



CVQA: Culturally-diverse Multilingual Visual Question Answering Benchmark

David Romero, Chenyang Lyu, Haryo Akbarianto Wibowo, Teresa Lynn, Injy Hamed, Aditya Nanda Kishore, Aishik Mandal, Alina Dragonetti, Artem Abzaliev, Atnafu Lambebo Tonja, et al.

► To cite this version:

David Romero, Chenyang Lyu, Haryo Akbarianto Wibowo, Teresa Lynn, Injy Hamed, et al.. CVQA: Culturally-diverse Multilingual Visual Question Answering Benchmark. 38th Conference on Neural Information Processing Systems, Dec 2024, Vancouver, Canada. pp.11479-11505. <hal-04872109>

HAL Id: hal-04872109

<https://hal.science/hal-04872109v1>

Submitted on 17 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

CVQA: Culturally-diverse Multilingual Visual Question Answering Benchmark

David Romero^{**}, Chenyang Lyu^{**}, Haryo Akbarianto Wibowo^{*}, Teresa Lynn, Injy Hamed, Aditya Nanda Kishore, Aishik Mandal, Alina Dragonetti, Artem Abzaliev, Atnafu Lambebo Tonja, Bontu Fufa Balcha, Chenxi Whitehouse, Christian Salamea, Dan John Velasco, David Ifeoluwa Adelani, David Le Meur, Emilio Villa-Cueva, Fajri Koto, Fauzan Farooqui, Frederico Belcavello, Ganzorig Batnasan, Gisela Vallejo, Grainne Caulfield, Guido Ivetta, Haiyue Song, Henok Biadgign Ademteu, Hernán Maina, Holy Lovenia, Israel Abebe Azime, Jan Christian Blaise Cruz, Jay Gala, Jiahui Geng, Jesus-German Ortiz-Barajas, Jinheon Baek, Jocelyn Dunstan, Laura Alonso Alemany, Kumaranage Ravindu Yasas Nagasinghe, Luciana Benotti, Luis Fernando D’Haro, Marcelo Viridiano, Marcos Estecha-Garitagoitia, Maria Camila Buitrago Cabrera, Mario Rodríguez-Cantelar, Mélanie Jouitteau, Mihail Mihaylov, Mohamed Fazli Mohamed Imam, Muhammad Farid Adilazuarda, Munkhjargal Gochoo, Munkh-Erdene Otgonbold, Naome Etori, Olivier Niyomugisha, Paula Mónica Silva, Pranjal Chitale, Raj Dabre, Rendi Chevi, Ruochen Zhang, Ryandito Diandaru, Samuel Cahyawijaya, Santiago Góngora, Soyeong Jeong, Sukannya Purkayastha, Tatsuki Kuribayashi, Thanmay Jayakumar, Tiago Timponi Torrent, Toqeer Ehsan, Vladimir Araujo, Yova Kementchedjhieva, Zara Burzo, Zheng Wei Lim, Zheng Xin Yong, Oana Ignat, Joan Nwatu, Rada Mihalcea, Thamar Solorio^{*}, and Alham Fikri Aji^{*}

^{*}Core Authors (MBZUAI)
www.cvqa-benchmark.org

AFRICA	ASIA	LATIN-AMERICA	EUROPE
🇳🇮 Nigeria - Igbo Category: Traditions / Art / History	🇯🇵 Japan - Japanese Category: Objects / Materials / Clothing	🇦🇷 Argentina - Spanish Category: Cooking and Food	🇫🇷 France - Breton Category: Sports and Recreation
 <p>Kedụ mmemme ndị a na-eme? (Which ceremony are they doing?)</p>	 <p>ここで一番最初に洗うのは体のどこですか? (What part of the body do you wash first here?)</p>	 <p>¿Para qué sirve la pala de hierro usada en este asado? (What is the iron shovel used for in this asado?)</p>	 <p>Petra a vank evit c'hoari ? (What is missing in order to play ?)</p>
a) I gba nkwa (Traditional marriage) ✓ b) Ncheta Omu (Birthday) ✗ c) Emume cheitaincy (Chieftaincy ceremony) ✗ d) Emume iri ji (New yam festival) ✗	a) 口 (Mouth) ✗ b) 足 (Leg) ✗ c) 右手 (Right hand) ✗ d) 左手 (Left hand) ✓	a) Poner las brasas debajo de la parrilla (To put the embers under the grill) ✓ b) Mover la carne (To move the meat) ✗ c) Poner tierra al fuego (To put dirt on the fire) ✗ d) Cortar la carne (To cut the meat) ✗	a) ar bidoc'hig (The piglet) ✗ b) ar c'hronometr (The stop watch) ✗ c) ar mestr (The master) ✓ d) ar pezh (The coin) ✗

Figure 1: We propose CVQA, a large-scale multilingual VQA benchmark, representing the cultures of 28 countries and 26 different languages across 10 diverse categories, comprising 9k samples.

Abstract

Visual Question Answering (VQA) is an important task in multimodal AI, and it is often used to test the ability of vision-language models to understand and reason on knowledge present in both visual and textual data. However, most of the current VQA models use datasets that are primarily focused on English and a few major world languages, with images that are typically Western-centric. While recent efforts have tried to increase the number of languages covered on VQA

^{*}Equal Contribution

datasets, they still lack diversity in low-resource languages. More importantly, although these datasets often extend their linguistic range via translation or some other approaches, they usually keep images the same, resulting in narrow cultural representation. To address these limitations, we construct CVQA², a new Culturally-diverse multilingual Visual Question Answering benchmark, designed to cover a rich set of languages and cultures, where we engage native speakers and cultural experts in the data collection process. As a result, CVQA includes culturally-driven images and questions from across 28 countries on four continents, covering 26 languages with 11 scripts, providing a total of 9k questions. We then benchmark several Multimodal Large Language Models (MLLMs) on CVQA, and show that the dataset is challenging for the current state-of-the-art models. This benchmark can serve as a probing evaluation suite for assessing the cultural capability and bias of multimodal models and hopefully encourage more research efforts toward increasing cultural awareness and linguistic diversity in this field.

1 Introduction

Visual Question Answering (VQA) [2, 43, 50] is a task that requires AI systems to answer textual questions based on a given context image. VQA serves as an essential measure for assessing the understanding and reasoning capabilities of Multimodal Large Language Models (MLLMs) across diverse images and texts. With the rapid development of MLLMs, significant improvements have been observed, including support for multiple languages [12, 5, 27, 45, 53]. However, there is still a lack of VQA benchmarks that capture a diverse set of languages and cultural contexts. Specifically, most VQA benchmarks only cover the English language [2, 33]. While some work has been undertaken on multilingual VQA, it either covers a limited set of popular languages or is producing questions via translation/generation of text from the original Western-centric images, thus failing to capture cultural nuances inherent in different languages [6, 44].

To address these limitations, we propose CVQA: a novel, large-scale, multilingual, culturally nuanced VQA benchmark that includes a diverse set of languages, including many that are underrepresented and understudied. CVQA follows the grassroots crowd-sourcing collaboration approaches taken by Masakhane for Africa [37], NusaCrowd for Indonesia [4], and AI4Bharat for India [20]. In our case, however, we collaborate across communities, rather than within one particular community, in order to maximize cultural and linguistic representation. Consequently, our data consists of 9k questions across 28 countries, covering 26 languages. We also sub-categorize CVQA based on Country-Language pairs, resulting in 33 distinct pairs, which is substantially more extensive than existing VQA benchmarks. Furthermore, each sample in CVQA falls into one of 10 diverse categories (see Table 1) and is annotated and validated by fluent speakers and those familiar with the respective cultures, ensuring high quality and diversity. Lastly, CVQA is written in both English and local languages, enabling us to benchmark multilingual MLLMs and English-only MLLMs.

In this study, we benchmark CVQA across various MLLMs and find that it presents a significant challenge for open MLLMs, which most of the time achieve no more than 50% accuracy. Additionally, we observe a notable degradation in model performance when questions are asked in native languages, particularly those in understudied languages such as Breton from France and Javanese from Indonesia, highlighting a significant gap in understanding multilingual prompts. We further conduct several ablation studies to analyze the models’ performance across different question categories, regions, languages, and image sources. Our contributions can be summarised as follows:

- First, we introduce CVQA, a new culturally diverse multilingual visual question answering dataset consisting of over 9,000 questions from across 28 countries and 26 languages.
- Second, we provide extensive documentation on our process to crowdsource such large dataset across numerous communities, including annotation guidelines.
- Finally, we provide an initial set of evaluations on this benchmark, to serve as a baseline for future research on vision-language models that are culturally diverse.

²<https://huggingface.co/datasets/afaji/cvqa>

We note that efforts to enhance cultural awareness in models are increasingly gaining attention. As such, our work contributes to the growing interest within the community and can encourage further initiatives to broaden the limited world view currently captured by MLLMs.

2 CVQA Data Collection

The construction of our CVQA dataset involved a detailed annotation process that aims at creating a culturally diverse and linguistically comprehensive dataset for Visual Question Answering. It is worth noting that, while defining culture is challenging, we follow Adilazuarda et al. [1] by using common-ground knowledge (e.g., information surrounding local dishes, history, places, etc. that is generally shared by the people within the region) as a proxy of culture. In this section, we now turn to outline the detailed procedures followed during the data collection and annotation phases.

2.1 Dataset Collection Design

Country-Language Pair Subset CVQA is a multilingual, multiple-choice locally-nuanced visual question-answering dataset. The format is similar to commonly used visual QA data such as VQA [2], VQA-2 [13] or GQA [17]. Yet, in contrast to them, we gathered images and created question-answer pairs based on the cultures of various locations. Moreover, for each location, the question-answer pairs were created in their respective local languages, along with parallel English translations. Some languages are shared across different locations (e.g., Mexico-Spanish vs Spain-Spanish), and vice-versa, different languages are shared across the same location (e.g., Indonesia-Indonesian vs Indonesia-Javanese). Therefore, to capture them, we group our CVQA dataset into several subsets based on this Country-Language pair, rather than simply on language or location only.

Annotators To elicit image collectors and annotation contributions to this project, we reached out to our network, which included both linguistic groups and NLP communities. Annotators needed to be fluent speakers of the language in question and be accustomed to the cultures of the locations for which they provided data. To promote data collection, contributors with significant contributions, either by contributing at least 100 validated question-answer pairs and/or managing several subsets, are rewarded as co-authors in this paper. The annotator demographic statistics can be seen in Figure 8, Appendix D. Our annotators are predominantly native speakers, with around 89% residing in the respective country for over 16 years. The age group distribution shows a significant concentration in the 18-30 age bracket, with about one-third female representation. Overall, the demographic profile highlights diversity in terms of age, with high levels of cultural familiarity and language proficiency.

Categories For the categorization of questions of our CVQA dataset, we incorporate 10 diverse categories to ensure a culturally-comprehensive representative set of visual questions, which are shown in Table 1. We mainly adopt the categorization from the OK-VQA dataset [33], with some modifications to fit the theme of our project. Specifically, the categories from the OK-VQA dataset used in our CVQA dataset are 1) to 7). We split the original category of *Geography, History, Language and Culture* into 2 separate categories of 8) and 9). In addition, we added a new category of 10) considering the effect that cultural icons and media have on everyday life.

Table 1: Categories in our Dataset. To save space in some of our results, we might refer them by shorthand version in brackets.

Category
1. Vehicles and Transportation (Vehicles)
2. Cooking and Food (Food)
3. People and Everyday Life (People)
4. Sports and Recreation (Sports)
5. Plants and Animals (Plants & Animals)
6. Objects, Materials, and Clothing (Objects)
7. Brands and Products (Brands)
8. Geography, Buildings, and Landmarks (Geography)
9. Tradition, Art, and History (Tradition)
10. Public Figure and Pop-Culture (Pop Culture)

2.2 Annotation Process

We developed concise annotation guidelines (in English) that are suitable for all Country-Language subset teams. Here we provide an overview of the key steps that annotators followed during the dataset creation process. The full guidelines are provided in Appendix A.

Image Selection and Preparation For each Country-Language pair, annotators were instructed to select images that depict diverse cultural aspects pertinent to their cultural backgrounds among one of




Figure 2: Statistics of the CVQA Benchmark

the 10 categories. We did not enforce balance across categories considering the different variations of cultural knowledge. We strongly recommend that annotators use their own personal images to avoid accidental data leakage from existing online sources. However, we noted that this request was not always possible, since some images are extremely hard to come by (e.g., photos of public figures or landmarks that are far from the annotator’s location). Therefore, we also allowed them to use images from our pre-defined list of open-use licensing sources³. For self-made images, we asked the annotators whether they were willing to make the image available for commercial or research purposes. For images from existing online sources, we applied the original license.

We requested annotators to avoid using sensitive images that would perpetuate stereotypes. In addition, the annotators were also requested to anonymize faces that were not public figures or fictional characters, as well as text that could reveal the answer to the accompanying questions. We also post-processed all images to remove all metadata such as geo-location, device type, and so on.

Question Creation The questions associated with each image had to be culturally relevant and formulated such that they would require the context of the image in order to be answerable. A maximum of three question-answer pairs could be provided for each image. Each question was accompanied by one correct answer and three distractors that were reasonably plausible, yet incorrect, thus forming a multiple-choice format.

While we follow the existing VQA benchmarks in terms of using a multiple-choice format, we are also aware that multiple-choice has some flaws when used to measure a model’s performance [41]. Hence, we made sure that CVQA is also convertible into free-text open-ended QA, by instructing the annotators to ensure that the question would be answerable even without the accompanying multiple choices (i.e., not through a deductive method). Moreover, to accommodate the multilingual aspect of the benchmark, each question-answer pair was created in the local language and manually translated into English.

Annotators were advised to create questions that promoted an understanding and appreciation of different cultures without perpetuating stereotypes. Typical questions ranged from simple identification queries (e.g., “What is the name of this food?”) to more complex ones involving multi-hop reasoning or local common-sense knowledge (e.g., “What is the color of the t-shirt the youngest member of this group is wearing?”).

Annotation Examples and Training The annotation guidelines provided multiple examples of well-formulated questions and answers to help guide annotator efforts (See Appendix A). These examples helped clarify the level of specificity and cultural relevance expected in the annotations. We provided a tutorial to annotators on how to edit and blur sensitive information in the images. To

³common.wikimedia.org, Flickr, GapMinder, Unsplash, Pixabay

confirm understanding, we spot-checked the annotators’ collected data throughout the annotation period and informed them if some of their data did not follow the guidelines.

Validation The last step in the CVQA data creation was the validation process. Each entry was validated by another annotator of the same Country-Language pair. The validators were instructed to ensure that each question followed the guidelines. Based on our spot-checking, common mistakes that we encouraged the validators to check were typos and grammatical mistakes, non-cultural questions, questions that could be answered without the image, as well as incorrectly-sourced images. More information on the annotation platform is provided in Appendix B.

2.3 Data Statistics

To ensure sufficient question variation, we set the minimum number of questions to be included in CVQA to be at least 200 questions per Country-Language subset. In the end, we gathered 9,044 total questions across all subsets. Some statistics of our collected data are shown in Table 2. Our CVQA covers a diverse set of languages and locations spread across the globe. We also capture languages written in various scripts. While Latin is the dominant script (used in 20 Country-Language pairs), the remaining scripts are diverse; covering Arabic, Amharic, Bengali, Chinese, Cyrillic, Hangul, Japanese, Perso-Arabic, Sinhalese, and Tamil. The Country-Language pairs and corresponding scripts are shown in Appendix E. CVQA covers several less commonly studied languages and regions, such as Ireland-Irish, Indonesia-Minangkabau, Indonesia-Javanese, France-Breton, Nigeria-Igbo and Mongolia-Mongolian.

Question distribution across the subset and categories are shown in Figure 2. Whether the image is coming from an external or personal source varies depending on the subset. We also note that the category with the most personal images is *Cooking and Food*, which we assume is due to the ease of obtaining such images. In contrast, the category with the least amount of personal images is *Public Figures and Pop Culture*, as it is less likely for people to have personal photos under this category.

To investigate the question variations, we categorize each question into question types of “what”, “how”, “why”, “where”, “who”, and “which” questions. We categorize the questions by simple string-matching performed on the English questions. While not perfect, we argue that this method should be able to capture the trend of the questions. As shown in Figure 2, the majority of the questions fall into “what” questions. Question distribution across different Country-Language pairs varies, with an interesting finding that India-Bengali has a lot of “how” questions. Across categories, perhaps unsurprisingly, the *Geography and Landmark* category has noticeably more “where” and “which” (e.g., in which city) questions, whereas the *Public Figure and Pop Culture* category has more “who” questions. By looking at the most frequently used words (Figure 7) across each category, we can see the general theme of the types of questions being asked. For example, questions in the *Cooking and Food* category often enquire about dish names, ingredients, or tastes.

Table 2: CVQA Data Statistics

No. of images	4,560
No. of questions	9,044
No. of countries	28
No. of languages	26
No. of Country-Language pairs	33
Avg. question per image	1.98
Avg. words per question	10.2
Avg. words per option	1.90

3 Experimental Setup

Models. To evaluate performance on our CVQA benchmark, we select a range of multimodal vision-language models with multilingual and monolingual English-only capabilities. For monolingual English-only models, we test CLIP [40] a contrastive-learning-based model, trained with approximately 400 million images and English-only text pairs from the web, where we use its *vit-large-patch14-336* version. We also use InstructBLIP(4.1B) [8], an English-only instruction-aware vision model based on BLIP-2 [24], trained with 13 held-in datasets covering different tasks in English. For multilingual models, we evaluate LLaVA-1.5 (7B) [22] based on Llama-2 [46], and mBLIP [12] a BLIP-2 based model that covers 96 languages (where we evaluate two model variations, mBLIP mT0-XL (4.9B) and mBLIP BLOOMZ (8.3B)). Lastly, we employ M-CLIP [5] a multilingual CLIP-based model that supports 68 languages, where we use its *XLNet-Roberta-Large-Vit-B-32*

version. We also evaluate the most advanced closed-source MLLMs, such as GPT-4o [36] and Gemini-1.5-Flash [45].

Evaluation Framework. We perform a zero-shot evaluation with two types of prompts, as follows: a location-aware prompt, which specifies the country, the question, and the options, (e.g., “*Location: {country}. Question: {question} Options: {options} Short Answer:*”); and a location-agnostic prompt, which follows the same template but does not specify the country in the prompt (e.g., “*Question: {question} Options: {options} Short Answer:*”). Additionally, due to the multilingual nature of CVQA, for each prompt, we evaluate using the English-only and local language question-option pairs. For the generative-based models, LLaVA, mBLIP and InstructBLIP, the image and the prompts are used as the input. The models then produce output probabilities and we treat the highest probability for the options (A,B,C,D) as the prediction (following MMLU [16]). On the other hand, for embedding-based models like CLIP and M-CLIP, we use the embedding-level similarity between the image and the combination of question and each answer candidate texts (*Question+Option-1,...,Question+Option-4*) to select the one with the highest similarity as the correct answer. We use accuracy to measure the performance, following the existing multiple-choice VQA tasks [2, 55].

4 Results

In this section, we discuss the performance of existing MLLMs on the CVQA benchmark.

Table 3: Average performance of MLLMs on our CVQA dataset with English prompts (EN) and local language prompts (LOC).

LLaVA-1.5-7B		M-CLIP		CLIP		mBLIP-mT0		mBLIP-BLOOMZ		InstructBLIP		Gemini-1.5-Flash		GPT-4o	
EN	LOC	EN	LOC	EN	LOC	EN	LOC	EN	LOC	EN	LOC	EN	LOC	EN	LOC
48.9	36.5	36.9	33.4	42.3	31.5	31.5	31.2	39.4	33.0	47.8	32.7	67.1	67.2	79.1	79.4

Table 4: LLaVA-1.5-7B and InstructBLIP results on various VQA datasets, where the results on the other datasets are taken from Liu et al. [26].

Model	VQA2 [13]	GQA [17]	VizWiz [15]	SciQA-IMG [28]	TextVQA [43]	CVQA (EN)	CVQA (LOC)
LLaVA-1.5-7B	78.5	62.0	50.0	66.8	58.2	48.9	36.5
InstructBLIP	-	49.2	34.5	60.5	50.1	47.8	32.7

Main Results The overall performance on our CVQA dataset of various open and closed-source MLLMs are shown in Table 3. Among open models, LLaVA-1.5-7B achieves the best performance, but still significantly lagging behind closed models by more than 10%. However, Table 4 shows that LLaVA-1.5-7B indeed achieves better performance on other established English VQA benchmarks, highlighting that culturally-specific questions that we collect in CVQA are challenging even for the best-performing open model (LLaVA-1.5-7B). The performance is even worse when the question is asked in local languages, emphasizing the models’ lower capability in handling non-English prompts.

The experimental results also highlight a substantial performance gap between open and closed-source MLLMs. Closed models like GPT-4o and Gemini-1.5-Flash demonstrate superior performance, with GPT-4o achieving the highest accuracy in both English (79.1%) and local language (79.4%) prompts. In contrast, open models like InstructBLIP and mBLIP-mT0 exhibit lower performance, particularly in local language prompts, indicating a need for more diverse training data and refined fine-tuning processes. While proprietary models show superior performance, it is hard to fully explain why, due to their closed nature. Additionally, their results are not reproducible. Therefore, we use open models in the rest of our experiments.

Performance per Country-Language. To see the capability of MLLMs in solving questions for each country and language, we report accuracy performance for Country-Language pairs in Figure 3. From this, we observe that all models struggle with questions in local languages, demonstrating the challenges for current MLLMs. In other words, across all models, their performance drops in local language questions compared to their performance in English questions. For instance, in the case of Brazil-Portuguese, LLaVA-1.5-7B achieved a score of 60.73% for English and 51.16% for

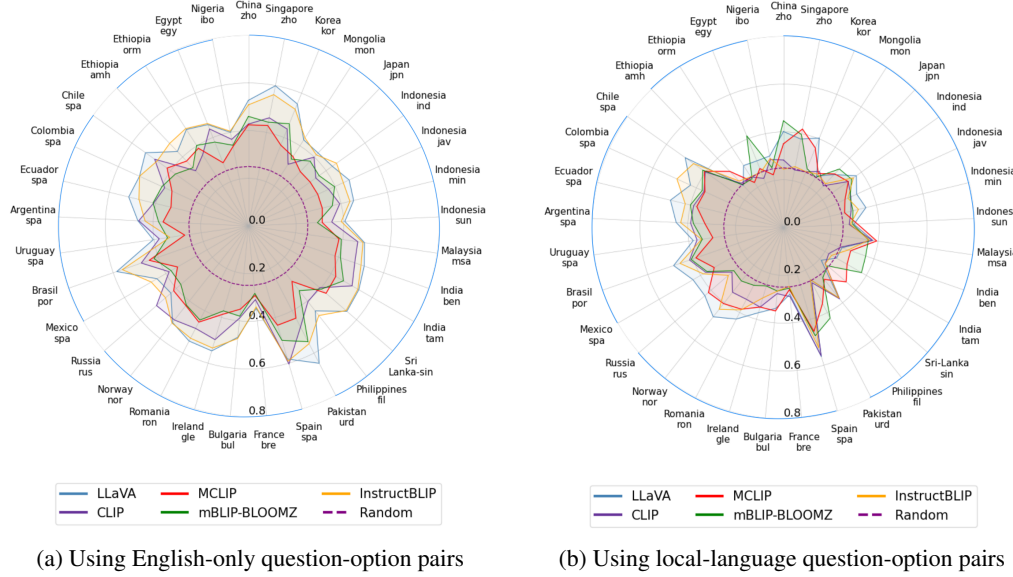


Figure 3: Model performance per Country-Language pair. The blue lines indicate separation by continent. All models show similar behaviour in the majority of cases, despite having different sizes.

Portuguese. Moreover, in Mongolia-Mongolian, all models struggled, with LLaVA-1.5-7B reaching only 40% for English and 27.62% for Mongolian, suggesting challenges in less resource-rich language environments. It is worth noting that, these multilingual MLLMs do not originally support some of the languages, which also explains the significant performance drop for those languages. In contrast, in languages that are more frequently studied in NLP and have more abundant training resources, the performance gap between English and local languages, such as Spanish, tends to be smaller [3].

Performance Across Categories. We show the breakdown performances of models per category in Table 5, where the categories themselves are described in Section 2.1. Note that the category *People and Everyday Life* consistently achieves the highest accuracy across most models, with LLaVA-1.5-7B obtaining 58.6% in English prompts. This can be possibly attributed to the extensive training data available for everyday human activity and interaction, which widely existed in many visual-related datasets. Conversely, the *Cooking & Food* and *Pop Culture* categories exhibit lower accuracy across models, especially in local language prompts. This demonstrates that the high diversity in food and pop culture across different cultures poses a great challenge for the generalization of MLLMs.

Table 5: Accuracy of models across categories. Per category, the best performing models on English (EN) and local language (LOC) question-option pairs are bolded and underlined, respectively.

Categories	LLaVA-1.5-7B		M-CLIP		CLIP		mBLIP-mT0		mBLIP-BLOOMZ		InstructBLIP	
	EN	LOC	EN	LOC	EN	LOC	EN	LOC	EN	LOC	EN	LOC
Brands	47.7	<u>36.5</u>	37.8	35.4	40.1	31.3	33.5	31.3	40.3	36.5	46.9	32.2
Food	45.9	<u>32.5</u>	33.1	29.1	40.8	30.9	27.8	26.9	37.4	30.7	43.8	32.0
Geography	46.1	<u>37.9</u>	36.5	35.6	41.0	32.7	31.1	32.6	35.1	32.6	43.6	32.5
Objects	50.5	<u>34.7</u>	36.4	32.7	39.0	24.2	32.4	30.9	42.3	31.8	52.0	31.6
People	58.6	<u>38.7</u>	45.7	38.2	44.9	36.2	34.4	34.1	47.1	37.6	58.3	38.1
Plants & Animals	54.8	<u>39.4</u>	40.6	31.9	45.9	27.9	35.4	35.9	46.4	38.2	56.8	34.8
Pop Culture	43.5	36.7	32.6	31.0	44.5	<u>37.1</u>	30.6	31.2	35.8	30.7	45.0	35.5
Sports	52.0	<u>40.0</u>	40.2	32.6	45.7	34.6	33.3	30.7	41.8	34.4	49.7	36.5
Tradition	48.8	<u>37.1</u>	36.6	34.9	42.4	31.8	31.6	32.1	39.0	32.9	46.5	30.5
Vehicles	50.9	<u>43.8</u>	38.4	39.4	40.5	32.6	37.4	35.7	42.1	33.3	52.6	33.3

Impact of External Image Source. The performance of various models on self-made versus web images is shown in Table 6. One of the interesting findings is the performance variability across image sources for different models. For self-made images, the accuracy of some models such as LLaVA-1.5-7B and CLIP tends to be lower compared to web images. For instance, LLaVA-1.5-7B

Table 6: Accuracy of different models divided by image source

Image Source	LLaVA-1.5-7B		M-CLIP		CLIP		mBLIP-mT0		mBLIP-BLOOMZ		InstructBLIP	
	EN	LOC	EN	LOC	EN	LOC	EN	LOC	EN	LOC	EN	LOC
Self-made Image	47.0	34.0	37.9	34.0	41.5	30.4	34.4	31.9	40.5	34.2	49.1	31.1
Web Image	48.0	37.0	36.8	33.6	44.1	33.1	34.0	32.8	39.6	33.5	46.7	33.1

Table 7: Location-aware and location-agnostic results

Prompt type	LLaVA-1.5-7B		M-CLIP		CLIP		mBLIP-mT0		mBLIP-BLOOMZ		InstructBLIP	
	EN	LOC	EN	LOC	EN	LOC	EN	LOC	EN	LOC	EN	LOC
Location-aware	48.9	36.5	36.9	33.4	42.3	31.5	31.5	31.2	39.4	33.0	47.8	32.7
Location-agnostic	47.7	35.4	37.5	33.8	43.1	31.5	34.2	32.2	39.8	33.6	47.5	32.1

achieves a 47.0% accuracy in English prompts on self-made images but slightly higher at 48.0% on web images. CLIP shows an accuracy of 44.1% in English prompts on web images compared to 41.5% on self-made images. While this trend is not consistent across the other models, the results still indicate that web images might be more representative of the data these models (such as LLaVA-1.5-7B and CLIP) were trained on, leading to better performance.

Location-Aware vs Location-Agnostic Prompt. The performance of the various models on location-aware versus location-agnostic prompts is shown in Table 7. While the inclusion of location information has a varied impact on different models, the overall difference between both prompt options is marginal, suggesting no significant effect of including location information on MLLMs.

Performance without Multiple Choice Options. Most of the evaluations we conduct on CVQA are under a multiple-choice setting. However, the multiple-choice setting is often brittle towards option ordering [38, 54], and not very natural with respect to real-world scenarios [30]. In this paragraph, we explore the model’s performance on CVQA in an open-ended QA setting. To evaluate in this setting, we prompt the models without giving them the options (e.g., “In which city is this monument located?”). Then, the answer is selected by choosing the model’s highest probability of generating the full answer phrase of one of the options [11] (e.g., Jakarta, Bandung, Bali, Surabaya). This way, it is robust towards ordering unlike predicting the answer letter (e.g., A), while also not giving the model multiple-choice options that can be indirectly used for deductive reasoning. Our result shows that LLaVA-1.5-7B achieved a noticeable performance drop when prompted without multiple choice, from 48.9% to just 30% average performance. This notes that in a more practical scenario, these models might be even more unreliable in cultural understanding.

5 Limitations

Our new benchmark dataset represents a diverse worldview through the inclusion of different languages and regions not covered in previous datasets. But we acknowledge that even CVQA is not comprehensive, as it covers only a fraction of the world’s languages and regions. CVQA also lacks an English-centric baseline, which could arguably provide an interesting comparison with the rest of the regions. Additionally, our data scale prevents using CVQA to train new models, limiting its use for benchmarking purposes only.

We note that each region has different characteristics of questions and difficulty—some regions are more likely to provide simpler, identity “what is” questions, whereas other regions might use questions that require deeper cultural knowledge. Therefore, comparing performance across languages/countries might not always be fair.

Culture is hard to define, and our CVQA ultimately serves only as a proxy to benchmark the model’s understanding of culture through local common knowledge. However, this by no means captures all cultural nuances [1]. Additionally, our location granularity captures country-level cultural knowledge. However, it might be interesting to capture cultural awareness at a more granular level, such as city-level, since each city might have variations in cultural common knowledge. Similarly, other demographic factors such as age might play a role in common knowledge.

In this section we discussed the following aspects: 1) the fact that this dataset cannot be considered as a comprehensive representation of the world languages and regions; 2) the different levels of question complexity; 3) a bounded definition of culture. While these limitations might be relevant, we consider them as plausible lines for future work and outside the scope of this initial effort.

6 Related Work

Substantial progress has been made in recent years on both datasets and methodologies for VQA [42, 23]. Since the introduction of early open-ended VQA datasets [2, 13], various formats like multiple-choice [49, 28], span extraction [34], and free-text generation [27] have been developed. Among these, multiple-choice datasets [28, 29, 52] are the most commonly used, likely due to their simplicity in evaluation and comparison. The development of these datasets has significantly accelerated research progress, serving as both training data and testbeds, especially the recently introduced ScienceQA [28] and MathVista [29] designed for evaluating MLLMs. The evolution of VQA methodologies has been revolutionary, transitioning from statistical machine learning [31] to neural-based methods [32, 17, 40], and advanced MLLMs [27, 36, 45] trained on massive multimodal data. Early VQA systems often required supervised learning and were limited to specific domains, but recent models like CLIP [40], LLaVA [27], and GPT-4V [36] are capable of zero-shot or few-shot learning, demonstrating strong performance. Despite this progress, significant limitations remain. Most VQA datasets focus primarily on English and a few major world languages [28, 29, 51], leading to language bias and under-representation of many languages and cultures. Additionally, the images in these datasets predominantly reflect Western scenes and styles, lacking the diversity needed to represent real-world scenarios across different cultures [7].

Some efforts have been made to create multilingual VQA datasets, such as FM-IQA [10], MCVQA [14], xGQA [39], MaXM [6], MTVQA [44], and MaRVL [25]. However, these datasets are still limited in terms of the number of languages and the cultural diversity of the images and questions, or being a translation of existing English data. On the other hand, there have been initiatives to create culturally-diverse datasets and benchmarks under text-only modality [35, 19, 48, 18, 9, 47, 21]. Our proposed benchmark aims to fill the gap that covers both textual and visual modality by creating a large-scale, culturally-and-linguistically diverse dataset that will enable the development of more inclusive and robust VQA models.

7 Conclusion

We proposed CVQA, a novel, human-written visual QA benchmark dataset that captures cultural nuances across a diverse set of languages and locations. CVQA encompasses 10 question categories, with each question written in both English and the native language. This allowed us to benchmark both multilingual visual models and English-only models. We provided insights into our dataset’s question types and commonly used terms for each category.

We then performed benchmarks on various visual models, including both multilingual and English-only models. Our benchmark demonstrated that CVQA presented challenges for open-source models. These models generally performed worse when queried in local languages compared to English, indicating poorer performance in handling multilingual queries. The performance is also considerably lower when we do not provide the multiple choice setting, which is a more realistic use case for this technology. We hope that publishing CVQA encourages the AI community to pay more attention to non-English-centric models and benchmarking, thereby advancing progress in multilingual, multimodal research.

References

- [1] M. F. Adilazuarda, S. Mukherjee, P. Lavania, S. Singh, A. Dwivedi, A. F. Aji, J. O’Neill, A. Modi, and M. Choudhury. Towards measuring and modeling "culture" in LLMs: A survey. *arXiv preprint arXiv:2403.15412*, 2024.
- [2] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh. VQA: Visual Question Answering. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 2425–2433, 2015.

- [3] Y. Bang, S. Cahyawijaya, N. Lee, W. Dai, D. Su, B. Wilie, H. Lovenia, Z. Ji, T. Yu, W. Chung, et al. A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*, 2023.
- [4] S. Cahyawijaya, H. Lovenia, A. F. Aji, G. Winata, B. Wilie, F. Koto, R. Mahendra, C. Wibisono, A. Romadhony, K. Vincentio, J. Santoso, D. Moeljadi, C. Wirawan, F. Hudi, M. S. Wicaksono, I. Parmonangan, I. Alfina, I. F. Putra, S. Rahmadani, Y. Oenang, A. Septiandri, J. Jaya, K. Dhole, A. Suryani, R. A. Putri, D. Su, K. Stevens, M. N. Nityasya, M. Adilazuarda, R. Hadiwijaya, R. Diandaru, T. Yu, V. Ghifari, W. Dai, Y. Xu, D. Damapusita, H. Wibowo, C. Tho, I. Karo Karo, T. Fatyanosa, Z. Ji, G. Neubig, T. Baldwin, S. Ruder, P. Fung, H. Sujaini, S. Sakti, and A. Purwarianti. NusaCrowd: Open source initiative for Indonesian NLP resources. In A. Rogers, J. Boyd-Graber, and N. Okazaki, editors, *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13745–13818, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.868. URL <https://aclanthology.org/2023.findings-acl.868>.
- [5] F. Carlsson, P. Eisen, F. Rekathati, and M. Sahlgren. Cross-lingual and multilingual CLIP. In *Proceedings of the 13th Language Resources and Evaluation Conference (LREC)*, pages 6848–6854, 2022.
- [6] S. Changpinyo, L. Xue, M. Yarom, A. Thapliyal, I. Szepkator, J. Amelot, X. Chen, and R. Soricut. MaXM: Towards multilingual visual question answering. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2667–2682, 2023.
- [7] X. Chen, H. Fang, T.-Y. Lin, R. Vedantam, S. Gupta, P. Dollár, and C. L. Zitnick. Microsoft COCO captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015.
- [8] W. Dai, J. Li, D. Li, A. M. H. Tiong, J. Zhao, W. Wang, B. Li, P. Fung, and S. Hoi. InstructBLIP: Towards general-purpose vision-language models with instruction tuning. In *Proceedings of NeurIPS*, 2023.
- [9] A. Dwivedi, P. Lavania, and A. Modi. Elicor: Corpus for analyzing llms for etiquettes. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6921–6931, 2023.
- [10] H. Gao, J. Mao, J. Zhou, Z. Huang, L. Wang, and W. Xu. Are you talking to a machine? dataset and methods for multilingual image question answering. In *Proceedings of NeurIPS*, 2015.
- [11] L. Gao, J. Tow, S. Biderman, S. Black, A. DiPofi, C. Foster, L. Golding, J. Hsu, K. McDonnell, N. Muenighoff, J. Phang, L. Reynolds, E. Tang, A. Thite, B. Wang, K. Wang, and A. Zou. A framework for few-shot language model evaluation, 2021. URL <https://doi.org/10.5281/zenodo.5371628>.
- [12] G. Geigle, A. Jain, R. Timofte, and G. Glavaš. mBLIP: Efficient bootstrapping of multilingual vision-LLMs. *arXiv preprint arXiv:2307.06930*, 2023.
- [13] Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh. Making the V in VQA matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6904–6913, 2017.
- [14] D. Gupta, P. Lenka, A. Ekbal, and P. Bhattacharyya. A unified framework for multilingual and code-mixed visual question answering. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 900–913, 2020.
- [15] D. Gurari, Q. Li, A. J. Stangl, A. Guo, C. Lin, K. Grauman, J. Luo, and J. P. Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3608–3617, 2018.
- [16] D. Hendrycks, C. Burns, S. Basart, A. Zou, M. Mazeika, D. Song, and J. Steinhardt. Measuring massive multitask language understanding. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.
- [17] D. A. Hudson and C. D. Manning. GQA: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6700–6709, 2019.
- [18] A. Jha, A. M. Davani, C. K. Reddy, S. Dave, V. Prabhakaran, and S. Dev. Seegull: A stereotype benchmark with broad geo-cultural coverage leveraging generative models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9851–9870, 2023.
- [19] A. Kabra, E. Liu, S. Khanuja, A. F. Aji, G. Winata, S. Cahyawijaya, A. Aremu, P. Ogayo, and G. Neubig. Multi-lingual and multi-cultural figurative language understanding. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8269–8284, 2023.

- [20] A. Kunchukuttan, D. Kakwani, S. Golla, G. N. C., A. Bhattacharyya, M. M. Khapra, and P. Kumar. Ai4bharat-indicnlp corpus: Monolingual corpora and word embeddings for indic languages, 2020.
- [21] W. Q. Leong, J. G. Ngui, Y. Susanto, H. Rengarajan, K. Sarveswaran, and W. C. Tjhi. Bhasa: A holistic southeast asian linguistic and cultural evaluation suite for large language models, 2023.
- [22] C. Li, C. Wong, S. Zhang, N. Usuyama, H. Liu, J. Yang, T. Naumann, H. Poon, and J. Gao. LLaVA-Med: Training a large language-and-vision assistant for biomedicine in one day. In *Proceedings of NeurIPS*, 2023.
- [23] C. Li, Z. Gan, Z. Yang, J. Yang, L. Li, L. Wang, and J. Gao. Multimodal foundation models: From specialists to general-purpose assistants. *Foundations and Trends® in Computer Graphics and Vision*, 16 (1-2):1–214, 2024.
- [24] J. Li, D. Li, S. Savarese, and S. Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, pages 19730–19742, 2023.
- [25] F. Liu, E. Bugliarello, E. M. Ponti, S. Reddy, N. Collier, and D. Elliott. Visually grounded reasoning across languages and cultures. In M.-F. Moens, X. Huang, L. Specia, and S. W.-t. Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10467–10485, Online and Punta Cana, Dominican Republic, Nov. 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.818. URL <https://aclanthology.org/2021.emnlp-main.818>.
- [26] H. Liu, C. Li, Y. Li, and Y. J. Lee. Improved baselines with visual instruction tuning. In *Proceedings of NeurIPS 2023 Workshop on Instruction Tuning and Instruction Following*, 2023.
- [27] H. Liu, C. Li, Q. Wu, and Y. J. Lee. Visual instruction tuning. In *Proceedings of NeurIPS*, 2023.
- [28] P. Lu, S. Mishra, T. Xia, L. Qiu, K.-W. Chang, S.-C. Zhu, O. Tafjord, P. Clark, and A. Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *Proceedings of NeurIPS*, 2022.
- [29] P. Lu, H. Bansal, T. Xia, J. Liu, C. Li, H. Hajishirzi, H. Cheng, K.-W. Chang, M. Galley, and J. Gao. MathVista: Evaluating mathematical reasoning of foundation models in visual contexts. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*, 2023.
- [30] C. Lyu, M. Wu, and A. F. Aji. Beyond probabilities: Unveiling the misalignment in evaluating large language models. *arXiv preprint arXiv:2402.13887*, 2024.
- [31] M. Malinowski and M. Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *Proceedings of NeurIPS*, 2014.
- [32] M. Malinowski, M. Rohrbach, and M. Fritz. Ask your neurons: A neural-based approach to answering questions about images. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1–9, 2015.
- [33] K. Marino, M. Rastegari, A. Farhadi, and R. Mottaghi. OK-VQA: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3195–3204, 2019.
- [34] M. Mathew, V. Bagal, R. P. Tito, D. Karatzas, E. Valveny, and C. V. Jawahar. InfographicVQA. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1697–1706, 2022.
- [35] A. Mukherjee, C. Raj, Z. Zhu, and A. Anastasopoulos. Global voices, local biases: Socio-cultural prejudices across languages. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15828–15845, 2023.
- [36] OpenAI. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*, 2023.
- [37] I. Orife, J. Kreutzer, B. Sibanda, D. Whitenack, K. Siminyu, L. Martinus, J. T. Ali, J. Abbott, V. Marivate, S. Kabongo, M. Meressa, E. Murhabazi, O. Ahia, E. van Biljon, A. Ramkilowan, A. Akinfaderin, A. Öktem, W. Akin, G. Kioko, K. Degila, H. Kamper, B. Dossou, C. Emezue, K. Ogueji, and A. Bashir. Masakhane – machine translation for africa, 2020.
- [38] P. Pezeshkpour and E. Hruschka. Large language models sensitivity to the order of options in multiple-choice questions. *arXiv preprint arXiv:2308.11483*, 2023.

- [39] J. Pfeiffer, G. Geigle, A. Kamath, J.-M. Steitz, S. Roth, I. Vulić, and I. Gurevych. xGQA: Cross-lingual visual question answering. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2497–2511, 2022.
- [40] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 8748–8763, 2021.
- [41] J. Robinson and D. Wingate. Leveraging large language models for multiple choice question answering. In *Proceedings of the 11th International Conference on Learning Representations (ICLR)*, 2023.
- [42] H. Sharma and A. S. Jalal. A survey of methods, datasets and evaluation metrics for visual question answering. *Image and Vision Computing*, 116:104327, 2021.
- [43] A. Singh, V. Natarajan, M. Shah, Y. Jiang, X. Chen, D. Batra, D. Parikh, and M. Rohrbach. Towards VQA models that can read. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8317–8326, 2019.
- [44] J. Tang, Q. Liu, Y. Ye, J. Lu, S. Wei, C. Lin, W. Li, M. F. F. B. Mahmood, H. Feng, Z. Zhao, Y. Wang, Y. Liu, H. Liu, X. Bai, and C. Huang. MTVQA: Benchmarking multilingual text-centric visual question answering. *arXiv preprint arXiv:2405.11985*, 2024.
- [45] G. Team. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- [46] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. C. Ferrer, M. Chen, G. Cucurull, D. Esiobu, J. Fernandes, J. Fu, W. Fu, B. Fuller, C. Gao, V. Goswami, N. Goyal, A. Hartshorn, S. Hosseini, R. Hou, H. Inan, M. Kardas, V. Kerkez, M. Khabsa, I. Kloumann, A. Korenev, P. S. Koura, M.-A. Lachaux, T. Lavril, J. Lee, D. Liskovich, Y. Lu, Y. Mao, X. Martinet, T. Mihaylov, P. Mishra, I. Molybog, Y. Nie, A. Poulton, J. Reizenstein, R. Rungta, K. Saladi, A. Schelten, R. Silva, E. M. Smith, R. Subramanian, X. E. Tan, B. Tang, R. Taylor, A. Williams, J. X. Kuan, P. Xu, Z. Yan, I. Zarov, Y. Zhang, A. Fan, M. Kambadur, S. Narang, A. Rodriguez, R. Stojnic, S. Edunov, and T. Scialom. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [47] B. Wang, Z. Liu, X. Huang, F. Jiao, Y. Ding, A. Aw, and N. F. Chen. SeaEval for multilingual foundation models: From cross-lingual alignment to cultural reasoning, 2024.
- [48] H. A. Wibowo, E. H. Fuadi, M. N. Nityasya, R. E. Prasajo, and A. F. Aji. Copal-id: Indonesian language reasoning with local culture and nuances. *arXiv preprint arXiv:2311.01012*, 2023.
- [49] L. Xu, H. Huang, and J. Liu. SUTD-TrafficQA: A question answering benchmark and an efficient network for video reasoning over traffic events. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9878–9888, 2021.
- [50] A. Yang, A. Miech, J. Sivic, I. Laptev, and C. Schmid. Just ask: Learning to answer questions from millions of narrated videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1686–1697, 2021.
- [51] W. Yu, Z. Yang, L. Li, J. Wang, K. Lin, Z. Liu, X. Wang, and L. Wang. MM-Vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*, 2023.
- [52] X. Yue, Y. Ni, K. Zhang, T. Zheng, R. Liu, G. Zhang, S. Stevens, D. Jiang, W. Ren, Y. Sun, C. Wei, B. Yu, R. Yuan, R. Sun, M. Yin, B. Zheng, Z. Yang, Y. Liu, W. Huang, H. Sun, Y. Su, and W. Chen. MMMU: A massive multi-discipline multimodal understanding and reasoning benchmark for expert AGI. *arXiv preprint arXiv:2311.16502*, 2023.
- [53] B. Zheng, B. Gou, J. Kil, H. Sun, and Y. Su. GPT-4V(ision) is a generalist web agent, if grounded. *arXiv preprint arXiv:2401.01614*, 2024.
- [54] C. Zheng, H. Zhou, F. Meng, J. Zhou, and M. Huang. Large language models are not robust multiple choice selectors. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*, 2024.
- [55] Y. Zhu, O. Groth, M. Bernstein, and L. Fei-Fei. Visual7W: Grounded Question Answering in Images. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

A Annotation Guidelines

Multilingual Multimodal Visual Question Answering

Benchmark: Annotation Guidelines

Introduction

This document provides guidelines for annotating images and corresponding questions and answers in multiple languages to create a culturally diverse and linguistically comprehensive multimodal QA benchmark.

Objective

To build a benchmark that represents a wide range of cultures and languages, to measure potential bias in visual AI models.

Guidelines for Contributors

Each region and language (eg. Ecuador-Spanish) will be represented by **at most 3 annotators, in which 1 will be the team lead**. Each person is expected to provide at least 100 visual questions to be considered as a co-author. The team lead will still have to provide questions, the only difference is that the team lead is responsible to find and to organize more annotators and will manage to contact and brief that annotator, if needed.

Image Selection:

- Contribute images that represent diverse cultural aspects that represent the specific cultural background you're contributing to. The image must fall into one of the categories below. **Pick one of the most relevant category (more later):**

Image Category *

- | | |
|--|--|
| <input type="checkbox"/> Vehicles and Transportation | <input type="checkbox"/> Brands, products, and companies |
| <input type="checkbox"/> Objects, materials, clothing | <input type="checkbox"/> Sports and recreation |
| <input type="checkbox"/> Cooking and food | <input type="checkbox"/> Traditions, art, and history |
| <input type="checkbox"/> Geography, buildings, and landmarks | <input type="checkbox"/> People and everyday life |
| <input type="checkbox"/> Plants and animal | <input type="checkbox"/> Public Figure and pop culture |
| <input type="checkbox"/> Other | |

- Images should be relevant to your culture/country.
- Ensure that **images are relevant to the questions being posed**. In other words, the image **is needed** to answer the question.

- If the image contains the answer's text, you can blur/crop the image so that the image does not contain the answer.
- Image source:
 - 1. Self/personal picture (**highly preferable**). You may ask your family/friend to donate their photos, if possible.
 - 2. We also accept external images from:
 - Flickr: <https://www.flickr.com/explore> (please make sure the associated license to the image is Creative Commons), this can be selected at the top left of Flickr ("Any License").
 - WikimediaCommons: https://commons.wikimedia.org/wiki/Main_Page (here you do not need to select any license for the images),
 - Unsplash: <https://unsplash.com/> (please make sure to search the image first and they select the license: Free). More details (Tutorial) at the end of this document.
 - Dollar Street: <https://www.gapminder.org/dollar-street> (here you do not need to select any license for the images), this webpage has images only from some countries, please make sure to select your country to find images if applicable.

More detailed instructions for each web page are shown at the end of this document.

- If you use an external image, you'll need to put the url of the original image.
- The image must be reasonable quality (not pixelated or blurry, can be understandable). You can upload images of any ratio as long as it is not too tall or wide (e.g.: don't submit panorama pictures).
- Do not show personally identifiable information (PII) such as faces, car plates, house addresses. Faces of public figures or fictional characters are ok. Also, **please be sure to blur text in the image that will leak the answer.** "PicdeFacer" can be used for blurring: <https://picdefacer.com/en/>. Tutorial on using PicdeFacer is shown at the end of this document.

Question and Answer Creation:

After finding the image, you must now formulate 1-3 questions + answers from that image. Specifically:

- The question must be answerable **only by looking at the image**.
- Ensure that the questions are culturally relevant and specific to the image content.
- Provide answers that are concise, accurate, and directly related to the question.

- You will also need to provide 1 correct option and 3 other incorrect options (distractors). For the distractors, choose options that are relevant, not obvious wrong answers.
- **The question must be answerable even without the multiple-choice.**
Example of the invalid question: (“What song is not performed by this musician” – not answerable if you don’t know the choices)
- Make sure the questions are **written fluently in both the local language and English**. Use a grammar checker if needed i.e. if you are not fluent in English.
- Be mindful of cultural sensitivities and avoid stereotyping or misrepresenting cultural aspects.
- Ensure there are **variations on your question**. Identity questions are fine, eg “What is this”, or “where is this”. But additionally adding more complex/difficult questions would be great. For example, multi-hop reasoning, counting, referencing, or questions that require local commonsense knowledge to be answered.

Category Definition

When selecting a category, pick one of the most relevant. Please follow the guideline:

- **Vehicles and Transportation:** Local public transport, local vehicles.

- **Objects, Materials, Clothing:** Questions about local/traditional clothes. Unique/local tools or items.

- **Cooking and Food:** Local dishes and food/drink. This category includes native fruits in the context of the image if that fruit is served as a food/drink.

- **Geography, Buildings, Landmarks:** Popular/common landmarks, local architecture/buildings. Local monuments.

- **Plants and Animals:** Plants and animals commonly found in the region.

- **Brands, Products, and Companies:** Questions about understanding local yet popular brands or companies. Even if the brand is about food/transportation, if the main focus of the question is the brand recognition itself, then it should be under this category.

- **Sports & Recreation:** Local sports and fun activities. Focuses on the activity itself rather than the location (in that case, it goes to the 'landmark' category).




- **Tradition, Art, History:** Local ceremonies/festivals/events, local dance/music, folklores. Historical artifacts.

- **People & Everyday Life:** Focuses on the people themselves: i.e., common habits/customs, common occupations and jobs, routine religious activities, everyday activities/routines.

- **Public Figures & Pop Culture:** Questions on the understanding of common public figures (e.g., politicians, artists, musicians, etc.). Common pop culture such as movies and games. If the category is still ambiguous to you, pick the one you think is the most appropriate.




Examples




Examples that can be improved




	<p>Make sure the image is needed to respond the question, example:</p> <ol style="list-style-type: none"> 1) ¿En qué mes se celebra esta fiesta? (In which month is this celebration held?) Correct 2) ¿En qué mes se celebra la fiesta de la "Mama Negra"? (In which month is the celebration of the "Mama Negra" held?) Wrong—As this question can be answered without looking at the image.
	<p>Make sure the question is not ambiguous:</p> <ol style="list-style-type: none"> 1) Where is this monument located? Wrong—Not specific, the answer could be a city, country, province, etc. 2) In which city is this monument located? Correct—specifically asking about the city
	<p>Make sure the question is not too vague:</p> <ol style="list-style-type: none"> 1) What is this? Question wording can be more specific 2) What is the name of this vehicle? Correct—specifically asking about the vehicle name.

Acceptable examples


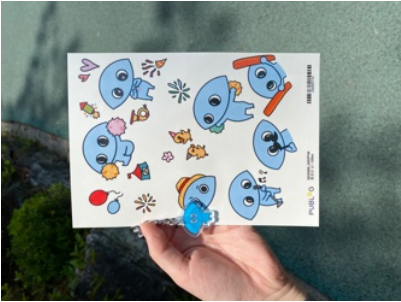

Image	Question
-------	----------

	<p>Category: Tradition / Art / History – Spanish/Mexico</p> <p>¿Qué se muestra en la imagen? (What is shown in the image?)</p> <p>A. el calendario azteca/ piedra del sol (the aztec calendar/ aztec sun stone) B. una serpiente azteca (an aztec serpent) C. coatlicue (coatlicue) D. tlálloc (tlaloc)</p> <p>¿En dónde es exhibe esta pieza? (Where is this piece exhibited?)</p> <p>A. En el museo nacional de antropología (In the National Museum of Atroplogy) B. en el castillo de Chapultepec (In the Chapultepec Castle) C. En el zócalo de la ciudad de Mexico (In the Mexico City zocalo) D. En Teotihuacan (In Teotihuacan)</p>
	<p>Category: Tradition/ Art / History – Igbo/Nigeria</p> <p>Kedu mmemme ndi a na-eme? (Which ceremony are they doing?)</p> <p>A. ! gba nkwu (Traditional marriage) B. Ncheta Omumu (Birthday) C. Emume cheiftaincy (Chieftaincy ceremony) D. Emume iri ji phuru (New yam festival)</p> <p>Kedu ebe a na-eme mmemme a? (Where is this ceremony held?)</p> <p>A. Ulo nna nwunye (The home of the bride's father) B. Ahia (The market) C. Ulo nna di (The home of the groom's father) D. Ulo nsọ (The church)</p>
	<p>Category: Tradition/ Art / History – Indonesian/Indonesia</p> <p>Pada tahun berapakah foto ini diambil? (In what year is this photo taken?)</p> <p>A. 2015 (2015) B. 2020 (2020) C. 2023 (2023) D. 2010 (2010)</p> <p>Apa nama pasukan yang ada di foto ini? (What is the name of the squad in this photo?)</p> <p>A. Paskibraka (Paskibraka) B. Brimob (Brimob) C. TNI (TNI) D. ABRI (ABRI)</p> <p>Apa tugas utama pasukan ini? (What is the main purpose of this squad?)</p> <p>A. Mengibarkan bendera (Hoisting the flag) B. Mengawal presiden (Escorting president)</p>

	<p>C. Menjaga keamanan (Maintaining security) D. Mengiringi pengantin (Accompanying the bride and groom)</p>
	<p>Category: Tradition, Art, History – Sundanese/Indonesia</p> <p>Naon kagunaan ieu hiji alat? (What is the use of this tool?)</p> <p>A. Alat musik (Musical instrument) B. Alat pertahanan diri (Self defence tool) C. Jemuran (Clothes drying equipment) D. Alat masak (Cooking tool) –</p> <p>Ieu hiji alat teh asalna ti propinsi mana di Indonesia? (This tool comes from which province in Indonesia?)</p> <p>A. Jawa Barat (West Java) B. Bali (Bali) C. Bengkulu (Bengkulu) D. Sumatra Barat (West Sumatra)</p>
	<p>Category: Plants and animals – Malay/Malaysia</p> <p>Apakah nama bunga dalam gambar ini? (What is the name of the flower in this picture?)</p> <p>A. Pakma (Rafflesia) B. Bunga raya (Hibiscus) C. Anggerik (Orchid) D. Bunga kertas (Bougainvillea)</p> <p>Di rantau Asia manakah bunga itu boleh ditemui? (In which region of Asia can the flower be found?)</p> <p>A. Asia Tenggara (Southeast Asia) B. Asia Timur (East Asia) C. Asia Selatan (South Asia) D. Asia Tengah (Central Asia)</p>
	<p>Category: People and everyday life - Javanese/Indonesia</p> <p>Opo arane wong seng nang tengah embong iki? (What is the term for the man in the middle of the road?)</p> <p>A. Polisi cepek (Polisi cepek) B. Tukang parkir (Parking assistance man) C. Mijjo (Grocery man) D. Tukang becak (Pedicap man)</p> <p>Opo seng dilakukno wong seng nang tengah dalam iku? (What does the man in the middle of the road do?)</p> <p>A. Ngatur prapatan (Managing the intersection) B. Njaluk donasi (Asking for donations) C. Ngawasi pelanggaran lalu lintas (Looking out for traffic violations) D. Nunjukno arah (Showing directions)</p>

	<p>Category: People and everyday life – Malay/Malaysian</p> <p>Roh manakah yang disembah dengan altar ini? (Which deity is worshiped on this altar?)</p> <p>A. Datuk Gong (Na Tuk Kong) B. Buddha (Buddha) C. Brahma (Brahma) D. Vishnu (Vishnu)</p> <p>Apakah agama yang diamalkan oleh pengguna altar ini? (What religion do the users of these altars practice?)</p> <p>A. Taoism (Taoisme) B. Buddha (Buddhisme) C. Islam (Islam) D. Hindu (Hinduisme)</p>
	<p>Category: People and Everyday Life – Indonesian/Indonesia</p> <p>Apa yang orang-orang ini lakukan? (What are these people doing?)</p> <p>A. Berwudhu (Performing ablution) B. Mandi (Taking a bath) C. Yoga (Yoga) D. Beribadah (Praying)</p> <p>Dimana biasanya orang-orang melakukan aktivitas di foto ini? (Where do people usually do the activity in this photo?)</p> <p>A. Masjid (Mosque) B. Gereja (Church) C. Pemandian umum (Public bath) D. Gym (Gym)</p>
	<p>Category: Cooking and Food – Tagalog/Philippines</p> <p>Anong tawag sa kakanin na ito? (What is the name of this rice cake?)</p> <p>A. Puto Bumbong (Puto Bumbong) B. Suman (Suman) C. Kutsinta (Kutsinta) D. Sapin-Sapin (Sapin-Sapin)</p> <p>Tuwing kailan ito madalas tinitinda sa Pilipinas? (When is this food usually sold in the Philippines?)</p> <p>A. Christmas Season (Christmas Season) B. Independence Day (Independence Day) C. Labor Day (Labor Day) D. National Heroes Day (National Heroes Day)</p> <p>Ano tawag dun sa brown? (What do you call the brown object?)</p>

	<p>A. Muscovado (Muscovado) B. Latik (Toasted coconut) C. Chocolate (Chocolate) D. Caramel (Caramel)</p>
	<p>Category: Object, Clothing, and Material – Korean/South Korea</p> <p>이런 종류의 요리에 사용되는 그릇을 무엇이라고 부르나요? (What is this type of bowl called in cooking?)</p> <p>A. 돌솥 (Dolsot) B. 복주머니 (Bokjumeoni) C. 냄비 (Pot) D. 팬 (Pan)</p> <p>그릇의 재질은 무엇인가요? (What is the material of the bowl?)</p> <p>A. 돌 (Stone) B. 도자기 (Ceramic) C. 유리 (Glass) D. 스테인리스 스틸 (Stainless Steel)</p>
	<p>Category: Landmark and building - Spanish/Ecuador</p> <p>¿Cómo se llama este monumento ubicado en Quito? (What is the name of this monument located in Quito?)</p> <p>A. Virgen de El Panecillo (The Virgin of El Panecillo) B. Manto de María (Manto de María) C. Mitad del mundo (Middle of the world) D. Cristo de la concordia (Christ of peace)</p>
	<p>Category: Landmark and building - Irish/Ireland</p> <p>Cén cathair ina bhfuil na dealbha seo? (In which city are these statues?)</p> <p>A. Cathair Bhaile Átha Cliath (Dublin City) B. Páras (Paris) C. Cathair Corcaigh (Cork City) D. Beirlín (Berlin)</p> <p>Cén eachtra stairiúil atá léirithe sna dealbha seo? (What historical event is depicted in these statues?)</p> <p>A. An Ghorta Mhór (The Great Famine) B. Éirí Amach 1916 (The 1916 Rising) C. Teitheadh na n-Iarlaí (The flight of the Earls) D. Cogadh 1835 (The 1835 war)</p> <p>Cén abhainn atá le taobh na ndealbh seo? (What river is beside these statues?)</p>

	<p>A. An Life (The Liffey) B. An tSionann (The Shannon) C. Abhainn an Rí (King's River) D. An Thames (The Thames)</p>
	<p>Category: Public Figure and Pop Culture – Korean/South Korea</p> <p>이 그룹의 리더는 무슨 색 티셔츠를 입고 있나요? (What t-shirt color is the leader of this group wearing?)</p> <p>A. 남색 (Navy) B. 초록색 (Green) C. 주황색 (Orange) D. 버건디색 (Burgundy)</p> <p>이 그룹의 가장 어린 멤버는 무슨 색 티셔츠를 입고 있나요? (What t-shirt color is the youngest member of this group wearing?)</p> <p>A. 회색 (Grey) B. 주황색 (Orange) C. 남색 (Navy) D. 초록색 (Green)</p>
	<p>Category: Brand/ Products – Korean/South Korea</p> <p>이 캐릭터의 이름은 무엇일까요? (What is the name of this character?)</p> <p>A. 납죽이 (Nubzuki) B. 어피치 (Apeach) C. 피카츄 (Pikachu) D. 스폰지밥 (SpongeBob)</p> <p>이 캐릭터를 나타내는 학교는 어디인가요? (What is the name of the school does this character represents?)</p> <p>A. 카이스트 (KAIST) B. 고려대학교 (Korea University) C. 서울대학교 (Seoul National University) D. 연세대학교 (Yonsei University)</p>
	<p>Category: Vehicle and Transportation – Spanish/Ecuador</p> <p>¿Cómo se llama este particular medio de transporte en Ecuador? (What is this particular means of transportation called in Ecuador?)</p> <p>A. Tricimoto (Tricycle) B. Motocicleta (Motorcycle) C. Carro (Car) D. Taxi (Taxi)</p> <p>¿En qué Ciudad del Ecuador transita este medio de transporte? (In which city of Ecuador does this means of transportation travel?)</p> <p>A. Atacames (Atacames) B. Ibarra (Ibarra) C. Azogues (Azogues) D. Mejía (Mejía)</p>

B Annotation Platform

The interface is divided into two main sections. The left section is for image upload and metadata. It includes a 'Browse Files' button with a cloud icon and the text 'Drag and drop files here'. Below this is a note about image quality and storage. There are two radio buttons for 'Is the image self-made/yours?' (Yes is selected). Below that are two radio buttons for 'If this is your own image, will you allow this image for commercial use?' (No is selected). There are two dropdown menus for 'Question Language' and 'Image Country Origin', both set to '-- SELECT --'. At the bottom is a section for 'Image Category' with checkboxes for various categories: Vehicles and Transportation, Objects / materials / clothing, Cooking and food, Geography / buildings / landmarks, Plants and animal, Brands / products / companies, Sports and recreation, Traditions / art / history, People and everyday life, and Public Figure and pop culture.

The right section is for question entry. It has a header 'Question #1' with a dropdown arrow. Below it are two text input fields: 'Question 1 (in local language) *' and 'Translated Question 1 (in English) *'. Further down are six more text input fields arranged in two columns and three rows, labeled 'Q1: Correct answer (in local language) *', 'Q1: Correct answer (in English) *', 'Q1: Wrong option 1 (in local language) *', 'Q1: Wrong option 1 (in English) *', 'Q1: Wrong option 2 (in local language) *', 'Q1: Wrong option 2 (in English) *', 'Q1: Wrong option 3 (in local language) *', and 'Q1: Wrong option 3 (in English) *'.

Figure 4: Annotation interface for inputting image and questions

	Submission Date	Submission ID	Image Upload	Is t...	If this is your...	If this is not your own i...
146	Mar 18, 2024	5865939274273446442		Yes	Yes, people can use t	
147	Mar 18, 2024	5865939274273557337		No		https://www.flickr.com/photos...
148	Mar 18, 2024	5865939274273646152		No		https://commons.wikimedia.org...
149	Mar 18, 2024	5865939274273647849				https://commons.wikimedia.org...
150	Mar 18, 2024	5865939274273855904			Yes, people can use t	
151	Mar 18, 2024	5865939274273979270			No, this image is rese	
152	Mar 18, 2024	5865939274274187387			Yes, people can use t	
153	Mar 18, 2024	5865939274274387192		Yes	Yes, people can use t	
154	Mar 18, 2024	5865939274274408912		No		https://flickr.com/photos/bona...
155	Mar 18, 2024	5865939274274558472		Yes	Yes, people can use t	
156	Mar 18, 2024	5865939274274664373		No		https://commons.wikimedia.org...
157	Mar 18, 2024	5865939274274801278		No		https://flic.kr/p/5jEX4d
158	Mar 18, 2024	5865939274274878679		Yes	Yes, people can use t	

Figure 5: Annotation interface for validation. Contributors can comment, edit, and star the entries

We use JotForm as our annotation platform. For question entry, contributors can upload and write questions in both languages in the form. The interface can be seen in Figure 4. During validation, contributors can see all the data submitted by other contributors (Figure 5). They can select the entry to see detailed preview of the entry (Figure 6). They can then either edit the data directly, provide comments, or confirm the data by starring the entry.


C Most-Frequent Words in the Questions

Figure 7 shows word clouds for the most frequent words in CVQA per category. We exclude stopwords as well as 'picture', 'photo', and 'image' from the list, since most questions contain these words. In this VQA context, we can treat them as stopwords.

3. Pretzel

Updated at May 21, 2024

Image Upload



Is the image self-made/yours?

Yes

If this is your own image, will you allow this image for commercial use?

Yes, people can use this image for commercial

If this is not your own image, enter the original image URL

Enter a text

Question Language

Minangkabau

Image Country Origin

Indonesia

Image Category

3. Pretzel

Updated at May 21, 2024

Cooking and food

Question 1 (in local language)

Apo namo jajanan pasa ko?

Translated Question 1 (in English)

What is the name of this traditional snack?

Qt: Correct answer (in local language)

Kue tambang

Qt: Correct answer (in English)

Tambang cake

Qt: Wrong option 1 (in local language)

Pretzel

Qt: Wrong option 1 (in English)

Pretzel

Qt: Wrong option 2 (in local language)

Onde-onde

Figure 6: During validation, contributors can preview the submission from other contributors

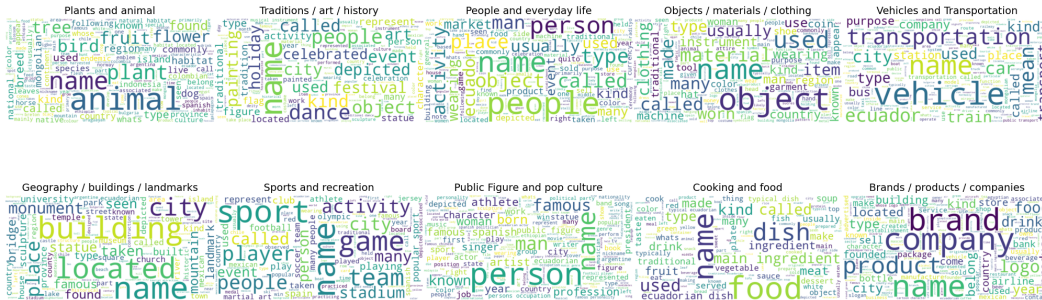


Figure 7: Word Cloud in CVQA per category

D CVQA Annotator Demographic

Figure 8 illustrates the demographic statistics of the annotators, based on an anonymous questionnaire we provided. At the time of writing, we have information for 36 out of 76 annotators. As such, this breakdown is a rough representation of the annotation group.




Figure 8: Annotator demographic statistics

E Country-Language Pairs and Scripts

In Table 8, we provide information on the script used in each Country-Language pair.

Country	Language	Script
<i>Africa</i>		
Egypt	Egyptian Arabic	Arabic
Ethiopia	Amharic	Amharic
Ethiopia	Oromo	Latin
Nigeria	Igbo	Latin
<i>Asia</i>		
China	Chinese	Chinese
India	Bengali	Bengali
India	Tamil	Tamil
Indonesia	Indonesian	Latin
Indonesia	Javanese	Latin
Indonesia	Minangkabau	Latin
Indonesia	Sundanese	Latin
Japan	Japanese	Japanese
South Korea	Korean	Hangul
Malaysia	Malay	Latin
Mongolia	Mongolian	Cyrillic
Pakistan	Urdu	Perso-Arabic
Philippines	Filipino	Latin
Singapore	Chinese	Chinese
Sri Lanka	Sinhala	Sinhalese
<i>Europe</i>		
Bulgaria	Bulgarian	Cyrillic
France	Breton	Latin
Ireland	Irish	Latin
Norway	Norwegian	Latin
Romania	Romanian	Latin
Russia	Russian	Cyrillic
Spain	Spanish	Latin
<i>Latin America</i>		
Argentina	Spanish	Latin
Brazil	Portuguese	Latin
Chile	Spanish	Latin
Colombia	Spanish	Latin
Ecuador	Spanish	Latin
Mexico	Spanish	Latin
Uruguay	Spanish	Latin

Table 8: The list of Country-Language pairs covered in CVQA and their corresponding scripts.