.

Zhang, Z.; Zhang, A.; Li, M.; and Smola, A. 2022. Automatic chain of thought prompting in large language models. *arXiv preprint arXiv:2210.03493*.

Zhou, D.; Schärli, N.; Hou, L.; Wei, J.; Scales, N.; Wang, X.; Schuurmans, D.; Cui, C.; Bousquet, O.; Le, Q. V.; et al. 2022. Least-to-Most Prompting Enables Complex Reasoning in Large Language Models. In *The Eleventh International Conference on Learning Representations*.