

Teaching LLMs to Abstain across Languages via Multilingual Feedback

Shangbin Feng¹ Weijia Shi¹ Yike Wang¹ Wenxuan Ding² Orevaoghene Ahia¹
Shuyue Stella Li¹ Vidhisha Balachandran³ Sunayana Sitaram³ Yulia Tsvetkov¹

¹University of Washington ²The University of Texas at Austin ³Microsoft Research

shangbin@cs.washington.edu

Abstract

Multilingual large language models (LLMs) often have knowledge disparities across languages, with larger gaps in under-resourced languages. Teaching LLMs to abstain in the face of knowledge gaps is thus a promising strategy to mitigate hallucinations in multilingual settings. However, previous studies on LLM abstention primarily focus on English; we find that directly applying these solutions beyond English results in up to 20.5% performance gaps between high and low-resource languages, potentially due to LLMs' drop in calibration and reasoning beyond a few resource-rich languages. To this end, we propose strategies to enhance LLM abstention by *learning from multilingual feedback*, where LLMs self-reflect on proposed answers in one language by generating multiple feedback items in related languages: we show that this helps identify the knowledge gaps across diverse languages, cultures, and communities. Extensive experiments demonstrate that our multilingual feedback approach outperforms various strong baselines, achieving up to 9.2% improvement for low-resource languages across three black-box and open models on three datasets, featuring open-book, closed-book, and commonsense QA. Further analysis reveals that multilingual feedback is both an effective and a more equitable abstain strategy to serve diverse language speakers, and cultural factors have great impact on language selection and LLM abstention behavior, highlighting future directions for multilingual and multi-cultural reliable language modeling.¹

1 Introduction

Large language models (LLMs) encode extensive information and aid knowledge-intensive tasks (Petroni et al., 2019; Brown et al., 2020; Yu et al., 2023). However, knowledge gaps and subsequent model hallucinations pose an everlasting challenge

¹Code and data are publicly available at <https://github.com/BunsenFeng/M-AbstainQA>.

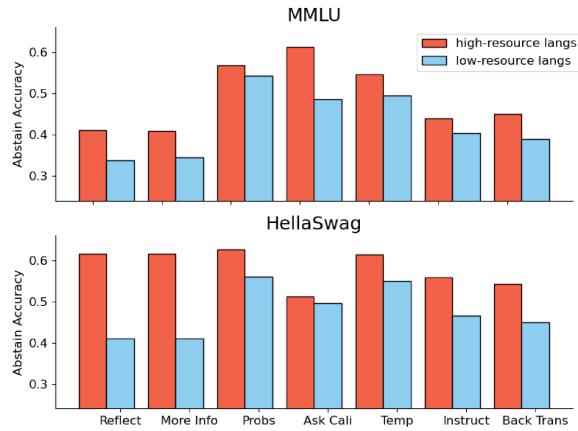


Figure 1: Average accuracy of abstention baselines in low- and high-resource languages with AYA-13B. **Existing abstain strategies drop by 8.4% on average when applied to QA in low-resource languages.**

that compromises LLM reliability (Lazaridou et al., 2021; Ji et al., 2023; Kumar et al., 2023; Mishra et al., 2024). A growing body of work seeks to enhance LLM reliability by teaching them to *abstain*, i.e., avoiding wrong answers in low-confidence scenarios to mitigate hallucinations and factual inaccuracies. While these studies put forward viable solutions, they are evaluated on English only (Gu and Hopkins, 2023; Varshney and Baral, 2023; Yang et al., 2023; Feng et al., 2024). However, the factuality of multilingual LLMs in low-resource languages is often worse (Zhang et al., 2023b; Lai et al., 2023; Kang et al., 2024), underserving diverse language speakers and communities. As such, there is an urgent need for robust abstaining strategies that work with the long tail of languages.

Developed and evaluated in English, are existing abstain approaches viable for low-resource languages? Drawing from Feng et al. (2024), we first evaluate seven existing abstain methods, spanning calibration (Jiang et al., 2021; Tian et al., 2023), prompting (Edunov et al., 2018; Kadavath et al., 2022; Feng et al., 2023), and training (Ouyang et al., 2022), on multilingual MMLU and Hel-

laswag datasets (Lai et al., 2023) featuring 8 high-resource and 7 low-resource languages. Figure 1 demonstrates that performance degrades by up to 12.8% and 20.5% for both datasets: while existing approaches perform well in English, they are limited by LLMs’ diminishing utility and calibration accuracy in low-resource languages, struggling to identify knowledge gaps and abstain accordingly. As a result, we ask: *how to identify knowledge gaps in LLMs and reliably abstain beyond English?*

To this end, we present the first study on multilingual LLM abstention and propose to teach LLMs to abstain by *generating and learning from multilingual feedback in related languages* (Figure 2). While the concept of generated feedback was previously demonstrated to improve reasoning and alignment in English-only scenarios (Du et al., 2023; Madaan et al., 2024), sampling diverse and high-quality feedback in low-resource languages is challenging due to LLMs’ diminishing utility in long-tail languages (Lai et al., 2023). Multilingual LLMs can leverage related languages to improve performance via transfer learning (Lin et al., 2019a; Pires et al., 2019; Asai et al., 2023; Tanwar et al., 2023), so we expect that generating feedback from *related languages* would help identify knowledge gaps across diverse domains and cultures. Therefore, we probe multilingual LLMs to provide feedback, on its proposed answer, in several *related* languages, where language relatedness is defined by linguistic typology, geography, or culture (Litell et al., 2017; Lin et al., 2019b; Sun et al., 2021). Together with the proposed answer and generated feedback from the most related languages, LLMs reason and self-reflect to make abstain decisions.

We evaluate baselines and our *multilingual feedback* approach using three open-source and proprietary LLMs (*GPT-4*, *Aya-13B*, and *ChatGPT*) on three datasets in open-domain, closed-book, and commonsense QA. Extensive experiments demonstrate that *multilingual feedback* consistently outperforms strong baselines across models and datasets, achieving up to 9.2% improvements of abstain accuracy for low-resource languages. Further analysis reveals that *multilingual feedback* presents a more equitable abstain strategy, highlighting culture as a driving factor in multilingual abstention. It impacts the optimal languages for feedback and LLMs’ performance gaps across diverse information domains.

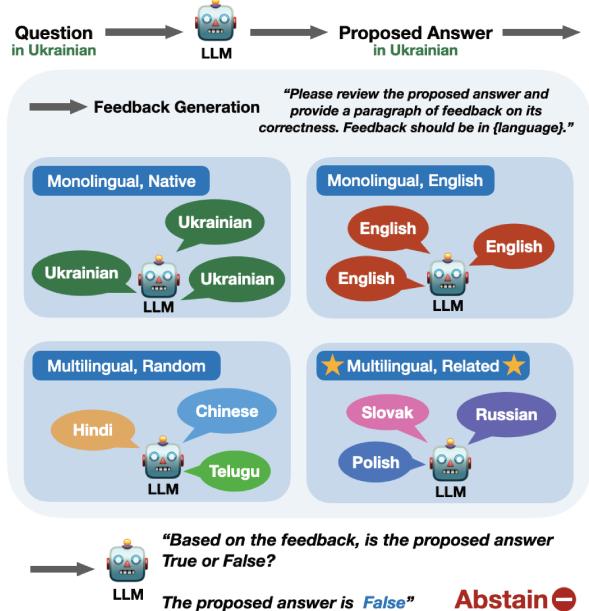


Figure 2: **Overview of abstaining via multilingual feedback.** LLMs generate feedback on its proposed answer in four language settings to make abstain decisions.

2 Methodology

Background We focus on teaching LLMs to **Abstain** in Question Answering (AbstainQA) (Feng et al., 2024): given a query q and an LLM, we aim to develop robust abstention strategies $f(q, \text{LLM}) \rightarrow \{\text{true}, \text{false}\}$. Ideally, the LLM abstains ($f = \text{true}$) when it would provide an incorrect answer and should not abstain ($f = \text{false}$) when it is capable of generating a correct answer (Feng et al., 2024). f should work for diverse languages of varying language families, resourceness levels, and speaker communities.

Since existing approaches to LLM abstention are limited by LLMs’ diminishing utility and calibration beyond English (Figure 1, §1), we propose to *teach LLMs to abstain via multilingual feedback*, hypothesizing that self-feedback about its proposed answer from related languages could help identify the blind spots across cultures, perspectives, and contexts. We present an overview in Figure 2.

Abstain with Multilingual Feedback LLMs take three steps to make an abstain decision:

- Given a question q , the LLM first generates a proposed answer: $a = \text{LLM}(q)$.
- We probe the LLM itself for multilingual feedback on its proposed answer, concretely $f_i = \text{LLM}(q, a | \ell_i)$ where feedback f_i is generated in language ℓ_i . We specifically use the prompt

“Please review the proposed answer and provide a paragraph of feedback on its correctness. Feedback should be in ℓ_i .” to elicit f_i .

- Finally, the LLM employs the multilingual feedback to reason and make an abstain decision: $\text{LLM}(\mathbf{q}, \mathbf{a}, \{f_1, f_2, \dots, f_k\}) \rightarrow \{\text{true}, \text{false}\}$. We specifically use the prompt “*Based on the feedback, is the proposed answer True or False?*” and abstain if the answer \mathbf{a} is deemed false.

Language Selection Contrary to English-only scenarios, it is often challenging to sample diverse and high-quality feedback in one low-resource language. We hypothesize that by generating feedback in related languages to the language of the question ℓ , LLMs could better identify internal knowledge gaps and patch the blind spots with information across varying cultures, perspectives, and more. We experiment with four modes of selecting feedback languages $\{\ell_1, \dots, \ell_k\}$.

- monolingual, native** (MONO-NATIVE): all feedback are sampled in the native language of the question: $\ell_1 = \dots = \ell_k = \ell$. This resembles the previous English-only setting where questions and feedback are in the same language (English).
- monolingual, English** (MONO-ENGLISH): regardless of the language of the question, all feedback are sampled in English: $\ell_1 = \dots = \ell_k = \text{English}$. This is because English is the highest-resource language and is often used as the source language in cross-lingual transfer (Conneau et al., 2018; Conneau and Lample, 2019; Hu et al., 2020; Wang et al., 2020b).
- multilingual, random** (MULTI-RANDOM): this is a control setting where we employ multiple languages for feedback generation, but the languages are randomly selected from a language pool \mathcal{L} : $\ell_i = \text{random_choice}(\mathcal{L})$.
- multilingual, related** (MULTI-RELATED): we propose to employ languages related to the language of the question ℓ for feedback generation. Concretely, we employ Lang2vec (Littell et al., 2017) to obtain the vector representation of a language v_ℓ^a in a linguistic attribute $a \in \mathcal{A}$.² We define the distance between a pair of languages

²Six attributes are considered in Lang2vec: syntactic, geographic, phonological, genetic, inventory, and featural.

as an average of distances across attributes:

$$\text{dist}(\ell, \ell') = \frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} \frac{(v_\ell^a)^T v_{\ell'}^a}{\|v_\ell^a\| \|v_{\ell'}^a\|}$$

The k languages with the least distance to ℓ are then selected for feedback generation: $\{\ell_1, \dots, \ell_k\} = \text{argmin}_{\ell' \in \mathcal{A}} \text{dist}(\ell, \ell')$. We employ $k = 3$ multilingual feedback by default.

3 Experiment Settings

Models We evaluate existing approaches and the four proposed monolingual/multilingual feedback strategies with three LLMs: *Aya-13B*, a specifically multilingual instruction-tuned model, *ChatGPT* and *GPT-4*, two general-purpose black-box LLMs. We employ greedy decoding for QA and making an abstain decision, and employ a temperature of 0.7 when sampling repeatedly (e.g., consistency-based baselines and feedback generation).

Datasets We evaluate with the Multilingual MMLU (M-MMLU) and Hellaswag (M-Hellaswag) datasets (Lai et al., 2023), featuring encyclopedic and commonsense knowledge. Originally in English, these QA problems were translated into 26 other languages through machine translation. These languages are characterized as 8 high-resource languages, 11 mid-resource languages, and 7 low-resource languages based on their proportion in pretraining data.³ We also present evaluation with Belebele (Bandarkar et al., 2023) in Appendix A, a multilingual reading comprehension dataset. For the three datasets, we create random splits with 200 instances for validation and 800 for test, with minor variation across languages due to data availability.

Baselines We compare with nine abstain baselines that could be adapted in multilingual settings: calibration-based PROBS (token probabilities), TEMP (Jiang et al., 2021), ASK CALI. (Tian et al., 2023); training-based INSTRUCT (Ouyang et al., 2022); prompting-based REFLECT (Kadavath et al., 2022), MOREINFO (Feng et al., 2023), BACKTRANS (Edunov et al., 2018); and consistency-based approaches SCTHRES. (Wang et al., 2022), CONFLICT (Feng et al., 2024). More details about the baselines are in Appendix B.

³Full language list in Appendix B.

Method	M-MMLU										M-Hellaswag									
	Avg-H	bn	ta	ne	ml	mr	te	kn	Avg-L	Avg-H	bn	ta	ne	ml	mr	te	kn	Avg-L		
AYA-13B																				
PROBS	.567	.551	.521	.519	.542	.564	.524	.574	.542	.626	.597	.567	.555	.547	.513	.560	.580	.560		
TEMP	.547	.515	.439	.485	.526	.547	.518	.432	.495	.614	.610	.556	.543	.489	.559	.527	.556	.549		
ASK CALI.	<u>.613</u>	.503	.494	.476	.474	.492	.490	.460	.486	.512	.510	.489	.508	.466	.496	.514	.490	.496		
INSTRUCT	.539	.441	.348	.412	.362	.417	.426	.419	.404	.559	.597	.421	.510	.333	.481	.442	.480	.466		
REFLECT	.410	.347	.300	.339	.336	.357	.335	.347	.337	<u>.615</u>	.489	.357	.448	.312	.437	.404	.426	.410		
MOREINFO	.409	.348	.299	.350	.357	.358	.337	.368	.345	<u>.615</u>	.497	.360	.444	.325	.441	.393	.413	.410		
BACKTRANS.	.450	.421	.333	.453	.346	.354	.411	.411	.390	.542	.571	.393	.484	.300	.487	.442	.474	.450		
SCTHRES.	.609	<u>.618</u>	.614	<u>.609</u>	<u>.610</u>	.600	.584	.610	<u>.607</u>	.532	.532	.443	.577	.543	<u>.572</u>	.589	.520	.539		
CONFLICT	.564	.567	.581	.568	.521	.568	.561	.582	.564	.536	.520	.546	.514	.559	.548	.553	.544	.540		
MONO-NATIVE	.512	.580	.515	.604	.529	.576	.533	.520	.551	.552	.578	.479	.452	.467	.481	.524	.526	.501		
MONO-ENGLISH	.611	.611	.607	.649	.460	.583	.594	.688	.599	.581	.513	.514	.503	.513	.506	<u>.565</u>	<u>.572</u>	.527		
MULTI-RANDOM	.540	.597	.615	.561	.524	.549	.628	.605	.583	.481	.403	.650	.497	.627	.565	<u>.565</u>	.553	.551		
MULTI-RELATED	.631	<u>.621</u>	.704	.595	.661	.590	.643	.628	.635	.603	.468	.636	.542	.693	.578	.558	.566	.577		
GPT-4																				
ASK CALI.	.432	.421	.404	.500	.598	.444	.450	.589	.487	.536	.342	.307	.461	.393	.452	.376	.304	.376		
INSTRUCT	.789	.566	.363	.493	.386	.556	.481	.465	.473	.656	.552	.186	.432	.160	.435	.272	.270	.330		
REFLECT	.686	.655	.585	.649	.528	.597	.519	.589	.589	.658	.545	.229	.561	.347	.571	.483	.408	.449		
MOREINFO	.694	.572	.711	.588	.677	.611	.558	.612	.619	.386	.461	.486	.555	.507	.584	.469	.543	.515		
BACKTRANS.	.764	.634	.563	.696	.535	.660	<u>.620</u>	.636	.621	.538	.522	.576	.564	.677	.558	.555	.582	.576		
SCTHRES.	.735	.541	.544	.596	.604	.650	.605	.598	.591	.759	.508	.679	.497	<u>.673</u>	.508	.528	.570	.566		
CONFLICT	.730	.555	.641	.589	.561	.629	.559	.590	.589	.639	.488	<u>.593</u>	.503	<u>.673</u>	.501	<u>.535</u>	.557	.550		
MONO-NATIVE	.728	.655	.548	.642	.567	.660	.589	.628	.613	.708	.558	.371	<u>.665</u>	.307	.597	.401	.447	.478		
MONO-ENGLISH	.789	.669	.541	.703	.543	.653	.550	.659	.617	.737	.584	.200	.613	.260	<u>.526</u>	.340	.421	.421		
MULTI-RANDOM	.698	.710	.570	.655	.567	.681	.581	.651	.631	.714	.532	.300	.606	.380	.532	.408	.441	.457		
MULTI-RELATED	.785	.752	.659	.730	.638	.674	.636	.659	.678	.722	.532	.543	.706	.647	.610	.531	<u>.572</u>	.592		

Table 1: Performance of [calibration](#), [training](#), [prompting](#), [consistency](#), and our proposed [feedback](#)-based approaches on two LLMs and two multilingual datasets. We employ the Abstain Accuracy metric, Avg-H and Avg-L denote average performance for high and low-resource languages, while we additionally present performance for the seven low-resource languages (Bengali, Tamil, Nepali, Malayalam, Marathi, Telugu, and Kannada). Best performance in **bold** and second-best in underline. Baselines that rely on token probabilities (e.g., [Probs](#)) are not compatible with GPT-4. **MULTI-RELATED achieves the best average performance in low-resource languages across all models and datasets, improving over baselines by up to 9.2%.**

Evaluation Metrics We use the Abstain Accuracy metric (A-Acc) proposed in [Feng et al. \(2024\)](#): LLMs should abstain when it would provide an incorrect answer and should not abstain when it would provide a correct answer, concretely $A\text{-Acc} = \frac{TP+TN}{TP+TN+FP+FN}$ and TP indicates the LLM should abstain and did. We additionally report other AbstainQA metrics (Reliable Accuracy, Effective Reliability) in Appendix A.

4 Results

We present the abstain accuracy results with two LLMs on two multilingual datasets in Table 1.

MULTI-RELATED achieves state-of-the-art performance. MULTI-RELATED achieves the highest average performance on low-resource languages (Avg-L) across all four model and dataset settings, improving over the strongest baseline by 4.9% on average. Out of the 7 low-resource lan-

guages, MULTI-RELATED achieves the best and top-2 performance in 3.25 and 4.75 languages on average. This improvement in low-resource languages comes with on-par performance in high-resource languages (Avg-H), outperforming baselines in 81% of the times across four (model, dataset) settings. This indicates that by generating and reflecting on multilingual feedback from related languages, LLMs greatly improve in identifying inherent knowledge gaps across languages.

Existing approaches greatly drop beyond high-resource languages. Ask for Calibration ([Tian et al., 2023](#)), an approach to solicit LLM confidence scores verbally, witness a 12.7% drop from high to low-resource languages ($0.613 \rightarrow 0.486$) on MMLU using AYA-13B. While it could generate meaningful confidence scores between 0 and 1 for high-resource languages, it collapses and repeatedly generate the same number (e.g., 0.8) for al-

most all questions in low-resource languages. Similar performance gaps and failure modes could be observed for previously strong approaches in English such as Instruction Tuning (35.3% drop, on average), Self-Reflect (33.3%), and SCthreshold (12.2%). In comparison, **MULT-RELATED** has a smaller drop of 8.5%: we further quantify the fairness of abstain strategies in Section 5.

Abstaining is a language-specific problem. Out of the seven low-resource languages, we observe that Tamil (ta) and Malayalam (ml) are consistently the most challenging languages across models, datasets, and approaches: an average performance of 0.484 and 0.492 is achieved on the two languages, while the global average for low-resource languages is 0.520. This could be attributed to their low representation in LLM pretraining data (Lai et al., 2023) and thus lower utility, meaning that there is no one-size-fits-all solution for abstaining across multilingual contexts and robust strategies should be language-specific. **MULTI-RELATED** takes linguistic knowledge into account by employing *related* languages for feedback generation, successfully achieving the best Avg-L performance across all models and datasets. We further study the utility of language relatedness in Section 5.

AYA-13B shows smaller gaps than GPT-4. While the performance of **MULTI-RELATED** is higher on GPT-4, the gap between low and high-resource languages is smaller with AYA-13B (1.7% vs. 16.9%). Since **MULTI-RELATED** specifically relies on generating and reasoning in multilingual contexts, the explicitly multilingual AYA-13B would be better than the general-purpose GPT-4 to this end. This motivates a potential collaboration between models: using a stronger general-purpose LLM for QA and a smaller but explicitly multilingual LLM for feedback generation. We further explore this in Section 5.

5 Analysis

MULTI-RELATED is more equitable. While we primarily focused on the performance gaps between high and low-resource languages in Section 4, measuring the fairness of a multilingual system goes beyond performance averages. Concretely, we follow Song et al. (2023) to measure utility and equity, indicating how well multilingual approaches serve diverse language speakers and performance

Method	Demo. (M_1, \uparrow)	Ling. (M_0, \uparrow)	Equity (G, \downarrow)
PROBS	0.5613	0.5632	0.0319
ASK CALI.	0.5976	0.5784	0.0488
INSTRUCT	0.4514	0.4280	0.0477
REFLECT	0.3983	0.3877	0.0460
BACKTRANS.	0.4342	0.4261	0.0517
SCTHRES.	0.5974	0.5916	0.0340
CONFLICT	0.5698	0.5630	0.0369
MONO-NATIVE	0.5181	0.5318	0.0472
MONO-ENGLISH	0.6038	0.5651	0.0564
MULTI-RANDOM	0.5442	0.5528	0.0390
MULTI-RELATED	0.6149	0.6027	0.0278

Table 2: Utility and equity metrics of abstain strategies, where \uparrow/\downarrow indicates that higher/lower values are desirable. Best performance in **bold** and second-best in underline. **MULTI-RELATED offers a fairer abstain strategy with higher utility and lower Gini coefficient.**

disparity across languages. For utility:

$$M_\tau = \sum_{\ell \in \mathcal{L}} d_\ell^\tau \cdot u_\ell, \quad d_\ell = \frac{n_\ell}{\sum_{\ell \in \mathcal{L}} n_\ell}$$

where u_ℓ denotes the utility/performance on language ℓ , n_ℓ denotes the number of native speakers, the exponential $\tau = 1$ indicates *demographic weighted* utility and $\tau = 0$ indicates *linguistic weighted* utility where all languages are treated as equals. For equity, performance on various languages are sorted in non-decreasing order ($u_i \leq u_{i+1}$) and the Gini coefficient is calculated:

$$G = \frac{1}{|\mathcal{L}|} \left(|\mathcal{L}| + 1 - 2 \frac{\sum_{i=1}^{|\mathcal{L}|} (|\mathcal{L}| + 1 - i) u_i}{\sum_{i=1}^{|\mathcal{L}|} u_i} \right)$$

where $|\mathcal{L}|$ indicates the total number of languages. The range of G is 0 to 1 and more equitable abstain strategies should have lower G values.

We present the demographic utility, linguistic utility, and equity metrics in Table 2. **MULTI-RELATED** outperforms baselines on both utility modes, while being more equitable across languages, evident in the 12.9% reduction in Gini Coefficient. On the contrary, MONO-ENGLISH have on-par demographic utility but worse linguistic utility and equity, indicating that generated feedback in English is unevenly helpful to other languages, whereas low-resource languages distant from English benefit much less.

MULTI-RELATED offers relevant, informative, and conflicting pieces of feedback. To better understand the quality and role of the generated

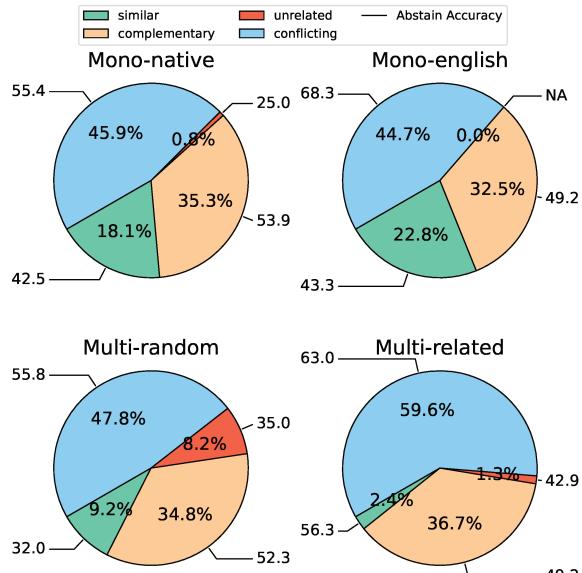


Figure 3: **GPT-4 evaluation of the role of the k pieces of feedback to a given question:** whether they are similar, complementary, conflicting, or unrelated. We also present the abstain accuracy in each feedback role below the role name, showing that conflicting feedback are best for LLM self-reflection while *Multi-related* has the most conflicting feedback scenarios.

feedback, we employ GPT-4 evaluation for large-scale automatic analysis. We compare the four feedback modes by first translating all feedback into English with GPT-4 to remove the confounding factor of language difference, then using LLM-as-a-judge methodology (Zheng et al., 2024) for pairwise comparison. Given a question, proposed answer, and a pair of two feedback from MULTI-RELATED and a baseline, GPT-4 evaluates which feedback is more *relevant* and *informative*. We present the win rate in Figure 4: MULTI-RELATED is consistently more relevant to the question across languages. While MONO-ENGLISH offers more informative feedback on high-resource languages, it degrades into the worst on low-resource languages while MULTI-RELATED becomes the best.

For feedback roles, we first conduct a manual examination to identify four potential roles of the k generated feedback: *similar*, *complementary*, *conflicting*, and *unrelated*. We then employ GPT-4 to evaluate the roles of k feedback, translated into English, for the same question. Figure 3 demonstrates that monolingual approaches result in 252.7% more similar and thus redundant feedback compared to multilingual settings, while MULTI-RANDOM result in greater unrelated feedback potentially due to the random selection of distant languages. In com-

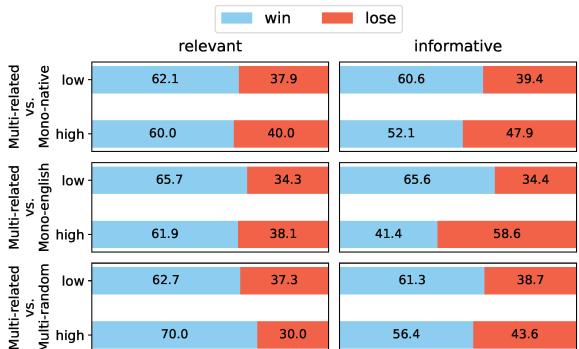


Figure 4: GPT-4 evaluation of feedback pairs to the same question, comparing MULTI-RELATED against other feedback settings to evaluate which produces more *relevant* and *informative* feedback.

parison, MULTI-RELATED produces 24.7% more conflicting scenarios where feedback disagree in content or conclusion: the abstain accuracy on *conflicting* scenarios are also the highest, indicating that LLMs face more knowledge conflicts (Xie et al., 2023; Wang et al., 2023b) by generating multiple feedback from related but different languages, which in turn aids self-reflection and making better-informed abstain decisions (Feng et al., 2024). We further present a qualitative analysis in Appendix A in addition to the automatic GPT-4 evaluation.

Culture is a driving factor in multilingual abstention. For MULTI-RELATED, we by default define language relatedness as the average of the six linguistic attributes in Lang2vec (Littell et al., 2017). (§2) We further investigate what aspects of “language relatedness” are most helpful for abstaining across multilingual contexts. Specifically, we additionally select related languages only by one of the six categories (e.g., syntactic or phonological relatedness). We introduce two additional settings: 1) LLMs are prompted to propose three related languages by themselves; 2) related languages in the same *culture* cluster according to the World Value Survey.⁴ We present the performance of various language relatedness settings in Table 3. We observe that geography and phonology are the most helpful linguistic attributes, while *culture*-informed language selection yields the best utility and equity results. This indicates that *multilingual feedback from languages of related socio-cultural backgrounds is most helpful for low-resource languages and overall fairness*.

To further investigate the impact of culture, we

⁴<https://www.worldvaluessurvey.org>

Setting	Avg-H	Avg-M	Avg-L	Equity (\downarrow)
DEFAULT	0.6411	0.5861	0.4432	0.0943
SYNTACTIC	0.6452	0.5843	0.4395	0.0967
GEOGRAPHIC	0.6589	0.5892	0.4712	0.0880
PHONOLOGICAL	<u>0.6499</u>	0.5881	<u>0.5017</u>	<u>0.0694</u>
GENETIC	0.6476	0.6024	0.4944	0.0839
INVENTORY	0.6371	0.5827	0.4356	0.0950
FEATURAL	0.6412	<u>0.6116</u>	0.4417	0.0916
LLM-GENERATED	0.6316	0.5929	0.4362	0.0981
CULTURE	0.6425	0.6202	0.5322	0.0438

Table 3: Performance averages for high, mid, and low-resource languages, as well as the equity metric G for various language relatedness settings. Best performance in **bold** and second-best in underline. **Culturally informed language selection is best for mid and low-resource languages and also more equitable.**

present the performance breakdown of various MMLU domains in Figure 6. We illustrate the 10 domains with the largest gaps between low- and high-resource languages and 10 domains with the least gaps. The largest gaps often come from west-centric topics such as “US history”, “European history”, and “US foreign policy”, while the smallest gaps are often on STEM domains that transcend socio-cultural contexts such as “logical fallacies”, “high school physics”, and “electrical engineering”. This again indicates that *culture* is a driving factor in multilingual abstention: improving LLM abstain capabilities is not only a technical problem but also a social-oriented one, where the existing West-centric LLMs (Naous et al., 2023) should better incorporate other cultures and perspectives for equitable improvements in factuality and reliability.

Abstain decisions are less transferable across unrelated and low-resource languages. One solution to multilingual abstain is to take the highest-resource language (e.g., English), make abstain decisions, and use that decision to abstain/generate in low-resource languages. However, to what extent do abstain decisions overlap across languages and thus transferable remains underexplored, which could not be taken for granted given the factuality variation across languages (Lai et al., 2023; Kang et al., 2024). To this end, we visualize the abstain overlap of parallel questions across various three-language groups in Figure 5, where overlapping parts indicate that MULTI-RELATED for 2 or 3 languages decided to abstain. For control group #1, the group of three related languages sees much greater overlap (74.5% 2+ overlap) than the three unrelated languages (48.1%). For control group

Method	High-Resource								
	ru	de	zh	fr	es	it	nl	vi	avg.
SELF	.818	.852	.794	.838	.823	.769	.831	.555	.785
OTHER	.733	.788	.762	.772	.815	.733	.808	.668	.760
Method	Low-Resource								
	bn	ta	ne	ml	mr	te	kn	/	avg.
SELF	.752	.659	.730	.638	.674	.636	.659	/	.678
OTHER	.788	.722	.735	.656	.669	.735	.697	/	.715

Table 4: Performance when using GPT-4 *itself* or the *other* AYA-13B multilingual LLM for feedback generation. The collaboration between a general-purpose LLM for QA and a smaller but more multilingual model for feedback generation benefits low-resource languages.

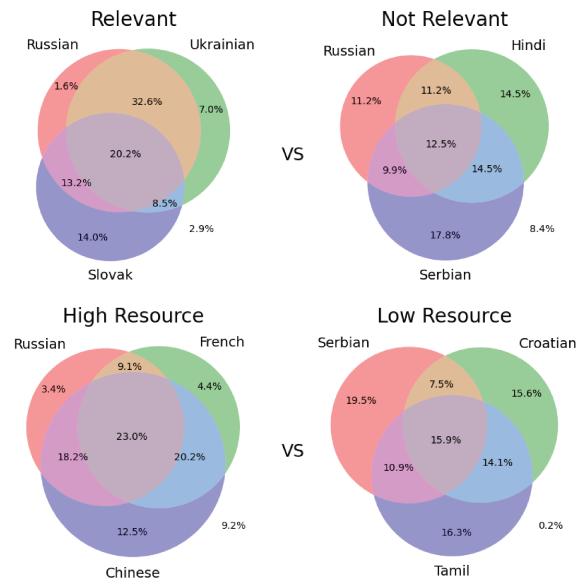


Figure 5: Overlap of abstain decisions made in different languages, where the overlap indicates that LLMs abstain in both/all three of the languages. **We find that abstain decisions are only somewhat transferrable between relevant and high-resource language clusters.**

#2, a group of three high-resource languages sees greater overlap (70.5%) than three low-resource languages (48.4%). These two findings together indicate that abstain decisions are only somewhat transferable in the case of high-resource closely related languages: however, many languages on the long tail are neither close to English nor well-represented in LLM training data, thus English-only abstain methods are not one-size-fit-all solutions and abstaining is a language-specific problem.

General-purpose LLMs could be supervised by a smaller but more multilingual model. Motivated by the finding that *GPT-4* has higher absolute performance but *Aya-13B* witnesses smaller gaps

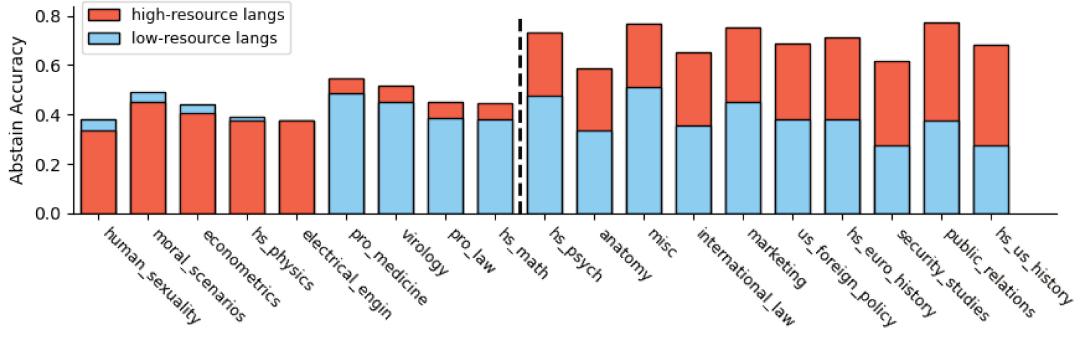


Figure 6: Abstain accuracy on various MMLU domains with high and low-resource languages: on the left we show the 10 domains with the least performance gaps and on the right we show the 10 with the most gaps. “hs” indicates high school. **While domains with the least gaps often feature STEM topics that are more objective, domains with the largest gaps are often driven by culture, especially West-centric social knowledge.**

with MULTI-RELATED (§4), we explore the collaboration between the two models: using *GPT-4* for question answering and *Aya-13B* for multilingual feedback generation. Table 4 demonstrates that while for high-resource languages this might be counterproductive, for low-resource languages it results in a 5.4% improvement on average. This indicates that when user queries come in low-resource and underrepresented languages, a smaller but explicitly multilingual model could be employed to supervise general-purpose black-box LLMs to improve abstaining and mitigate hallucinations.

6 Related Work

Teaching LLMs to Abstain Existing works focus on various types of approaches: *Calibration-based* approaches focus on eliciting the confidence levels of LLMs with token probabilities (Sun et al., 2022; Zhou et al., 2023a; Liu et al., 2023a) or semantic markers of uncertainty (Kuhn et al., 2022; Zhou et al., 2023b, 2024), where previous research evaluate (Radford et al., 2019; Ahuja et al., 2022; Liang et al., 2023; Tao et al., 2023; He et al., 2023) and improve (Kamath et al., 2020; Desai and Durrett, 2020; Jagannatha and Yu, 2020; Kong et al., 2020; Jiang et al., 2021; Lin et al., 2022a; Mielke et al., 2022) calibration for various tasks (Wang et al., 2020a; Stengel-Eskin and Van Durme, 2023; Kalai and Vempala, 2023; Zablotskaia et al., 2023). *Prompting-based* approaches induce self-reflection by including none-of-the-above options (Kadavath et al., 2022), prompt to self-correct reasoning (Kim et al., 2024; Shinn et al., 2023; Huang et al., 2023a; Chen et al., 2023), ask for additional information (Feng et al., 2023), and more (Wang et al., 2023a; Si et al., 2023). *Training-*

based approaches aim to adapt LLMs for abstention with linear probing (Slobodkin et al., 2023; Azaria and Mitchell, 2023), training an extra module (Cobbe et al., 2021), or alignment objectives (Zhang et al., 2023a; Yang et al., 2023; Sun et al., 2023; Bashlovkina et al., 2023; Huang et al., 2023b; Liu et al., 2023b). *Consistency-based* (Wang et al., 2022; Cole et al., 2023) and *collaboration-based* approaches (Feng et al., 2024) are also explored to gauge LLM confidence through output variation from a single model or knowledge variation across multiple models. However, most existing approaches were proposed and evaluated with English only, while Figure 1 shows that multilingual abstention poses new challenges to existing solutions and leads to performance gaps based on language resourcesness. To mitigate this gap, we propose to teach LLMs to abstain by learning from multilingual feedback, where diverse feedback are generated in related languages to enhance reliability and make trustworthy abstain decisions.

Multilingual Factuality While early factuality studies were mostly conducted in English (Huang et al., 2023b; Zhang et al., 2023c; Ji et al., 2023), understanding and mitigating hallucinations beyond English is increasingly important for LLMs to equitably serve diverse populations (Liu et al., 2022; Lai et al., 2023; Xu et al., 2023a,b; Qi et al., 2023; Schott et al., 2023; Kang et al., 2024; Gao et al., 2024). In multilingual summarization, metrics and evaluations are proposed to quantify factual errors and utility (Aharoni et al., 2022; Qiu et al., 2023; Clark et al., 2023). In machine translation, faithfulness across diverse languages is also a critical concern (Lee et al., 2018; Raunak et al., 2021; Xu et al., 2023c; Dale et al.,

2023a,b). A diverse range of models (Lin et al., 2022b; Muennighoff et al., 2023; Lai et al., 2023; Üstün et al., 2024), datasets (Artetxe et al., 2020; Clark et al., 2020; Longpre et al., 2021; Chalkidis et al., 2022; Gehrmann et al., 2022; Ebrahimi et al., 2022; Li et al., 2022; Asai et al., 2023; Ogundepo et al., 2023; Ahuja et al., 2023; Wang et al., 2024), and studies on multilingual transfer (Lin et al., 2019a; Pires et al., 2019; Wu and Dredze, 2019; Karthikeyan et al., 2019; Wu et al., 2022; Fujinuma et al., 2022; Üstün et al., 2022; Schmidt et al., 2022; Asai et al., 2023; Philippy et al., 2023; Tanwar et al., 2023; Reusens et al., 2023; Li et al., 2024b; Gao et al., 2024) also contribute to the improvement of LLM factuality and utility beyond English. In this work, we present the first study on LLM abstaining in multilingual contexts and make an important step toward improving the reliability of multilingual LLMs and mitigating hallucinations.

7 Conclusion

We propose to improve the reliability of multilingual LLMs by abstaining via *multilingual feedback*, where LLMs generate feedback to their proposed answer in related languages for self-reflection. Extensive experiments demonstrate that *multilingual feedback* achieves up to 9.2% improvement against baselines across models and datasets, while presenting a more equitable solution to multilingual abstention. Further analysis reveals that abstention is a language-specific problem, that multilingual feedback in related languages both improves the accuracy of abstention and calibrates the fairness across higher- and lower-resource languages, and that cultural relatedness is an important factor in the utility and equity of abstention, highlighting that multilingual modeling is not only a technical problem but also a social-oriented one.

Limitations

Our study of teaching LLMs to abstain focuses on the *knowledge perspective*, i.e., LLMs should abstain when their parametric knowledge is insufficient to provide a correct answer. However, the abstain problem also has implications from the *safety perspective* (Huang et al., 2023b; Liu et al., 2023b). We envision future methodologies and evaluations that tackle both directions of the abstain problem across diverse language contexts.

Our approach, teaching LLMs to abstain via multilingual feedback, involves sampling multiple

feedback from related languages to promote self-reflection. This sampling introduces minor randomness in LLMs’ abstain decisions (Appendix A). In addition, it would incur greater inference costs than the most simple prompting approaches, but is also not the most expensive (Feng et al., 2024). When a black-box LLM with hundreds of billions of parameters is served behind an API call, our approach enables the incorporation of one extra multilingual 7B model for stronger reliability (Table 4) and does not add much to the overall cost.

Ethics Statement

While abstaining in multilingual contexts is a technical problem, we discover the role of culture in AbstainQA and that west-centric LLMs (Naous et al., 2023; Li et al., 2024a; Rao et al., 2024) are hindering progress on equitable LLM abstention (§5). This encourages research at the intersection of multilingualism and culture (Choenni et al., 2024). We envision future work on not only proposing technical solutions to the abstain problem, but also improving the representation of diverse values, perspectives, and cultures in LLMs.

Acknowledgements

We gratefully acknowledge support from the National Science Foundation under CAREER Grant No. IIS2142739, and NSF grants No. IIS2125201 and IIS2203097. This work was also supported in part by gift funding from Google and MSR.

References

Roe Aharoni, Shashi Narayan, Joshua Maynez, Jonathan Herzig, Elizabeth Clark, and Mirella Lapata. 2022. mface: Multilingual summarization with factual consistency evaluation. *arXiv preprint arXiv:2212.10622*.

Kabir Ahuja, Harshita Diddee, Rishav Hada, Milli- cent Ochieng, Krithika Ramesh, Prachi Jain, Akshay Nambi, Tanuja Ganu, Sameer Segal, Mohamed Ahmed, Kalika Bali, and Sunayana Sitaram. 2023. MEGA: Multilingual evaluation of generative AI. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.

Kabir Ahuja, Sunayana Sitaram, Sandipan Dandapat, and Monojit Choudhury. 2022. On the calibration of massively multilingual language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.

Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. On the cross-lingual transferability of monolingual representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637.

Akari Asai, Sneha Kudugunta, Xinyan Velocity Yu, Terra Blevins, Hila Gonen, Machel Reid, Yulia Tsvetkov, Sebastian Ruder, and Hannaneh Hajishirzi. 2023. Buffet: Benchmarking large language models for few-shot cross-lingual transfer. *arXiv preprint arXiv:2305.14857*.

Akari Asai, Xinyan Yu, Jungo Kasai, and Hanna Hajishirzi. 2021. One question answering model for many languages with cross-lingual dense passage retrieval. *Advances in Neural Information Processing Systems*, 34:7547–7560.

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.

Lucas Bandarkar, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer, and Madian Khabsa. 2023. The belebele benchmark: a parallel reading comprehension dataset in 122 language variants. *arXiv preprint arXiv:2308.16884*.

Vasilisa Bashlovkina, Zhaobin Kuang, Riley Matthews, Edward Clifford, Yennie Jun, William W Cohen, and Simon Baumgartner. 2023. Trusted source alignment in large language models. *arXiv preprint arXiv:2311.06697*.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*.

Ilias Chalkidis, Tommaso Pasini, Sheng Zhang, Letizia Tomada, Sebastian Schwemer, and Anders Søgaard. 2022. FairLex: A multilingual benchmark for evaluating fairness in legal text processing. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Jiefeng Chen, Jinsung Yoon, Sayna Ebrahimi, Sercan Arik, Tomas Pfister, and Somesh Jha. 2023. Adaptation with self-evaluation to improve selective prediction in LLMs. In *Findings of the Association for Computational Linguistics: EMNLP 2023*.

Rochelle Choenni, Anne Lauscher, and Ekaterina Shutova. 2024. The echoes of multilinguality: Tracing cultural value shifts during lm fine-tuning. *arXiv preprint arXiv:2405.12744*.

Elizabeth Clark, Shruti Rijhwani, Sebastian Gehrmann, Joshua Maynez, Roee Aharoni, Vitaly Nikolaev, Thibault Sellam, Aditya Siddhant, Dipanjan Das, and Ankur Parikh. 2023. SEAHORSE: A multilingual, multifaceted dataset for summarization evaluation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.

Jonathan H Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. Tydi qa: A benchmark for information-seeking question answering in typologically diverse languages. *Transactions of the Association for Computational Linguistics*, 8:454–470.

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

Jeremy Cole, Michael Zhang, Daniel Gillick, Julian Eisenschlos, Bhuwan Dhingra, and Jacob Eisenstein. 2023. Selectively answering ambiguous questions. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.

Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. *Advances in neural information processing systems*, 32.

Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. Xnli: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485.

David Dale, Elena Voita, Loïc Barrault, and Marta R Costa-jussà. 2023a. Detecting and mitigating hallucinations in machine translation: Model internal workings alone do well, sentence similarity even better. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 36–50.

David Dale, Elena Voita, Janice Lam, Prangthip Hansanti, Christophe Ropers, Elahe Kalbassi, Cynthia Gao, Loïc Barrault, and Marta Costa-jussà. 2023b. HalOmi: A manually annotated benchmark for multilingual hallucination and omission detection in machine translation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.

Shrey Desai and Greg Durrett. 2020. Calibration of pre-trained transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 295–302, Online. Association for Computational Linguistics.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. 2023. Improving factuality and reasoning in language models through multiagent debate. *arXiv preprint arXiv:2305.14325*.

Abteen Ebrahimi, Manuel Mager, Arturo Oncevay, Vishrav Chaudhary, Luis Chiruzzo, Angela Fan, John Ortega, Ricardo Ramos, Annette Rios, Ivan Vladimir Meza Ruiz, Gustavo Giménez-Lugo, Elisabeth Mager, Graham Neubig, Alexis Palmer, Rolando Coto-Solano, Thang Vu, and Katharina Kann. 2022. AmericasNLI: Evaluating zero-shot natural language understanding of pretrained multilingual models in truly low-resource languages. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. Understanding back-translation at scale. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500.

Shangbin Feng, Weijia Shi, Yuyang Bai, Vidhisha Balachandran, Tianxing He, and Yulia Tsvetkov. 2023. Knowledge card: Filling llms’ knowledge gaps with plug-in specialized language models. In *The Twelfth International Conference on Learning Representations*.

Shangbin Feng, Weijia Shi, Yike Wang, Wenxuan Ding, Vidhisha Balachandran, and Yulia Tsvetkov. 2024. Don’t hallucinate, abstain: Identifying llm knowledge gaps via multi-llm collaboration. *arXiv preprint arXiv:2402.00367*.

Yoshinari Fujinuma, Jordan Boyd-Graber, and Katharina Kann. 2022. Match the script, adapt if multilingual: Analyzing the effect of multilingual pretraining on cross-lingual transferability. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Changjiang Gao, Hongda Hu, Peng Hu, Jiajun Chen, Jixing Li, and Shujian Huang. 2024. Multilingual pre-training and instruction tuning improve cross-lingual knowledge alignment, but only shallowly. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*.

Sebastian Gehrmann, Abhik Bhattacharjee, Abinaya Mahendiran, Alex Wang, Alexandros Papangelis, Aman Madaan, Angelina Mcmillan-major, Anna Shvets, Ashish Upadhyay, Bernd Bohnet, Bingsheng Yao, Bryan Wilie, Chandra Bhagavatula, Chaobin You, Craig Thomson, Cristina Garbacea, Dakuo Wang, Daniel Deutsch, Deyi Xiong, Di Jin, Dimitra Gkatzia, Dragomir Radev, Elizabeth Clark, Esin Durmus, Faisal Ladhak, Filip Ginter, Genta Indra Winata, Hendrik Strobelt, Hiroaki Hayashi, Jekaterina Novikova, Jenna Kanerva, Jenny Chim, Jiawei Zhou, Jordan Clive, Joshua Maynez, João Sedoc, Juraj Juraska, Kaustubh Dhole, Khyathi Raghavi Chandu, Laura Perez Beltrachini, Leonardo F . R. Ribeiro, Lewis Tunstall, Li Zhang, Mahim Pushkarna, Mathias Creutz, Michael White, Mihir Sanjay Kale, Moussa Kamal Eddine, Nico Daheim, Nishant Subramani, Ondrej Dusek, Paul Pu Liang, Pawan Sasanka Ammanamanchi, Qi Zhu, Ratish Puduppully, Reno Kriz, Rifat Shahriyar, Ronald Cardenas, Saad Mhamood, Salomey Osei, Samuel Cahyawijaya, Sanja Štajner, Sébastien Montella, Shailza Jolly, Simon Mille, Tahmid Hasan, Tianhao Shen, Tosin Adewumi, Vikas Raunak, Vipul Raheja, Vitaly Nikolaev, Vivian Tsai, Yacine Jernite, Ying Xu, Yisi Sang, Yixin Liu, and Yufang Hou. 2022. GEMv2: Multilingual NLG benchmarking in a single line of code. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*.

Zhengyao Gu and Mark Hopkins. 2023. On the evaluation of neural selective prediction methods for natural language processing. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Guande He, Peng Cui, Jianfei Chen, Wenbo Hu, and Jun Zhu. 2023. Investigating uncertainty calibration of aligned language models under the multiple-choice setting. *arXiv preprint arXiv:2310.11732*.

Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *International Conference on Machine Learning*, pages 4411–4421. PMLR.

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

Xiaowei Huang, Wenjie Ruan, Wei Huang, Gao Jin, Yizhen Dong, Changshun Wu, Saddek Bensalem, Ronghui Mu, Yi Qi, Xingyu Zhao, Kaiwen Cai, Yang-hao Zhang, Sihao Wu, Peipei Xu, Dengyu Wu, André Freitas, and Mustafa A. Mustafa. 2023b. A survey of safety and trustworthiness of large language models through the lens of verification and validation. *ArXiv*, abs/2305.11391.

Zhiqi Huang, Puxuan Yu, and James Allan. 2023c. Improving cross-lingual information retrieval on low-resource languages via optimal transport distillation. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, pages 1048–1056.

Abhyuday Jagannatha and Hong Yu. 2020. Calibrating structured output predictors for natural language processing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.

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.

Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. 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.

Saurav Kadavath, Tom Conerly, Amanda Askell, T. J. Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zachary 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, John Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom B. Brown, Jack Clark, Nicholas Joseph, Benjamin Mann, Sam McCandlish, Christopher Olah, and Jared Kaplan. 2022. Language models (mostly) know what they know. *ArXiv*, abs/2207.05221.

Adam Tauman Kalai and Santosh S Vempala. 2023. Calibrated language models must hallucinate. *arXiv preprint arXiv:2311.14648*.

Amita Kamath, Robin Jia, and Percy Liang. 2020. Selective question answering under domain shift. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5684–5696, Online. Association for Computational Linguistics.

Haoqiang Kang, Terra Blevins, and Luke Zettlemoyer. 2024. Comparing hallucination detection metrics for multilingual generation. *arXiv preprint arXiv:2402.10496*.

K Karthikeyan, Zihan Wang, Stephen Mayhew, and Dan Roth. 2019. Cross-lingual ability of multilingual bert: An empirical study. In *International Conference on Learning Representations*.

Geunwoo Kim, Pierre Baldi, and Stephen McAleer. 2024. Language models can solve computer tasks. *Advances in Neural Information Processing Systems*, 36.

Lingkai Kong, Haoming Jiang, Yuchen Zhuang, Jie Lyu, Tuo Zhao, and Chao Zhang. 2020. Calibrated language model fine-tuning for in- and out-of-distribution data. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1326–1340, Online. Association for Computational Linguistics.

Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2022. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. In *The Eleventh International Conference on Learning Representations*.

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.

Viet Lai, Chien Nguyen, Nghia Ngo, Thuat Nguyen, Franck Dernoncourt, Ryan Rossi, and Thien Nguyen. 2023. Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 318–327.

Angeliki Lazaridou, Adhi Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Gimenez, Cyprien de Masson d’Autume, Tomas Kuncicky, Sebastian Ruder, et al. 2021. Mind the gap: Assessing temporal generalization in neural language models. *Advances in Neural Information Processing Systems*, 34:29348–29363.

Katherine Lee, Orhan Firat, Ashish Agarwal, Clara Fan-njiang, and David Sussillo. 2018. Hallucinations in neural machine translation.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kütter, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *NeurIPS*.

Cheng Li, Damien Teney, Linyi Yang, Qingsong Wen, Xing Xie, and Jindong Wang. 2024a. Culturepark: Boosting cross-cultural understanding in large language models. *arXiv preprint arXiv:2405.15145*.

Chong Li, Shaonan Wang, Jiajun Zhang, and Chengqing Zong. 2024b. Improving in-context learning of multilingual generative language models with cross-lingual alignment. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*.

Mingqi Li, Fei Ding, Dan Zhang, Long Cheng, Hongxin Hu, and Feng Luo. 2022. Multi-level distillation of semantic knowledge for pre-training multilingual language model. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.

Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. 2023. Holistic evaluation of language models. *Transactions on Machine Learning Research*.

Sheng-Chieh Lin, Amin Ahmad, and Jimmy Lin. 2023. mAggretreiver: A simple yet effective approach to zero-shot multilingual dense retrieval. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.

Stephanie Lin, Jacob Hilton, and Owain Evans. 2022a. Teaching models to express their uncertainty in words. *arXiv preprint arXiv:2205.14334*.

Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuhui Chen, Daniel Simig, Myle Ott, Naman Goyal, Shruti Bhosale, Jingfei Du, Ramakanth

Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona Diab, Veselin Stoyanov, and Xian Li. 2022b. Few-shot learning with multilingual generative language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.

Yu-Hsiang Lin, Chian-Yu Chen, Jean Lee, Zirui Li, Yuyan Zhang, Mengzhou Xia, Shruti Rijhwani, Junxian He, Zhisong Zhang, Xuezhe Ma, Antonios Anasatopoulos, Patrick Littell, and Graham Neubig. 2019a. Choosing transfer languages for cross-lingual learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.

Yu-Hsiang Lin, Chian-Yu Chen, Jean Lee, Zirui Li, Yuyan Zhang, Mengzhou Xia, Shruti Rijhwani, Junxian He, Zhisong Zhang, Xuezhe Ma, et al. 2019b. Choosing transfer languages for cross-lingual learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3125–3135.

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.

Linlin Liu, Xin Li, Ruidan He, Lidong Bing, Shafiq Joty, and Luo Si. 2022. Enhancing multilingual language model with massive multilingual knowledge triples. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.

Xin Liu, Muhammad Khalifa, and Lu Wang. 2023a. Litcab: Lightweight calibration of language models on outputs of varied lengths. *arXiv preprint arXiv:2310.19208*.

Yang Liu, Yuanshun Yao, Jean-Francois Ton, Xiaoying Zhang, Ruocheng Guo, Hao Cheng, Yegor Klochkov, Muhammad Faaiz Taufiq, and Hang Li. 2023b. Trustworthy llms: a survey and guideline for evaluating large language models’ alignment. In *Socially Responsible Language Modelling Research*.

Shayne Longpre, Yi Lu, and Joachim Daiber. 2021. Mkqa: A linguistically diverse benchmark for multilingual open domain question answering. *Transactions of the Association for Computational Linguistics*, 9:1389–1406.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. 2024. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36.

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

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. *arXiv preprint arXiv:2401.06855*.

Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, et al. 2023. Crosslingual generalization through multitask finetuning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15991–16111.

Tarek Naous, Michael J Ryan, Alan Ritter, and Wei Xu. 2023. Having beer after prayer? measuring cultural bias in large language models. *arXiv preprint arXiv:2305.14456*.

Odunayo Ogundepo, Tajuddeen Gwadabe, Clara Rivera, Jonathan H Clark, Sebastian Ruder, David Adelani, Bonaventure Dossou, Abdou Diop, Claytone Sika-sote, Gilles Hacheme, et al. 2023. Cross-lingual open-retrieval question answering for african languages. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14957–14972.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.

Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473.

Fred Philippy, Siwen Guo, and Shohreh Haddadan. 2023. Towards a common understanding of contributing factors for cross-lingual transfer in multilingual language models: A review. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. How multilingual is multilingual BERT? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.

Jirui Qi, Raquel Fernández, and Arianna Bisazza. 2023. Cross-lingual consistency of factual knowledge in multilingual language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.

Yifu Qiu, Yftah Ziser, Anna Korhonen, Edoardo Ponti, and Shay B Cohen. 2023. Detecting and mitigating hallucinations in multilingual summarisation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8914–8932.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.

Abhinav Rao, Akhila Yerukola, Vishwa Shah, Katharina Reinecke, and Maarten Sap. 2024. Normad: A benchmark for measuring the cultural adaptability of large language models. *arXiv preprint arXiv:2404.12464*.

Vikas Raunak, Arul Menezes, and Marcin Junczys-Dowmunt. 2021. The curious case of hallucinations in neural machine translation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1172–1183.

Manon Reusens, Philipp Borchert, Margot Mieskes, Jochen De Weerdt, and Bart Baesens. 2023. Investigating bias in multilingual language models: Cross-lingual transfer of debiasing techniques. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.

Fabian David Schmidt, Ivan Vulić, and Goran Glavaš. 2022. Don’t stop fine-tuning: On training regimes for few-shot cross-lingual transfer with multilingual language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.

Tim Schott, Daniel Furman, and Shreshta Bhat. 2023. Polyglot or not? measuring multilingual encyclopedic knowledge in foundation models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.

Tianhao Shen, Mingtong Liu, Ming Zhou, and Deyi Xiong. 2022. Recovering gold from black sand: Multilingual dense passage retrieval with hard and false negative samples. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.

Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Richard James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2024. REPLUG: Retrieval-augmented black-box language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik R Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Chenglei Si, Weijia Shi, Chen Zhao, Luke Zettlemoyer, and Jordan Boyd-Graber. 2023. Getting more out of mixture of language model reasoning experts. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8234–8249.

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.

Yueqi Song, Simran Khanuja, Pengfei Liu, Fahim Faisal, Alissa Ostapenko, Genta Winata, Alham Aji, Samuel Cahyawijaya, Yulia Tsvetkov, Antonios Anastasopoulos, et al. 2023. Globalbench: A benchmark for global progress in natural language processing. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14157–14171.

Elias Stengel-Eskin and Benjamin Van Durme. 2023. Calibrated interpretation: Confidence estimation in semantic parsing. *Transactions of the Association for Computational Linguistics*, 11:1213–1231.

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 *The 16th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*.

Meiqi Sun, Wilson Yan, Pieter Abbeel, and Igor Mordatch. 2022. Quantifying uncertainty in foundation models via ensembles. In *NeurIPS 2022 Workshop on Robustness in Sequence Modeling*.

Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. 2023. Aligning large multimodal models with factually augmented rlhf. *arXiv preprint arXiv:2309.14525*.

Eshaan Tanwar, Subhabrata Dutta, Manish Borthakur, and Tanmoy Chakraborty. 2023. Multilingual LLMs are better cross-lingual in-context learners with alignment. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Linwei Tao, Younan Zhu, Haolan Guo, Minjing Dong, and Chang Xu. 2023. A benchmark study on calibration. In *The Twelfth International Conference on Learning Representations*.

Nandan Thakur, Jianmo Ni, Gustavo Hernandez Abrego, John Wieting, Jimmy Lin, and Daniel Cer. 2024. Leveraging LLMs for synthesizing training data across many languages in multilingual dense retrieval. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*.

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.

Ahmet Üstün, Viraat Aryabumi, Zheng-Xin Yong, Wei-Yin Ko, Daniel D’souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, et al. 2024. Aya model: An instruction finetuned open-access multilingual language model. *arXiv preprint arXiv:2402.07827*.

Ahmet Üstün, Arianna Bisazza, Gosse Bouma, Gertjan van Noord, and Sebastian Ruder. 2022. Hyper-X: A unified hypernetwork for multi-task multilingual transfer. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.

Neeraj Varshney and Chitta Baral. 2023. Post-abstention: Towards reliably re-attempting the abstained instances in QA. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Bin Wang, Zhengyuan Liu, Xin Huang, Fangkai Jiao, Yang Ding, AiTi Aw, and Nancy Chen. 2024. SeaEval for multilingual foundation models: From cross-lingual alignment to cultural reasoning. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*.

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

Shuo Wang, Zhaopeng Tu, Shuming Shi, and Yang Liu. 2020a. On the inference calibration of neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3070–3079, Online. Association for Computational Linguistics.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*.

Yike Wang, Shangbin Feng, Heng Wang, Weijia Shi, Vidhisha Balachandran, Tianxing He, and Yulia Tsvetkov. 2023b. Resolving knowledge conflicts in large language models. *arXiv preprint arXiv:2310.00935*.

Zirui Wang, Zachary C. Lipton, and Yulia Tsvetkov. 2020b. On negative interference in multilingual models: Findings and a meta-learning treatment. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4438–4450. Association for Computational Linguistics.

Spencer Whitehead, Suzanne Petryk, Vedaad Shakib, Joseph Gonzalez, Trevor Darrell, Anna Rohrbach, and Marcus Rohrbach. 2022. Reliable visual question answering: Abstain rather than answer incorrectly. In *European Conference on Computer Vision*, pages 148–166. Springer.

John Wieting, Jonathan Clark, William Cohen, Graham Neubig, and Taylor Berg-Kirkpatrick. 2023. Beyond contrastive learning: A variational generative model for multilingual retrieval. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Linjuan Wu, Shaojuan Wu, Xiaowang Zhang, Deyi Xiong, Shizhan Chen, Zhiqiang Zhuang, and Zhiyong Feng. 2022. Learning disentangled semantic representations for zero-shot cross-lingual transfer in multilingual machine reading comprehension. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Shijie Wu and Mark Dredze. 2019. Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.

Jian Xie, Kai Zhang, Jiangjie Chen, Renze Lou, and Yu Su. 2023. Adaptive chameleon or stubborn sloth: Revealing the behavior of large language models in knowledge conflicts. In *The Twelfth International Conference on Learning Representations*.

Fangyuan Xu, Weijia Shi, and Eunsol Choi. 2024. RE-COMP: Improving retrieval-augmented LMs with context compression and selective augmentation. In *The Twelfth International Conference on Learning Representations*.

Haoran Xu, Weiting Tan, Shuyue Li, Yunmo Chen, Benjamin Van Durme, Philipp Koehn, and Kenton Murray. 2023a. Condensing multilingual knowledge with lightweight language-specific modules. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.

Shaoyang Xu, Junzhuo Li, and Deyi Xiong. 2023b. Language representation projection: Can we transfer factual knowledge across languages in multilingual language models? In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.

Weijia Xu, Sweta Agrawal, Eleftheria Briakou, Marianna J Martindale, and Marine Carpuat. 2023c. Understanding and detecting hallucinations in neural machine translation via model introspection. *Transactions of the Association for Computational Linguistics*, 11:546–564.

Yuqing Yang, Ethan Chern, Xipeng Qiu, Graham Neubig, and Pengfei Liu. 2023. Alignment for honesty. *arXiv preprint arXiv:2312.07000*.

Michihiro Yasunaga, Armen Aghajanyan, Weijia Shi, Richard James, Jure Leskovec, Percy Liang, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2023. Retrieval-augmented multimodal language modeling.

Jifan Yu, Xiaozhi Wang, Shangqing Tu, Shulin Cao, Daniel Zhang-Li, Xin Lv, Hao Peng, Zijun Yao, Xiaohan Zhang, Hanming Li, et al. 2023. Kola: Carefully benchmarking world knowledge of large language models. In *The Twelfth International Conference on Learning Representations*.

Polina Zablotskaia, Du Phan, Joshua Maynez, Shashi Narayan, Jie Ren, and Jeremiah Liu. 2023. On uncertainty calibration and selective generation in probabilistic neural summarization: A benchmark study. In *Findings of the Association for Computational Linguistics: EMNLP 2023*.

Hanning Zhang, Shizhe Diao, Yong Lin, Yi R Fung, Qing Lian, Xingyao Wang, Yangyi Chen, Heng Ji, and Tong Zhang. 2023a. R-tuning: Teaching large language models to refuse unknown questions. *arXiv preprint arXiv:2311.09677*.

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. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.

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

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.

Han Zhou, Xingchen Wan, Lev Proleev, Diana Mincu, Jilin Chen, Katherine Heller, and Subhrajit Roy. 2023a. Batch calibration: Rethinking calibration for in-context learning and prompt engineering. *arXiv preprint arXiv:2309.17249*.

Kaitlyn Zhou, Jena D Hwang, Xiang Ren, and Maarten Sap. 2024. Relying on the unreliable: The impact of language models' reluctance to express uncertainty. *arXiv preprint arXiv:2401.06730*.

Kaitlyn Zhou, Dan Jurafsky, and Tatsunori Hashimoto. 2023b. Navigating the grey area: How expressions of uncertainty and overconfidence affect language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5506–5524, Singapore. Association for Computational Linguistics.

A Analysis (cont.)

MULTI-RELATED helps abstaining in cross-lingual retrieval. When retrieval corpora are not readily available in low-resource languages, cross-lingual retrieval (Asai et al., 2021; Shen et al., 2022; Huang et al., 2023c; Wieting et al., 2023; Lin et al., 2023; Thakur et al., 2024) is often necessary for retrieval-augmented LLMs (Lewis et al., 2020; Shi et al., 2024; Yasunaga et al., 2023; Xu et al., 2024), where user queries are translated to high-resource languages and retrieval is performed with that language. We investigate whether our multilingual feedback approach works in this setting: we use English Wikipedia for retrieval⁵ and prepend back-translated paragraphs before the query from the seven low-resource languages. We evaluate various abstain approaches with CHATGPT and present performance in Figure 7. Our proposed *multilingual feedback* approach outperforms baselines for six of the seven low-resource languages, by 6.9% on average. This indicates that our *Multi-related* approach could also improve multilingual LLM reliability in retrieval-augmented settings.

FP and FN False positives refer to cases where the LLM should be able to provide the correct answer but abstained, while false negatives are cases where the LLM did not abstain but generated an incorrect answer. We present the false positive and false negative rates of MULTI-RELATED in Figure 8: we find that on high-resource languages, LLMs tend to be more “confident” and the FN is usually higher; for low-resource languages, LLMs tend to be more “conservative” and the FP is usually higher. We argue that having a high FP for low-resource languages is desirable since LLM has diminishing factuality on the long tail of languages, thus LLMs should be more cautious and abstain more.

Correlation between QA Performance and Abstain Performance We present the question answering accuracy as well as the abstain accuracy across various languages in Figure 9. We find that there is no lock-step synchronization between the two metrics, indicating that abstaining is an independent problem to question answering that needs further studies.

Another Dataset: Belebele Belebele (Bandarkar et al., 2023) is a multilingual reading comprehension benchmark featuring parallel questions

Method	Avg-H	Avg-M	Avg-L
PROBS	0.8028	<u>0.7550</u>	0.3604
TEMP	0.5988	0.5583	0.4331
ASK CALI.	0.4370	0.4336	<u>0.6163</u>
INSTRUCT	<u>0.8036</u>	0.7534	0.3704
REFLECT	0.5814	0.5231	0.4429
MOREINFO	0.7888	0.7430	0.3814
BACKTRANS.	0.6711	0.6366	0.4396
SCTHRES.	0.8136	0.7688	0.4634
CONFLICT	0.7072	0.6926	0.5741
MONO-NATIVE	0.6955	0.6774	0.4784
MONO-ENGLISH	0.6048	0.5891	0.6009
MULTI-RANDOM	0.7161	0.6807	0.4804
MULTI-RELATED	0.7906	0.7445	0.6386

Table 5: Performance averages for high, mid, and low-resource languages on Belebele (Avg-H, -M, and -L).

across 122 languages and variants. We evaluate baselines and our feedback-based approaches on Belebele and present the results in Table 5. MULTI-RELATED achieves the best performance on low-resource language, while falling behind the strongest baselines in this reading comprehension setting. This motivates using different methodologies for abstention in different language contexts.

Working Examples We conduct qualitative analysis to validate the generated feedback and abstain decisions. We specifically present several working examples in Tables 14, 15, and 16.

Standard Deviation Since MULTI-RELATED samples feedback from multiple languages, this sampling introduces randomness in the feedback content and potentially different abstain decisions. We re-run MULTI-RELATED three times with temperature $\tau = 0.7$, and we find that the standard deviation across runs is 0.0227, 0.0198, and 0.0086 for high, mid, and low-resource languages, indicating that the abstain performance is largely stable.

AbstainECE Aside from a binary decision of abstaining or answering, the probabilities of the abstain decision token (True/False) could be employed as an indicator for probabilistic abstention. We present the AbstainECE metric (Feng et al., 2024) in Table 6, which demonstrates that MONO-ENGLISH and MULTI-RELATED are stronger while the latter is best for low-resource languages. We envision improving LLM calibration with multilingual contexts could also help.

Number of feedback We employ from one to five pieces of feedback in MULTI-RANDOM and

⁵We retrieve Wikipedia with the WikiSearch API.

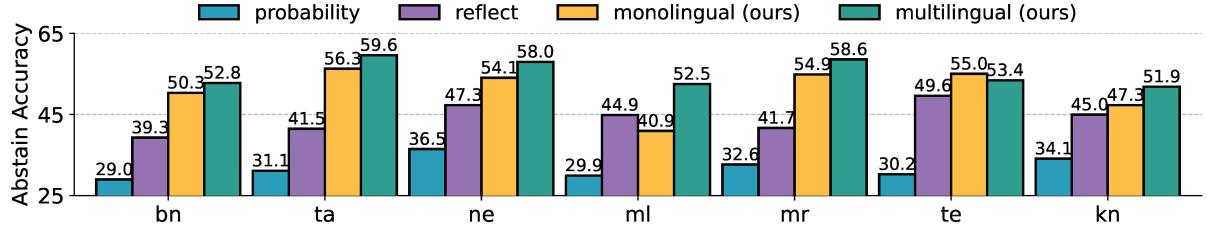


Figure 7: Abstain accuracy in the cross-lingual retrieval setting, where English Wikipedia is employed for retrieval to aid QA in low-resource languages. Multilingual feedback consistently produces more accurate abstain decisions in six of the seven low-resource languages.

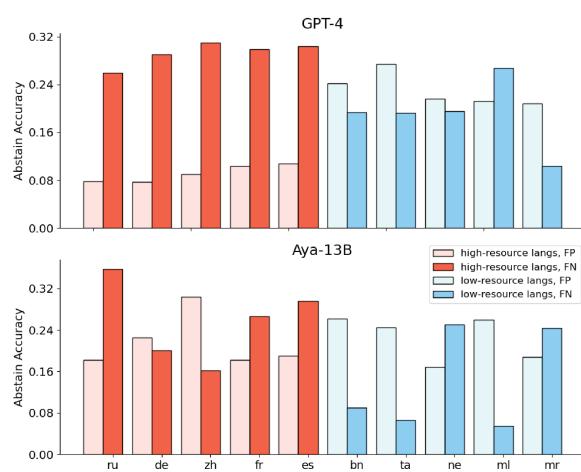


Figure 8: False positives and false negatives for MULTI-RELATED across low- and high-resource languages.

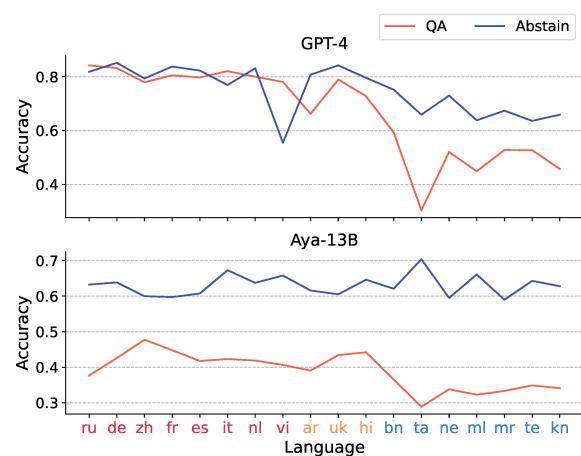


Figure 9: QA accuracy and abstain accuracy on the M-MMLU dataset with two LLMs. There is no lockstep synchronization between the two performance metrics, indicating that abstention is an independent research question. High-, mid-, and low-resource languages are labeled in red, orange, and blue colors, respectively.

Method	Avg-H	Avg-M	Avg-L
MONO-NATIVE	0.4594	0.4630	0.4276
MONO-ENGLISH	0.4410	0.4314	0.4114
MULTI-RANDOM	0.4713	0.4829	0.4475
MULTI-RELATED	0.4426	0.4476	0.3990

Table 6: AbstainECE averages for high, mid, and low-resource languages on Belebele (Avg-H, -M, and -L), the lower the better.

	1	2	3	4	5
Avg-H	0.5768	0.5878	0.5397	0.5856	0.5703
Avg-M	0.5484	0.5528	0.5435	0.5556	0.5501
Avg-L	0.4688	0.4959	0.5825	0.5003	0.5004

Table 7: Abstain accuracy with one to five feedback(s) with MULTI-RANDOM, AYA-13B, and M-MMLU.

	high	mid	low	equity (↓)
default	0.6411	0.5861	0.4432	0.0943
reverse	0.6285	0.5297	0.4184	0.1294
self-included	0.6055	0.5638	0.4429	0.0743
lang var.	0.5790	0.4488	0.3726	0.0964

Table 8: Ablation settings of MULTI-RELATED, with ChatGPT and M-MMLU.

report performance in Table 7. We find that performance saturates with three pieces of feedback, while it also fluctuates across language resource-ness levels.

Other Metrics In addition to abstain accuracy, we additionally report two more AbstainQA metrics, reliable accuracy (Feng et al., 2024) and effective reliability (Whitehead et al., 2022; Si et al., 2023) in Table 9. MULTI-RELATED achieves the best performance in four of the six settings.

Ablation Settings We present several ablation setting of MULTI-RELATED: 1) *reverse*, where the most distant languages are selected for feedback generation; 2) *self-included*, where the language of

	Reliable Accuracy			Effective Reliability		
	Avg-H	Avg-M	Avg-L	Avg-H	Avg-M	Avg-L
PROBS	0.4772	0.4800	0.3868	-0.0269	-0.0232	-0.1220
TEMP	0.4635	0.4651	0.3694	-0.0472	-0.0442	-0.1705
ASK CAL.	0.5297	0.5207	0.4012	0.0209	0.0143	-0.1069
INSTRUCT	0.4255	0.4256	0.3477	-0.1412	-0.1404	-0.2538
REFLECT	0.4069	0.4019	0.3364	-0.1843	-0.1944	-0.3260
BACKTRANS.	0.4277	0.4198	0.3517	-0.1292	-0.1445	-0.2648
SCTHRES.	0.5389	0.5254	0.3878	0.0260	0.0172	-0.0547
CONFLICT	0.4585	0.4559	0.3636	-0.0316	-0.0338	-0.1004
MONO-NATIVE	0.4333	0.4437	0.3973	-0.0705	-0.0604	-0.0821
MONO-ENGLISH	0.4796	0.4594	0.3884	-0.0242	-0.0406	-0.0695
MULTI-RANDOM	0.4565	0.4376	0.3640	-0.0344	-0.0416	-0.0656
MULTI-RELATED	0.5402	0.4973	0.4474	0.0279	-0.0077	-0.0289

Table 9: Other AbstainQA metrics with AYA-13B and MMLU.

the question is also employed to generate feedback; 3) *lang var.*, where the feedback content stays the same but translated to different related languages. Table 8 demonstrates that the default setting often works best, while including the original language for feedback generation could be beneficial for certain cases.

Another Interpretation of Abstain Overlap
For Figure 5, another way is to compare the proportion of consensus, where LLMs abstain for zero or all three of languages. In this definition, the same conclusion still holds: the first control group has 23.1% vs. 20.9%, while the second control group has 32.1% vs. 16.2%.

Randomness in Sampling Feedback We randomly sample feedback sets with temperature of 1 and repeat for 3 runs. If the LLM abstains/answers in all 3 runs, then it is deemed consistent; 1:2 and 2:1 scenarios are then deemed as inconsistent. We present results in Table 10, showing that learning to abstain from multilingual feedback is largely consistent.

	bn	ml	mr	ne	ta	te
MMLU, consistent	103	101	99	117	114	95
MMLU, inconsistent	42	26	35	31	21	34
consistent rate %	71.03	79.53	73.88	79.05	84.44	73.64
Hellaswag, consistent	115	136	115	131	122	122
Hellaswag, inconsistent	39	14	39	24	18	25
consistent rate %	74.68	90.67	74.68	84.52	87.14	82.99

Table 10: Consistency when repeating the feedback sampling for three times.

B Experiment Details

Dataset Details We employ M-MMLU, M-Hellaswag (Lai et al., 2023), and Belebele (Bandardarkar et al., 2023) as evaluations of multilingual AbstainQA. Specifically, we adopt 26 languages beyond English: 8 high-resource (Russian, ru; German, de; Chinese, zh; French, fr; Spanish, es; Italian, it; Dutch, nl; Vietnamese, vi), 11 mid-resource languages (Indonesian, id; Arabic, ar; Hungarian, hu; Romanian, ro; Danish, da; Slovak, sk; Ukrainian, uk; Catalan, ca; Serbian, sr; Croatian, hr; Hindi, hi), and 7 low-resource languages (Bengali, bn; Tamil, ta; Nepali, ne; Malayalam, ml; Marathi, Mr; Telugu, te; Kannada, kn). We follow the definition of language resourceness based on pretraining data frequency in Lai et al. (2023). We randomly sample 200 questions for validation and 800 questions for test from each language, with minor variation across languages based on data availability.

Model Details We employ the “CohereForAI/aya-101” model checkpoint on Huggingface for AYA-13B, and the Azure OpenAI API checkpoint of “gpt4” for GPT-4, and the “GPT-3.5-TURBO-INSTRUCT” model checkpoint on OpenAI API for CHATGPT.

Baseline Details We refer readers to Feng et al. (2024) for a complete description of baselines. For the additional BACKTRANSLATION baseline, we translate the question to English and make an abstain decision in English, then use that abstain decision for other languages.

GPT-4 Evaluation Details For quality evaluation, we employ “*Question: <question> Proposed Answer: <answer> Feedback 1: <feedback> Feedback 2: <feedback> Which feedback is more relevant to the question?*” and “*Question: <question> Proposed Answer: <answer> Feedback 1: <feedback> Feedback 2: <feedback> Which feedback is more informative?*”. For role evaluation, we employ “*Question: <question> Proposed Answer: <answer> Feedback 1: <feedback> Feedback 2: <feedback> Feedback 3: <feedback> What is the relationship among the three feedbacks? A. similar B. complementary C. conflicting D. unrelated Relationship:*”.

Implementation Details We present the related languages employed for feedback generation in the Language Relatedness study (§5) in Tables 11, 12, and 13.

```
default: { "en": ["German", "Dutch", "French"], "ru": ["Ukrainian", "Romanian", "Catalan"], "de": ["Dutch", "English", "French"], "zh": ["Arabic", "Slovak", "Danish"], "fr": ["Catalan", "German", "Spanish"], "es": ["Catalan", "Romanian", "French"], "it": ["Catalan", "Romanian", "Ukrainian"], "nl": ["German", "Italian", "Ukrainian"], "vi": ["Indonesian", "English", "Bengali"], "id": ["Vietnamese", "Catalan", "Russian"], "ar": ["Chinese", "Slovak", "Danish"], "hu": ["Romanian", "German", "French"], "ro": ["Catalan", "Italian", "Spanish"], "da": ["Slovak", "Dutch", "Ukrainian"], "sk": ["Chinese", "Arabic", "Danish"], "uk": ["Russian", "Italian", "Croatian"], "ca": ["Romanian", "Spanish", "Italian"], "sr": ["Slovak", "Danish", "Croatian"], "hr": ["Ukrainian", "Italian", "Dutch"], "hi": ["Bengali", "Talugu", "Marathi"], "bn": ["Hindi", "Telugu", "Nepali"], "ta": ["Malayalam", "Marathi", "Kannada"], "ne": ["Kanaada", "Telugu", "Hindi"], "ml": ["Tamil", "Marathi", "Kannada"], "mr": ["Tamil", "Malayalam", "Hindi"], "te": ["Kannada", "Tamil", "Nepali"], "kn": ["Telugu", "Malaayalam", "Tamil"] }
```

```
syntactic: "en": ["Spanish", "German", "French"], "ru": ["Ukrainian", "German", "Spanish"], "de": ["Dutch", "English", "Russian"], "zh": ["Arabic", "Slovak", "Hungarian"], "fr": ["Spanish", "English", "German"], "es": ["English", "French", "Russian"], "it": ["Catalan", "Romanian", "Dutch"], "nl": ["German", "Italian", "Danish"], "vi": ["Indonesian", "English", "French"], "id": ["Vietnamese", "English", "Italian"], "ar": ["Chinese", "Slovak", "Hungarian"], "hu": ["Russian", "Italian", "Romanian"], "ro": ["Italian", "Ukrainian", "Spanish"], "da": ["Dutch", "German", "French"], "sk": ["Chinese", "Arabic", "Hungarian"], "uk": ["Russian", "Italian", "Romanian"], "ca": ["Italian", "Dutch", "Romanian"], "sr": ["Catalan", "Ukrainian", "German"], "hr": ["Serbian", "Vietnamese", "Danish"], "hi": ["Kannada", "Russian", "Ukrainian"], "bn": ["Marathi", "Hindi", "Tamil"], "ta": ["Telugu", "Kannada", "Marathi"], "ne": ["Kannada", "Telugu", "Hindi"], "ml": ["Telugu", "Kannada", "Tamil"], "mr": ["Tamil", "Bengali", "Telugu"], "te": ["Tamil", "Nepali", "Kannada"], "kn": ["Tamil", "Nepali", "Hindi"],
```

```
featural: "en": ["German", "Russian", "French"], "ru": ["Romanian", "Ukrainian", "English"], "de": ["English", "French", "Dutch"], "zh": ["Arabic", "Slovak", "English"], "fr": ["German", "English", "Russian"], "es": ["English", "Russian", "French"], "it": ["Dutch", "Romanian", "Ukrainian"], "nl": ["German", "Italian", "English"], "vi": ["Indonesian", "English", "French"], "id": ["Vietnamese", "Catalan", "English"], "ar": ["Chinese", "Slovak", "English"], "hu": ["Rominian", "English", "Russian"], "ro": ["Russian", "Italian", "Hungarian"], "da": ["Serbian", "English", "Russian"], "sk": ["Chinese", "Arabic", "English"], "uk": ["Russian", "Italian", "Romanian"], "ca": ["Italian", "Dutch", "Romanian"], "sr": ["Danish", "Russian", "Spanish"], "hr": ["Catalan", "English", "Russian"], "hi": ["Bengali", "Nepali", "Telugu"], "bn": ["Hindi", "Nepali", "Telugu"], "ta": ["Malayalam", "Marathi", "Telugu"], "ne": ["Hindi", "Bengali", "Marathi"], "ml": ["Tamil", "Marathi", "Kannada"], "mr": ["Tamil", "Nepali", "Malayalam"], "te": ["Hindi", "Bengali", "Tamil"], "kn": ["Hindi", "Tamil", "Nepali"],
```

```
genetic: "en": ["German", "Dutch", "Danish"], "ru": ["Ukrainian", "Slovak", "Serbian"], "de": ["Dutch", "English", "Danish"], "zh": ["English", "Russian", "German"], "fr": ["Spanish", "Catalan", "Italian"], "es": ["Catalan", "Romanian", "French"], "it": ["Romanian", "Catalan", "Spanish"], "nl": ["German", "English", "Danish"], "vi": ["English", "Russian", "German"], "id": ["English", "Russian", "German"], "ar": ["English", "Russian", "German"], "hu": ["English", "Russian", "German"], "ro": ["Spanish", "Italian", "Catalan"], "da": ["German", "English", "Dutch"], "sk": ["Russian", "Ukrainian", "Serbian"], "uk": ["Russian", "Slovak", "Serbian"], "ca": ["Spanish", "Romanian", "Italian"], "sr": ["Croatian", "Russian", "Ukrainian"], "hr": ["Serbian", "Russian", "Slovak"], "hi": ["Bengali", "Marathi", "German"], "bn": ["Hindi", "Marathi", "English"], "ta": ["Malayalam", "Kannada", "Telugu"], "ne": ["English", "Russian", "German"], "ml": ["Tamil", "Kannada", "Telugu"], "mr": ["Hindi", "Bengali", "Russian"], "te": ["Tamil", "Malayalam", "Kannada"], "kn": ["Malayalam", "Tamil", "Telugu"],
```

Table 11: Related languages across different method settings, part 1.

geographic: "en": ["French", "Dutch", "Danish"], "ru": ["English", "German", "French"], "de": ["French", "Italian", "Dutch"], "zh": ["English", "Russian", "German"], "fr": ["English", "German", "Spanish"], "es": ["French", "Catalan", "English"], "it": ["German", "French", "Hungarian"], "nl": ["English", "German", "French"], "vi": ["Indonesian", "Bengali", "Nepali"], "id": ["Vietnamese", "Bengali", "Tamil"], "ar": ["English", "Russian", "German"], "hu": ["German", "Italian", "Romanian"], "ro": ["German", "Italian", "Hungarian"], "da": ["English", "German", "French"], "sk": ["German", "Italian", "Hungarian"], "uk": ["German", "Hungarian", "Romanian"], "ca": ["French", "Spanish", "Italian"], "sr": ["German", "Italian", "Hungarian"], "hr": ["German", "Italian", "Hungarian"], "hi": ["Nepali", "Marathi", "Telugu"], "bn": ["Nepali", "Vietnamese", "Hindi"], "ta": ["Malayalam", "Marathi", "Telugu"], "ne": ["Hindi", "Bengali", "Vietnamese"], "ml": ["Tamil", "Marathi", "Telugu"], "mr": ["Hindi", "Tamil", "Malayalam"], "te": ["Hindi", "Tamil", "Malayalam"], "kn": ["Tamil", "Malayalam", "Marathi"],

inventory: "en": ["German", "Marathi", "Telugu"], "ru": ["Ukrainian", "Croatian", "Romanian"], "de": ["Dutch", "French", "English"], "zh": ["Arabic", "Danish", "Slovak"], "fr": ["Hungarian", "Dutch", "German"], "es": ["Hungarian", "German", "Indonesian"], "it": ["Catalan", "Romanian", "Ukrainian"], "nl": ["German", "French", "Hungarian"], "vi": ["English", "Dutch", "German"], "id": ["Catalan", "Croatian", "Romanian"], "ar": ["Chinese", "Danish", "Slovak"], "hu": ["French", "Romanian", "Italian"], "ro": ["Ukrainian", "Catalan", "Italian"], "da": ["Chinese", "Arabic", "Slovak"], "sk": ["Chinese", "Arabic", "Danish"], "uk": ["Romanian", "Russian", "Italian"], "ca": ["Indonesian", "Italian", "Romanian"], "sr": ["Chinese", "Arabic", "Danish"], "hr": ["Catalan", "Indonesian", "Hungarian"], "hi": ["Telugu", "Bengali", "Nepali"], "bn": ["Telugu", "Nepali", "Hindi"], "ta": ["Kannada", "Malayalam", "Marathi"], "ne": ["Marathi", "Bengali", "Kannada"], "ml": ["Kannada", "Marathi", "Tamil"], "mr": ["Kannada", "Malayalam", "Nepali"], "te": ["Hindi", "Bengali", "Nepali"], "kn": ["Malayalam", "Marathi", "Tamil"],

phonological: "en": ["Indonesian", "Russian", "Catalan"], "ru": ["Catalan", "Hungarian", "Hindi"], "de": ["French", "Hungarian", "English"], "zh": ["Italian", "Dutch", "Arabic"], "fr": ["German", "Hungarian", "Hindi"], "es": ["English", "Russian", "Catalan"], "it": ["Chinese", "Dutch", "Arabic"], "nl": ["Chinese", "Italian", "Arabic"], "vi": ["Indonesian", "English", "Russian"], "id": ["English", "Russian", "Catalan"], "ar": ["Chinese", "Italian", "Dutch"], "hu": ["Russian", "Catalan", "German"], "ro": ["Russian", "Catalan", "German"], "da": ["Chinese", "Italian", "Dutch"], "sk": ["Chinese", "Italian", "Dutch"], "uk": ["Chinese", "Italian", "Dutch"], "ca": ["Russian", "Hungarian", "Hindi"], "sr": ["Spanish", "Chinese", "Italian"], "hr": ["Chinese", "Italian", "Dutch"], "hi": ["Russian", "Catalan", "French"], "bn": ["Telugu", "Kannada", "Russian"], "ta": ["Chinese", "Italian", "Dutch"], "ne": ["Romanian", "Telugu", "Kannada"], "ml": ["Chinese", "Italian", "Dutch"], "mr": ["Chinese", "Italian", "Dutch"], "te": ["Kannada", "Russian", "Catalan"], "kn": ["Kannada", "Russian", "Catalan"],

Table 12: Related languages across different method settings, part 2.

WVS: "en": ["English", "English", "English"], "ru": ["Ukrainian", "Romanian", "Russian"], "de": ["German", "Dutch", "Danish"], "zh": ["Chinese", "Chinese", "Chinese"], "fr": ["French", "Slovak", "Hungarian"], "es": ["French", "Slovak", "Hungarian"], "it": ["French", "Slovak", "Hungarian"], "nl": ["German", "Dutch", "Danish"], "vi": ["Vietnamese", "Vietnamese", "Vietnamese"], "id": ["Indonesian", "Indonesian", "Indonesian"], "ar": ["Arabic", "Hindi", "Bengali"], "hu": ["French", "Slovak", "Hungarian"], "ro": ["Ukrainian", "Romanian", "Russian"], "da": ["German", "Dutch", "Danish"], "sk": ["French", "Slovak", "Hungarian"], "uk": ["Ukrainian", "Romanian", "Russian"], "ca": ["Catalan", "Catalan", "Catalan"], "sr": ["Serbian", "Serbian", "Serbian"], "hr": ["French", "Slovak", "Hungarian"], "hi": ["Arabic", "Hindi", "Bengali"], "bn": ["Arabic", "Hindi", "Bengali"], "ta": ["Arabic", "Hindi", "Bengali"], "ne": ["Arabic", "Hindi", "Bengali"], "ml": ["Arabic", "Hindi", "Bengali"], "mr": ["Arabic", "Hindi", "Bengali"], "te": ["Arabic", "Hindi", "Bengali"], "kn": ["Arabic", "Hindi", "Bengali"],

LLM-generated: "en": ["Frisian", "Dutch", "German"], "ru": ["Belarusian", "Ukrainian", "Rusyn"], "de": ["Dutch", "Luxembourgish", "Yiddish"], "zh": ["Cantonese", "Shanghainese", "Hokkien"], "fr": ["Italian", "Spanish", "Catalan"], "es": ["Portuguese", "Catalan", "Italian"], "it": ["Sicilian", "Neapolitan", "Tuscan"], "nl": ["Afrikaans", "Frisian", "German"], "vi": ["Muong", "Khmer", "Mon"], "id": ["Malay", "Minangkabau", "Javanese"], "ar": ["Hebrew", "Aramaic", "Amharic"], "hu": ["Finnish", "Estonian", "Udmurt"], "ro": ["Italian", "Spanish", "French"], "da": ["Swedish", "Norwegian", "Icelandic"], "sk": ["Czech", "Polish", "Slovene"], "uk": ["Russian", "Belarusian", "Rusyn"], "ca": ["Occitan", "Spanish", "French"], "sr": ["Croatian", "Bosnian", "Montenegrin"], "hr": ["Serbian", "Bosnian", "Montenegrin"], "hi": ["Urdu", "Punjabi", "Bengali"], "bn": ["Assamese", "Odia", "Maithili"], "ta": ["Kannada", "Telugu", "Malayalam"], "ne": ["Maithili", "Bhojpuri", "Awadhi"], "ml": ["Tamil", "Tulu", "Kannada"], "mr": ["Konkani", "Sanskrit", "Gujarati"], "te": ["Kannada", "Tamil", "Malayalam"], "kn": ["Telugu", "Tamil", "Tulu"],

Table 13: Related languages across different method settings, part 3.

Question: 非营利组织活动报告表的主要目的是提供相关信息给其 A: 资源提供者 B: 管理者 C: 受益人 D: 州监管机构

(translated: The main purpose of the non-profit organization activity report form is to provide relevant information to its A: resource providers. B: managers. C: beneficiaries. D: state regulatory agencies.)

Proposed Answer: C

Feedback 1:

في حين أن تقديم المعلومات للمستفيدين هو هدف رئيسي لتقرير الأنشطة لمنظمة غير ربحية، إلا أنه ليس الهدف الوحيد. يهدف التقرير أيضاً إلى تعزيز الشفافية، جذب المתרبيين والشركاء، وتقدير الأداء لتحقيق التحسين المستمر. من المهم تحديد المستفيدين، الذين قد يشملون الأعضاء، المتربيين، الشركاء، والمجتمع العام، حيث أن لكل فئة احتياجات معلوماتية مختلفة. يجب أن يكون التقرير شاملًا، يغطي المعلومات المالية، إنجازات البرنامج، التحديات المستقلة، والخطط، ويجب أن يتم تقديمها بفعالية باستخدام لغة واضحة وملخصات ونقاط رئيسية. هذا النهج لا يقتصر على تقديم المعلومات فحسب، بل يعزز أيضًا مصداقية المنظمة ويزيد من الثقة بين أصحاب المصلحة.

(translated: While providing information to beneficiaries is a key goal of a non-profit organization's activity report, *it is not the sole purpose*. The report also aims to enhance transparency, attract donors and partners, and evaluate performance for continuous improvement. Specifying the beneficiaries, who may include members, donors, partners, and the general community, is essential as each group has different informational needs.)

Feedback 2:

Okrem informovania prijímateľov správa často slúži aj na preukádzanie transparentnosti organizácie voči donorom, regulačným orgánom, dobrovoľníkom a verejnosti. Pomáha tiež pri hodnotení efektívnosti a dosiahnutých výsledkov, čo môže podporiť ďalší rozvoj organizácie a jej projektov. Celkovo je teda cieľom správy nielen informovať, ale aj budovať dôveru a zodpovednosť voči všetkým zainteresovaným stranám.

(translated: In addition to informing beneficiaries, *the report often serves to demonstrate the organization's transparency to donors*, regulatory bodies, volunteers, and the public. It also helps in evaluating the effectiveness and outcomes achieved, which can support the further development of the organization and its projects. Overall, the report's purpose is not only to inform but also to build trust and accountability towards all stakeholders involved.)

Feedback 3:

非营利组织的主要目的之一是,向其受益人提供有关组织活动的反馈信息,主要依据组织的运作特征. 受益人包括项目组织的利益相关者,如基金管理人,或组织会员,如捐赠者。非营利组织的负责人应该遵守对组织的监管规定。

(translated: One of the main purposes of non-profit organizations is to provide feedback to their beneficiaries about the organization's activities, based primarily on the operational characteristics of the organization. Beneficiaries include stakeholders of project organizations, such as fund managers, or *organizational members, such as donors*. The leaders of non-profit organizations should comply with regulatory requirements for the organization.)

Abstain: Yes (correct answer is A)

Table 14: Working example one, where the three pieces of feedback unanimously point out the error in the answer.

Question: 以下哪个提供了有关地球核心、地幔和下地壳结构的大部分信息？A: 测量地球磁场的强度和波动B: 潜流岩的检测C: 从深部钻孔中收集样本D: 研究通过地球的地震波传播的速度和路径

(translated: Which of the following provides most of the information about the structure of the Earth's core, mantle, and lower crust? A: Measuring the strength and fluctuations of the Earth's magnetic field B: Detecting plumes of molten rock C: Collecting samples from deep drilling D: Studying the speed and paths of seismic waves passing through the Earth)

Proposed Answer: C

Feedback 1:

وتعتبر الاستكشافات الجيولوجية إحدى الطرق الرئيسية في المعرفة بشكل كبير عن كيفية تفاعل طبقات الخارج والداخل الجيولوجي مع بعضها البعض، باستخدام تقنيات الاستكشاف الجيولوجي (أو الاستكشاف الجيولوجي) وخاصة استكشاف الغلاف الخارجي للقشرة الأرضية ((وهذا ما يحدث في الحفر العميق)) والحفاظ على النموذج الجيولوجي المألوف للعين البشرية.

(translated: Geological exploration is considered one of the main ways to significantly increase our knowledge about how the layers of the Earth's exterior and interior interact with each other, using geological exploration techniques, especially exploring the outer shell of the Earth's crust (as in deep drilling) and maintaining the familiar geological model for the human eye.)

Feedback 2:

Materiály z inžinierskej činnosti sa dokážu zaujímať o štruktúru zeme vďaka spôsobu, akým tvoria základne zemského plášťa, tiež o spôsobe používania jadrovej energie, uvoľňovania energie z dopadu gravitačných silov. Predpokladáme pôsobenie gravitácie na štruktúru centrálneho a periférneho plášťa, rovnako ako aj gravitačné účinky uvoľnenej energie medzi nimi. Na základe výsledkov dokážeme predpokladať, ako sa skladá zloženie povrchovej vrstvy, ako ovplyvňujú polárnu tepelnú uhl'ovodíkovú vrstvu, ako ovplyvňuje povrchovú vlhkosť, a ako ovplyvňuje teplotnú pozíciu rôznych vrstiev.

(translated: Materials from engineering activities can be interested in the structure of the earth due to the way they form the foundations of the earth's mantle, as well as the methods of using nuclear energy and releasing energy from the impact of gravitational forces. We assume the influence of gravity on the structure of the central and peripheral mantle, as well as the gravitational effects of the released energy between them. Based on the results, we can predict how the composition of the surface layer is formed, how it affects the polar thermal hydrocarbon layer, how it affects surface moisture, and how it affects the temperature position of different layers.)

Feedback 3:

Undersøgelse af jordens magnetfelt og dens bevægelser, samt måling af jordoverfladens temperatur og fugtighed er de mest anvendte metoder. Undersøgelsen af jordens magnetfelt og dens bevægelser er ikke alene vigtigt for forståelsen af jordens magnetfelt, men også påvirker andre områder, såsom jordskælv og vulkaner, hvor man kan studere jordens magnetfelt, mens man i samme tid observerer andre systemer, såsom atmosfære og havanlæg.

(translated: The investigation of the Earth's magnetic field and its movements, as well as the measurement of surface temperature and humidity, are the most commonly used methods. The study of the Earth's magnetic field and its movements is not only important for understanding the Earth's magnetic field, but also affects other areas, such as earthquakes and volcanoes, where one can study the Earth's magnetic field while observing other systems, such as the atmosphere and oceanic structures, at the same time.)

Abstain: Yes (correct answer is D)

Table 15: Working example two, where there is a conflict among the three feedback.

Question: 关于总统提名大会代表，以下哪个陈述是正确的？A: 代表通常更有可能在某个时候注册为第三方投票者。B: 大多数普通人比代表更有教育背景。C: 代表通常对政治没有太大兴趣。D: 通常，普通人的意识形态比代表更少。

(translated: Which statement is correct about delegates to the presidential nominating convention? A: Delegates are more likely to register as third-party voters at some point. B: Most ordinary people have a higher level of education than delegates. C: Delegates generally have less interest in politics. D: Typically, ordinary people have less ideology than delegates.)

Proposed Answer: C

Feedback 1:

الناس العاديون معقدون ويتختلفون في مشاعرهم ومزاجهم. لديهم أيضًا الكثير من التفكير السلبي ويكونون أقل تحفيراً سياسياً. في المقابل، يلجأ هذا النوع من الأشخاص إلى التصويت شخصياً وتلبيд المرشحين لشنغ بعض المقاعد كافية في الانتخابات الرئاسية. تظهر الدراسة أن غالبية الأفراد المشاركين في العملية الانتخابية يجدون هذا الوضع مثيراً للاهتمام من حيث التعاطف والتقبل. يختلف الوضع من حيث التفكير السلبي. عادةً ما يعتقد الناس أن هذا الوضع لم يكن بسبب مشاعرهم تجاه السياسة، بل بسبب نقص مشاركتهم.

(translated: *Ordinary people are complex and differ in their feelings and moods.* They also have a lot of negative thinking and are *less politically motivated*. In contrast, this type of people resort to voting personally and endorse candidates to hold some seats as a minority in the presidential candidate. The study shows that the majority of individuals participating in the electoral process find this situation interesting in terms of empathy and embrace. It differs in terms of negative thinking. People usually believe that this situation was not due to their feelings towards politics, but due to their lack of participation.)

Feedback 2:

Zvyčajne je obyčajný človek, ktorý podporuje kandidátov z celého štátu, zvyčajne ľudom plne neznámy. Umožňujú voličom hlasovať za svojho kandidáta a predstavujú ho podľa svojho názoru. Keďže na zvolenie kandidáta je potrebný nejaký druh spoločného zvolenia, získali ďalší hlas členovia zvoliteľskej delegácie a títo sú usilujúci o účasť na zvolovaní. O takéto pôsobenie sa postará delegát zvoliteľskej delegácie, ktorý je ľudom plne neznámy.

(translated: Usually, an ordinary person who supports candidates from all over the state, usually to people completely unknown. They allow voters to vote for their candidate and represent him according to their opinion. Since some kind of common election is needed to elect a candidate, additional votes were obtained by members of the elector's delegation, who are seeking to participate in the election. A delegate of the elector's delegation, who is completely unknown to people, will take care of this action.)

Feedback 3:

Delegererne fra staten har ofte mere viden om politik end de fleste almindelige mennesker.
(translated: The state's delegates often have *more knowledge about politics than most ordinary people*.)

Abstain: Yes (correct answer is D)

Table 16: Working example three, where there is a conflict among the three feedback.