

CausalAbstain: Enhancing Multilingual LLMs with Causal Reasoning for Trustworthy Abstention

Yuxi Sun^{1*}, Aoqi Zuo^{2*}, Wei Gao³, Jing Ma^{1†}

¹ Department of Computer Science, Hong Kong Baptist University

² School of Mathematics and Statistics, The University of Melbourne

³ School of Computing and Information Systems, Singapore Management University

csyxsun@comp.hkbu.edu.hk, azuo@student.unimelb.edu.au

weigao@smu.edu.sg, majing@comp.hkbu.edu.hk

Abstract

Large Language Models (LLMs) often exhibit knowledge disparities across languages. Encouraging LLMs to *abstain* when faced with knowledge gaps is a promising strategy to reduce hallucinations in multilingual settings. Current abstention strategies for multilingual scenarios primarily rely on generating feedback in various languages using LLMs and performing self-reflection. However, these methods can be adversely impacted by inaccuracies and biases in the generated feedback. To address this, from a causal perspective, we introduce *CausalAbstain*, a method that helps LLMs determine whether to utilize multiple generated feedback responses and how to identify the most useful ones. Extensive experiments demonstrate that *CausalAbstain* effectively selects helpful feedback and enhances abstention decisions with interpretability in both native language (CASUAL-NATIVE) and multilingual (CAUSAL-MULTI) settings, outperforming strong baselines on two benchmark datasets covering encyclopedic and commonsense knowledge QA tasks. Our code and data are open-sourced at <https://github.com/peachch/CausalAbstain>.

1 Introduction

Large language models (LLMs) demonstrate impressive capabilities in encoding vast amounts of information and supporting knowledge-intensive tasks (Petroni et al., 2021; Yang et al., 2024a; Yu et al., 2024). However, hallucinations and bias (Mishra et al., 2024; Kumar et al., 2023; Ji et al., 2023) can arise when knowledge is missing or inaccurate, posing challenges to the reliability of LLMs (Feng et al., 2024b). A promising approach, which has inspired recent research, is to teach LLMs to *abstain* — *avoid incorrect answers in low-confidence scenarios to mitigate hallucina-*

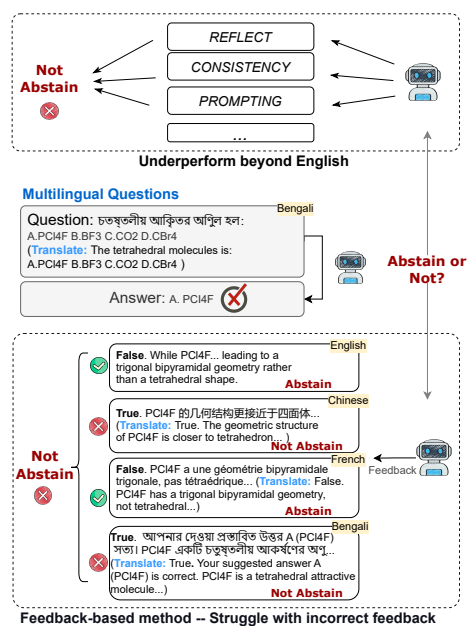


Figure 1: Prior approaches often underperform beyond English, particularly in low-resource languages. While the feedback-based method is well-suited for multilingual queries, some generated feedback may be incorrect (e.g., Chinese and Bengali), leading to incorrect decisions, i.e., failing to abstain from a wrong answer (A). The English translations (labeled as *Translate*:) are provided for reading.

tions and factual inaccuracies (Madhusudhan et al., 2024; Feng et al., 2024c,c).

Existing studies have explored abstention strategies in English (Varshney and Baral, 2023; Feng et al., 2024c; Yang et al., 2024b), but the ability of LLMs to abstain in multilingual settings remains largely underexplored (Feng et al., 2024b). The factual accuracy of multilingual LLMs, particularly in low-resource languages, is often poorer (Zhang et al., 2023a; Kang et al., 2024), limiting their global applicability and leaving many regions underserved. Therefore, developing robust abstention strategies across languages is essential to enhance

*Equal contribution.

†Corresponding author.

ing the reliability and inclusivity of LLMs.

Current *training-based* (Azaria and Mitchell, 2023; Slobodkin et al., 2023) and *calibration-based* (Zhou et al., 2024; Liu et al., 2024b; Tian et al., 2023) strategies require a hold-out dataset for hyperparameter tuning and determining confidence score thresholds, which may hinder generalization across different knowledge domains. *Prompt-based* abstention strategies (Kadavath et al., 2022; Edunov et al., 2018), such as prompting LLMs to evaluate their own answers before abstaining or generating additional knowledge prior to answering, have been studied. Similarly, *self-consistency* method (Wang et al., 2023a) encourages LLMs to reflect on their responses cooperatively. These approaches have been shown to underperform in low-resource languages, demonstrated by Feng et al. (2024b). To address this limitation, they propose a *feedback-based method* to teach LLMs to abstain by generating and reflecting on relevant feedback, making it more adaptable to multiple languages.

However, all of these methods still raise concerns because: 1) They rely solely on LLMs evaluating their own generated texts, which can be affected by hallucinations and potential bias (Xie et al., 2024; Ji et al., 2023); 2) The generated information or feedback may be irrelevant or unreliable, particularly in different language resources, leading to negative impacts. Figure 1 illustrates this issue: When LLM reviews its previous answer and provides feedback in different languages, including native language¹, some feedback is of low quality or even incorrect in certain languages, ultimately influencing the final abstention decision (e.g., the Chinese feedback in the figure). This raises a crucial research question: *How can we determine whether to utilize LLM-generated feedback and how to identify the helpful ones from multiple feedback sources?*

Inspired by the ability of the Structural Causal Model (SCM) (Pearl et al., 2000) to measure path-specific causal effects, we propose a causal approach to assess the impact of feedback on a model’s proposed answer. This allows us to determine whether to utilize the generated feedback for answer refinement and to identify the most helpful feedback to make a final abstention decision. Specifically, we probe a multilingual LLM to provide feedback on its proposed answer in both its

native language (CASUAL-NATIVE) and multiple related languages² (CASUAL-MULTI). By generating feedback over multiple runs, the model can produce diverse knowledge and varying assessments of the answer. We then apply causal inference to evaluate the causal effect of the generated feedback, determining whether it meaningfully enhances the abstention decision. Additionally, since different languages have varying degrees of representation in LLM pre-training data — some being severely underrepresented (Lai et al., 2023b), we propose candidate aggregation to come up with a more robust multilingual abstention strategy, which employs a voting mechanism to ensure a reliable abstention decision across multiple languages.

Our contributions are mainly three-fold:

- We present the first study on causal abstention in multilingual LLMs and propose *CausalAbstain*, a training-free approach that helps multilingual LLMs identify incorrect or biased feedback and abstain accordingly.
- Our method integrates feedback from both the native language and multiple related languages (i.e., CASUAL-NATIVE and CASUAL-MULTI), leveraging path-specific causal effect measurement to guide feedback selection.
- We evaluate CASUAL-MULTI and CASUAL-NATIVE on two datasets covering encyclopedic and commonsense knowledge QA. Experimental results demonstrate that our approach outperforms state-of-the-art baselines in abstention effectiveness.

2 Causality with Feedback

In this section, we introduce the fundamental concepts of causal inference, providing the background that supports *CausalAbstain* in §3.

2.1 Causal Graph

Causal inference is typically performed using the Structural Causal Model (Pearl et al., 2000). A key component of SCM is the causal graph, a directed acyclic graph (DAG) that represents causal relationships between variables. We denote the causal graph as $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where \mathcal{V} represents the set of variables in the graph and \mathcal{E} denotes the set of causal relationships between them. To address

¹The same language as the original answer is referred to as “native language”.

²Following the definition of language relatedness by Sun et al. (2021), which considers cultural, geographical, and typological factors.

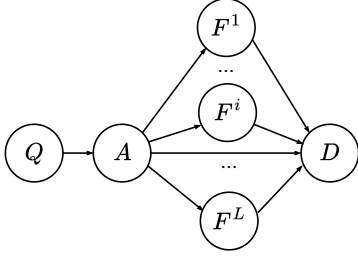


Figure 2: The causal graph of the question-answering problem with feedback.

the question answering with feedback problem, we construct a specific causal graph in Figure 2.

In the abstention task, the question (Q) and the originally proposed answer (A) together form the input to the LLM, which then generates the feedback F^i in language i . In Figure 2, a directed edge from one variable to another indicates a causal relationship, meaning the former influences the latter. The *likelihood distribution of the final abstention decision*, denoted as D , is influenced by both A and the generated feedback F^i . In the absence of feedback, D is directly decided by A . Additionally, since multilingual feedback can affect D , multiple causal paths may exist from A to D , mediated by different feedback variables in $\{F^i | 1 \leq i \leq L\}$, where L represents the number of languages considered for feedback³.

2.2 Causal Effect

Based on the SCM, various causal effects can be measured to quantify the impact of a treatment on an outcome. In Figure 2, A serves as the treatment for D . There are two distinct types of effects between A and the final decision D : the natural direct effect (NDE) (e.g., $A \rightarrow D$) and the indirect effect mediated by the generated feedback (e.g., $A \rightarrow F^i \rightarrow D$). The total effect (TE) of A on D can be decomposed into NDE and total indirect effect (TIE), as follows:

$$\text{TE} = \text{NDE} + \text{TIE}. \quad (1)$$

If we denote the TIE with respect to each feedback F^i as TIE^i , then

$$\text{TIE} = \sum_{i=1}^L \text{TIE}^i. \quad (2)$$

³We do not model a direct causal link from Q to F^i , as the question alone cannot trigger feedback on the answer—abstention feedback is only meaningful in the context of the answer. The same reasoning applies to the absence of a direct edge from Q to D .

NDE is computed by fixing the feedback mediations and comparing the potential outcomes with and without A :

$$\text{NDE} = \mathbb{E}[D(Q, A) - D_0], \quad (3)$$

where \mathbb{E} represents the expectation operator, and $D(Q, A)$ represents the final decision based solely on the original answer given the question, meaning it corresponds exactly to the original answer. D_0 serves as the baseline for D , which is typically modeled as a binomial distribution with a binomial probability 0.5. The indirect effect mediated by feedback F^i is computed by contrasting the potential outcomes with and without the feedback F^i , given the presence of A :

$$\text{TIE}^i = \mathbb{E}[D(Q, A, F^i) - D(Q, A)]. \quad (4)$$

3 CausalAbstain

Ideally, feedback can provide valuable insights to improve an answer to a question. However, not all feedback is helpful. As illustrated in Figure 1, feedback may sometimes be less informative, introduce erroneous information, or even be incorrect, which can negatively impact the abstention performance. To address this challenge, we propose a causality-based abstention mechanism to determine whether to incorporate feedback and how to effectively utilize it in deriving the final decision.

3.1 Task Formulation

We focus on assisting multilingual LLMs in abstaining during question answering (Feng et al., 2024c). The LLM is first prompted with a given query Q , which may be in different languages, and provides the answer $A = \text{LLM}(Q)$. The LLM is then tasked with providing an abstention decision based on the proposed answer and query: $f(Q, A) \rightarrow \{Abstain, Not Abstain\}$. For our setting, the prompt to elicit LLMs’ abstention decision directly is: “Please review the correctness of proposed answer *True* or *False* directly.” The LLM is expected to *Abstain* when it is likely to provide an incorrect answer and *Not Abstain* when it is expected to answer correctly.

When incorporating multilingual feedback F^i , where $F^i = \text{LLM}(Q, A|i)$ represents the feedback generated in language i , the feedback is elicited using the prompt (Feng et al., 2024b): “Please review the proposed answer and provide a paragraph of feedback on its correctness. Feedback should be in

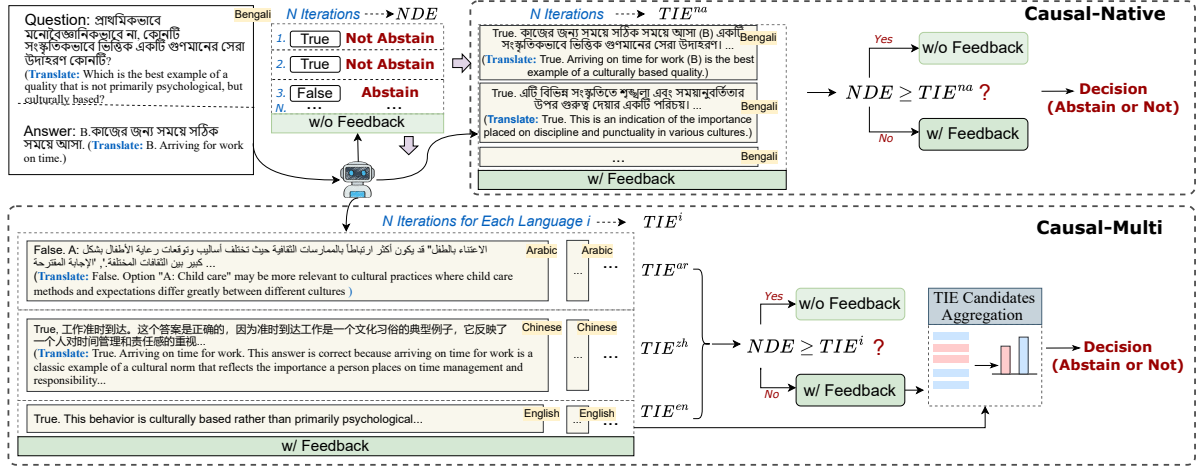


Figure 3: The framework of *CausalAbstain*. We propose an abstention strategy to determine whether to adhere to or abstain from the proposed answer (i.e., the answer B) based on the feedback in the native language (CAUSAL-NATIVE) and related languages (CAUSAL-MULTI), by leveraging natural direct effect (NDE) and total indirect effect (TIE) to assess the impact of feedback on the final abstention decision. We provide the English translations (Translate:) for reading.

language i ". The abstention decision is then made based on the multilingual feedback, with the decision function updated to $f(Q, A, F^1, F^2 \dots F^L) \rightarrow \{Abstain, Not\ Abstain\}$. Specifically, for the feedback in each language, the prompt is: "Based on the feedback for measuring the correctness of the answer, is the proposed answer *True* or *False*?"

Our goal is to develop a robust abstention strategy f by leveraging causality. The framework of *CausalAbstain* is illustrated in Figure 3, which consists of methods for incorporating native-language (i.e., CAUSAL-NATIVE) and multilingual (i.e., CAUSAL-MULTI) feedback, presented in §3.2 and §3.3, respectively.

3.2 CAUSAL-NATIVE

For the native language, we propose CAUSAL-NATIVE to determine whether to adhere to the proposed answer or abstain from it based on the monolingual feedback. The causal graph for our problem setting follows Figure 2, where we consider two specific types of causal relationships between the question Q , the original answer A , and the final decision D : 1) NDE implied by $A \rightarrow D$. It quantifies the direct impact of A on D in the absence of feedback. 2) TIE^{na} implied by causal path $A \rightarrow F^{na} \rightarrow D$, which is mediated by feedback F^{na} in the native language. It captures the extent to which feedback contributes to the overall causal effect of A on the final decision D .

Given the general Equations (3) and (4), we use the Jensen-Shannon Divergence (JSD) in practice

to compare the different outcome distributions in these equations. For the expectation operation, given each query-answer pair, we iteratively generate N feedback instances in the native language to approximate F^{na} , denoted as \hat{F}^{na} . We then use it to compute $\hat{D}(Q, A, \hat{F}^{na})$. Additionally, we prompt without feedback in N iterations to approximate $\hat{D}(Q, A)$. For the j -th feedback instance, we denote the corresponding decision as $\hat{D}_j(Q, A, \hat{F}_j^{na})$. All feedback are thus collected as $\hat{D}(Q, A, \hat{F}^{na}) = \{\hat{D}_j(Q, A, \hat{F}_j^{na}) | 1 \leq j \leq N\}$. Therefore, the practical computation of NDE and TIE^{na} is formulated as follows:

$$NDE = JSD(\hat{D}(Q, A), D_0), \quad (5)$$

$$TIE^{na} = JSD(\hat{D}(Q, A, \hat{F}^{na}), \hat{D}(Q, A)). \quad (6)$$

For any given example, our causal abstention strategy f determines whether to incorporate feedback based on the relationship between NDE and TIE^{na} , as illustrated in Figure 3. If $NDE \geq TIE^{na}$, the indirect effect is minimal, indicating that the feedback has little to no impact on improving the response. In this case, the abstention decision is made without considering the feedback and thus majority voted by $\hat{D}(Q, A)$. Conversely, if $NDE < TIE^{na}$, the indirect effect is substantial, suggesting that the feedback plays a significant role in refining the response. In this case, the abstention decision is made considering the feedback and is voted by majority of $\hat{D}(Q, A, \hat{F}^{na})$ over N native feedback instances.

3.3 CAUSAL-MULTI

Previous research has demonstrated that multiple feedback items in related languages can help LLMs identify knowledge gaps across diverse languages, cultures, and communities (Feng et al., 2024b). However, the challenge remains in effectively determining whether to utilize multilingual feedback. To address this, we formulate the abstention strategy CAUSAL-MULTI for multilingual settings.

The calculation of NDE and TIEⁱ for feedback in language i follows the same approach as described in §3.2. And the abstention strategy f is: If $\text{NDE} \geq \text{TIE}^i$ for feedback in any languages i , it implies that the feedback can be disregarded, and the original answer can be used as the final answer — similar to the native language scenario. Conversely, if there exists any case where $\text{NDE} < \text{TIE}^i$ in some language i , it suggests that the feedback could play a crucial role, such as identifying key nuances or providing knowledge that significantly enhances the response.

Considering the inconsistent representation of different languages, especially low-resource ones in the pre-training data of LLMs (Lai et al., 2023b), we propose a robust multilingual abstention strategy to mitigate potential biases introduced by feedback in a specific language. To achieve this, we introduce an aggregated voting mechanism across feedback from all languages. Specifically, we consider L languages and denote the abstention decision provided by the feedback in language i as \hat{D}^i , where $\hat{D}^i = \{\hat{D}_1^i, \hat{D}_2^i, \dots, \hat{D}_N^i\}$ and N is the number of feedback instances. The final decision is derived from the set $\{\hat{D}^1, \hat{D}^2, \dots, \hat{D}^L\}$, with a majority voting crossing over the $N \times L$ choices to determine whether to abstain or not.

4 Evaluation

4.1 Experimental Settings

Datasets & Languages. We evaluate *CausalAbstain* using the Hellaswag (M-Hellaswag) (Lai et al., 2023b) and Multilingual MMLU (M-MMLU) datasets, which focus on QA using general and commonsense knowledge in multiple languages. Following Lai et al. (2023a,b), we categorize languages based on their data ratios in the pre-training corpus. We randomly sample 500 instances for testing and 200 for validation from high-/medium-/low-resource languages. Additionally, we adopt the related language settings from Feng et al. (2024b) and set the number of iterations to

$N = 3$. The analysis of iteration times and language relatedness is provided in Appendix A. Further details on the datasets and language lists are described in Appendix B.

LLMs. We conduct experiments with three LLMs: ChatGPT and GPT-4o (two general black-box LLMs with strong multilingual capabilities (Bang et al., 2023; OpenAI, 2023)); Aya-13B (a multilingual open-source model).

Baselines We compare *CausalAbstain* against existing approaches that are adaptable to multiple languages: *Calibration-based*: ASK CALI (Tian et al., 2023); *Prompting-based*: REFLECT and MOREINFO (Kadavath et al., 2022; Feng et al., 2024a); *Consistency-based*: CONFLICT (Feng et al., 2024c); and *Feedback-based*: MULTI-RELATED (Feng et al., 2024b).

Metrics. We follow the Abstain Accuracy metrics, defined as $\frac{TP+TN}{TP+TN+FP+FN}$ proposed by Feng et al. (2024c), which evaluates whether the abstention decisions are correct. That is, an LLM should abstain when it would produce an incorrect answer and should not abstain when it would give a correct answer. The term $TP+TN$ represents cases where the LLM makes correct abstention decisions, including 1) the answer is correct, and the model does not abstain; 2) the answer is incorrect, and the model abstains.

Answer Likelihood Distribution. Since LLMs are black-box models (Gat et al., 2023; Cheng et al., 2024), directly conducting causal inference with parameterization of D is challenging. Thus, for each query-answer pair, we repeatedly generate N samples as a representative dataset to approximate distribution D . The likelihood distribution \hat{D} can be modeled as a binomial distribution, where the probability of $\hat{D} = 1$ is computed using a softmax function applied to the average indicator values, \bar{I}_1 and \bar{I}_0 , of \hat{D}_j , representing the decision corresponding to the j -th feedback: $P(\hat{D} = 1) = \frac{\exp(\bar{I}_1)}{\exp(\bar{I}_1) + \exp(\bar{I}_0)}$, where $\bar{I}_1 = \frac{1}{N} \sum_{j=1}^N \mathbb{I}(\hat{D}_j = 1)$, $\bar{I}_0 = \frac{1}{N} \sum_{j=1}^N \mathbb{I}(\hat{D}_j = 0)$, and \mathbb{I} is the indicator function, while N is both the number of iterations and the sample size.

4.2 Experimental Results

In Table 1, we present the Abstain Accuracy results for three LLMs evaluated on two multilingual datasets, M-MMLU and M-Hellaswag.

Method	M-MMLU									M-Hellaswag								
	zh	it	ar	id	bn	te	ne	kn	Overall	zh	it	ar	id	bn	te	ne	kn	Overall
AYA-13B																		
ASK CALI.	<u>.552</u>	.534	.593	.547	.587	.424	.477	.539	.532	.511	.504	.520	.541	.481	.538	.498	.496	.511
REFLECT	.478	.474	.469	.463	.370	.395	.320	.418	.423	.481	.483	.507	.481	.463	.502	.472	.509	.477
MOREINFO	.500	.479	.500	.435	.348	.429	.353	.411	.431	.489	.457	.533	.519	.506	.552	.429	.452	.489
CONFLICT	.470	.547	.456	.461	.550	.497	.550	.500	.504	.485	<u>.538</u>	.515	.511	.472	.484	.498	.496	.501
MULTI-RELATED	.555	<u>.525</u>	<u>.560</u>	.510	.555	<u>.524</u>	.555	.609	.549	.516	.517	.511	.455	.580	<u>.557</u>	.494	.592	.525
CAUSAL-NATIVE	.482	<u>.538</u>	.510	.475	.550	.497	.475	.507	.504	.593	.534	<u>.580</u>	<u>.613</u>	.679	.457	<u>.552</u>	<u>.657</u>	<u>.564</u>
CAUSAL-MULTI	.512	<u>.538</u>	.530	<u>.535</u>	<u>.574</u>	.544	<u>.545</u>	<u>.579</u>	<u>.545</u>	<u>.588</u>	.564	.589	.626	<u>.670</u>	.574	.597	.685	.612
GPT-4o																		
ASK CALI.	.208	.205	.293	.208	.462	.413	.419	.531	.343	.377	.192	.336	.265	.351	.524	.395	.487	.366
REFLECT	.503	.526	.503	.552	.462	.465	.493	.527	.504	.279	.256	.368	.381	.383	.513	.413	.526	.390
MOREINFO	.768	<u>.833</u>	.682	.779	.613	.597	<u>.669</u>	.543	.686	.617	.660	.704	.671	.589	.551	.516	.441	.594
CONFLICT	.665	.692	.623	.734	.503	.519	.487	.565	.599	.416	.417	.507	.471	.610	.585	.561	.599	.521
MULTI-RELATED	<u>.770</u>	.850	<u>.740</u>	<u>.800</u>	.720	.605	.675	<u>.639</u>	<u>.725</u>	.766	.769	.776	.761	<u>.747</u>	.626	.703	.691	.730
CAUSAL-NATIVE	.805	.756	.760	.805	<u>.737</u>	.610	.608	.623	.713	.792	<u>.833</u>	.730	<u>.800</u>	.746	.721	<u>.741</u>	.631	<u>.749</u>
CAUSAL-MULTI	.765	.830	.760	.795	.740	.660	.675	.671	.738	<u>.772</u>	.846	<u>.743</u>	.806	.753	.700	.748	.664	.754
CHATGPT																		
ASK CALI.	.440	.385	.535	.427	.618	.622	<u>.568</u>	.575	.522	.455	.500	.590	.558	.537	.570	.631	.675	.565
REFLECT	.543	.457	.470	.566	.516	<u>.588</u>	.550	.534	.528	.560	.468	.498	.498	.498	.628	.592	<u>.592</u>	.542
MOREINFO	.560	.611	.451	.534	.300	.401	.320	.390	.446	.524	.517	.445	.476	.403	.385	.343	.386	.435
CONFLICT	.522	.534	.522	.530	.608	.562	.599	<u>.549</u>	<u>.553</u>	.567	.641	<u>.594</u>	.588	.597	<u>.611</u>	<u>.597</u>	.561	.595
MULTI-RELATED	.565	.625	.520	.570	.505	.490	.485	.404	.521	<u>.571</u>	<u>.645</u>	.585	.571	<u>.567</u>	.529	.455	.461	.548
CAUSAL-NATIVE	<u>.570</u>	.590	.480	.600	.493	.565	.455	.503	.532	.558	.568	.585	<u>.592</u>	.424	.515	.394	.425	.508
CAUSAL-MULTI	.575	<u>.620</u>	<u>.523</u>	<u>.595</u>	.570	.581	.568	.500	.567	.593	.658	.641	.631	.558	.547	.519	.513	<u>.583</u>

Table 1: Performance comparison of *CausalAbstain* in native language (CAUSAL-NATIVE) and multi-language (CAUSAL-MULTI) settings, against the calibration, prompting, consistency, and feedback-based baselines on two datasets. We present the performance of the Abstain Accuracy metric from high-/medium-resource languages (Chinese, Italian, Arabic, and Indonesian) to low-resource languages (Bengali, Telugu, Nepali, and Kannada). Overall denotes average performance for all languages. The best and second-best results are highlighted in **bold** and underline, respectively.

CAUSAL-MULTI achieves state-of-the-art performance. Our proposed CAUSAL-MULTI outperforms the the strongest baseline in 4 out of 6 settings (across three models and two datasets), achieving an average accuracy improvement of 3.5% over the best-competing method. Across the 8 different resource languages, CAUSAL-MULTI rank the first in 3.3 languages and second in 2.5 languages on average. In contrast, MULTI-RELATED is a method that leverages multiple feedback sources using LLMs without selecting feedback. While it performs well in high- and medium-resource languages with GPT-4o, it struggles in low-resource languages (e.g., Nepali; Kannada) within ChatGPT. This performance drop could be attributed to GPT-4o’s stronger ability to handle unreliable feedback in high-resource and medium-resource languages, suggesting that merely utilizing LLMs to assess all feedback does not always yield optimal results. These findings indicate that leveraging causal effects to filter helpful multilingual feedback can greatly enhance LLMs’ performance in different languages, especially in settings

where resource languages vary (see §5).

Multilingual feedback outperforms monolingual feedback. We find that CAUSAL-NATIVE performs better with stronger LLMs, particularly in high- and medium-resource languages. It achieves the second-best performance with GPT-4o in M-Hellaswag but declines with ChatGPT (dropping from 58.3% to 50.8% when compared with CAUSAL-MULTI). This indicates that multilingual feedback may be more beneficial for smaller LLMs. Additionally, the same abstention strategies yield varying results in different languages within the same LLM, which can be attributed to the ratio variations in pre-training data of different languages (Lai et al., 2023b). We further explore the impact of language relatedness of CAUSAL-MULTI in Appendix A.

4.3 Ablative Study

Whether to incorporate feedback in the abstention decision is determined by comparing NDE and TIEⁱ. To evaluate the effectiveness of this comparison strategy in determining the importance of

Ablative Settings	M-MMLU									M-Hellaswag								
	AYA-13B									CHATGPT								
	zh	it	ar	id	bn	te	ne	kn	Overall	zh	it	ar	id	bn	te	ne	kn	Overall
CAUSAL-MULTI	.512	.538	.530	.535	.574	.544	.545	.579	.545	.588	.564	.589	.626	.670	.574	.597	.685	.612
1) ignore feedback entirely	.465	.470	.530	.510	.515	.551	.465	.486	.499	.580	.534	.532	.545	.696	.561	.600	.571	.577
2) consider feedback only	.512	.542	.495	.500	.560	.503	.540	.523	.522	.589	.564	.572	.587	.670	.574	.597	.685	<u>.605</u>
3) w/o comparison	.491	.542	.515	.530	.555	.510	.460	.555	.520	.606	.572	.576	.626	.666	.556	.611	.685	.612
4) w/o aggregation	.577	.547	.490	.500	.555	.476	.574	.569	<u>.536</u>	.476	.500	.502	.605	.679	.583	.477	.600	.553

Table 2: Ablation study of CAUSAL-MULTI on two benchmarks with ChatGPT and Aya-13B. OVERALL denotes the average performance for all resource languages. The best and second-best results of Overall are highlighted in **bold** and underline, respectively.

feedback, we conduct ablation experiments under four different settings: **1)** majority voting based on $\hat{D}(Q, A)$ only (ignoring feedback entirely); **2)** majority voting by $\hat{D}(Q, A, \hat{F}^i)$ for all i only (considering feedback only); **3)** majority voting by combining both $\hat{D}(Q, A)$ and $D(Q, A, \hat{F}^i)$ (w/o comparison); **4)** majority voting based on $\hat{D}(Q, A, \hat{F}^i)$ only when $NDE < TIE^i$ (w/o aggregating over all languages).

Multilingual feedback evaluated with the causal framework enhances performance. When ignoring feedback entirely, the two LLMs demonstrate their poorest performance, with a maximum drop of 9.3% (from 57.9% to 48.6% in Kannada) and an average decrease of 5% and 4.1% across both LLMs. When LLMs make the abstention decisions without comparison (i.e., combine both $\hat{D}(Q, A)$ and $\hat{D}(Q, A, \hat{F}^i)$), we observe an average performance drop of 3.7%, with a maximum decrease of 8% in Telugu (from 54.5% to 46%). The performance of considering feedback only drops maximum of 2.3% in ChatGPT, suggesting that some feedback may introduce bias and negatively influence the performance.

Our aggregated voting strategy across related languages outperforms voting solely on language i where $TIE^i > NDE$. While in some languages, such as Nepali (57.4% on M-MMLU) and Telugu (58.3% on M-Hellaswag), providing the abstain decision based on $TIE^i > NDE$ without aggregation performs slightly better, we observe considerable performance drops in other languages, such as 8.7% in Arabic (Aya-13B) and 6.9% in Indonesian (ChatGPT) on M-Hellaswag. This varia-

tion is likely due to inconsistent language representation in the pre-training data of LLMs (Lai et al., 2023b). CAUSAL-MULTI utilizes the aggregated voting mechanism, successfully achieving greater robustness in different languages. We further provide a case where relying solely on the language with $TIE^i > NDE$ (as seen in the feedback and abstention decision for Dutch) results in an incorrect final abstention decision in Table 13.

5 Analysis

Case Study. We represent **three types of cases:** **1)** Without feedback (Table 3): LLMs can make the correct abstention decision without any feedback. However, biased feedback may negatively impact the final decision; **2)** Feedback-driven (Table 3): Feedback plays a crucial role in influencing the abstention decision and helps the model achieve correct conclusion; **3)** Multilingual feedback scenarios (Table 13): Helpful feedback from related languages leads to the correct abstention decision, whereas relying solely on feedback in the native language proves ineffective.

Performance in different resource languages. In Figure 4, we illustrate the distribution of NDE and TIE in low-, high- and medium-resource languages with ChatGPT and GPT-4o. We can see from the results of the left and right columns (ChatGPT vs. GPT-4o), GPT-4o exhibits a higher proportion of NDE and a lower proportion of TIE. This suggests that more powerful LLMs tend to have higher confidence in their proposed answers, leading to greater consistency when reviewing their own responses. Furthermore, when comparing

<p>Question: 集体安全的含义是什么? A.多个国家共同行动进行自卫的权利B.盟友为维护受害国家的权利进行自卫的权利C.经联合国安理会授权使用武装力量D.经联合国大会授权的维和行动(Translate: What is the meaning of collective security? A. The right of several countries to act together in self-defense B. The right of allies to defend themselves in order to protect the rights of the victim country C. The use of armed forces authorized by the UN Security Council D. Peacekeeping operations authorized by the UN General Assembly)</p> <p>Proposed Answer: C.经联合国安理会授权使用武装力量(Translate: C. The use of armed forces authorized by the UN Security Council)</p> <p>No Feedback: Please review the correctness of proposed answer True or False directly.. Iterate 1: True, Iterate 2: True, Iterate 3: True. -> NDE = 0.0285</p>	
<p>Feedback in native language: Please review the proposed answer and provide a paragraph of feedback on its correctness. Feedback should be in Chinese.</p> <p>Iterate 1: 你选择的答案C, 并不完全符合“集体安全”的定义...”... (Translate: False. The answer you chose, C does not fully meet the definition of "collective security". The concept of collective security usually refers to ... Therefore, the most appropriate answer is A, which means "the right of multiple countries to act together in self-defense"...)</p>	<p>Iterate 2: True. 答案的确是C... (Translate: True. The answer is indeed C. The meaning of collective security refers to ensuring the security of all countries in the international system through multilateral cooperation and coordination with international organizations (such as the United Nations). If a country is invaded or threatened... Therefore, option C is correct...)</p>
<p>Iterate 3: 所提出的答案C是正确的... (Translate: True, proposed answer C is correct .. The UN Security Council plays a central role in collective security mechanisms, particularly in authorizing the use of armed force. Therefore, option C accurately reflects the core mechanism of collective security...)</p>	
<p>Iterate 1: False, Iterate 2: True, Iterate 3: True. -> TIE = 0.0123</p> <p>CausalAbstain: NDE > TIE, then do not utilize feedback. The final causal abstain decision voted by "No Feedback" is: Do Not Abstain. (The correct answer is C.)</p>	
<p>Question: 对于复数 $z = a + bi$, 已知a和b为正整数, 并且 $z^3 + z^2 + z$ 是实数。找出最小可能的b值。 A. 10 B.1 C.6 D.11 (Translate: For a complex number $z = a + bi$, given that a and b are positive integers, and $z^3 + z^2 + z$ is a real number. Find the smallest possible value of b. A. 10 B.1 C.6 D.11)</p> <p>Proposed Answer: C.6</p> <p>No Feedback: Please review the correctness of proposed answer True or False directly. Iterate 1: True, Iterate 2: False, Iterate 3: True. -> NDE = 0.0034</p>	
<p>Feedback in native language: Please review the proposed answer and provide a paragraph of feedback on its correctness. Feedback should be in Chinese.</p> <p>Iterate 1: False. 对于复数 $z = a + bi$, 选项D是正确的, 而提议的答案C是不正确的。(Translate: False. For the complex number $z = a + bi$, it is known that a and b are positive integers, and $z^3 + z^2 + z$ is a real number. We need to find the smallest b value that satisfies the condition. From the above derivation... When $a = 6$, $b=11$ is an integer. Therefore, the minimum b value is 11. So option D is correct and proposed answer C is incorrect.)</p>	<p>Iterate 2: False. 你的答案是C, 但我们需要验证它是否正确...答案C是不正确的。(Translate: Your answer is C, but we need to verify whether it is correct. We consider the complex number $z = a + bi$, where a and b are positive integers, and $z^3 + z^2 + z$ is a real number... Therefore, answer C is incorrect. Please recheck the calculation process, the correct answer should be b = 1)</p>
<p>Iterate 3: 你所提出的答案C (6) 是正确的... (Translate: Your answer C (6) is correct. To verify this, we need to consider the condition that the complex number $z = a + bi$, that is, $z^3 + z^2 + z$ must be a real number. After calculation and analysis, we find that this condition can be satisfied when $b=6$. Therefore, option C is the correct answer. Your answer is correct...)</p>	
<p>Iterate 1: False, Iterate 2: False, Iterate 3: True. -> TIE = 0.0137</p> <p>CausalAbstain: NDE < TIE, then choose to use the feedback. The final causal abstain decision voted by "Feedback in native language" is: Abstain. (The correct answer is D.)</p>	

Table 3: Two examples in the native language (Chinese). In the first one, *CausalAbstain* chooses not to use feedback and answers correctly. In the second one, it effectively decides to abstain proposed answer with feedback, while the LLM gives an incorrect answer without feedback. We condense the Chinese expressions and primarily showcase the English translations (Translate:) for readability. The full context is provided in Table 11 and 12.

Method	M-MMLU			M-Hellaswag		
	Aya H/M	ChatGPT L	GPT-4o L	Aya H/M	ChatGPT L	GPT-4o L
ASK CALI	.557	.507	.447	.596	.229	.456
REFLECT	.471	.376	.509	.547	.521	.487
MOREINFO	.479	.385	.539	.353	.766	.606
CONFLICT	.484	.524	.527	.580	.679	.519
MULTI-RELATED	.538	.561	.570	.471	.790	.660
CAUSAL-NATIVE	.501	.507	.560	.504	.782	.644
CAUSAL-MULTI	.529	.561	.578	.555	.788	.687
					.592	.632
					.631	.534
						.792
						.716

Table 4: The average performance of high-/medium-resource (H/M) and low-resource languages (L).

the top and bottom rows (language resource levels), both LLMs show a higher proportion of NDE in high- and medium-resource languages, indicating greater confidence in their answers for well-represented languages. Additionally, Table 4 reveals that the average performance gap between different resource languages is smaller in Aya-13B. This suggests that Aya-13B, being explicitly multilingual, performs more consistently among various languages compared to the general-purpose GPT-4o and ChatGPT.

CAUSAL-RANDOM and CAUSAL-ENGLISH. We employ multiple languages for feedback generation, but the languages are randomly selected from a language pool, which is named CAUSAL-RANDOM. As shown in Table 5, we observe that overall, CAUSAL-MULTI performs better than

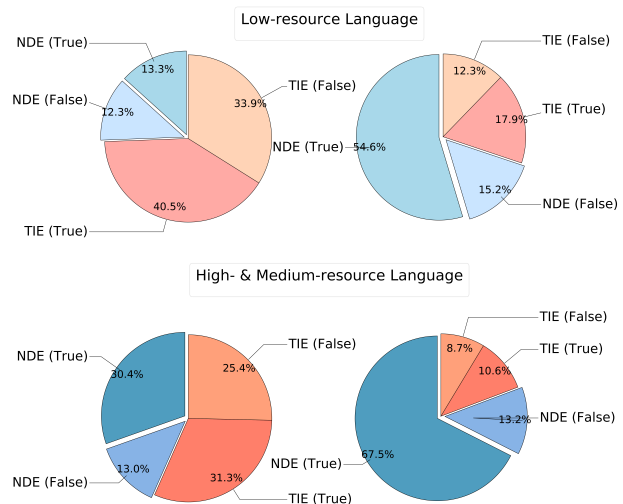


Figure 4: The distribution and accuracy rate of NDE and TIE with ChatGPT (left column) and GPT-4o (right column) in different resource languages.

CAUSAL-RANDOM and CAUSAL-NATIVE, which showcase the importance of choosing the related languages. This implies that the source of performance improvement comes from both the languages selected and the CausalAbstain method. As shown in Table 5, in GPT-4o, CAUSAL-ENGLISH underperforms CAUSAL-NATIVE. In ChatGPT, the overall performance of CAUSAL-ENGLISH is better than CAUSAL-NATIVE, which is similar to the findings in previous research (Shi et al., 2022; Zhang

CAUSAL-RANDOM										
LLMs	zh	it	ar	id	bn	te	ne	kn	Overall	Overall (CAUSAL-MULTI)
ChatGPT	.584	.615	.598	.596	.571	.524	.493	.513	.562	.583
GPT-4o	.746	.807	.710	.780	.714	.653	.748	.621	.722	.754
Aya-13B	.525	.508	.510	.600	.666	.520	.582	.647	.570	.612
CAUSAL-ENGLISH										
LLMs	zh	it	ar	id	bn	te	ne	kn	Overall	Overall (CAUSAL-NATIVE)
ChatGPT	.571	.602	.681	.570	.588	.601	.532	.561	.588	.508
GPT-4o	.785	.794	.690	.800	.688	.666	.716	.644	.723	.749
Aya-13B	.549	.559	.515	.575	.675	.515	.611	.616	.577	.564

Table 5: The results of CAUSAL-RANDOM and CAUSAL-ENGLISH in M-Hellaswag. The overall performance of CASUAL-MULTI and CAUSAL-NATIVE is presented for comparison purposes.

et al., 2023b) that indicate ChatGPT performs better in English than other languages.

6 Related Work

Causal for LLMs LLMs can significantly benefit from causality, as it enhances their ability to understand and reason about cause-and-effect relationships within data (Liu et al., 2024a). A large body of research has explored this (Li et al., 2021; Jin et al., 2024a; Wang et al., 2023b; Jin et al., 2024b; Jiang et al., 2024; Chen et al., 2024), including identifying knowledge bias pre-trained in LLMs that can lead to incorrect answers and hallucinations (Zhang et al., 2024; Wu et al., 2024); improving LLMs performance in specific tasks, such as visual question answering (Zhao et al., 2023; Zang et al., 2023).

Abstention. Existing abstention strategies can be categorized as follows: *Calibration-based* strategies aim to extract confidence score to gauge uncertainty (Tian et al., 2023; Kuhn et al., 2023). *Prompt-based* strategies employ instructions to induce reflection and determine whether the generated answer is reliable (Kadavath et al., 2022; Edunov et al., 2018; Si et al., 2023).

Consistency-based and *collaboration-based* strategies involve multiple LLMs reflecting on their answers cooperatively (Feng et al., 2024c; Wang et al., 2023a). Recent *feedback-based* methods adapt to multilingual settings. However, these approaches either underperform outside of English or struggle with generated feedback quality. To address this gap, we propose a causal method that identifies helpful feedback, improving multilingual abstention.

7 Conclusion and Future Work

We propose a novel causal method, *CausalAbstain*, to assist LLMs in abstaining from providing incorrect answers. Extensive experiments on two datasets demonstrate that our method, CAUSAL-MULTI, achieves state-of-the-art performance in multilingual settings, while CAUSAL-NATIVE exhibits competitive performance in monolingual settings, suggesting the effectiveness of abstention based on causality. In future work, we will explore additional contextual factors, such as model uncertainty and external knowledge sources, to further improve abstention decisions across diverse tasks and languages.

Limitation

Our approach, CAUSAL-MULTI, leverages multilingual feedback to help LLMs make abstention decisions. While it requires prompt LLMs multiple times, leading to higher inference costs compared to simpler prompting approaches, it is still not the most computationally expensive method compared with (Feng et al., 2024c). To mitigate the cost, we introduce CAUSAL-NATIVE, which reduces the number of prompting requests while maintaining competitive performance. We compare the LLMs’ overhead during inference using different abstention strategies, shown in Table 9.

Furthermore, our causal graph in Figure 2 provides a general framework for formulating the abstention task. Future work could explore incorporating latent variables, such as LLM biases, to further refine the approach.

Ethics Statement

Our study on mitigating bias and hallucinations in LLMs acknowledges the ethical implications of data-driven biases in AI, particularly their impact on performance. All experiments were conducted using publicly available datasets, and no human participants were involved.

Acknowledgement

This work is partially supported by Tencent Rhino-Bird Focused Research Program (Value-aligned Credible Large Language Model) and RMGS project (Artificial Intelligence and Big Data Analytics for Social Good).

References

- Amos Azaria and Tom Mitchell. 2023. [The internal state of an LLM knows when it’s lying](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 967–976, Singapore. Association for Computational Linguistics.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. 2023. A multi-task, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*.
- Hang Chen, Bingyu Liao, Jing Luo, Wenjing Zhu, and Xinyu Yang. 2024. [Learning a structural causal model for intuition reasoning in conversation](#). *IEEE Transactions on Knowledge and Data Engineering*, 36(7):3210–3223.
- Jiale Cheng, Xiao Liu, Kehan Zheng, Pei Ke, Hongning Wang, Yuxiao Dong, Jie Tang, and Minlie Huang. 2024. [Black-box prompt optimization: Aligning large language models without model training](#). *Preprint*, arXiv:2311.04155.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation at scale](#). *Preprint*, arXiv:1808.09381.
- Shangbin Feng, Weijia Shi, Yuyang Bai, Vidhisha Balachandran, Tianxing He, and Yulia Tsvetkov. 2024a. [Knowledge card: Filling llms’ knowledge gaps with plug-in specialized language models](#). *Preprint*, arXiv:2305.09955.
- Shangbin Feng, Weijia Shi, Yike Wang, Wenxuan Ding, Orevaoghene Ahia, Shuyue Stella Li, Vidhisha Balachandran, Sunayana Sitaram, and Yulia Tsvetkov. 2024b. [Teaching LLMs to abstain across languages via multilingual feedback](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4125–4150, Miami, Florida, USA. Association for Computational Linguistics.
- Shangbin Feng, Weijia Shi, Yike Wang, Wenxuan Ding, Vidhisha Balachandran, and Yulia Tsvetkov. 2024c. [Don’t hallucinate, abstain: Identifying LLM knowledge gaps via multi-LLM collaboration](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14664–14690, Bangkok, Thailand. Association for Computational Linguistics.
- Yair Gat, Nitay Calderon, Amir Feder, Alexander Chapanin, Amit Sharma, and Roi Reichart. 2023. [Faithful explanations of black-box nlp models using llm-generated counterfactuals](#). *Preprint*, arXiv:2310.00603.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. [Survey of hallucination in natural language generation](#). *ACM Computing Surveys*, 55(12):1–38.
- Haitao Jiang, Lin Ge, Yuhe Gao, Jianian Wang, and Rui Song. 2024. [Llm4causal: Democratized causal tools for everyone via large language model](#). *Preprint*, arXiv:2312.17122.
- Zhijing Jin, Yuen Chen, Felix Leeb, Luigi Gresele, Ojasv Kamal, Zhiheng Lyu, Kevin Blin, Fernando Gonzalez Adauto, Max Kleiman-Weiner, Mrinmaya Sachan, and Bernhard Schölkopf. 2024a. [Cladder: Assessing causal reasoning in language models](#). *Preprint*, arXiv:2312.04350.
- Zhijing Jin, Jiarui Liu, Zhiheng Lyu, Spencer Poff, Mrinmaya Sachan, Rada Mihalcea, Mona Diab, and Bernhard Schölkopf. 2024b. [Can large language models infer causation from correlation?](#) *Preprint*, arXiv:2306.05836.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom Brown, Jack Clark, Nicholas Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Kaplan. 2022. [Language models \(mostly\) know what they know](#). *Preprint*, arXiv:2207.05221.
- Haoqiang Kang, Terra Blevins, and Luke Zettlemoyer. 2024. [Comparing hallucination detection metrics for multilingual generation](#). *Preprint*, arXiv:2402.10496.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. [Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation](#). *Preprint*, arXiv:2302.09664.
- Sachin Kumar, Vidhisha Balachandran, Lucille Njoo, Antonios Anastasopoulos, and Yulia Tsvetkov. 2023. [Language generation models can cause harm: So what can we do about it? an actionable survey](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3299–3321, Dubrovnik, Croatia. Association for Computational Linguistics.
- Viet Dac Lai, Nghia Trung Ngo, Amir Pouran Ben Veyseh, Hieu Man, Franck Dernoncourt, Trung Bui, and Thien Huu Nguyen. 2023a. [Chatgpt beyond english: Towards a comprehensive evaluation of large language models in multilingual learning](#). *arXiv preprint arXiv:2304.05613*.
- Viet Dac Lai, Chien Van Nguyen, Nghia Trung Ngo, Thuat Nguyen, Franck Dernoncourt, Ryan A. Rossi, and Thien Huu Nguyen. 2023b. [Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback](#). *Preprint*, arXiv:2307.16039.

- Zhongyang Li, Xiao Ding, Kuo Liao, Bing Qin, and Ting Liu. 2021. [Causalbert: Injecting causal knowledge into pre-trained models with minimal supervision](#). *Preprint*, arXiv:2107.09852.
- Patrick Littell, David R. Mortensen, Ke Lin, Katherine Kairis, Carlisle Turner, and Lori Levin. 2017. [URIEL and lang2vec: Representing languages as typological, geographical, and phylogenetic vectors](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 8–14, Valencia, Spain. Association for Computational Linguistics.
- Xiaoyu Liu, Paiheng Xu, Junda Wu, Jiaxin Yuan, Yifan Yang, Yuhang Zhou, Fuxiao Liu, Tianrui Guan, Hao-liang Wang, Tong Yu, Julian McAuley, Wei Ai, and Furong Huang. 2024a. [Large language models and causal inference in collaboration: A comprehensive survey](#). *Preprint*, arXiv:2403.09606.
- Xin Liu, Muhammad Khalifa, and Lu Wang. 2024b. [Litcab: Lightweight language model calibration over short- and long-form responses](#). *Preprint*, arXiv:2310.19208.
- Nishanth Madhusudhan, Sathwik Tejaswi Madhusudhan, Vikas Yadav, and Masoud Hashemi. 2024. [Do llms know when to not answer? investigating abstention abilities of large language models](#). *Preprint*, arXiv:2407.16221.
- Abhika Mishra, Akari Asai, Vidhisha Balachandran, Yizhong Wang, Graham Neubig, Yulia Tsvetkov, and Hannaneh Hajishirzi. 2024. [Fine-grained hallucination detection and editing for language models](#). *Preprint*, arXiv:2401.06855.
- R OpenAI. 2023. [Gpt-4 technical report](#). arxiv 2303.08774. *View in Article*, 2(5).
- Judea Pearl et al. 2000. [Models, reasoning and inference](#). Cambridge, UK: CambridgeUniversityPress, 19.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim Rocktäschel, and Sebastian Riedel. 2021. [KILT: a benchmark for knowledge intensive language tasks](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2523–2544, Online. Association for Computational Linguistics.
- Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung, Yi Tay, Sebastian Ruder, Denny Zhou, et al. 2022. [Language models are multilingual chain-of-thought reasoners](#). *arXiv preprint arXiv:2210.03057*.
- Chenglei Si, Weijia Shi, Chen Zhao, Luke Zettlemoyer, and Jordan Boyd-Graber. 2023. [Getting more out of mixture of language model reasoning experts](#). *Preprint*, arXiv:2305.14628.
- Aviv Slobodkin, Omer Goldman, Avi Caciularu, Ido Dagan, and Shauli Ravfogel. 2023. [The curious case of hallucinatory \(un\)answerability: Finding truths in the hidden states of over-confident large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3607–3625, Singapore. Association for Computational Linguistics.
- Jimin Sun, Hwijeen Ahn, Chan Young Park, Yulia Tsvetkov, and David R. Mortensen. 2021. [Cross-cultural similarity features for cross-lingual transfer learning of pragmatically motivated tasks](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2403–2414, Online. Association for Computational Linguistics.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher Manning. 2023. [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*, pages 5433–5442, Singapore. Association for Computational Linguistics.
- Neeraj Varshney and Chitta Baral. 2023. [Post-abstention: Towards reliably re-attempting the abstained instances in qa](#). *Preprint*, arXiv:2305.01812.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023a. [Self-consistency improves chain of thought reasoning in language models](#). *Preprint*, arXiv:2203.11171.
- Zhaowei Wang, Quyet V. Do, Hongming Zhang, Jiayao Zhang, Weiqi Wang, Tianqing Fang, Yangqiu Song, Ginny Y. Wong, and Simon See. 2023b. [Cola: Contextualized commonsense causal reasoning from the causal inference perspective](#). *Preprint*, arXiv:2305.05191.
- Junda Wu, Tong Yu, Xiang Chen, Haoliang Wang, Ryan Rossi, Sungchul Kim, Anup Rao, and Julian McAuley. 2024. [DeCoT: Debiasing chain-of-thought for knowledge-intensive tasks in large language models via causal intervention](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14073–14087, Bangkok, Thailand. Association for Computational Linguistics.
- Jian Xie, Kai Zhang, Jiangjie Chen, Renze Lou, and Yu Su. 2024. [Adaptive chameleon or stubborn sloth: Revealing the behavior of large language models in knowledge conflicts](#). *Preprint*, arXiv:2305.13300.
- Jian Yang, Xinyu Hu, Gang Xiao, and Yulong Shen. 2024a. [A survey of knowledge enhanced pre-trained language models](#). *ACM Transactions on Asian and Low-Resource Language Information Processing*.

Yuqing Yang, Ethan Chern, Xipeng Qiu, Graham Neubig, and Pengfei Liu. 2024b. [Alignment for honesty](#). *Preprint*, arXiv:2312.07000.

Jifan Yu, Xiaozhi Wang, Shangqing Tu, Shulin Cao, Daniel Zhang-Li, Xin Lv, Hao Peng, Zijun Yao, Xiaohan Zhang, Hanming Li, Chunyang Li, Zheyuan Zhang, Yushi Bai, Yantao Liu, Amy Xin, Nianyi Lin, Kaifeng Yun, Linlu Gong, Jianhui Chen, Zhili Wu, Yunjia Qi, Weikai Li, Yong Guan, Kaisheng Zeng, Ji Qi, Hailong Jin, Jinxin Liu, Yu Gu, Yuan Yao, Ning Ding, Lei Hou, Zhiyuan Liu, Bin Xu, Jie Tang, and Juanzi Li. 2024. [Kola: Carefully benchmarking world knowledge of large language models](#). *Preprint*, arXiv:2306.09296.

Chuanqi Zang, Hanqing Wang, Mingtao Pei, and Wei Liang. 2023. Discovering the real association: Multimodal causal reasoning in video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19027–19036.

Shuo Zhang, Liangming Pan, Junzhou Zhao, and William Yang Wang. 2024. [The knowledge alignment problem: Bridging human and external knowledge for large language models](#). *Preprint*, arXiv:2305.13669.

Xiang Zhang, Senyu Li, Bradley Hauer, Ning Shi, and Grzegorz Kondrak. 2023a. [Don’t trust ChatGPT when your question is not in English: A study of multilingual abilities and types of LLMs](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7915–7927, Singapore. Association for Computational Linguistics.

Xiang Zhang, Senyu Li, Bradley Hauer, Ning Shi, and Grzegorz Kondrak. 2023b. [Don’t trust chatgpt when your question is not in english: a study of multilingual abilities and types of llms](#). *arXiv preprint arXiv:2305.16339*.

Shitian Zhao, Zhuowan Li, Yadong Lu, Alan Yuille, and Yan Wang. 2023. [Causal-cog: A causal-effect look at context generation for boosting multi-modal language models](#). *Preprint*, arXiv:2312.06685.

Han Zhou, Xingchen Wan, Lev Proleev, Diana Mincu, Jilin Chen, Katherine Heller, and Subhrajit Roy. 2024. [Batch calibration: Rethinking calibration for in-context learning and prompt engineering](#). *Preprint*, arXiv:2309.17249.

Appendix

A Analysis

The ablative study of GPT-4o. We present the ablation study of GPT-4o, as shown in Table 6. Similar trends are observed in the ablation experiments on ChatGPT and Aya-13B. GPT-4o demonstrates the poorest average performance when feedback is entirely ignored, with a drop of 7.6% on Avg-L (from 53.4% to 45.8% in M-Hwllaswag). When GPT-4o makes the abstention decisions without comparison (i.e., combining both $\hat{D}(Q, A)$ and $\hat{D}(Q, A, \hat{F}^i)$), we observe an average performance drop of 3.8% in M-MMLU and 2.4% in M-Hellaswag on Avg-L. Additionally, the absence of feedback results in relatively good performance for relatively high-resource languages, aligning with the observations in Figure 4.

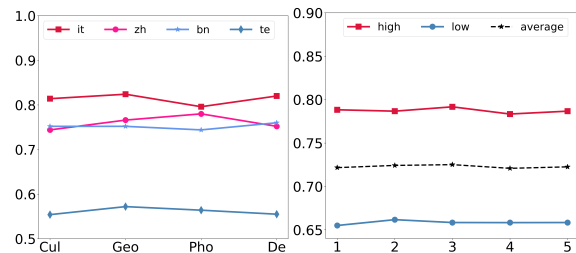


Figure 5: The impact of language relatedness (different settings) on specific languages, and the effect of iteration count on high- and low-resource languages.

The influence of language relatedness and iteration count.

Following the six linguistic attributes identified in Lang2vec (i.e., syntactic, geographic, phonological, genetic, inventory, and featural (Littell et al., 2017)), Feng et al. (2024b) analyzed language relationships based on these attributes. They also introduced the *culture* attribute, which defines related languages within the same cultural cluster according to the World Values Survey. Their findings indicate that *geography*, *phonology* and *culture* are the most influential attributes. Following (Feng et al., 2024b), the related languages are shown in Table 8. To assess the impact of related languages on CAUSAL-MULTI, we conducted an experiment comparing four different settings (culture, geography, phonology, default) as shown in Figure 5. Our observations revealed that performance varied across different language queries depending on the choice of related languages. Consequently, we utilized 100 held-out sets in our experiments to determine the most suitable related

GPT-4o											
M-MMLU											
Ablative Settings	Avg-HM	zh	it	ar	id	bn	te	ne	kn	Avg-L	Overall
CAUSAL-MULTI	.788	.765	.830	.760	.795	.740	.670	.675	.671	.686	.738
1) ignoring feedback entirely	.761	.755	.830	.740	.720	.670	.531	.595	.636	.608	.685
2) considering feedback only	.781	.760	.815	.765	.785	.725	.639	.685	.630	.670	.726
3) w/o comparison	.786	.760	.825	.760	.800	.735	.578	.655	.623	.648	.717
4) w/o aggregating across all languages	.783	.755	.825	.760	.790	.715	.660	.690	.630	.674	.728
M-Hellaswag											
Ablative Settings	Avg-HM	zh	it	ar	id	bn	te	ne	kn	Avg-L	Overall
CAUSAL-MULTI	.620	.593	.658	.598	.630	.558	.547	.519	.513	.534	.577
1) ignoring feedback entirely	.570	.571	.559	.593	.557	.476	.457	.472	.425	.458	.514
2) considering feedback only	.620	.593	.658	.593	.630	.558	.547	.519	.513	.534	.576
3) w/o comparison	.601	.571	.615	.580	.639	.554	.502	.497	.486	.510	.556
4) w/o aggregating across all languages	.570	.580	.576	.563	.562	.584	.524	.493	.530	.533	.552

Table 6: Ablation study of CAUSAL-MULTI on two benchmarks using GPT-4o. The results include the average performance for high- and medium-resource languages (Avg-HM), low-resource languages (Avg-L), and overall performance across all languages.

Method	LLaMa								Phi									
	zh	it	ar	id	bn	te	ne	kn	Overall	zh	it	ar	id	bn	te	ne	kn	Overall
ASK CALI	.646	.442	.695	.662	.269	.388	.304	.473	.485	.381	.423	.490	.435	.442	.442	.514	.467	.449
REFLECT	.419	.519	.570	<u>.546</u>	.531	.504	.480	.527	.512	.652	.603	.570	<u>.578</u>	.455	.450	.480	.457	<u>.531</u>
MOREINFO	.536	.506	.490	.474	<u>.497</u>	.380	.514	.473	.484	.484	.468	.543	.487	.407	.543	.601	.504	.505
CONFLICT	.510	<u>.532</u>	.536	.494	.421	<u>.527</u>	.514	.434	.496	.568	.526	.457	.435	.476	<u>.504</u>	.527	.527	.503
MULTI-RELATED	.497	<u>.532</u>	.477	.481	.476	.466	.460	.442	.479	.497	.532	.563	.455	.497	.442	<u>.581</u>	.543	.514
CAUSAL-NATIVE	.574	.474	.536	.532	.455	.480	<u>.580</u>	<u>.542</u>	<u>.522</u>	.516	.608	.589	.584	.441	.410	.445	.500	.512
CAUSAL-MULTI	<u>.612</u>	.589	<u>.576</u>	.523	.496	.581	.581	.558	.565	<u>.574</u>	.647	.609	.584	.482	.441	.506	<u>.531</u>	.547

Table 7: The abstention accuracy in M-MMLU in LLaMa and Phi, the **bold** and underline represent the best and second performance, respectively.

languages for each specific language.

Additionally, we conduct experiments on the different settings of iteration time from 1-5 with a held-out set (100 instances in the validation set), as shown in Figure 5. While variations in the number of iterations have a minimal impact on performance, we have set the iteration count to three, which may yield a slight improvement on average.

The comparison between LLMs’ answer accuracy and the abstention accuracy. In Figure 6, we observe that there is no synchronized correlation between the two performance metrics. This suggests that abstention is a distinct research question, similar to the observation made by (Feng et al., 2024b) in other abstention strategy settings.

Results on LLaMa and Phi. We provide additional results within LLaMa (LLaMa3.2) and Phi (Phi4). As shown in the following Table 7, we can observe that CAUSAL-MULTI outperforms all the baselines with LLaMa and Phi. Meanwhile, CAUSAL-NATIVE shows the second-best performance within LLaMa.

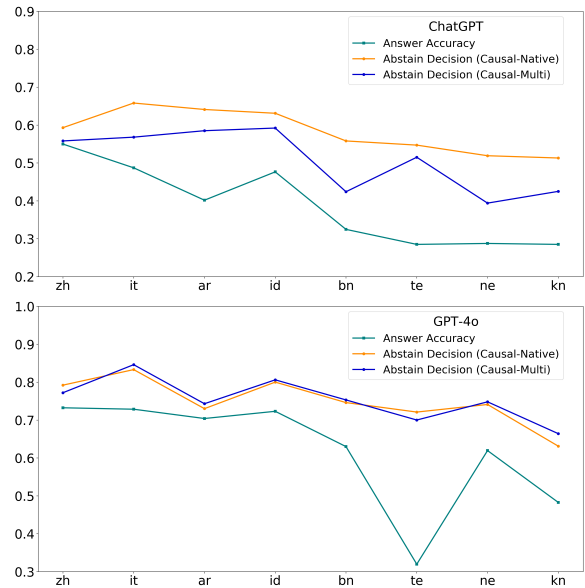


Figure 6: Comparison of the accuracy of LLMs’ proposed answers and their abstain decisions.

Related language settings: "zh": ["Chinese", "Chinese", "Chinese", "English", "Russian", "German", "Italian", "Dutch", "Arabic", "Arabic", "Slovak", "Danish"], "it": ["French", "Slovak", "Hungarian", "German", "French", "Hungarian", "Chinese", "Dutch", "Arabic", "Catalan", "Romanian", "Ukrainian"], "id": ["Indonesian", "Indonesian", "Indonesian", "Vietnamese", "Bengali", "Tamil", "English", "Russian", "Catalan", "Vietnamese", "Catalan", "Russian"], "ar": ["Arabic", "Hindi", "Bengali", "English", "Russian", "German", "Chinese", "Italian", "Dutch", "Chinese", "Slovak", "Danish"], "bn": ["Arabic", "Hindi", "Bengali", "Nepali", "Vietnamese", "Hindi", "Telugu", "Kannada", "Russian", "Hindi", "Telugu", "Nepali"], "ne": ["Arabic", "Hindi", "Bengali", "Hindi", "Vietnamese", "Kanaada", "Telugu", "Hindi", "Romanian", "Telugu", "Kannada"], "te": ["Arabic", "Hindi", "Bengali", "Hindi", "Tamil", "Malayalam", "Kannada", "Russian", "Catalan", "Kannada", "Tamil", "Nepali"], "kn": ["Arabic", "Hindi", "Bengali", "Tamil", "Malayalam", "Marathi", "Kannada", "Russian", "Catalan", "Telugu", "Malayalam", "Tamil"]

Table 8: Language settings.

B Experiment Details

Dateset. We use M-MMLU and MHel-laswag (Lai et al., 2023b) and follow Lai et al. (2023a,b) in categorizing languages based on their data ratios in the pre-training corpus. A language is classified as high-resource (HRL) if its data ratio exceeds 1%, medium-resource (MRL) if it falls between 0.1% and 1%, and low-resource (LRL) if it is below 0.1%. Specifically, we include languages spanning from low to high resource levels: Bengali (bn), Telugu (te), Nepali (ne), Kannada (kn), Arabic (ar), Indonesian (id), Chinese (zh), and Italian (it).

Model details. We use the Azure OpenAI API checkpoint "gpt-4o-2024-08-06" for GPT-4o, and the "GPT-3.5-Turbo" model checkpoint from OpenAI for ChatGPT. Additionally, we utilize the "CohereForAI/aya-101" model checkpoint from Huggingface for Aya-13B. Following Feng et al. (2024c), we conduct the experiments on the baselines.

Method	# LLM inference requests
MOREINFO	2
ASK CALI	6
REFLECT	2
MULTI-RELATED	5
COOP-SELF	8
CAUSAL-NATIVE	4
CAUSAL-MULTI	10

Table 9: The # of LLM inference requests for each query with different abstention strategies.

LLMs overhead with different abstention strategies. For efficiency, as shown in the following Table 10, the computational efficiency of the proposed *CausalAbtsain* is $\mathcal{O}(N)$, where N is the number of iterations. Larger N may lead to higher costs. In addition, various abstention mechanisms have differing computational overheads, with the most

N	# LLM inference requests of CAUSAL-NATIVE	# LLM inference requests of CAUSAL-MULTI
1	2	4
2	3	7
3	4	10
4	5	13
5	6	16

Table 10: The # of LLM inference requests for each query with CAUSAL-NATIVE and CAUSAL-MULTI, where N is the number of iterations.

significant cost being the number of LLM inference requests for each query. A comparison of these mechanisms, including ours, is provided in Table 9, following the approach outlined by Feng et al. (2024c).

Question: 对于复数 $z = a + bi$, 已知 a 和 b 为正整数, 并且 $z^3 + z^2 + z$ 是实数。找出最小可能的 b 值。A. 10 B.1 C.6 D.11 (Translate: For a complex number $z = a + bi$, given that a and b are positive integers, and $z^3 + z^2 + z$ is a real number. Find the smallest possible value of b . A. 10 B.1 C.6 D.11)

Proposed Answer: C.6

No Feedback: Please review the proposed answer True or False directly.

Iterate 1: True,

Iterate 2: False,

Iterate 3: True. -> NDE = 0.0034

Feedback in native language:

Please review the proposed answer and provide a paragraph of feedback on its correctness. Feedback should be in language i .

Iterate 1: False. 对于复数 $z = a + bi$, 已知 a 和 b 为正整数, 并且 $z^3 + z^2 + z$ 是实数。我们需要找到满足条件的最小 b 值。从上面的推导, 我们有虚部为零的条件: $b(1 + 2a + 3a^2 - b^2) = 0$ 因为 b 为正整数, 所以: $1 + 2a + 3a^2 = b^2$ 我们尝试不同的 a 值以找到最小的 b 值: $a = 1$ 时: ... $a = 2$ 时: ... $a = 6$ 时, $b=11$ 是整数。因此, 最小的 b 值为 11。所以选项 D 是正确的, 而提议的答案 C 是不正确的。(Translate: False. For the complex number $z = a + bi$, it is known that a and b are positive integers, and $z^3 + z^2 + z$ is a real number. We need to find the smallest b value that satisfies the condition. From the above derivation, we have the condition that the imaginary part is zero: $b(1 + 2a + 3a^2 - b^2) = 0$ Because b is a positive integer, so: $1 + 2a + 3a^2 = b^2$ We try different a values to find the smallest b value: When $a = 1$:... When $a = 2$:... When $a = 6$, $b=11$ is an integer. Therefore, the minimum b value is 11. So option D is correct and proposed answer C is incorrect.) ✓

Iterate 2: False。你的答案是 C, 但我们需要验证它是否正确。我们考虑复数 $z = a + bi$, 其中 a 和 b 是正整数, 并且 $z^3 + z^2 + z$ 是实数。我们可以将 z 分别代入中, 展开并分离出实部和虚部。设虚部为零, 通过计算可以得到相应的条件。经过详细推导和计算, 最小可能的 b 值应为 1。因此, 答案 C 是不正确的。(Translate: Your answer is C, but we need to verify whether it is correct. We consider the complex number $z = a + bi$, where a and b are positive integers, and $z^3 + z^2 + z$ is a real number. We can substitute z into it separately, expand it and separate the real part and the imaginary part. Let the imaginary part be zero, and the corresponding conditions can be obtained through calculation. After detailed derivation and calculation, the minimum possible b value should be 1. Therefore, answer C is incorrect. Please recheck the calculation process, the correct answer should be $b = 1$) ✓

Iterate 3: 你所提出的答案 C (6) 是正确的。为了验证这一点, 我们需要考虑复数 $z = a + bi$ 的条件, 即 $z^3 + z^2 + z$ 必须是实数。经过计算和分析, 我们发现当 $b=6$ 时, 这个条件可以满足。因此, 选项 C 是正确答案。你的解答准确无误。(Translate: Your answer C (6) is correct. To verify this, we need to consider the condition that the complex number $z = a + bi$, that is, $z^3 + z^2 + z$ must be a real number. After calculation and analysis, we find that this condition can be satisfied when $b=6$. Therefore, option C is the correct answer. Your answer is correct...) ✗

Based on the feedback for measuring the correctness of the answer, is the proposed answer True or False?

Iterate 1: False,

Iterate 2: False,

Iterate 3: True.

-> TIE = 0.0137

CausalAbstain: NDE < TIE, then choose to use the feedback. The final causal abstain decision from TIE is: Abstain. (The correct answer is D.)

Table 11: An example in the native language setting demonstrates that *CausalAbstain* effectively uses the feedback to decide to abstain, while LLM gives an incorrect answer without feedback.

Question: 集体安全的含义是什么? A.多个国家共同行动进行自卫的权利B.盟友为维护受害国家的权利进行自卫的权利C.经联合国安理会授权使用武装力量D.经联合国大会授权的维和行动

(**Translate:** What is the meaning of collective security? A. The right of several countries to act together in self-defense B. The right of allies to defend themselves in order to protect the rights of the victim country C. The use of armed forces authorized by the UN Security Council D. Peacekeeping operations authorized by the UN General Assembly)

Proposed Answer: C

No Feedback:

Please review the proposed answer True or False directly.

Iterate 1: **True**,

Iterate 2: **True**,

Iterate 3: **True**.

-> **NDE = 0.0285**

Feedback in native language:

Please review the proposed answer and provide a paragraph of feedback on its correctness. Feedback should be in Chinese.

Iterate 1: 你选择的答案C，即“经联合国安理会授权使用武装力量”，并不完全符合“集体安全”的定义。集体安全的概念通常指的是多个国家为了维护共同的安全利益而采取联合行动，防止和应对对其中任何一个国家的侵略或威胁。因此，最合适的答案是A，即“多个国家共同行动进行自卫的权利”... (**Translate:** False. The answer you chose, C, which means "the use of armed force authorized by the UN Security Council", does not fully meet the definition of "collective security". The concept of collective security usually refers to joint actions taken by multiple countries to safeguard common security interests and prevent and respond to aggression or threats against any of them. Therefore, the most appropriate answer is A, which means "the right of multiple countries to act together in self-defense"...)

Iterate 2: True。答案的确是C。集体安全的含义是指在国际体系中，通过多边合作和国际组织（如联合国）的协调，确保各国的安全。如果某一国家受到侵略或威胁，其他国家将在集体安全的框架内采取行动来恢复和平与安全。联合国安理会有权根据《联合国宪章》第七章授权使用武装力量，以维持或恢复国际和平与安全。因此，选项C是正确的... (**Translate:** True. The answer is indeed C. The meaning of collective security refers to ensuring the security of all countries in the international system through multilateral cooperation and coordination with international organizations (such as the United Nations). If a country is invaded or threatened, other countries will act within the framework of collective security to restore peace and security. The United Nations Security Council has the authority to authorize the use of armed force under Chapter 7 of the United Nations Charter to maintain or restore international peace and security. Therefore, option C is correct...)

✓

Iterate 3: 所提出的答案C是正确的。集体安全的含义是指各国通过共同协作，采取集体行动来应对对国际和平与安全的威胁。联合国安理会在集体安全机制中起着核心作用，特别是在授权使用武装力量方面。因此，选项C“经联合国安理会授权使用武装力量”准确地反映了集体安全的核心机制... (**Translate:** True, proposed answer C is correct. The meaning of collective security refers to countries working together and taking collective actions to respond to threats to international peace and security. The UN Security Council plays a central role in collective security mechanisms, particularly in authorizing the use of armed force. Therefore, option C "use of armed force authorized by the United Nations Security Council" accurately reflects the core mechanism of collective security...)

Based on the feedback for measuring the correctness of the answer, is the proposed answer True or False?

Iterate 1: **False**,

Iterate 2: **True**,

Iterate 3: **True**.

-> **TIE = 0.0123**

CausalAbstain: NDE > TIE, then abstain from utilizing feedback. The final causal abstain decision from NDE is: Do Not Abstain. (The correct answer is C.)

Table 12: An example in the native language demonstrates that *CausalAbstain* correctly addresses the query without utilizing feedback. Additionally, some feedback exhibits biased content.

Question: 以下哪个最不可能减少样本调查中的偏差? A.首次未回应调查的人后继续跟进B.以中立的方式提问, 避免影响回答C.使用分层随机抽样而不是简单随机抽样D.随机选择样本

(**Translate:** Which of the following is least likely to reduce bias in sample surveys? A. Follow up with those who did not respond to the survey for the first time B. Ask questions in a neutral manner to avoid influencing the answers C. Use stratified random sampling instead of simple random sampling D. Select samples randomly)

Proposed Answer: D.随机选择样本

No Feedback: Please review the proposed answer True or False directly.

Iterate 1: **True**, Iterate 2: **False**, Iterate 3: **False**. → **NDE = 0.0034**

Feedback in native language:

Please review the proposed answer and provide a paragraph of feedback on its correctness. Feedback should be in Chinese.

提议的答案是正确的。随机选择样本 (选项D) 是减少样本调查中偏差的常见方法之一, 因为它可以确保每个个体都有相同的被选中的机会, 从而使样本更具代表性。相反, A、B、C选项都是具体措施, 旨在进一步减少偏差。例如, A通过跟进未回应调查的人, 可以减少非响应偏差; B通过以中立的方式提问, 可以减少测量偏差; C通过使用分层随机抽样, 可以确保不同子群体在样本中的代表性。因此, D选项最不可能减少样本调查中的偏差。

(**Translate:** The proposed answer is correct. Randomly selecting a sample (option D) is one of the common ways to reduce bias in sample surveys because it ensures that every individual has the same chance of being selected, making the sample more representative. In contrast, options A, B, and C are specific measures designed to further reduce bias. For example, A can reduce non-response bias by following up with people who did not respond to the survey; B, it can reduce measurement bias by asking questions in a neutral way; C, by using stratified random sampling, it can ensure that different subgroups are represented in the sample. Therefore, option D is least likely to reduce bias in sample surveys)

Iterate 1: **True**, Iterate 2: **True**, Iterate 3: **False**. → **TIE = 0.0137**

CausalAbstain: NDE < TIE, then choose to use the feedback. The final causal abstain decision from TIE is: Not Abstain. (The correct answer is C.) → **[WRONG]**

Feedback in related language:

Please review the proposed answer and provide a paragraph of feedback on its correctness. Feedback should be in language i.

English: The proposed answer, D, is incorrect. Randomly choosing samples (D) is actually a fundamental method for reducing bias in sample surveys, as it ensures that every individual in the population has an equal chance of being selected. The correct answer should be A. Following up with individuals who did not initially respond (A) can help to reduce nonresponse bias, while asking neutral questions (B) avoids influencing answers, and using stratified random sampling (C) ensures representation from different subgroups. Therefore, A, B, and C are all methods that help reduce bias, making D the least likely to reduce bias.

Iterate 1: **False**, Iterate 2: **False**, Iterate 3: **True**. → **TIE_en = 0**

Italian: La risposta proposta, D, "randomly choosing a sample," non è corretta. La scelta casuale del campione è una tecnica fondamentale per ridurre il bias nel campionamento. Le altre opzioni, come seguire le persone che inizialmente non hanno risposto (A), usare un linguaggio neutrale nelle domande (B), e usare il campionamento stratificato (C), sono tutte strategie che aiutano a ridurre il bias. Pertanto, la risposta più corretta sarebbe quella che non contribuisce a ridurre il bias, ma la risposta D contribuisce effettivamente a ridurre il bias.

(**Translate:** The proposed answer, D, "randomly choosing a sample," is incorrect. Random sampling is a key technique for reducing sampling bias. Other options, such as following up with people who initially did not respond (A), using neutral language in questions (B), and using stratified sampling (C), are all strategies that help reduce bias. Therefore, the most correct answer would be the one that does not help reduce the bias, but answer D actually helps reduce the bias.)

Iterate 1: **False**, Iterate 2: **True**, Iterate 3: **False**. → **TIE_it = 0**

Dutch: De voorgestelde antwoordoptie D is inderdaad correct. Het willekeurig kiezen van een steekproef (optie D) is een gebruikelijke methode om een representatieve steekproef te verkrijgen en kan helpen om vooringenomenheid te verminderen. Echter, in vergelijking met de andere opties, is het minder specifiek gericht op het verminderen van bepaalde soorten bias die kunnen optreden in een steekproefonderzoek. Optie A (het opvolgen van niet-reagerende personen) helpt om non-respons bias te verminderen. Optie B (neutrale vragen stellen) helpt om respons bias te verminderen. Optie C (gestratificeerde steekproeftrekking) helpt om variabiliteit binnen subgroepen te beheersen en kan preciezere schattingen geven. Daarom is D inderdaad het minst waarschijnlijk om bias te verminderen in vergelijking met de andere opties.

(**Translate:** The suggested answer option D is indeed correct. Random sampling (option D) is a common method of obtaining a representative sample and can help reduce bias. However, compared to the other options, it is less specifically aimed at reducing certain types of bias that can occur in a sample survey. Option A (following up on non-responders) helps to reduce non-response bias. Option B (asking neutral questions) helps to reduce response bias. Option C (stratified sampling) helps control variability within subgroups and can provide more precise estimates. Therefore, D is indeed the least likely to reduce bias compared to the other options.)

Iterate 1: **True**, Iterate 2: **False**, Iterate 3: **True**. → **TIE_du = 0.0137**

Based on the feedback for measuring the correctness of the answer, is the proposed answer True or False?

CausalAbstain: NDE not ≤ all TIE (NDE < TIE_du), then voted the feedback and answer of multilingual TIE, the final causal abstain decision (five False, four True): Abstain. (The correct answer is C.) → **[CORRECT]**

Table 13: An example where feedback in the native language led to an incorrect decision, while multilingual feedback proved effective. We present one of the feedback details from N iterations.