

# 000 MMA-ASIA: A MULTILINGUAL AND MULTI- 001 MODAL ALIGNMENT FRAMEWORK FOR CULTURALLY- 002 GROUNDED EVALUATION 003

004 **Anonymous authors**  
005

006 Paper under double-blind review  
007

## 008 ABSTRACT 009

010 Large language models (LLMs) are now used worldwide, yet their multimodal  
011 understanding and reasoning often degrade outside Western, high-resource settings.  
012 A critical, yet underexplored challenge is assessing whether these models  
013 possess genuine “cultural awareness” that is consistent across different modalities  
014 (text, vision, speech) and languages. We propose MMA-ASIA, a comprehen-  
015 sive framework to evaluate LLMs’ cultural awareness with a focus on Asian con-  
016 texts. MMA-ASIA centers on a human-curated, multilingual, and multimodally  
017 aligned multiple-choice benchmark covering 8 Asian countries and 10 languages,  
018 comprising 27,000 questions; over 79% require multi-step reasoning grounded  
019 in cultural context, moving beyond simple memorization. Crucially, this is the  
020 first dataset aligned at the input level across three modalities: text, image (vi-  
021 sual question answering), and speech. This enables direct tests of cross-modal  
022 transfer. Building on this benchmark, we propose a five-dimensional evalua-  
023 tion protocol that measures – (i) cultural-awareness disparities across countries,  
024 (ii) cross-lingual consistency, (iii) cross-modal consistency, (iv) cultural knowl-  
025 edge generalization, and (v) grounding validity. To ensure rigorous assessment,  
026 a Cultural Awareness Grounding Validation Module detects “shortcut learning”  
027 by checking whether the requisite cultural knowledge supports correct answers.  
028 Finally, through comparative model analysis, attention tracing, and an innovative  
029 Vision-ablated Prefix Replay (VPR) method, we probe why models diverge across  
030 languages and modalities, offering actionable insights for building culturally reli-  
031 able multimodal LLMs.  
032

## 033 1 INTRODUCTION

034 Large language and vision–language models are being increasingly deployed across various cul-  
035 tures and languages. Yet, their behavior remains uneven; performance is strongest in high-resource,  
036 Western contexts and degrades in non-Western settings, particularly across Asia (Chiu et al., 2025;  
037 Romero et al., 2024; Vayani et al., 2025; Wang et al., 2024; Myung et al., 2025; Ng et al., 2025).  
038 As multimodal, multilingual models proliferate (Chen et al., 2024a; Bai et al., 2023; Jiang et al.,  
039 2023; OpenAI et al., 2024; Touvron et al., 2023; Romanou et al., 2024), a fundamental question  
040 arises: Does a model possess a coherent underlying understanding of a culture, or does it merely  
041 exhibit fragmented knowledge depending on the input modality and language? While recent works  
042 have probed cultural knowledge in text or images separately, they fail to capture the consistency  
043 of cultural awareness—defined here as the model’s ability to provide stable, reasoned answers to  
044 semantically equivalent inputs regardless of whether the prompt is textual, visual, or spoken.  
045

046 In this paper, we investigate (i) cultural awareness consistency, defined as the extent to which a  
047 model gives stable answers to semantically equivalent inputs when the representation (text, im-  
048 age+question, or spoken question) or the language changes; (ii) cultural awareness grounding, de-  
049 fined as whether correct answers rely on appropriate cultural signals rather than exploitable short-  
050 cuts; and (iii) cultural awareness generalization, defined as whether a model that has access to  
051 the relevant cultural knowledge can perform the required reasoning within those cultural contexts  
052 Balepur et al. (2024); Molfese et al. (2025); Zheng et al. (2023); Kreutzer et al. (2025). Neverthe-  
053 less, conducting such evaluations presents significant challenges. Existing culture-centric datasets  
(e.g., Myung et al. (2025); Wang et al. (2024)) frequently suffer from two key limitations: (i) insuffi-  
cient alignment of instances across modalities, (ii) inadequate representation of low-resource Asian  
languages. Furthermore, evaluation processes are easily hacked through memorization or elimina-

054  
 055 **Table 1: Comparison of existing culture-related benchmark datasets with MMA-Asia.**   and  denote the the text, image and speech modalities, respectively. “MLA”, “MDA” and “RI” denote “multilingual alignment”, “multimodal alignment” and “Response interpretability”, respectively.  
 056  
 057  
 058

Benchmark	MLA	MDA	RI	# of Ctries	# of Langs	Modality	Multi-step reasoning	Question forms	Total samples	Total domains
SeaEval	✗	✗	✗	4	1		✓	Diverse	415	-
CLICk	✗	✗	✗	1	2		-	Diverse	1,995	11
BLEnD	✓	✗	✗	16	13		-	Fixed	52,557	6
Culturalbench	✗	✗	✗	45	1		-	Fixed	1,696	17
CVQA	✓	✗	✗	30	31	 	-	Fixed	10,000	10
CulturalbenchVQA	✗	✗	✗	11	1	 	-	Diverse	2,378	5
ALM-bench	✓	✗	✗	73	100	 	-	Diverse	22,763	19
Md3	✗	✗	✗	3	1		-	Diverse	3,689	2
MULTI-AUDIOJAIL	✓	✗	✗	6	6		-	Diverse	102,720	-
Ours	✓	✓	✓	8	10	  	✓	Diverse	27,000	9

071  
 072 **tion in multiple choice questions (MCQs), which bypass the genuine reasoning capabilities (Wang**  
 073 **et al., 2025; Hartmann et al., 2023; Liu et al., 2025; Yong et al., 2025). As a result, we still lack a**  
 074 **principled way to separate actual cultural competence from artifacts.**

075 **To address this research gap, we introduce MMA-ASIA<sup>1</sup>, an explainable evaluation framework for**  
 076 **Asian cultural knowledge. MMA-ASIA aligns tri-modal items (textual question, image+question,**  
 077 **and Text-to-Speech (TTS)-spoken question) with identical semantics and provides parallel local-**  
 078 **language and English versions authored by native experts across 8 countries and 10 languages to**  
 079 **make a comprehensive evaluation. The framework measures five axes: (1) cultural awareness disparity,**  
 080 **(2) cross-modal consistency, (3) cross-lingual consistency, (4) cultural knowledge generalization**  
 081 **under held-out regimes, and (5) grounding validation via targeted ablations and negative controls.**  
 082 **Using MMA-ASIA, we evaluate 15 multilingual and multimodal LLMs(e.g., GPT-4o, Qwen,**  
 083 **Llama). We find that (i) accuracy drops markedly in low-resource Asian languages compared to En-**  
 084 **glish, (ii) cross-modal consistency lags text-only performance, indicating incomplete transfer from**  
 085 **language to vision and speech, and (iii) grounding controls reduce a non-trivial fraction of apparent**  
 086 **“wins,” revealing shortcut use. We also analyze multi-step, culture-specific reasoning errors and**  
 087 **where visual or linguistic cues fail to connect. We summarize our contributions as follows:**

- 088 **• Aligned tri-modal, multilingual benchmark.** We release **MMA-ASIA** with 27,000 multilingual  
 089 multimodal questions authored by in-country experts across 8 countries and 10 languages.
- 090 **• Five-axis evaluation protocol.** We formalize cultural awareness *consistency* (cross-modal/cross-  
 091 lingual) and *grounding* with negative controls and ablations, plus generalization tests under held-  
 092 out themes/countries. We provide reference implementations and CI-tested evaluation scripts.
- 093 **• Extensive baselines and analyses.** We report zero-shot baselines for 14 model families (mul-  
 094 tilingual LLMs and VLMs), including common-support subsets, and diagnose failure modes by  
 095 modality, language, and reasoning step count.

## 096 2 RELATED WORK

097 **Cultural knowledge in text.** Recent benchmarks assess culture-specific knowledge via MCQs  
 098 (Kim et al., 2024; Wang et al., 2024; Susanto et al., 2025; Myung et al., 2025; Chiu et al., 2025; Cho  
 099 et al., 2025), consistently showing (i) performance gaps favoring English/high-resource settings and  
 100 (ii) sensitivity to formatting. However, most lack aligned multimodal counterparts to test cross-  
 101 modal cultural understanding.

102 **Cultural perception in images (VQA).** Multilingual VQA datasets use community-sourced im-  
 103 ages and questions (Romero et al., 2024; Nayak et al., 2024; Vayani et al., 2025), revealing vision-  
 104 language gaps and language sensitivity. Yet they typically lack text/speech-parallel versions of  
 105 identical items, making it hard to isolate whether failures stem from cultural knowledge, visual  
 106 grounding, or language handling.

107 <sup>1</sup>We will release the data, splits, prompts, decoding settings, and per-item metadata (e.g., knowledge points,  
 108 reasoning tags) for benchmarking, reproducing, and future extensions.

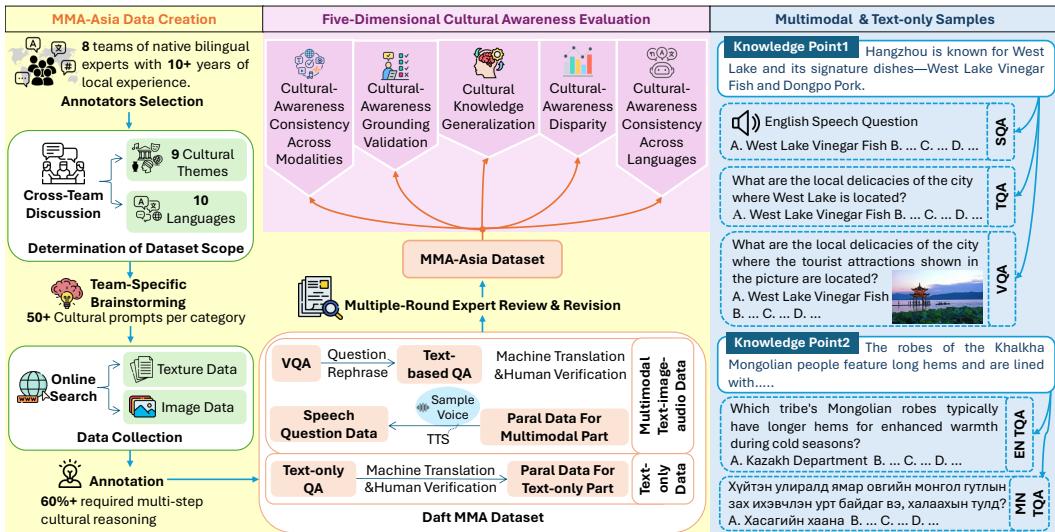


Figure 1: An overview of the MMA-Asia evaluation framework: data creation pipeline (yellow), representative dataset samples (pink), and evaluation dimensions (blue).

**Speech and accent robustness.** Speech datasets reveal substantial accent-related biases (Eisenstein et al., 2023; Roh et al., 2025), and perturbations can drastically change outcomes. However, they rarely evaluate cultural knowledge directly or analyze trimodal alignment and consistency.

**Consistency and grounding.** MCQ performance may reflect shortcuts rather than grounded reasoning (Balepur et al., 2024; Molfese et al., 2025; Zheng et al., 2023; Aakanksha et al., 2024; Kirk et al., 2024). Benchmarks seldom include negative controls or report cross-lingual/cross-modal consistency—essential for distinguishing cultural competence from artifacts.

**How is MMA-ASIA different from others?** Existing datasets evaluate cultural understanding within a single modality or language at a time, without tightly *aligning* instances across modalities and languages, and no built-in *grounding controls*. MMA-ASIA addresses these gaps by: (i) providing semantically aligned tri-modal items (text, image+question, speech) in parallel local-language and English versions; (ii) adopting *cross-modal* and *cross-lingual* consistency as prime metrics; and (iii) integrating targeted ablations and negative controls to test whether answers rely on the intended cultural signal rather than shortcuts. This design enables clearer attribution of failure modes: knowledge vs. language vs. modality, and more reliable measurement of cultural awareness in multimodal, multilingual models. We summarize key differences among representative datasets in Table 1.

### 3 BENCHMARK CONSTRUCTION

MMA-ASIA was collaboratively constructed by research teams from eight countries: China, Singapore, Japan, South Korea, Mongolia, Vietnam, Indonesia, and India. The pipeline comprised five stages: (i) annotator selection, (ii) selection of representative cultural themes and languages, (iii) collection of text and image materials, (iv) question authoring and annotation by country sub-teams, and (v) human review and revision for quality and cultural representativeness. For the definition of cultural themes, we followed the framework proposed by Adilazuarda et al. (2024).

#### 3.1 ANNOTATOR SELECTION

An in-country expert team curated each national subset. All team members were native speakers of the local language and proficient in English. Annotators had lived in the respective country for more than ten years, ensuring deep familiarity with local cultural contexts. Detailed annotator information is in Appendix A.1. Before annotation began, we held project-wide briefings to explain the scope and requirements. We also distributed a detailed English annotation guideline (see Appendix A.3).

#### 3.2 CULTURAL THEMES AND LANGUAGES

Based on established taxonomies from related work Romero et al. (2024); Myung et al. (2025), we constructed an initial set of cultural categories. Through rigorous selection criteria, we retained only highly representative categories, excluding those with insufficient cultural specificity (e.g., "Science and Technology") or negligible cross-national variation (e.g., "Plants and Ani-

162        **mals”).** Through collaborative discussions, we finalized 9 cultural themes, *Daily life habits/Culture*,  
 163        *Food/Cuisine*, *Transportation*, *Buildings*, *History*, *Geographical location and climate*, *Education*,  
 164        *Fashion/Clothing* and *Language/Race*. We balanced the number of questions across themes as much  
 165        as possible. Each national subset includes the country’s official language(s) and English. For India,  
 166        we selected Hindi as the representative language due to its large speaker base among the 22 official  
 167        languages. For Singapore, we included all four official languages: English, Chinese, Malay, and  
 168        Tamil. In total, MMA-ASIA covers ten languages; full details are available in Appendix A.2.

### 169        3.3 TEXT AND IMAGE DATA COLLECTION

170        For both the text-only and multimodal tracks, teams generated at least 56 keywords or short phrases  
 171        per category as *cultural prompts*. If a category could not supply enough prompts, the shortfall was  
 172        filled using prompts from other categories. These prompts were designed to capture both diversity  
 173        and geographic breadth, reducing the risk of homogenizing Asian cultures. Using these prompts,  
 174        team members retrieved relevant texts and images from the web and extracted short passages to  
 175        serve as the basis for question authoring. When source content was ambiguous, we cross-checked  
 176        with multiple references to ensure authenticity. All images were obtained from Creative Commons  
 177        (CC)-licensed resources (details in Appendix A.2).

### 178        3.4 QUESTION CREATION AND ANNOTATION

179        Teams used collected materials to create multi-choice QA data, requiring at least 60% of questions  
 180        to involve multi-step cultural reasoning. Multi-step reasoning questions require sequential derivation  
 181        and/or synthesis from multiple independent knowledge components, not just single-fact recall  
 182        or paraphrase (detailed examples in Appendix A.2). Question templates are not fixed for preserving  
 183        variety in question styles. Each national subset has two components: a *multimodal* and a *text-only*  
 184        component. All QA data were authored in official local languages and translated to English using  
 185        Claude 4 for Tamil and GPT-4o for other languages. All translations underwent manual verification;  
 186        mistranslations were corrected, and for terms lacking standard English equivalents, we applied  
 187        phonetic transliteration or adopted the locally prevalent rendering.

188        **Multimodal Component.** Annotators created VQA items where the correct answer requires visual  
 189        understanding. Each VQA question was rephrased into a semantically-equivalent text-only MCQ.  
 190        The answer options and the correct answer were kept unchanged. We provided both the original and  
 191        rephrased items in English and the local language. We also generated speech inputs by converting  
 192        the text to audio using high-quality TTS systems (Appendix A.4 provides TTS toolkit and speech  
 193        data building details). For Spoken QA, we considered two configurations: (i) converting only the  
 194        question stem to speech while keeping textual options, and (ii) converting both the stem and the  
 195        options to speech. To preserve comparability with VQA under controlled variables and to reduce  
 196        ambiguity introduced by fully spoken options, our main experiments adopt the “spoken stem +  
 197        textual options” configuration across five evaluation dimensions. Results for the fully spoken setting  
 198        (spoken stem and options) on the test set are reported in Appendix A.10 for reference. To reflect  
 199        accent effects, we produced English speech in both accent-neutral and locally accented versions.

200        **Text-only Component.** This component contains questions that were not suitable for pairing with an  
 201        image or are inherently text-based. We applied the same requirement that at least 60% of questions  
 202        involve multi-step reasoning. All questions were created in multiple languages.

203        For each question, annotators additionally identified the requisite *knowledge points*. As shown in  
 204        Fig. 1, these denote the minimal information necessary to arrive at the correct answer, usually  
 205        summarized in a few concise sentences. These knowledge points are included with the dataset and  
 206        support the evaluation of whether model explanations reflect culturally authentic reasoning.

207        For each question, annotators identified the requisite knowledge points, the minimal information  
 208        needed for the correct answer, typically summarized in a few sentences (shown in Fig. 1). These  
 209        knowledge points are used for model’s cultural awareness grounding validation (Section 4.3).

### 210        3.5 HUMAN REVIEW AND REVISION

211        After each country team completed a draft, in-country linguists conducted quality reviews. The  
 212        review covered: ambiguity in wording, accuracy of English translations, clarity and fluency of the  
 213        speech data, completeness of knowledge points, and appropriateness of answer options. Teams  
 214        revised their subsets based on this feedback, yielding the final high-quality release.

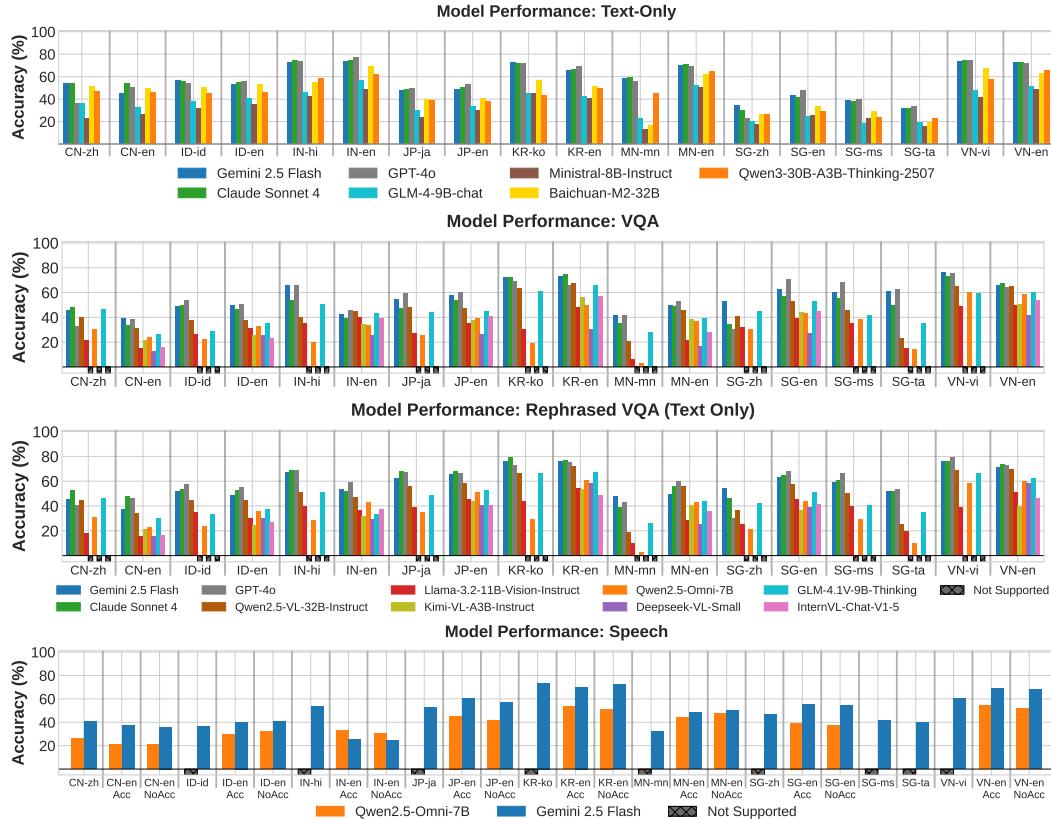


Figure 2: **Performance of LLMs on MMA-Asia across modalities.** Exact values are provided in Appendix A.8. For each country and modality, the dataset contains 500 questions presented in multiple languages. The vertical axis reports Accuracy (%), defined as the number of items where the model’s chosen option exactly matches the correct option, divided by 500. The x-axis label  $\{\text{Country}\}-\{\text{Language}\}$  denotes the cultural dataset for  $\{\text{Country}\}$ , presented in  $\{\text{Language}\}$ .

## 4 EXPERIMENTS

We evaluate existing LLMs on the MMA-ASIA benchmark. Unless stated otherwise, all runs are *zero-shot* with a unified prompt template whose language matches the question language (prompts and experimental settings are in Appendix A.6). We report results along five dimensions and, for each, analyze the factors that drive performance. We access three closed-source models: GPT-4o (OpenAI et al., 2024), Claude-Sonnet-4, and Gemini 2.5 Pro (Team, 2025a), and eleven open-source multilingual or multimodal models, including the Qwen (Bai et al., 2025; Team, 2025b), LLaMA (Touvron et al., 2023), and GLM (GLM et al., 2024; Team et al., 2025c) families. Models and tasks are detailed in Appendix A.5. For models without multilingual support, we report English-only scores for comparability. For speech evaluation, we include only models that accept *speech tokens* directly; models that require intermediate automatic speech recognition (ASR) are excluded.

### 4.1 MAIN RESULTS

Figure 2 summarizes accuracy on MMA-ASIA. Nearly all state-of-the-art models score below 80% and most below 50%, highlighting the benchmark’s difficulty. Closed-source models outperform open-source models; even the strongest open-source family (Qwen) trails the closed-source average by more than 10 percentage points. Performance varies with (i) the resource level of the language and (ii) the evaluation modality. The following subsections analyze: (1) cultural-awareness disparities across countries and languages, (2) cross-lingual and cross-modal consistency, (3) cultural-awareness grounding, and (4) cultural knowledge generalization.

### 4.2 CULTURAL AWARENESS DISPARITY

**Across countries (language factor controlled).** To isolate country effects, we compare each model’s scores across countries within a fixed modality and, for each country, retain the model’s

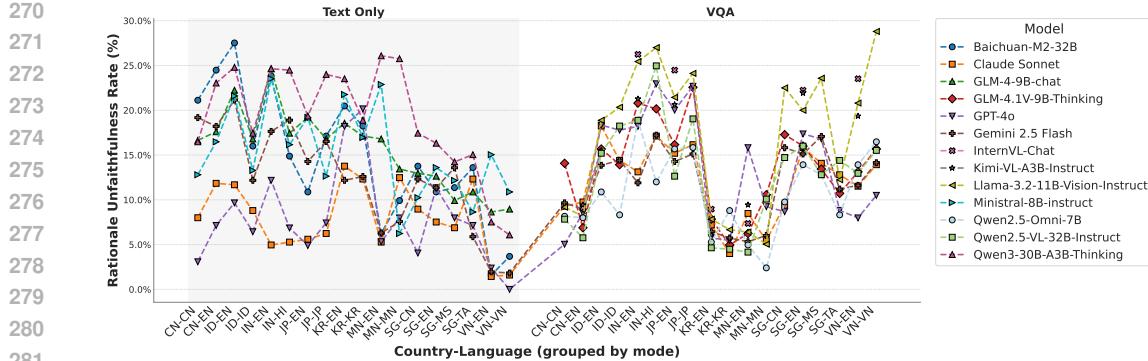


Figure 3: **Rationale Unfaithfulness Rates of LLMs across text-only and VQA.** Similar trends are observed for Rephrase VQA and Spoken QA; detailed results are provided in Appendix A.14.

best score over its available languages (Figure 2, Tables 13, 14, 15, 16). For example, in the text-only setting, GPT-4o scores 71.4% in Korean and 68.8% in English; we use the precision in English to represent the awareness of Korean culture in GPT-4o, to avoid confounding cultural competence with language proficiency. Models show higher awareness for Korean culture on average (63.98% across 4 modalities), plausibly due to global diffusion and richer data availability Jang et al. (2024); Dal Yong (2018). Remarkably, Vietnamese culture (62.96%) is on par with Korean, likely reflecting Vietnam’s high social-media penetration (79.8% of the population) and thus large volumes of user-generated content (DataReportal et al., 2025). In contrast, China and India exhibit larger gaps, consistent with multilayered cultural forms and greater regional heterogeneity. Mongolia trails further, consistent with low-resource language settings and sparser training corpora.

**Across languages (within culture).** English prompts often outperform low-resource languages, reflecting the breadth of English corpora and limited cross-lingual transfer (Hu et al., 2025; Zheng et al., 2025b;a). This advantage diminishes or reverses for medium/high-resource local languages (e.g., Chinese, Japanese), where culture-specific terms and proper names are well represented locally but rare in English corpora, hurting retrieval and grounding. For example, “*乌拉* (Wuhu),” a lineage among Uyghur ancestral groups, lacks a standard English equivalent; transliteration is rare and ambiguous in English data. Thus, when the model is competent in the relevant local language, using that language can yield better cultural grounding than English.

**Across modalities (holding language fixed).** We observe a consistent ordering: *text-only* > *VQA* > *spoken QA*. Data availability follows the same order (text ≈ image–text ≈ raw speech). Speech adds uncertainty (noise, homophony), and many architectures encode modalities separately and then fuse downstream, introducing alignment/compression losses that widen gaps. Interestingly, in Speech, Qwen and Gemini outperform their standard English baselines in 6 and 5 country-specific cultural settings, respectively (Figure 2). Accents appear to serve as a prior cue for specific cultures, enhancing the models’ accuracy on corresponding tasks. We attribute this to the co-occurrence of accents and their related cultural content within the data (see Appendix A.12 for detailed analysis). From the foregoing analysis, it is evident that LLMs in Asian cultural contexts also display cultural and modality biases shaped by data distributions; furthermore, given the limited effectiveness of cross-lingual cultural knowledge transfer, English cannot be assumed to perform reliably better on culture-related tasks. In contrast, in the speech modality, accents, often treated as noise, paradoxically serve as effective cultural cues that activate relevant context and improve performance. In speech, accents, typically considered noise, actually serve as effective cultural cues that activate relevant context and improve performance.

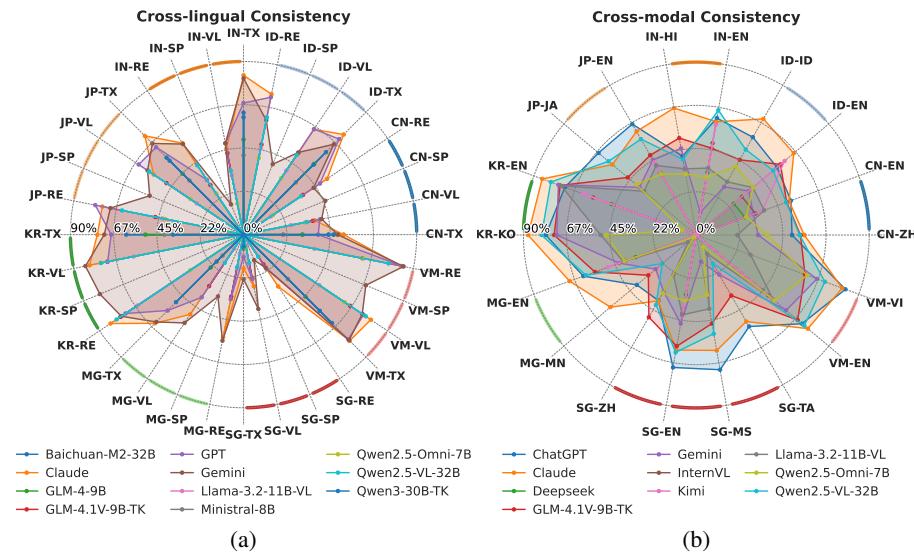
Meanwhile, further analysis of how different cultural categories affect model performance, as well as the relationship between cultural categories and language are provided in Sec A.7.

### 4.3 CULTURAL AWARENESS GROUNDING VALIDATION

MCQs are convenient but can be solved via shortcuts (e.g., option elimination) rather than grounded knowledge (Myung et al., 2025; Romero et al., 2024; Wang et al., 2024; 2025; Hartmann et al., 2023). We adopt two measures: (i) **Retained background knowledge:** Each item accompanied by its *knowledge points* (supporting evidence). (ii) **Explainable responses:** During testing, models must provide a textual rationale for their choice.

324  
 325 Table 2: Evaluated under the MCQ+Explanation metric, Text-only modality performance is reported  
 326 as Accuracy (%), defined as the number of items where the model’s choice exactly matches the  
 327 correct option, divided by 500.

Model	CN-en	ID-en	IN-en	JP-en	KR-en	MN-en	SG-en	VN-en
Gemini 2.5 Flash	27.2	32.1	56.0	33.9	53.0	63.7	32.1	71.0
Claude Sonnet 4	41.8	43.1	69.2	45.0	52.2	65.7	34.3	70.8
GPT-4o	43.1	46.1	65.0	48.3	50.6	64.1	36.3	69.4
GLM-4-9B-chat	15.0	27.2	32.4	14.4	24.2	35.8	12.2	42.4
Mistral-8B-Instruct	9.5	14.0	25.5	11.0	19.0	27.9	11.6	33.6
Baichuan-M2-32B	25.3	25.7	45.2	29.9	31.1	55.4	22.7	61.2
Qwen3-30B-A3B	22.8	21.4	37.6	18.6	26.3	38.9	12.9	58.3



353 Figure 4: (a) **Cross-lingual consistency with fixed country and modality** and (b) **cross-modal**  
 354 **consistency with fixed language and country.** TX/VL/RE/SP represent text QA, visual QA,  
 355 rephrase QA, and speech QA.

356 We use a LLM-as-Judge approach to verify whether, *given a correct answer*, the model’s explanation  
 357 matches the item’s knowledge points. To reduce variability across judges, we require explanations  
 358 in English. The human consistency checks, the multimodel consistency evaluation, and all LLM-  
 359 as-Judge parameter settings are described in Appendix A.13. Figure 3 reports the *Rationale Un-*  
 360 *faithfulness Rate* (RUR) for text-only and VQA items, defined as the proportion of correct answers  
 361 whose explanations contradict or omit the required knowledge. Despite strong overall accuracy,  
 362 proprietary models (Claude, GPT, and Gemini) still show RUR values between 5% and 20%. The  
 363 issue is more pronounced for open-source models: Llama-3.2-11B-Vision-Instruct and Qwen3-30B-  
 364 A3B-Thinking have the highest RUR, with Llama’s rate particularly elevated on non-English inputs.  
 365 Qwen3-30B-A3B-Thinking often produces long explanations that contain hallucinations and some-  
 366 times derives the correct option from premises that contradict the ground truth, indicating reliance  
 367 on generic heuristics rather than culturally grounded reasoning.

368 These results suggest that MCQs alone do not reliably measure cultural understanding (especially  
 369 for open-source models) and should be augmented with culture-grounded verification, including  
 370 knowledge-point checks and rationale assessment. Table 2 reports the Text-Only results of the  
 371 MCQ+Explanation evaluation, where model performance is assessed jointly on the selected  
 372 options and the generated explanations. This explanation-based adjudication removes spurious correct  
 373 answers and yields a more faithful estimate of the models’ cultural awareness. The results for the  
 374 other modalities are provided in Sec. A.9.

#### 375 4.4 CULTURAL-AWARENESS CONSISTENCY ACROSS LANGUAGES

376 Cross-lingual cultural-awareness consistency is defined as the degree to which a model gives consis-  
 377 tent outputs to semantically equivalent prompts posed in different languages, irrespective of answer  
 378 correctness, which can be expressed as Eq. 1 (Wang et al., 2024).

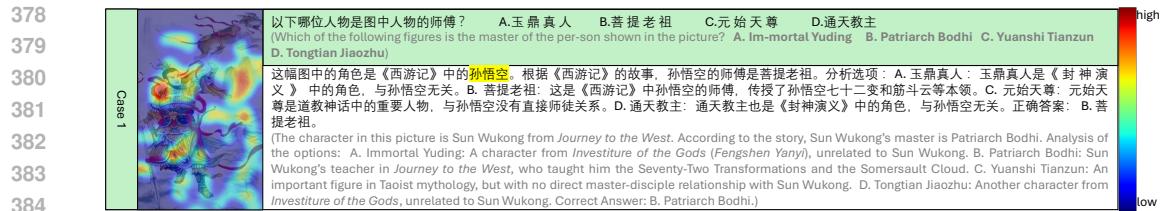


Figure 5: **Attention heatmap visualization over image regions during incorrect model answers.**  
Color scale from blue (low) to red (high) indicates increasing model attention.

$$\text{Consistency}_s = \frac{1}{N} \sum_{i=1}^N \frac{1}{\binom{m_i}{s}} \sum_{S \subseteq L_i, |S|=s} \mathbf{1} \left( \left| \{a_i^\ell : \ell \in S\} \right| = 1 \right). \quad (1)$$

where  $L_i$  denotes the set of languages available for question  $i$ , and  $m_i = |L_i|$  is its cardinality;  $S \subseteq L_i$  with  $|S| = s$  denotes any size- $s$  language subset;  $a_i^\ell$  is the model’s answer to question  $i$  when prompted in language  $\ell$ ;  $\mathbf{1}(\cdot)$  is the indicator function; and  $\left| \{a_i^\ell : \ell \in S\} \right| = 1$  asserts that all answers within  $S$  are identical. If  $m_i < s$ , there are no valid subsets and item  $i$  contributes zero.

As shown in Figure 4a, contemporary multilingual LLMs exhibit weak cross-lingual consistency on culturally grounded tasks in both text-only and VQA settings. The weakness is most evident for language pairs with large resource gaps. For Mongolian culture, the disparity between Mongolian and English yields only 65.2% consistency for Claude, while all open-source models remain below 50%. By contrast, Korean culture shows higher consistency, plausibly reflecting the global diffusion of contemporary Korean media and the resulting multilingual exposure to related knowledge (Jang et al., 2024; Dal Yong, 2018). Consistency also declines sharply as the number of evaluated languages increases. For Singapore-related items, when Chinese, English, Tamil, and Malay are assessed jointly, the maximum consistency does not exceed 45% (Gemini on VQA), despite relatively high pairwise values of 60.60% (EN–TA), 64.20% (EN–MS), and 55.20% (EN–ZH). In some culturally challenging cases, visual context can partially bridge languages: for Indian culture, GLM-4.1 achieves 44.20% cross-lingual consistency on Hindi VQA, which is 13.2 points higher than its rephrased text-only counterpart, although both remain low.

From these observations, we observe that the consistency between languages depends on the data resources and cultural exposure. Resource asymmetry degrades consistency, whereas cultural prominence helps. Consistency decays nonlinearly as more languages are considered, and strong pairwise agreement does not guarantee multi-language coherence. Visual cues can narrow gaps in certain settings, but are insufficient to overcome the structural limitations of low-resource languages.

#### 4.5 CULTURAL-AWARENESS CONSISTENCY ACROSS MODALITIES

Cross-modal cultural-awareness consistency evaluates whether a model gives the same output for semantically equivalent queries presented in different modalities. Figure 4b shows that, across the eight Asian countries, the pattern largely matches the cross-lingual case: averaged over models, cross-lingual consistency is 48%, and cross-modal consistency rarely exceeds 67%. This gap indicates asymmetric transfer of cultural knowledge across modalities. Under low-resource language settings, almost all models struggle to maintain stable cross-modal answers. Within the same national context, medium- to high-resource local languages typically yield higher cross-modal consistency than English. To examine the observed ordering *text-only* > VQA > *spoken QA* (see Section 4.2), we conduct a detailed error analysis with Qwen2.5-VL-32B-Instruct. Because spoken QA adds additional complexities (noise, accents, intonation), our analysis in this section focuses on VQA versus text-only performance; we leave a fuller study of speech to future work.

We specifically isolate instances where the model succeeded with pure text input but failed in the VQA context. Apart from the most common errors arising from a lack of understanding of culture-related image contexts, our analysis reveals two additional predominant categories of errors:

**The pitfall of prompt-guided selective attention.** Models often tend to focus predominantly on explicitly mentioned objects in prompts, whereas cultural VQA requires a more nuanced ability to identify culture-specific visual cues within images. To validate whether the models’ focuses are truly culture-specific, we extract visual evidence using answer-conditioned multi-layer Grad-CAM (Selvaraju et al., 2019), interpolating and mapping the resulting heatmaps back to the original image for visualization. Specifically, given an image  $v$ , a textual prompt  $x$  and an answer  $y_{a:b} = (y_a, \dots, y_b)$  autoregressively produced by a model with parameters  $\theta$ , we define an answer-conditioned objective on the log-likelihood of the answer tokens as Eq. 2. Here we use the token-sum objective. For mem-

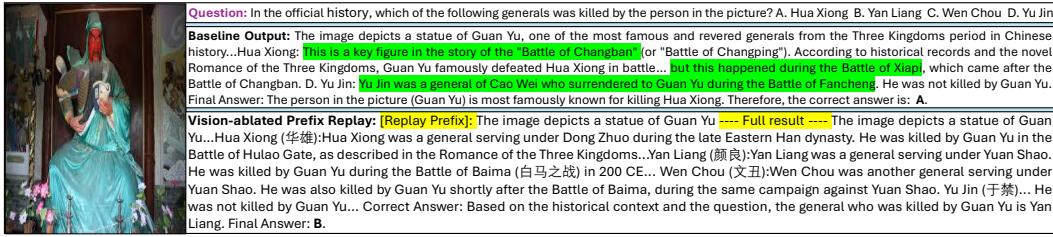


Figure 6: Comparison of baseline vs. “VPR” outputs, with hallucinations highlighted in green.

ory efficiency, only the visual tower is set to require gradients. We denote the forward activation of the  $\ell$ -th visual block as  $A(\ell) \in \mathbb{R}^{C \times H \times W}$ , the gradient of the block is  $\mathbf{G}^{(\ell)}$ . A per-layer Grad-CAM is built via channelwise inner product followed by ReLU (Eq. 3).

$$J_{\text{sum}} = - \sum_{t=a}^b \log p_{\theta}(y_t \mid v, x, y_{<t}) \quad (2) \quad \mathbf{M}^{(\ell)} = \text{ReLU} \left( \sum_{c=1}^C \mathbf{G}_c^{(\ell)} \odot \mathbf{A}_c^{(\ell)} \right) \in \mathbb{R}^{H \times W} \quad (3)$$

After lightweight smoothing and normalization of  $\mathbf{M}^{(\ell)}$ , the resulting  $\hat{\mathbf{M}}^{(\ell)}$  then undergoes cross-layer aggregation and bilinear interpolation upsampling on the  $k$  blocks at the end of the visual tower, as defined by Eqs. 4 and 5, to the original image resolution. For more robust results, we specifically choose the last 3 blocks and do cross-layer aggregation by averaging them.

$$\mathbf{M} = \text{Agg}_{\ell \in \mathcal{L}} (\hat{\mathbf{M}}^{(\ell)}) \quad (4) \quad \tilde{\mathbf{M}} = \text{Bilinear}(\mathbf{M}, H_{\text{img}}, W_{\text{img}}) \in [0, 1]^{H_{\text{img}} \times W_{\text{img}}} \quad (5)$$

This phenomenon is clearly revealed by the heatmaps in Figure 5, showing how the model focuses on subjects explicitly mentioned in the prompt when answering questions. However, this selective attention can unfortunately lead the model to overlook other critically important local details within the image, consequently resulting in erroneous inferences. In Case 1 (Fig. 5), when the model is presented with the question “Which of the following figures is the master of the person shown in the picture?” in Chinese, its attention is predominantly drawn to the figure on the left side of the image, who is explicitly referenced in the prompt. However, the model overlooks the Howling Celestial Dog in the upper-right corner of the image, a crucial clue for identifying Erlang Shen. Additional attention-visualization examples and analyses are given in Appendix A.15.

**Image tokens contribute to reasoning hallucinations.** We find that visual content increases reasoning hallucinations in models compared to text-only QA, despite accurate image recognition capabilities. In Figure 6, while the baseline model correctly identifies “Guan Yu”, it still produces multiple reasoning hallucinations (highlighted in green). However, text-only Rephrase VQA queries show no such hallucinations (Appendix A.11), indicating that reasoning errors likely originate from image tokens and suggest modality-specific bias in multimodal reasoning. To test this hypothesis, we propose Vision-ablated Prefix Replay (VPR), which generates image descriptions then removes visual conditions while maintaining fixed prefixes for subsequent reasoning (details in Appendix A.11). In Figure 6, VPR conditions the model’s reasoning on the generated caption ‘The image depicts a statue of Guan Yu’ while discarding visual tokens, eliminating hallucinations and producing the correct answer. Furthermore, we select 50 questions where VQA answers contained hallucinations but corresponding text-only queries were correct. VPR eliminates hallucinations and produced correct answers for 19 of these cases (38%), supporting our hypothesis.

Multimodal models demonstrate significant inconsistency in cultural awareness across modalities, indicating flawed cultural knowledge transfer. In VQA, this deficiency stems from two core issues: “*selective attention pitfall*” where models over-focus on text-prompted subjects while missing key visual cues, and visual token-induced “*reasoning hallucinations*”.

#### 4.6 CULTURAL KNOWLEDGE GENERALIZATION

Prior work (Balepur et al., 2024; Molfese et al., 2025; Zheng et al., 2023) suggests that scaling increases LLMs’ factual memory but not genuine logical generalization. To distinguish whether cultural multi-step reasoning errors stem from knowledge gaps or generalization deficits, we conduct a deconstruction study. We decompose each question into atomic sub-questions testing single knowledge points, which models answer first (see Table 23 for an example). We then evaluate the original question under two conditions: (i) with in-context “sub-question → model answer” pairs, and (ii) from scratch. If models solve all subquestions but fail the original question from scratch, they possess the knowledge but cannot transfer it, indicating generalization deficits. If they err on sub-questions, failure likely reflects missing culture-specific knowledge.

486  
 487 Table 3: Cultural knowledge generalization performance of different models under different lan-  
 488 guage settings. We define Successful Correction as correct sub-questions and final answers; Inte-  
 489 gration Failure as incorrect final synthesis despite correct sub-questions; and Sub-question Failure  
 490 as errors in the initial reasoning stage.

491 Language	492 Model	493 Successful Correction (%)	494 Integration Failure (%)	495 Sub-question Failure (%)
492 <b>English</b>	Claude	35.3	2.6	62.1
	Qwen	18.5	10.7	70.8
	Mistral	18.9	1.4	80.7
	GLM	30.9	2.0	67.1
496 <b>Chinese</b>	Claude	41.1	3.5	55.4
	Qwen	35.5	2.6	61.9
	Mistral	25.2	1.0	73.8
	GLM	40.5	2.8	56.7
499 <b>Indonesian</b>	Claude	20.2	9.2	70.6
	Qwen	13.4	3.7	82.9
	Mistral	4.4	1.1	94.5
	GLM	6.5	5.3	88.2

503 We investigate the failure modes of Claude Sonnet 4, GLM-4-9B-chat, Mistral-8B-Instruct, and  
 504 Qwen3-30B-A3B-Thinking on multi-step reasoning tasks across English, Chinese, and Indonesian.  
 505 Table 10 summarizes the number of failure cases for each model in each language. Our experi-  
 506 ment utilizes problems where each model initially failed, decomposing them into 3–8 atomic sub-  
 507 questions with ground-truth answers. An LLM (GPT-4o) judges the correctness of each step. The  
 508 results in Table 3 reveal distinct failure patterns influenced significantly by language resource avail-  
 509 ability and model architecture.

510 **Knowledge Availability vs. Integration:** In high-resource languages, errors primarily stem from  
 511 integration failures rather than knowledge deficits. Post-decomposition, Claude and GLM achieve  
 512 successful correction rates of 41.1% and 40.5% in Chinese, demonstrating adequate cultural knowl-  
 513 edge but impaired single-step activation. Conversely, Indonesian performance is bottlenecked by  
 514 fundamental knowledge gaps, with sub-question failure rates reaching 94.5% (Mistral) and 88.2%  
 515 (GLM), confirming that decomposition cannot remedy training data deficiencies.

516 **The “Lost in Integration” Phenomenon:** Beyond knowledge gaps, integration failures, where  
 517 models answer sub-questions correctly but fail final synthesis, constitute a significant error source.  
 518 This is pronounced in Qwen’s English performance: despite reasonable sub-question accuracy, inte-  
 519 gration failure reaches 10.7% (versus ~2% for GLM and Mistral), showing reasoning misalignment  
 520 where isolated fact retrieval succeeds but cross-lingual synthesis fails. Claude exhibits similar pat-  
 521 terns in Indonesian (9.2% integration failure), indicating that logical synthesis degrades faster than  
 522 fact recall in lower-resource languages. Cross-lingual cultural knowledge transfer remains challeng-  
 523 ing regardless of model scale. In high-resource settings, closed-source models (Claude) demon-  
 524 strate broad knowledge coverage but limited cultural generalization, while open-source models face  
 525 knowledge gaps and poor transferability. In low-resource languages, closed-source models show  
 526 greater degradation in cultural generalization, whereas knowledge deficits account for over 88.5%  
 527 of open-source model failures, underscoring their inadequate low-resource language training.

## 528 5 CONCLUSION

529 We introduced MMA-ASIA, a tri-modal (text, image, speech), multilingual benchmark and frame-  
 530 work for evaluating cultural awareness in LLMs across 8 Asian countries and 10 languages. Our  
 531 contributions include an aligned, human-curated dataset with substantial multi-step reasoning, a  
 532 five-dimensional protocol that measures accuracy, cross-lingual and cross-modal consistency, cul-  
 533 tural knowledge generalization, and grounding validity, and analysis tools that reveal shortcut use.  
 534 Results show persistent data-driven cultural bias, uneven cross-lingual transfer, and fragile multi-  
 535 modal reasoning (selective visual attention and image-induced hallucinations). At the same time,  
 536 accented speech can act as a useful cultural cue. Models also struggle to integrate known facts  
 537 into multi-step reasoning, indicating a generalization bottleneck. We argue for consistency- and  
 538 grounding-aware evaluation, as well as methods that strengthen cross-modal alignment and broaden  
 539 high-quality coverage in low-resource languages. MMA-ASIA provides data, protocols, and base-  
 lines to track progress toward culturally reliable multimodal LLMs.

540 **ICLR ETHICS STATEMENT**

541 **CONTRIBUTE TO SOCIETY AND TO HUMAN WELL-BEING**

543 • We acknowledge that people worldwide are stakeholders in computing, and we will use our  
544 skills for the benefit of society, its members, and the natural environment.

545 • We strive to minimise negative consequences, including threats to health, safety, personal  
546 security, and privacy. Our dataset is created solely from cultural facts and excludes any  
547 negative or false information.

548 • When the interests of multiple groups conflict, we give increased attention and priority to  
549 the needs of those who are less advantaged.

550 • In creating the dataset, we deliberately account for cultural diversity and base all content  
551 strictly on verifiable cultural facts.

553 **UPHOLD HIGH STANDARDS OF SCIENTIFIC EXCELLENCE**

554 • We are committed to open enquiry, intellectual rigour, integrity, and collaboration.

555 • We report accurately and honestly. We do not make false or misleading claims, fabricate or  
556 falsify data, or misrepresent results. Methods and results are presented in a transparent and  
557 reproducible manner.

558 • All content is developed from publicly available data; no personal private information is  
559 collected or disclosed. This work involves no interventions on human subjects and poses  
560 no risk of harm; accordingly, ethical approval from an institutional review board was not  
561 required.

562 • We acknowledge all contributions to the research and comply with agreements related to  
563 intellectual property, publication, and authorship.

564 **AVOID HARM**

566 • Our questions are created from publicly available, verifiable cultural facts. The dataset  
567 contains no pornographic or violent content, nor any racist or discriminatory material.

568 **BE HONEST, TRUSTWORTHY, AND TRANSPARENT**

570 • We describe the professional backgrounds of contributors in Section A, and all researchers  
571 adhere to principles of honesty, trustworthiness, and transparency.

572 **BE FAIR AND TAKE ACTION NOT TO DISCRIMINATE**

573 • Dataset creators are required to avoid content that is racist or contains hate speech. Each  
574 contributor's work is assessed fairly.

576 **RESPECT THE WORK REQUIRED TO PRODUCE NEW IDEAS AND ARTEFACTS**

577 • We properly cite all prior research, tools, models, and commercial software that we use or  
578 build upon. We treat all project members with respect and equity.

579 • All images are used under appropriate Creative Commons licences and solely for research  
580 purposes; any subsequent release will strictly follow the relevant licence terms.

581 • Except for the image subset, all dataset content is authored by our team, and no infringement  
582 is involved.

583 **RESPECT PRIVACY**

585 • The dataset content is derived from publicly available and verifiable cultural facts.

586 • We apply post-processing to images to protect any potential personal privacy.

587 • We do not disclose any private information of team members or contributors.

589 **HONOUR CONFIDENTIALITY**

590 • This work is not subject to any non-disclosure agreements. Upon paper acceptance, we will  
591 openly release the dataset.

594 REPRODUCIBILITY STATEMENT  
595

596 DID YOU REPORT THE NUMBER OF PARAMETERS IN THE MODELS USED, THE TOTAL  
597 COMPUTATIONAL BUDGET (E.G., GPU HOURS), AND COMPUTING INFRASTRUCTURE USED?  
598

600 Yes, all this information can be found in Appendix A.5, A.6, A.4.  
601

602 DID YOU DISCUSS THE EXPERIMENTAL SETUP, INCLUDING HYPERPARAMETER SEARCH AND  
603 BEST-FOUND HYPERPARAMETER VALUES?  
604

605 Yes, all this information can be found in Appendix A.6.  
606

607 DID YOU REPORT DESCRIPTIVE STATISTICS ABOUT YOUR RESULTS (E.G., ERROR BARS  
608 AROUND RESULTS, SUMMARY STATISTICS FROM SETS OF EXPERIMENTS), AND IS IT  
609 TRANSPARENT WHETHER YOU ARE REPORTING THE MAX, MEAN, ETC. OR JUST A SINGLE  
610 RUN?  
611

612 Yes, it can be found in Section 4 and Appendix A.6, A.8, A.11, A.15, A.10 in the paper.  
613

614 IF YOU USED EXISTING PACKAGES (E.G., FOR PREPROCESSING, FOR NORMALIZATION, OR  
615 FOR EVALUATION, SUCH AS NLTK, SPACY, ROUGE, ETC.), DID YOU REPORT THE  
616 IMPLEMENTATION, MODEL, AND PARAMETER SETTINGS USED?  
617

618 Yes, you can find it in the Appendix A.4, A.6.  
619

620 DID YOU USE HUMAN ANNOTATORS (E.G., CROWDWORKERS) OR RESEARCH  
621 WITH HUMAN PARTICIPANTS?  
622

623 Yes.  
624

625 DID YOU REPORT THE FULL TEXT OF INSTRUCTIONS GIVEN TO PARTICIPANTS, INCLUDING  
626 E.G., SCREENSHOTS, DISCLAIMERS OF ANY RISKS TO PARTICIPANTS OR ANNOTATORS, ETC.?  
627

628 Yes, you can find it in the Appendix A.3.  
629

630 DID YOU REPORT INFORMATION ABOUT HOW YOU RECRUITED (E.G., CROWDSOURCING  
631 PLATFORM, STUDENTS) AND PAID PARTICIPANTS, AND DISCUSS IF SUCH PAYMENT IS  
632 ADEQUATE GIVEN THE PARTICIPANTS' DEMOGRAPHIC (E.G., COUNTRY OF RESIDENCE)?  
633

634 Yes, you can find it in the Appendix A.1.  
635

636 DID YOU DISCUSS WHETHER AND HOW CONSENT WAS OBTAINED FROM PEOPLE WHOSE DATA  
637 YOU'RE USING/CURATING? FOR EXAMPLE, IF YOU COLLECTED DATA VIA CROWDSOURCING,  
638 DID YOUR INSTRUCTIONS TO CROWDWORKERS EXPLAIN HOW THE DATA WOULD BE USED?  
639

640 Yes. The purpose of the dataset for research was clearly defined at the beginning of the paper.  
641

642 DID YOU REPORT THE BASIC DEMOGRAPHIC AND GEOGRAPHIC CHARACTERISTICS OF THE  
643 ANNOTATOR POPULATION THAT IS THE SOURCE OF THE DATA?  
644

645 Yes, you can find it in Appendix A.1.  
646

647

648 REFERENCES  
649

650 Aakanksha, Arash Ahmadian, Beyza Ermis, Seraphina Goldfarb-Tarrant, Julia Kreutzer, Marzieh  
651 Fadaee, and Sara Hooker. The multilingual alignment prism: Aligning global and local pref-  
652 erences to reduce harm. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Pro-  
653 ceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp.  
654 12027–12049, Miami, Florida, USA, November 2024. Association for Computational Linguis-  
655 tics. doi: 10.18653/v1/2024.emnlp-main.671. URL <https://aclanthology.org/2024.emnlp-main.671/>.

656 Muhammad Farid Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Shivdutt Singh,  
657 Alham Fikri Aji, Jacki O’Neill, Ashutosh Modi, and Monojit Choudhury. Towards measur-  
658 ing and modeling “culture” in LLMs: A survey. In Yaser Al-Onaizan, Mohit Bansal, and  
659 Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Nat-  
660ural Language Processing*, pp. 15763–15784, Miami, Florida, USA, November 2024. Asso-  
661 ciation for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.882. URL  
662 <https://aclanthology.org/2024.emnlp-main.882/>.

663 Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang  
664 Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, local-  
665 ization, text reading, and beyond, 2023. URL <https://arxiv.org/abs/2308.12966>.

666 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,  
667 Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan,  
668 Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng,  
669 Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report, 2025.  
670 URL <https://arxiv.org/abs/2502.13923>.

671 Nishant Balepur, Abhilasha Ravichander, and Rachel Rudinger. Artifacts or abduction: How do  
672 LLMs answer multiple-choice questions without the question? In Lun-Wei Ku, Andre Martins,  
673 and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Com-  
674 putational Linguistics (Volume 1: Long Papers)*, pp. 10308–10330, Bangkok, Thailand, August  
675 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.555. URL  
676 <https://aclanthology.org/2024.acl-long.555/>.

677 Zhe Chen, Jiannan Wu, Wenhui Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong  
678 Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. Internvl:  
679 Scaling up vision foundation models and aligning for generic visual-linguistic tasks, 2024a. URL  
680 <https://arxiv.org/abs/2312.14238>.

681 Zhe Chen, Jiannan Wu, Wenhui Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong  
682 Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning  
683 for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer  
684 Vision and Pattern Recognition*, pp. 24185–24198, 2024b.

685 Yu Ying Chiu, Liwei Jiang, Bill Yuchen Lin, Chan Young Park, Shuyue Stella Li, Sahithya Ravi,  
686 Mehar Bhatia, Maria Antoniak, Yulia Tsvetkov, Vered Shwartz, and Yejin Choi. Culturalbench:  
687 A robust, diverse, and challenging cultural benchmark by human-ai culturalteaming, 2025. URL  
688 <https://arxiv.org/abs/2410.02677>.

689 Seungho Cho, Changgeon Ko, Eui Jun Hwang, Junmyeong Lee, Huije Lee, and Jong C. Park.  
690 Language over content: Tracing cultural understanding in multilingual large language models,  
691 2025. URL <https://arxiv.org/abs/2510.16565>.

692 Coqui.ai. Tts: A deep learning toolkit for text-to-speech. <https://github.com/coqui-ai/TTS>, 2025. Accessed: 2025-09-24.

693 Jin Dal Yong. An analysis of the korean wave as transnational popular culture: North american youth  
694 engage through social media as tv becomes obsolete. *International Journal of Communication*,  
695 12:404–422, 2018.

696 DataReportal, Kepios, We Are Social, and Meltwater. Digital 2025: Vietnam, 02 2025. URL  
697 <https://datareportal.com/reports/digital-2025-vietnam>. Accessed 21  
698 Sep 2025.

699

702 Zhihao Du, Qian Chen, Shiliang Zhang, Kai Hu, Heng Lu, Yexin Yang, Hangrui Hu, Siqi Zheng, Yue  
 703 Gu, Ziyang Ma, et al. Cosyvoice: A scalable multilingual zero-shot text-to-speech synthesizer  
 704 based on supervised semantic tokens. *arXiv preprint arXiv:2407.05407*, 2024.

705

706 Zhihao Du, Changfeng Gao, Yuxuan Wang, Fan Yu, Tianyu Zhao, Hao Wang, Xiang Lv, Hui Wang,  
 707 Xian Shi, Keyu An, et al. Cosyvoice 3: Towards in-the-wild speech generation via scaling-up and  
 708 post-training. *arXiv preprint arXiv:2505.17589*, 2025.

709

710 Jacob Eisenstein, Vinodkumar Prabhakaran, Clara Rivera, Dorottya Demszky, and Devyani Sharma.  
 711 Md3: The multi-dialect dataset of dialogues, 2023. URL <https://arxiv.org/abs/2305.11355>.

712

713 Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu  
 714 Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadai Sun, Jiajie Zhang, Jiale Cheng,  
 715 Jiayi Gui, Jie Tang, Jing Zhang, Juanzi Li, Lei Zhao, Lindong Wu, Lucen Zhong, Mingdao Liu,  
 716 Minlie Huang, Peng Zhang, Qinkai Zheng, Rui Lu, Shuaiqi Duan, Shudan Zhang, Shulin Cao,  
 717 Shuxun Yang, Weng Lam Tam, Wenyi Zhao, Xiao Liu, Xiao Xia, Xiaohan Zhang, Xiaotao Gu,  
 718 Xin Lv, Xinghan Liu, Xinyi Liu, Xinyue Yang, Xixuan Song, Xunkai Zhang, Yifan An, Yifan  
 719 Xu, Yilin Niu, Yuantao Yang, Yueyan Li, Yushi Bai, Yuxiao Dong, Zehan Qi, Zhaoyu Wang,  
 720 Zhen Yang, Zhengxiao Du, Zhenyu Hou, and Zihan Wang. Chatglm: A family of large language  
 721 models from glm-130b to glm-4 all tools, 2024.

722

723 Vansh Gupta, Sankalan Pal Chowdhury, Vilém Zouhar, Donya Rooein, and Mrinmaya Sachan.  
 724 Multilingual performance biases of large language models in education, 2025. URL <https://arxiv.org/abs/2504.17720>.

725

726 Valentin Hartmann, Anshuman Suri, Vincent Bindschaedler, David Evans, Shruti Tople, and Robert  
 727 West. Sok: Memorization in general-purpose large language models, 2023. URL <https://arxiv.org/abs/2310.18362>.

728

729 Peng Hu, Sizhe Liu, Changjiang Gao, Xin Huang, Xue Han, Junlan Feng, Chao Deng, and Shujian  
 730 Huang. Large language models are cross-lingual knowledge-free reasoners, 2025. URL <https://arxiv.org/abs/2406.16655>.

731

732 Moonkyoung Jang, Dokyung Kim, and Hyunmi Baek. More than just a fan: the influence of k-pop  
 733 fandom on the popularity of k-drama on a global ott platform. *Applied Economics Letters*, 31(2):  
 734 152–157, 2024.

735

736 Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chap-  
 737 lot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier,  
 738 Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril,  
 739 Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023. URL <https://arxiv.org/abs/2310.06825>.

740

741

742 Eunsu Kim, Juyoung Suk, Philhoon Oh, Haneul Yoo, James Thorne, and Alice Oh. CLICK: A bench-  
 743 mark dataset of cultural and linguistic intelligence in Korean. In Nicoletta Calzolari, Min-Yen  
 744 Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (eds.), *Proceedings*  
 745 *of the 2024 Joint International Conference on Computational Linguistics, Language Resources*  
 746 *and Evaluation (LREC-COLING 2024)*, pp. 3335–3346, Torino, Italia, May 2024. ELRA and  
 747 ICCL. URL <https://aclanthology.org/2024.lrec-main.296/>.

748

749 Jaehyeon Kim, Jungil Kong, and Juhee Son. Conditional variational autoencoder with adversarial  
 750 learning for end-to-end text-to-speech, 2021. URL <https://arxiv.org/abs/2106.06103>.

751

752 Hannah Rose Kirk, Alexander Whitefield, Paul Röttger, Andrew Bean, Katerina Margatina, Juan  
 753 Ciro, Rafael Mosquera, Max Bartolo, Adina Williams, He He, Bertie Vidgen, and Scott A. Hale.  
 754 The prism alignment dataset: What participatory, representative and individualised human feed-  
 755 back reveals about the subjective and multicultural alignment of large language models, 2024.  
 URL <https://arxiv.org/abs/2404.16019>.

756 Julia Kreutzer, Eleftheria Briakou, Sweta Agrawal, Marzieh Fadaee, and Kocmi Tom. *Déjà vu:*  
 757 Multilingual llm evaluation through the lens of machine translation evaluation, 2025. URL <https://arxiv.org/abs/2504.11829>.

758

759 Chaoqun Liu, Wenzuan Zhang, Jiahao Ying, Mahani Aljunied, Anh Tuan Luu, and Lidong Bing.  
 760 *SeaExam and SeaBench: Benchmarking LLMs with local multilingual questions in Southeast*  
 761 *Asia.* In Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.), *Findings of the Association for Compu-*  
 762 *tational Linguistics: NAACL 2025*, pp. 6119–6136, Albuquerque, New Mexico, April 2025. As-  
 763 *sociation for Computational Linguistics.* ISBN 979-8-89176-195-7. doi: 10.18653/v1/2025.findi-  
 764 *ngs-naacl.341.* URL <https://aclanthology.org/2025.findings-naacl.341/>.

765

766 Francesco Maria Molfese, Luca Moroni, Luca Gioffré, Alessandro Scirè, Simone Conia, and  
 767 Roberto Navigli. Right answer, wrong score: Uncovering the inconsistencies of llm evaluation in  
 768 multiple-choice question answering. *arXiv preprint arXiv:2503.14996*, 2025.

769

770 Junho Myung, Nayeon Lee, Yi Zhou, Jiho Jin, Rifki Afina Putri, Dimosthenis Antypas, Hsu-  
 771 vas Borkakoty, Eunsu Kim, Carla Perez-Almendros, Abinew Ali Ayele, Víctor Gutiérrez-  
 772 Basulto, Yazmín Ibáñez-García, Hwaran Lee, Shamsuddeen Hassan Muhammad, Kiwoong Park,  
 773 Anar Sabuhi Rzayev, Nina White, Seid Muhie Yimam, Mohammad Taher Pilehvar, Nedjma  
 774 Ousidhoum, Jose Camacho-Collados, and Alice Oh. Blend: A benchmark for llms on every-  
 775 day knowledge in diverse cultures and languages, 2025. URL <https://arxiv.org/abs/2406.09948>.

776

777 Shravan Nayak, Kanishk Jain, Rabiul Awal, Siva Reddy, Sjoerd van Steenkiste, Lisa Anne Hen-  
 778 dricks, Karolina Stańczak, and Aishwarya Agrawal. Benchmarking vision language models for  
 779 cultural understanding, 2024. URL <https://arxiv.org/abs/2407.10920>.

780

781 Raymond Ng, Thanh Ngan Nguyen, Yuli Huang, Ngee Chia Tai, Wai Yi Leong, Wei Qi Leong,  
 782 Xianbin Yong, Jian Gang Ngui, Yosephine Susanto, Nicholas Cheng, Hamsawardhini Rengarajan,  
 783 Peerat Limkonchotiwat, Adithya Venkatadri Hulagadri, Kok Wai Teng, Yeo Yeow Tong, Bryan  
 784 Siow, Wei Yi Teo, Wayne Lau, Choon Meng Tan, Brandon Ong, Zhi Hao Ong, Jann Railey  
 785 Montalan, Adwin Chan, Sajeban Antonyrex, Ren Lee, Esther Choa, David Ong Tat-Wee, Bing  
 786 Jie Darius Liu, William Chandra Tjhi, Erik Cambria, and Leslie Teo. Sea-lion: Southeast asian  
 787 languages in one network, 2025. URL <https://arxiv.org/abs/2504.05747>.

788

789 OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan  
 790 Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-  
 791 Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol,  
 792 Alex Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Con-  
 793 neau, Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian,  
 794 Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein,  
 795 Andrew Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrey  
 796 Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia,  
 797 Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben  
 798 Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake  
 799 Samic, Bob McGrew, Bobby Spero, Bogo Giertler, Bowen Cheng, Brad Lightcap, Brandon  
 800 Walkin, Brendan Quinn, Brian Guaraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo  
 801 Lugaresi, Carroll Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li,  
 802 Chan Jun Shern, Channing Conger, Charlotte Burette, Chelsea Voss, Chen Ding, Cheng Lu,  
 803 Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim,  
 804 Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley  
 805 Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler,  
 806 Daniel Levin, Daniel Levy, David Carr, David Farhi, David Mely, David Robinson, David Sasaki,  
 807 Denny Jin, Dev Valladares, Dimitris Tsipras, Doug Li, Duc Phong Nguyen, Duncan Findlay,  
 808 Edede Oiwoh, Edmund Wong, Ehsan Asdar, Elizabeth Proehl, Elizabeth Yang, Eric Antonow,  
 809 Eric Kramer, Eric Peterson, Eric Sigler, Eric Wallace, Eugene Brevdo, Evan Mays, Farzad Kho-  
 rasani, Felipe Petroski Such, Filippo Raso, Francis Zhang, Fred von Lohmann, Freddie Sulit,  
 Gabriel Goh, Gene Oden, Geoff Salmon, Giulio Starace, Greg Brockman, Hadi Salman, Haiming  
 Bao, Haitang Hu, Hannah Wong, Haoyu Wang, Heather Schmidt, Heather Whitney, Heewoo Jun,  
 Hendrik Kirchner, Henrique Ponde de Oliveira Pinto, Hongyu Ren, Huiwen Chang, Hyung Won  
 Chung, Ian Kivlichan, Ian O’Connell, Ian O’Connell, Ian Osband, Ian Silber, Ian Sohl, Ibrahim

810 Okuyucu, Ikai Lan, Ilya Kostrikov, Ilya Sutskever, Ingmar Kanitscheider, Ishaan Gulrajani, Ja-  
 811 cob Coxon, Jacob Menick, Jakub Pachocki, James Aung, James Betker, James Crooks, James  
 812 Lennon, Jamie Kiros, Jan Leike, Jane Park, Jason Kwon, Jason Phang, Jason Teplitz, Jason Wei,  
 813 Jason Wolfe, Jay Chen, Jeff Harris, Jenia Varavva, Jessica Gan Lee, Jessica Shieh, Ji Lin, Jiahui  
 814 Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joanne Jang, Joaquin Quinonero Candela, Joe Beutler, Joe  
 815 Landers, Joel Parish, Johannes Heidecke, John Schulman, Jonathan Lachman, Jonathan McKay,  
 816 Jonathan Uesato, Jonathan Ward, Jong Wook Kim, Joost Huizinga, Jordan Sitkin, Jos Kraaijeveld,  
 817 Josh Gross, Josh Kaplan, Josh Snyder, Joshua Achiam, Joy Jiao, Joyce Lee, Juntang Zhuang,  
 818 Justyn Harriman, Kai Fricke, Kai Hayashi, Karan Singh, Katy Shi, Kavin Karthik, Kayla Wood,  
 819 Kendra Rimbach, Kenny Hsu, Kenny Nguyen, Keren Gu-Lemberg, Kevin Button, Kevin Liu, Kiel  
 820 Howe, Krithika Muthukumar, Kyle Luther, Lama Ahmad, Larry Kai, Lauren Itow, Lauren Work-  
 821 man, Leher Pathak, Leo Chen, Li Jing, Lia Guy, Liam Fedus, Liang Zhou, Lien Mamitsuka,  
 822 Lilian Weng, Lindsay McCallum, Lindsey Held, Long Ouyang, Louis Feuvrier, Lu Zhang, Lukas  
 823 Kondraciuk, Lukasz Kaiser, Luke Hewitt, Luke Metz, Lyric Doshi, Mada Aflak, Maddie Simens,  
 824 Madelaine Boyd, Madeleine Thompson, Marat Dukhan, Mark Chen, Mark Gray, Mark Hudnall,  
 825 Marvin Zhang, Marwan Aljubeh, Mateusz Litwin, Matthew Zeng, Max Johnson, Maya Shetty,  
 826 Mayank Gupta, Meghan Shah, Mehmet Yatbaz, Meng Jia Yang, Mengchao Zhong, Mia Glaese,  
 827 Mianna Chen, Michael Janner, Michael Lampe, Michael Petrov, Michael Wu, Michele Wang,  
 828 Michelle Fradin, Michelle Pokrass, Miguel Castro, Miguel Oom Temudo de Castro, Mikhail  
 829 Pavlov, Miles Brundage, Miles Wang, Minal Khan, Mira Murati, Mo Bavarian, Molly Lin, Murat  
 830 Yesildal, Nacho Soto, Natalia Gimelshein, Natalie Cone, Natalie Staudacher, Natalie Summers,  
 831 Natan LaFontaine, Neil Chowdhury, Nick Ryder, Nick Stathas, Nick Turley, Nik Tezak, Niko Fe-  
 832 lix, Nithanth Kudige, Nitish Keskar, Noah Deutsch, Noel Bundick, Nora Puckett, Ofir Nachum,  
 833 Ola Okelola, Oleg Boiko, Oleg Murk, Oliver Jaffe, Olivia Watkins, Olivier Godement, Owen  
 834 Campbell-Moore, Patrick Chao, Paul McMillan, Pavel Belov, Peng Su, Peter Bak, Peter Bakkum,  
 835 Peter Deng, Peter Dolan, Peter Hoeschele, Peter Welinder, Phil Tillet, Philip Pronin, Philippe  
 836 Tillet, Prafulla Dhariwal, Qiming Yuan, Rachel Dias, Rachel Lim, Rahul Arora, Rajan Troll, Ran-  
 837 dall Lin, Rapha Gontijo Lopes, Raul Puri, Reah Miyara, Reimar Leike, Renaud Gaubert, Reza  
 838 Zamani, Ricky Wang, Rob Donnelly, Rob Honsby, Rocky Smith, Rohan Sahai, Rohit Ramchan-  
 839 dani, Romain Huet, Rory Carmichael, Rowan Zellers, Roy Chen, Ruby Chen, Ruslan Nigmat-  
 840 ullin, Ryan Cheu, Saachi Jain, Sam Altman, Sam Schoenholz, Sam Toizer, Samuel Miserendino,  
 841 Sandhini Agarwal, Sara Culver, Scott Ethersmith, Scott Gray, Sean Grove, Sean Metzger, Shamez  
 842 Hermani, Shantanu Jain, Shengjia Zhao, Sherwin Wu, Shino Jomoto, Shirong Wu, Shuaiqi Xia,  
 843 Sonia Phene, Spencer Papay, Srinivas Narayanan, Steve Coffey, Steve Lee, Stewart Hall, Suchir  
 844 Balaji, Tal Broda, Tal Stramer, Tao Xu, Tarun Gogineni, Taya Christianson, Ted Sanders, Tejal  
 845 Patwardhan, Thomas Cunningham, Thomas Degry, Thomas Dimson, Thomas Raoux, Thomas  
 846 Shadwell, Tianhao Zheng, Todd Underwood, Todor Markov, Toki Sherbakov, Tom Rubin, Tom  
 847 Stasi, Tomer Kaftan, Tristan Heywood, Troy Peterson, Tyce Walters, Tyna Eloundou, Valerie Qi,  
 848 Veit Moeller, Vinnie Monaco, Vishal Kuo, Vlad Fomenko, Wayne Chang, Weiyi Zheng, Wenda  
 Zhou, Wesam Manassra, Will Sheu, Wojciech Zaremba, Yash Patil, Yilei Qian, Yongjik Kim,  
 Youlong Cheng, Yu Zhang, Yuchen He, Yuchen Zhang, Yujia Jin, Yunxing Dai, and Yury Malkov.  
 Gpt-4o system card, 2024. URL <https://arxiv.org/abs/2410.21276>.

849 Xingyu Ren, Alexandros Lattas, Baris Gecer, Jiankang Deng, Chao Ma, and Xiaokang Yang. Facial  
 850 geometric detail recovery via implicit representation. In *2023 IEEE 17th International Conference*  
 851 *on Automatic Face and Gesture Recognition (FG)*, 2023.

852 Jaechul Roh, Virat Shejwalkar, and Amir Houmansadr. Multilingual and multi-accent jailbreaking  
 853 of audio llms, 2025. URL <https://arxiv.org/abs/2504.01094>.

854 Angelika Romanou, Negar Foroutan, Anna Sotnikova, Zeming Chen, Sree Harsha Nelaturu, Shiv-  
 855 alika Singh, Rishabh Maheshwary, Micol Altomare, Mohamed A. Haggag, Snegha A, Alfonso  
 856 Amayuelas, Azril Hafizi Amirudin, Viraat Aryabumi, Danylo Boiko, Michael Chang, Jenny  
 857 Chim, Gal Cohen, Aditya Kumar Dalmia, Abraham Diress, Sharad Duwal, Daniil Dzenhaliou,  
 858 Daniel Fernando Erazo Florez, Fabian Farestam, Joseph Marvin Imperial, Shayekh Bin Islam,  
 859 Perttu Isotalo, Maral Jabbarishvili, Börje F. Karlsson, Eldar Khalilov, Christopher Klamm, Fa-  
 860 jri Koto, Dominik Krzemiński, Gabriel Adriano de Melo, Syrielle Montariol, Yiyang Nan, Joel  
 861 Niklaus, Jekaterina Novikova, Johan Samir Obando Ceron, Debjit Paul, Esther Ploeger, Jebish  
 862 Purbey, Swati Rajwal, Selvan Sunitha Ravi, Sara Rydell, Roshan Santhosh, Drishti Sharma, Mar-  
 863 jana Prifti Skenduli, Arshia Soltani Moakhar, Bardia Soltani Moakhar, Ran Tamir, Ayush Kumar

864 Tarun, Azmine Toushik Wasi, Thenuka Ovin Weerasinghe, Serhan Yilmaz, Mike Zhang, Imanol  
 865 Schlag, Marzieh Fadaee, Sara Hooker, and Antoine Bosselut. Include: Evaluating multilingual  
 866 language understanding with regional knowledge, 2024. URL <https://arxiv.org/abs/2411.19799>.  
 867

868 David Romero, Chenyang Lyu, Haryo Akbarianto Wibowo, Teresa Lynn, Injy Hamed, Aditya Nanda  
 869 Kishore, Aishik Mandal, Alina Dragonetti, Artem Abzaliev, Atnafu Lambebo Tonja, Bontu Fufa  
 870 Balcha, Chenxi Whitehouse, Christian Salamea, Dan John Velasco, David Ifeoluwa Adelani,  
 871 David Le Meur, Emilio Villa-Cueva, Fajri Koto, Fauzan Farooqui, Frederico Belcavello, Gan-  
 872 zorig Batnasan, Gisela Vallejo, Grainne Caulfield, Guido Ivetta, Haiyue Song, Henok Bi-  
 873 adglign Ademtew, Hernán Maina, Holy Lovenia, Israel Abebe Azime, Jan Christian Blaise  
 874 Cruz, Jay Gala, Jiahui Geng, Jesus-German Ortiz-Barajas, Jinheon Baek, Jocelyn Dunstan,  
 875 Laura Alonso Alemany, Kumaranage Ravindu Yaras Nagasinghe, Luciana Benotti, Luis Fer-  
 876 nando D’Haro, Marcelo Viridiano, Marcos Estecha-Garitagoitia, Maria Camila Buitrago Cabr-  
 877 era, Mario Rodríguez-Cantelar, Mélanie Jouitteau, Mihail Mihaylov, Mohamed Fazli Mohamed  
 878 Imam, Muhammad Farid Adilazuarda, Munkhjargal Gochoo, Munkh-Erdene Otgonbold, Naome  
 879 Etori, Olivier Niyomugisha, Paula Mónica Silva, Pranjal Chitale, Raj Dabre, Rendi Chevi,  
 880 Ruochen Zhang, Ryandit Diandaru, Samuel Cahyawijaya, Santiago Góngora, Soyeong Jeong,  
 881 Sukannya Purkayastha, Tatsuki Kuribayashi, Teresa Clifford, Thanmay Jayakumar, Tiago Tim-  
 882 poni Torrent, Toqueer Ehsan, Vladimir Araujo, Yova Kementchedjhieva, Zara Burzo, Zheng Wei  
 883 Lim, Zheng Xin Yong, Oana Ignat, Joan Nwatu, Rada Mihalcea, Thamar Solorio, and Alham Fikri  
 884 Aji. Cvqa: Culturally-diverse multilingual visual question answering benchmark, 2024. URL  
 885 <https://arxiv.org/abs/2406.05967>.  
 886

886 Ashwin Sankar, Srija Anand, Praveen Varadhan, Sherry Thomas, Mehak Singal, Shridhar Kumar,  
 887 Deovrat Mehendale, Aditi Krishana, Giri Raju, and Mitesh Khapra. Indicvoices-r: Unlocking  
 888 a massive multilingual multi-speaker speech corpus for scaling indian tts. *Advances in Neural  
 889 Information Processing Systems*, 37:68161–68182, 2024.

890 Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh,  
 891 and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localiza-  
 892 tion. *International Journal of Computer Vision*, 128(2):336–359, October 2019. ISSN 1573-1405.  
 893 doi: 10.1007/s11263-019-01228-7. URL <http://dx.doi.org/10.1007/s11263-019-01228-7>.  
 894

895 Yosephine Susanto, Adithya Venkatadri Hulagadri, Jann Railey Montalan, Jian Gang Ngui, Xian Bin  
 896 Yong, Weiqi Leong, Hamsawardhini Rengarajan, Peerat Limkonchotiwat, Yifan Mai, and  
 897 William Chandra Tjhi. Sea-helm: Southeast asian holistic evaluation of language models, 2025.  
 898 URL <https://arxiv.org/abs/2502.14301>.  
 899

900 Gemini Team. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long  
 901 context, and next generation agentic capabilities, 2025a. URL <https://arxiv.org/abs/2507.06261>.  
 902

903 Kimi Team, Angang Du, Bohong Yin, Bowei Xing, Bowen Qu, Bowen Wang, Cheng Chen, Chenlin  
 904 Zhang, Chenzhuang Du, Chu Wei, Congcong Wang, Dehao Zhang, Dikang Du, Dongliang Wang,  
 905 Enming Yuan, Enzhe Lu, Fang Li, Flood Sung, Guangda Wei, Guokun Lai, Han Zhu, Hao Ding,  
 906 Hao Hu, Hao Yang, Hao Zhang, Haoning Wu, Haotian Yao, Haoyu Lu, Heng Wang, Hongcheng  
 907 Gao, Huabin Zheng, Jiaming Li, Jianlin Su, Jianzhou Wang, Jiaqi Deng, Jiezhong Qiu, Jin Xie,  
 908 Jinhong Wang, Jingyuan Liu, Junjie Yan, Kun Ouyang, Liang Chen, Lin Sui, Longhui Yu, Meng-  
 909 fan Dong, Mengnan Dong, Nuo Xu, Pengyu Cheng, Qizheng Gu, Runjie Zhou, Shaowei Liu,  
 910 Sihan Cao, Tao Yu, Tianhui Song, Tongtong Bai, Wei Song, Weiran He, Weixiao Huang, Weixin  
 911 Xu, Xiaokun Yuan, Xingcheng Yao, Xingzhe Wu, Xinhao Li, Xinxing Zu, Xinyu Zhou, Xinyuan  
 912 Wang, Y. Charles, Yan Zhong, Yang Li, Yangyang Hu, Yanru Chen, Yejie Wang, Yibo Liu, Yibo  
 913 Miao, Yidao Qin, Yimin Chen, Yiping Bao, Yiqin Wang, Yongsheng Kang, Yuanxin Liu, Yuhao  
 914 Dong, Yulun Du, Yuxin Wu, Yuzhi Wang, Yuzi Yan, Zaida Zhou, Zhaowei Li, Zhejun Jiang,  
 915 Zheng Zhang, Zhilin Yang, Zhiqi Huang, Zihao Huang, Zijia Zhao, Ziwei Chen, and Zongyu Lin.  
 916 Kimi-vl technical report, 2025a. URL <https://arxiv.org/abs/2504.07491>.  
 917

M2 Team, Chengfeng Dou, Chong Liu, Fan Yang, Fei Li, Jiyuan Jia, Mingyang Chen, Qiang Ju,  
 918 Shuai Wang, Shunya Dang, Tianpeng Li, Xiangrong Zeng, and et al. Yijie Zhou. Baichuan-m2:

918 Scaling medical capability with large verifier system, 2025b. URL <https://arxiv.org/abs/2509.02208>.

919

920

921 Qwen Team. Qwen3 technical report, 2025b. URL <https://arxiv.org/abs/2505.09388>.

922

923 V Team, Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, Haomiao Tang, Jiale

924 Cheng, Ji Qi, Junhui Ji, Lihang Pan, Shuaiqi Duan, Weihan Wang, Yan Wang, Yean Cheng,

925 Zehai He, Zhe Su, Zhen Yang, Ziyang Pan, Aohan Zeng, Baoxu Wang, Bin Chen, Boyan Shi,

926 Changyu Pang, Chenhui Zhang, Da Yin, Fan Yang, Guoqing Chen, Jiazheng Xu, Jiale Zhu, Jiali

927 Chen, Jing Chen, Jinhao Chen, Jinghao Lin, Jinjiang Wang, Junjie Chen, Leqi Lei, Letian Gong,

928 Leyi Pan, Mingdao Liu, Mingde Xu, Mingzhi Zhang, Qinkai Zheng, Sheng Yang, Shi Zhong,

929 Shiyu Huang, Shuyuan Zhao, Siyan Xue, Shangqin Tu, Shengbiao Meng, Tianshu Zhang, Tianwei

930 Luo, Tianxiang Hao, Tianyu Tong, Wenkai Li, Wei Jia, Xiao Liu, Xiaohan Zhang, Xin Lyu,

931 Xinyue Fan, Xuancheng Huang, Yanling Wang, Yadong Xue, Yanfeng Wang, Yanzi Wang, Yifan

932 An, Yifan Du, Yiming Shi, Yiheng Huang, Yilin Niu, Yuan Wang, Yuanchang Yue, Yuchen Li,

933 Yutao Zhang, Yuting Wang, Yu Wang, Yuxuan Zhang, Zhao Xue, Zhenyu Hou, Zhengxiao Du,

934 Zihan Wang, Peng Zhang, Debing Liu, Bin Xu, Juanzi Li, Minlie Huang, Yuxiao Dong, and Jie

935 Tang. Glm-4.5v and glm-4.1v-thinking: Towards versatile multimodal reasoning with scalable

936 reinforcement learning, 2025c. URL <https://arxiv.org/abs/2507.01006>.

937

938 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée

939 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Ar-

940 mand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation

941 language models, 2023. URL <https://arxiv.org/abs/2302.13971>.

942

943 Ashmal Vayani, Dinura Dissanayake, Hasindri Watawana, Noor Ahsan, Nevasini Sasikumar, Omkar

944 Thawakar, Henok Biadgign Ademtew, Yahya Hmaiti, Amandeep Kumar, Kartik Kuckreja,

945 Mykola Maslych, Wafa Al Ghallabi, Mihail Mihaylov, Chao Qin, Abdelrahman M Shaker, Mike

946 Zhang, Mahardika Krisna Ihsani, Amiel Esplana, Monil Gokani, Shachar Mirkin, Harsh Singh,

947 Ashay Srivastava, Endre Hamerlik, Fathinah Asma Izzati, Fadillah Adamsyah Maani, Sebastian

948 Cavada, Jenny Chim, Rohit Gupta, Sanjay Manjunath, Kamila Zhumakhanova, Feno Herini-

949 aina Rabevohitra, Azril Amirudin, Muhammad Ridzuan, Daniya Kareem, Ketan More, Kun-

950 yang Li, Pramesh Shakya, Muhammad Saad, Amirkouya Ghasemaghaei, Amirkbek Djanibekov,

951 Dilshod Azizov, Branislava Jankovic, Naman Bhatia, Alvaro Cabrera, Johan Obando-Ceron,

952 Olympiah Otieno, Fabian Farestam, Muztoba Rabbani, Sanoojan Baliah, Santosh Sanjeev, Ab-

953 duragim Shtanchaev, Maheen Fatima, Thao Nguyen, Amrin Kareem, Toluwani Aremu, Nathan

954 Xavier, Amit Bhatkal, Hawau Toyin, Aman Chadha, Hisham Cholakkal, Rao Muhammad Anwer,

955 Michael Felsberg, Jorma Laaksonen, Thamar Solorio, Monojit Choudhury, Ivan Laptev, Mubarak

956 Shah, Salman Khan, and Fahad Khan. All languages matter: Evaluating lmms on culturally

957 diverse 100 languages, 2025. URL <https://arxiv.org/abs/2411.16508>.

958

959 Bin Wang, Zhengyuan Liu, Xin Huang, Fangkai Jiao, Yang Ding, AiTi Aw, and Nancy F. Chen.

960 Saeval for multilingual foundation models: From cross-lingual alignment to cultural reasoning,

961 2024. URL <https://arxiv.org/abs/2309.04766>.

962

963 Xinyi Wang, Antonis Antoniades, Yanai Elazar, Alfonso Amayuelas, Alon Albalak, Kexun Zhang,

964 and William Yang Wang. Generalization v.s. memorization: Tracing language models' capabili-

965 ties back to pretraining data, 2025. URL <https://arxiv.org/abs/2407.14985>.

966

967 Zhiyu Wu, Xiaokang Chen, Zizheng Pan, Xingchao Liu, Wen Liu, Damai Dai, Huazuo Gao, Yiyang

968 Ma, Chengyue Wu, Bingxuan Wang, Zhenda Xie, Yu Wu, Kai Hu, Jiawei Wang, Yaofeng Sun,

969 Yukun Li, Yishi Piao, Kang Guan, Aixin Liu, Xin Xie, Yuxiang You, Kai Dong, Xingkai Yu,

970 Haowei Zhang, Liang Zhao, Yisong Wang, and Chong Ruan. Deepseek-vl2: Mixture-of-experts

971 vision-language models for advanced multimodal understanding, 2024. URL <https://arxiv.org/abs/2412.10302>.

972

973 Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang

974 Fan, Kai Dang, Bin Zhang, Xiong Wang, Yunfei Chu, and Junyang Lin. Qwen2.5-omni technical

975 report, 2025. URL <https://arxiv.org/abs/2503.20215>.

976

977 Weihao Xuan, Rui Yang, Heli Qi, Qingcheng Zeng, Yunze Xiao, Aosong Feng, Dairui Liu, Yun

978 Xing, Junjue Wang, Fan Gao, Jinghui Lu, Yuang Jiang, Huitao Li, Xin Li, Kunyu Yu, Ruihai

972 Dong, Shangding Gu, Yuekang Li, Xiaofei Xie, Felix Juefei-Xu, Foutse Khomh, Osamu Yoshie,  
 973 Qingyu Chen, Douglas Teodoro, Nan Liu, Randy Goebel, Lei Ma, Edison Marrese-Taylor, Shijian  
 974 Lu, Yusuke Iwasawa, Yutaka Matsuo, and Irene Li. MMLU-ProX: A multilingual benchmark for  
 975 advanced large language model evaluation. In Christos Christodoulopoulos, Tanmoy Chakraborty,  
 976 Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods*  
 977 *in Natural Language Processing*, pp. 1513–1532, Suzhou, China, November 2025. Association  
 978 for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.  
 979 79. URL <https://aclanthology.org/2025.emnlp-main.79/>.

980 Zheng-Xin Yong, Beyza Ermis, Marzieh Fadaee, Stephen H. Bach, and Julia Kreutzer. The state of  
 981 multilingual llm safety research: From measuring the language gap to mitigating it, 2025. URL  
 982 <https://arxiv.org/abs/2505.24119>.

983 Zhi Zhang, Srishti Yadav, Fengze Han, and Ekaterina Shutova. Cross-modal information flow in  
 984 multimodal large language models, 2025. URL <https://arxiv.org/abs/2411.18620>.

985 Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. Large language models are  
 986 not robust multiple choice selectors. *arXiv preprint arXiv:2309.03882*, 2023.

987 Weihua Zheng, Xin Huang, Zhengyuan Liu, Tarun Kumar Vangani, Bowei Zou, Xiyan Tao, Yuhao  
 988 Wu, Ai Ti Aw, Nancy F. Chen, and Roy Ka-Wei Lee. Adamcot: Rethinking cross-lingual factual  
 989 reasoning through adaptive multilingual chain-of-thought, 2025a. URL <https://arxiv.org/abs/2501.16154>.

990 Weihua Zheng, Roy Ka-Wei Lee, Zhengyuan Liu, Kui Wu, AiTi Aw, and Bowei Zou. Ccl-xcot: An  
 991 efficient cross-lingual knowledge transfer method for mitigating hallucination generation, 2025b.  
 992 URL <https://arxiv.org/abs/2507.14239>.

## A APPENDIX

### A.1 ANNOTATOR DEMOGRAPHIC

Our annotation team comprises members from eight different countries. All team members are native speakers of their local languages and proficient in English, with professional backgrounds in natural language processing or speech processing. Each annotator has lived in the respective country for more than ten years. Table 4 presents the annotators’ details and professional backgrounds; to protect privacy, we replace personal names with numeric identifiers within each country. After each team completes the first round of annotations, in-country linguistic experts conduct a data review. Table 5 lists the language experts’ information; likewise, we anonymize personal names.

### A.2 DATA STATISTICS

Under the MMA-ASIA framework, the dataset covers 8 countries and 10 languages, with each country’s split presented in both English and its local language, totaling 27,000 questions. Table 6 lists the countries included in our dataset and their corresponding local language(s). Over 79% of all items are multi-step cultural reasoning questions. We define a multi-step cultural reasoning item as one whose solution requires sequential derivation and/or synthesis from at least two independent knowledge components, rather than mere recall or paraphrase of a single cultural fact. Table 7 presents an example multi-step question with its analysis. The proportion of multi-step items by country and modality is shown in Figure 8. The dataset spans nine categories—*Daily Life/Culture, Food/Cuisine, Transportation, Buildings, History, Geographical Location & Climate, Education, Fashion/Clothing, and Language/Ethnicity*—with per-country category distributions summarized in Figure 7.

The following outlines the potential contents included in each category.

- **Daily Life/Culture:** Covers everyday etiquette, customs, and values—greeting practices, daily routines and social decorum, family and community interactions, as well as festivals and folk practices. It may also include consumption and leisure preferences, common life scenarios, and behavioral norms.
- **Food/Cuisine:** Regional cuisines, representative dishes and ingredients, cooking techniques, dietary taboos and table manners, utensils, and dining settings. Also includes festive foods, street snacks, and regional taste differences.

1026

1027

1028

1029

1030

1031

Table 4: Data Annotator Demographics and Skills

Country	ID	Gender	Age	Education	English Prof.	Local Lang. Prof.	Professional Background
China	1	Male	28	Ph.D.	Proficient	Proficient	NLP
	2	Male	25	Master Deg.	Proficient	Proficient	NLP
	3	Male	30	Ph.D.	Proficient	Proficient	NLP
	4	Male	25	Master Deg.	Familiar	Fluent	NLP
	5	Male	24	Master Deg.	Familiar	Proficient	NLP
Japan	1	Male	23	Master Deg.	Proficient	Proficient	SP
	2	Male	35	Ph.D.	Fluent	Proficient	SP
Mongolia	1	Male	26	Ph.D.	Proficient	Proficient	SP
	2	Male	25	Master Deg.	Proficient	Proficient	SP
Korea	1	Male	31	Ph.D.	Proficient	Proficient	NLP
	2	Male	26	Master Deg.	Proficient	Proficient	NLP
	3	Female	25	Master Deg.	Proficient	Proficient	NLP
India	1	Male	25	Bachelor's Deg.	Proficient	Proficient	NLP
	2	Male	31	Ph.D.	Proficient	Proficient	NLP
Vietnam	1	Male	21	Bachelor's Deg.	Proficient	Proficient	SP
	2	Male	20	Bachelor's Deg.	Proficient	Proficient	SP
	3	Male	21	Bachelor's Deg.	Proficient	Proficient	SP
	4	Male	21	Bachelor's Deg.	Proficient	Proficient	SP
	5	Male	29	Master Deg.	Proficient	Proficient	SP
Indonesia	1	Female	21	Bachelor's Deg.	Proficient	Proficient	NLP
	2	Female	26	Ph.D.	Proficient	Proficient	NLP
Singapore	1	Female	26	Ph.D.	Proficient	Proficient	NLP
	2	Female	18	Bachelor's Deg.	Proficient	Proficient	NLP
	3	Male	18	Bachelor's Deg.	Proficient	Proficient	NLP
	4	Male	21	Bachelor's Deg.	Proficient	Proficient	NLP
	5	Female	28	Ph.D.	Proficient	Proficient	SP

Note: NLP: Natural Language Processing; SP: Speech Processing.

1059

1060

1061

1062

1063

1064

Table 5: Data Reviewer Demographics and Skills

Country	ID	Gender	Age	Education	English Prof.	Local Lang. Prof.
China	1	Female	30	Master Deg.	Proficient	Proficient
Japan	1	Female	32	Master Deg.	Proficient	Proficient
Mongolia	1	Male	25	Master Deg.	Proficient	Proficient
Korea	1	Male	54	Bachelor's Deg.	Proficient	Proficient
India	1	Male	28	Master Deg.	Proficient	Proficient
Vietnam	1	Male	27	Master Deg.	Proficient	Proficient
Indonesia	1	Female	28	Bachelor's Deg.	Proficient	Proficient
Singapore	1	Female	26	Bachelor's Deg.	Proficient	Proficient
	2	Female	22	Bachelor's Deg.	Proficient	Proficient
	3	Male	30	Ph.D.	Proficient	Proficient
	4	Male	26	Ph.D.	Proficient	Proficient

1079

1080

Table 6: Local Languages by Country

1081

1082

1083

1084

1085

1086

1087

1088

1089

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

1100

1101

1102

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

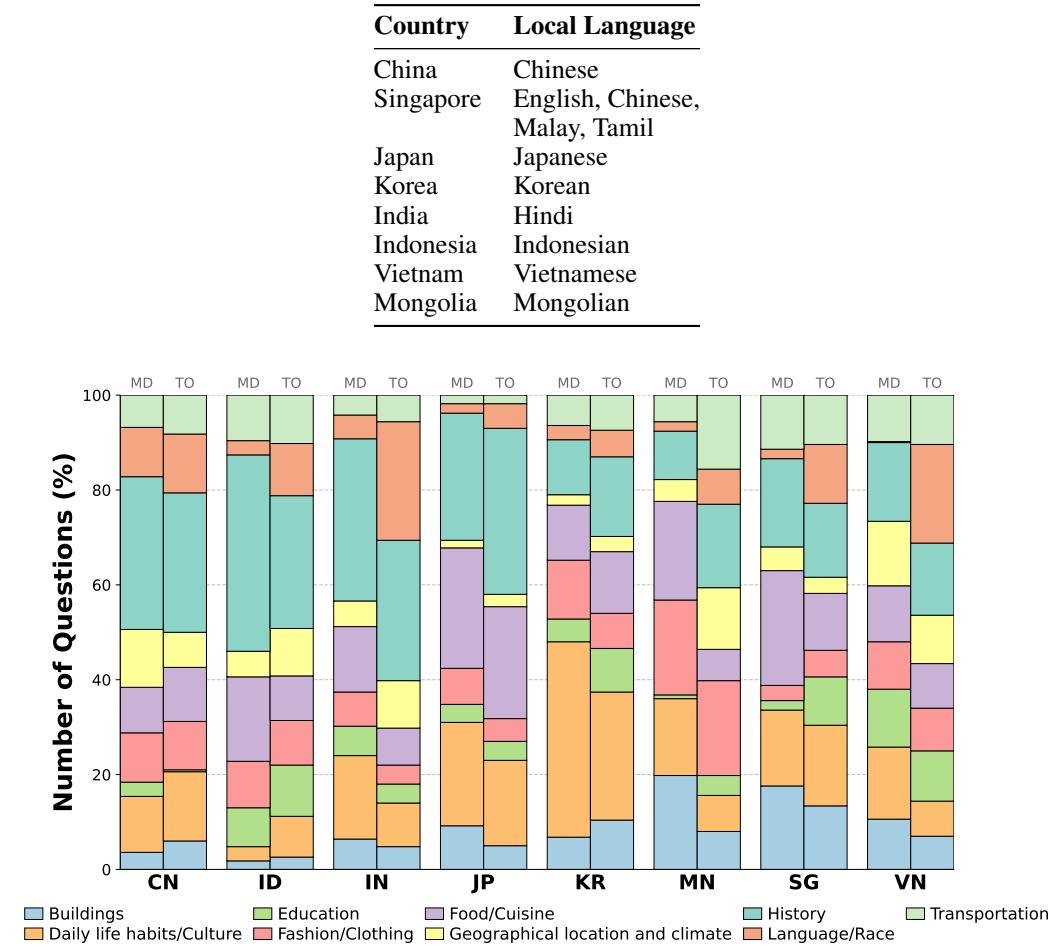


Figure 7: Distribution of question categories across countries and modalities

- **Transportation:** Transportation across historical periods and the evolution of vehicles, along with regional differences in modes of transport. May also cover landmark transit systems and commuting culture.
- **Buildings:** Traditional and modern architectural styles; religious and public buildings; housing forms and materials; city skylines and the preservation of historic districts. Can also address symbolic meanings in architecture and region-specific structural features.
- **History:** Major historical periods and events, notable figures and heritage sites, and how historical memory shapes contemporary society and culture. May also include colonial/independence histories and cultural change driven by migration and war.
- **Geographical Location & Climate:** Landforms and terrain, climate zones and seasonal variation, natural resources, and ecosystems. Extends to lifestyle, clothing, and dietary adaptations shaped by geography and climate.
- **Education:** Renowned national works of literature, art, or music, as well as the structure of the education system and pathways to advancement, including stories surrounding prestigious institutions.
- **Fashion/Clothing:** Traditional attire and its ceremonial contexts; modern dress styles and aesthetic trends; accessories and color preferences; occupational/school uniforms and seasonal clothing. May also discuss cultural symbolism embedded in garments.

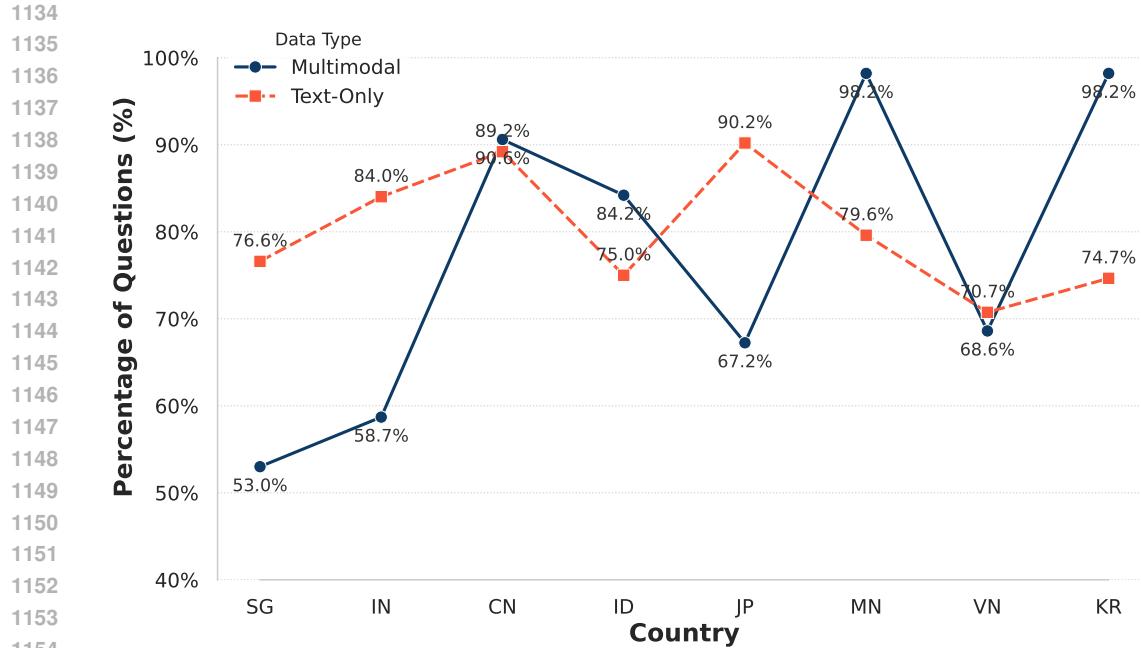


Figure 8: Multi-step reasoning question proportions across countries

- **Language/Ethnicity:** Code-switching (where applicable); official and commonly used languages; dialects and accent features; writing systems and naming conventions; multi-ethnic compositions and cultural practices. Also includes politeness strategies in language and norms of cross-group communication.

Table 7: An example of multi-step reasoning question

Question	An ancient tower became famous due to a poem by the Tang Dynasty poet Cui Hao. In which period was it proposed to be constructed with iron materials? (A) 16th year of Guangxu, Qing Dynasty (B) 1st year of Jiaqing, Qing Dynasty (C) 7th year of Tongzhi, Qing Dynasty (D) 8th year of the Republic of China
Reasoning Decomposition	<p><b>Step 1: Identify the tower.</b> The poem is Cui Hao's "Yellow Crane Tower," so the tower is <i>Yellow Crane Tower</i>.</p> <p><b>Step 2: Recall historical events.</b> In the 10th year of Guangxu (1884) the tower was destroyed by fire. In the 16th year of Guangxu (1890), Zhang Zhidong (Governor-General of Hubei and Hunan) first proposed rebuilding the tower using iron materials.</p> <p><b>Step 3: Match with the options.</b> A (1890): <b>Correct</b>, matches the historical fact. B (1796): Incorrect, too early. C (1868): Incorrect, before the fire. D (1919): Incorrect, after the Qing Dynasty and not the first proposal.</p> <p><b>Final Answer: A. 16th year of Guangxu, Qing Dynasty.</b></p>

Figure 9 and Figure 10 show some samples of our dataset.

### A.3 ANNOTATION GUIDELINE

The content of the guideline distributed to annotators is shown in Figure 11. To minimize heuristic cues arising from non-cultural knowledge, we add a consistency constraint on distractors: they must belong to the same category as the correct option and closely resemble it in observable attributes and semantic representation. We also encourage each team to uncover cultural elements unique to their own country, rather than focusing only on widely known aspects. For the Language category, if code-switching is prevalent in the annotators' country, we strongly encourage including such language-assessment examples in the Text-Only portion of the dataset. All content involving racism or hate speech is prohibited from inclusion in our dataset.

	Knowledge Point	Source Question	English Question
1188			
1189	The Mogao Caves, also known as the Thousand Buddha Grottoes, feature the "Transformation Tableau of the Medicine Buddha Sutra" in Cave 220, which depicts scenes from the High Tang period, and the "Mural of Zhang Yichao's Army in Procession" in Cave 156, which showcases scenes from the Late Tang period.	某石窟有千佛洞之称, 其在唐代有以下哪些壁画作品? A. 《药师经变画》 B. 《炽盛光佛图》 C. 《张议潮统军出行图》 D. 《五台山图》	Which mural paintings from the Tang Dynasty are found in a certain grotto known as the Thousand Buddha Caves?", A. Transformation Tableau of the Medicine Buddha Sutra B. Painting of Tejaprabhā Buddha C. Mural of Zhang Yichao's Army in Procession D. Mount Wutai
1190			
1191			
1192			
1193			
1194			
1195			
1196	Gyeongbokgung Palace was built with Bukaksan as its main mountain, strategically placing its buildings on a spacious site with Gwanghwamun as its main gate, opening onto a wide boulevard that formed the center of Hanyang, the capital of the Joseon Dynasty. Gyeongbokgung was destroyed during the Imjin War in 1592, the 25th year of King Seonjo's reign.	백악산을 주산으로 넓은 지형에 건물을 배치하고 정문 앞으로 넓은 육조거리가 펼쳐진 한양의 중심이 1592년 소실된 계기는 무엇인가? A. 6.25 전쟁 B. 화재 발생 C. 임진왜란 D. 을미사변	What was the cause of the destruction in 1592 of the central area of Hanyang, where buildings were laid out across a wide terrain with Bugaksan as the main mountain and a broad Yukjo Street stretching out in front of the main gate? A. Korean War B. Fire outbreak C. Imjin War D. Eulmi Incident
1197			
1198			
1199			
1200			
1201			
1202			
1203			
	Knowledge Point	Source Question	English Question
1204	Char Kway Teow: This stir-fried flat noodle dish often contains prawns. It is typically served with lime. Fried Hokkien Mee: a stir-fried dish of yellow noodles and rice vermicelli cooked in a rich stock made from pork bones and prawn heads. It is almost always served with a wedge of calamansi lime on the side.	பின்வரும் பொதுவான சிங்கப்பூர் நூட்டுலஸ் உணவுகளில் எது அடிக்கடி இராவல் மற்றும் எலுமிச்சை சேர்க்கப்பட்டிருக்கும்? A. சார்க்டோ டியோ B. லெக்சா C. வறுத்த லூராக்கியன் மீ D. கிரே:பிள் நூட்டுலஸ்	Which of the following common Singaporean noodle dishes often include prawns and lime? A. Char kway teow B. Leksa C. Fried Hokkien Mee D. Crayfish Noodles
1205			
1206			
1207			
1208			
1209			
1210			
1211	The word "Gamcha" is Bengali/Assamese word which comes from two very simple and commonly used Bengali/Assamese words, (ଗା) ga which means "Body", and (ମୁଖ) mucha which means "wipe". It is often used as traditional full sized handkerchief. And it looks like a towel and acts as one too in case of necessity, it acts as a mask which means dust and pollution stays away from you, it also acts as sun protection so that you don't get much tanned or skin burn.	ரவி தெகிண ஏதியா மே உணக்டிர்஧ீய ஜலவாயு கே லீ அனக்லித் பார்ப்ரிக் பரிசான கா அடியன் கர ரஹ ஹ. உந்ஹோன் பார்சம் பங்கல் மே யாமீன் புரூஷ கே ஏக ஸமூஹ கோ பஸினா பாஷனே, ஧ை மே சிர டகனே ஓர் யான் தக கி அஸ்஥ாயீ மாஸ்க் கே ரப மே ஹ்லக் ஸ்தி கப்பட கா உபயோக கரதே ஹே தேயா வஹ ஸ்ம்பவத: கிஸ பார்ப்ரிக் வஸ்து கா உல்லைக் கர ரஹ ஹ? A. குர்தா B. ஧ாதி C. மாஸ்க D. ஶர்வானி	Ravi is studying traditional garments in South Asia adapted for tropical climates. He notices a group of rural men in West Bengal using a light cotton cloth for wiping sweat, covering their heads in the sun, and even as a makeshift mask. Which traditional item is he likely referring to? A. Kurta B. Dhoti C. Gamcha D. Sherwani
1212			
1213			
1214			
1215			
1216			
1217			
1218			
1219			
1220			
1221			
1222			
1223			
1224			
1225			

Figure 9: Text-Only Question Examples

#### A.4 DETAILS OF THE TTS TOOL AND PROCEDURE FOR BUILDING THE SPEECH DATA

Before generating the English audio, we standardized the input text through normalization of numbers and symbols, handling of abbreviations and special terms, and sentence segmentation.

To ensure high-quality speech synthesis, we employed CosyVoice(Du et al., 2025) for English audio generation. This tool supports voice cloning from sampled speakers, producing speech that preserves the timbre and accent of the reference voice. We collected representative recordings from native speakers across eight countries to capture diverse accents for speech synthesis. For standard English, we adopted CosyVoice’s built-in default English voice (English female voice). For non-English languages, CosyVoice was also used to generate Chinese, Japanese, and Korean audio. In addition, we employed in-house high-quality TTS systems built by different speech processing teams for Vietnamese, Tamil, Mongolian, and Malay, while the Coqui-ai TTS toolkit (Coqui.ai, 2025) was used for Indonesian and Hindi.

Each generated audio sample was individually verified. When errors occurred—such as inappropriate pauses, missing segments, or mispronunciations—we first adjusted the input text and resynthesized the audio, as TTS systems are often highly sensitive to textual variations. If repeated corrections still failed, we resorted to manual re-recording. Unlike in other language tasks, our requirement here was not fluency or naturalness, but rather clear articulation of the questions and answer options.

1242	Knowledge Point	Image	Question
1243	Multicolor woodblock printing enabled full-color ukiyo-e and mass popularity; in Edo, portraitists like Sharaku and Toyokuni thrived, and yakusha-e and bijin-ga became bestsellers. And most yakusha-e are a subset of kabuki-e.		Question: この画像のような、絵をなんというでしょう？ A.美人絵 B.墨絵 C.歌舞伎絵 D.役者絵  Rephrased Question: 浮世絵で歌舞伎役者や当時の人気俳優の姿を描いた絵を何というでしょう？
1244	<b>Tugu Proklamasi</b> (the Proclamation Monument) is a monument commemorating the Proclamation of Independence of the Republic of Indonesia. It stands within the Proclamation Park complex on Jalan Proklamasi in Central Jakarta. The house where the proclamation was read was demolished in the 1960s.		Di lokasi pada gambar dulu ada sebuah rumah; kapan rumah itu dirobohkan? A. Tepat setelah hari kemerdekaan Indonesia B. Tahun 1960-an C. Tahun 1970-an D. Awal tahun 2000  Rephrased Question: Kapan rumah yang terletak di tugu proklamasi tersebut dibongkar?

Figure 10: VQA/Rephrase VQA Question Examples

All audio files were standardized in WAV format with a 16 kHz sampling rate. And our speech generation uses two input types.

1. Only the question stem from the *Rephrase VQA (text-only)* item;
2. The entire *Rephrase VQA (text-only)* item, including the question and its answer options.

Table 8 summarizes the models used for speech generation across different languages.

Table 8: Models used for speech data generation

Language	Model
English with accent, Korean, Japanese, Chinese	CosyVoice2-0.5B (Du et al., 2025)
English without accent	CosyVoice-300M (Du et al., 2024)
Vietnamese, Tamil, Mongolian, Malay	Internally developed TTS systems
Indonesian	Indonesian TTS (Kim et al., 2021)
Hindi	AI4Bharat Indic-TTS (Sankar et al., 2024)

## A.5 MODEL SELECTION

**Open-source multilingual text-only LLMs.** We evaluate Qwen3-30B-A3B-Thinking-2507 Team (2025b), Baichuan-M2-32B Team et al. (2025b), GLM-4-9B-ChatGLM et al. (2024), and Ministrals-8B-Instruct Jiang et al. (2023) for their multilingual capabilities.

**Open-source multimodal LLMs.** We evaluate nine vision–language (and omni) models on image and text: Qwen2.5-VL-32B-Instruct (Bai et al., 2025), Llama-3.2-11B-Vision-Instruct (Touvron et al., 2023), Kimi-VL-A3B-Instruct (Team et al., 2025a), DeepSeek-VL-Small (Wu et al., 2024), GLM-4.1V-9B-Thinking (Team et al., 2025c), InternVL-Chat-V1-5 (Chen et al., 2024b), and Qwen2.5-Omni-7B (Xu et al., 2025). Kimi-VL-A3B-Instruct, DeepSeek-VL-Small, and InternVL-Chat-V1-5 are evaluated in English only; Qwen2.5-VL-32B-Instruct, Llama-3.2-11B-Vision-Instruct, and GLM-4.1V-9B-Thinking are evaluated in multiple languages. Qwen2.5-Omni-7B is evaluated across image, text, and speech in a multilingual setting. [We summarize the information of the LLMs evaluated in our experiments in Table 9.](#)

## A.6 PROMPT TEMPLATES AND EXPERIMENTAL SETTINGS

**Prompt templates for evaluation tasks across modalities.** Table 11 presents the English prompts used in our evaluations across different modalities. When the query switches to another language, the

1296  
1297  
1298  
1299  
1300  
1301  
1302  
1303

1304  
1305  
1306  
1307  
1308  
1309  
1310  
1311  
1312  
1313  
1314  
1315  
1316  
1317  
1318  
1319  
1320  
1321  
1322  
1323  
1324  
1325  
1326  
1327  
1328  
1329  
1330  
1331  
1332  
1333  
1334  
1335  
1336  
1337  
1338  
1339  
1340  
1341  
1342  
1343  
1344  
1345  
1346  
1347  
1348  
1349

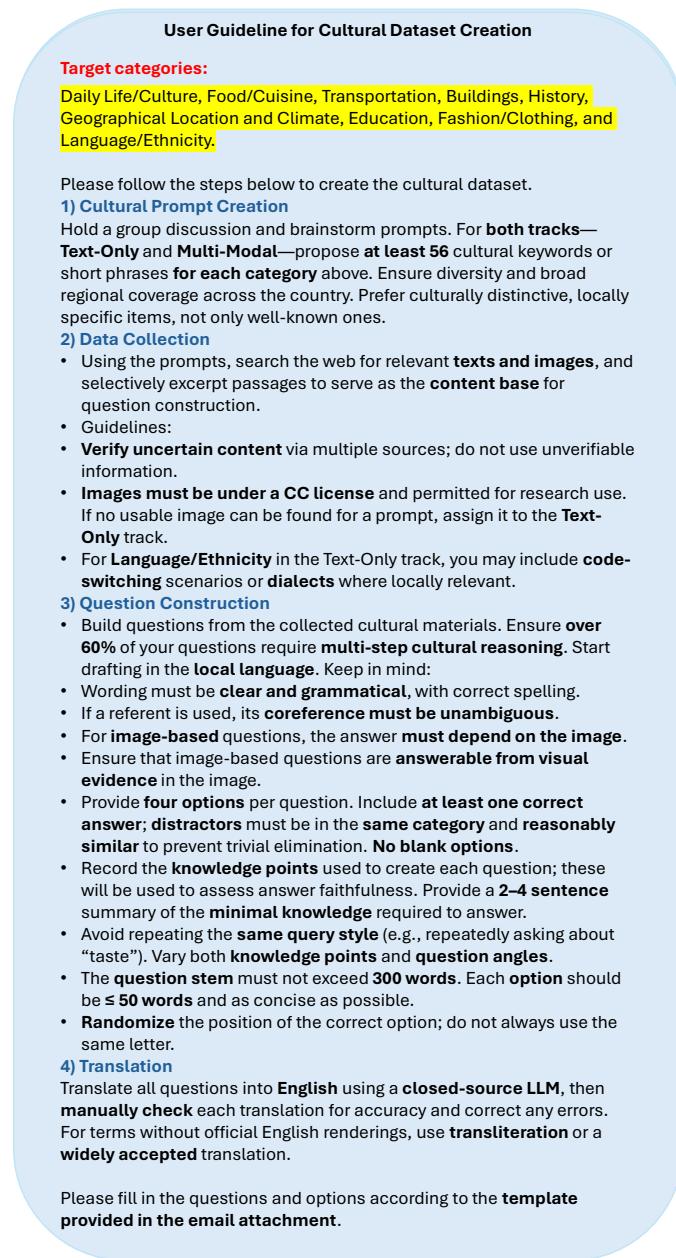


Figure 11: Annotator Guideline

corresponding translated version of the prompt will be used to ensure input-language consistency. Table 12 presents the prompts used to invoke closed-source model APIs for translation, answer-consistency evaluation, and answer extraction.

**Experiments setting.** All evaluations in this work are conducted in the zero-shot setting, using single-turn inference for each model on an NVIDIA H100 80G. For image inputs that exceed a model’s maximum allowable resolution, we proportionally downscale the image until it is under 1megapixel before testing. We decode with greedy search (no sampling; `do_sample=False`, `num_beams=1`), so `temperature`/`top-p`/`top-k` are not used; the maximum output length is set to 2048 tokens to ensure reproducibility. GPT-4o and Gemini 2.5 Flash are accessed via Open-

Table 9: Overview of selected LLMs

Model	Type	Language Coverage	Modalities
<b>Qwen2.5-VL-32B-Instruct</b>	Vision-Language (VL)	Multilingual	Image + Text
<b>Llama-3.2-11B-Vision-Instruct</b>	Vision-Language (VL)	Multilingual	Image + Text
<b>Kimi-VL-A3B-Instruct</b>	Vision-Language (VL)	English only	Image + Text
<b>DeepSeek-VL2-Small</b>	Vision-Language (VL)	English only	Image + Text
<b>GLM-4.1V-9B-Thinking</b>	Vision-Language (VL)	Multilingual	Image + Text
<b>InternVL-Chat-V1-5</b>	Vision-Language (VL)	English only	Image + Text
<b>Qwen2.5-Omni-7B</b>	Omni-modal	Multilingual	Image + Text + Speech
<b>Qwen3-30B-A3B-Thinking-2507</b>	Text-only LLM	Multilingual	Text
<b>Baichuan-M2-32B</b>	Text-only LLM	Multilingual	Text
<b>GLM-4-9B-Chat</b>	Text-only LLM	Multilingual	Text
<b>Minstral 8B-Instruct</b>	Text-only LLM	Multilingual	Text
<b>GPT-4o</b>	Vision-Language (VL)	Multilingual	Image + Text
<b>Claude-Sonnet-4</b>	Vision-Language (VL)	Multilingual	Image + Text
<b>Gemini 2.5 Flash</b>	Omni-modal	Multilingual	Image + Text + Speech

Table 10: Number of Incorrect Cases Across Different Models and Languages.

Model	Number of Incorrect Cases		
	English	Chinese	Indonesian
GLM-4-9B-chat	337	320	312
Mistral-8B-Instruct	370	384	341
Claude Sonnet 4	232	231	222
Qwen3-30B-A3B	271	265	276

Router API platform, and Claude via the Anthropic API. For all closed-source models, we set `temperature=0` to minimize randomness and improve reproducibility. Our speech inputs are no longer than 30 seconds and sampled at 16 kHz—well below Gemini 2.5 Flash’s maximum speech-input duration and Qwen2.5-Omni-7B’s maximum input token limit. So we do not perform any input-length processing.

#### A.7 PERFORMANCE VARIATION ACROSS DATA CATEGORIES

Table 11: English prompts for different modals.

Modal	Prompt Template
Text-Only/Rephrase	Please answer the following culture-related question.\n{question}\n{options}\nThis is a multiple-choice question. Please first return all possible option letters, then explain your choice in English.
VQA	Based on the image, please answer the following culture-related question.\n{question}\n{options}\nThis is a multiple-choice question. Please first return all possible option letters, then explain your choice in English.
Speech question & text options	This is a culture-related question.\nBased on the question mentioned in this audio, please choose the correct answers from the following provided options.\n{options}\nThis is a multiple-choice question. Please first return all possible option letters, then explain your choice in English.
Speech question & options	This is a culture-related question. Based on the question and options mentioned in this audio, please choose the correct options. This is a multiple-choice question. Please first return all possible option letters, then explain your choice in English.

1404  
1405 Table 12: Prompts Used with Closed-Source APIs for Translation, Answer Extraction, and An-  
1406 swer-Reference Consistency Checking  
1407  
1408

1409 <b>Task</b>	1410 <b>Prompt</b>
1411 Translation	1412 Translate the following sentence into English. \n{Input sentence}\nThis is 1413 a multiple-choice question in the cultural domain of {Country}. Pay partic- 1414 ular attention to the terms in the input and use their official translations; if 1415 no official translation exists, you may use transliteration. Ensure accuracy, 1416 faithfulness, and fluency. Return only the translation; do not include any ad- 1417 ditional hints or analysis.
1418 LLM-as-Judge	1419 You are an expert evaluator. Your task is to determine if the 'Model An- 1420 swer' correctly and completely incorporates the information from the 1421 'Knowledge Point'. Knowledge Point:\n{knowledge_point}\nModel An- 1422 swer\n{model_answer}\n\n. If the Model Answer correctly and completely 1423 includes the information from the Knowledge Point, or if the model's re- 1424 sponse aligns with the Knowledge Point's content and viewpoint. And the 1425 Model Answer doesn't contain any factual error, Answer 'yes'. Otherwise 1426 Answer 'no'. Please only return 'yes' or 'no'.
1427 Answer Extraction	1428 This is a model's response to a multiple-choice question. First, understand 1429 the text, then extract the model's chosen options, returning only the option 1430 letters (e.g., A, B, C). Do not include the option content. Output the result in 1431 a format like [A, B]. If the response does not contain any final choice, return 1432 NA.

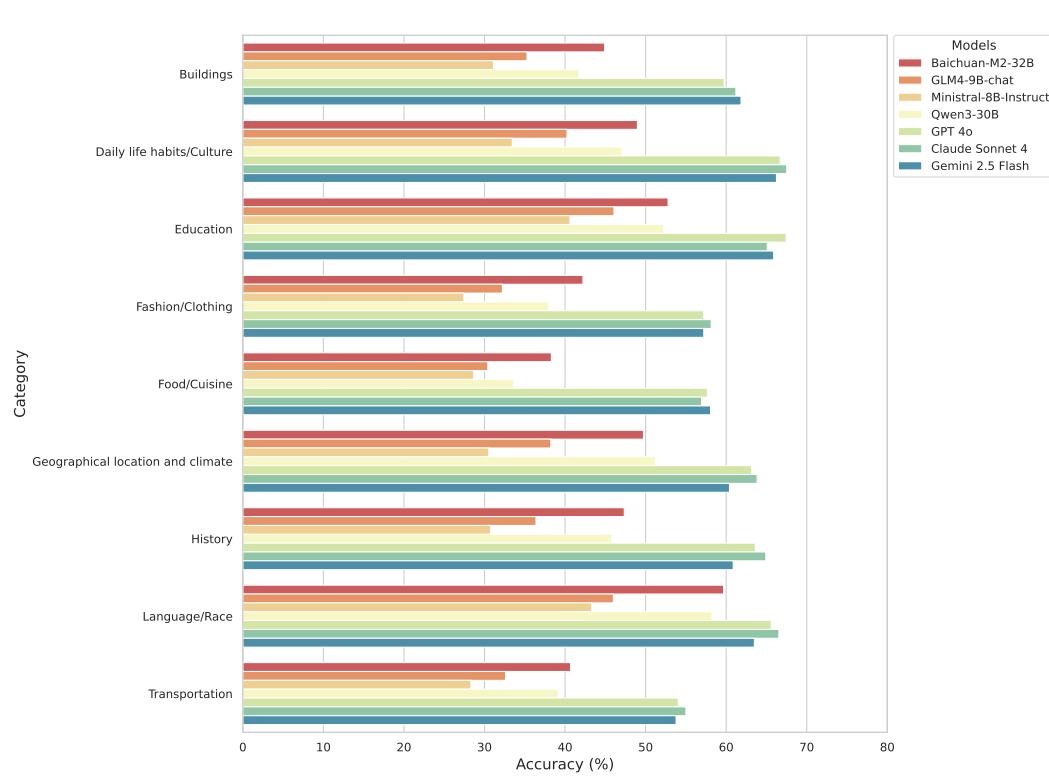
1429 To investigate how different cultural categories affect model performance, we evaluate models by  
1430 category under each modality, taking the average accuracy across all languages within a modality as  
1431 the category's accuracy. As shown in Figures 12, 13, and 14, under the Text-Only and VQA modalities,  
1432 models consistently struggle more with the "Fashion/Clothing" and "Transportation" categories  
1433 compared to others, while performing strongly on "Daily life habits/Culture" and "Education." In  
1434 the speech modality, models show slightly weaker performance on "Food/Cuisine" and "History,"  
1435 but still maintain strong results on "Daily life habits/Culture" and "Education".

1436 These patterns likely reflect both data and task asymmetries across categories. Daily life  
1437 habits/culture and education rely more on high-frequency, broadly documented facts and relatively  
1438 coarse-grained reasoning, making them easier for models trained on abundant, well-aligned multi-  
1439 lingual data. By contrast, fashion/clothing and transportation encode highly time-sensitive and  
1440 region-specific concepts (e.g., changing trends, local garment names, route and line names), while  
1441 food/cuisine and history also depend on rare, culturally bound proper nouns; in the speech modality,  
1442 models must directly map variable acoustic realizations of these low-frequency terms to the correct  
1443 concepts, further depressing performance in these categories.

1444 We further explore the relationship between cultural types and languages based on Figure 16. When  
1445 Mongolian is employed as the input language, the model exhibits markedly superior performance  
1446 in the education category relative to other categories. Conversely, when Hindi, Vietnamese, and  
1447 Tamil are utilized as input languages, the model demonstrates suboptimal performance in the Fash-  
1448 ion/Clothing, Geographical location and climate, and education categories, respectively.

1449 We posit that the underlying cause of this phenomenon may be attributed to the fact that for low-  
1450 resource languages such as Mongolian, casual text from the internet (e.g., social media, forums) is  
1451 relatively scarce. High-quality corpora for such languages predominantly originate from government  
1452 documents, digitized textbooks, academic papers, or Wikipedia. This compels the model to "absorb"  
1453 a substantial volume of formal, educational texts during the training phase. Consequently, the model  
1454 exhibits a form of "overfitting" with exceptionally high performance when processing the education  
1455 category. In contrast, Hindi internet data may be dominated by news, politics, religion, or literature.  
1456 Contemporary fashion and clothing-related content on the Indian internet likely exists primarily  
1457 in English (or Hinglish, a Hindi-English code-mixed variety). If the model is trained exclusively  
1458 on pure Hindi, it consequently lacks the vocabulary and contextual framework necessary to describe  
1459 "fashion." Similarly, if Vietnamese training corpora lack specific technical descriptions of geography  
1460 and climate, the model will exhibit domain-specific knowledge gaps. Tamil, being widely used

1458 across diverse regions including Sri Lanka, India, and Singapore, encounters substantial educational  
 1459 disparities across these countries, and such regional conflicts may impede the model's ability to  
 1460 adequately fit educational domain knowledge.



1486 Figure 12: Performance of different models on different categories in the text-only dataset.  
 1487

#### 1489 A.8 PERFORMANCE OF LLMs ON MMA-ASIA ACROSS MODALITIES

1490 The exact data corresponding to the bar chart in Section 4.2, Figure 2 are presented in Tables 13, 14,  
 1491 15, 16.

1493 Table 13: Text-only modality performance with exact numbers, measured by Accuracy (%): the  
 1494 number of items where the model's choice exactly matches the correct option, divided by 500. “-  
 1495 -” means “not support”. The better-performing result among different languages within the same  
 1496 country is **bolded**.

Model	CN-zh	CN-en	ID-id	ID-en	IN-hi	IN-en	JP-ja	JP-en	KR-ko
Gemini 2.5 Flash	<b>53.6</b>	45.4	<b>56.2</b>	53.2	73.0	<b>73.6</b>	47.4	<b>48.2</b>	72.2
Claude Sonnet 4	<b>53.8</b>	53.6	<b>55.6</b>	54.8	<b>74.4</b>	74.2	48.8	<b>50.6</b>	71.4
GPT-4o	36.0	<b>50.2</b>	53.6	<b>55.8</b>	73.2	<b>77.2</b>	49.8	<b>53.2</b>	71.4
GLM-4-9B-chat	<b>36.0</b>	32.6	37.6	<b>49.4</b>	45.6	<b>57.0</b>	30.2	<b>33.6</b>	45.4
Mistral-8B-Instruct	23.2	<b>26.0</b>	31.8	<b>35.6</b>	42.2	<b>49.0</b>	23.6	<b>30.2</b>	44.8
Baichuan-M2-32B	<b>51.2</b>	49.8	50.6	<b>53.2</b>	54.6	<b>69.0</b>	39.6	<b>40.8</b>	56.4
Qwen3-30B-A3B	<b>47.0</b>	45.8	44.8	<b>46.2</b>	58.6	<b>62.2</b>	<b>38.4</b>	38.0	43.4
Model	KR-en	MN-mn	MN-en	SG-zh	SG-en	SG-ms	SG-ta	VN-vi	VN-en
Gemini 2.5 Flash	65.2	58.2	<b>70.0</b>	34.0	<b>43.4</b>	39.0	31.4	<b>73.6</b>	73.0
Claude Sonnet 4	66.0	59.4	<b>71.0</b>	29.8	<b>41.8</b>	37.8	31.6	<b>74.4</b>	72.2
GPT-4o	68.8	55.4	<b>69.4</b>	22.4	<b>47.6</b>	39.4	33.2	<b>74.8</b>	71.8
GLM-4-9B-chat	42.8	23.0	<b>52.6</b>	20.2	<b>24.8</b>	18.4	19.4	47.6	<b>51.0</b>
Mistral-8B-Instruct	40.8	13.0	<b>50.8</b>	17.2	<b>25.2</b>	22.4	15.4	41.8	<b>48.6</b>
Baichuan-M2-32B	51.6	16.8	<b>61.6</b>	26.6	<b>33.6</b>	28.8	19.4	<b>67.6</b>	62.8
Qwen3-30B-A3B	<b>49.8</b>	44.6	<b>65.0</b>	26.4	<b>29.2</b>	23.6	22.8	57.8	<b>65.8</b>

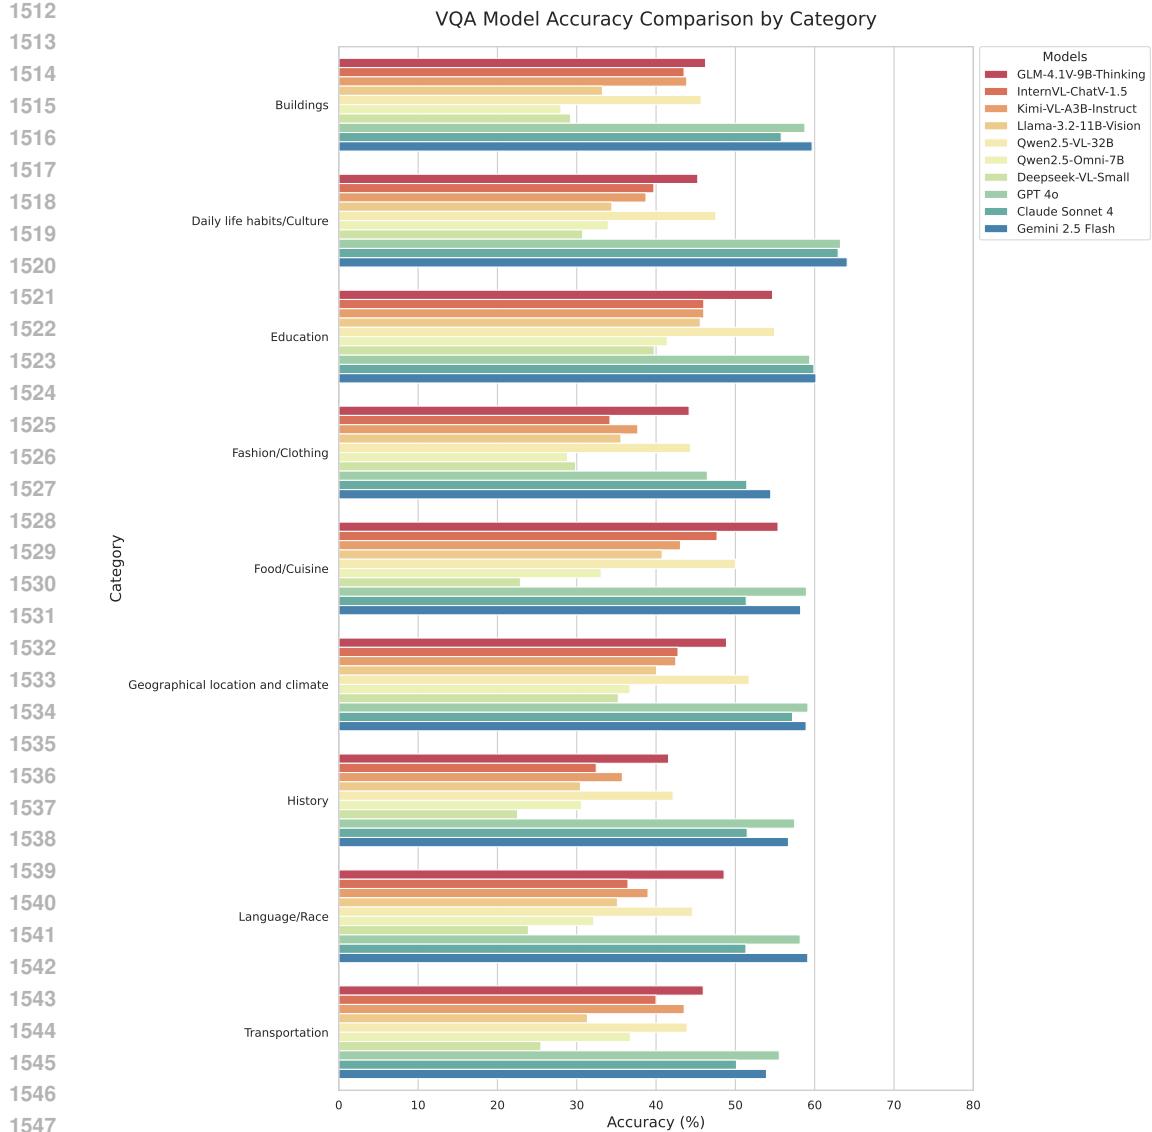


Figure 13: Performance of different models on different categories in the VQA dataset.

#### A.9 JOINT MCQ ANSWERING AND EXPLANATION PERFORMANCE OF LLMs ON MMA-ASIA ACROSS MODALITIES

The model performance based on both the MCQ choices and the model-generated explanations are provided in Tables 2, 17, 18, 19.

## A.10 RESULTS FOR FULLY SPOKEN QUESTION AND ANSWERING

We considered two configurations when constructing the TTS-Spoken QA dataset: (i) converting only the question stem to speech while keeping the answer options as text, and (ii) converting both the stem and the options to speech. To preserve comparability with VQA under controlled variables and to minimize ambiguity introduced by fully spoken options, our main experiments adopt the “spoken stem + textual options” setting across five evaluation dimensions. Results for the fully spoken setting (spoken stem and spoken options) on the test set are provided in Figure 17 and Table 20 for reference. We find that converting both the question and options to speech leads to a significant performance drop compared with the “spoken question + textual options” configuration, indicating that spoken options introduce greater uncertainty than the spoken question itself. This warrants further investigation in future work.

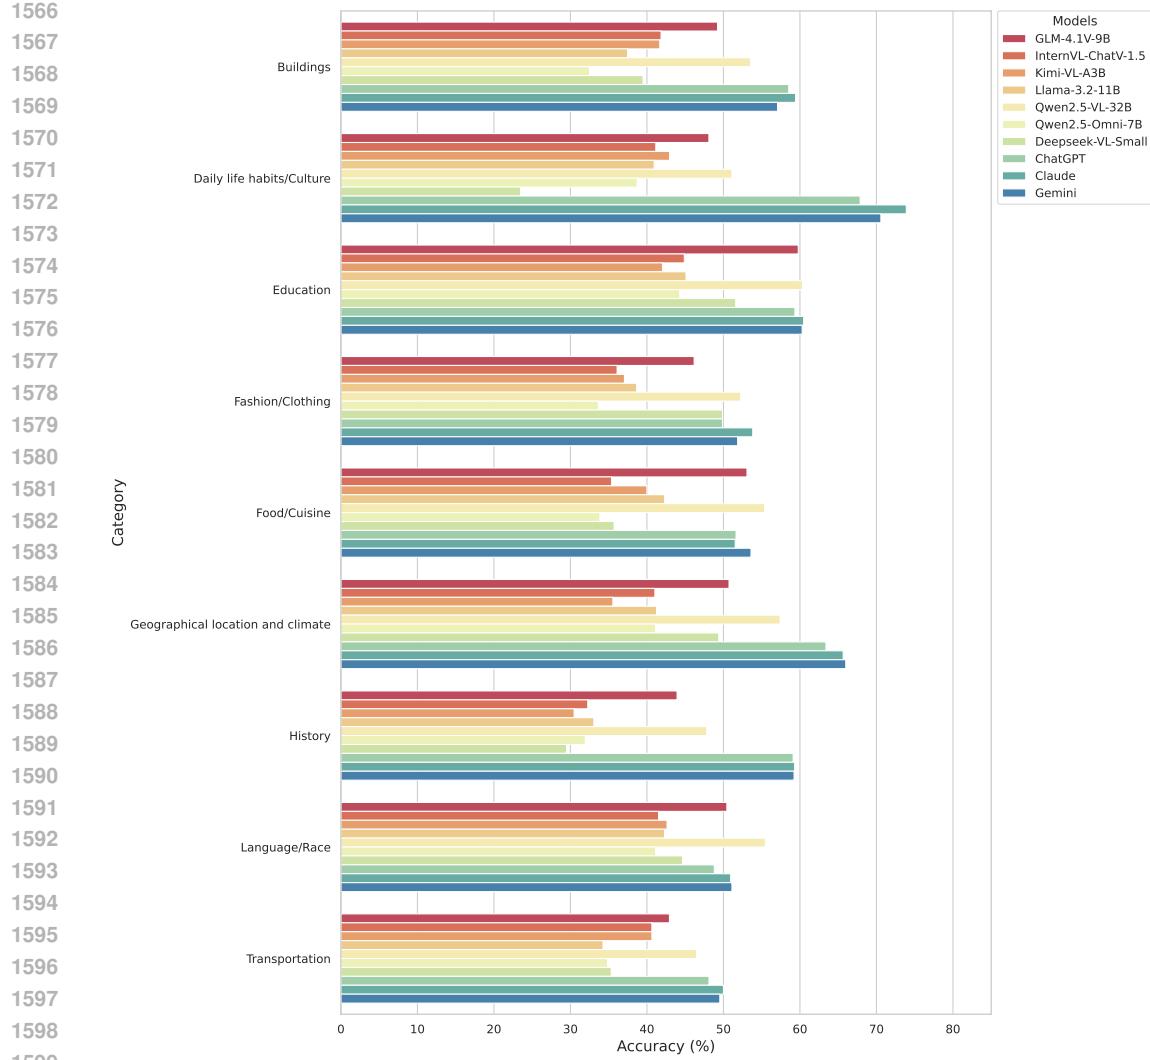


Figure 14: Performance of different models on different categories in the rephrased VQA dataset.

### A.11 VISION-ABLATED PREFIX REPLAY

We found that visual content increases reasoning hallucinations in models compared to text-only QA, despite accurate image recognition capabilities. To validate our hypothesis, we propose a “Vision-ablated Prefix Replay” (VPR) method. This method enables a model with parameters  $\theta$  to first describe the image contents based on image  $\mathbf{x}^{\text{img}}$  and text prompt  $\mathbf{x}^{\text{text}}$ . After this initial description, we structurally ablate the visual condition and fix the prefix  $\hat{\mathbf{S}}_{1:n}$  for subsequent reasoning generation. This evaluates the marginal contribution of visual conditions to reasoning. Specifically, assuming the model completes image description within the first  $n$  tokens, we remove visual conditions starting from the  $(n+1)$ -th token and use only the text prompt and generated tokens as prefix. The joint probability distribution of the subsequent sequence  $\mathbf{S}_{>n} = (s_{n+1}, \dots, s_T)$  can be expressed as:

$$p_{\theta}(\mathbf{S}_{>n} \mid \mathbf{x}_{\text{text}}, \emptyset, \hat{\mathbf{S}}_{1:n}) = \prod_{t=n+1}^T p_{\theta}(s_t \mid \mathbf{x}_{\text{text}}, \emptyset, \hat{\mathbf{S}}_{1:n}, \mathbf{s}_{n+1:t-1}). \quad (6)$$

Previous work has explored related ideas. For example, Zhang et al. (2025) attempt to suppress visual leakage by blocking the attention paths to image-token positions during decoding. However, this cannot fully eliminate the influence of visual content: in a causal language model, the information of earlier image tokens is encoded into subsequent question tokens, so residual visual

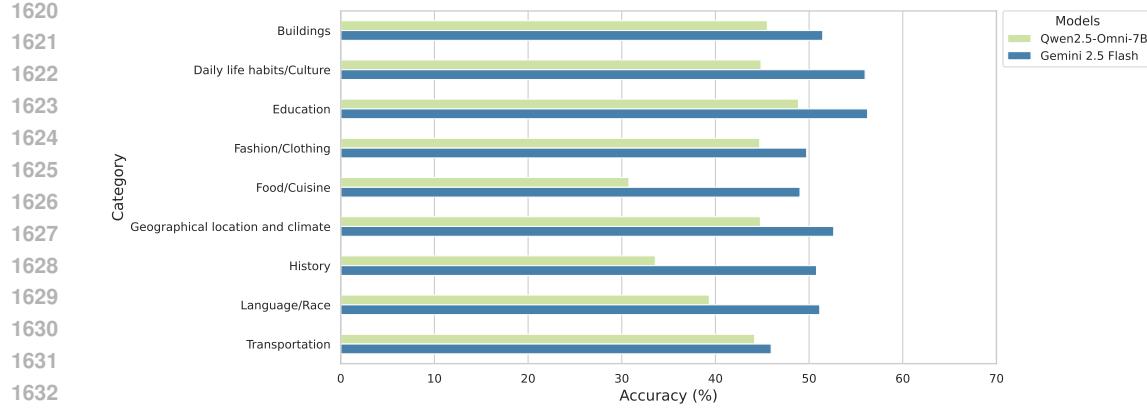


Figure 15: Performance of different models on different categories in the Speech dataset.

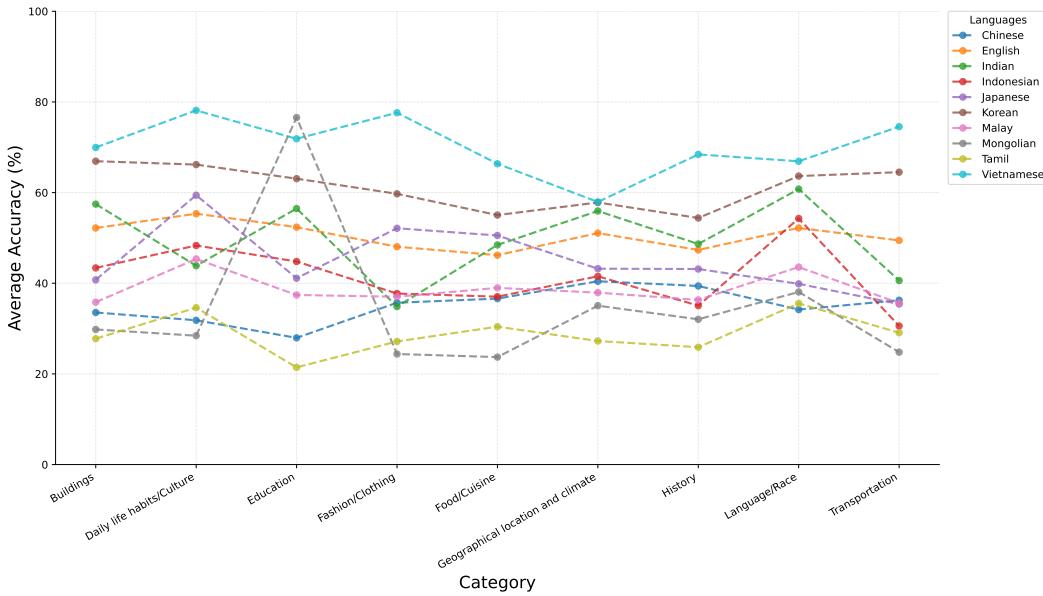


Figure 16: Average performance of the models on different question categories under different language settings.

information remains even when attention to image tokens is blocked. In contrast, our method first elicits a textual description of the image, then removes the visual input and recomputes the representations of the prefix tokens, thereby purging visual information and enabling a more precise assessment of the image content’s contribution to the model’s reasoning process.

For the question in Figure 6, the completed responses under different generation modes are provided in Table 21.

#### A.12 ANALYSIS OF SPEECH AS A CULTURAL PRIOR

Compared to images and text, speech input introduces greater uncertainty through environmental noise, homophony, and accents—with accents closely tied to cultural context. Our research reveals that accents function beyond mere noise. Testing synthetic speech in standard English versus multiple national accents, we found Qwen and Gemini outperformed their standard English baselines in 6 and 5 country-specific cultural settings (Figure 2), respectively. Notably, Qwen achieved 2.8% and 3.6% accuracy gains for Indonesian and Japanese accents (Table 16). We attribute this to systematic co-occurrence of accented English with country-specific entities and contexts in training corpora, enabling accents to serve as cultural and lexical priors during inference. Our findings demonstrate

Table 14: VQA modality performance with exact numbers, measured by Accuracy (%): the number of items where the model’s choice exactly matches the correct option, divided by 500. “-” means “not support”. The better-performing result among different languages within the same country is bolded.

Model	CN-zh	CN-en	ID-id	ID-en	IN-hi	IN-en	JP-ja	JP-en	KR-ko
Gemini 2.5 Flash	<b>45.6</b>	38.8	48.6	<b>49.8</b>	<b>65.4</b>	42.4	54.2	<b>58.0</b>	72.0
Claude Sonnet 4	<b>47.6</b>	33.6	<b>49.8</b>	46.4	<b>54.0</b>	39.0	47.4	<b>53.6</b>	71.8
GPT-4o	32.8	<b>38.4</b>	<b>53.6</b>	50.2	<b>65.8</b>	45.6	59.4	<b>60.0</b>	<b>69.2</b>
Qwen2.5-VL-32B	<b>40.2</b>	30.8	37.6	<b>37.8</b>	40.2	<b>45.0</b>	<b>47.6</b>	46.8	63.2
Llama-3.2-11B-Vision	<b>21.6</b>	15.2	26.2	<b>31.2</b>	35.4	<b>40.2</b>	27.2	<b>35.2</b>	30.0
Kimi-VL-A3B-instruct	-	<b>21.2</b>	-	<b>25.4</b>	-	<b>34.4</b>	-	<b>37.2</b>	-
Qwen2.5-Omni-7B	<b>30.6</b>	23.6	21.8	<b>32.4</b>	19.4	<b>33.6</b>	25.6	<b>39.2</b>	18.6
Deepseek-VL-Small	-	<b>12.2</b>	-	<b>25.4</b>	-	<b>25.0</b>	-	<b>26.0</b>	-
GLM-4.1V-9B-Thinking	<b>46.4</b>	26.2	28.6	<b>34.8</b>	<b>50.8</b>	43.4	44.0	<b>44.8</b>	61.0
InternVL-Chat-V1.5	-	<b>15.4</b>	-	<b>23.0</b>	-	<b>39.2</b>	-	<b>40.6</b>	-
Model	KR-en	MN-mn	MN-en	SG-zh	SG-en	SG-ms	SG-ta	VN-vi	VN-en
Gemini 2.5 Flash	<b>72.8</b>	41.2	<b>49.2</b>	53.0	<b>62.2</b>	59.8	60.8	<b>76.6</b>	65.4
Claude Sonnet 4	<b>74.4</b>	35.0	<b>49.0</b>	34.6	<b>57.2</b>	55.2	49.8	<b>73.2</b>	67.4
GPT-4o	65.8	41.4	<b>53.0</b>	30.2	<b>70.6</b>	68.4	62.4	<b>75.8</b>	63.8
Qwen2.5-VL-32B	<b>67.0</b>	20.6	<b>45.2</b>	40.8	<b>52.8</b>	45.6	22.6	<b>65.0</b>	64.6
Llama-3.2-11B-Vision	<b>47.8</b>	5.8	<b>21.6</b>	32.2	<b>39.4</b>	35.0	15.0	48.6	<b>49.6</b>
Kimi-VL-A3B-instruct	<b>56.0</b>	-	-	38.6	<b>43.6</b>	-	-	-	<b>50.0</b>
Qwen2.5-Omni-7B	<b>49.2</b>	3.2	<b>36.6</b>	30.0	<b>43.4</b>	38.6	14.2	<b>59.8</b>	58.2
Deepseek-VL-Small	<b>30.0</b>	-	-	16.4	<b>27.2</b>	-	-	-	<b>41.6</b>
GLM-4.1V-9B-Thinking	<b>65.4</b>	28.0	<b>39.4</b>	44.8	<b>52.8</b>	41.8	34.8	59.4	<b>59.8</b>
InternVL-Chat-V1.5	<b>56.6</b>	-	-	28.0	<b>44.8</b>	-	-	-	<b>54.0</b>

Table 15: Rephrase VQA (Text-Only) modality performance with exact numbers, measured by Accuracy (%): the number of items where the model’s choice exactly matches the correct option, divided by 500. “-” means “not support”. The better-performing result among different languages within the same country is bolded.

Model	CN-zh	CN-en	ID-id	ID-en	IN-hi	IN-en	JP-ja	JP-en	KR-ko
Gemini 2.5 Flash	<b>45.2</b>	37.2	<b>52.0</b>	48.8	<b>67.0</b>	53.8	62.4	<b>65.6</b>	<b>76.4</b>
Claude Sonnet 4	<b>52.8</b>	48.0	<b>53.8</b>	52.6	<b>68.6</b>	51.8	67.8	<b>68.2</b>	<b>79.6</b>
GPT-4o	41.0	<b>46.2</b>	<b>57.8</b>	55.2	<b>69.0</b>	59.6	<b>67.0</b>	66.4	73.2
Qwen2.5-VL-32B	<b>45.0</b>	34.4	<b>45.0</b>	44.4	<b>51.1</b>	47.0	56.0	<b>58.4</b>	66.8
Llama-3.2-11B-Vision	<b>18.4</b>	15.6	<b>35.2</b>	30.4	<b>39.8</b>	36.6	39.2	<b>45.6</b>	43.8
Kimi-VL-A3B-Instruct	-	<b>21.6</b>	-	<b>24.6</b>	-	<b>32.2</b>	-	<b>43.8</b>	-
Qwen2.5-Omni-7B	<b>31.4</b>	22.8	23.4	<b>35.8</b>	28.4	<b>42.8</b>	35.0	<b>51.4</b>	29.6
Deepseek-VL-Small	-	<b>15.6</b>	-	<b>30.6</b>	-	<b>29.2</b>	-	<b>40.4</b>	-
GLM-4.1V-9B-Thinking	<b>46.4</b>	30.6	<b>37.6</b>	33.6	<b>51.4</b>	33.8	49.0	<b>53.0</b>	66.6
InternVL-Chat-V1.5	-	<b>16.8</b>	-	<b>27.0</b>	-	<b>37.4</b>	-	<b>40.9</b>	-
Model	KR-en	MN-mn	MN-en	SG-zh	SG-en	SG-ms	SG-ta	VN-vi	VN-en
Gemini 2.5 Flash	76.4	47.8	<b>49.2</b>	54.6	<b>63.6</b>	59.6	52.4	<b>76.0</b>	71.6
Claude Sonnet 4	77.0	38.8	<b>56.0</b>	46.0	<b>65.2</b>	61.0	51.8	<b>76.0</b>	73.6
GPT-4o	<b>75.2</b>	43.2	<b>59.8</b>	30.2	<b>68.2</b>	66.4	53.6	<b>79.4</b>	72.8
Qwen2.5-VL-32B	<b>72.0</b>	19.0	<b>56.2</b>	36.4	<b>57.6</b>	50.8	25.4	69.2	<b>70.0</b>
Llama-3.2-11B-Vision	<b>54.6</b>	10.2	<b>29.0</b>	25.4	<b>45.2</b>	40.2	20.0	39.4	<b>51.2</b>
Kimi-VL-A3B-Instruct	<b>53.8</b>	-	-	<b>40.4</b>	36.6	-	-	-	<b>40.0</b>
Qwen2.5-Omni-7B	<b>60.6</b>	2.6	<b>43.4</b>	21.4	<b>43.6</b>	29.4	9.8	58.4	<b>60.0</b>
Deepseek-VL-Small	<b>58.8</b>	-	-	25.2	<b>39.2</b>	-	-	-	<b>58.4</b>
GLM-4.1V-9B-Thinking	<b>67.0</b>	26.6	<b>43.8</b>	42.0	<b>51.0</b>	40.4	34.8	<b>66.2</b>	62.8
InternVL-Chat-V1.5	<b>48.6</b>	-	-	36.0	<b>41.6</b>	-	-	-	<b>46.0</b>

that accents can function as valuable cultural cues rather than simply noise sources for model exploitation.

1728  
 1729 Table 16: Speech modality (speech question & text options) performance with exact numbers, mea-  
 1730 sured by Accuracy (%): the number of items where the model’s choice exactly matches the correct  
 1731 option, divided by 500. “-” means “not support”. The better-performing result among different lan-  
 1732 guages within the same country is **bolded**.

Model	CN-zh	CN-en	CN-en	ID-id	ID-en	ID-en	IN-hi	IN-en	IN-en
	-	Acc	NoAcc	-	Acc	NoAcc	-	Acc	NoAcc
Qwen2.5-Omni-7B	<b>26.2</b>	21.2	21.0	29.4	29.4	<b>32.0</b>	33.2	<b>33.2</b>	30.4
Gemini 2.5 Flash	<b>40.6</b>	37.6	36.0	36.6	40.2	<b>40.4</b>	<b>53.4</b>	25.4	24.2
Model	JP-ja	JP-en	JP-en	KR-ko	KR-en	KR-en	MN-mn	MN-en	MN-en
	-	Acc	NoAcc	-	Acc	NoAcc	-	Acc	NoAcc
Qwen2.5-Omni-7B	<b>45.4</b>	41.8	-	<b>53.6</b>	51.4	-	43.8	<b>47.8</b>	-
Gemini 2.5 Flash	52.4	<b>60.2</b>	57.0	<b>73.6</b>	70.2	72.0	32.0	48.0	<b>50.4</b>
Model	SG-zh	SG-en	SG-en	SG-ms	SG-ta	VN-vi	VN-en	VN-en	VN-en
	-	Acc	NoAcc	-	-	-	Acc	NoAcc	
Qwen2.5-Omni-7B	<b>39.4</b>	39.4	37.0	-	-	-	<b>54.6</b>	52.0	
Gemini 2.5 Flash	47.0	<b>55.2</b>	54.4	41.4	39.6	60.6	<b>68.6</b>	67.8	

1744  
 1745 Table 17: Evaluated under the MCQ+Explanation metric, VQA modality performance with ex-  
 1746 act numbers, measured by Accuracy (%): the number of items where the model’s choice exactly  
 1747 matches the correct option, divided by 500.

Model	CN-en	ID-en	IN-en	JP-en	KR-en	MN-en	SG-en	VN-en
Gemini 2.5 Flash	29.8	35.9	30.5	43.7	66.4	43.9	47.1	53.8
Claude Sonnet 4	23.8	28.2	25.9	38.4	67.2	40.5	41.4	55.9
GPT-4o	29.9	31.9	27.4	40.0	60.1	37.2	53.2	55.8
Qwen2.5-VL-32B	25.0	22.6	26.1	34.2	62.4	41.0	36.8	51.6
Llama-3.2-11B-Vision	7.0	12.3	14.8	13.8	40.0	15.2	19.4	28.8
Kimi-VL-A3B-instruct	11.8	9.9	13.1	16.6	49.9	29.2	21.7	30.6
Qwen2.5-Omni-7B	15.6	21.5	14.7	23.5	43.9	31.6	29.5	44.3
GLM-4.1V-9B-Thinking	19.3	19.1	22.6	28.6	57.6	33.2	36.8	46.7
InternVL-Chat-V1.5	6.3	7.6	13.0	16.1	47.6	20.6	22.6	30.5

1757  
 1758 Table 18: Evaluated under the MCQ+Explanation metric, Rephrase VQA (Text-Only) modality per-  
 1759 formance with exact numbers, measured by Accuracy (%): the number of items where the model’s  
 1760 choice exactly matches the correct option, divided by 500.

Model	CN-en	ID-en	IN-en	JP-en	KR-en	MN-en	SG-en	VN-en
Gemini 2.5 Flash	30.2	27.1	27.9	58.3	75.1	45.0	53.7	62.1
Claude Sonnet 4	39.5	37.2	43.3	<b>59.7</b>	75.1	51.7	55.0	63.0
GPT-4o	39.8	35.7	48.4	61.1	68.8	53.5	62.6	66.4
Qwen2.5-VL-32B	28.0	27.1	30.0	47.5	70.1	51.1	44.3	57.8
Llama-3.2-11B-Vision	7.6	13.8	17.6	24.5	52.6	21.6	27.1	35.5
Kimi-VL-A3B-Instruct	12.5	12.3	15.2	26.5	51.7	32.2	23.5	28.5
Qwen2.5-Omni-7B	18.0	25.6	25.5	39.4	59.5	39.6	35.3	49.9
GLM-4.1V-9B-Thinking	22.9	19.8	20.8	42.6	64.6	36.3	39.3	51.8
InternVL-Chat-V1.5	9.3	11.0	16.9	21.2	45.4	28.0	25.1	31.3

1770  
 1771 Table 19: Evaluated under the MCQ+Explanation metric, Speech modality (speech question & text  
 1772 options) performance with exact numbers, measured by Accuracy (%): the number of items where  
 1773 the model’s choice exactly matches the correct option, divided by 500.

Model	CN-en	ID-en	IN-en	JP-en	KR-en	MN-en	SG-en	VN-en
Qwen2.5-Omni-7B	21.0	32.0	30.4	41.8	51.4	47.8	37.0	52.0
Gemini 2.5 Flash	36.0	40.4	24.2	57.0	72.0	50.4	54.4	67.8

### A.13 CONSISTENCY ANALYSIS AND HYPERPARAMETERS SETTING FOR LLM-AS-JUDGE

Before adopting the LLM-as-judge paradigm, we conducted a small-scale study to assess human-  
 model agreement and inter-model agreement. In addition, we further investigated the accuracy of  
 using the LLM-as-Judge method in English versus non-English languages.

Table 20: Speech modality (speech question & options) performance with exact numbers, measured by Accuracy (%): the number of items where the model’s choice exactly matches the correct option, divided by 500. “-” means “not support”. The better-performing result among different languages within the same country is **bolded**.

Model	CN-zh	CN-en	CN-en	ID-id	ID-en	ID-en	IN-hi	IN-en	IN-en
	-	Acc	NoAcc	-	Acc	NoAcc	-	Acc	NoAcc
Qwen2.5-Omni-7B	<b>28.4</b>	14.2	15.2	-	23.8	<b>27.0</b>	-	<b>22.2</b>	20.8
Gemini 2.5 Flash	<b>27.4</b>	13.8	15.2	25.0	<b>31.0</b>	27.0	<b>34.6</b>	34.2	29.4
Model	JP-ja	JP-en	JP-en	KR-ko	KR-en	KR-en	MN-mn	MN-en	MN-en
	-	Acc	NoAcc	-	Acc	NoAcc	-	Acc	NoAcc
Qwen2.5-Omni-7B	-	<b>35.2</b>	32.8	-	34.0	<b>35.0</b>	-	23.6	<b>38.4</b>
Gemini 2.5 Flash	35.0	<b>42.8</b>	36.2	38.6	20.8	<b>49.4</b>	10.6	18.2	<b>28.6</b>
Model	SG-zh	SG-en	SG-en	SG-ms	SG-ta	VN-vi	VN-en	VN-en	VN-en
	-	Acc	NoAcc	-	-	-	Acc	NoAcc	
Qwen2.5-Omni-7B	<b>27.0</b>	26.0	23.6	-	-	-	<b>39.2</b>	37.0	
Gemini 2.5 Flash	28.8	<b>42.2</b>	37.6	16.0	15.8	<b>52.0</b>	46.2	45.8	

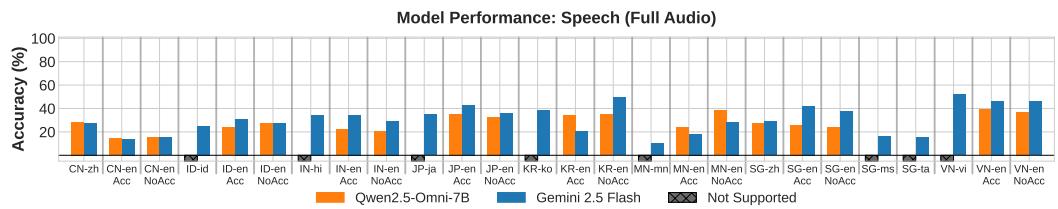


Figure 17: Speech modality (speech question & options) performance with exact numbers, measured by Accuracy (%): the number of items where the model’s choice exactly matches the correct option, divided by 500.

**Human–model agreement.** We sampled 50 items from the dataset, each comprising a multi-step reasoning question and its decomposed sub-questions. Three annotators independently judged whether the model’s answer was semantically consistent with the gold answer for each (binary: *yes/no*). For each sub-question, the human judgment was determined by majority vote. We then queried **Claude Sonnet 4**, **GPT-4o**, and **Gemini 2.5 Flash** via API to obtain their judgments on the same items. An item was counted as *consistent* for a model only if the model’s judgments for all sub-questions and the final question matched the human judgments. Results showed human–model agreement of **98%** for GPT-4o, **98%** for Claude, and **96%** for Gemini. Considering cost, we selected **GPT-4o** as the primary judge.

**Inter-model agreement.** Given the binary nature of the task and the observed human–model agreement rates (98%, 98%, 96%), the conservative lower-bound on inter-model agreement is **96%**. We therefore conclude that a single, top-performing judge model suffices for our setting, and cross-model adjudication is unnecessary.

**Human–model agreement rates under different language settings.** Prior work has extensively evaluated the performance of different LLMs on the same tasks across different languages, revealing significant gaps between English and low-resource language settings Gupta et al. (2025); Romanou et al. (2024); Xuan et al. (2025). (Xuan et al., 2025) utilize the MMLU-ProX benchmark to demonstrate that leading models, including GPT-4o and Gemini, suffer a massive accuracy drop, up to 30%, when transitioning from English to low-resource languages. Complementing this, (Romanou et al., 2024) highlight a cultural gap in the INCLUDE benchmark, where a lack of regional knowledge accounts for nearly 40% of model failures in non-English contexts. Furthermore, in the educational domain, (Gupta et al., 2025) report that GPT-4o and Gemini exhibit significant performance biases, struggling with complex pedagogical tasks such as feedback generation in languages like Telugu and Farsi compared to their English performance. We have also conducted experiments on using LLMs as evaluators in non-English scenarios. We compared the accuracy of using LLMs to evaluate outputs in English, Malay, and Tamil, with this evaluation performed on 50 semantically identical questions across different languages. We used human judgment as ground truth to calculate each model’s Human–Model consistency across different languages. As shown in Table 22, model evaluation accuracy in non-English scenarios is significantly lower than in English scenarios, espe-

cially in low-resource languages. Therefore, our evaluation experiments for model’s explanation are currently conducted only in English settings to ensure accurate reflection of model capabilities.

#### A.14 RATIONALE UNFAITHFULNESS RATES ACROSS REPHRASE VQA AND SPOKEN QA

The results for LLMs’ Rationale Unfaithfulness Rates (RUR) across Rephrase VQA and Speech are shown in Fig. 18. We observe patterns consistent with the Text-Only and VQA modalities: closed-source models generally have lower RURs than open-source models, though they still fall within the 5%–20% range. Among open-source models, Llama shows a markedly higher RUR on non-Spanish languages than on Spanish, which we attribute to linguistic bias stemming from the disproportionately large share of Spanish in Llama’s training data relative to other languages.

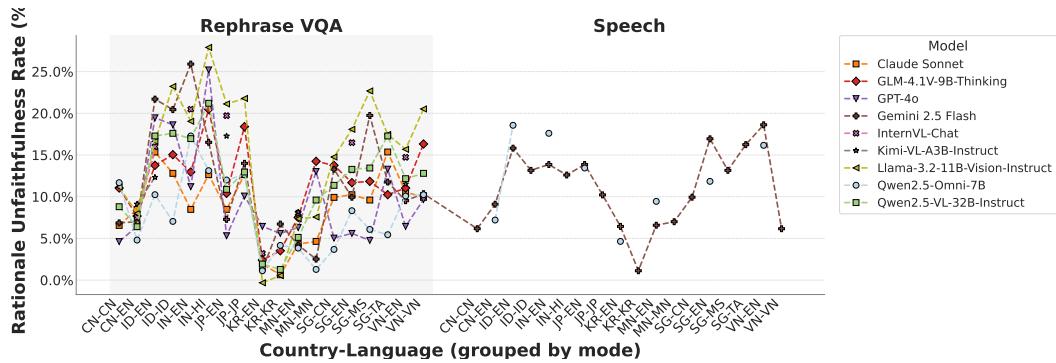


Figure 18: Rationale Unfaithfulness Rates of LLMs across Rephrase VQA and Speech.

#### A.15 ATTENTION HEATMAPS FOR INCORRECT MODEL PREDICTIONS

Figures 19 shows attention heatmaps over image regions when the model produces incorrect answers. In Case 2, the model concentrates on the subject’s clothing and surrounding scene but overlooks the footwear beside the person and the key text in the upper-left corner that identifies the subject. Similarly, in Figure 3 the model focuses excessively on attire and playing posture while ignoring the crucial detail that the instrument’s body is wrapped in snakeskin. All of these support our finding: the model’s excessive focus on items mentioned in the prompt causes it to overlook critical cultural details.

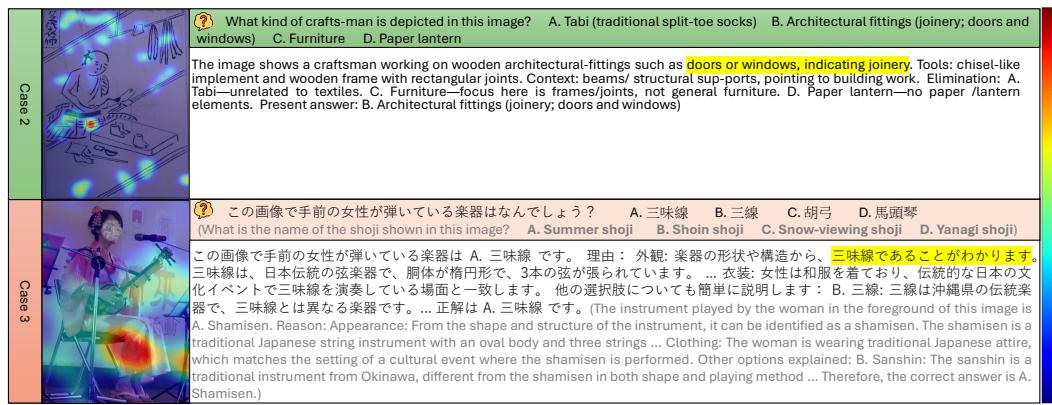


Figure 19: Attention heatmap visualization over image regions during incorrect model answers. Color scale from blue (low) to red (high) indicates increasing model attention.

#### A.16 ERROR CASE STUDY

As illustrated in Figure 20, when queried with “What is the mountain in the picture called?”, the model exhibits hallucinations under both English and Chinese prompts; however, the Chinese reasoning more closely approximates the ground truth. When the VQA query is reformulated into a

corresponding text-only question, errors under the English prompt are primarily attributed to cross-lingual alignment failures. Specifically, the model erroneously maps “the world’s number one mountain” to the Chinese phrase “天下第一奇山”, resulting in an incorrect prediction. Conversely, the Chinese prompt elicits the correct answer and explanation. This observation substantiates our hypothesis that image tokens play a more significant role in inducing hallucinations during reasoning. Furthermore, for non-English languages with adequate representation, employing the native language for culture-specific queries consistently enhances performance.

Figure 21 depicts the divergence in model responses across Text-Only, VQA, and Speech modalities for the same inquiry. VQA results indicate that the model lacks specific visual grounding for “Chongqing Guoyuan Port,” rendering it unable to recognize the location. Consequently, it relies on generic scene-level inference, which leads to erroneous outputs. In the Speech modality, failures stem from token misinterpretation or transcription errors—for instance, the phonetic misrecognition of “Chongqing Guoyuan” as “Chongqing Guoyan”, which similarly yields an incorrect response. In other words, the lack of relevant image training data and the disturbances introduced by speech input both make the model more prone to errors in the VQA and speech modalities than in the text-only modality.

Question	Model Answer	CA
What is the mountain in the picture called? A. The world’s most unique mountain B. The world’s number one mountain C. Central Mountain D. The most revered of the Five Great Mountains	The mountain in the image is Mount Huangshan (Yellow Mountain), famous for its granite peaks, ancient pine trees, and sea of clouds. None of the options directly name Mount Huangshan. However, if the question is about reverence and cultural significance, D. The most revered of the Five Great Mountains could be seen as the closest, as it usually refers to Mount Tai (Taishan), the most sacred of the Five Great Mountains in Chinese culture. If the image is indeed Huangshan, strictly speaking, none of the options are fully correct.	A
图中的大山被称为什么? (What is the mountain in the picture called?) A. 天下第一奇山 (The world’s most unique mountain) B. 天下第一山 (The world’s number one mountain) C. 中岳 (Central Mountain) D. 五岳独尊 (The most revered of the Five Great Mountains)	A, D 这张图片展示的是中国著名的黄山。黄山以奇松、怪石、云海、温泉和冬雪“五绝”闻名,被誉为“天下第一奇山”(选项A)。同时,黄山也有“五岳归来不看山,黄山归来不看岳”的美誉,因此也被称为“五岳独尊”(选项D)。(This image shows China’s famous Mount Huangshan. Huangshan is renowned for its “five wonders”—strange pines, peculiar rocks, seas of clouds, hot springs, and winter snow—and is therefore known as “the world’s most unique mountain” (Option A). It is also praised with the saying “After visiting the Five Great Mountains, one need not see other mountains; after visiting Huangshan, one need not see the Five Great Mountains,” so it is also called “the most revered of the Five Great Mountains” (Option D).)	A
What are the other names for Huangshan? A. The world’s most unique mountain B. The world’s number one mountain C. Central Mountain D. The most revered of the Five Great Mountains	The correct answers are A. The world’s most unique mountain and B. The world’s number one mountain. Explanation: Huangshan (黄山), also known as Yellow Mountain, is one of China’s most famous and iconic mountain ranges. It is celebrated for its stunning natural beauty, including granite peaks, ancient pine trees, and seas of clouds. Over the centuries, Huangshan has inspired many Chinese poets, painters, and philosophers. A. The world’s most unique mountain: Huangshan is often considered one of the most unique mountains in the world because of its extraordinary landscape and cultural significance. Its beauty is regarded as unparalleled, so this title is appropriate. B. The world’s number one mountain: In Chinese culture, Huangshan is sometimes called “the world’s number one mountain” (天下第一奇山), highlighting its exceptional status and reputation. C. Central Mountain: This usually refers to Mount Song (嵩山), one of the Five Great Mountains of China in Henan Province. It is not another name for Huangshan. D. The most revered of the Five Great Mountains: This title is typically used for Mount Tai (泰山), which has great historical and spiritual significance. Huangshan is not one of the Five Great Mountains. Therefore, the correct options are A and B.	A
图中的大山被称为什么? (What are the other names for Huangshan?) A. 天下第一奇山 (The world’s most unique mountain) B. 天下第一山 (The world’s number one mountain) C. 中岳 (Central Mountain) D. 五岳独尊 (The most revered of the Five Great Mountains)	正确答案: A 解释: 黄山是中国著名的山岳,以奇松、怪石、云海、温泉和冬雪闻名,被誉为“天下第一奇山”(选项A)。被称为“天下第一山”的是泰山(选项B)。选项C“中岳”是嵩山的别称,与黄山无关。选项D“五岳独尊”是泰山的荣誉称号,也与黄山无关。(Correct answer: A) Explanation: Mount Huangshan is a famous mountain in China, known for its unique pine trees, oddly shaped rocks, seas of clouds, hot springs, and winter snow. It is honored as “the world’s most unique mountain” (Option A). The title “the world’s number one mountain” refers to Mount Tai (Option B). Option C, “Central Mountain,” is another name for Mount Song and has nothing to do with Huangshan. Option D, “the most revered of the Five Great Mountains,” is also an honorary title for Mount Tai and is unrelated to Huangshan.)	A

Figure 20: A comparison of the model’s responses to the question “What are the other names for Huangshan?” across languages and modalities. The first two examples are VQA questions in English and Chinese, and the last two are text-only questions in English and Chinese.

### A.17 SUB-QUESTION DECOMPOSITION

Table 23 illustrates examples of deconstructing multi-step reasoning questions into single-fact verification subquestions. We enumerate all single-fact subquestions embedded in each multi-step item to assess the model’s generalization.

### A.18 THE USE OF LARGE LANGUAGE MODELS

In this work, LLMs are used for three purposes: (1) we use GPT-4o and Claude Sonnet 4 to translate created datasets to the English version; (2) we use GPT-4o to evaluate whether model responses match the reference answers in Section 4.3 and Section 4.6; and (3) we use GPT-4o to polish the manuscript.

### A.19 DATA SENSITIVITY AND PRIVACY DISCUSSION

In this section, we discuss data sensitivity and privacy for three modalities.

**Text.** We collect culture-related textual facts from the public internet and filter out any content containing racism or hate speech. Our “knowledge points” are a few sentences manually summarized by annotators, and all questions are authored from scratch, so no infringement issues are involved.

1944	Question	Model Answer	CA
1945	What multimodal transport functions are demonstrated by Chongqing Guoyuan Port?	The possible options are A, B, C. Chongqing Guoyuan Port can demonstrate multimodal transport functions including water - to - water transfer, rail - water intermodal transport, and road - water intermodal transport. It doesn't have air - water intermodal transport function. So the answer is A, B, and C. If you have any other questions about Chinese culture or anything else, feel free to let me know.	A, B, C
1946	A. Water-to-water transfer B. Rail-water intermodal transport C. Road-water intermodal transport D. Air-water intermodal transport		
1947	What multimodal transport functions are reflected in the transportation hub shown in the picture?	The possible options are B, C. The image shows a large number of containers, which are typically used for rail and road transport. There's also a waterway nearby, indicating rail-water intermodal transport and road-water intermodal transport. So, the correct answer is B and C. What do you think about this? Do you have any other questions related to this?	A, B, C
1948	A. Water-to-water transfer B. Rail-water intermodal transport C. Road-water intermodal transport D. Air-water intermodal transport		
1949			
1950			
1951			
1952	 Audio Content: (What multimodal transport functions are demonstrated by Chongqing Guoyuan Port?) A. Water-to-water transfer B. Rail-water intermodal transport C. Road-water intermodal transport D. Air-water intermodal transport	The correct options are B, C, and D. Chongqing Guoyuan Port demonstrates rail - water intermodal transport, road - water intermodal transport, and air - water intermodal transport. Well, that's the answer. If you have any other questions or want to discuss more about multimodal transport, feel free to let me know.	A, B, C
1953			
1954			
1955			

Figure 21: Comparison of model responses across text-only, VQA, and speech modalities to the English question “What multimodal transport functions are demonstrated by Chongqing Guoyuan Port?”. The first, second, and third examples correspond to text-only, VQA, and speech, respectively, with red highlights marking misinterpretations of the speech input.

1960

1961

**Images.** We use only images under Creative Commons licenses and strictly for research purposes. We apply an automated face-blurring tool (Ren et al., 2023) to protect privacy, followed by manual review to catch any misses or false positives. Please note that faces of historical figures or cartoon characters are not masked. Our images do not contain pornography, violence, or other harmful content.

**Speech.** We obtained voice samples from individuals outside the data team for speech synthesis, without disclosing any personal information. Consent was obtained prior to recording, and the audio is used solely for research.

1969

#### A.20 ANALYSIS OF CULTURAL AND LINGUISTIC BIASES

Building upon the phenomena observed in our experiments, we conducted a systematic and in-depth analysis of model biases. Given the intrinsic inextricability of culture and language, we structured our analysis along two dimensions: **Language Bias** and **Culture Bias**, while also exploring the underlying causes of these phenomena.

**Language Bias** Experimental results reveal a discernible bias across different models regarding linguistic scripts. Specifically, the Llama-3.2-11B model tends to exhibit superior performance for the same cultural context when prompts are formulated in Latin-based languages (e.g., English, Indonesian, and Malay), whereas performance degrades noticeably in non-Latin languages.

As illustrated in Fig.22, when presented with the prompt “*In the city where Qian Xuesen’s ancestral home is located, you hear two people having a conversation: A: Qin ah zong ah peng! B: Qin ah zong ah peng! What are they doing?*” in English and Chinese respectively, the model correctly identifies the answer and provides an explanation in the English context. However, in the Chinese context, the model ignores the critical premise regarding “*the city where Qian Xuesen’s ancestral home is located*,” resulting in severe hallucinations.

Similarly, Qwen3-30B-A3B-Thinking-2507 demonstrates robust performance on non-low-resource languages but experiences a sharp performance decline in Mongolian and Tamil. Specifically, under identical cultural contexts, performance drops by 20.4% and 6.4% respectively compared to English prompts, highlighting a bias in the distribution of training languages.

Concurrently, we observed that while errors occur in non-low-resource languages, they primarily manifest as inference errors regarding adherence to problem premises. Conversely, in low-resource languages, model outputs often show weak correlation with the question, frequently devolving into mere elaborations of the options. This indicates a marked deficiency in contextual integration for low-resource languages, limiting the models’ capacity for cultural comprehension and analysis.

These phenomena not only uncover biases in language distribution under identical cultural contexts but also expose the limitations of cross-lingual knowledge transfer within multilingual Large Language Models (LLMs). Since the next-token prediction task serves as the primary pre-training objective, the majority of data relies on predicting the subsequent token within a monolingual sequence, lacking explicit objectives for cross-lingual semantic space alignment. Consequently, knowledge tends to remain siloed within single languages, restricting cross-lingual transfer and contextual grounding in low-resource languages. Future research should explore methods to leverage

1998 knowledge inherent in high-resource languages to aid understanding and generation in low-resource  
 1999 languages via semantic alignment.

2000 **Culture Bias** Prior literature has established that most models, due to training data distribution,  
 2001 exhibit stronger capabilities in mainstream Western cultures while underperforming in Asian contexts.  
 2002 In our study, we observed similar cultural biases across eight Asian nations. Our analysis  
 2003 reveals that models are most attuned to Korean culture (63.98%), supported by global diffusion and  
 2004 data richness Jang et al. (2024); Dal Yong (2018). Vietnamese culture (62.96%) shows similarly  
 2005 high performance, hypothetically driven by the abundance of digital content resulting from the coun-  
 2006 try's 79.8% social media penetration rate (DataReportal et al., 2025). However, distinct challenges  
 2007 emerge for China and India, where larger performance gaps consistent with cultural and regional  
 2008 heterogeneity are observed. The results for Mongolia fall further behind, corroborating the impact  
 2009 of data scarcity in low-resource language domains.

2010 In terms of data distribution, cultural popularity correlates positively with the quality and quantity of  
 2011 training data (e.g., Korea, Vietnam). However, cultural diversity and complexity serve as significant  
 2012 confounding factors. Although Chinese culture possesses a vast volume of data, its high complexity  
 2013 and regional diversity make it difficult for models to fit; for instance, the distinction between North-  
 2014 ern “Tangyuan” and Southern “Yuanxiao” is subtle yet non-negligible. Similarly, India encompasses  
 2015 22 official languages, with distinct regional customs for the same festivals. Conversely, Mongolia  
 2016 exhibits a significant “long-tail effect”: it is not only a typical low-resource language domain but also  
 2017 suffers from data sparsity exacerbated by the concurrent use of Cyrillic and Traditional Mongolian  
 2018 scripts.

2019 From a model training perspective, the local languages of Indonesia, Vietnam, and Malaysia share  
 2020 the Latin alphabet with English, facilitating the cross-lingual transfer of English knowledge during  
 2021 training. Furthermore, shared tokens reduce the compression rate of these languages in the vocabu-  
 2022 lary, thereby enhancing model performance.

2023 These findings highlight that cultural bias in current models is jointly driven by uneven data distribu-  
 2024 tion and cultural-linguistic complexity (e.g., China, India, Mongolia), beyond simple “data volume”  
 2025 effects seen in Korea and Vietnam. Going forward, a key direction is to curate higher-quality, re-  
 2026 gionally diverse data and develop training strategies that better support cross-script transfer and  
 2027 fine-grained, intra-cultural distinctions, especially for low-resource and long-tail cultures.

2028

2029

2030

2031

2032

2033

2034

2035

2036

2037

2038

2039

2040

2041

2042

2043

2044

2045

2046

2047

2048

2049

2050

2051

2052  
2053  
2054  
2055  
2056  
2057  
2058  
2059  
2060  
2061  
2062  
2063

Question	Category	Language	Model Answer
In the city where Qian Xuesen's ancestral home is located, you hear two people having a conversation: A: Qin ah zong ah peng! B: Qin ah zong ah peng! What are they doing? A. Sawing wood B. Tug of war C. Fluffing cotton D. Rock-paper-scissors	Language/Race	English	In the conversation, the two people are saying "Qin ah zong ah peng!" which is a Wu Chinese phrase used as a signal to start a game of rock-paper-scissors.
在钱学森祖籍所在的城市，你听到两个人正在对话： 甲：琴啊棕啊棚！ 乙：琴啊棕啊棚！ 请问，他们这是在做什么呢？ A.锯木头 B.拔河 C.弹棉花 D.石头剪刀布	Language/Race	Chinese	他们这是在弹棉花。"琴啊棕啊棚"是一句传统的中国谐音谜语，谐音为"弹棉花"。
What characteristics of Mongolian people's production and daily life do traditional saddles reflect? A. Agriculture-based, horse tack primarily used for farm work. B. Nomadic lifestyle and military requirements are equally emphasized. C. Urban trade is frequent, and saddlery primarily focuses on aesthetics. D. Relying on donkey and mule transport, saddle culture had limited influence.	Transportation	English	Possible Option Letters: B Explanation: Traditional Mongolian saddles mainly reflect two key aspects of Mongolian life: <b>Nomadic lifestyle</b> Mongolians have long lived as nomads on the grasslands and relied heavily on horses for moving, herding animals, and daily travel. Their saddles are strong, stable, and suitable for long-distance riding across rough terrain. <b>Military needs</b> The Mongols were famous warriors. Horses were essential in battle, so the saddles were designed to help riders stay steady and fight while riding, especially for shooting arrows on horseback. Why the other options are wrong: Option A: The saddle reflects a mobile, horse-based lifestyle, not mainly agriculture. Option C: It is about practicality and function, not urban trade or decoration. Option D: Mongolians mainly relied on horses, not donkeys or mules. So Option B is the most accurate choice.
Уламжлалт эмээл нь Монгол ард түмний үйлдвэрлэл, өдөр тутмын амьдралын ямар шинж чанаруудыг тусгасан байдаг вэ? А.Хөдөө аж ахуй нь голлон, морин тоног төхөөрөмж нь ихэвчлэн талбайн ажилд ашигладаг. Б. Нүүдэлчин амьдрал ба цэргийн хэрэгээ тэнцүү эрхэмлэгддэг. С. Хотын худалдаа ихтэй, эмээл гой сайхныг эрхэмлэдэг. Д. Илжиг, пуусанд тулгуурласан тээвэр, эмээлний соёлын нөлөө хягаарлагддаг.	Transportation	Mongolian	Буцах боломжтой сонголтууд: А, В, D Тайлбарлал: А. Хөдөө аж ахуй нь голлон, морин тоног төхөөрөмж нь ихэвчлэн талбайн ажилд ашиглагддаг. * Уламжлалт эмээл нь Монгол ард түмний үйлдвэрлэл нь хөдөө аж ахуйг эрхэмлэдэг байсан бөгөөд энэ нь талбайн ажил болон морины ашиглалтан хөдөлмөрөөр илэрдэг. Морин тоног төхөөрөмжийг талбайн ажилд голчлон ашигладаг. Зарин тохиолдод морийн нүүцдэл, өдөр тутмын хөдөлгөөн, аж ахуйн бусад үйл ажиллагатай холбогтой ажилд ашигладаг. Б. Нүүдэлчин амьдрал ба цэргийн хэрэгээ тэнцүү эрхэмлэгддэг * Монгол ард түмний үйлдвэрлэл нь нүүдэгчийн амьдралдаа тулгуурладаг. Нүүдэлчин амьдрал нь тал нутгаар нүүч, мал малах, байршилаа байгаа верчлэх заргатай холбогтой. Үүнтэй зэрэгцэн, Монголчууд цэргийн хэрэгээг маш иңдэрт тавьдаг байсан. Цэргийн хэрэгээ нь нүүдэлчин амьдралтай салшгүй холбогтой бөгөөд энэ нь Монгол ард түмний үйлдвэрлэлийн үндсэн шинж чанаруудын нэг юм. Д. Илжиг, пуусанд тулгуурласан тээвэр, эмээлний соёлын нөлөө хягаарлагддаг * Монгол ард түмний үйлдвэрлэлд илжиг, пуусанд тулгуурласан тээврийг ашиглаж явдал байсан ч энэ нь үндсэн гол хэсэг нь биш юм. Илжиг, пуусын тээвэр нь тодорхой нөхцөлд ашиглагддаг ч мөртой харьцуулхад хягаарлагддаг нөлөөтэй. Энэ нь эмээлний соёл Монголын нүүдэгчин, морин соёлд илүү гүн шингэнсийг харуулдаг. Сонголт: А, В, D

2090  
2091  
2092  
2093  
2094  
2095  
2096  
2097  
2098  
2099  
2100  
2101  
2102  
2103  
2104  
2105

Figure 22: **Language and Cultural Bias Cases.** The first and second examples are the English and Chinese versions of the same question; the third and fourth examples are the English and Mongolian versions of the same question.

Table 21: Model outputs across different modes for the same question.

Mode	Answer
Rephrase VQA baseline	<p>Answer: B. Yan Liang. Explanation: In historical records and the classic Chinese novel <i>Romance of the Three Kingdoms</i>, Guan Yu is known for his legendary bravery and martial prowess. One of his most famous feats occurred during the Battle of Hulao Gate (虎牢关之战) in 194 CE, which took place during the late Eastern Han dynasty. Here's the context: <b>Hua Xiong</b>: During the Battle of Hulao Gate, Hua Xiong was a general serving under Dong Zhuo, a powerful warlord who controlled the Han imperial court. Hua Xiong challenged the coalition forces led by Yuan Shu, Yuan Shao, and others. Guan Yu volunteered to fight Hua Xiong and defeated him in battle. This event is well-known as one of Guan Yu's early displays of martial skill. <b>Yan Liang and Wen Chou</b>: After defeating Hua Xiong, Guan Yu continued to distinguish himself in battle. He later encountered two other generals, Yan Liang and Wen Chou, who were also serving under Dong Zhuo. In the same battle, Guan Yu fought and killed both Yan Liang and Wen Chou. This event further solidified his reputation as a formidable warrior. <b>Yu Jin</b>: Yu Jin was a general who served under Cao Cao, a rival of Guan Yu's lord, Liu Bei. While Yu Jin and Guan Yu did face each other in battle, Yu Jin was not killed by Guan Yu. Instead, Yu Jin surrendered to Guan Yu during the Battle of Fancheng (樊城之战) in 219 CE, after being trapped by water flooding caused by Guan Yu. Why the Correct Answer is B (Yan Liang): Yan Liang was explicitly mentioned in historical records and the <i>Romance of the Three Kingdoms</i> as one of the generals killed by Guan Yu during the Battle of Hulao Gate. The other options do not align with the historical context: Hua Xiong was defeated but not killed by Guan Yu in the same battle. Wen Chou was also killed by Guan Yu, but the question asks for a single correct answer, and Yan Liang is more commonly associated with this specific event. Yu Jin was not killed by Guan Yu; he surrendered instead. Thus, the correct answer is B. Yan Liang.</p>
VQA baseline	<p>The image depicts a statue of Guan Yu, one of the most famous and revered generals from the Three Kingdoms period in Chinese history. Guan Yu is known for his loyalty, bravery, and righteousness, and he is often associated with the Shu Han kingdom. Key Points to Consider: Guan Yu's Achievements: Guan Yu is celebrated for his military prowess and is known for several significant victories during the Three Kingdoms era. <b>Hua Xiong</b> (华雄): This is a key figure in the story of the "Battle of Changban" (or "Battle of Changping"). According to historical records and the novel Romance of the Three Kingdoms, Guan Yu famously defeated Hua Xiong in battle. Hua Xiong was a general serving under Dong Zhuo, and his defeat marked a turning point in the coalition against Dong Zhuo. Analysis of the Options: A. Hua Xiong: This is correct. Hua Xiong was indeed killed by Guan Yu during the Battle of Changban. This event is well-documented in both historical records and the Romance of the Three Kingdoms. B. Yan Liang (颜良): While Guan Yu did kill Yan Liang, this occurred later in his career, during the Battle of Baima (White Horse). However, the question specifically asks about the person in the picture, and the most iconic victory associated with Guan Yu is his defeat of Hua Xiong. C. Wen Chou (文丑): Wen Chou was another general defeated by Guan Yu, but this happened during the battle of Xiapi, which came after the Battle of Changban. D. Yu Jin (于禁): Yu Jin was a general of Cao Wei who surrendered to Guan Yu during the Battle of Fancheng. He was not killed by Guan Yu. Final Answer: The person in the picture (Guan Yu) is most famously known for killing Hua Xiong. Therefore, the correct answer is: A.</p>
Vision-ablated Prefix Replay	<p>The image depicts a statue of Guan Yu, one of the most famous and revered generals from the Three Kingdoms period in Chinese history. Guan Yu is known for his loyalty, bravery, and martial prowess. Analysis of the question: We need to determine which general was killed by Guan Yu based on historical records. <b>Hua Xiong</b> (华雄): Hua Xiong was a general serving under Dong Zhuo during the late Eastern Han dynasty. He was killed by Guan Yu in the Battle of Hulao Gate, as described in the Romance of the Three Kingdoms. This event is one of the most iconic moments in Guan Yu's early career. <b>Yan Liang</b> (颜良): Yan Liang was a general serving under Yuan Shao. He was killed by Guan Yu during the Battle of Baima (白马之战) in 200 CE. This battle was part of the larger conflict between Cao Cao and Yuan Shao. <b>Wen Chou</b> (文丑): Wen Chou was another general serving under Yuan Shao. He was also killed by Guan Yu shortly after the Battle of Baima, during the same campaign against Yuan Shao. <b>Yu Jin</b> (于禁): Yu Jin was a general serving under Cao Cao. He was not killed by Guan Yu but rather surrendered to Guan Yu during the Battle of Fancheng in 219 CE when Guan Yu's forces surrounded him. Yu Jin later returned to Cao Cao's service after Guan Yu's defeat. Correct Answer: Based on the historical context and the question, the general who was killed by Guan Yu is Yan Liang. Final Answer: B.</p>

2160  
 2161  
 2162  
 2163  
 2164  
 2165  
 2166  
 2167  
 2168  
 2169

2170 Table 22: Human–model agreement rates for different models used as LLM-as-judge under different  
 2171 language settings.

Language	GPT-4o	Claude Sonnet 4	Gemini 2.5 Flash
English	<b>98%</b>	<b>98%</b>	<b>96%</b>
Malay	90%	88%	92%
Tamil	76%	80%	72%

2172  
 2173  
 2174  
 2175  
 2176  
 2177  
 2178  
 2179  
 2180  
 2181  
 2182  
 2183  
 2184  
 2185  
 2186  
 2187  
 2188  
 2189  
 2190  
 2191  
 2192  
 2193  
 2194  
 2195  
 2196

2197 Table 23: Sub-question decomposition of multi-step reasoning questions.  
 2198

Multi-Step Question	Sub-Question	Answer for Sub-Question
Which of the following buildings belongs to the same ethnic characteristic architecture as the stilted building? A. Tulou B. Moxiaolou C. Diaofang D. Yaodong	What ethnic group is associated with the stilted building?	The Tuja ethnic group.
	What ethnic group is associated with the Tulou?	Hakka
	What ethnic group is associated with the Moxiaolou?	The Tuja ethnic group.
	...	...

2205  
 2206  
 2207  
 2208  
 2209  
 2210  
 2211  
 2212  
 2213