CaLMQA: Exploring culturally specific long-form question answering across 23 languages

Anonymous ACL submission

Abstract

Despite rising interest in measuring the cultural knowledge of large language models (LLMs), their ability to generate *long-form* answers to culturally specific questions remains unexplored. To fill this gap, we perform the first study of long-form cultural QA by introducing CALMQA, a dataset of 2.6K questions across 23 different languages that contains two high-level question categories: *culturally specific* and *culturally agnostic*. To collect culturally-specific questions, which we define as those more likely to be asked by people from cultures associated with the question's language, we crawl naturally-occurring questions from community web forums and also hire native speakers of under-resourced, rarely-studied languages such as Fijian and Kirundi to write questions. We evaluate many aspects of LLM-generated long-form answers, finding that (1) for many languages, even the best models make critical surface-level errors (e.g., answering in the wrong language, repetition), especially for low-resource languages; and (2) answers to culturally-specific questions contain more factual errors than answers to culturally-agnostic questions. In addition to these insights, our paper develops data collection and evaluation methodologies for cultural long-form QA that we hope will facilitate future research.

1 Introduction

While large language models (LLMs) are increasingly employed by users across the world, growing evidence reveals significant gaps in their understanding of diverse cultures (Tao et al., 2024; Li et al., 2024) and a persistent bias toward Western-centric perspectives (Palta and Rudinger, 2023; Durmus et al., 2024; AlKhamissi et al., 2024; Naous et al., 2024a). Existing research on these shortcomings has been limited to prompting LLMs about cultural knowledge in English (Rao et al., 2024; Kirk et al., 2024) or using short-answer and multiple-choice QA formats (Etxaniz et al., 2024; Shi et al., 2024). Unfortunately, these settings do not enable us to examine how well LLMs can generate *long-form* text about culturally specific topics.

In this work, we bridge this gap by investigating how well LLMs can answer *culturally specific* questions posed in the language of the target culture, which demand nuanced, long-form responses. We define ***culturally specific*** questions as those more likely to be asked by individuals from cultures closely tied to the language of the question, such as "Kuber iki umwami wa mbere w'uburundi yitwa Ntare?" (Kirundi).[1] We contrast the quality of an LLM's answers to these questions with its answers to ***culturally agnostic*** questions that have consistent answers across many cultures (e.g., "Why is smoking bad for the heart?").

We present **C**ultur**a**l **L**ong-form **M**ultilingual **Q**uestion **A**nswering (CALMQA), the first multilingual long-form QA dataset with a focus on culturally specific questions. Our dataset consists of 2.6K questions (1.5K culturally specific and 1.1K culturally agnostic) spanning 23 high- to low-resource languages[2] and a broad range of topics (Figure 1).[3] CALMQA's culturally specific questions include both naturally-occurring questions collected from community web forums as well as questions in low-resource languages written by hired native speakers. We create the culturally agnostic questions by translating a seed set of 51 English questions into the 22 other languages, following common practice in prior work (Artetxe et al., 2020; Lewis et al., 2020; Longpre et al., 2021; Alonso et al., 2024).

Evaluation of multilingual long-form QA is challenging: lexical metrics used for short-form QA do

[1] English translation: "Why was the first king of Burundi called Ntare (Lion)?"

[2] We define low-resource languages as those not supported by Google Translate or lacking their own Wikipedia resource as of June 2024.

[3] 976 of the 1.5K culturally specific questions also contain reference answers.

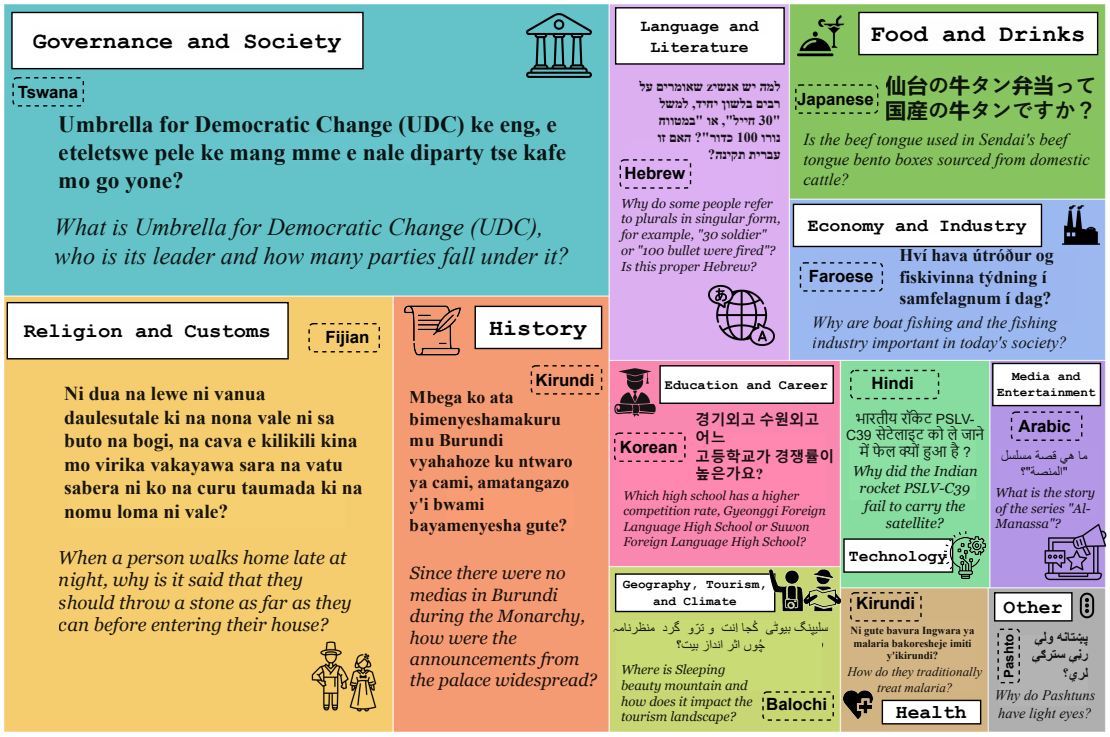

Figure 1: Distribution of topics in CALMQA, with box size indicating the frequency of each topic. Each topic is accompanied by an example and its English translation. Table 8 contains descriptions of the topics, and §B describes our topic classification method.

not correlate with human preferences in long-form QA (Krishna et al., 2021; Xu et al., 2023) or transfer from English to other languages (Kang et al., 2024; Koto et al., 2021; Min et al., 2023; Song et al., 2024). We thus adopt a multi-aspect evaluation protocol including (1) surface-level measures of language identification and repetition; (2) automatic factuality metrics run on translated answers; and (3) human evaluations from native speakers.

Our experiments show that seven popular LLMs, including closed models such as CLAUDE-3-OPUS, GEMINI-1.5-PRO and GPT-4O, suffer from basic surface-level issues, especially on low-resource languages (e.g., none of them reliably generates text in Afar). Additionally, open-weight models such as MIXTRAL-8X22B and LLAMA-3-70B often apologize or generate text in English when prompted with non-English questions. We observe that factuality of LLM-generated culturally specific answers is significantly lower than culturally agnostic answers, underscoring the importance of studying the former category. Factuality drops considerably on low-resource languages, with GPT-4-TURBO and GPT-4O generally performing best on average.

We further conduct a human evaluation on a subset of the data (spanning five languages) for the best-performing models. Native speakers rate and rank answers from different LLMs, and an analysis of their annotations reveals that omissions and factuality issues are strong predictors of answer quality ratings. This human evaluation also supports our automatic factuality evaluation in that culturally agnostic questions are twice as likely to receive higher ratings than culturally specific questions, regardless of the generation model. Taken as a whole, our work establishes a foundation for studying culturally specific long-form question answering in both high and low-resource languages. We will release the CALMQA dataset and code for research purposes.

2 CaLMQA: Cultural Long-form Multilingual Question Answering

Each of the 2.6K examples in CALMQA consists of (1) a question written in one of 23 languages, (2) its English translation, (3) a tag indicating whether the question is culturally specific, and (4) an optional human-written reference answer (for high- and mid-resource languages). This section describes CALMQA's collection process and its statistics.

2.1 Data Collection

We collect our dataset through two processes. For high- and mid-resource languages, we follow prior

work (Fan et al., 2019) and collect questions from community Q&A forums. For low resource languages where such web content is scarce, we hire freelancers to write cultural-specific questions.

Culturally specific questions for high- and mid-resource languages: Many countries have their own community forums where people can exchange information, similar to Quora, Reddit or StackExchange in English. We collect culturally specific questions from these websites via a crowdsourcing process: first, we ask crowdworkers from different countries[4] to provide a link to a community web forum in their language that contains many complex questions that cover a diverse range of topics. Next, we ask workers to provide examples of culturally specific questions and real users' answers from the identified websites along with their English translations. We manually review provided examples and websites, using Google Translate to get English translations of website content. Overall, this process yielded 923 questions across 11 languages with answers at a cost of $1427 USD (Table 3, left).

Culturally specific questions for low-resource languages: Unlike existing LFQA datasets, we also include *low-resource languages* in our study, of which we select twelve to include in CALMQA (Table 3, right). We chose languages with scarce online resources that are not well-studied in prior work, but for which we could also find at least one annotator bilingual in English. We hired 29 native speakers (1-3 annotators per language, depending on their availability) on Upwork,[6] each of whom received guidelines, took a paid ($7 USD) comprehension test, and then wrote culturally specific questions with English translations for $0.65-1.00 USD per question. As having them write answers for all of these languages was prohibitively expensive, we collected answers only for Kirundi along with their English translations ($2 USD per question, $106 USD total). This process yielded a total of 548 questions with English translations at a cost of $833 USD. The entire protocol was reviewed and deemed *exempt* by an Institutional Review Board. Please refer to §A.3 for more details.

[4]We hired workers on the Prolific[5] platform, selecting English-proficient crowdworkers whose native language matches the target language. They selected culturally specific questions with their answers from the website(s) in their native language and were paid $0.65-1.33 USD per question.

[6]https://www.upwork.com/

Collecting culturally agnostic questions in 23 languages: To conduct a controlled experiment contrasting LLM performance on questions that do and do not require cultural knowledge in these languages, we also collect long-form questions that can be answered (1) without culturally specific knowledge and (2) without using technical jargon. They differ from our previously-collected questions in terms of topics, focusing more on the scientific domain. We manually select 51 English questions fulfilling these criteria from the `r/explainlikeimfive` subreddit, following prior work on LFQA (Fan et al., 2019). Then, we use GPT-4-TURBO to translate these questions into all other languages, excluding Balochi, Fijian and Kirundi due to poor translation quality. For these three languages, we pay workers from the low-resource culturally specific collection process to manually write translations. Overall, we obtain 1173 *culturally agnostic* questions (51 questions in English translated further into 22 other languages).

Quality control: We set up a multi-stage quality control process for the collection of culturally specific questions. First, we ran a qualification task for all crowdworkers to ensure they understood our definition of *culturally specific* and *long-form* questions. Each collected question and answer was manually reviewed by at least author to ensure it met our guidelines; workers provided clarifications or replaced unsuitable questions as needed. Please refer to §A.2 for more details about the guidelines and qualification tasks.

2.2 Dataset Analysis

Table 1 and Table 3 summarize the statistics CALMQA's 2.6K culturally specific and culturally agnostic questions. We measure the length of questions with bytes (Clark et al., 2020) as token count is not comparable across languages due to different compression rates (Ahia et al., 2023). High- and mid-resource language questions are generally longer than low-resource language questions, excepting Arabic and Balochi: this phenomenon occurs because high-resource languages generally consist of a question and accompanying context, an artifact of questions from online Q&A forums. By contrast, our low-resource language questions are manually written by workers who are told not to contextualize their questions; see Table 5 for examples. Finally, we categorize CALMQA's culturally specific questions based on their topic by first man-

Language	# Q	Q. Bytes (avg/std)	A. Bytes (avg/std)
English	78	275.7 / 189.0	674.1 / 475.9
Arabic	85	74.3 / 61.3	2105.0 / 2378.6
Chinese	75	193.4 / 329.5	588.8 / 939.7
German	96	304.6 / 227.4	1169.0 / 744.7
Hebrew	96	142.5 / 84.2	2043.6 / 1934.9
Hindi	91	122.4 / 52.8	3618.8 / 1867.1
Hungarian	75	301.1 / 279.8	379.3 / 333.2
Japanese	75	512.0 / 359.3	920.6 / 637.1
Korean	75	126.3 / 138.7	1008.6 / 936.3
Russian	75	310.3 / 438.3	4546.7 / 5067.9
Spanish	102	429.9 / 271.1	852.0 / 817.9
Total	923	253.5 / 279.1	1640.8 / 2291.3

Language	# Q	Q. Bytes (avg/std)	A. Bytes (avg/std)
Afar	25	43.7 / 16.5	N/A
Balochi	65	122.7 / 52.4	N/A
Faroese	30	47.8 / 16.6	N/A
Fijian	75	75.0 / 36.9	N/A
Hiligaynon	65	93.4 / 39.1	N/A
Kirundi	53	64.6 / 21.2	557.2 / 160.9
Papiamento	10	66.8 / 28.5	N/A
Pashto	75	64.8 / 26.9	N/A
Samoan	25	51.2 / 19.3	N/A
Tongan	10	81.2 / 19.2	N/A
Tswana	65	87.2 / 43.4	N/A
Wolof	50	45.3 / 18.9	N/A
Total	548	75.1 / 41.3	557.2 / 160.9

Table 1: Data statistics of high- & mid-resource language (left) and low-resource language (right) culturally specific questions. We report the number of bytes in the UTF-8 encoding. Collecting answers for all low-resource languages would be prohibitively expensive; we collect them for Kirundi only. See Table 3 for culturally agnostic questions.

ually curating a set of categories based on a random selection of 25 questions and then classifying the remainder of the questions using GPT-4-Turbo (see Appendix §B for more). Figure 1 shows a treemap of the question categories and provides an example question from each category.

3 Evaluating LLMs on CALMQA

We generate answers to CALMQA's questions using seven state-of-the-art large language models. Then, we evaluate these answers for both surface-level properties (e.g., repetition) and factuality using automatic metrics. Finally, we perform a fine-grained human evaluation of generated answers in a subset of five languages. Overall, our evaluation shows that all LLMs struggle to answer culturally specific questions compared to culturally agnostic ones, especially in low-resource languages.

Models: We evaluate four closed-source LLMs (Claude-3-Opus, Gemini-1.5-Pro, GPT-4-Turbo, GPT-4o (Anthropic, 2024; Gemini Team, 2024; OpenAI, 2024a,b) and three open-weights LLMs (Aya-13B, Llama-3-70B, Mixtral-8x22B (Üstün et al., 2024; AI@Meta, 2024; Mistral AI, 2024). More model details are in Appendix Table 10.

Answer Generation: Each model is prompted with a question from our dataset in a zero-shot setup without additional instructions. We believe this approach accurately mirrors how humans typically seek information. A model prompted in a given language is expected to return an answer in that same language. We use greedy decoding and limit outputs to 2048 tokens. The total cost of API calls is $530 USD.[7]

3.1 Automatic Evaluation

Since common QA metrics like BLEU (Papineni et al., 2002) or ROUGE (Lin, 2004a) do not correlate well with human judgement for long-form QA (Xu et al., 2023; Krishna et al., 2021), we (1) identify answers with surface level issues (e.g., incorrect language), and (2) extract and verify factual claims from the remaining answers using the VeriScore pipeline of Song et al. (2024).

Identifying surface-level issues: Useful answers must be in the correct language (i.e., the language of the question) and free from word or phrase repetition. We start by detecting answers in the wrong language using a pipeline that combines `polyglot`[8] and `langid`[9] (see §C for details).[10] Then, we identify responses with repetitions by employing `tiktoken`[11] with the `o200_base` encoding and flagging any answer in which a sequence of 20 tokens is repeated four or more times.[12]

[7] We note the total cost of calls for each model as follows: Gemini-1.5-Pro $17 USD, GPT-4o $40 USD, GPT-4-Turbo $80, Llama-3-70B and Mixtral-8x22B $4 USD, and Claude-3-Opus $390 USD.

[8] https://pypi.org/project/polyglot/

[9] https://pypi.org/project/py3langid/

[10] Balochi, Kirundi, Papiamento, and Hiligaynon are excluded due to low language identification accuracy.

[11] https://github.com/openai/tiktoken

[12] Gemini-1.5-Pro often returned an API error for questions in low-resource languages; we mark all of these as in-

Evaluating factuality: After filtering out the surface-level issues, we proceed to evaluate the factuality of the remaining answers, including for each model only languages where at least 50% of the answers remained after filtering.[13] To evaluate factuality of long-form texts, FACTSCORE (Min et al., 2023) verifies automatically extracted claims against retrieved evidence, and recent work expands this to multilingual texts by translating the non-English responses into English (Shafayat et al., 2024). Following this, we translate our answers into English using GPT-4O. Then, we apply the claim extraction and verification pipeline introduced in Song et al. (2024), which improves the robustness of FACTSCORE by focusing exclusively on verifiable, non-trivial claims and using Google Search to obtain evidence.[14][15] Finally, for every valid answer (i.e., answer without surface-level issues), we obtain a list of claims, corresponding search results, and faithfulness labels (supported or unsupported); see Figure 7 for more details.

3.1.1 Results of automatic evaluation

Table 2 reports micro-averaged automatic metrics of each model on culturally agnostic and culturally specific sets, respectively. We summarize the main findings below:

Open-weight models produce many surface-level errors: Overall, the performance of open-weight models is suboptimal, with AYA-13B producing the most repetitions (52.3%), and LLAMA-3-70B outputting answers in an incorrect language – mostly English – in 78.6% of cases. GEMINI-1.5-PRO, on the other hand, returns API errors frequently enough to be unusable for some mid- and low-resource languages.

Surface-level errors are independent of question type: Both culturally specific and agnostic questions suffer relatively equally from surface-level issues, with closed LLMs producing fewer such errors. The language most affected by surface-level issues is Afar, with only 11.8%–25.5% answers being free of surface-level errors. Notably, Afar lacks its own Wikipedia, and ultimately no model is able to effectively generate text in this language. See Appendix Figure 6 for a per-language breakdown of surface-level issues.

Answers to culturally agnostic questions are more factually precise: We conclude by looking at the precision of the extracted factual claims (see Figure 2). Overall, generated answers to culturally agnostic questions tend to be more factually precise (31%–76%) than answers to culturally specific questions (18%–68%).[16] For culturally specific questions, all models perform visibly better for English and Hindi, while questions about Tswana culture yield lowest performance (26% – 30%). Finally, we note that the overall model performance for culturally specific questions is low, with the best model being GPT-4O at 53.2% precision, and the worst being AYA-13B with only 34.0% accuracy (see Table 2).

Limitations of automatic evaluation: Our automatic evaluations focus on surface level issues and factuality, which while important do not allow us to assess the overall quality of the generated answers. We are unable to automatically evaluate aspects such as relevance to the question, content selection, and the presentation of information.

3.2 Human Evaluation

These limitations of automatic evaluation motivates us to conduct a human evaluation with native speakers. While expensive, human evaluation provides a more comprehensive assessment; thus, we perform human evaluation for a subset of five languages: Kirundi, Fijian, Hindi, German, and English.

Evaluation setup: We evaluate answers produced by CLAUDE-3-OPUS, GPT-4-TURBO, and MIXTRAL-8X22B.[17] For each language we sampled 10 culturally specific and 10 culturally agnostic questions.[18] We recruit native speakers of each language via both Prolific and Upwork, all of whom participated in the question collection process as well, paying $7.50 USD per question and an additional $8.00 USD for reviewing the guidelines for a total cost of $720 USD. Annotators are presented with a question, reference answer (if available), and

valid.

[13] We employ this strategy to make sure that a sufficient number of answers remains for evaluation.

[14] We use Google Search via the Serper API at a total cost of $330 USD.

[15] VeriScore's claim extraction and verification models were run on 1xA40 GPU for 24h.

[16] Overall, all models generate a similar number of factual claims on average for both culturally specific and culturally agnostic questions, with the former yielding slightly lower mean claim counts (see Figure 8).

[17] GPT-4O was not available at the time of this evaluation.

[18] For culturally specific questions, annotators selected 10 questions they were confident in answering accurately. For culturally agnostic questions, we supplied annotators with bullet-point answers in English.

MODEL	Surface Level			Factuality		
	Wrong Language (%) ↓	Repetitions (%) ↓	W/o Issues (%) ↑	Factual Precision (%) ↑	# of Langs Eval for Fact. Pr. ↑	Correct/Total Claims ↑
Culturally Agnostic Questions						
CLAUDE-3-OPUS	4.3	0.6	95.0	68.2	18	10,301/15,115
GPT-4-TURBO	3.6	3.0	93.4	73.0	18	16,966/23,228
GPT-4O	2.7	7.5	90.4	74.3	17	14,858/19,998
GEMINI-1.5-PRO	0.3*	0.1	58.2	75.0	10	10,676/14,243
MIXTRAL-8X22B	33.2	10.7	57.2	69.0	10	5,579/8,090
AYA-13B	25.2	67.5	27.1	52.4	2	121/231
LLAMA-3-70B	76.0	0.5	23.5	70.7	5	4,952/7,009
Culturally Specific Questions						
CLAUDE-3-OPUS	2.1	0.1	97.8	46.2	18	8,709/18,839
GPT-4-TURBO	2.8	0.9	96.3	51.7	18	14,371/27,796
GPT-4O	2.2	2.0	95.8	53.2	18	13,372/25,120
GEMINI-1.5-PRO	0.1*	0.1	57.8	51.5	10	9,447/18,340
MIXTRAL-8X22B	19.9	10.1	70.4	49.7	10	5,905/11,883
AYA-13B	14.3	51.8	42.7	34.0	7	336/987
LLAMA-3-70B	73.2	0.7	26.1	51.1	4	3,935/7,701

Table 2: Model performance on culturally agnostic (upper part) and culturally specific (lower part) aggregated across languages. Due to language identification errors, we exclude Balochi, Kirundi, Papiamento, and Hiligaynon from surface-level metrics. *GEMINI-1.5-PRO returned API errors for 44.4% (agnostic) / 27.5% (specific) of answers, which likely obscures surface-level errors that it makes. Factual precision is only computed over languages where at least 50% of model answers lack surface-level issues.

answers generated by the three models in random order. For each candidate answer, they are asked to: (1) identify whether it is in the correct language, (2) mark minor and major errors (e.g., grammar),[19] (3) evaluate factuality, (4) note significant omissions, (5) comment on the overall answer's quality (Figure 3), and (6) rate it on a 5-point scale (excellent, good, average, poor, unusable). Finally, annotators rank the three answers from best to worst and provide a free-form explanation for their ranking. We provide details of the workflow in Figure 12 and §D. The study was reviewed by the Institutional Review Board and received a *non-human subject* determination.

3.2.1 Results of human evaluation

Looking at the overall answer ratings, human annotators prefer GPT-4-TURBO's answers, followed by those of CLAUDE-3-OPUS and then lastly MIXTRAL-8X22B (Figure 4). To confirm, we fit a cumulative link mixed model (`clmm()`) for predicting ratings from models (Appendix Table 13), with annotators nested within language included as a random effect.[20] We find that a MIXTRAL-8X22B answer has an 88% chance of having a lower rating than a CLAUDE-3-OPUS answer ($p<.001$) and a 94% chance of having a lower rating than a GPT-4-TURBO answer ($p<.001$). Also, a CLAUDE-3-OPUS answer has a 30% chance of having a lower rating than a GPT-4-TURBO answer ($p<.001$).

Answer ratings are lower for culturally specific questions: Figure 4 suggests that LLMs generate worse answers for culturally specific questions than for culturally agnostic questions. To check this, we fit a cumulative link mixed model for predicting ratings from question type (Table 15), with annotators nested within language included as a random effect. We see that an answer to a culturally agnostic question has a 67% chance of having a higher rating than an answer to a culturally specific question ($p<.001$). CLAUDE-3-OPUS's performance drop on culturally specific questions is notable: its answer to a culturally specific question has an 80% chance of receiving a lower rating compared to a culturally agnostic question ($p<.001$).

Factuality and omission issues are strong predictors of answer rating: To determine which variables of this experiment (e.g., model, question type, factuality issues, omissions) correlate with answer rating, we fit cumulative link mixed

[19] This step was included to help the annotators visualize any issues with the answer. Hence, we did not require annotators to classify errors beyond a simple minor vs major distinction.

[20] We use `clmm` from the `ordinal` package (Christensen, 2023) because of the ordinal nature of our response variable (ratings) and repeated measures, with annotators rating each model multiple times for different questions.

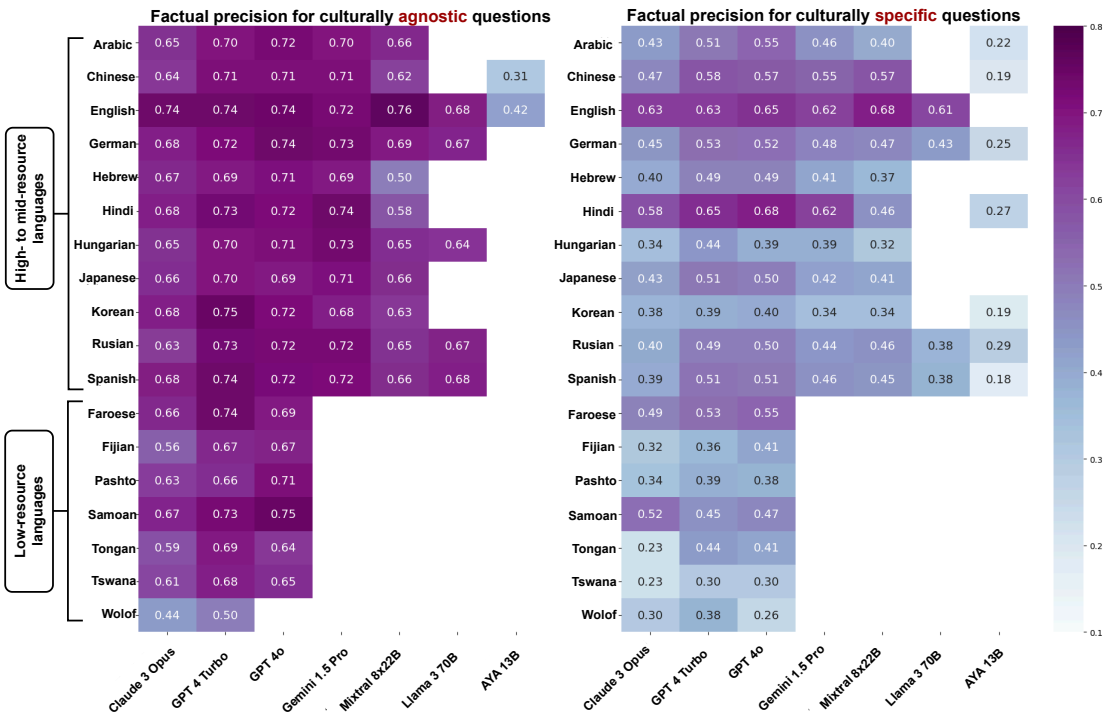

Figure 2: Factual precision for answers without surface-level issues. The left heatmap shows the results for culturally agnostic questions while the right heatmap shows the results for culturally specific questions. Factual precision degrades on culturally specific questions, especially for low-resource languages.

models for predicting the rating, with each variable being used as the sole predictor of a separate model. Table 18 shows R^2 values of these models. We observe high marginal R^2 for the factuality issues model (R^2=0.560) and the omissions model (R^2=0.740), indicating that these factors are strong predictors of answer rating.[21]

3.2.2 Analyzing annotator comments

We can obtain finer-grained insights about answer quality by analyzing comments provided by the annotators. For each comment field, we iteratively develop an annotation schema to analyze the submitted comments. We manually annotate the comments using these schemas, and link these results to the corresponding ratings and scores provided by the annotators. Annotation schemas and evaluation results are in Appendix §D.

Factuality errors are more frequent for culturally specific answers: All 12 issues regarding incorrect dates, entity and events (e.g., "It is mentioned that Nifty was launched in 1995 but it was actually launched in 1996.") occur in culturally specific answers, likely due to a greater prevalence of dates, entities and events in culturally specific questions about topics like HISTORY than culturally agnostic topics like HEALTH AND WELLNESS.

GPT-4-TURBO answers rank first due to content. We analyze the reasons mentioned for ranking each model's answers as best. Having good CONTENT (e.g. due to being complete; see Table 23 for a description) is a reason for GPT-4-TURBO being chosen as best answer 51% of the time (e.g. "*Answer 1 (GPT-4-TURBO) is the perfect answer and and explains all the points needed to understand how to play the game 'Teen Patti'.*"). In the culturally agnostic setting, where CLAUDE-3-OPUS and GPT-4-TURBO perform comparably, more GPT-4-TURBO wins (48%) are attributed to CONTENT than CLAUDE-3-OPUS wins (32%). The full result can be found in §D (Table 19).

4 Related Work

Cultural knowledge in LLMs: Liu et al. (2024), a recent survey of cultural NLP, taxonomizes cultural NLP works into 10 categories including values (Naous et al., 2024b; Jin et al., 2023; Tao et al., 2024), norms and morals (Shi et al., 2024; Rao et al., 2024), and knowledge. Cultural knowledge has been explored through the creation of

[21] In case of mixed effects models, *marginal* R^2 refers to the proportion of variance explained by the fixed effects (predictors) alone.

	Question	Generated Answer	Model	Comment
Kirundi	**Kubera iki cazira ko umwana arya mugahinga ku mwaka mushasha?** *Why was it forbidden for kids to eat at their neighbors' house at the new year?*	Mu muco nyarwanda, Kubera iki cazira ko umwana arya mugahinga ku mwaka mushasha bifite impamvu zikurikira:1. Mugahinga ni ibiryo by'abakuru: Mu muco nyarwanda, mugahinga ufatwa nk'ibiryo by'abakuru kubera ko ari ibiryo bikomeye (...)	Claude 3 Opus	The answer is not answering the question. It talks about something unrelated to the question asked. It's really off-topic. It is also written in Kinyarwanda and talks about Rwanda from the very beginning. The first three words: "Mu muco nyarwanda..." are translated as "In Rwandan culture...", which shows that it is purely about Rwanda, not Burundi.
Hindi	**जनरल और ओबीसी में क्या अंतर है?** *What is difference between general and OBC?*	General (G) और Obstetrics and Gynecology (OB/GYN) दोनों अलग-अलग प्रकार के चिकित्सक हैं और उनके बीच कई अंतर हैं।\n\n1. विशेषता: General चिकित्सक सामान्य रोगों और बीमारियों के लिए प्रदर्शित होते हैं, जबकि OB/GYN चिकित्सक स्त्री के योनि संबंधी बीमारियों और गर्भावस्था से संबंधित बीमारियों के लिए प्रदर्शित होते हैं। (...)	Mixtral 8x22B	The entire answer is wrong. The question is to state the difference between General and OBC category of people (based on caste and society status), but the answer is about general and gynecology practitioners (doctors). It is irrelevant to the question. The answer is written partially in Hindi and English. It is not clear. It is fluent. It sounds artificial.
German	**Was ist Überhangmandat? Gegoogelt, aber nichts passendes gefunden. Ist zu schwer erklärt.** What is an overhang seat? Googled it, but couldn't find anything suitable. It's explained too complicatedly.	Ein Überhangmandat ist ein Begriff aus dem Wahlrecht, der vor allem in Verhältniswahlsystemen mit Direktkandidaten eine Rolle spielt, wie zum Beispiel in Deutschland bei Bundestagswahlen. Ich versuche es einfach zu erklären:\n\nStellen wir uns vor, eine Partei tritt in einem Wahlgebiet an, in dem direkt gewählte Abgeordnete (durch Erststimme) und über Listen gewählte Abgeordnete (durch Zweitstimme) in ein Parlament einziehen. (...)	GPT-4 Turbo	Very helpful, has a good mixture of theory and a well explained example. Complete, clear and fluent. Definitely it would go for "human-like".

Figure 3: Examples of comments on LLM-generated answers written by human annotators.

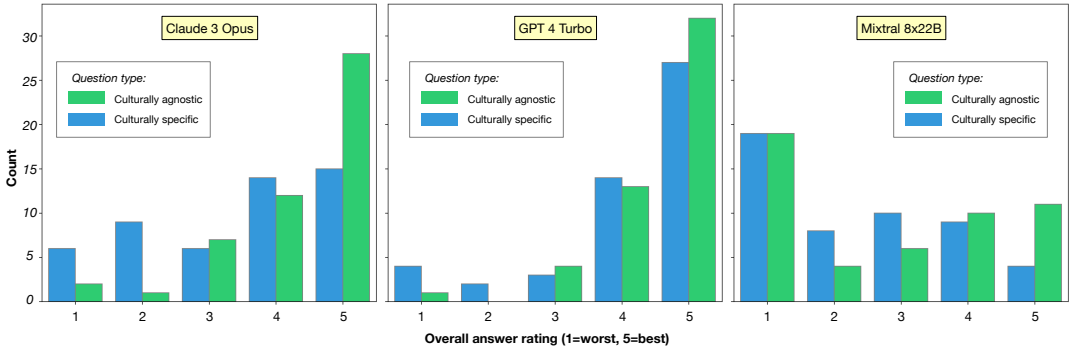

Figure 4: Distribution of human ratings of answers by model and question type. Each model generates 50 answers per question type. Humans give higher ratings for culturally agnostic answers, especially for CLAUDE-3-OPUS.

knowledge bases (Fung et al., 2024; Nguyen et al., 2022) as well as benchmarks and datasets for tasks like probing (Keleg and Magdy, 2023; Yin et al., 2022; Zhou et al., 2024) and short-form QA. Cultural knowledge QA works include non-English information-seeking benchmarks in the style of MMLU (Hendrycks et al., 2020; Kim et al., 2024; Koto et al., 2024a; Li et al., 2023; Koto et al., 2023; Son et al., 2024), common sense datasets (Myung et al., 2024; Wibowo et al., 2023; Koto et al., 2024b) and evaluations (Shen et al., 2024). We are not aware of any existing LFQA datasets of cultural knowledge; surveys of cultural NLP (Liu et al., 2024; Hershcovich et al., 2022) do not mention any such datasets.

Evaluation of Long-Form QA: Evaluating long-form QA (LFQA) remains challenging. Lexical metrics like ROUGE (Lin, 2004b) and BERTScore (Zhang et al., 2019) show poor correlation with human ratings (Krishna et al., 2021; Xu et al., 2023; Cambazoglu et al., 2021; Chen et al., 2019). Later work (Bhat et al., 2023) focuses on English and employs GPT-4 to validate answers and shows better correlation with human judgments. More recent studies on factuality break answers into claims and verify them against knowledge sources (Min et al., 2023; Wei et al., 2024; Song et al., 2024). Shafayat et al. (2024) extend this approach to multilingual answers by translating them into English and validating the obtrained scores. Our pipeline builds on this work by translating responses into English and using VERISCORE (Song et al., 2024) pipeline to validate the claims.

5 Conclusion & Future Work

We introduce CALMQA, the first multilingual long-form QA dataset, which contains 2.6K culturally specific and agnostic questions across 23 high- and low-resource languages. Our evaluation of seven state-of-the-art LLMs reveals that culturally specific questions are more difficult for models than culturally agnostic ones, evidenced by lower factuality and human ratings. Furthermore, we observe critical surface-level issues (wrong language, repetition) in all models, especially for low-resource languages. Our results stress the importance of diversifying pre- and post-training datasets to emphasize cultural knowledge acquisition, which can help improve culturally specific QA. Also, improving cross-lingual transfer to address data scarcity may help for underrepresented languages like Afar.

Limitations

While we strive to cover as many aspects of the cultures represented in CALMQA as possible, we acknowledge that it is not feasible to encompass every cultural nuance. Additionally, for low-resource languages, we employed workers to manually write questions, which impacts scalability. Finally, our culturally agnostic questions are translations from English performed by GPT-4-TURBO, and thus may not match the quality of human translations.

Our automatic evaluation relies on surface-level measures such as language detection and token repetitions. While this approach allows us to determine that current LLMs still struggle with producing outputs in the correct language and without repetitions, it does not assess the relevance or fluency of outputs that lack these surface-level issues. This underscores the need for comprehensive metrics to evaluate overall answer quality in multilingual LFQA, which we leave to future work.

We assess factuality of model generated answers by translating them into English and extracting verifiable claims and validating them against evidence retrieved through web searches. However, this evaluation is influenced by three factors: (1) the quality of translation, (2) the quantity of extracted claims and (3) the availability of relevant online evidence. While we do not observe any evident issues with our pipeline during data inspection, it is possible then these factors influenced the results.

Our human evaluation uses 100 questions across 5 languages to demonstrate that models struggle to generate well-written, factual, and complete answers in non-English languages. Large-scale human evaluation is time-consuming and prohibitively expensive, and finding workers proficient in low-resource languages presented a significant challenge, constraining our evaluation efforts. However, we have shown that we can statistically justify various insights about LLM multilingual capabilities with our scale of data.

Ethical Considerations

The protocols for data collection and human evaluation described in this paper were reviewed and deemed *exempt* by the Institutional Review Board. All annotators provided informed consent for the use and publication of their annotations and collected questions. They were compensated fairly for their work, with their preferred rates respected for both the question collection and evaluation tasks.

References

Orevaoghene Ahia, Sachin Kumar, Hila Gonen, Jungo Kasai, David R. Mortensen, Noah A. Smith, and Yulia Tsvetkov. 2023. Do all languages cost the same? tokenization in the era of commercial language models. *ArXiv*, abs/2305.13707.

AI@Meta. 2024. Llama 3 model card.

Badr AlKhamissi, Muhammad ElNokrashy, Mai Alkhamissi, and Mona Diab. 2024. Investigating cultural alignment of large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12404–12422, Bangkok, Thailand. Association for Computational Linguistics.

Iñigo Alonso, Maite Oronoz, and Rodrigo Agerri. 2024. Medexpqa: Multilingual benchmarking of large language models for medical question answering.

Anthropic. 2024. The Claude 3 Model Family: Opus, Sonnet, Haiku. Technical report, Anthropic. Accessed: 2024-05-23.

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*. Association for Computational Linguistics.

Meghana Moorthy Bhat, Rui Meng, Ye Liu, Yingbo Zhou, and Semih Yavuz. 2023. Investigating answerability of llms for long-form question answering. *ArXiv*, abs/2309.08210.

Berkant Barla Cambazoglu, Valeriia Bolotova-Baranova, Falk Scholer, Mark Sanderson, Leila Tavakoli, and W. Bruce Croft. 2021. Quantifying human-perceived answer utility in non-factoid question answering. *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*.

Anthony Chen, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. Evaluating question answering evaluation. In *Conference on Empirical Methods in Natural Language Processing*.

Rune H. B. Christensen. 2023. *ordinal—Regression Models for Ordinal Data*. R package version 2023.12-4.

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.

Esin Durmus, Karina Nguyen, Thomas I. Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, Liane Lovitt, Sam McCandlish, Orowa Sikder, Alex Tamkin, Janel Thamkul,

Jared Kaplan, Jack Clark, and Deep Ganguli. 2024. Towards measuring the representation of subjective global opinions in language models.

Julen Etxaniz, Gorka Azkune, Aitor Soroa, Oier Lopez de Lacalle, and Mikel Artetxe. 2024. Bertaqa: How much do language models know about local culture?

Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. Eli5: Long form question answering.

Yi Fung, Ruining Zhao, Jae Doo, Chenkai Sun, and Heng Ji. 2024. Massively multi-cultural knowledge acquisition & lm benchmarking.

Gemini Team. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Xiaodong Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *ArXiv*, abs/2009.03300.

Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. 2022. Challenges and strategies in cross-cultural nlp. *ArXiv*, abs/2203.10020.

Jiho Jin, Jiseon Kim, Nayeon Lee, Haneul Yoo, Alice Oh, and Hwaran Lee. 2023. Kobbq: Korean bias benchmark for question answering. *Transactions of the Association for Computational Linguistics*, 12:507–524.

Haoqiang Kang, Terra Blevins, and Luke Zettlemoyer. 2024. Comparing hallucination detection metrics for multilingual generation.

Amr Keleg and Walid Magdy. 2023. Dlama: A framework for curating culturally diverse facts for probing the knowledge of pretrained language models.

Eunsu Kim, Juyoung Suk, Philhoon Oh, Haneul Yoo, James Thorne, and Alice Oh. 2024. Click: A benchmark dataset of cultural and linguistic intelligence in korean. *ArXiv*, abs/2403.06412.

Hannah Rose Kirk, Alexander Whitefield, Paul Röttger, Andrew Bean, Katerina Margatina, Juan Ciro, Rafael Mosquera, Max Bartolo, Adina Williams, He He, Bertie Vidgen, and Scott A. Hale. 2024. The prism alignment project: What participatory, representative and individualised human feedback reveals about the subjective and multicultural alignment of large language models.

Fajri Koto, Nurul Aisyah, Haonan Li, and Timothy Baldwin. 2023. Large language models only pass primary school exams in indonesia: A comprehensive test on indommlu. In *Conference on Empirical Methods in Natural Language Processing*.

Fajri Koto, Jey Han Lau, and Timothy Baldwin. 2021. Evaluating the efficacy of summarization evaluation across languages. *ArXiv*, abs/2106.01478.

Fajri Koto, Haonan Li, Sara Shatnawi, Jad Doughman, Abdelrahman Boda Sadallah, Aisha Alraeesi, Khalid Almubarak, Zaid Alyafeai, Neha Sengupta, Shady Shehata, Nizar Habash, Preslav Nakov, and Timothy Baldwin. 2024a. Arabicmmlu: Assessing massive multitask language understanding in arabic. *ArXiv*, abs/2402.12840.

Fajri Koto, Rahmad Mahendra, Nurul Aisyah, and Timothy Baldwin. 2024b. Indoculture: Exploring geographically-influenced cultural commonsense reasoning across eleven indonesian provinces. *ArXiv*, abs/2404.01854.

Kalpesh Krishna, Aurko Roy, and Mohit Iyyer. 2021. Hurdles to progress in long-form question answering.

Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020. MLQA: Evaluating cross-lingual extractive question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7315–7330, Online. Association for Computational Linguistics.

Haonan Li, Yixuan Zhang, Fajri Koto, Yifei Yang, Hai Zhao, Yeyun Gong, Nan Duan, and Tim Baldwin. 2023. Cmmlu: Measuring massive multitask language understanding in chinese. *ArXiv*, abs/2306.09212.

Huihan Li, Liwei Jiang, Nouha Dziri, Xiang Ren, and Yejin Choi. 2024. Culture-gen: Revealing global cultural perception in language models through natural language prompting.

Chin-Yew Lin. 2004a. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

Chin-Yew Lin. 2004b. Rouge: A package for automatic evaluation of summaries. In *Annual Meeting of the Association for Computational Linguistics*.

Chen Cecilia Liu, Iryna Gurevych, and Anna Korhonen. 2024. Culturally aware and adapted nlp: A taxonomy and a survey of the state of the art.

Shayne Longpre, Yi Lu, and Joachim Daiber. 2021. Mkqa: A linguistically diverse benchmark for multilingual open domain question answering.

Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation.

Mistral AI. 2024. Cheaper, Better, Faster, Stronger. Technical report, Mistral AI. Accessed: 2024-06-19.

Junho Myung, Nayeon Lee, Yi Zhou, Jiho Jin, Rifki Afina Putri, Dimosthenis Antypas, Hsuvas Borkakoty, Eunsu Kim, Carla Perez-Almendros, Abinew Ali Ayele, Víctor Gutiérrez-Basulto, Yazmín Ibáñez-García, Hwaran Lee, Shamsuddeen Hassan Muhammad, Kiwoong Park, Anar Sabuhi Rzayev, Nina White, Seid Muhie Yimam, Mohammad Taher Pilehvar, Nedjma Ousidhoum, Jose Camacho-Collados, and Alice Oh. 2024. Blend: A benchmark for llms on everyday knowledge in diverse cultures and languages.

Tarek Naous, Michael Ryan, Alan Ritter, and Wei Xu. 2024a. Having beer after prayer? measuring cultural bias in large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16366–16393, Bangkok, Thailand. Association for Computational Linguistics.

Tarek Naous, Michael J. Ryan, Alan Ritter, and Wei Xu. 2024b. Having beer after prayer? measuring cultural bias in large language models.

Tuan-Phong Nguyen, Simon Razniewski, Aparna S. Varde, and Gerhard Weikum. 2022. Extracting cultural commonsense knowledge at scale. *Proceedings of the ACM Web Conference 2023*.

OpenAI. 2024a. Gpt-4 technical report.

OpenAI. 2024b. Model release blog: GPT-4o. Technical report, OpenAI. Accessed: 2024-05-23.

Shramay Palta and Rachel Rudinger. 2023. FORK: A bite-sized test set for probing culinary cultural biases in commonsense reasoning models. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9952–9962, Toronto, Canada. Association for Computational Linguistics.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

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

Sheikh Shafayat, Eunsu Kim, Juhyun Oh, and Alice Oh. 2024. Multi-fact: Assessing factuality of multilingual llms using factscore.

Siqi Shen, Lajanugen Logeswaran, Moontae Lee, Honglak Lee, Soujanya Poria, and Rada Mihalcea. 2024. Understanding the capabilities and limitations of large language models for cultural commonsense. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5668–5680, Mexico City, Mexico. Association for Computational Linguistics.

Weiyan Shi, Ryan Li, Yutong Zhang, Caleb Ziems, Chunhua yu, Raya Horesh, Rogério Abreu de Paula, and Diyi Yang. 2024. Culturebank: An online community-driven knowledge base towards culturally aware language technologies.

Guijin Son, Hanwool Lee, Sungdong Kim, Seungone Kim, Niklas Muennighoff, Taekyoon Choi, Cheonbok Park, Kang Min Yoo, and Stella Biderman. 2024. Kmmlu: Measuring massive multitask language understanding in korean. *ArXiv*, abs/2402.11548.

Yixiao Song, Yekyung Kim, and Mohit Iyyer. 2024. Veriscore: Evaluating the factuality of verifiable claims in long-form text generation.

Yan Tao, Olga Viberg, Ryan S. Baker, and Rene F. Kizilcec. 2024. Cultural bias and cultural alignment of large language models.

Maxim Tkachenko, Mikhail Malyuk, Andrey Holmanyuk, and Nikolai Liubimov. 2020-2022. Label Studio: Data labeling software. Open source software available from https://github.com/heartexlabs/label-studio.

Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Jie Huang, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. Long-form factuality in large language models.

Haryo Akbarianto Wibowo, Erland Hilman Fuadi, Made Nindyatama Nityasya, Radityo Eko Prasojo, and Alham Fikri Aji. 2023. Copal-id: Indonesian language reasoning with local culture and nuances. *ArXiv*, abs/2311.01012.

Fangyuan Xu, Yixiao Song, Mohit Iyyer, and Eunsol Choi. 2023. A critical evaluation of evaluations for long-form question answering.

Da Yin, Hritik Bansal, Masoud Monajatipoor, Liunian Harold Li, and Kai-Wei Chang. 2022. GeoMLAMA: Geo-diverse commonsense probing on multilingual pre-trained language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2039–2055, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *ArXiv*, abs/1904.09675.

Li Zhou, Taelin Karidi, Nicolas Garneau, Yong Cao, Wanlong Liu, Wenyu Chen, and Daniel Hershcovich. 2024. Does mapo tofu contain coffee? probing llms for food-related cultural knowledge. *ArXiv*, abs/2404.06833.

Ahmet Üstün, Viraat Aryabumi, Zheng-Xin Yong, Wei-Yin Ko, Daniel D'souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, Freddie Vargus, Phil Blunsom, Shayne Longpre, Niklas Muennighoff, Marzieh Fadaee, Julia Kreutzer, and Sara Hooker. 2024. Aya model: An instruction finetuned open-access multilingual language model.

A Data Collection

This appendix provides extra details about the processes used to collect data for CALMQA. §A.1 describes the identification of websites used for data collection. §A.2 outlines the data collection methods for high- and mid-resource languages, and §A.3 details the data collection process for low-resource languages. Table 5 contains example entries from the dataset. Table 3 and Table 6 provide more details on the number of questions and languages included in the dataset.

A.1 Website Survey

We conducted a survey to find websites with non-English cultural questions. The instructions outlined the survey's goal, defined a good website, and specified what constitutes a culturally specific question. Our criteria for a good website included:

- At least 500 answered "good" questions (as defined below). Websites could contain other questions as we could filter them out.
- Most questions and answers should be in a non-English language.
- Questions should cover a diverse range of topics, not just one or two broad areas (e.g., fashion, technology).
- The website should contain culturally specific questions not found on English websites or in English QA datasets.
- The website should have a large community of contributors with many questions answered.

The survey evolved through an iterative process of piloting and refining based on the results.

Survey participants were English-proficient crowdworkers on the Prolific platform (https://www.prolific.com), whose native language was not English. The survey took about 10 minutes to complete, and we paid $10 for valid responses, totaling $510. We considered a response valid if it showed a good-faith effort, even if the website was of insufficient quality or duplicated in another response. From 51 responses, we obtained 4 websites used for question collection. Some websites were rejected despite having good questions because the proportion of good to bad questions was too low for feasible collection. Remaining websites were identified by the authors. See Table 4 for the full list of websites employed.

A.2 High- and Mid-Resource Culturally Specific Questions

Culturally specific questions in high-resource languages were collected by workers from the websites in Table 4. Each collector was required to read guidelines, pass a guidelines understanding test and complete a test pilot of 5 questions in order to qualify for the main task. This protocol was reviewed by the Institutional Review Board.

Guidelines We provided a guidelines slideshow detailing the rules for selecting questions. The main rules for questions where:

1. The question should require long answer.
2. The question should be culture specific.
3. A native speaker would ask this [question].
4. The question should be objective.
5. Questions should not need pictures/links.

Guidelines Understanding Test Our guidelines understanding test consisted of a form consisting of 11 multiple-selection multiple-choice graded questions. The first question assessed question was "Which of these are listed as **important rules for questions** in the guidelines? (you should select all correct answers)", which required showing understanding of long-form culturally specific information-seeking questions. The remaining 10 questions were curated examples of questions that each may or may not have had issues. Test takers were required to select all the reasons why a question was not suitable according to the guidelines, or select that the question was suitable. We reviewed test results manually, and accordingly chose which workers to pass. We provided passing workers with the test answers, so that they could learn from their mistakes. We paid workers $3.33 USD for completing the test.

LANGUAGE	CULTURALLY SPECIFIC			CULTURALLY AGNOSTIC		
	# Q	Q. BYTES (AVG/STD)	A. BYTES (AVG/STD)	# Q	Q. BYTES (AVG/STD)	A. BYTES (AVG/STD)
HIGH- & MID-RESOURCE LANGUAGES						
English	78	275.7 / 189.0	674.1 / 475.9	51	67.1 / 31.7	N/A
Arabic	85	74.3 / 61.3	2105.0 / 2378.6	51	108.7 / 56.4	N/A
Chinese	75	193.4 / 329.5	588.8 / 939.7	51	68.1 / 31.4	N/A
German	96	304.6 / 227.4	1169.0 / 744.7	51	82.2 / 39.8	N/A
Hebrew	96	142.5 / 84.2	2043.6 / 1934.9	51	93.0 / 42.9	N/A
Hindi	91	122.4 / 52.8	3618.8 / 1867.1	51	184.2 / 90.3	N/A
Hungarian	75	301.1 / 279.8	379.3 / 333.2	51	82.3 / 38.2	N/A
Japanese	75	512.0 / 359.3	920.6 / 637.1	51	104.3 / 50.6	N/A
Korean	75	126.3 / 138.7	1008.6 / 936.3	51	93.0 / 43.3	N/A
Russian	75	310.3 / 438.3	4546.7 / 5067.9	51	134.6 / 70.8	N/A
Spanish	102	429.9 / 271.1	852.0 / 817.9	51	83.6 / 36.1	N/A
LOW-RESOURCE LANGUAGES						
Afar	25	43.7 / 16.5	N/A	51	81.1 / 39.8	N/A
Balochi	65	122.7 / 52.4	N/A	51	96.1 / 48.5	N/A
Faroese	30	47.8 / 16.6	N/A	51	75.1 / 34.5	N/A
Fijian	75	75.0 / 36.9	N/A	51	92.5 / 40.6	N/A
Hiligaynon	65	93.4 / 39.1	N/A	51	83.6 / 39.7	N/A
Kirundi	53	64.6 / 21.2	557.2 / 160.9	51	88.2 / 43.1	N/A
Papiamento	10	66.8 / 28.5	N/A	51	74.1 / 35.3	N/A
Pashto	75	64.8 / 26.9	N/A	51	118.1 / 55.6	N/A
Samoan	25	51.2 / 19.3	N/A	51	80.5 / 37.6	N/A
Tongan	10	81.2 / 19.2	N/A	51	102.4 / 47.9	N/A
Tswana	65	87.2 / 43.4	N/A	51	88.8 / 43.4	N/A
Wolof	50	45.3 / 18.9	N/A	51	78.2 / 44.1	N/A

Table 3: Combined data statistics for culturally specific and culturally agnostic questions. For each language, we report the number of questions (# Q), average and standard deviation of question bytes (Q. Bytes) and answer bytes (A. Bytes) in UTF-8 encoding. Answer bytes for culturally agnostic questions are not available, and are marked as N/A.

A.3 Low-Resource Culturally Specific Questions

Questions for low-resource languages were collected by hiring native speakers proficient in English through Upwork. They were paid $0.65 to $1.00 USD per submitted question with its English translation. Annotators were required to read the guidelines and complete a short comprehension task, for which they were paid $7 USD. Additionally, answers to all Kirundi questions were paid $2 USD per answer. This protocol was reviewed by the Institutional Review Board.

Annotators were instructed to write up to 25 questions in their native language along with English translations, ensuring the questions met the following criteria:

- The question requires a long-form answer (at least 3-4 sentences).
- The question is culturally specific, meaning it is more likely to be asked in the region where the language is spoken.
- The question is something a native speaker of the language might ask.
- The question has an objective answer (i.e., not based on opinions).

B Question Categorization

In this section we describe the process of categorizing all questions into predefined set of categories.

Method We selected 25 random culturally specific questions from the dataset. We manually created a list of broad categories with descriptions and examples, and then 2 authors independently applied the categorization on the 25 questions. We reviewed disagreements and accordingly refined the categories. Then we used GPT-4o to categorize using the prompts in Table 7, with temperature set to

WEBSITE	LANGUAGE	ISO	# Q	URL
Ejaba	Arabic	ar	29	https://www.ejaba.com/
Ujeeb	Arabic	ar	56	https://ujeeb.com/
Zhihu	Chinese	zh	75	https://www.zhihu.com/
Reddit ELI5	English	en	78	https://www.reddit.com/r/explainlikeimfive/
Gutefrage	German	de	96	https://www.gutefrage.net/
Quora	Hebrew	he	96	https://he.quora.com
Let's Diskuss	Hindi	hi	91	https://hi.letsdiskuss.com/
Gyakori kérdések	Hungarian	hu	75	https://www.gyakorikerdesek.hu/
Yahoo Japan	Japanese	ja	17	https://chiebukuro.yahoo.co.jp/
OKWave	Japanese	ja	58	https://okwave.jp/
Naver	Korean	ko	75	https://kin.naver.com/qna/
Yandex	Russian	ru	75	https://yandex.ru/q/
Todoexpertos	Spanish	es	102	https://www.todoexpertos.com/

Table 4: Websites from which cultural questions were obtained, with the number of questions retrieved by website. Multiple websites were used for a given language if workers were struggling with a given website.

0.0. After minor clarifications to category descriptions, we found that GPT-4O produced adequate categories for all 25 questions. We consequently used the model to categorize all of CALMQA. Our final categories, with descriptions and examples, can be found in Table 8.

Results Figure 5 shows the number of questions by category and language. We observe that one of RELIGION, BELIEFS, CUSTOMS, AND TRADITIONS, GOVERNANCE AND SOCIETY and HISTORY is the top category of almost every language (the exceptions being English and Korean). Furthermore, RELIGION, BELIEFS, CUSTOMS, AND TRADITIONS is the predominantly the top category for low-resource languages (10/12). This difference is likely due to the question collection process for low-resource languages.

C Automatic Evaluation

In this section of the appendix we present the details of automatic evaluation. All evaluated models are listed in Table 10. Examples of the model tendencies detected by automatic evaluation are in Table 11.

Language accuracy Figure 9 displays the percentage of responses each model generated in the correct language, independent of correctness or fluency of the answer. We used `polyglot` (https://pypi.org/project/polyglot/) and `langid` (https://pypi.org/project/py3langid/) for language identification, choosing them based on their performance for specific languages. This identification was also applied to the questions to estimate its performance across languages. Our pipeline accurately recognized 100% of instances in 14 languages. For other languages, accuracy typically remained above 90%, with Fijian at 98.67%, Russian at 97.33%, Tongan at 96.92%, Samoan at 92.00%, and Wolof at 90% (see Table 9). However, identification accuracy for Kirundi was notably lower at 35.85%, as the libraries frequently misclassified it as the closely related Kinyarwanda. The automatic identification process failed entirely for Balochi, Hiligaynon, and Papiamento, which is reflected in seemingly low performance for these languages across all the models.

Repetitions Figure 10 illustrates the percentage of responses affected by repetitions, analyzed

FIELD	CONTENT
Language	English
Question (Original)	Why is the game development industry laying off so many people? It feels like every day I see a post about some big game development studio laying off a large chunk of their workforce. What are the causes of this?
Question (English)	Why is the game development industry laying off so many people? It feels like every day I see a post about some big game development studio laying off a large chunk of their workforce. What are the causes of this?
Culturally Specific	True
Answer	There's two main reasons that are impacting not just the game dev industry but others as well. The first reason is that there's been a lot of corporate acquisitions lately like Microsoft buying Activision-Blizzard (AB). When that happens you get a lot of redundant positions that the parents company can absorb, which means there's going to be layoffs in the future. As an example, say there's a department at AB in charge of hiring / Talent acquisition. Microsoft already has a department that does that, so most of those people (minus a few token employees who know the AB company culture) will be let go. The second reason is tied around interest rates. That is an ELI5 in itself, but basically for most of the last decade it was super cheap to get money and use that for projects. Now, it's significantly more expensive to get a loan for money, so companies are less willing to spend money and/or need to cut back because they were supported by cheap money. So as an example, say you're a small game dev in 2013. You can probably get a loan from a Venture Capital (VC) fund for almost zero interest, and there's tons of VCs available to offer you money. Because of that, it's easy to get money to fund your game studio. Fast forward to 2023, and now there are fewer VCs and they charge more interest for their loans. This means you either can't fund your studio or you need to make cuts to survive on the money you generate on your own.
Language	Hungarian
Question (Original)	Hogyan lehet kikeverni a Horthy-kori sisakok színét?
Question (English)	How can you mix the color of helmets from the Horthy era?
Culturally Specific	True
Answer	M35 tábori sisak. Neten rákeresve találsz róla képeket. Nem kell megijedni a sok árnyalattól, annak idején sem volt tökéletesen egységes. Ez egy zöldesbarna szín. Talán a RAL6025 áll a legközelebb hozzá. Festékboltban kikeverik géppel. Nem drágább, mint külön megvenni hozzá egy egy dobozzal a festéket és kevergetni. Ecsettel festették az eredetit.
Language	Fijian
Question (Original)	Na cava na vuna era vinakata kina na Nasi ni veiyanuyanu mera sa lesi i Viti Levu?
Question (English)	What is the reason the nurses from the outer islands want to be assigned to Viti Levu?
Culturally Specific	True
Answer	N/A
Language	Kirundi
Question (Original)	Ni kubera iki twama dukeneye gushira ama aprikasiyo ku gihe? Hoba iki iyo tutabikoze?
Question (English)	Why do we need to constantly do software updates? What happens if I don't?
Culturally Specific	False
Answer	N/A

Table 5: Examples of entries in CALMQA. Metadata like source are omitted here.

LANGUAGE	ISO	FAMILY	Branch	MORPHOLOGY	ORDER	SCRIPT	Region	SPEAKERS
				High- & Mid-Resource				
Arabic	ar	Afro-Asiatic	Semitic	fusional	SVO	Arabic alphabet	Arab world	720M
Chinese	zh	Sino-Tibetan	Sinitic	analytic	SVO	Hanzi	Mainland China, Taiwan, Singapore	1.38B
English	en	Indo-European	Germanic	analytic	SVO	Latin	World-wide	1.5B
German	de	Indo-European	Germanic	fusional	SVO	Latin	Germany, Austria, Switzerland, etc.	133M
Hebrew	he	Afro-Asiatic	Semitic	fusional	SVO	Hebrew script	Israel	9.3M
Hindi	hi	Indo-European	Indo-Iranian	fusional	SOV	Devanagari	India	610M
Hungarian	hu	Uralic	Finno-Ugric	agglutinative	SVO	Latin	Hungary	13M
Japanese	ja	Japonic	Japanese	agglutinative	SOV	Kanji, Kana	Japan	123M
Korean	ko	Koreanic	Korean	agglutinative	SOV	Hangul	Korea	82M
Russian	ru	Indo-European	Balto-Slavic	fusional	SVO	Cyrillic	Russia, Russian-speaking world	255M
Spanish	es	Indo-European	Italic	fusional	SVO	Latin	Spain, Central and South Americas, the US	559M
				Low-Resource				
Afar	aa	Afro-Asiatic	Cushitic	agglutinative	SOV	Latin	Ethiopia, Djibouti, Eritrea	2.6M
Balochi	bal	Indo-European	Indo-Iranian	agglutinative	SOV	Balochi Standard Alphabet	Pakistan, Iran, Afghanistan	8.8M
Faroese	fo	Indo-European	Germanic	fusional	SVO	Latin	Faroe Islands, Denmark	69K
Fijian	fj	Austronesian	Malayo-Polynesian	agglutinative	VOS	Latin	Fiji	640K
Hiligaynon	hil	Austronesian	Malayo-Polynesian	analytic	VSO	Latin	Philippines	9.1M
Kirundi	rn	Niger-Kongo	Atlantic–Congo	agglutinative	SVO	Latin	Burundi	12-13M
Papiamento	pap	Portuguese-based creole	Afro-Portuguese	analytic	SVO	Latin	Aruba, Curaçao, Bonaire	300K
Pashto	ps	Indo-European	Indo-Iranian	fusional	SOV	Pashto alphabet	Afghanistan, Pakistan and Iran	58.8M
Samoan	sm	Austronesian	Malayo-Polynesian	analytic	VSO	Latin	Samoa	510K
Tongan	to	Austronesian	Polynesian	agglutinative	VSO	Latin	Tonga	187K
Tswana	tn	Niger-Kongo	Atlantic–Congo	agglutinative	SVO	Latin	Botswana, South Africa, Zimbabwe	13.9M
Wolof	wo	Niger-Kongo	Atlantic–Congo	agglutinative	SVO	Latin primarily	Senegal	12.3M

Table 6: Linguistic and usage information of the languages in the CALMQA dataset

LANGUAGE	PROMPT
English	You are categorizing questions about different cultures into specific categories. Your task is to assign one category to each question. Here are the available categories: education and career – *[Education and Career Description]*. Example: *[Education and Career Example]* *[categories]...* health and wellness – *[Health and Wellness Description]*. Example: *[Health and Wellness Example]* other – *[Other Description]* Here is the question to categorize: <original_question>*[question]*</original_question> Categorize this question into one of the categories. Output your choice in the following format: <category>category name</category> Your choice:
All except English	You are categorizing questions about different cultures into specific categories. Your task is to assign one category to each question. Here are the available categories: education and career – *[Education and Career Description]*. Example: *[Education and Career Example]* *[categories]...* health and wellness – *[Health and Wellness Description]*. Example: *[Health and Wellness Example]* other – *[Other Description]* Here is the question to categorize: <original_question>*[question]*</original_question> <translation>*[translation]*</translation> Categorize this question into one of the categories. Output your choice in the following format: <category>category name</category> Your choice:

Table 7: Prompts used with GPT-4o to categorize questions. Strings in the form *[form]* are placeholders that are replaced at runtime. The categories used are in Table 8.

Category	Description	Example
Education and Career	Questions related to school, education system, jobs and career paths. Includes developing new skills for new jobs.	*Why do young children drop out of school?*
Governance and Society	Questions about laws, governance and policies, as well as politics and social issues.	*What are the reasons why Japan cannot have casinos?*
Geography, Tourism, and Climate	Questions concerning the geography, climatic conditions, environmental factors of a region, tourism and travelling.	*What is the significance of Gorée Island?*
Technology	Questions about the technology, technological advancements, uses of technology and digital innovation.	*Are stores that accept VISA debit cards marked VISA? Or is it a store with a VISA PLUS mark?*
Economy and Industry	Questions regarding modern-day economic practices, key industries, trade, and economic development.	*Why is our country not developing like others?*
Media and Entertainment	Questions about the media and entertainment specific to the region.	*Why are trademarks obscured on broadcast?*
Food and Drinks	Questions related to culinary traditions, typical foods and beverages, preparation methods, culinary practices, and cultural significance of meals.	*Why is bread with sausage called a sandwich in Russia?*
History	Questions about historical events, significant figures, and important periods that have shaped a culture.	*How did the Bujumbura market burn?*
Language, Art and Literature	Questions about the language, dialects, as well as art forms and literary works/traditions.	*Why is the Balochi Language categorized into its three main dialects?*
Religion, Beliefs, Customs, and Traditions	Questions regarding religious practices, beliefs, rituals, customs, traditions, and holiday.	*In the Islamic religious teaching, what's the meaning of the seven tens?*
Health and Wellness	Questions related to traditional and modern health practices, public health issues, and well-being.	*Why methadone? What makes it "better" than other opioids for maintenance therapy or tapering off another drug?*
Other	Questions that do not fit neatly into the above categories.	

Table 8: Categories of questions in CALMQA.

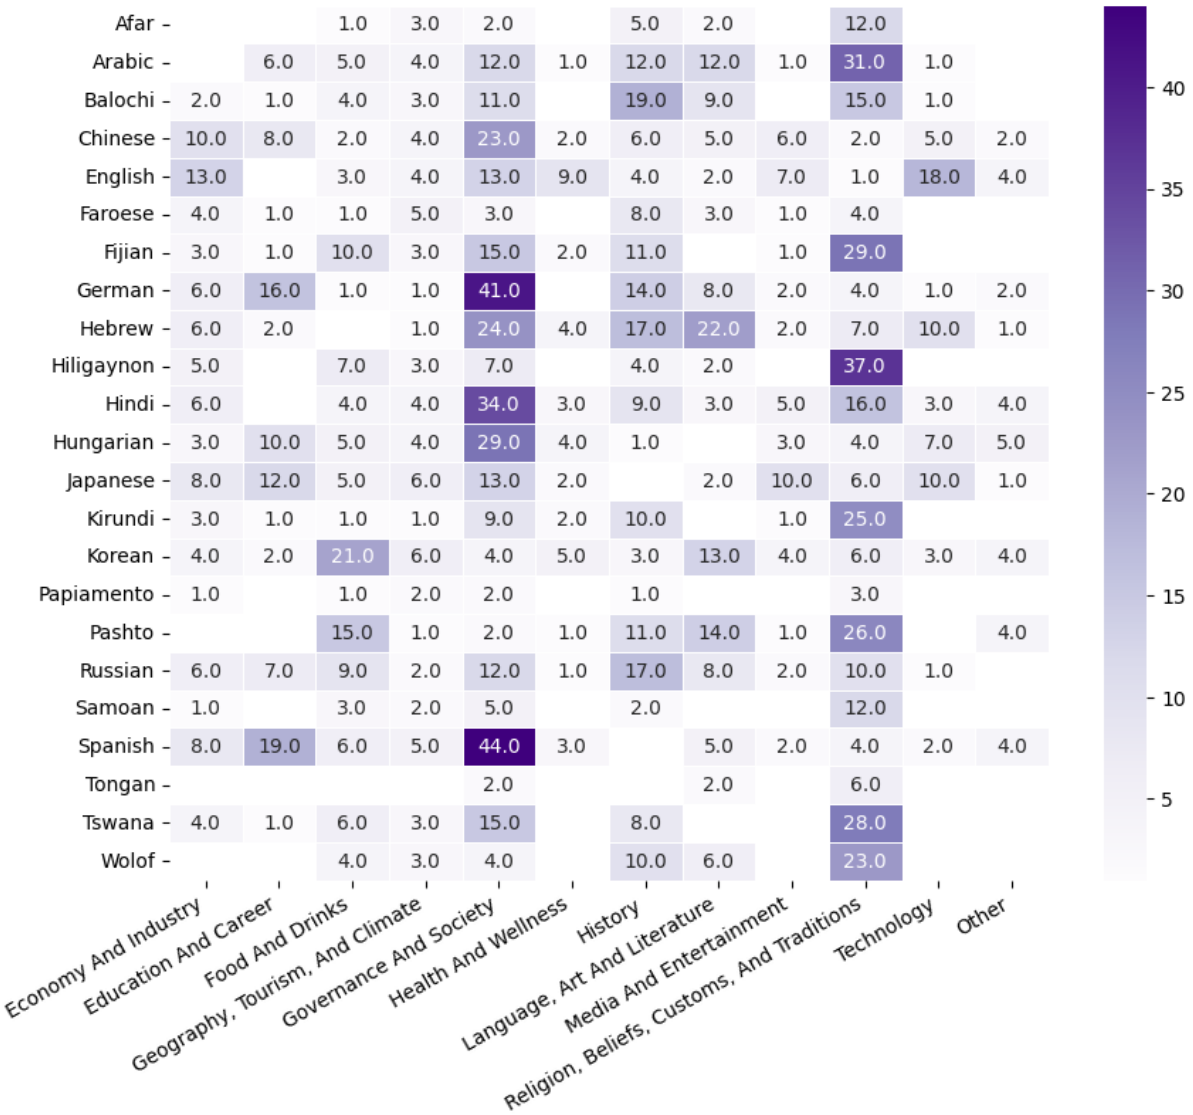

Figure 5: Number of questions by category and language.

LANGUAGE	LANG ID	ACCURACY (%)
Afar	aa	100.00
Arabic	ar	100.00
Balochi	bal	0.00
Chinese	zh	100.00
English	en	100.00
Faroese	fo	100.00
Fijian	fj	98.67
German	de	100.00
Hebrew	he	100.00
Hiligaynon	hil	0.00
Hindi	hi	100.00
Hungarian	hu	100.00
Japanese	ja	100.00
Kirundi	rn	35.85
Korean	ko	100.00
Papiamento	pap	0.00
Pashto	ps	100.00
Russian	ru	97.33
Samoan	sm	92.00
Spanish	es	100.00
Tongan	to	100.00
Tswana	tn	96.92
Wolof	wo	90.00

Table 9: Accuracy of the language detection pipeline on the test set made from questions in the given language. Note that the language detection libraries are often more accurate on longer texts (i.e., texts longer than the length of a single question).

by language across different models. To identify these repetitions, we employed `tiktoken` (https://github.com/openai/tiktoken) with the `o200_base` encoding. We specifically identified instances where at least 20 consecutive tokens were repeated at least four times within an answer.

Further analysis of models' surface-level errors We further analyzed the responses for specific textual indicators. Detected patterns in model responses are presented with examples in Table 11.

Our textual analysis demonstrates issues in MIXTRAL-8X22B responses for low-resource languages. 31.47% of MIXTRAL-8X22B responses to questions in low-resource languages contain phrases like "sorry", "apologize" or "understand" (e.g., *"I'm sorry for any confusion, but it seems you're using a language that I'm not currently able to understand or translate."*). MIXTRAL-8X22B responses to questions in high-resource languages do not contain these apology-related keywords, revealing an inability to answer the question specifically in low-resource languages. The apologetic textual markers were seen in less than 1% of other models' responses except for LLAMA-3-70B's, where they were present in 14.74% of low-resource and 10.48% of high-resource language answers.

Textual indicators also uncover deficiencies in LLAMA-3-70B responses. Notably, 37.87% of responses from LLAMA-3-70B explicitly mention the English name of the language (e.g., *"I see you're speaking in Balochi!"*), indicating that although the system recognizes the language of the question, it nonetheless responds in English. This is in contrast to MIXTRAL-8X22B, which does so in 7.21% of responses, GPT-4-TURBO at 1.84%, and less than 1% for other models. Additionally, approximately 19.71% of LLAMA-3-70B responses include terms like "translate" or "translation" (e.g., *"I apologize, but I'm having trouble understanding your question. Could you please rephrase or translate your question into a language I can understand, such as English?"*), where the system either declines to answer (with or without apology), requests an En-

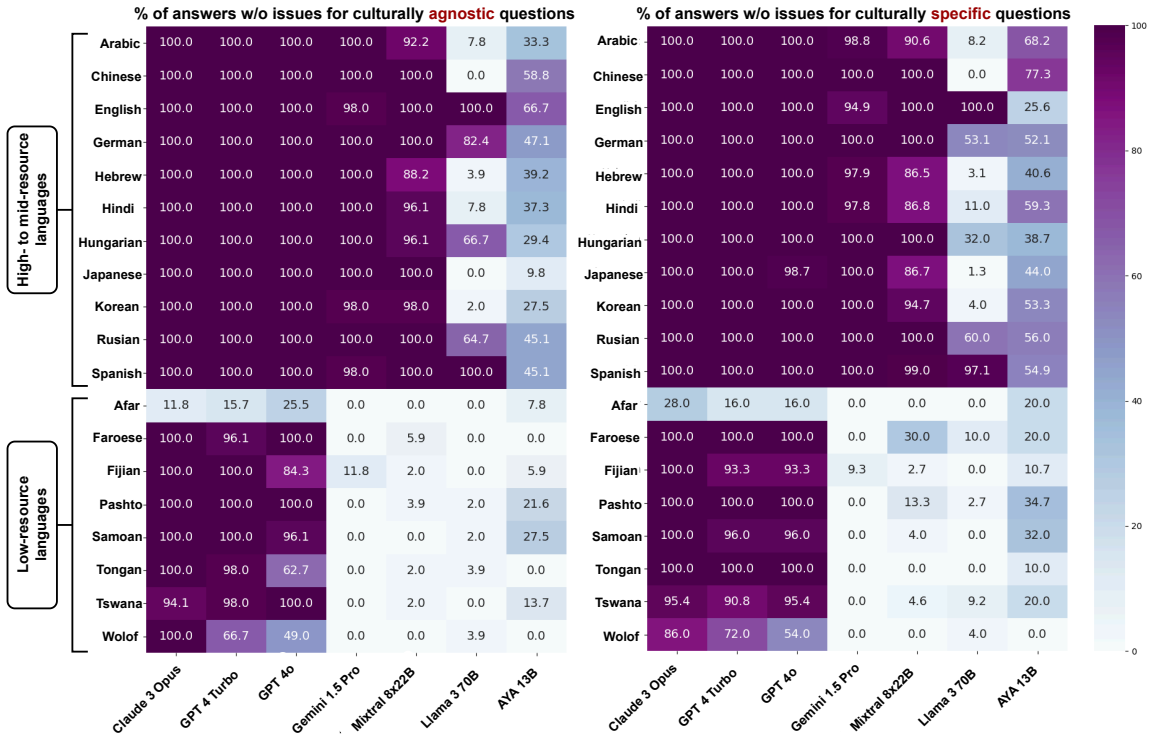

Figure 6: Percentage of model answers without surface issues per language. The left heatmap shows the results for culturally agnostic questions while the right heatmap shows the results for culturally specific questions.

glish translation, or provides a translation itself. In comparison, 8.43% of MIXTRAL-8X22B responses exhibit similar behavior, with less than 1% for other models. Lastly, we observed an unusually high proportion of emojis in responses generated by LLAMA-3-70B, with 17.54% containing at least one emoji.

Claim extraction and verification pipeline We first translated the answers into English with GPT-4o. Then we extract claims using a finetuned Mistral 7B model and use them to query Serper API for evidence. Then we prompt a finetuned Mistral 7B model for verification. Both models were introduced in Song et al. (2024). The pipeline is visualized in Figure 7.

Mean claim count We report the mean claim count by model, language of the question, and question type in Figure 8. We exclude all answers with surface-level issues as well as languages for which the model produced less than 50% of valid answers (i.e., answers without identified surface level issues).

Answer statistics: We compute the lengths of generated answers using `tiktoken` with the `o200k_base` encoding. Table 12 presents statistics for the length of answers generated by each model. To account for variations in token count due to the language of generation and the presence of repetitions, we provide separate statistics for all answers and for those produced in the correct language without repetitions. Finally, we provide the percentage of answers produced in English for a non-English question in Figure 11.

D Human Evaluation

In this section, we present the details of human evaluation.

Evaluation Task The evaluation was conducted using `LabelStudio` (Tkachenko et al., 2020-2022). On the UI, annotators were presented with a question, a gold answer (if applicable), and three competitive answers in random order. The annotation process for each answer involved: (1) marking any mistakes,[22] (2) stating whether the answer is in the correct language, (3) evaluating factual accuracy, (4) noting any content omissions, (5) commenting on the overall quality of each answer, (6) rating each answer on a 5-point scale (excellent, good, average, poor, unusable). Upon completing the ratings, annotators ranked the three answers from best

[22]This step was included to help the annotators visualize any issues with the answer.

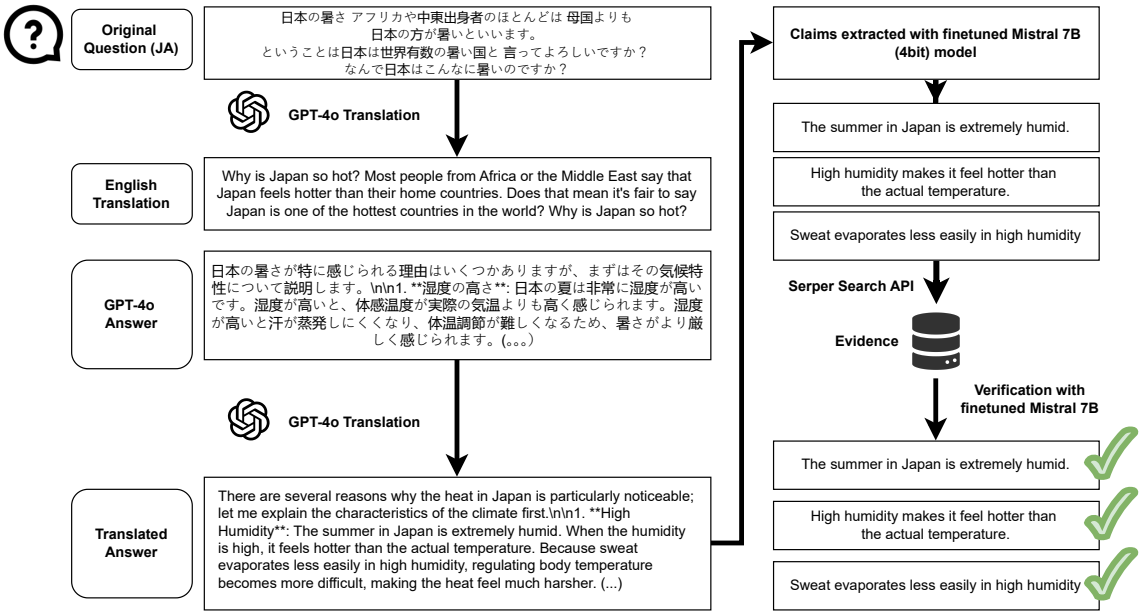

Figure 7: Claim extraction and verification pipeline. Example showing extraction and verification of claims for a question and answer in Japanese. English translations were obtained with GPT-4o. Only part of the answer is provided for readability.

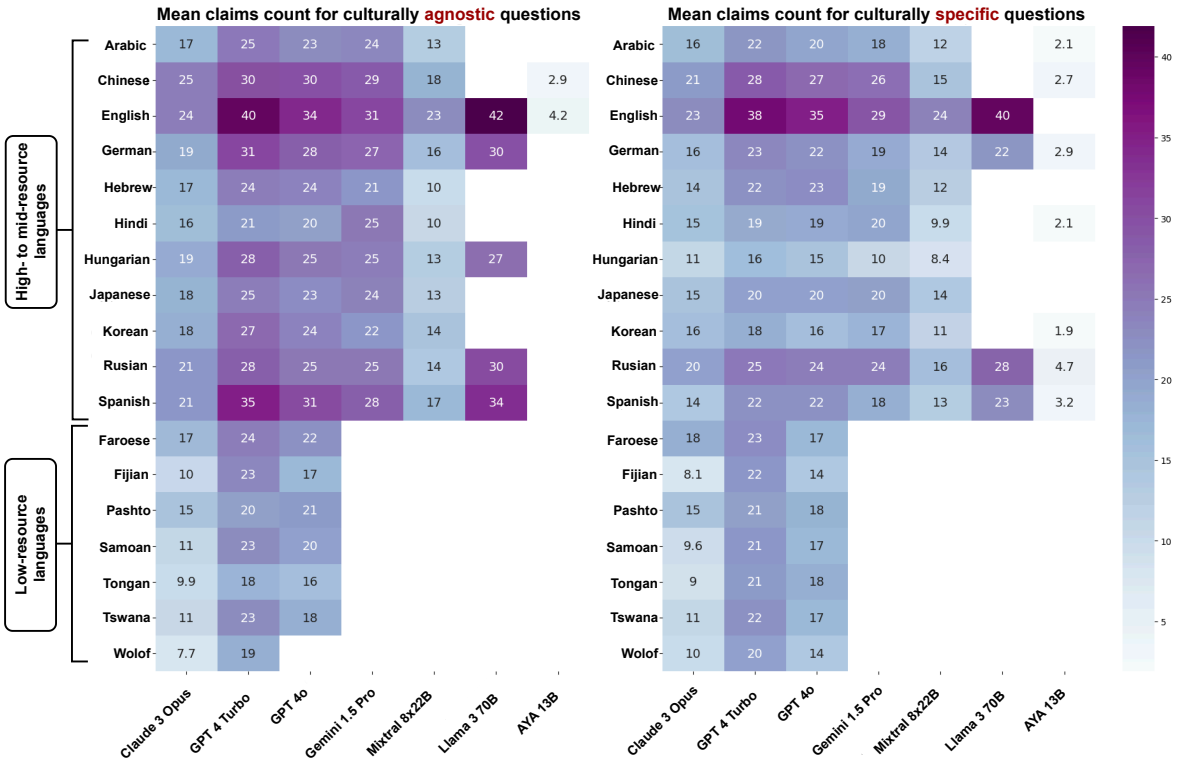

Figure 8: Mean claim count for answers without surface-level issues. The left heatmap shows the results for culturally agnostic questions while the right heatmap shows the results for culturally specific questions. Only languages where at least 50% of answers were free from surface-level issues are included.

Figure 9: Percentage of responses generated in the correct language by model for culture specific questions. Blank cells for GEMINI-1.5-PRO indicate languages where the API returned an error message. It is important to note that language detection libraries performed poorly for Balochi, Hiligaynon, and Papiamento. Additionally, detection accuracy for Kirundi was compromised, with instances of Kirundi being incorrectly identified as Kinyarwanda. Please see Table 9 for details.

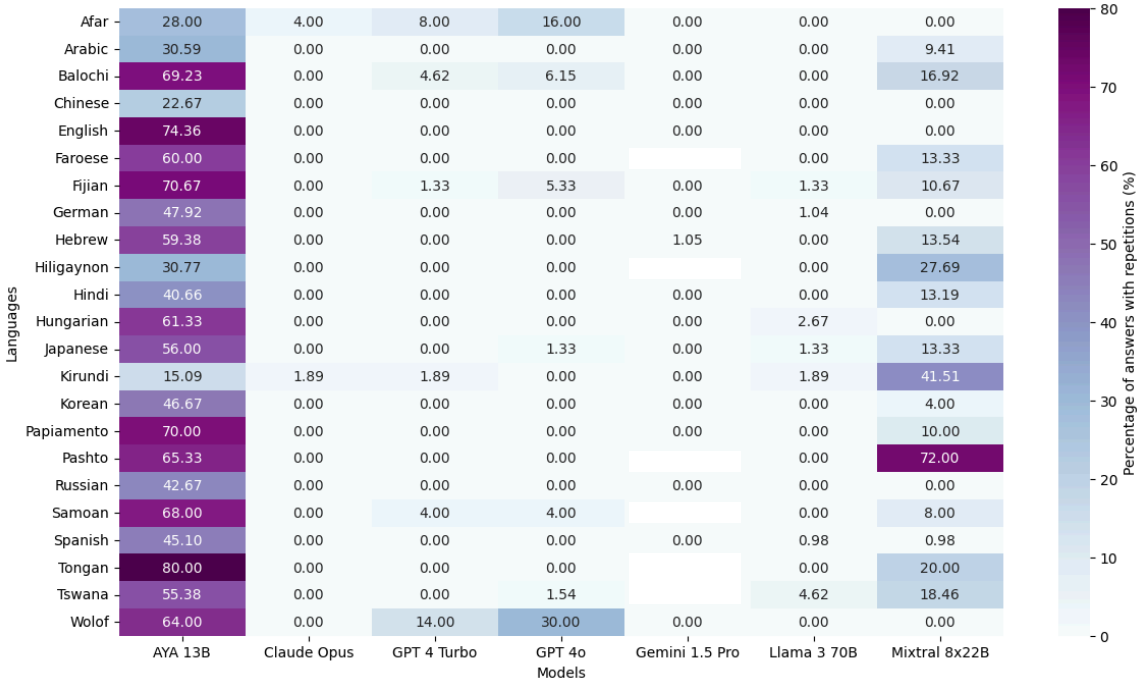

Figure 10: Percentage of answers with identified repetitions for culturally specific questions

MODEL	CHECKPOINT	AVAIL.	# PARAM
GPT-4-TURBO	`gpt-4-0125-preview`	**closed**	?
GPT-4O	`gpt-4o-2024-05-13`	**closed**	?
CLAUDE-3-OPUS	`claude-3-opus-20240229`	**closed**	?
GEMINI-1.5-PRO	`gemini-1.5-pro-preview-0514`	**closed**	?
LLAMA-3-70B	`Meta-Llama-3-70B-Instruct` x	**open-weight**	70B
MIXTRAL-8X22B	`Mixtral-8x22B-Instruct-v0.1` x	**open-weight**	8x22B
AYA-13B	`aya-101` x	**open-weight**	13B

Table 10: General information about models we evaluate using CALMQA.

ISSUE	QUESTION	GENERATION
The model answers a non-English question in English	A BKV járatain egy (külföldi) nyugdíjas miként igazolja, hogy nyugdíjas a sofőrnél? ... *How can a (foreign) senior citizen prove their eligibility for free travel on Budapest's public transport (BKV) to the driver?* ...	I understand your concern! As a foreign pensioner, your mother is entitled to free travel on BKV (Budapest public transportation) services, but she needs to prove her pensioner status to the driver or inspector. According to the BKV website, foreign pensioners can use their national pensioner ID or a document that proves their pensioner status, along with a valid ID (e.g., passport). The document should be in English or have an English translation. Here are some options your mother can consider: ... [Hungarian/LLAMA-3-70B]
A sequence of at least 20 tokens is repeated at least 4 times	Ke eng se se dirang gore boroko jwa go thibelwa go tsoga se dire? *So what's the science behind sleep paralysis?*	Go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore go a kgonega gore ... [Tswana/AYA-13B]
The model apologizes and is unable to answer the question	Duni daddarba leh yan iyyaan duma duni core kee akah yan iyya? *How do they know how hot the earth's core is?*	I'm sorry for any confusion, but it seems like the text you've provided is in a language that I don't currently support. I can understand and generate text in several languages, but I need to know which one you're using to provide an accurate response. If you're able to translate your question into English, I'd be more than happy to help! [Afar/MIXTRAL-8X22B]
The model recognizes the language of the text	Umugirwa wo gufuba umwami wagenda gute mu Burundi? *How was the Kings' funeral done in Burundi?*	A question in Kirundi! Umugirwa wo gufuba umwami wagenda gute mu Burundi? Translation: "What is the significance of the drum in Burundian culture?" In Burundian culture, drums (umugirwa) play a significant role in various aspects of life, including: ... [Kirundi/LLAMA-3-70B]
The model translates the question into English	Mahimo bala kontrolon ukon himuon sang tawo ang tiempo? Kon mahimo, hasta diin kag paano? *Can weather be controlled or man made? If so, to what extent and how?*	What a fascinating question! In English, your question translates to: "Can we control time like we control a person? If so, up to what extent and how?" Unfortunately, the answer is no, we cannot control time like we control a person. Time is a fundamental concept that governs the universe, and it is not a physical entity that can be manipulated or controlled. ... [Hiligaynon/LLAMA-3-70B]

Table 11: Examples of patterns detected in the automatic evaluation of model responses to questions in CALMQA.

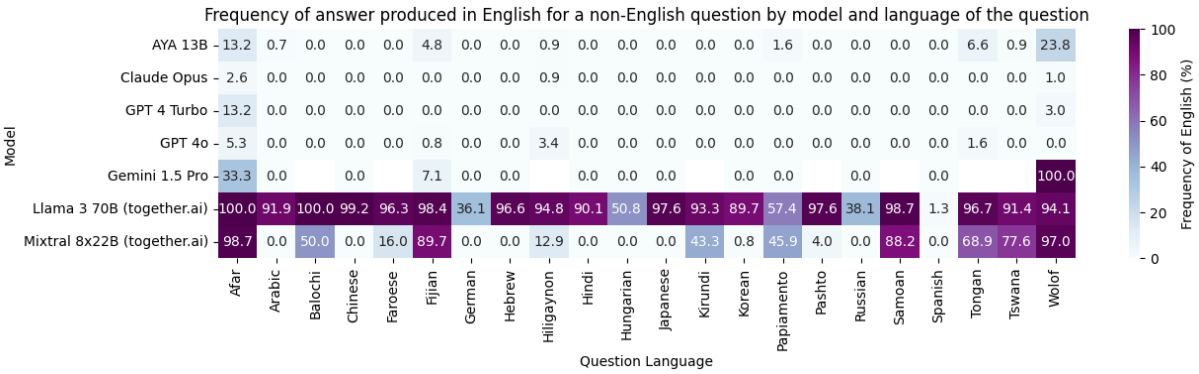

Figure 11: Percentage of answers produced in English for a non-English question by model which produced the answer

	All Data			Correct Lang / No Repetitions		
MODEL	MEAN	MEDIAN	STD	MEAN	MEDIAN	STD
CLAUDE-3-OPUS	296.4	293	88.9	302.2	297	79.2
GPT-4-TURBO	472.6	482	155.2	468.9	477	147.2
GPT-4O	446.6	425	268	434.9	430	184.8
GEMINI-1.5-PRO	265.6	270	247.1	421.6	421	177.7
AYA-13B	867.2	345	894.7	53.7	34	64.4
LLAMA-3-70B	395.9	410	171.4	478.7	484	138.8
MIXTRAL-8X22B	305.3	237	281.9	255.4	252	114

Table 12: Mean, median, and standard deviation of token counts in answers generated by different models. To account for variations in token count due to the language of generation and the presence of repetitions, we provide separate statistics for all answers and for answers produced in the correct language without repetitions. Token counts were computed using `tiktoken` with the `o200k_base` encoding.

to worst and provided a free-form explanation for their ranking. Figure 12 illustrates the overall flow of the evaluation task. The study was submitted for the review to Institutional Review Board and received a *non-human subject* determination.

Guidelines and Consent We provided human evaluation guidelines, describing how to use the interface (including a tutorial video) and explaining each of the steps in the annotation process. The guidelines link to the consent form.

Data Human evaluation was done for answers generated by CLAUDE-3-OPUS, GPT-4-TURBO, and MIXTRAL-8X22B for questions in English, German, Hindi, Fijian and Kirundi. For culturally specific questions, annotators chose 10 questions in their language that they felt confident they knew the answer to. For culturally agnostic questions, we sampled 10 English culturally agnostic questions, and used the original English and the translations into the 4 other languages. We provided annotators with bullet-point answers in English for the culturally agnostic questions.

Workers and Cost German and Hindi annotators were recruited via Prolific, while Fijian and Kirundi annotators were recruited via Upwork. English annotations were performed by one of the authors. All annotators were native speakers of their respective languages and had participated in the question collection. Each question took approximately 20–40 minutes to evaluate, with annotators receiving compensation of $7.50 USD per question and an additional $8.00 USD for reviewing the guidelines, totaling $158 USD per language. The overall cost of the evaluation amounted to approximately $720 USD.[23]

Results Figure 13 and Figure 14 show the results of annotation for whether the answer was generated in the same language as the question (see Table 20 for detailed counts). Figure 15 and Figure 16 display the annotations of the severity of factual issues in each answer (see Table 21 for detailed counts). Figure 17 and Figure 18 present the annotations of the severity of omissions in each answer (see Table 22 for detailed counts). Figure 19 and Figure 20 show the rankings of the models for both culturally specific and culturally agnostic questions. Figure 4 showsh ratings by model by question type. Finally, Figure 21 shows the distributions of scores assigned for each model by the question type and language of generation.

Statistical analysis We conducted a statistical analysis using the `clmm()` function from the `ordinal` package in R. Each model was fitted with the ordinal ratings (1–5) as the response variable and different predictors, allowing for random intercepts for annotators. Table 15 shows the results of a model with question type (either culturally specific or culturally agnostic) as the predictor. Table 16 presents the results of an analysis with model type, question type, and their interaction as predictors, complemented by Table 17, which shows the results of a post-hoc analysis. Finally, Table 18 displays the R^2 values for models with different predictors, namely model type, question type, omission ratings, factuality ratings, and language accuracy

[23] We also covered Upwork charges which the platform impose on the freelancers.

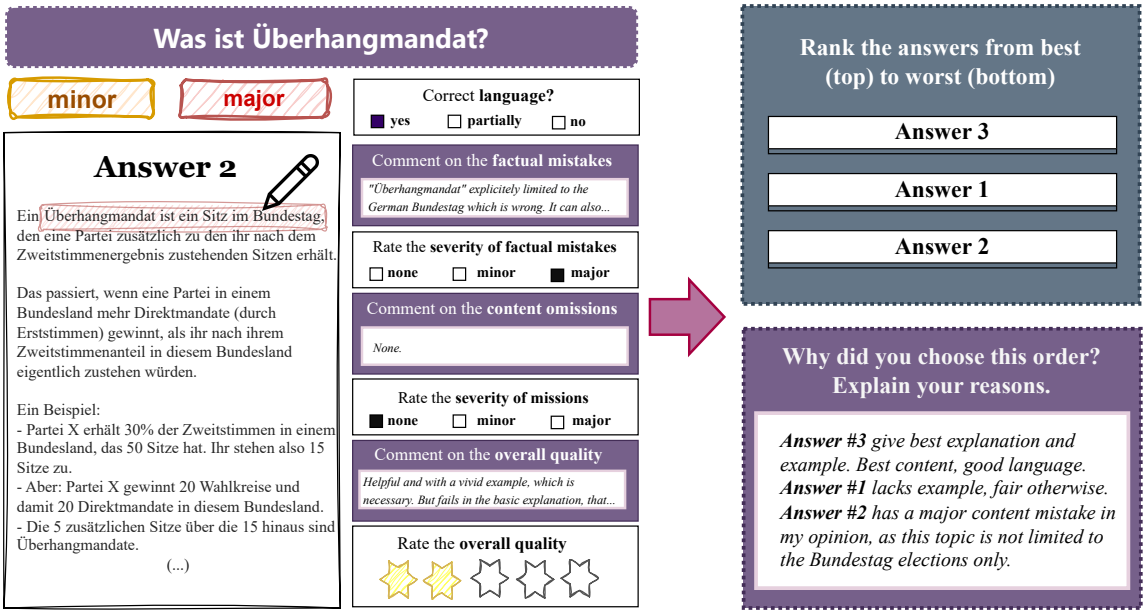

Figure 12: Our human evaluation pipeline. The annotator has to first read the answer, mark and classify all the mistake, and then comment and rate different properties of the answer. Once they have completed evaluating all three answers they are asked to rank them with respect to each other and provide a justification for the ranking. The example shows a culturally specific questions and one answer in German. The answer was produced by CLAUDE-3-OPUS.

ratings.

Analysis of the annotations We conducted manual analyses of the comments provided by the annotators. For each analysis, we iteratively designed an annotation schema to analyze the submitted comments. Table 23 describes the annotation schema used for analyzing the comments on model ranking (i.e., the annotator's reason for ranking a model 1st, 2nd, or 3rd). The results of this analysis are presented in Table 24. Table 25 describes the categories used to analyze the comments on factuality. The results of this analysis are presented in Table 26. Finally, Table 27 describes the categories used to analyze the general comments left by the annotators for each answer. The results of this analysis are presented in Table 28 and Table 29.

D.1 Additional Insights

We capture here insights gained from analyzing human evaluation that we could not fit in the main body of text.

Enumerating facts makes responses seem less human-like. German and Hindi annotators remarked about the presence fact enumerations (often in the form of dot points) for some model answers. For German, the enumeration structure made responses seem artificial (e.g. *'Again very AI made structure. "here are common methods" and a following enumeration plus the asterisk titles...'*). For Hindi, listing facts makes the responses not seem human-like, though not necessarily like an AI either (e.g. *'The answer is just stating points on why is smoking harmful, so it neither sounds human-like nor artificial.'*). More broadly, the fact enumeration structure was described negatively in 5 responses, neutrally in 18 responses and positively in 2.

GPT-4-TURBO made the most grammar/spelling errors. 9 out of 12 spelling and grammar issues were noted for GPT-4-TURBO responses. 8 of these issues occurred for Fijian (e.g. 'There is a minor error, and the system might have spelled "nodra"" incorrectly. However, the language content is relevant so the rating is 4 out of 5, and it sounds like a human.') and the last was in German ('Defninetly helpful, complete and clear. Also fluent. One spelling mistake found: Zusammengefasend is no German word should be "zusammengefasst" or similar. But that could be a human-alike typo.'). This mistakes were present in otherwise mostly positive responses, suggesting that the issues were not due to lack of language understanding. We suspect that this phenomenon may be the result of a tokenizer issue.

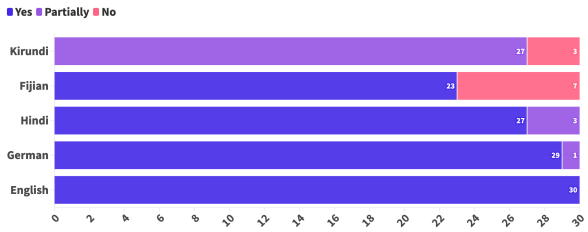

Figure 13: Annotations on Language Correctness for Culturally Specific Questions

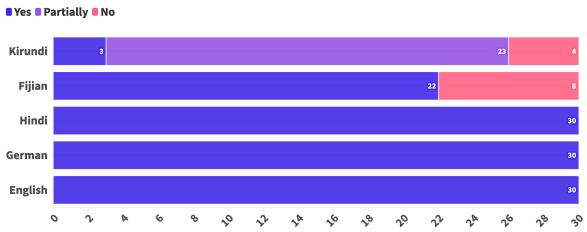

Figure 14: Annotations on Language Correctness for Agnostic-Specific Questions

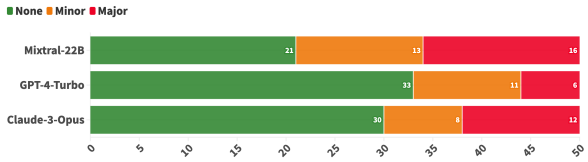

Figure 15: Factuality issues as assessed by the annotators by model for **culturally specific** questions

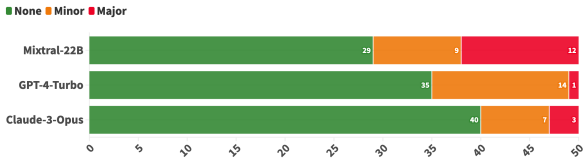

Figure 16: Factuality issues as assessed by the annotators by model for **culturally agnostic** questions

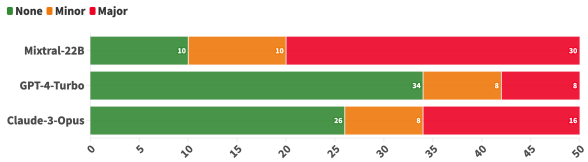

Figure 17: Omissions as assessed by the annotators by model for **culturally specific** questions

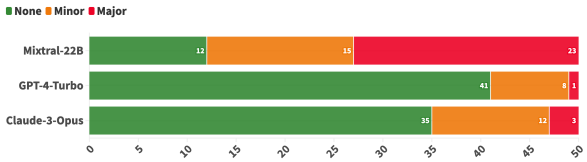

Figure 18: Omissions as assessed by the annotators by model for **culturally agnostic** questions

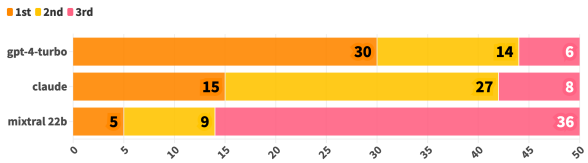

Figure 19: Number of times each model was ranked as *first*, *second*, and *last* for **culturally specific** questions.

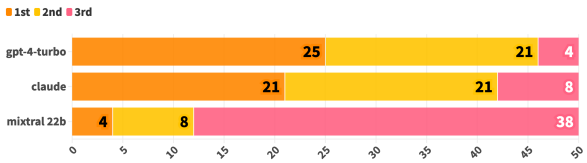

Figure 20: Number of times each model was ranked as *first*, *second*, and *last* for **culturally agnostic** questions.

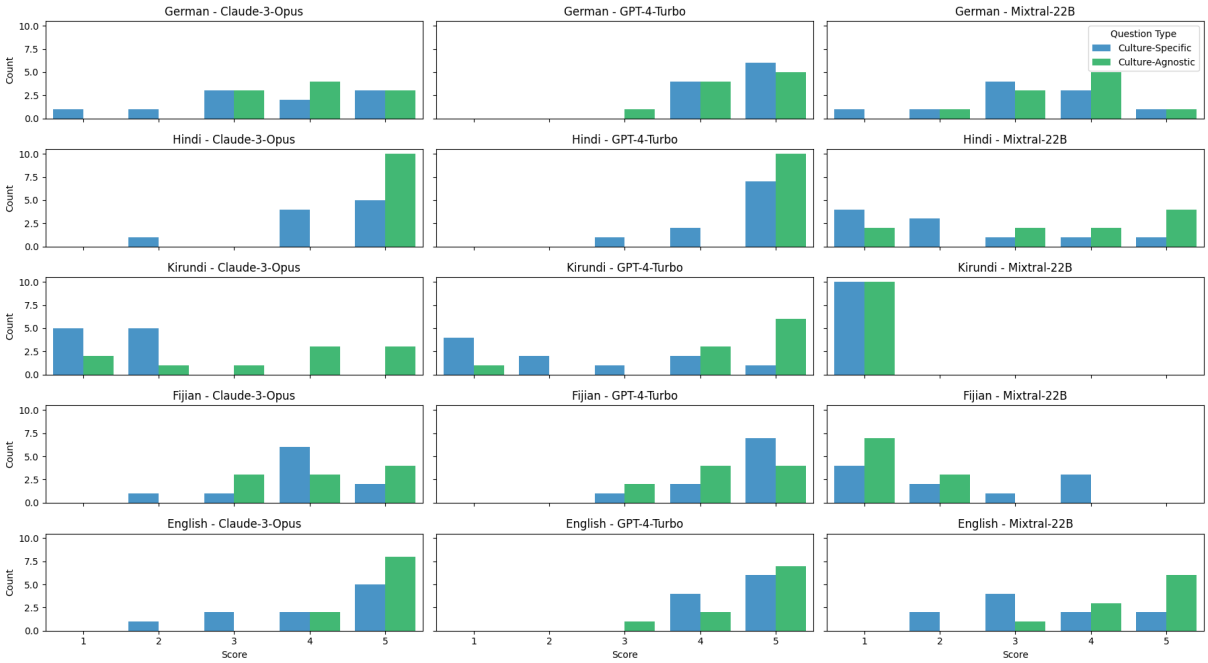

Figure 21: Scores distribution by language and model for Culturally Specific and Culturally Agnostic questions

Formula	`rating` ∼ `model` + (1 | `language/annotator`)			
	Random Effects			
Group	Name	Variance	Std. Dev.	
language	(Intercept)	0.7175	0.847	
Number of groups:	5			
	Fixed Effects			
Coefficient	Estimate	Std. Error	z value	Pr(> |z|)
GPT-4-TURBO	0.8635	0.2885	2.993	0.00276 **
MIXTRAL-8X22B	-1.9493	0.2844	-6.854	7.18e-12 ***
Significance codes:	0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1			

Table 13: Results of cumulative link mixed model with ordinal ratings as the response variable and model as the predictor.

Contrast	**Estimate**	**SE**	**df**	***z*-ratio**	***p*-value**
CLAUDE-3-OPUS – GPT-4-TURBO	-0.863	0.288	Inf	-2.993	**0.0078**
CLAUDE-3-OPUS – MIXTRAL-8X22B	1.949	0.284	Inf	6.854	**<.0001**
GPT-4-TURBO – MIXTRAL-8X22B	2.813	0.315	Inf	8.936	**<.0001**
P value adjustment: Tukey method for comparing a family of 3 estimates					

Table 14: Post-hoc analysis for the model in Table 13. Tests performed using the `emmeans` library in R.

Formula	`rating` $\sim$ `type` + (1 | `language/annotator`)			
	Random Effects			
Group	Name	Variance	Std. Dev.	
annotator	(Intercept)	0.9418	0.9705	
Number of Groups: annotator	5			
	Fixed Effects			
Coefficient	Estimate	Std. Error	z value	Pr($>\|z\|$)
Culturally Agnostic	0.7259	0.2192	3.312	0.000926***
Significance codes:	0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1			

Table 15: Results of cumulative link mixed model with ordinal ratings as the response variable and question type (culturally specific vs culturally agnostic) as the predictor.

Formula	`rating` $\sim$ `model` * `type` + (1 | `language/annotator`)				
Random effects:					
Groups	**Name**	**Variance**	**Std.Dev.**		
Language	(Intercept)	0.7566	0.8698		
Annotator:Language	(Intercept)	0.7566	0.8698		
Fixed effects:					
Coefficient	**Estimate**	**Std. Error**	***z*-value**	**Pr(>|z|)**	
GPT-4-TURBO	1.1679	0.3868	3.020	0.002531	**
MIXTRAL-8X22B	-1.7013	0.3853	-4.415	1.01e-05	***
Culturally Agnostic	1.3561	0.3956	3.428	0.000607	***
GPT-4-TURBO:Culturally Agnostic	-0.6182	0.5849	-1.057	0.290489	
MIXTRAL-8X22B:Culturally Agnostic	-0.7163	0.5467	-1.310	0.190109	

Table 16: Cumulative link mixed model fitted with the Laplace approximation fitted with `clmm()` in R. The response variable is the ratings (an ordinal variable on a 5-point scale), with predictors being model (CLAUDE-3-OPUS, GPT-4-TURBO, or MIXTRAL-8X22B) and question type (culturally specific and culturally agnostic). Annotator nested within language is included as a random effect. The baseline model is CLAUDE-3-OPUS and the baseline question type is culturally specific. Model's conditional R^2 is 0.497 (including random effects) and marginal R^2 is 0.266 (only fixed effects). Please refer to Table 17 for post-hoc analysis.

Contrast	Estimate	SE	df	z-ratio	p-value
Spec. CLAUDE-3-OPUS- Agn. CLAUDE-3-OPUS	-1.356	0.396	Inf	-3.428	**0.0091**
Spec. CLAUDE-3-OPUS- Spec. GPT-4-TURBO	-1.168	0.387	Inf	-3.020	**0.0380**
Spec. CLAUDE-3-OPUS- Agn. GPT-4-TURBO	-1.906	0.424	Inf	-4.492	**0.0001**
Spec. CLAUDE-3-OPUS- Spec. MIXTRAL-8X22B	1.701	0.385	Inf	4.415	**0.0002**
Spec. CLAUDE-3-OPUS- Agn. MIXTRAL-8X22B	1.061	0.374	Inf	2.835	0.0687
Agn. CLAUDE-3-OPUS- Spec. GPT-4-TURBO	0.188	0.412	Inf	0.457	1.0000
Agn. CLAUDE-3-OPUS- Agn. GPT-4-TURBO	-0.550	0.442	Inf	-1.242	1.0000
Agn. CLAUDE-3-OPUS- Spec. MIXTRAL-8X22B	3.057	0.429	Inf	7.123	**<.0001**
Agn. CLAUDE-3-OPUS- Agn. MIXTRAL-8X22B	2.418	0.414	Inf	5.842	**<.0001**
Spec. GPT-4-TURBO- Agn. GPT-4-TURBO	-0.738	0.436	Inf	-1.694	1.0000
Spec. GPT-4-TURBO- Spec. MIXTRAL-8X22B	2.869	0.420	Inf	6.836	**<.0001**
Spec. GPT-4-TURBO- Agn. MIXTRAL-8X22B	2.229	0.404	Inf	5.514	**<.0001**
Agn. GPT-4-TURBO- Spec. MIXTRAL-8X22B	3.607	0.462	Inf	7.800	**<.0001**
Agn. GPT-4-TURBO- Agn. MIXTRAL-8X22B	2.967	0.445	Inf	6.669	**<.0001**
Spec. MIXTRAL-8X22B- Agn. MIXTRAL-8X22B	-0.640	0.383	Inf	-1.669	1.0000

Table 17: Post-hoc analysis for the model in Table 16 with Bonferroni adjustment. Spec. refers to culturally specific questions while Agn. refers to culturally agnostic questions. Tests performed using the `emmeans` library in R.

Predictor	Conditional R^2	Marginal R^2
Model	0.214	0.189
Omission	0.752	0.740
Factuality	0.614	0.560
Language Acc.	0.339	0.327
Q-Type	0.093	0.061
Model * Q-Type	0.497	0.266

Table 18: Conditional and Marginal R^2 values for different predictors. We fit cumulative link mixed models (`clmm()` in R) with *ratings* as the response variable and different predictors. All models included random intercepts for annotators. Omission, Factuality, and Language Accuracy were treated as ordinal variables (no issues > minor issues > major issues), whereas Q-Type and Model are categorical variables with two and three levels respectively. The last model was fitted with the interaction between the Model and the Q-Type. The Conditional R^2 refers to the variance explained by both fixed effects (predictors) and random effects (annotators), while Marginal R^2 refers to the variance explained by fixed effects only.

Model	Win Rate	Reason	# (*spec./agn.*)	Comment
GPT-4-Turbo	55%	Content	27 / 24	*Answer 1 (GPT-4-Turbo) is the perfect answer and and explains all the points needed to understand how to play the game 'Teen Patti'.*
		Language/ Format	11 / 7	*Answer 3 (GPT-4-Turbo) is very well structured and easy to follow. It covers all the information as well.*
		Factuality/ Relevance	10 / 3	*A3 (GPT-4-Turbo) is more factual than A1 (Claude-3-Opus) and A2 (Mixtral-8x22B).*
Claude-3-Opus	36%	Content	6 / 16	*Answer 3 (Claude-3-Opus) covers the topic in its entirety and hence is ranked 1st.*
		Language/ Format	5 / 13	*Answer 2 (Claude-3-Opus) is more readable because the information is listed as points.*
		Factuality/ Relevance	6 / 6	*A3 (Claude-3-Opus) is more detailed and factual than A1 (GPT-4-Turbo) and A2 (Mixtral-8x22B).*
Mixtral-8x22B	9%	Content	2 / 1	*A1 (Mixtral-8x22B) is better explained than A2 (Claude-3-Opus) and A3 (GPT-4-Turbo).*
		Language/ Format	3 / 2	*All answers have equal quality content, so they are distinguished by their structure/verbosity. Answer 3 (Mixtral-8x22B) has a very natural structure. Answer 2 (GPT-4-Turbo) and answer 1 (Claude-3-Opus) have redundancies and answer 1's are slightly worse.*
		Factuality/ Relevance	1 / 0	*Answer 2 (Mixtral-8x22B) was slightly more specific to Western Europe than Answer 1 (GPT-4-Turbo), but both were roughly equal in quality.*

Table 19: Win rates of the three models in human-evaluated 3-way comparisons of answers for 100 questions. Reasons behind the annotators' decisions are provided, with separate reason counts for *culturally specific* and *culturally agnostic* questions. A breakdown of reasons into finer-grained categories is provided in Table 23.

Language	Model	Culturally Specific			Culturally Agnostic		
		Yes	Partially	No	Yes	Partially	No
German	Claude-3-Opus	10	0	0	10	0	0
	GPT-4-Turbo	10	0	0	10	0	0
	Mixtral-8x22B	9	1	0	10	0	0
Hindi	Claude-3-Opus	9	1	0	10	0	0
	GPT-4-Turbo	10	0	0	10	0	0
	Mixtral-8x22B	8	2	0	10	0	0
Kirundi	Claude-3-Opus	0	10	0	1	9	0
	GPT-4-Turbo	0	10	0	1	9	0
	Mixtral-8x22B	0	7	3	1	5	4
Fijian	Claude-3-Opus	10	0	0	10	0	0
	GPT-4-Turbo	10	0	0	10	0	0
	Mixtral-8x22B	3	0	7	2	0	8
English	Claude-3-Opus	10	0	0	10	0	0
	GPT-4-Turbo	10	0	0	10	0	0
	Mixtral-8x22B	10	0	0	10	0	0

Table 20: Count of instances generated in the language of the question by model and question-type, and the language being evaluated

Language	Model	Culturally Specific			Culturally Agnostic		
		None	Minor	Major	None	Minor	Major
German	Claude-3-Opus	8	1	1	10	0	0
	GPT-4-Turbo	8	2	0	10	0	0
	Mixtral-8x22B	8	2	0	9	1	0
Hindi	Claude-3-Opus	7	2	1	10	0	0
	GPT-4-Turbo	7	3	0	10	0	0
	Mixtral-8x22B	1	4	5	6	2	2
Kirundi	Claude-3-Opus	0	1	9	3	4	3
	GPT-4-Turbo	0	4	6	3	6	1
	Mixtral-8x22B	0	0	10	0	0	10
Fijian	Claude-3-Opus	7	3	0	8	2	0
	GPT-4-Turbo	8	2	0	3	7	0
	Mixtral-8x22B	5	5	0	5	5	0
English	Claude-3-Opus	8	1	1	9	1	0
	GPT-4-Turbo	10	0	0	9	1	0
	Mixtral-8x22B	7	2	1	9	1	0

Table 21: Factuality issues in model generation by model, question type and language of the question

LANGUAGE	MODEL	CULTURALLY SPECIFIC			CULTURALLY AGNOSTIC		
		None	Minor	Major	None	Minor	Major
German	CLAUDE-3-OPUS	6	1	3	7	3	0
	GPT-4-TURBO	6	4	0	8	2	0
	MIXTRAL-8X22B	1	7	2	3	6	1
Hindi	CLAUDE-3-OPUS	8	1	1	10	0	0
	GPT-4-TURBO	9	0	1	10	0	0
	MIXTRAL-8X22B	3	1	6	4	4	2
Kirundi	CLAUDE-3-OPUS	0	0	10	4	3	3
	GPT-4-TURBO	2	2	6	6	3	1
	MIXTRAL-8X22B	0	0	10	0	0	10
Fijian	CLAUDE-3-OPUS	6	3	1	6	4	0
	GPT-4-TURBO	8	1	1	7	3	0
	MIXTRAL-8X22B	3	0	7	0	0	10
English	CLAUDE-3-OPUS	6	3	1	8	2	0
	GPT-4-TURBO	9	1	0	10	0	0
	MIXTRAL-8X22B	3	2	5	5	5	0

Table 22: Count of omission issues by severity type, model, and language for culturally specific and culturally agnostic questions

TYPE	DESCRIPTION	EXAMPLE
	Content: Completeness/Explanation	
COMPLETENESS	The answer was perceived as complete.	*Answer 3 (Claude-3-Opus) covers the topic in its entirety and hence is ranked 1st.* [Hindi]
EXPLANATION/EXAMPLES	The answer included useful explanation and/or examples.	*A3 (Gpt-4-Turbo) is better explained than A1 and A2.* [Fijian]
DETAILS/BACKGROUND	The answer included details and/or necessary background.	*Answer 1 (GPT-4-Turbo) and 2 (Claude-3-Opus) are similar but answer 1 has detailed information about the methods to measure body mass compared to 2.* [Hindi]
GENERAL	The answer was general, which was appropriate for the given question.	*Answer 3 (GPT-4-Turbo) is perfect. The writing style of the answer is the best compared to the other answers. For instance, it mentions the timeline in general rather than pointing out exact years of the event (In my opinion, different sources and online transcripts have a little variation in years in terms of history so it is the best to keep it general).* [Hindi]
	Language/Presentation	
LANGUAGE	The answer was fluent/used better language or was less AI-like.	*Both answer 1 (GPT-4-Turbo) and 3 (Claude-3-Opus) are good. Answer 1 sounds more human-like which is why it is ranked 1st.* [Hindi]
STRUCTURE	The structure of the answer was better.	*All answers have equal quality content, so they are distinguished by their structure/verbosity. Answer 3 (Mixtral-22B) has a very natural structure.* [English]
SIMPLE/CLEAR/SPECIFIC	The answer was clear and/or simple, to the point.	*Answer 3 (GPT-4-Turbo) is slightly clearer than answer 1 (Claude-3-Opus).* [English]
SUCCINCT	The answer was succinct.	*All three answers are complete by content in my view (good answer requires more, but that is more than question covers). Answer two (Mixtral-22B) I regard the best, as the density of content in a few lines is awesome - in most of the other questions, "death by long text and details" is valid, here I opt for short and good.* [German]
	Factuality/Correctness	
FACTUAL	The answer is better in terms of factuality.	*A3 (GPT-4-Turbo) is more factual than A1 (Claude-3-Opus) and A2 (Mixtral-22B).* [Fijian]
RELEVANT	The answer is the most relevant to the question. Often mentioned when other answers were irrelevant.	*Answer 3 (GPT-4-Turbo) is placed in the first position because it is relevant though not specific to Burundi. It could be used if it were specific. It has some important information (...).* [Kirundi]
NO ISSUES	There were no apparent issues in the answer.	*Answer 3 (Claude-3-Opus) had no notable issues.* [English]
PARTIAL ANSWER	The answer at least partially addresses the question (while other answers may be refusals, repetitions, or simply irrelevant/wrong).	*Answer 3 (GPT-4-Turbo) is the only one that tries to answer the question. The other 2 just point out differences between the medications.* [English]

Table 23: Categories used for analysis of reasons for specific ranking of the answers

		Completeness - Explanation				Language - Presentation				Factuality - Correctness			
Model	Type	Complete	Explanation	Details	General	Language	Structure	Simple	Succinct	Factual	Relevant	No issues	Partial ans
CLAUDE-3-OPUS	*Spec.*	4	0	2	0	1	2	1	1	1	1	2	2
GPT-4-TURBO	*Spec.*	11	5	9	2	3	2	4	2	4	1	2	3
MIXTRAL-8X22B	*Spec.*	1	1	0	0	1	0	1	1	1	0	0	0
CLAUDE-3-OPUS	*Agn.*	11	0	4	1	2	1	7	3	1	1	4	0
GPT-4-TURBO	*Agn.*	14	4	6	0	1	4	1	1	0	0	2	1
MIXTRAL-8X22B	*Agn.*	0	0	1	0	0	1	0	1	0	0	0	0

Table 24: Count of different reasons mentioned by the annotator for ranking each model's answer as the best out of three. Note that in some cases more than one reason might have been give by the annotator. *Spec.* refers to Culturally Specific questions, while *Agn.* refers to Culturally Agnostic questions.

TYPE	DESCRIPTION	COMMENT EXAMPLE (LANGUAGE/MODEL)
	Direct Factual Errors	
DATE	Issues involving incorrect temporal references.	*It is mentioned that Nifty was launched in 1995 but it was actually launched in 1996.* [Hindi/GPT-4-TURBO]
ENTITY	Incorrect entity such as a person, place, or organization.	*Almost everything is incorrect because the answer states that Ntare Rugamba is the person who accepted to die in the place of the king, while Ntare Rugamba is the king who ruled before the King Mwezi Gisabo.* [Kirundi/GPT-4-TURBO]
EVENT	Errors in the details or occurrence of events.	*It says that Aurangzeb got the mosque built at the place, however this claim is very strong. He got the temple destroyed but it is not sure if he got the mosque built, as it was a decade after the demolition of the temple.* [Hindi/MIXTRAL-8X22B]
REASON	Incorrect reasons or causative explanations for events or situations.	*"Ni o lobika na ligamu, o sa vakalevutaka na kena yawa mai yalomu" means folding your elbow increase the distance from your spirit"* [Fijian/GPT-4-TURBO]
	Contextual and Logical Errors	
SCOPE	Errors involving the incorrect extent or range of a fact.	*"Überhangmandat" explicitely limited to the German Bundestag which is wrong. It can also apply for regional votes for a single state for instance.* [German/CLAUDE-3-OPUS]
ILLOGICAL	Statements that are logically inconsistent or defy common sense.	*Yes, there is an error in the phrases "e rawa ni vakavuna na turu ni wai ni mata ena vuku ni kena vakacacani na mata ni mata" means the cutting of onion can damage the eye of the eye sounds grammatically incorrect.* [Fijian/GPT-4-TURBO]
NEGATION	Incorrect use of negation or statements that present the opposite of what is true.	*Answer implies that internists are not PCPs* [English/CLAUDE-3-OPUS]
PLACEMENT	Inaccuracies resulting from the improper contextual placement of information.	*Putting "7. secondhand smoke" as a dot point incorrectly suggests that it is a reason why smoking is bad for the heart* [English/CLAUDE-3-OPUS]
	Interpretative and Relevance Errors	
IRRELEVANT OR HALLUCINATED	All information in the answer is irrelevant or is fabricated.	*The entire answer is wrong. The question is to state the difference between General and OBC category of people (based on caste and society status), but the answer is about general and gynecology practitioners (doctors). It is irrelevant to the question.* [Hindi/MIXTRAL-8X22B]
DIFFERENT CULTURE	The answer includes references to the wrong culture.	*The answer given is incorrect because it is not specific to Burundi. It is in Kinyarwanda and talks about Rwanda.* [Kirundi/CLAUDE-3-OPUS]
UNCLEAR OR UNDOCUMENTED	Facts that are not clearly expressed or lack sufficient documentation or sourcing.	*The phrase "leweka" is not clear in the content. This issue might have a typo error.* [Fijian/CLAUDE-3-OPUS]
NOT AN ANSWER	The model refuses to answer usually producing the response in English sometimes accompanied with an attempt to translate the question.	*Incorrect language. (refusal to answer produced in English)* [Fijian/MIXTRAL-8X22B]

Table 25: Categories used for analysis of comments on the factuality of the answers

Model	Type	Unclear/ Undoc.	Date	Entity	Event	Reason	Scope
Claude-3-Opus	*Spec.*	1	1	2	1	0	2
GPT-4-Turbo	*Spec.*	0	1	1	0	0	1
Mixtral-8x22B	*Spec.*	1	1	2	2	1	1
Claude-3-Opus	*Agn.*	2	0	0	0	0	0
GPT-4-Turbo	*Agn.*	5	0	0	0	1	0
Mixtral-8x22B	*Agn.*	2	0	0	0	1	0
Model	**Type**	**Illogical**	**Placement**	**Negation**	**Irrelevant or Hallucinated**	**Different Culture**	**Not an Answer**
Claude-3-Opus	*Spec.*	4	0	1	6	4	0
GPT-4-Turbo	*Spec.*	2	1	0	0	7	0
Mixtral-8x22B	*Spec.*	3	0	2	7	0	8
Claude-3-Opus	*Agn.*	1	1	0	6	0	0
GPT-4-Turbo	*Agn.*	1	1	1	4	0	0
Mixtral-8x22B	*Agn.*	0	0	1	5	0	10

Table 26: Count of different types of factuality issues mentioned by annotators in their comments. The issues are presented by question type (*culturally specific* or *culturally agnostic*) and by model which generated the answer. The taxonomy used for this annotation can be found in Table 25.

CATEGORY	DESCRIPTION	COMMENT EXAMPLE (LANGUAGE/MODEL)
	Content Issues	
UNHELPFULNESS	Annotator cannot discern the question's answer from the provided answer text	*This answer is not very helpful because it not specific to Burundi. Of course it contains some relevant information, but it lacks specificity.* [Kirundi/GPT-4-TURBO]
IRRELEVANCE	Contains information not needed to answer the question	*The answer is irrelevant to the question and it isn't helpful. It explains a completely different thing (explains the importance of blood sample rather than explaining why the blood sample is taken from the elbow area).* [Hindi/MIXTRAL-8X22B]
UNDERSPECIFICITY	Ideas are not explained in sufficient detail	*Some instances of generic reasons without much detail about how it applies in Western Europe in particular.* [English/GPT-4-TURBO]
INCOMPLETENESS	Extra information is needed to understand the answer	*The points mentioned in the answer are true but they do not cover the entirety of the topic. ... So it could be said that the answer is subjective rather than complete and clear.* [Hindi/GPT-4-TURBO]
REPETITIVENESS	Some phrase or idea is repeated	*The phrases "Sa vakamuria na noda vakarau ni dautama o ira na iTaukei vulagi....." is repeated in almost every line...* [Fijian/MIXTRAL-8X22B]
GRAMMAR/SPELLING	Grammar or spelling mistakes are present	*There is a minor error, and the system might have spelled "nodra" incorrectly.* [Fijian/CLAUDE-3-OPUS]
	Content Merits	
HELPFULNESS	Information needed to answer the question is present and understandable	*It is helpful because it tells us the reasons for updating apps: Security, Good performance and New features.* [Kirundi/GPT-4-TURBO]
COMPLETENESS	Answer does not need extra information to be understood	*It is complete, even makes an arc from outer core to inner core and a camparison to the sun surface temperature. So it covers more for what has been asked for but I like this additional information.* [German/GPT-4-TURBO]
	Presentation Issues	
ARTIFICIALITY	Perceived by the annotator to be written by AI	*Structure with asterix and enumeration very artificial / non human like.* [German/GPT-4-TURBO]
NON-FLUENCY	Reading the text is awkward/difficult for the annotator	*...Also Argument 5 is "estimation" and this is again mentioned in the summary -> a bit stange and hence not fluent.* [German/CLAUDE-3-OPUS]
INCLARITY	Answer ideas or terms are difficult to understand	*Fluent and complete, but not clear, as complex word i.e. Zytokinen (Cytokines) are not explained.* [German/CLAUDE-3-OPUS]
	Presentation Merits	
NATURALITY	Perceived by the annotator to be human-written/natural	*Complete, but so different to the other answers. most human-like answer as being relatively general to a more or less vague question.* [German/CLAUDE-3-OPUS]
FLUENCY	Written in a smooth manner and easy to read	*There are no inaccuracies in the answer which makes it helpful. It makes use of easier vocabulary which sounds fluent.* [Hindi/CLAUDE-3-OPUS]
CLARITY	Ideas are expressed in an easy-to-understand manner	*The answer is complete and helpful. It is very clear because the information is subdivided into general and OBC sections and it is easy to follow.* [Hindi/GPT-4-TURBO]

Table 27: Categories used for the analysis of annotators' general comments on the quality of answers

MODEL	TYPE	UNHELPFULNESS*	INCOMPLETENESS*	ARTIFICIALITY*	NON-FLUENCY*	INCLARITY*
CLAUDE-3-OPUS	*Spec.*	9	13	4	4	10
GPT-4-TURBO	*Spec.*	6	7	4	2	5
MIXTRAL-8X22B	*Spec.*	20	17	19	10	15
CLAUDE-3-OPUS	*Agn.*	3	5	5	1	1
GPT-4-TURBO	*Agn.*	1	4	11	1	5
MIXTRAL-8X22B	*Agn.*	16	14	17	8	9
MODEL	TYPE	IRRELEVANCE	UNDERSPECIFITY	REPETITIVENESS	GRAMMAR/SPELLING	
CLAUDE-3-OPUS	*Spec.*	1	1	0	1	
GPT-4-TURBO	*Spec.*	4	1	0	2	
MIXTRAL-8X22B	*Spec.*	4	7	4	1	
CLAUDE-3-OPUS	*Agn.*	1	1	2	1	
GPT-4-TURBO	*Agn.*	0	0	0	7	
MIXTRAL-8X22B	*Agn.*	3	1	4	0	

Table 28: Counts of different types of issues noted in annotators' comments about general answer quality. The issues are presented by question type (*culturally specific* or *culturally agnostic*) and by model which generated the answer. The taxonomy used for this annotation can be found in Table 27. Our UI suggested to annotators to make comments (positive or negative) about categories marked with *.

MODEL	TYPE	HELPFULNESS	COMPLETENESS	NATURALITY	FLUENCY	CLARITY
CLAUDE-3-OPUS	*Spec.*	24	19	24	22	22
GPT-4-TURBO	*Spec.*	24	20	27	22	19
MIXTRAL-8X22B	*Spec.*	13	7	18	15	12
CLAUDE-3-OPUS	*Agn.*	26	20	30	22	30
GPT-4-TURBO	*Agn.*	29	24	26	17	25
MIXTRAL-8X22B	*Agn.*	16	12	22	13	16

Table 29: Counts of different types of merits noted in annotators' comments about general answer quality. The merits are presented by question type (*culturally specific* or *culturally agnostic*) and by model which generated the answer. The taxonomy used for this annotation can be found in Table 27. Our UI suggested to annotators to make comments (positive or negative) about all these categories.