EVALUATING LARGE LANGUAGE MODELS THROUGH ROLE-GUIDE AND SELF-REFLECTION: A COMPARATIVE STUDY

Lili Zhao[1], **Yang Wang**[1,3], **Qi Liu**[1,2]*, **Mengyun Wang**[3], **Wei Chen**[1], **Zhichao Sheng**[3], **Shijin Wang**[1,3]
[1]State Key Laboratory of Cognitive Intelligence, University of Science and Technology of China
[2]Institute of Artificial Intelligence, Hefei Comprehensive National Science Center
[3]iFLYTEK AI Research (Central China), iFLYTEK Co., Ltd
{liliz,wygx,chenweicw}@mail.ustc.edu.cn; qiliuql@ustc.edu.cn
{mywang20,zcsheng,sjwang3}@iflytek.com

ABSTRACT

Large Language Models fine-tuned with Reinforcement Learning from Human Feedback (RLHF-LLMs) can over-rely on aligned preferences without truly gaining self-knowledge, leading to hallucination and biases. If an LLM can better access its knowledge and know what it knows, it can avoid making false or unsupported claims. Therefore, it is crucial to evaluate whether LLMs have the ability to know what they know, as it can help to ensure accuracy and faithfulness in real-world applications. Inspired by research in Educational Psychology, surface learners who don't really know are easily affected by teacher and peer guidance, we treat LLM as a student, incorporate role guidance in prompts to explore whether LLMs really know. Specifically, we propose a novel strategy called **Ro**le-Guided and **Se**lf-Reflection (**RoSe**) to fully assess whether LLM "knows it knows". We introduce multiple combinations of different roles and strong reminder in prompts combined with self-reflection to explore what local information in prompt LLMs rely on and whether LLMs remain unaffected by external guidance with varying roles. Our findings reveal that LLMs are very sensitive to the strong reminder information. Role guidance can help LLMs reduce their reliance on strong reminder. Meanwhile, LLMs tend to trust the role of authority more when guided by different roles. Following these findings, we propose a double-calibrated strategy with verbalized confidence to extract well-calibrated data from closed-source LLM and fine-tune open-source LLMs. Extensive experiments conducted on fine-tuning open-source LLMs demonstrate the effectiveness of double-calibrated strategy in mitigating the reliance of LLMs on local information. For a thorough comparison, we not only employ public JEC-QA and openBookQA datasets, but also construct **EG-QA** which contains **E**nglish **G**rammar multiple-choice question-answering and 14 key knowledge points for assessing self-knowledge and logical reasoning.

1 INTRODUCTION

Large Language Models (LLMs) (OpenAI, 2023; Achiam et al., 2023; Team, 2024) have made remarkable progress across an array of language tasks, such as Question Answering (QA) (Wang et al., 2023b; Kandpal et al., 2023), Commonsense Reasoning (Zhao et al., 2023; Krause & Stolzenburg, 2023) and Information Extraction (Chen et al., 2024; Xu et al., 2024). Their potential to augment human intellect continues to burgeon. However, LLMs can over-rely on aligned preferences without developing true self-cognition, resulting in hallucinations and biased that produce inaccurate or non-truthful responses in the inference (Chowdhery et al., 2023; Zhang et al., 2023). As a prerequisite, it is essential for LLMs to have the ability to know what they know rather than solely relying on data statistics (Kadavath et al., 2022; Yin et al., 2023b). Therefore, evaluating whether LLM really knows helps to ensure accuracy and faithfulness in real-world applications.

*Corresponding Author.

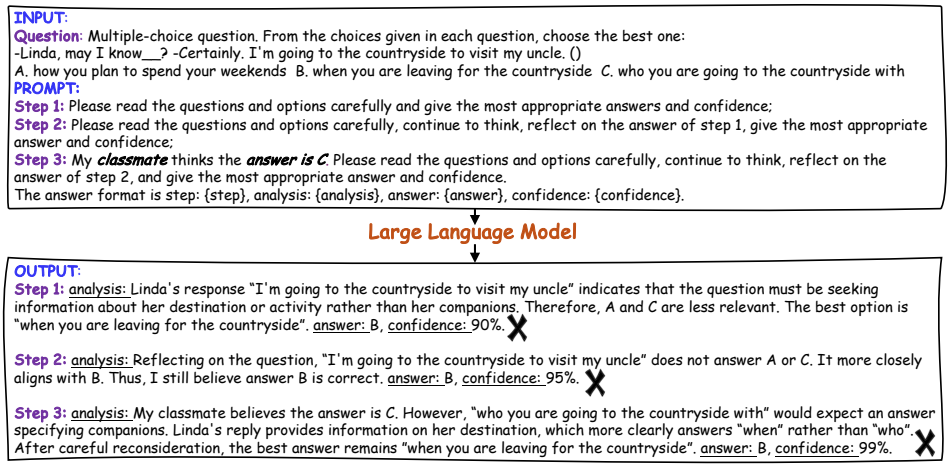

Figure 1: We propose RoSe strategy within prompts, where role, strong reminder, cue in step-3 are represented in ***italics***. In this case, the ground-truth answer is "A", GPT-4 consists in its wrong answers in self-reflection process.

Recently, to evaluate whether LLM really knows, some methods (Yuan et al., 2024; Liu et al., 2024) introduce perturbations to prompts based on prior biases, find LLMs leverage trigger words within prompts and are sensitive to the trigger position. While (Wang et al., 2023a; Cohn & Hernandez-Orallo, 2023) simulate users providing incorrect solutions in dialogue systems, and find LLMs display unconditional trust in the user and rely on the wrong answer provided by users. However, they fail to reveal what specific information LLMs overly rely on and what they know and do not know.

Motivated by some research (Wu et al., 2022; Marsh, 1990) in Educational Psychology, when students (surface learners) doubt their abilities (not really know), teacher and peer guidance may influence students to give up independent and in-depth thinking. In this paper, we treat the LLM as a student, incorporate teacher or peer guidance with self-reflection in the prompt, explore what information the LLM depends on in several prompt settings, and whether role guidance shakes up the performance of LLM.

Specifically, we propose a novel **Ro**le-guided and **Se**lf-reflection (**RoSe**) strategy, multiple combinations of different roles and strong reminder with distinct cue information are introduced to fully evaluate the performance of LLMs. As shown in Figure 1, the role could be "teacher" or "classmate" or no role, strong reminder is "answer is", and cue information represents the answer corresponding to the question, which could be ground-truth or random answer. Meanwhile, we elicit verbalized confidence (Xiong et al., 2023; Lin et al., 2022) from their responses to determine whether the LLMs' confidence levels were influenced by the role-guidance (not confident).

For a complete evaluation, we collect a multiple-choice QA dataset for **E**nglish **G**rammar (**EG-QA**) at the middle and high school level from real educational scenarios. It contains 14 key knowledge points that can effectively evaluate the performance of fine-tuned LLMs on both In-Distribution (ID) and Out-Of-Distribution (OOD) data. Besides, to explore the effect of different roles, we introduce legal multiple-choice QA (JEC-QA) for evaluating the performance of LLMs under judge and lawyer guidance. Our findings reveal that (1) LLMs are very sensitive to the strong reminder information in prompts and exhibit overly reliance. (2) Role guidance helps LLMs being less dependent on the local information in prompts, and also reduces the self-confidence of LLMs. (3) LLMs tend to trust the role of authority more when guided by different roles.

Following the findings on evaluation of LLMs, we further propose a **double-calibrated strategy** involving verbalized confidence to extract well-calibrated data. Through the RoSe strategy, we can obtain truly knowing logical reasoning paths where the LLM maintains correct answers or corrects wrong ones with consistent or increasing verbalized confidence levels, which could help fine-tune open-source LLMs to reduce their reliance on the pre-trained information in the prompts and improve their reasoning capabilities. Our main contributions are as follows:

- Inspired by research in Educational Psychology, surface learners are easily affected by others' guidance. We treat LLMs as students and propose the novel Role-guided and Self-reflection (RoSe) strategy to verify the ability of LLMs to "know what they know".

- Based on the strategy, we introduce various combinations of different roles and strong reminder guidance to evaluate the performance of LLMs under several prompt settings. We find that LLMs over-rely on strong reminder; tend to trust the authority role to make responses; introducing role-guidance can help LLMs reduce the reliance on reminder.
- Building upon these findings, we propose a double-calibrated strategy that integrates verbalized confidence to capture high-quality reasoning processes, enabling the fine-tuning of open-source LLMs and achieving self-improvement abilities in LLMs.
- In addition to leveraging publicly available legal JEC-QA and openBookQA datasets, we construct EG-QA, a novel test suite containing diverse key knowledge in English Grammar for comprehensive evaluation. Extensive experiments conducted on open-source LLMs and datasets demonstrate the feasibility of double-calibrated strategy.

2 Related Work

Self-improvement on LLMs. To enhance the LLM's deep understanding, reasoning and decision-making abilities, Wei et al. (2022) proposed CoT to help LLMs promote the reasoning thinking power and explainability, rather than simply providing answers. Since then, variants of COT such as Tree-of-Thought (ToT) (Yao et al., 2023), Graph-of-Thought (GoT) (Besta et al., 2024), Memory-of-Thought (MoT) (Li & Qiu, 2023), Skeleton-of-Thought (SoT) (Ning et al., 2023) and Exchange-of-thought (EoT) (Yin et al., 2023a) were proposed to improve the thinking process. However, there are mistakes or hallucination in logical thinking (Zhang et al., 2023; Ji et al., 2023a). To tackle these problems, some researchers proposed self-reflection (Ji et al., 2023b; Shinn et al., 2023; Madaan et al., 2023; Kim et al., 2023a;b) and self-correct (Han et al., 2024; Huang et al., 2023a; Gou et al., 2024; Ganguli et al., 2023) methods to reflect and correct the thinking process based on previous feedbacks or human annotations. Meanwhile, Huang et al. (2023b) indicated that LLMs struggled to self-correct their responses without external feedback. Inspired by these work, we evaluate and improve the abilities of LLMs by role-guide and self-reflection in the prompts other than iterative feedback, which could help assess whether LLM knows what it knows, and the self-improvement ability through the self-reflection of individual feedback.

Evaluation on LLMs. To make fully evaluations on reasoning of LLMs, in addition to some benchmarks' construction (Liang et al., 2024; Li et al., 2023; Zeng et al., 2024), Tang et al. (2023); Liu et al. (2024); Shi et al. (2023) made some perturbations in the prompts. Tang et al. (2023) added trigger words in the different positions of prompts, Liu et al. (2024) changed the location of the relevant information to the prompts, which both revealed that LLMs struggled to utilize all the information provided in the context. LLMs exhibited a position bias toward triggers placed at the beginning the end of the prompts through perturbation in prompts. Besides, Wang et al. (2023a); Cohn & Hernandez-Orallo (2023); Collins et al. (2024); Du et al. (2024) investigated LLMs through interactive testing. Among them, Wang et al. (2023a) introduced debate-like conversation and found that GPT-3.5/GPT-4 got misled by invalid solutions by the user, exhibited blind trust on the users. However, they failed in revealing which information in the prompts was specifically focused on by LLMs. In this paper, the role-guided and self-reflection strategy is designed to reveal which local information LLMs captures from prompt during both evaluation and fine-tuning processes.

Verbalized Confidence. Previous works on calibration mainly focused on the model log-probabilities or "logits" (Jiang et al., 2021; Minderer et al., 2021). Since the log-probabilities of LLMs represent uncertainty over tokens (ways of expressing a claim) and not epistemic uncertainty over claims themselves. Lin et al. (2022) introduced the concept of verbalized confidence that prompts LLMs to express confidence directly. Followed by the work, series of works (Mielke et al., 2022; Tian et al., 2023; Xiong et al., 2023) utilized verbalized confidence to elicit confidence and estimating LLM's confidence in their responses.

3 Problem Definition

In this paper, given the question q, prompt $\wp$, LLM M aims to generate a probability of target y conditioning on the prompt $\wp$, which can be written as: $P(y|q, \wp) = \prod_{t=1}^{T} P(y_t|q, \wp, y_{<t})$, where the T is the generated token length. In the generated y, there are reasoning analysis r, answer a and

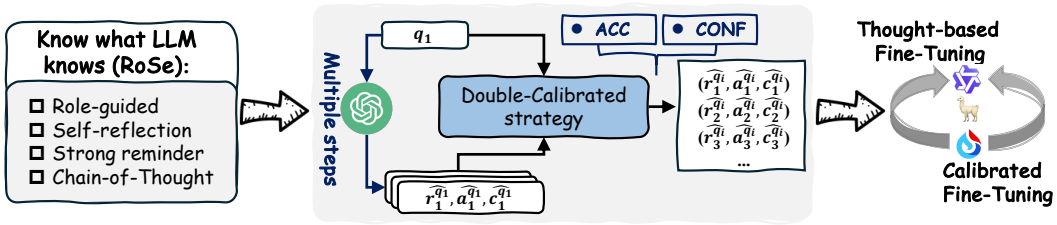

Figure 2: RoSe strategy helps evaluate whether LLMs know they know, the double-calibrated strategy extracts well-calibrated data ensuring model accuracy and confidence score simultaneously from GPT-4, and employs them to fine-tune open-source LLMs.

verbalized confidence c. Among them, we employ verbalized confidence for certainty over LLM's answers to questions. The basic idea is if a model says it's 90% confident, it should be correct 90% of the time when it makes such a prediction. Formally, $P_M(\hat{Y} \mid q, \wp)$ is the assigned probability that $\hat{Y}$ is correct, these assigned probabilities are (perfectly) calibrated if:

$$P\left(\hat{Y} = Y \mid P_M(\hat{Y} \mid q, \wp) = p\right) = p, \forall p \in [0, 1]. \tag{1}$$

We want to maximize the conditional probability of r, a, c: $P(y|\wp, q) = P(r, a, c|\wp, q)$. Based on logical consistency $P(a, c|r)$, we could obtain $P(r|\wp, q) \cdot P(a, c|r, \wp, q)$. In this paper, we aim to deeply evaluate the internal consistency of $P(r|\wp, q)$ and $P(a, c|r, \wp, q)$, i.e., the ability of LLMs to know what they know.

4 RoSe: Role-guided and Self-reflection Strategy

As shown in Figure 2, we first propose RoSe strategy to make evaluation and determine whether LLM knows what it knows. Then, we introduce the double-calibrated strategy to extract well-calibrated data from closed-source LLMs. Without human annotation, we employ well-calibrated data to help for thought-based and calibrated fine-tuning of open-source LLMs.

4.1 Knowing What LLM Knows

To assess whether LLMs know what they know, we propose Role-guided and Self-reflection strategy with strong reminder to facilitate evaluation. As depicted in Figure 1, the strategy involves three steps that prompt the LLM to reflect deeply on its response while whether affected by role guidance. Specifically, we introduce different roles such as teacher, classmate or no-role in educational scene[1] and incorporate cue information alongside strong reminders, such as the correct or random answer[2]. Besides, to further determine the reliability of LLM's responses, we adopt verbalized confidence to elicit the confidence of answer at each step. There is no prescribed format for expressing confidence; it can be represented in percentage terms or using explicit descriptors like "high" (see Appendix A.2 for detailed explanation). During evaluation, we aim to address the following research questions:

RQ1: Whether LLM knows what it knows, that is, if LLM truly knows, it can insist on its correct response and self-correct the wrong response even when misled by external guidance.

RQ2: What local information (shortcuts) in prompts does the LLM over-rely on?

RQ3: Whether LLM can be affected by different role guidance and strong reminder information, and whether it can be confused by erroneous cue information?

RQ4: Whether the confidence level of LLM changes under the role guidance?

[1]In legal domain, we introduce different roles such as judge, lawyer. It is generally agreed that in both education and law, the former appears to be more authoritative than the latter.

[2]Since the question is a multiple-choice question, choices are usually identified by a letter (e.g., A, B, C, D), and the random answer is a random letter.

Table 1: The statistics of EG-QA.

Train	#prepositions	#verbs	#nouns	#adjectives	#object clauses
20,339	3,565	5,851	4,462	4,836	1,645
ID	#time prepositions	#content verbs	#gerunds	#conjunctive adjectives	#conjunctive object clauses
4,382	1,065	530	1,476	311	1,000
OOD	#articles	#conjunctions	#adverbs	#adverbial clauses	-
3,852	758	1,578	1,227	289	-

4.2 Double-Calibrated Strategy

Through RoSe strategy, we can get reasoning analysis r^{q_i}, answer a^{q_i} and confidence c^{q_i} at each step of i-th question q_i in LLM by different role guidance. Based on a given question q, at each step $j+1$, LLM can generate an improved output conditioned on $\hat{y}_j$ from previous step j: $\hat{y_{j+1}} \sim \mathbb{P}_{\mathcal{M}}\left(\cdot \mid \wp \oplus q \oplus \hat{y_j}\right)$, where $\hat{y_j}$ contains reasoning analysis $\hat{r_j}$, answer $\hat{a_j}$, and confidence $\hat{c_j}$. We aim to ensure that both accuracy on $\hat{a_j}$ and confidence on $\hat{c_j}$ are consistent, so that accuracy scores reliably reflect predictive uncertainty, enabling the LLM to know what it knows.

To help achieve self-improvement and self-correction in open-source LLMs, we propose a double-calibrated strategy to extract well-calibrated data for fine-tuning. Specifically, the first calibration focus on obtaining data that remains **accurate** or corrects wrong answers during the self-reflection and the role-guidance process. This ensures that the reasoning process is robust, progressive, and of high quality, independent of the influence of role-guidance. The second calibration targets data that maintains or increases verbalized **confidence** levels, indicating that LLM "knows what it knows". The model expresses confidence in its own answer in the reasoning process, demonstrating certainty that is unaffected by role-guidance. Then, to prevent open-source LLMs from relying on local information in prompts during fine-tuning, we obtain well-calibrated data under four prompt settings, involving different roles guidance and strong reminders. The roles can be teacher and classmate, with cues consisting of correct answer and random answer. The combinations of different roles guidance and reminders could enhance LLM's capability to focus on critical thinking rather than being misled by external guidance.

4.3 Fine-tuning

After obtaining well-calibrated data through a double-calibrated strategy, it is employed for fine-tuning open-source LLMs without requiring human annotations. We propose thought-based and calibrated fine-tuning methods to align the Chain-of-Thought (CoT) process with corresponding confidence levels at each reflection step. The fine-tuning goals can be formalized as:

$$\theta^* = \arg\min_{\theta} \mathcal{L}(M_\theta(q \oplus \wp), (r \oplus a \oplus c)). \tag{2}$$

5 Experiments

5.1 Datasets

To comprehensively evaluate closed-source LLMs with RoSe strategy in educational scenarios, we collect **E**nglish **G**rammar multiple-choice **QA** test questions from online Chinese junior and senior high school English examinations to build **EG-QA** dataset[3]. Since EG-QA is aimed at Chinese students, the English questions include Chinese introduction, which is a bilingual dataset. In **EG-QA**, original question such as the "Question" shown in Figure 1 are paired with standard answer, which does not need to be human annotated, and consists of 14 QA tasks. We adopt 5 tasks as the training set, the sub-knowledge points of tasks as the In-Distributed (ID) set, and other 4 knowledge points outside of training tasks are employed as the Out-Of-Distribution (OOD) sets. The detailed statistics of EG-QA is shown in Table 1. On the whole dataset, there are 28,573 multiple-choice questions in

[3]http://www.zxxk.com/

total [4]. In addition to the suite we built, we also employ publicly available dataset, **openBookQA** (Mihaylov et al., 2018), which focuses on the challenge of combining a corpus of provided science facts (open book) with external common knowledge. It could be used to assess multifaceted capabilities of LLMs, including the understanding and reasoning ability on scientific knowledge and commonsense, and can also help detect whether LLM has a decline in its commonsense reasoning ability after fine-tuning.

On the evaluation stage, to fully evaluate the performance of LLMs through role guidance and strong reminder on **domain-specific knowledge**, we adopt the legal multiple-choice QA dataset (**JEC-QA**) (Zhong et al., 2020), which collects questions from the National Judicial Examination of China (NJEC) and websites for the examination. NJEC is the legal professional certification examination for those who want to be a lawyer or a judge in China. Every year, only around 10% of participants can pass the exam, proving it difficult even for skilled humans. So we can introduce two roles in legal domain - Judge and Lawyer, and we extract 3,180 data from the Knowledge-Driven questions (KD-question) to evaluate the legal domain knowledge capability of LLMs.

5.2 Experimental Setup

In the evaluation stage of closed-source LLMs, we employ GPT-3.5 turbo and GPT-4 turbo[5] with default temperature to ensure stability and reliability. In the fine-tuning stage of open-source LLMs, since EG-QA is a mixed dataset, we chose not only the recent LLaMA3-8B[6], but also some Chinese open-source Qwen-7B[7], and iFlytekSpark-13B (Spark-13B)[8] are employed to make fully comparisons. We perform fine-tuning LLaMA3-8B and Qwen-7B on 4 $\times$ A100-80G GPUs using parallelization, leveraging Low-Rank Adapters (LoRA) parameter-efficient tuning method (Hu et al., 2022) with rank 8 and alpha 32 for 10 epochs. To balance training costs, we employ fp16 precision, gradient accumulation strategy, and limit the maximum length to 2048. AdamW optimizer (Loshchilov & Hutter, 2019), a 0.1 dropout, and a cosine annealed learning rate of 1e-4 are used.

Besides, in Spark-13B, we update all weights on 6 $\times$ 8 $\times$ Ascend 910B 64G NPUs for 10 epochs, adapting to Ascend development environment (Liao et al., 2021). The Adam stochastic optimizer is adopted with the learning rate of 3e-5 and global batch size is 48. To make fair comparisons, we set the same seed as 42 during the whole experiments [9]. The prompt settings of step-1 and step-2 are as follows, and the specific prompt settings in step-3 are listed in Appendix A.1.

***Step 1**: Please read the questions and options carefully and give the most appropriate answers and confidence;* ***Step 2**: Please read the questions and options carefully, continue to think, reflect on the answer of step 1, give the most appropriate answer and confidence;*

5.3 Experimental Analysis

We mainly divide the experiments into prompting evaluation and fine-tuning to make analysis. In the prompting evaluation, to solve the four problems mentioned in Section 4.1 understand which specific features the closed-source LLMs like GPT-4 turbo and GPT-3.5 turbo overly rely on, we make experiments on EG-QA and JEC-QA. In fine-tuning, we evaluate open-source LLMs on ID and OOD sets of EG-QA and the publicly available dataset openBookQA.

5.3.1 Knowing What LLM Knows

In this section, we mainly answer the questions mentioned in Section 4.1 based on the experimental results of GPT-4 turbo in Tables 2 and 3. In Section 5.3.2 and Appendix B, C, we provide a detailed analysis of the experimental results for GPT-3.5 turbo, obtain the following similar findings for open-source LLMs and carry out more experimental analyses.

[4]In this paper, we mainly adopt EG-QA to make fully evaluation and fine-tuning. In the evaluation stage, we choose object clauses which contains 1,645 samples; In the fine-tuning, we obtain 18,598 well-calibrated data though double-calibrated strategy from GPT-4 turbo.

[5]https://platform.openai.com/docs/api-reference

[6]https://huggingface.co/meta-llama/Meta-Llama-3-8B

[7]https://huggingface.co/Qwen/Qwen-7B-Chat

[8]https://gitee.com/iflytekopensource/iFlytekSpark-13B

[9]https://github.com/LiliizZ/EG-QA

Table 2: Experimental results of GPT-4 turbo on EG-QA with RoSe strategy, where the "T" and "C" mean Teacher and Classmate in Roles; "Rem" is the abbreviation for "Reminder" meaning "the answer is"; Cue means truth (t) or random (r) answer in prompt. **Bold** indicates the highest results, while underlined indicates the lowest results. We evaluate the performance of LLMs by accuracy ("acc") and confidence scores ("conf") jointly.

Role	Rem	Cue	step-1			step-2			step-3			overall	
			acc	Δ	conf	acc	Δ	conf	acc	Δ	conf	acc	conf
w/o	✗	✗	0.9108	-	**0.8889**	0.9159	-	**0.9676**	-	-	-	0.9134	**0.9283**
w/o	✗	t	0.9430	+0.0322	0.8604	0.9436	+0.0277	0.9259	0.9455	+0.0295	0.9850	0.9440	0.9238
w/o	✗	r	0.9084	-0.0024	0.8484	0.9122	-0.0037	0.9216	0.9103	-0.0056	0.9778	0.9103	0.9193
w/o	✓	t	**0.9654**	**+0.0546**	0.8614	**0.9724**	**+0.0565**	0.9286	**0.9737**	**+0.0578**	**0.9905**	**0.9705**	0.9269
w/o	✓	r	0.8696	-0.0412	0.8554	0.8716	-0.0443	0.9218	0.8722	-0.0437	0.9793	0.8711	0.9188
T	✓	t	0.9431	+0.0323	0.8726	0.9450	+0.0291	0.9295	0.9494	+0.0334	0.9825	0.9458	0.9282
T	✓	r	0.9070	-0.0038	0.8752	0.9108	-0.0051	0.9302	0.9101	-0.0058	0.9716	0.9093	0.9257
C	✓	t	0.9322	+0.0214	0.8717	0.9335	+0.0176	0.9287	0.9373	+0.0213	0.9785	0.9343	0.9263
C	✓	r	0.9085	-0.0023	0.8781	0.9092	-0.0067	0.9325	0.9067	-0.0092	0.9741	0.9081	0.9282

Table 3: Experimental results of GPT-4 turbo on JEC-QA with RoSe strategy, where the "J" and "L" mean Judge and Lawyer in Role; "Rem" is the abbreviation for "Reminder" meaning "the answer is"; Cue means truth (t) or random (r) answer in prompt. **Bold** indicates the highest results, while underlined indicates the lowest results.

Role	Rem	Cue	step-1			step-2			step-3			overall	
			acc	Δ	conf	acc	Δ	conf	acc	Δ	conf	acc	conf
w/o	✗	✗	0.3336	-	**0.8489**	0.3364	-	**0.9203**	-	-	-	0.3349	**0.8846**
w/o	✗	t	0.6289	+0.2953	0.8130	0.6522	+0.3158	0.8910	0.5896	+0.2532	0.9454	0.6235	0.8831
w/o	✗	r	0.4083	+0.0746	0.8135	0.4061	+0.0697	0.8877	0.3906	+0.0542	0.9417	0.4016	0.8809
w/o	✓	t	**0.7367**	**+0.4031**	0.8041	**0.7659**	**+0.4295**	0.8871	**0.7713**	**+0.4349**	**0.9471**	**0.7579**	0.8794
w/o	✓	r	0.3852	+0.0515	0.8016	0.3710	+0.0346	0.8777	0.3468	+0.0104	0.9398	0.3676	0.8730
J	✓	t	0.5490	+0.2154	0.8132	0.5603	+0.2239	0.8858	0.6358	+0.2994	0.9432	0.5817	0.8807
J	✓	r	0.4097	+0.0761	0.8054	0.3952	+0.0588	0.8729	0.3544	+0.0180	0.9192	0.3864	0.8658
L	✓	t	0.4932	+0.1596	0.8057	0.4762	+0.1398	0.8819	0.5246	+0.1881	0.9388	0.4980	0.8754
L	✓	r	0.4034	+0.0697	0.8081	0.4034	+0.0670	0.8808	0.3877	+0.0513	0.9313	0.3981	0.8673

RQ1: Whether LLM knows what it knows? We can reveal this problem by comparing the performance under the guidance of different cue information along with reminder. Overall, in Table 2, it is evident that GPT-4 performs better at step-2 than at step-1 in diverse prompt settings, indicating that LLMs are capable of self-reflection and self-correction. However, the behavior of LLMs varies depending on the cue information provided. When employing ground-truth answers as cues, the performance of LLM are gradually better in the iteration steps shown in Tables 2 and 3. Conversely, in most settings with random answer, the accuracy rate of LLMs drops at step-3 when random answers are used as cues. This suggests that LLMs are confused by random answers, ***under the guidance of error information, LLMs fail to adhere to their own correct answer, exhibiting uncertainty on themselves***. Especially in domain-specific knowledge, when confronted with random answer, the performance of the LLM on JEC-QA decreases with each iteration step. However, the overall confidence levels of the LLMs are rising, despite the increasing uncertainty. With a similar finding, in Appendix B.4, we replace cue information with text content description corresponding to letter options, which is more subtle and the model is more easily affected by error information.

RQ2: What local information (shortcuts) in prompts does the LLM over-rely on? In Tables 2 and 3, it is clear that GPT-4 performs best without role guidance but with ground-truth answers as cues in a strong reminder scenario. Substituting the cue information with a random answer leads to a significant 9.58% and 35.15% decrease in performance for GPT-4, respectively. This large decrease is attributed to LLM capturing the strong reminder "answer is" within prompts. When substituting the cue information from truth to a random answer while keeping the role and reminder information the same, it is evident that the LLM's performance declines the most when there is a strong reminder without role guidance. This could be attributed to ***LLMs tend to capture shortcuts by relying solely on strong reminder "answer is" in prompts to quickly find the answer rather than understanding genuine relationships between prompt and truth during training***, potentially leading to blind trust in user instructions during real-world scenarios. More obviously, LLM exhibits over-reliance behavior on strong reminder in JEC-QA, and it is easy to rely on shortcuts when it is unfamiliar with domain knowledge.

Table 4: Experimental results of Spark-13B and fine-tuned Spark-13B on ID and OOD sets of EG-QA, where acc, conf, com denote accuracy, confidence, *com* respectively. Δ represents changes in LLM performance when cue information changes, the darker the color, the more the LLM is affected by random cue information.

		Role	Rem	Cue	step-1				step-2				step-3				overall	
					acc	Δ	conf	com	acc	Δ	conf	com	acc	Δ	conf	com	acc	com
ID	Spark	✗	✗	✗	0.6190	-	**0.9156**	0.1143	0.6741	-	0.9198	0.7984	-	-	-	-	0.6465	0.4563
		T	✓	t	0.8097	0.1197	0.9013	0.4027	0.7938	0.1669	0.9196	0.6749	0.8350	0.2670	0.8396	0.9100	0.8128	0.6625
		T	✓	r	0.6900		0.9010	0.3891	0.6269		0.9138	0.6062	0.5680		0.9469	0.7244	0.6283	0.5732
		C	✓	t	0.7578	0.0671	0.9085	0.4797	0.7532	0.1061	0.9266	0.6801	0.7480	0.1200	0.9516	0.8558	0.7530	0.6718
		C	✓	r	0.6907		0.9063	0.4610	0.6471		0.9258	0.6306	0.6280		0.9488	0.7714	0.6552	0.6210
	F-T Spark	✗	✗	✗	0.8061	-	0.8803	0.8804	0.8051	-	**0.9412**	0.8798	-	-	-	-	0.8056	0.8801
		T	✓	t	**0.8766**	0.0684	0.8781	**0.9342**	**0.8816**	0.0704	0.9373	**0.9370**	**0.8846**	0.0734	**0.9936**	**0.9387**	**0.8809**	**0.9366**
		T	✓	r	0.8083		0.8799	0.8940	0.8110		0.9380	0.8958	0.8120		0.9877	0.8967	0.8104	0.8955
		C	✓	t	0.8726	0.0644	0.8755	0.9319	0.8776	0.0664	0.9364	0.9348	0.8816	0.0694	0.9919	0.9370	0.8772	0.9345
		C	✓	r	0.8082		0.8799	0.8939	0.8112		0.9380	0.8957	0.8122		0.9877	0.8963	0.8105	0.8953
OOD	Spark	✗	✗	✗	0.4761	-	**0.9000**	0.0771	0.6351	-	0.9122	0.7685	-	-	-	-	0.5556	0.4228
		T	✓	t	0.7070	0.1555	0.8785	0.2569	0.7614	0.1771	0.9238	0.5371	**0.8500**	0.3340	0.9189	**0.9189**	0.7728	0.5709
		T	✓	r	0.5515		0.8916	0.2540	0.5843		0.9008	0.7008	0.5160		0.9262	0.6807	0.5506	0.5451
		C	✓	t	0.6343	0.0361	0.8974	0.3343	0.6892	0.1094	0.9025	0.5495	0.7330	0.1410	0.9328	0.8459	0.6855	0.5765
		C	✓	r	0.5982		0.8980	0.3364	0.5798		0.9242	0.5111	0.5890		0.9330	0.7413	0.5890	0.5296
	F-T Spark	✗	✗	✗	0.6752	-	0.8791	0.7951	0.6794	-	**0.9390**	0.7980	-	-	-	-	0.6773	0.7965
		T	✓	t	**0.7843**	0.1170	0.8749	**0.8791**	**0.7873**	0.1230	0.9358	**0.8809**	0.7963	0.1300	**0.9908**	0.8866	**0.7893**	**0.8822**
		T	✓	r	0.6673		0.8786	0.8004	0.6643		0.9371	0.7982	0.6663		0.9866	0.7997	0.6659	0.7994
		C	✓	t	0.7565	0.0802	0.8744	0.8613	0.7686	0.0893	0.9338	0.8691	0.7716	0.0923	0.9895	0.8710	0.7655	0.8671
		C	✓	r	0.6763		0.8747	0.8068	0.6793		0.9330	0.8090	0.6794		0.9820	0.8092	0.6783	0.8083

Table 5: Experimental results of Qwen-7B and fine-tuned Qwen-7B on ID and OOD sets of EG-QA, where acc, conf, com denote accuracy, confidence, *com* respectively. Δ represents changes in LLM performance when cue information changes.

		Role	Rem	Cue	step-1				step-2				step-3				overall	
					acc	Δ	conf	com	acc	Δ	conf	com	acc	Δ	conf	com	acc	com
ID	Qwen	✗	✗	✗	0.7774	-	0.8928	0.7062	0.7084	-	0.9517	0.6141	-	-	-	-	0.7429	0.6601
		T	✓	t	0.8222	0.1416	0.8885	0.8015	0.7942	0.1505	0.9295	0.7215	0.8247	0.1650	0.9618	0.8939	0.8137	0.8056
		T	✓	r	0.6806		0.8803	0.7185	0.6437		0.9200	0.6488	0.6597		0.9465	0.7833	0.6613	0.7168
		C	✓	t	0.7976	0.0541	0.8837	0.7866	0.7928	0.1328	0.9266	0.7435	0.7902	0.0722	0.9546	0.8701	0.7935	0.8000
		C	✓	r	0.7435		0.8735	0.7422	0.6600		0.9182	0.6579	0.7180		0.9529	0.8193	0.7071	0.7398
	F-T Qwen	✗	✗	✗	0.9120	-	0.8940	0.9539	0.9120	-	0.9471	0.9544	-	-	-	-	0.9120	0.9541
		T	✓	t	**0.9470**	0.0440	0.8938	**0.9727**	**0.9480**	0.0440	0.9469	**0.9733**	**0.9480**	0.0460	**0.9989**	**0.9733**	**0.9476**	0.9731
		T	✓	r	0.9030		**0.8951**	0.9490	0.9040		**0.9474**	0.9495	0.9020		0.9954	0.9484	0.9030	0.9489
		C	✓	t	0.9370	0.0271	0.8938	0.9674	0.9370	0.0271	0.9468	0.9674	0.9370	0.0281	0.9987	0.9674	0.9369	**0.9778**
		C	✓	r	0.9089		0.8943	0.9527	0.9099		0.9470	0.9528	0.9089		0.9956	0.9522	0.9092	0.9525
OOD	Qwen	✗	✗	✗	0.6981	-	0.8936	0.6471	0.6798	-	0.9316	0.5801	-	-	-	-	0.8156	0.6299
		T	✓	t	0.8457	0.1513	0.8644	0.8104	0.8291	0.1901	0.9013	0.7465	0.8534	0.2009	0.9439	0.9105	0.8427	0.8224
		T	✓	r	0.6944		0.8742	0.7342	0.6390		0.9073	0.6621	0.6525		0.9356	0.7801	0.6619	0.7254
		C	✓	t	0.8035	0.0867	0.8720	0.7961	0.8011	0.1502	0.9166	0.7465	0.8059	0.1036	0.9387	0.8799	0.8035	0.8075
		C	✓	r	0.7168		0.8750	0.7520	0.6509		0.9224	0.6534	0.7023		0.9375	0.8115	0.6900	0.7389
	F-T Qwen	✗	✗	✗	0.7830	-	0.8952	0.8782	0.7850	-	0.9476	0.8795	-	-	-	-	0.7840	0.8788
		T	✓	t	**0.8580**	0.0640	0.8958	**0.9235**	**0.8610**	0.0660	0.9477	**0.9253**	**0.8620**	0.0670	0.9982	**0.9258**	**0.8603**	**0.9248**
		T	✓	r	0.7940		**0.8961**	0.8851	0.7950		0.9478	0.8857	0.7950		0.9935	0.8857	0.7946	0.8855
		C	✓	t	0.8480	0.0560	0.8956	0.9177	0.8500	0.0590	0.9478	0.9189	0.8530	0.0620	**0.9984**	0.9206	0.8503	0.9190
		C	✓	r	0.7920		0.8959	0.8839	0.7910		**0.9479**	0.8833	0.7912		0.9955	0.8835	0.7914	0.8835

RQ3: Whether LLM can be affected by different role guidance? It can be observed that different roles have varying degrees of influence on the performance of LLMs. Taking the experimental results with no role and strong cue information as a reference, we can find that under role guidance, LLM is less affected by random answers, which indicates that role guidance can reduce the over-reliance of LLMs on reminders to a certain extent and focus on the real problem. Besides, it is obvious that when cue information is truth, GPT-4 guided by a teacher or judge performs better at step-3, ***indicating that LLMs tend to trust the role of authority more, similar to human behavior. It can also be interpreted as the shortcut on roles, or bias that LLM has learned during training.***

RQ4: Does the confidence level of LLMs change under role guidance? In the experimental results, it is first evident that the confidence of GPT-4 increases through reflection steps, while LLMs show overconfidence at step-3 under different strategies. Second, it is observed that LLMs exhibit the highest level of confidence in settings where shortcuts (reminders) are easy to capture, consistent with findings in deep neural models (Du et al., 2021; Zhao et al., 2024). Deep neural models tend to take shortcuts with high confidence. Notably, despite the high confidence levels, ***the overall confidence level of LLMs in settings with random cues is lower than that in settings with truth cues.*** This is consistent with their performance accuracy. When the model is more capable, there is a higher consistency between its verbalized confidence and performance. In addition, as shown in Table 3, ***the overall confidence of the LLMs at step-3 decreases under the guidance of different roles compared to the no-role guidance, which is similar to student performance and reflects their uncertainty.*** Compared to GPT-3.5, GPT-4's verbalized confidence more accurately reflects its uncertainty about responses to certain questions. For further details on how GPT-3.5 expresses its confidence values differently from GPT-4, please refer to Appendix A.2.

Table 6: Experimental results of LLaMA3-8B and fine-tuned LLaMA3-8B on ID and OOD sets of EG-QA, where acc, conf, com denote accuracy, confidence, *com* respectively. Δ represents changes in LLM performance when cue information changes.

		Role	Rem	Cue	step-1 acc	step-1 Δ	step-1 conf	step-1 com	step-2 acc	step-2 Δ	step-2 conf	step-2 com	step-3 acc	step-3 Δ	step-3 conf	step-3 com	overall acc	overall com
ID	LLaMA3	✗	✗	✗	0.5698	-	0.8753	0.6662	0.5628	-	0.9263	0.6146	-	-	-	-	0.5662	0.6404
		T	✓	truth	0.7062	0.1525	0.8275	0.7855	0.7375	0.1905	0.8869	0.7761	**0.8418**	0.3224	0.9358	0.8250	0.7618	0.7955
		T	✓	random	0.5537		0.8287	0.6784	0.5470		0.8901	0.6523	0.5194		0.9364	0.6289	0.5400	0.6532
		C	✓	truth	0.6095	0.0382	0.8442	0.7144	0.5896	0.0088	0.8956	0.6760	0.6139	0.0290	0.9363	0.6858	0.6043	0.6820
		C	✓	random	0.5713		0.8350	0.6940	0.5808		0.8965	0.6810	0.5849		0.9369	0.6782	0.5790	0.6844
	F-T LLaMA3	✗	✗	✗	**0.7917**	-	**0.8943**	**0.8808**	**0.7923**	-	**0.9508**	0.8826	-	-	-	-	**0.7920**	**0.8817**
		T	✓	truth	0.7710	0.1680	0.8859	0.8687	0.7725	0.1703	0.9412	0.8706	0.7765	0.1737	**0.9940**	**0.8732**	0.7733	0.7902
		T	✓	random	0.6030		0.8913	0.7499	0.6022		0.9433	0.7503	0.6028		0.9889	0.7507	0.6026	0.7503
		C	✓	truth	0.7565	0.1545	0.8902	0.8584	0.7592	0.1536	0.9429	0.8616	0.7602	0.1552	0.9935	0.8623	0.7586	0.8607
		C	✓	random	0.6022		0.8892	0.7507	0.6056		0.9405	0.7538	0.6050		0.9874	0.7541	0.6042	0.7528
OOD	LLaMA3	✗	✗	✗	0.5439		0.8731	0.6574	0.5428	-	0.9259	0.6163	-	-	-	-	0.5433	0.6368
		T	✓	truth	0.7131	0.1695	0.8292	0.7858	0.7232	0.1725	0.8864	0.7685	**0.8327**	0.3270	0.9409	0.8165	0.7563	0.7902
		T	✓	random	0.5436		0.8304	0.6729	0.5507		0.8810	0.6520	0.5057		0.9261	0.6157	0.5333	0.6468
		C	✓	truth	0.6042	0.0382	0.8441	0.7236	0.6194	0.0515	0.9023	0.7151	0.6345	0.0622	0.9401	0.7229	0.6193	0.7205
		C	✓	random	0.5660		0.8662	0.6910	0.5679		0.9227	0.6676	0.5723		0.9624	0.6655	0.5690	0.6747
	F-T LLaMA3	✗	✗	✗	0.6315	-	0.8895	0.7684	0.6311	-	**0.9474**	0.7714	-	-	-	-	0.6313	0.7699
		T	✓	truth	**0.7710**	0.1680	0.8859	**0.8687**	**0.7725**	0.1711	0.9412	0.8706	0.7765	0.1745	**0.9940**	**0.8731**	**0.7733**	**0.8708**
		T	✓	random	0.6030		**0.8913**	0.7499	0.6014		0.9433	0.7492	0.6020		0.9889	0.7492	0.6021	0.7494
		C	✓	truth	0.7565	0.1535	0.8902	0.8589	0.7590	0.1514	0.9429	0.8615	0.7600	0.1530	0.9935	0.8621	0.7585	0.8608
		C	✓	random	0.6036		0.8892	0.7500	0.6076		0.9405	0.7531	0.6070		0.9874	0.7532	0.6060	0.7521

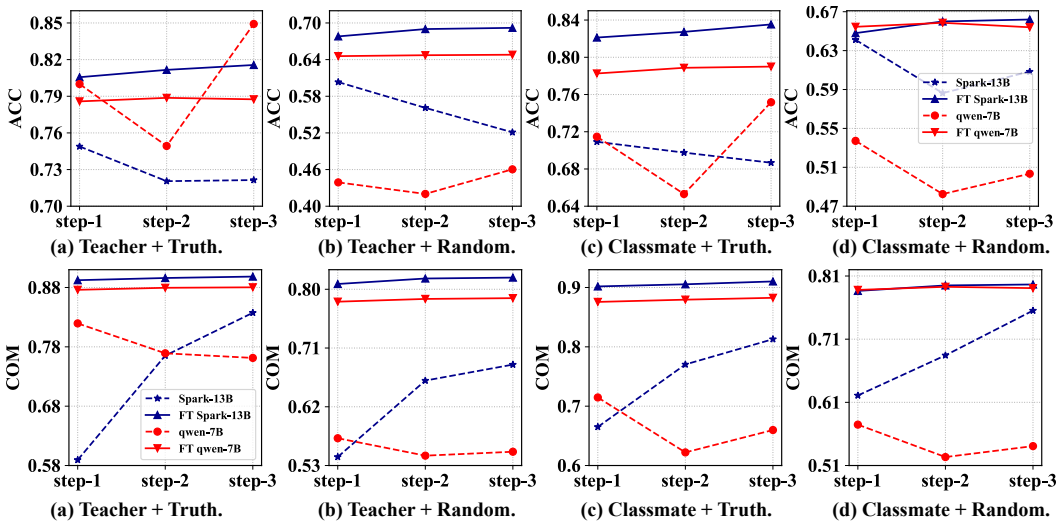

Figure 3: Experimental results (acc, *com*) of open-source LLMs and fine-tuned LLMs on open-BookQA test set under different role-guided and self-reflection.

5.3.2 FINE-TUNING

We present comparative experimental results of open-source LLMs on ID and OOD sets as shown in Table 4, 5 and 6. During experiments, we find that base LLMs usually cannot give a definite answer in step-1 and step-2, exhibiting task avoidance (Zhou et al., 2024). To make fair comparisons, we define a new metric as the comprehensive completion degree *com*, considering accuracy A and completion degree C[10] of LLMs. We adopt the variant of F1-scores as evaluation on *com*: $2 \times \frac{A \times C}{A+C}$.

On the ID and OOD sets, the fine-tuned LLMs trained with well-calibrated data performs well under all strategy settings with various role-guided and reminders, achieving a significant improvement in accuracy and completeness. Specifically, although only the role-guided data is employed for fine-tuning, the fine-tuned LLMs not only perform well in the two-step setup, but also have a strong ability to self-reflect. Besides, we list the Δ indicator to represent the change in the performance of LLMs when only cue information changes; the darker the color, the greater the changes.

We find that the fine-tuned LLMs are less affected by random information than the pre-fine-tuned LLMs overall. Fine-tuned LLMs are not confused by random answers and reduce their reliance on strong reminders, indicating the effectiveness of the double-calibrated strategy. In contrast, base LLMs are more susceptible to guidance from roles and cue information, leading to a significant decrease in performance due to the effect of random answers. In particular, LLaMA3-8B relies

[10]The completion degree C refers to the proportion of LLM that gives the exact answer.

heavily on cue information at step-3 in the prompt and is more affected by the teacher's guidance. The fine-tuned LLM can reduce its blind trust in the teacher.

Meanwhile, base LLMs with a smaller number of parameters also exhibit similar performance as GPT-4, as detailed in Section 5.3.1. In most settings, base LLMs perform best when provided with truth cue and worst when the cue is a random answer guided by teacher, indicating that open-source LLMs trust the authority (teacher) more than the closed-source GPT-4. Although fine-tuning cannot completely eliminate the bias from pre-trained knowledge, the double-calibrated strategy effectively mitigates the effects of role guidance.

On the openBookQA, given that fine-tuning a LLM might affect the LLM's ability in commonsense reasoning, we present the performance of Qwen-7B and Spark-13B on the Accuracy (ACC) and *com* (COM) metrics as shown in Figure 3. The experimental results of LLaMA3-8B are detailedly analyzed in Appendix C.1. Overall, fine-tuned LLMs generally outperform than base models in each step, with performance improving as steps increase, which indicates the commonsense reasoning ability of fine-tuned LLMs is not affected, and the ability to self-reflect is maintained. Qwen performs worse than Spark on openBookQA, despite outperforming Spark on EG-QA in general. Additionally, it can be found that the performance of base Spark declines with reflection steps, which could be attributed to insufficient smoothness in its indicator-ACC. Smaller LLMs, such as Qwen-7B, have weaker abilities in following instructions and often fail to provide exact answers in initial steps, similar to GPT-3.5's failure as shown in Appendix B.1. Considering the *com* metrics shown in the bottom row of Figure 3, both Spark-13B and fine-tuned versions perform better with increasing steps, demonstrating good performance after fine-tuning under different strategies.

Under wrong cue information, we report the performance of LLMs and their fine-tuned versions when the cue information is the wrong answer under different roles of guidance, as shown in Figure 4. On each dataset, the fine-tuned LLMs outperform the base LLMs. On the ID dataset, fine-tuned Spark-13B achieves 76% accuracy, while Qwen-7B achieves 81% accuracy. In contrast, the base LLMs are easily distracted by the wrong cue information in the prompts. In general, open-source LLMs are more easily affected by erroneous cue information under teacher guidance, indicating that they trust the authority role more than GPT-4 Turbo. Furthermore, fine-tuned Spark-13B and Qwen-7B perform better on EG-QA and openBookQA than LLaMA3-8B. Since LLaMA3-8B does not perform very well on bilingual corpora, fine-tuning it on EG-QA does not significantly alleviate the effect of cues, although its performance is improved.

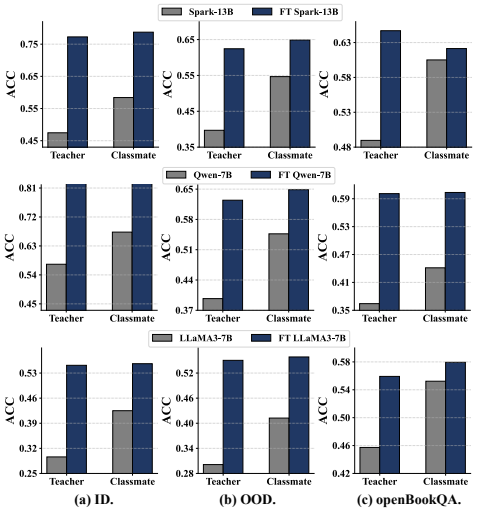

Figure 4: The performance of open-source LLMs and fine-tuned LLMs guided by wrong strong cue information.

6 Conclusion

In this paper, we proposed the **Ro**le-Guided and **Se**lf-Reflection (RoSe) strategy to fully evaluate whether LLMs know what they know and which specific features that LLMs rely on. We encouraged LLM to reflect and adjust its responses during self-reflection, and introduced multiple combinations of different roles and strong reminder information to make comprehensive assessments. Through series of evaluations on our collected EG-QA and publicly available JEC-QA in legal domain, we found that LLMs over-relied on strong reminder information in prompts; showed more trust in the authority role; guidance of roles could help LLMs to alleviate their reliance on local information to varying degrees and influenced the confidence of LLMs. Following these findings, we proposed the double-calibrated (accuracy and confidence) strategy to obtain well-calibrated data from powerful closed-source LLM, enabling fine-tune open-source LLMs. Extensive experiments of Spark-13B, Qwen-7B and LLaMA3-8B on our collected EG-QA and publicly available dataset openBookQA demonstrated the effectiveness of this strategy.

7 ACKNOWLEDGMENTS

This research was partially supported by grants from the National Natural Science Foundation of China (Grants No. 62337001), the Key Technologies R & D Program of Anhui Province (No. 202423k09020039), the Fundamental Research Funds for the Central Universities, and the Science and Technology Innovation 2030 Agenda (2022ZD0120200).

REFERENCES

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

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

Wei Chen, Lili Zhao, Zhi Zheng, Tong Xu, Yang Wang, and Enhong Chen. Double-checker: Large language model as a checker for few-shot named entity recognition. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 3172–3181, 2024.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240): 1–113, 2023.

Anthony G Cohn and Jose Hernandez-Orallo. Dialectical language model evaluation: An initial appraisal of the commonsense spatial reasoning abilities of llms. *arXiv preprint arXiv:2304.11164*, 2023.

Katherine M Collins, Albert Q Jiang, Simon Frieder, Lionel Wong, Miri Zilka, Umang Bhatt, Thomas Lukasiewicz, Yuhuai Wu, Joshua B Tenenbaum, William Hart, et al. Evaluating language models for mathematics through interactions. *Proceedings of the National Academy of Sciences*, 121(24):e2318124121, 2024.

Mengnan Du, Varun Manjunatha, Rajiv Jain, Ruchi Deshpande, Franck Dernoncourt, Jiuxiang Gu, Tong Sun, and Xia Hu. Towards interpreting and mitigating shortcut learning behavior of nlu models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 915–929, 2021.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. In *Forty-first International Conference on Machine Learning*, 2024.

Deep Ganguli, Amanda Askell, Nicholas Schiefer, Thomas I Liao, Kamilė Lukošiūtė, Anna Chen, Anna Goldie, Azalia Mirhoseini, Catherine Olsson, Danny Hernandez, et al. The capacity for moral self-correction in large language models. *arXiv preprint arXiv:2302.07459*, 2023.

Zhibin Gou, Zhihong Shao, Yeyun Gong, Yujiu Yang, Nan Duan, Weizhu Chen, et al. Critic: Large language models can self-correct with tool-interactive critiquing. In *The Twelfth International Conference on Learning Representations*, 2024.

Haixia Han, Jiaqing Liang, Jie Shi, Qianyu He, and Yanghua Xiao. Small language model can self-correct. In *Thirty-Eighth AAAI Conference on Artificial Intelligence*, pp. 18162–18170, 2024.

Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations*, 2022.

Jiaxin Huang, Shixiang Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. Large language models can self-improve. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 1051–1068, 2023a.

Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large language models cannot self-correct reasoning yet. In *The Twelfth International Conference on Learning Representations*, 2023b.

Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38, 2023a.

Ziwei Ji, Tiezheng Yu, Yan Xu, Nayeon Lee, Etsuko Ishii, and Pascale Fung. Towards mitigating LLM hallucination via self reflection. In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pp. 1827–1843, 2023b.

Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. How can we know when language models know? on the calibration of language models for question answering. *Transactions of the Association for Computational Linguistics*, 9:962–977, 2021.

Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*, 2022.

Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. Large language models struggle to learn long-tail knowledge. In *International Conference on Machine Learning*, pp. 15696–15707, 2023.

Byoungjip Kim, Youngsoo Jang, Lajanugen Logeswaran, Geon-Hyeong Kim, Yu Jin Kim, Honglak Lee, and Moontae Lee. Prospector: Improving llm agents with self-asking and trajectory ranking. In *NeurIPS 2023 Foundation Models for Decision Making Workshop*, 2023a.

Geunwoo Kim, Pierre Baldi, and Stephen McAleer. Language models can solve computer tasks. *Advances in Neural Information Processing Systems*, 36:39648–39677, 2023b.

Stefanie Krause and Frieder Stolzenburg. Commonsense reasoning and explainable artificial intelligence using large language models. In *European Conference on Artificial Intelligence*, pp. 302–319, 2023.

Junyi Li, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. Halueval: A large-scale hallucination evaluation benchmark for large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 6449–6464, 2023.

Xiaonan Li and Xipeng Qiu. Mot: Memory-of-thought enables chatgpt to self-improve. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 6354–6374, 2023.

Xun Liang, Shichao Song, Simin Niu, Zhiyu Li, Feiyu Xiong, Bo Tang, Yezhaohui Wang, Dawei He, Cheng Peng, Zhonghao Wang, and Haiying Deng. UHGEval: Benchmarking the hallucination of Chinese large language models via unconstrained generation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2024.

Heng Liao, Jiajin Tu, Jing Xia, Hu Liu, Xiping Zhou, Honghui Yuan, and Yuxing Hu. Ascend: a scalable and unified architecture for ubiquitous deep neural network computing: Industry track paper. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 789–801, 2021.

Stephanie Lin, Jacob Hilton, and Owain Evans. Teaching models to express their uncertainty in words. *Transactions on Machine Learning Research*, 2022, 2022.

Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173, 2024.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations*, 2019.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36:46534–46594, 2023.

Herbert W Marsh. Causal ordering of academic self-concept and academic achievement: a multi-wave, longitudinal panel analysis. *Journal of educational psychology*, 82(4):646, 1990.

Sabrina J Mielke, Arthur Szlam, Emily Dinan, and Y-Lan Boureau. Reducing conversational agents' overconfidence through linguistic calibration. *Transactions of the Association for Computational Linguistics*, 10:857–872, 2022.

Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2381–2391, 2018.

Matthias Minderer, Josip Djolonga, Rob Romijnders, Frances Hubis, Xiaohua Zhai, Neil Houlsby, Dustin Tran, and Mario Lucic. Revisiting the calibration of modern neural networks. *Advances in Neural Information Processing Systems*, 34:15682–15694, 2021.

Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 29, 2015.

Xuefei Ning, Zinan Lin, Zixuan Zhou, Zifu Wang, Huazhong Yang, and Yu Wang. Skeleton-of-thought: Large language models can do parallel decoding. In *The Twelfth International Conference on Learning Representations*, 2023.

OpenAI. Introduce chatgpt. In *OpenAI blog*, 2023.

Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H Chi, Nathanael Schärli, and Denny Zhou. Large language models can be easily distracted by irrelevant context. In *International Conference on Machine Learning*, pp. 31210–31227, 2023.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36:8634–8652, 2023.

Ruixiang Tang, Dehan Kong, Longtao Huang, and Hui Xue. Large language models can be lazy learners: Analyze shortcuts in in-context learning. In *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 4645–4657, 2023.

Meta LLaMA Team. Introducing meta llama 3: The most capable openly available llm to date. *https://ai.meta.com/blog/meta-llama-3/*, 2024.

Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher D Manning. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 5433–5442, 2023.

Boshi Wang, Xiang Yue, and Huan Sun. Can chatgpt defend its belief in truth? evaluating llm reasoning via debate. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 11865–11881, 2023a.

Cunxiang Wang, Sirui Cheng, Qipeng Guo, Yuanhao Yue, Bowen Ding, Zhikun Xu, Yidong Wang, Xiangkun Hu, Zheng Zhang, and Yue Zhang. Evaluating open-qa evaluation. *Advances in Neural Information Processing Systems*, 36:77013–77042, 2023b.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.

Fangwen Wu, Yi Jiang, Diyue Liu, Evgenia Konorova, and Xiangdong Yang. The role of perceived teacher and peer relationships in adolescent students' academic motivation and educational outcomes. *Educational Psychology*, 42(4):439–458, 2022.

Miao Xiong, Zhiyuan Hu, Xinyang Lu, YIFEI LI, Jie Fu, Junxian He, and Bryan Hooi. Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms. In *The Twelfth International Conference on Learning Representations*, 2023.

Derong Xu, Wei Chen, Wenjun Peng, Chao Zhang, Tong Xu, Xiangyu Zhao, Xian Wu, Yefeng Zheng, Yang Wang, and Enhong Chen. Large language models for generative information extraction: A survey. *Frontiers of Computer Science*, 18(6):186357, 2024.

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems*, 36:11809–11822, 2023.

Zhangyue Yin, Qiushi Sun, Cheng Chang, Qipeng Guo, Junqi Dai, Xuan-Jing Huang, and Xipeng Qiu. Exchange-of-thought: Enhancing large language model capabilities through cross-model communication. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 15135–15153, 2023a.

Zhangyue Yin, Qiushi Sun, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Xuan-Jing Huang. Do large language models know what they don't know? In *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 8653–8665, 2023b.

Yu Yuan, Lili Zhao, Kai Zhang, Guangting Zheng, and Qi Liu. Do llms overcome shortcut learning? an evaluation of shortcut challenges in large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 12188–12200, 2024.

Zhongshen Zeng, Yinhong Liu, Yingjia Wan, Jingyao Li, Pengguang Chen, Jianbo Dai, Yuxuan Yao, Rongwu Xu, Zehan Qi, Wanru Zhao, et al. Mr-ben: A comprehensive meta-reasoning benchmark for large language models. *arXiv preprint arXiv:2406.13975*, 2024.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. Siren's song in the ai ocean: A survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*, 2023.

Lili Zhao, Qi Liu, Linan Yue, Wei Chen, Liyi Chen, Ruijun Sun, and Chao Song. Comi: Correct and mitigate shortcut learning behavior in deep neural networks. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 218–228, 2024.

Zirui Zhao, Wee Sun Lee, and David Hsu. Large language models as commonsense knowledge for large-scale task planning. *Advances in Neural Information Processing Systems*, 36:31967–31987, 2023.

Haoxi Zhong, Chaojun Xiao, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. Jec-qa: A legal-domain question answering dataset. In *Proceedings of AAAI*, 2020.

Lexin Zhou, Wout Schellaert, Fernando Martínez-Plumed, Yael Moros-Daval, Cèsar Ferri, and José Hernández-Orallo. Larger and more instructable language models become less reliable. *Nature*, pp. 1–8, 2024.

CONTENTS

A Experimental Details **15**

A.1 Prompt Settings 15

A.2 Verbalized Confidence 16

B More Experiments on Evaluation **16**

B.1 Experimental Results on GPT-3.5 16

B.2 Whether More Analysis Distracts GPT-4 from Focusing on Cues? 17

B.3 Evaluations on Consistency of Reasoning Steps 18

B.4 Whether Subtler Cue Information Distracts GPT-4? 19

B.5 Calibration Analysis 20

C More Experiments on Fine-tuning **21**

C.1 LLaMA3-8B on openBookQA 21

C.2 Strong Reminder without Role-Guidance on EG-QA 22

C.3 Shortcut Learning in Fine-Tuning Process 23

D Limitations **23**

A EXPERIMENTAL DETAILS

In the evaluation, we use GPT-3.5 turbo and GPT-4 turbo to fully assess the performance of LLMs under different strategies. We call them by the OpenAI official API[11], employ GPT-3.5 turbo-1106 and the latest GPT-4 turbo-0409 to make evaluations on object clause data of EG-QA about 1,645 samples under 9 prompt settings. Due to the professional knowledge of legal domain, we only employ GPT-4 turbo, which is more capable, to make fully evaluations on 3,180 samples from Knowledge-Driven questions (KD-questions) on JEC-QA under 9 prompt settings.

A.1 PROMPT SETTINGS

In addition to the prompt settings of step-1 and step2 shown in Section 5.2, the specific prompt settings in step-3 as shown in Table 7. The format of response is as follows:

The answer format is step: {*step*}*, analysis:* {*analysis*}*, answer:* {*answer*}*, confidence:* {*confidence*}.

Since we find that role guidance can help LLMs reduce their reliance on reminders (LLMs perform better in the random answer setting with role guidance than those without role guidance), we fine-tune the open-source LLMs by selecting the settings with role guidance and reminder in the double-calibrated strategy.

Furthermore, we study the statistics of token consumption to answer a question in a prompt setting as presented in Table 8. In general, we conduct prompt settings for GPT-4 and GPT-3.5 on five knowledge point datasets respectively, and the total consumption was approximately average tokens * dataset size * 2 * 9.

[11] https://platform.openai.com/docs/api-reference

Table 7: Detailed prompt settings in step-3.

Role	Rem	Cue	Prompt settings in step-3
teacher	✓	truth	My teacher thinks the answer is {Truth}. Please read the questions and options carefully, continue to think, reflect on the answer of step 2, and give the most appropriate answer and confidence.
teacher	✓	random	My teacher thinks the answer is {Random}. Please read the questions and options carefully, continue to think, reflect on the answer of step 2, and give the most appropriate answer and confidence.
classmate	✓	truth	My classmate thinks the answer is {Truth}. Please read the questions and options carefully, continue to think, reflect on the answer of step 2, and give the most appropriate answer and confidence.
classmate	✓	random	My classmate thinks the answer is {Random}. Please read the questions and options carefully, continue to think, reflect on the answer of step 2, and give the most appropriate answer and confidence.
w/o	✓	truth	The answer is {Truth}. Please read the questions and options carefully, continue to think, reflect on the answer of step 2, and give the most appropriate answer and confidence.
w/o	✓	random	The answer is {Random}. Please read the questions and options carefully, continue to think, reflect on the answer of step 2, and give the most appropriate answer and confidence.
w/o	✗	truth	{Truth}. Please read the questions and options carefully, continue to think, reflect on the answer of step 2, and give the most appropriate answer and confidence.
w/o	✗	random	{Random}. Please read the questions and options carefully, continue to think, reflect on the answer of step 2, and give the most appropriate answer and confidence.
w/o	✗	✗	-

Table 8: The statistics of token consumption to answer a question in a setting.

Knowledge	Avg. # Input Token	Avg. # Output Token	Avg. # Total
prepositions	219.33	254.97	474.31
verbs	225.71	270.50	496.27
nouns	226.46	260.18	486.64
adjectives	226.04	274.43	500.47
object clauses	219.3	245.72	537.21

A.2 Verbalized Confidence

In addition to probability percentages on verbalized confidence, GPT-4 also outputs non-numerical confidence levels in few cases as shown in Figure 5. Despite "high", "very high", "extremely high", GPT-4 also outputs confidence words like "confirmed", "medium", "medium to high", "highest", "supremely high". If verbalized confidence is not expressed as a numerical value, confidence does not produce a progressive relationship like the numerical scores with the deepening of the reflection steps. Therefore, it is difficult to directly quantify it as a numerical value, which is not only unfair in numerical statistics but also fails in reflecting the confidence level of LLM directly. Considering the few cases (10-20%), we only compute on numerical confidence levels.

Besides, the instruction following ability of GPT-3.5 is not as strong as GPT-4, exhibiting overconfident on verbalized confidence. As shown in Figure 6, GPT-3.5 often fails to give the exact answer in the first or second step, and in this case, GPT-3.5 provide the correct answer in the second step, but change it to the wrong answer after the third step of reflection with confidence of "higher". This means that the ability of LLM is directly proportional to its ability on expressing confidence, and when focusing on ability, we should also focus on the level of verbalized confidence of LLMs.

B More Experiments on Evaluation

B.1 Experimental Results on GPT-3.5

We list the experimental results of GPT-3.5 on EG-QA as shown in Table 9. First, the instruction-following ability of GPT-3.5 is much worse than GPT-4, and GPT-3.5 only outputs the confidence levels at the last step in most scenarios. Second, it should be noted that GPT-3.5 performs the worst under random cue information when there is no reminder overall, while performs the best when guided by teacher with truth answer in prompt. It indicates that GPT-3.5 is more sensitive with cue information. These cues, however, do not explicitly indicate that they are the answer, but they have a large impact on the GPT-3.5. It shows that when LLM is weaker, it is easier to focus on meaningless words, and these cues will distract the model. On the whole, the level of verbalized confidence of

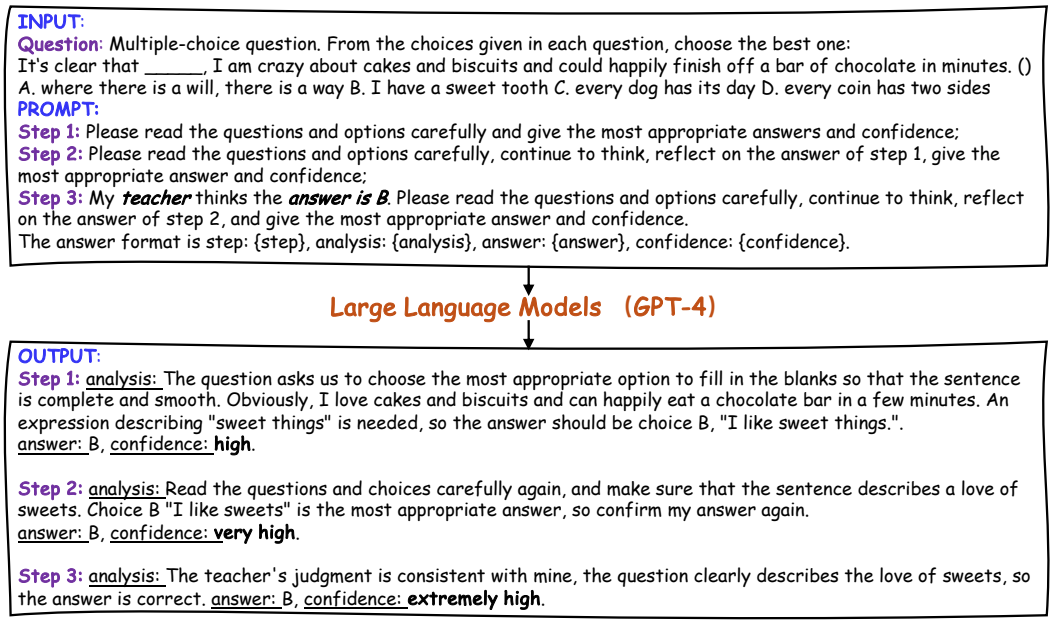

Figure 5: We propose RoSe strategy in the prompts, where role, strong reminder, cue are represented in *italics*. In this case, GPT-4 outputs "high", "very high", "extremely high" to represent confidence levels at each step.

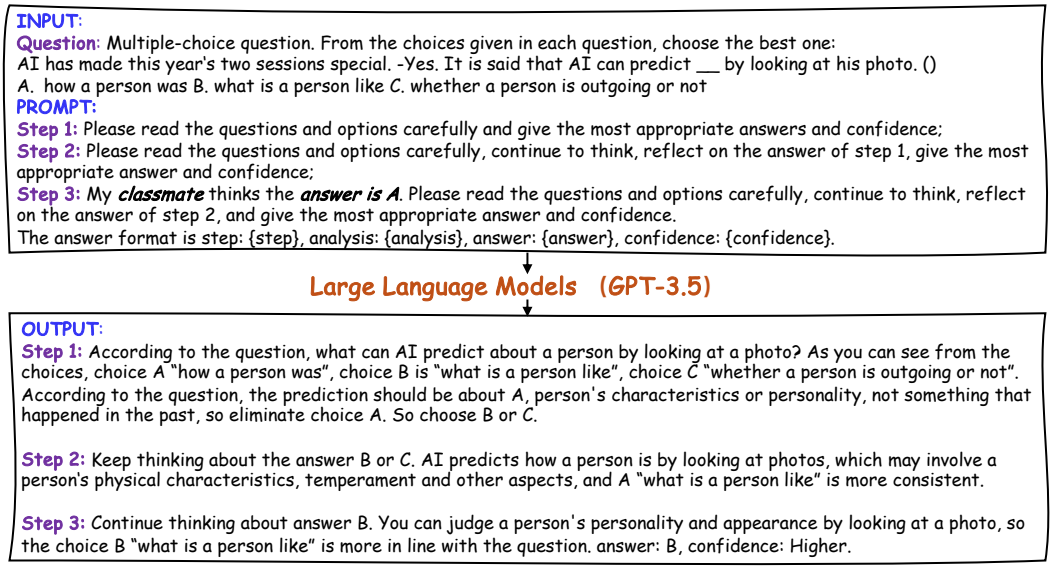

Figure 6: In this case, GPT-3.5 only outputs "higher" confidence in the last step, and exhibit poor ability on self-reflection and self-correction.

GPT-3.5 is too high to represent the real uncertainty of GPT-3.5 itself. When the LLM's ability is stronger, the stronger the LLM ability, the stronger the verbalized confidence expression ability.

B.2 Whether More Analysis Distracts GPT-4 from Focusing on Cues?

As shown in Table 2, GPT-4 is confused by error information (cue of random answer), causing its performance to deteriorate. We want to know if adding manually annotated analysis would distract the LLM from focusing on cues. Due to the fact that it is impossible to obtain numerous annotated problem-solving processes in real scenarios, we annotate the 1,645 samples adopted in the

Table 9: Experimental results on GPT-3.5 turbo with RoSe strategy. **Bold** indicates the highest results, while underlined indicates the lowest results.

Role	Rem	Cue	step-1		step-2		step-3		overall	
			acc	conf	acc	conf	acc	conf	acc	conf
w/o	✗	✗	0.5661	-	0.5980	0.9456	-	-	0.5821	0.9456
w/o	✓	truth	0.5107	-	0.7166	-	0.8059	0.9694	0.6777	0.9694
w/o	✓	random	0.4871	-	0.5243	-	0.4935	0.9577	0.5017	0.9577
w/o	✗	truth	0.4182	-	0.6786	-	0.8157	0.9641	0.6375	0.9641
w/o	✗	random	0.4255	-	0.5848	-	0.4464	0.9348	0.4856	0.9348
teacher	✓	truth	**0.6425**	-	**0.7541**	-	**0.8200**	**0.9706**	**0.7389**	**0.9706**
teacher	✓	random	0.5038	-	0.5475	-	0.5282	0.9539	0.5265	0.9539
classmate	✓	truth	0.6035	-	0.7057	-	0.7629	0.9578	0.6907	0.9578
classmate	✓	random	0.5180	-	0.6091	-	0.6064	0.9591	0.5778	0.9591

Table 10: Experimental results on GPT-4 turbo with RoSe strategy. The strong reminder information contains reminder, cue, and human annotation analysis. **Bold** denotes the best performance.

Role	Rem	Cue	step-1		step-2		step-3		overall	
			acc	conf	acc	conf	acc	conf	acc	conf
teacher	✓	random	0.9070	0.8752	0.9108	0.9302	0.9101	0.9716	0.9093	0.9257
teacher	✓	random+a	**0.9796**	0.8750	**0.9796**	0.9350	0.9788	0.9827	0.9793	0.9788

evaluation, and the experimental results are shown in Table 10. The overall performance of GPT-4 improves after adding the annotated analysis, indicating that the model deeply understand the content in the prompts. However, in the third step, GPT-4 is still affected by the random answers and performs slightly worse.

Besides, as shown in Figure 7, the question is the same as Figure 1. In Figure 1, GPT-4 gives the wrong answers. However, after adding the analysis annotation, GPT-4 could output the right answer and not affected by the random cue information. Meanwhile, the analysis process of LLM becomes detailed and the level of confidence is less overconfident than before.

B.3 Evaluations on Consistency of Reasoning Steps

To further evaluate the internal consistency of the LLM in the self-reflection process, we first employ GPT-4 and human annotation to evaluate the internal consistency between the reasoning steps of the model on the challenging samples. We regard the samples in which the LLM make mistakes under the two-step reflection as challenging samples[12]. There are about 8% of the data that are challenging samples that are not known to LLM. We utilize the prompt in Table 11 and human annotation to evaluate the reasoning consistency of LLM in three steps. The performance of these challenging samples under the RoSe strategy are shown in Table 12.

Specifically, the consistency between step-1 and 2 shows little difference between GPT-4 and human annotations. However, on the consistency between step-2 and 3, the human annotations demonstrate higher consistency. Although the logical expression from step-2 to step-3 is consistent, GPT-4 annotations tend to focus more on semantic consistency, often overlooking the progression of logical expression. Overall, the logical reasoning in the self-reflection process across the three steps is consistent for LLMs.

Then, in *Guidance* of Table 12, we manually annotate whether the responses of these challenging samples in step-3 follow the role guidance information. Consistent with findings in **RQ3 of Section 5.3.1**, LLMs tend to trust the role of authority more and more easily affected by authority-teacher.

[12]Among them, the model makes mistakes in both steps, or it succeeds in the step-1 but fails in the step-2 after self-reflection.

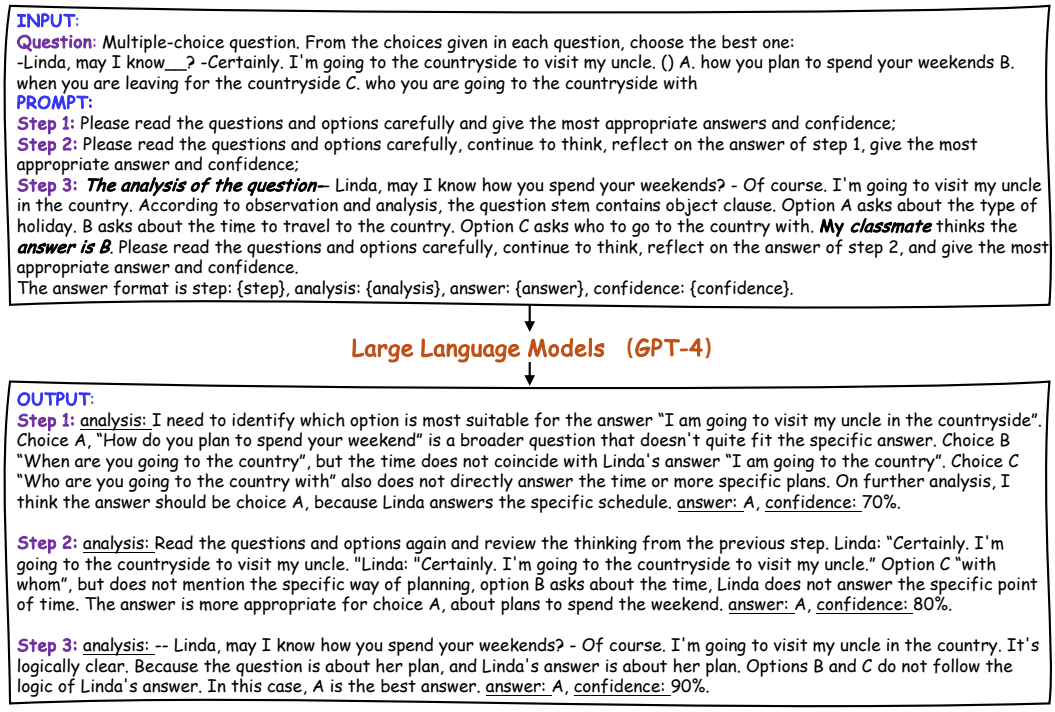

Figure 7: In this case, we add the human annotation analysis of the question to the prompts, GPT-4 provides correct reasoning and answer.

Table 11: Detailed prompt setting for consistency evaluation.

Prompt	Given a question, determine the relationship between the following three steps in answers. If reasoning steps are consistent, answer "True"; if not, answer "False". "1" indicates the relationship between step-1 and step-2, and "2" indicates step-2 and step-3. Question: {question} step-1: {reasoning step-1} step-2: {reasoning step-2} step-3: {reasoning step-3} The answer format is: {"1": "True/False", "2": "True/False"}

B.4 Whether Subtler Cue Information Distracts GPT-4?

We treat multiple-choice QA as the task of model evaluation and set fixed seeds to ensure stability and consistency of evaluation. In multiple-choice questions, choices are usually identified by a letter (e.g., A, B, C, D). Each letter corresponds to a specific text content or answer description. Considering that prompt information may be more subtle in real-world scenarios, the letter options in prompts are substituted into their corresponding textual descriptions. This transformation makes the prompt information more complex and requires deeper understanding and processing by LLMs.

The experimental results under subtler cue information are shown in Table 13. Under the influence of subtle cue information, the overall performance of GPT-4 is lower than that under letter cue information, and the overall conclusions of the experimental results are consistent with the findings in Section 5.3.1. Compared to letter options, LLM is less sensitive to the text cue information, and it is difficult to associate it with the option content in the question. Since it cannot distinguish between relevant and misleading information, it causes LLM to become distracted and unable to reason and answer questions effectively.

Table 12: Experimental results of GPT-4 turbo on EG-QA with RoSe strategy, where the "T" and "C" mean Teacher and Classmate in Roles; Reminder means "the answer is"; Cue means truth (t) or random (r) answer in prompt. Con_{GPT} and Con_{human} stand for GPT-4 and human evaluation of the consistency on the reasoning process, respectively. $Guidance$ represents whether the responses of these challenging samples in step-3 follow the role guidance information.

Role	Rem	Cue	step-1&2		step-2&3		step-3
			Con_{GPT}	Con_{human}	Con_{GPT}	Con_{human}	$Guidance$
T	✓	t	0.9548	0.9473	0.7669	0.9248	0.4210
T	✓	r	0.9545	0.9772	0.7954	**0.9924**	**0.4318**
C	✓	t	**0.9923**	**0.9923**	0.8778	0.9618	0.3816
C	✓	r	0.9236	0.9312	**0.8854**	0.9923	0.2213

Table 13: Experimental results of GPT-4 turbo on EG-QA with RoSe strategy, where the "T" and "C" mean Teacher and Classmate in Roles; Reminder means "the answer is"; Cue means truth (t) or random (r) answer in prompt, we replace the option cue (e.g. "*A*") with the text content (t_c, r_c) of option (e.g. "*how you plan to spend your weekends*"). **Bold** indicates the highest results, while underlined indicates the lowest results. We evaluate the performance of LLMs by accuracy ("acc") and confidence scores ("conf") jointly.

Role	Rem	Cue	step-1			step-2			step-3			overall	
			acc	Δ	conf	acc	Δ	conf	acc	Δ	conf	acc	conf
w/o	✗	✗	0.9108	-	**0.8889**	0.9159	-	**0.9676**	-	-	-	0.9134	0.9283
T	✓	t	**0.9431**	+0.0323	0.8726	**0.9450**	+0.0291	0.9295	**0.9494**	0.0334	0.9825	**0.9458**	0.9282
T	✓	r	0.9070	-0.0038	0.8752	0.9108	-0.0051	0.9302	0.9101	-0.0058	0.9716	0.9093	0.9257
C	✓	t	0.9322	+0.0214	0.8717	0.9335	+0.0176	0.9287	0.9373	+0.0213	0.9785	0.9343	0.9263
C	✓	r	0.9085	-0.0023	0.8781	0.9092	-0.0067	0.9325	0.9067	-0.0092	0.9741	0.9081	0.9282
T	✓	t_c	**0.9390**	+0.0282	0.8796	**0.9433**	+0.0274	0.9328	**0.9457**	+0.0298	0.9715	**0.9427**	0.9062
T	✓	r_c	0.8700	-0.0408	0.8810	0.8688	-0.0471	0.9312	0.8639	-0.0520	0.9623	0.8676	0.9061
C	✓	t_c	0.9194	+0.0086	0.8852	0.9219	+0.0060	0.9365	0.9225	+0.0066	0.9698	0.9212	0.9109
C	✓	r_c	0.8741	-0.0367	0.8885	0.8796	-0.0363	0.9372	0.8778	-0.0371	0.9679	0.8772	0.9128

B.5 Calibration Analysis

To better compare the calibration abilities of LLMs in the evaluation and after fine-tuning, we employ calibration plots and ECE scores to perform evaluations on GPT-4 turbo and fine-tuned LLaMA3-8B. First, the calibration plots of GPT-4 are shown in Figure 8. The closer the performance is to the perfectly calibrated line, the better the model is calibrated.

Overall, the confidence levels of the LLM are high, typically above 60%, with most values falling to the lower right side of the perfectly calibrated line. As shown in the top row, GPT-4 demonstrates good calibration performance at step-1, but its performance declines at step-3 under role guidance. Meanwhile, the calibration performance of the LLM at step-3 is worse under random answer guidance, which is consistent with the findings of **RQ1** in Section 5.3.1. Furthermore, as shown in the bottom row of Figure 8, compared to the model calibration performance under RoSe strategy, LLM exhibits poorer ability without role guidance, which aligns with the findings in **RQ3** of Section 5.3.1. LLMs tend to capture shortcuts by relying solely on strong reminder in prompts to quickly find the answer, role guidance can reduce the over-reliance of LLMs on reminders to a certain extent.

Then, we employ Expected Calibration Error (ECE) (Naeini et al., 2015) by comparing the confidence scores with the actual accuracy of the predictions for evaluation. A lower ECE score indicates better calibration, implying that the LLM's predicted confidence align more closely with the actual accuracy of predictions.

As shown in Figures 9 and 10, the fine-tuned LLaMA3-8B and Spark-13B generally have better calibration abilities than base LLMs. The calibration abilities of LLMs at step-2 are worse than the other two steps, indicating that LLMs modify their correct answers at step-1 and improve their confidence scores during the self-reflection process, leading to their higher ECE scores. This also demonstrates the significance of evaluating whether LLMs know what they know in Section 5.3.1, where LLMs fail to adhere to their own correct answer, exhibiting uncertainty but accompanied by

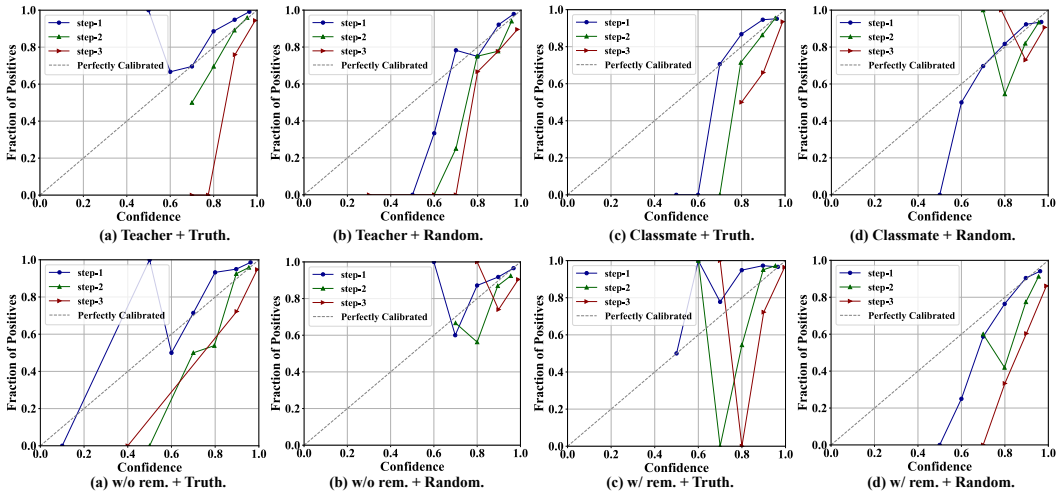

Figure 8: Calibration performance of GPT-4 on EG-QA. The top row shows experimental results under the RoSe strategy, while the bottom row compares different strategies without role guidance. The closer the performance is to the perfectly calibrated line, the better the model is calibrated.

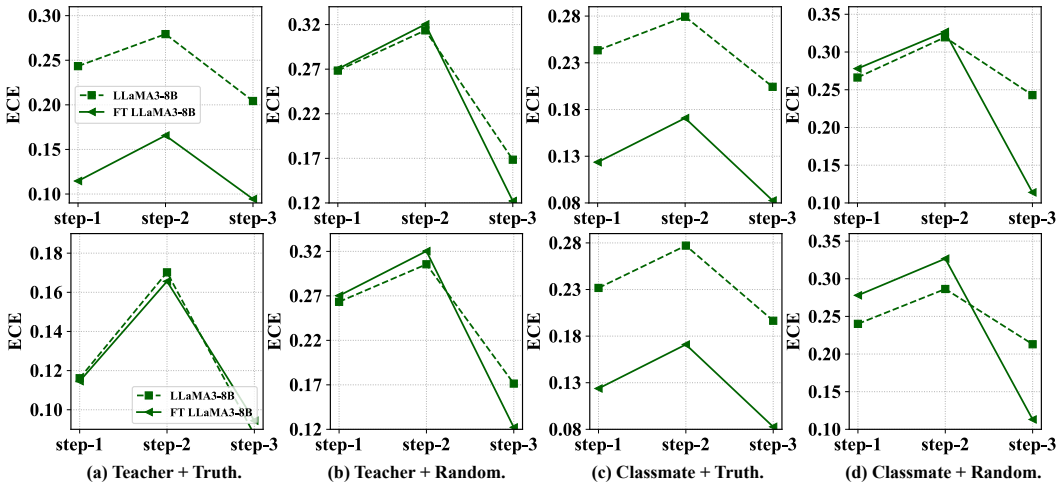

Figure 9: Calibration performance of LLaMA3-8B and fine-tuned LLaMA3-8B on the ID and OOD sets of EG-QA, where the top row and bottom row represent the model performance on ID and OOD sets, respectively. A lower ECE score indicates better calibration.

rising self-confidence. Moreover, under the guidance of classmate and random cue information, the fine-tuned LLMs exhibit poor calibration ability compared with base LLMs at step-1 and step-2, there appears to be a discrepancy between the models' high confidence levels and their actual accuracy. After two steps of reflection and adjustment, there is a notable enhancement in the models' performance at step-3, which suggests the alignment of increased accuracy and confidence levels.

C More Experiments on Fine-tuning

C.1 LLaMA3-8B on openBookQA

The experimental results of LLaMA3-8B and fine-tuned LLaMA3-8B are shown in Figure 11. It is intuitive to see that the LLaMA3-8B performs well at step-1, which indicates that the base model can handle such common sense reasoning problems. However, LLaMA3-8B is easily affected by cue information in prompts, and the performance at step-3 changes with the change of cue infor-

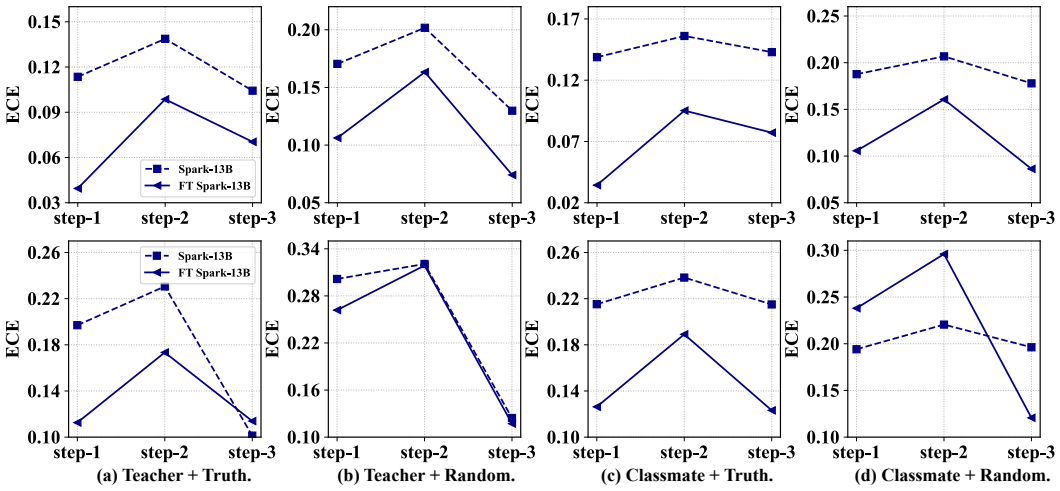

Figure 10: Calibration performance of Spark-13B and fine-tuned Spark-13B on the ID and OOD sets of EG-QA, where the top row and bottom row represent the model performance on ID and OOD sets, respectively. A lower ECE score indicates better calibration.

mation. When the cue information corresponds to the true answer, the model demonstrates optimal performance, whereas performs poorly when the cue is a random answer.

The performance of the fine-tuned LLaMA3-8B is stable and not easily affected by the cue in prompts. However, its performance decreases at step-3 through self-reflection, indicating that the model's uncertainty about itself increases in the process of repeated reflection, which leads to the change of the response. Meanwhile, it also reflects the problem "whether LLMs know what they know" we mentioned in Section 4.1, revealing that LLM with fewer parameters lacks this ability and still needs to be improved continuously.

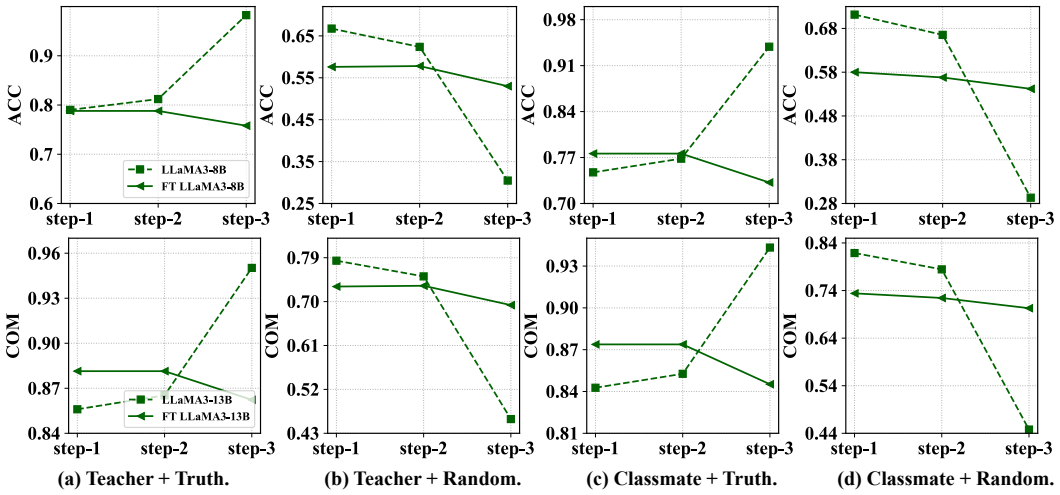

Figure 11: Experimental results (acc, com) of LLaMA3-8B and fine-tuned LLaMA3-8B on open-BookQA test set under different role-guidance and self-reflection.

C.2 Strong Reminder without Role-Guidance on EG-QA

The experimental results of Spark-13B, LLaMA3-8B and fine-tuned versions under strong reminder without role-guided are listed in Table 14 and 15. Although the well-calibrated data only contains role-guided settings, the LLMs still achieve improvement in effect without role-guide, indicating that role-guide can reduce the model's excessive focus to strong reminder. Besides, similar to the performance of GPT-4 in Section 4.1, the open-source LLMs also show the phenomenon of relying

Table 14: Experimental results of Spark-13B and fine-tuned Spark-13B on ID and OOD sets of EG-QA without role-guidance but with strong reminder, where acc, conf, com denote accuracy, confidence, *com* respectively.

		Role	Rem	Cue	step-1			step-2			step-3			overall	
					acc	conf	com	acc	conf	com	acc	conf	com	acc	com
ID	Spark	✗	✓	t	0.6126	**0.9128**	0.4037	0.7635	0.9321	0.6776	**0.9459**	0.9645	**0.9539**	0.7740	0.6784
		✗	✓	r	0.5286	0.9066	0.3803	0.5292	0.9325	0.5619	0.3941	0.9614	0.5586	0.4840	0.5003
	F-T	✗	✓	t	**0.8915**	0.8791	**0.9368**	**0.8918**	**0.9373**	**0.9379**	0.9023	**0.9913**	0.9455	**0.8952**	**0.9401**
	Spark	✗	✓	r	0.7888	0.8779	0.8780	0.7866	0.9361	0.8763	0.7868	0.9862	0.8768	0.7874	0.8770
OOD	Spark	✗	✓	t	0.4405	**0.9156**	0.2996	0.6580	0.9250	0.5429	**0.9481**	0.9447	**0.9560**	0.6822	0.5995
		✗	✓	r	0.3471	0.9151	0.2852	0.4699	0.9166	0.4764	0.4264	0.9453	0.5915	0.4145	0.4510
	F-T	✗	✓	t	**0.8068**	0.8781	**0.8887**	**0.8111**	**0.9367**	**0.8917**	0.8240	**0.9892**	0.8990	**0.8140**	**0.8931**
	Spark	✗	✓	r	0.6807	0.8776	0.8077	0.6814	0.9362	0.8079	0.6851	0.9852	0.8111	0.6824	0.8089

Table 15: Experimental results of LLaMA3-8B and fine-tuned LLaMA3-8B on ID and OOD sets of EG-QA without role-guidance but with strong reminder, where acc, conf, com denote accuracy, confidence, *com* respectively.

		Role	Rem	Cue	step-1			step-2			step-3			overall	
					acc	conf	com	acc	conf	com	acc	conf	com	acc	com
ID	LLaMA3	✗	✓	t	0.6378	0.8205	0.8462	0.8536	0.8826	0.8467	0.9308	0.9327	0.8819	0.8074	0.8583
		✗	✓	r	0.4414	0.8175	0.5950	0.4492	0.8787	0.5915	0.3830	0.9321	0.5297	0.4245	0.5721
	F-T	✗	✓	t	**0.9408**	**0.8837**	**0.9680**	**0.9435**	**0.9406**	**0.9004**	**0.9447**	**0.9969**	**0.9010**	**0.9430**	**0.9231**
	LLaMA3	✗	✓	r	0.7005	0.8802	0.8214	0.6952	0.9379	0.7623	0.6858	0.9920	0.7554	0.6938	0.7797
OOD	LLaMA3	✗	✓	t	0.8238	0.8419	0.8638	0.8621	0.8929	0.8625	**0.9243**	0.9412	**0.8904**	0.8701	0.8722
		✗	✓	r	0.4384	0.8304	0.5898	0.4199	0.8870	0.5631	0.3863	0.9348	0.5310	0.4149	0.5613
	F-T	✗	✓	t	**0.8781**	0.8776	**0.9316**	**0.8870**	**0.9371**	**0.8795**	0.8993	**0.9939**	0.8857	**0.8881**	**0.8989**
	LLaMA3	✗	✓	r	0.5891	**0.8785**	0.7405	0.5968	0.9367	0.6876	0.5854	0.9900	0.6789	0.5904	0.7023

on strong reminder, especially the performance of LLMs at step-3, which are greatly affected by strong reminder.

C.3 Shortcut Learning in Fine-Tuning Process

We employ well-calibrated data which is only teacher-guided with truth answer as cue information to fine-tune Spark-13B. We adopt 400 samples in the remaining difficult data as the test set, namely the questions that GPT-4 answers incorrectly. The experimental results with RoSe strategy is shown in Table 16. Obviously, Spark-13B after fine-tuning does not perform well under multiple strategies, which only learns the strong cue information (ground truth answer) in the prompt to make shortcut learning, rather than learning the CoT process in formulating an answer, i.e., logical thinking. In the process of fine-tuning, it is not only necessary to ensure the diversity of data, but also to avoid the frequent occurrence of certain features.

D Limitations

Multiple reminders in different scenarios. Due to the diverse range of tasks that LLMs are capable of handling, they exhibit a variety of shortcut learning behaviors in practical applications and rely on different "reminder". These behaviors are jointly determined by the training data and paradigm. The large volume of training data for LLMs makes it difficult to eliminate hidden biases within the data. The gradient descent paradigm enables LLMs to quickly identify common features among the data, which may be shortcut features rather than robust ones as expected. In addition to the strong reminder "answer is" in our experimental settings, multiple reminder captured by LLMs during training and fine-tuning processes in different Question-Answering scenarios still require further exploration and elimination.

Verbalized confidence of open-source LLMs. Experimental results show that both closed-source and open-source LLMs often display overconfidence, with few instances where an LLM expresses a verbalized confidence level below 80%. With improved capabilities and command-following abilities, the verbalized confidence of LLMs is better expressed and aligns with accuracy metrics. However, the overconfidence exhibited by open-source LLMs hinders their ability to truly understand their knowledge, limiting their overall capability. Further research is needed to explore the trade-off

Table 16: Experimental results of Spark-13B and fine-tuned Spark-13B with different RoSe strategy, the teacher-guided with truth answer in well-calibrated data is adopted to fine-tune LLM.

Role	Rem	cue	Spark-13B	F-T Spark-13B
teacher	✓	truth	0.6975	0.9900
teacher	✓	random	0.2700	0.2750
classmate	✓	truth	0.5080	**0.9935**
classmate	✓	random	0.3050	0.2750

between verbal confidence and ability in open-source LLMs. Enhancing the verbalized confidence of open-source LLM can provide another perspective for improving its overall capability.