

Challenging Multimodal LLMs with African Standardized Exams: A Document VQA Evaluation

Victor Olufemi Oreoluwa Babatunde Kausar Moshood Emmanuel Bolarinwa

LyngualLabs

{victor, oreoluwa, kausar, emmanuel}@lynguallabs.org

Abstract

Despite rapid advancements in multimodal large language models (MLLMs), their ability to process low-resource African languages in document-based visual question answering (VQA) tasks remains limited. This paper evaluates three state-of-the-art MLLMs—GPT-4o, Claude-3.5 Haiku, and Gemini-1.5 Pro—on WAEC/NECO standardized exam questions in Yoruba, Igbo, and Hausa. We curate a dataset of multiple-choice questions from exam images and compare model accuracies across two prompting strategies: (1) using English prompts for African language questions, and (2) using native-language prompts. While GPT-4o achieves over 90% accuracy for English, performance drops below 40% for African languages, highlighting severe data imbalance in model training. Notably, native-language prompting improves accuracy for most models, yet no system approaches human-level performance, which reaches over 50% in Yoruba, Igbo, and Hausa. These findings emphasize the need for diverse training data, fine-tuning, and dedicated benchmarks that address the linguistic intricacies of African languages in multimodal tasks, paving the way for more equitable and effective AI systems in education.

1 Introduction

The rapid advancements in artificial intelligence (AI) have led to the emergence of multimodal large language models (MLLMs) capable of processing and understanding both textual and visual information (Peng et al., 2023; Ahuja et al., 2024). Notable examples include OpenAI’s GPT, Anthropic’s Claude and Google’s Gemini. These models exhibit impressive capabilities in interpreting combined visual-textual inputs, allowing them to extract text from images and answer questions about that content. However, their ability to accurately process text from images in low-resource languages remains an open question (Adelani et al.,

2025). Despite the progress in multilingual NLP, most state-of-the-art models are primarily trained on high-resource languages, resulting in suboptimal performance for many African languages. Low-resource languages are severely underrepresented in the datasets used to train and evaluate MLLMs (Joshi et al., 2020; Adelani et al., 2025), and African languages such as Yoruba, Igbo, and Hausa present unique linguistic and orthographic challenges that differ significantly from dominant languages on which these models are typically trained. (Orife et al., 2020). The scarcity of high-quality training data for these languages exacerbates the performance disparity between high- and low-resource languages (Nayak et al., 2024). Recent benchmarks confirm that multimodal models perform very well on English but struggle on many African languages due to data limitations (Ahuja et al., 2024)

1.1 WAEC and NECO: Importance in West African Education

The West African Examinations Council (WAEC) and the National Examinations Council (NECO) play a crucial role in standardized education across West Africa. These organizations administer high-stakes secondary school examinations that assess students’ proficiency in core subjects, including language proficiency in English as well as indigenous languages like Yoruba, Igbo, and Hausa. WAEC and NECO exams serve as key determinants of academic progression, influencing university admissions and career opportunities. Standardized exams such as WAEC and NECO provide an objective measure of students’ knowledge, making them an appealing testbed for evaluating AI models’ natural language understanding abilities in structured educational contexts. By leveraging real-world examination content, this study assesses whether state-of-the-art MLLMs can process structured educational material effectively in African languages.

1.2 Research Objectives

This study aims to systematically evaluate the performance of GPT-4o, Claude-3.5 Haiku and Gemini-1.5 Pro in natural language comprehension for African languages by addressing the following objectives:

- **Objective 1:** Assess the ability of multimodal LLMs to accurately extract and process text from WAEC/NECO examination images.
- **Objective 2:** Compare performance under different prompt languages, analyzing whether using English vs. native-language prompts affects answer accuracy.

2 Related Works

2.1 Multimodal Large Language Models and Their Capabilities

Multimodal large language models integrate multiple data modalities, such as text and images, to enhance comprehension and reasoning (Peng et al., 2023; Ahuja et al., 2024). These models build on advances in vision-language pre-training that combine visual encoders with language models (Radford et al., 2021; Li et al., 2023). State-of-the-art MLLMs have achieved impressive performance on many text-based visual tasks, including image captioning, document understanding, and visual question answering. In general, these models perform well on tasks in high-resource languages. However, studies have shown that their effectiveness diminishes significantly in low-resource languages such as Yoruba, Igbo, and Hausa (Adelani et al., 2025; Schneider and Sitaram, 2024)). For instance, the IrokoBench evaluation found a substantial drop in GPT-4o’s performance on African language understanding compared to English. Similarly, a culturally diverse VQA benchmark (Romero et al., 2024) demonstrated that even powerful vision-language models fail to generalize across linguistically diverse or culturally unfamiliar inputs. (Zhang et al., 2023) introduced M3Exam, a multilingual, multimodal exam benchmark, and reported major performance discrepancies between high-resource and low-resource languages. While current MLLMs can process Latin-script inputs with high accuracy, they struggle with the complex morphology and orthographic variations present in many African languages (Liu et al., 2023). This gap underscores that simply scaling to multimodal inputs is not sufficient for broad multilingual competency.

2.2 Challenges in Multilingual NLP for Low-Resource African Languages

The lack of training data remains a fundamental challenge in multilingual NLP research, particularly for African languages (Adelani et al., 2025). Unlike English or other widely spoken languages, Yoruba, Igbo, and Hausa have relatively limited corpora and annotated datasets available for training or fine-tuning large models. This data scarcity negatively impacts model performance on both text-only and multimodal tasks (Schneider and Sitaram, 2024)). Even large multilingual language models like XLM-R (Conneau et al., 2020) or BLOOM struggle on African languages that were underrepresented in their training data. In addition, many African languages have unique linguistic properties – for example, tonal phonology and extensive use of diacritics in Yoruba, or complex noun classes in some Bantu languages – which prove difficult for pre-trained LLMs to handle. These orthographic and grammatical nuances are often lost or misinterpreted by models not specifically adapted to them (Orife et al., 2020). Recent studies such as (Nayak et al., 2024) highlight that vision-language models exhibit poor understanding of culturally or linguistically specific content, reinforcing the importance of developing benchmarks that reflect real-world linguistic diversity. There have been efforts to bolster NLP for African languages – for example, the Masakhane project’s participatory approach to machine translation (Nekoto et al., 2020) and the creation of language-specific models like AfriBERTa (Ogueji et al., 2021), but these are text-only initiatives. Until similar resources and benchmarks are created on the multimodal front, AI models will continue to exhibit biases favoring high-resource languages over under-represented ones (Ahuja et al., 2024; Lu et al., 2024). Our work addresses this gap by providing a focused evaluation on Yoruba, Igbo, and Hausa, thereby pushing towards more inclusive multimodal model development.

2.3 Optical Character Recognition and Text Processing in Multimodal AI

Optical Character Recognition (OCR) plays a crucial role in multimodal AI by enabling models to extract text from images. However, existing OCR-focused evaluations – for example, the comprehensive OCRBench suite (Li et al., 2023) – indicate that current MLLMs often struggle with non-Latin

scripts, accented characters, and handwritten text. Even models like GPT-4 and Claude that are adept at English OCR can falter when faced with, say, a scanned Hausa document or Yoruba text with tone marks. Kosmos-2, a recent grounded multimodal model (Peng et al., 2023), introduced new capabilities for aligning text with visual regions, but its effectiveness on low-resource African languages remains unclear, as it was primarily evaluated on mainstream languages and object-centric tasks. In our study, we do not explicitly re-evaluate OCR accuracy at the character level; instead, we assess how well multimodal LLMs handle the output of OCR in a downstream task – specifically, answering multiple-choice questions based on exam images. By leveraging real WAEC/NECO exam questions, our evaluation extends prior research and provides new insights into OCR performance within an African educational context. In particular, our results can reveal whether state-of-the-art models accurately interpret the extracted text (including any diacritics or uncommon characters) and use it correctly to select answers. This complements existing OCR benchmarks by focusing on end-to-end comprehension: from image to extracted text to answer selection.

2.4 Standardized Exam Benchmarks in AI Research

Standardized exams have become a widely adopted benchmark for evaluating AI models. The structured format of exam questions—where each item follows a consistent style and has a known correct answer offers a controlled environment for assessing an AI’s reading comprehension, reasoning, and problem-solving abilities. Several recent studies have used exam-based benchmarks to evaluate large language models. For example, M3Exam (Zhang et al., 2023) compiles real multilingual exam questions and shows that GPT-4o and similar models perform well on high-resource languages but struggle on under-represented languages. Similarly, the MEGEVERSE benchmark (Ahuja et al., 2024) evaluated LLMs across 83 languages and highlighted substantial performance gaps in low-resource linguistic settings. Our study follows a similar methodology of exam-driven evaluation but narrows the focus specifically to structured educational content in popular Nigerian languages. By concentrating on WAEC/NECO multiple-choice questions in Yoruba, Igbo, and Hausa, we provide an in-depth look at model capabilities in a context

that had not been examined in prior multilingual benchmarks. This approach also complements efforts like Hendrycks et al.(2021)’s MMLU, which included a broad range of subjects and some languages: we add the dimension of image-based text understanding in an educational assessment scenario.

3 Methodology

3.1 Dataset Curation

The dataset for this study was curated from past WAEC and NECO examination questions in Yoruba, Igbo, Hausa, and English. We targeted multiple-choice questions (MCQs) from recent years to ensure a representative sample of modern usage. The curation process involved several steps:

3.1.1 Data Collection

We obtained past examination papers from students and bookshops that sell educational materials. However, acquiring exam questions for language subjects (Hausa, Igbo, Yoruba) online proved extremely challenging, if not nearly impossible, due to their limited availability compared to more widely documented subjects. To ensure a sizable dataset in each target language, we focused on examination papers from the years 2008–2024.

3.1.2 Question Segmentation

Each question‘ was manually cropped from scanned examination sheets to isolate it as an individual image. This ensured that each image contained exactly one question for the model to answer, standardizing the input format. Only multiple-choice questions were included to maintain a uniform evaluation style.

3.1.3 Answer Key Verification

Many exams came with official answer keys, which we treated as gold-standard answers. For questions lacking official keys (or in cases where only the exam paper was available), we consulted linguistic and subject matter experts fluent in Yoruba, Igbo, or Hausa to determine the correct answer. These expert-verified answers were cross-checked to ensure accuracy.

3.1.4 Categorization

Each question was labeled by language (Yoruba, Igbo, Hausa, or English) and by exam year. This allows us to perform year-wise or language-wise

analysis of the results. We ensured a roughly balanced number of questions per language where possible, though English had a naturally larger pool. This dataset of exam questions forms the basis for our evaluation. By using real educational content, we ensure that the evaluation is grounded in tasks that have practical importance and linguistic richness. Table 1 below summarizes the provisional composition of the dataset:

Year	English	Yoruba	Igbo	Hausa
2008	0	20	0	0
2009	0	16	0	0
2010	0	19	0	0
2011	0	20	0	0
2012	0	17	0	0
2013	0	20	0	0
2014	0	20	0	0
2015	0	19	0	0
2016	0	20	0	0
2017	0	0	0	0
2018	29	19	0	0
2019	30	19	0	0
2020	30	19	0	0
2021	60	38	24	0
2022	60	40	45	20
2023	60	36	45	20
2024	0	36	0	36
Total	269	378	114	76

Table 1: Dataset composition by year and language. WAEC and NECO Combined

3.2 Model Selection and Evaluation Criteria

We selected three state-of-the-art multimodal LLMs for benchmarking: GPT-4o (OpenAI), Claude-3.5 Haiku (Anthropic), and Gemini-1.5 pro (Google DeepMind). These models although uneven in sizes were chosen due to their cutting-edge performance and diverse origins (industry leaders in AI). We accessed GPT-4o, Claude-3.5 Haiku, and Gemini-1.5 Pro via their official API endpoints. While other emerging models (such as Mistral) could be considered, we limited our testing to these three due to time and resource constraints. Our evaluation was based on two primary criteria:

- **Answer Accuracy:** The percentage of questions for which the model’s answer matched the expert-verified correct answer. This is a di-

rect measure of performance on the multiple-choice questions.

- **Language-wise Performance:** We compare accuracy across the four languages (English, Yoruba, Igbo, Hausa) to identify any performance disparities.

3.3 Experimental Setup

We designed a uniform evaluation pipeline and prompting strategy to ensure a fair comparison between models. Key aspects of the experimental setup are outlined below:

3.3.1 Prompting Strategy

We employed two query strategies for each question image:

1. An English-prompted query.
2. A native-language-prompted query.

In the English prompt condition, the model was instructed in English (e.g., “Analyze the image and answer the question”) while being given an image containing a Yoruba/Igbo/Hausa question. In the native prompt condition, we translated the instruction into the question’s language (Yoruba, Igbo, or Hausa) so that the model received the prompt in the same language as the question. This allows us to test whether prompting in the local language improves understanding or not. Each model thus answers every question twice: once with an English prompt and once with a native-language prompt.

3.3.2 Temperature Setting

While most API parameters were left at their defaults, we explicitly set the temperature to 0.1 to ensure minimal randomness and greater response consistency across model runs. This controlled setting ensures that each model selects the most probable answer rather than generating diverse outputs.

3.3.3 Multiple-Choice Answering Format

To reduce variability in how models produce answers and to minimize open-ended generation issues (e.g., hallucinations), we prompted the models to choose one of the options A, B, C, D, E for each question. We adopted a best-practice format inspired by prior VQA benchmarks (e.g., CVQA): the model is instructed to assign a probability score to each option and then select the option with the highest probability. In practice, we implemented

this by instructing the model to output a JSON object with scores, which forces the model to make a single choice.

3.3.4 Human Evaluation via Independent NLP Community

A crucial component of our methodology involved human evaluation using participants from an independent NLP community. They have a network of linguistics enthusiasts, including students with a keen interest in language processing. We engaged a subset of students from this community to answer the same multiple-choice exam questions that were presented to the AI models. The students were selected based on their fluency in Yoruba, Igbo, or Hausa, but they were not necessarily language experts. Their responses provide a useful baseline to compare human vs. model performance on these questions. (This human study was conducted with appropriate consent and is intended for qualitative comparison, not as a rigorous benchmark.)

3.3.5 Prompt Template

We crafted a consistent system message for all models, emphasizing the task and format. Below is a simplified example of the prompt content used (shown here in English for brevity):

System Prompt:

“You are a knowledgeable assistant for answering exam questions. Carefully read the question in the image and evaluate each of the four choices. Provide the answer by indicating the option (A, B, C, D, or E) with the highest probability of being correct, along with probability scores for each option in JSON format.”

User Prompt:

“Analyze the following question image and determine the correct answer (A, B, C, D, or E). Respond in JSON with your probabilities for each option.”

For native-language trials, the prompts were translated appropriately (e.g., to Yoruba). All models were thus given a very similar cue and format requirement, to the extent their API allowed system instructions.

3.4 Evaluation Metric

We used a strict accuracy metric for each model’s responses. A model receives a score of 1 for a question if its highest-probability choice matches the

correct answer, and 0 otherwise. We then compute overall accuracy as well as per-language accuracy.

The above methodology enables a controlled and fair evaluation of each model’s ability to interpret exam images and answer questions in multiple languages. All model outputs and metadata are logged for analysis.

4 Results

The evaluation results provide insights into the performance of GPT-4o, Gemini-1.5 Pro, and Claude-3.5 Haiku on multiple-choice exam questions in Yoruba, Hausa, Igbo, and English. We analyze accuracy under two prompting conditions:

1. Prompting in English.
2. Prompting in the respective African language.

We also compare the models’ performance to human baseline scores.

4.1 Model Performance Across Languages

The Table 2 below presents the accuracy scores for each model across different languages and prompt conditions:

Prompt	GPT-4o Accuracy	Gemini-1.5 Pro Accuracy	Claude-3.5 Haiku Accuracy
Yoruba Exam Questions			
Yoruba Prompt	32.80% (124/378)	29.63% (112/378)	26.72% (101/378)
English Prompt	31.74% (121/378)	33.86% (128/378)	25.92% (98/378)
Hausa Exam Questions			
English Prompt	39.47% (30/76)	36.84% (28/76)	28.95% (22/76)
Hausa Prompt	43.42% (33/76)	44.74% (34/76)	23.68% (18/76)
English Exam Questions			
English Prompt	90.33% (243/269)	73.61% (198/269)	55.39% (149/269)
Yoruba Prompt	79.55% (214/269)	72.49% (195/269)	39.03% (105/269)
Hausa Prompt	80.30% (216/269)	72.86% (196/269)	40.89% (110/269)
Igbo Prompt	81.04% (218/269)	72.12% (194/269)	36.43% (98/269)
Igbo Exam Questions			
English Prompt	27.19% (31/114)	31.58% (36/114)	18.42% (21/114)
Igbo Prompt	28.95% (33/114)	35.96% (41/114)	23.68% (27/114)

Table 2: Accuracy scores for GPT-4o, Gemini, and Claude across different languages and prompt conditions.

4.2 Key Observations

- **Higher Accuracy in English:** As expected, models performed significantly better on English-only questions, with GPT-4o achieving the highest accuracy (90.33%), followed by Gemini-1.5 Pro (73.61%) and Claude -3.5 Haiku (55.39%). This confirms that the models handle high-resource languages much better than low-resource ones.
- **Effect of Prompting English Questions in African Languages:** Interestingly, when English questions were prompted in Yoruba,

Hausa, and Igbo, accuracy dropped compared to using English prompts. GPT-4o’s accuracy dropped from 90.33% (English prompt) to 79.55% (Yoruba prompt), 80.30% (Hausa prompt), and 81.04% (Igbo prompt). Gemini-1.5 Pro and Claude-3.5 Haiku showed similar trends, highlighting how translation and linguistic context impact comprehension.

- **Native Language Prompts Improve Accuracy:** For Yoruba, Hausa, and Igbo, prompting the model in the native language generally resulted in higher accuracy than when the prompt was in English. The effect was particularly noticeable in Hausa (e.g., GPT-4o: 43.42% Hausa-prompted vs. 39.47% English-prompted).

4.3 Comparison with Human Performance

We also compared model results with human performance, where participants from an independent NLP community answered the same exam questions. The results are presented in Table 3 below:

Language	Human Accuracy
Hausa	68.0%
Igbo	52.3%
Yoruba	56.0%

Table 3: Comparison of human accuracy on multiple-choice exam questions across three African languages.

Human accuracy was significantly higher than all model performances across the three African languages, reinforcing that even non-expert humans outperform state-of-the-art AI models on structured educational tasks in Yoruba, Igbo, and Hausa.

These results provide strong evidence of the performance gap between AI models and human linguistic abilities, particularly in low-resource African languages.

5 Discussion

The results indicate several key trends regarding multimodal LLMs’ performance in African languages. Below, we discuss the implications of these findings and analyze potential causes and areas for improvement.

5.1 Performance Disparities Across Languages

Our findings confirm that AI models struggle significantly with low-resource languages, particularly in the context of document VQA for standardized exams. The sharp decline in accuracy from 90.33% in English (GPT-4o) to below 40% for Yoruba, Hausa, and Igbo underscores the severe data imbalance in model training. Despite advancements in multilingual AI, African languages remain underrepresented in training datasets, leading to weaker comprehension and reasoning abilities when processing structured educational assessments. This highlights the critical need for more inclusive AI models capable of handling the complexities of standardized exam content in low-resource languages.

5.2 Effect of Native vs. English Prompting

A critical takeaway from the study is that models perform better when prompted in the same language as the question. This trend was particularly consistent for GPT-4o and partially observed in Gemini and Claude. The improvement suggests that prompting in the target language helps models better interpret syntactic and semantic nuances, likely because it reduces the additional complexity of cross-language interpretation.

However, this trend was not uniform across all languages and models. For example, Gemini performed slightly better when prompted in English for Yoruba questions, suggesting that some models may rely on English as an anchor for reasoning. This discrepancy warrants further investigation into the internal translation and tokenization processes of multimodal LLMs.

5.3 Human vs. AI Performance Gap

Human participants vastly outperformed all models, with an average accuracy of 56.0% (Yoruba), 52.3% (Igbo), and 68% (Hausa). This performance gap is expected, but its magnitude is striking, especially considering that the human evaluators were not expert linguists, but student enthusiasts from an independent NLP community. The disparity suggests that AI models lack fundamental linguistic and contextual understanding needed for structured educational tasks in African languages.

This reinforces the need for more diverse and representative training datasets to improve multimodal AI comprehension of African languages. It also suggests the potential for fine-tuning or adapta-

tion strategies to boost model performance in these languages.

6 Conclusion

This study evaluated the performance of state-of-the-art multimodal large language models—GPT-4o, Gemini, and Claude—on structured educational tasks in Yoruba, Igbo, and Hausa. Our findings reveal several key takeaways:

- **Performance disparities exist across languages**, with models performing significantly better in English than in African languages for VQA in standardized exams.
- **Prompting in native languages improves model accuracy**, particularly for GPT-4o, suggesting a need for further multilingual optimization.
- **AI models still lag behind human performance**, with human participants from an independent NLP community achieving much higher accuracy than all models.

These results highlight the challenges of multimodal NLP for African languages and emphasize the need for greater linguistic inclusivity in AI training data and model design.

7 Future Work

Future research can build upon this study by expanding and improving multimodal datasets for African languages, ensuring high-quality resources that help bridge performance gaps. Fine-tuning large language models (LLMs) with domain-specific data could further enhance their comprehension and reasoning capabilities in these languages. Additionally, the development of standardized evaluation benchmarks for African multimodal NLP would enable systematic model comparisons. Investigating OCR accuracy for African scripts is another crucial area, as many languages have unique orthographic systems that present distinct challenges. Finally, broadening the scope beyond Yoruba, Igbo, and Hausa to include a wider range of African languages would provide a more comprehensive understanding of NLP challenges on the continent.

References

- David Ifeoluwa Adelani, Jessica Ojo, Israel Abebe Azime, Jian Yun Zhuang, Jesujoba O. Alabi, Xuanli He, Millicent Ochieng, Sara Hooker, Andiswa Bukula, En-Shiun Annie Lee, Chiamaka Chukwunke, Happy Buzaaba, Blessing Sibanda, Godson Kalipe, Jonathan Mukiibi, Salomon Kabongo, Foutse Yuehgoh, Mmasibidi Setaka, Lolwethu Ndolela, Nkiruka Odu, Rooweither Mabuya, Shamsuddeen Hassan Muhammad, Salomey Osei, Sokhar Samb, Tadesse Kebede Guge, Tombekai Vangoni Sherman, and Pontus Stenetorp. 2025. [Irokobench: A new benchmark for african languages in the age of large language models](#). *Preprint*, arXiv:2406.03368.
- Sanchit Ahuja, Divyanshu Aggarwal, Varun Gumma, Ishaan Watts, Ashutosh Sathe, Millicent Ochieng, Rishav Hada, Prachi Jain, Mohamed Ahmed, Kalika Bali, and Sunayana Sitaram. 2024. [MEGAVERSE: Benchmarking large language models across languages, modalities, models and tasks](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2598–2637, Mexico City, Mexico. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). *Preprint*, arXiv:2009.03300.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. [Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models](#). *Preprint*, arXiv:2301.12597.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). *Preprint*, arXiv:2304.08485.
- Yujie Lu, Dongfu Jiang, Wenhui Chen, William Yang Wang, Yejin Choi, and Bill Yuchen Lin. 2024. [Wildvision: Evaluating vision-language models in the wild with human preferences](#). *Preprint*, arXiv:2406.11069.

- Shravan Nayak, Kanishk Jain, Rabiul Awal, Siva Reddy, Sjoerd Van Steenkiste, Lisa Anne Hendricks, Karolina Stanczak, and Aishwarya Agrawal. 2024. [Benchmarking vision language models for cultural understanding](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5769–5790, Miami, Florida, USA. Association for Computational Linguistics.
- Wilhelmina Nekoto, Vukosi Marivate, Tshinondiwa Matsila, Timi Fasubaa, Taiwo Fagbohunge, Solomon Oluwole Akinola, Shamsuddeen Muhammad, Solomon Kabongo Kabenamualu, Salomey Osei, Freshia Sackey, Rubungo Andre Niyongabo, Ricky Macharm, Perez Ogayo, Orevaoghene Ahia, Musie Meressa Berhe, Mofetoluwa Adeyemi, Masabata Mokgesi-Seling, Lawrence Okegbemi, Laura Martinus, Kolawole Tajudeen, Kevin Degila, Kelechi Ogueji, Kathleen Siminyu, Julia Kreutzer, Jason Webster, Jamiil Toure Ali, Jade Abbott, Iro Orife, Ignatius Ezeani, Idris Abdulkadir Dangan, Herman Kamper, Hady Elshahar, Goodness Duru, Ghollah Kioko, Murhabazi Espoir, Elan van Biljon, Daniel Whitenack, Christopher Onyefuluchi, Chris Chinenye Emezue, Bonaventure F. P. Dossou, Blessing Sibanda, Blessing Bassey, Ayodele Olabiyi, Arshath Ramkilowan, Alp Öktem, Adewale Akinfaderin, and Abdallah Bashir. 2020. [Participatory research for low-resourced machine translation: A case study in African languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2144–2160, Online. Association for Computational Linguistics.
- Kelechi Ogueji, Yuxin Zhu, and Jimmy Lin. 2021. [Small data? no problem! exploring the viability of pretrained multilingual language models for low-resourced languages](#). In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 116–126, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Iro Orife, David I. Adelani, Timi Fasubaa, Victor Williamson, Wuraola Fisayo Oyewusi, Olamilekan Wahab, and Kola Tubosun. 2020. [Improving yorùbá diacritic restoration](#). *Preprint*, arXiv:2003.10564.
- Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. 2023. [Kosmos-2: Grounding multimodal large language models to the world](#). *Preprint*, arXiv:2306.14824.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). *Preprint*, arXiv:2103.00020.
- 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 Biadglign Ademtew, 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-Garitagotia, Maria Camila Buitrago Cabrera, Mario Rodríguez-Cantelar, Mélanie Joutiteau, 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, Teresa Clifford, 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, Tamar Solorio, and Alham Fikri Aji. 2024. [Cvqa: Culturally-diverse multilingual visual question answering benchmark](#). *Preprint*, arXiv:2406.05967.
- Florian Schneider and Sunayana Sitaram. 2024. [M5 – a diverse benchmark to assess the performance of large multimodal models across multilingual and multicultural vision-language tasks](#). *Preprint*, arXiv:2407.03791.
- Wenxuan Zhang, Sharifah Mahani Aljunied, Chang Gao, Yew Ken Chia, and Lidong Bing. 2023. [M3exam: A multilingual, multimodal, multilevel benchmark for examining large language models](#). *Preprint*, arXiv:2306.05179.