



## Where Are We?

# Evaluating LLM Performance on African Languages

Ife Adebara<sup>ξ</sup> Hawau Olamide Toyin<sup>Ω</sup> Nahom Tesfu Ghebremichael<sup>Ω</sup>

AbdelRahim Elmadany<sup>ξ</sup> Muhammad Abdul-Mageed<sup>ξ,Ω,λ</sup>

<sup>ξ</sup>The University of British Columbia <sup>Ω</sup>MBZUAI <sup>λ</sup>Invertible AI

{ife.adebara, a.elmadany, muhammad.mageed}@ubc.ca



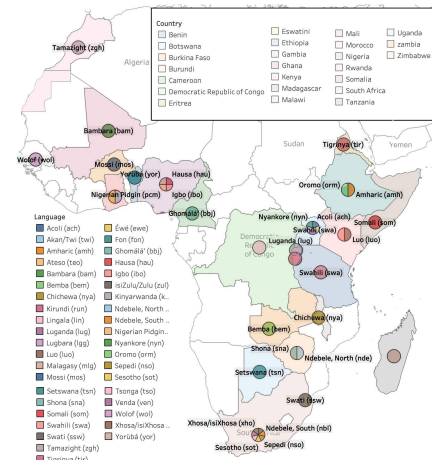
(a) Language policies for education across Africa.



(b) Official national language policies across Africa.



(c) Official regional language policies across Africa. *None* means no additional policy is available on a regional level.



(d) We collect SAHARA to empirically evaluate LLM performance on African languages, allowing us to demonstrate with evidence how current language policies directly impact progress in the field.

Figure 1: Maps of Africa showing the languages covered in this work and the language policies across the continent. The term *Indigenous* language is broadly defined here as one that is ‘native’ to the area (Walsh, 2005). The term *Both* in maps (a), (b), and (c) refers to *Indigenous* and *foreign* languages combined. Knowledge of what is an Indigenous language is based exclusively on *Ethnologue* (<https://www.ethnologue.com>).

## Abstract

Africa’s rich linguistic heritage remains under-represented in NLP, largely due to historical policies that favor foreign languages and create significant data inequities. In this paper, we integrate theoretical insights on Africa’s lan-

guage landscape with an empirical evaluation using SAHARA— a comprehensive benchmark we curate from large-scale, publicly accessible datasets capturing the continent’s linguistic diversity. By systematically assessing the performance of leading large language mod-

els (LLMs) on SAHARA, we demonstrate how policy-induced data variations directly impact model effectiveness across African languages.<sup>1</sup> Our findings reveal that while models perform reasonably well on few languages, many Indigenous languages remain marginalized due to sparse data. Leveraging these insights, we offer actionable recommendations for policy reforms and inclusive data practices. Overall, our work underscores the urgent need for a dual approach—combining theoretical understanding with empirical evaluation—to foster linguistic diversity in AI for African communities.

## 1 Introduction

The digital landscape of Africa is as diverse as its cultures and languages, yet the continent’s rich linguistic heritage remains largely underrepresented in NLP research. Despite significant transformations over the past decade driven by advances in large language models (LLMs) and a growing interest in linguistic diversity, most efforts have concentrated on a handful of widely spoken languages such as Swahili, Afrikaans, and Hausa, while many Indigenous languages continue to be sidelined due to the paucity of accessible, high-quality datasets. This oversight, deeply rooted in historical language policies that have neglected robust data collection and accessibility, has far-reaching implications for the development of language technology across Africa (Eberhard et al., 2021). Recent initiatives, including the work of Masakhane (Dione et al., 2023b; Adelani et al., 2021c, 2023, 2024b; Bandarkar et al., 2024a) and the efforts of platforms like HuggingFace, have begun to address these challenges by curating language-specific datasets and adapting transfer learning techniques for low-resource African languages. Models such as Cheetah (Adebara et al., 2024a), Serengeti (Adebara et al., 2023), Toucan (Elmadany et al., 2024), Afri-XLMR (Alabi et al., 2022), AfriTeVa (Oladipo et al., 2023; Jude Ogundepo et al., 2022), mBERT (Devlin et al., 2018), XLM-R (Conneau et al., 2020), and LLaMA (Touvron et al., 2023) have leveraged multilingual data to enhance performance on these languages, yet progress remains uneven—highlighting that the advancements achieved are primarily concentrated on a select few.

In this study, we undertake a comprehensive empirical evaluation of leading LLMs on an extensive

benchmark that we collect using mostly existing datasets to unravel how current African language policies, by dictating data availability and accessibility, directly shape model performance across the continent’s diverse linguistic spectrum. By systematically analyzing the distribution of existing datasets across languages and performance patterns among various LLMs, we not only pinpoint which languages benefit from ample data resources but also illuminate the underlying factors that contribute to their relative success. Our investigation reveals that the disparities in NLP outcomes are closely tied to policy-driven data inequities, offering concrete evidence of the need for more inclusive, forward-thinking language policies. These insights provide actionable recommendations for enhancing dataset creation and model training, ultimately aiming to bridge the digital divide and foster linguistic diversity in artificial intelligence for African communities.

Overall, we make the following contributions: **(1) Benchmark with wide, diverse coverage.** We introduce *Sahara*, a comprehensive benchmark assembled from large-scale, publicly accessible, and inclusive datasets that capture the rich linguistic diversity of Africa. This benchmark enables a systematic evaluation of data availability and model performance across a broad range of African languages. **(2) Dynamic leaderboard.** We develop a dynamic leaderboard built on best design practices, providing a transparent and continually updated platform to track progress and benchmark innovations in African NLP. **(3) Empirical evaluation.** We perform an extensive empirical assessment of a wide spectrum of models thereby establishing robust baselines and revealing performance disparities linked to data resource availability. **(4) Policy-driven and actionable insights.** We uncover clear links between current language policies, data availability, and model performance, demonstrating how policy-induced data inequities shape the effectiveness of language models across Africa. Our analysis provides recommendations for future research and development by informing policy reforms that foster more inclusive language technologies tailored to the needs of African communities. We now discuss language policies in Africa.

## 2 Language Policies in Africa

Across Africa, the predominant response to the continent’s multilingual landscape has been the

<sup>1</sup>SAHARA is publicly available at <https://github.com/UBC-NLP/sahara>.

adoption of a foreign language for official functions. However, in some instances, select Indigenous African languages receive official recognition at regional or national levels or within educational contexts (Petzell, 2012; Foster, 2021; Ouane and Glanz, 2010).<sup>2</sup> For instance in Nigeria, English is the official language, with only three out of 512 Indigenous languages recognized at the regional level. Similarly, Ghana designates English as its official language while also recognizing ten of its 73 Indigenous languages for institutional use. In Tanzania, Swahili is the sole Indigenous official language among 118 languages, alongside English. Kenya grants official status to 12 of its 61 languages, while South Africa recognizes 12 out of 20 Indigenous languages as institutional languages (Adebara and Abdul-Mageed, 2022). Figures 1(a), 1(b), 1(c), and 1(d) illustrate the distribution of language policies across the continent.

Even when Indigenous languages are granted official status alongside foreign languages, they often serve a symbolic rather than functional role. For instance, although Kiswahili is recognized as an official language by the African Union, its website and official document releases remain in English and French. A similar pattern emerges in education systems: where Indigenous languages are used, their role is typically confined to early childhood education and is usually paired with a foreign language rather than serving as the sole medium of instruction (Petzell, 2012; Foster, 2021; Ouane and Glanz, 2010). These language policies significantly shape language use across various domains, including newspapers, radio, television, and social media. Due to the dominance of foreign languages in official contexts, major newspapers and government publications are predominantly produced in languages such as English, French, and Portuguese, thereby limiting access for speakers of Indigenous languages (ResCue and Agbozo, 2021). Although some Indigenous languages are used on television and radio, foreign languages tend to dominate national broadcasts, especially in political discourse

and formal news reporting (Cheo et al., 2023; Myers and Harford, 2020; National Communications Authority, 2016).

Social media presents a more flexible environment where users can communicate in Indigenous languages; however, platform support remains uneven (Sunday et al., 2018; Molale and Mpofu, 2023; ResCue and Agbozo, 2021). Many African languages are underrepresented in digital spaces, lacking features such as text input options, spell checkers, or automated translation services. Consequently, users often resort to code-switching between Indigenous and dominant foreign languages, reinforcing existing linguistic divides. These trends underscore how official language policies shape broader language use, ultimately contributing to the marginalization of Indigenous languages in both public and digital communication spheres.

The analysis of language policies in Africa reveals how historical choices have led to significant data imbalances for Indigenous languages. Motivated by this, we built our benchmark, SAHARA to evaluate how these policies directly impact data availability, coverage, and the performance of LLMs on African languages. In the following section, we detail the construction of SAHARA and demonstrate its role in linking language policy to empirical outcomes in NLP.

### 3 Related Work

The development of NLP for African languages is hindered by data scarcity, policy-induced disparities, and suboptimal model performance. Recent initiatives have sought to mitigate these challenges by curating diverse datasets, rigorously evaluating multilingual models, and advocating for more inclusive language policies. For example, several initiatives have contributed to expanding datasets for African (Yimam et al., 2020; Muhammad et al., 2022; Aliyu et al., 2022; Muhammad et al., 2023b,a; Ilevbare et al., 2024; Elmadany et al., 2024). Some notable ones like Adelani et al. (2021c, 2022b) introduced MasakhaNER, a named entity recognition dataset for 10 African languages, highlighting the limitations in existing resources. Similarly, IROKOBench (Adelani et al., 2024b) evaluated LLMs on African languages, revealing stark performance gaps between high- and low-resource languages. Adebara et al. (2022) developed AfroLID, a neural language identification model covering 517 African languages, showing

<sup>2</sup>An *Indigenous* language is defined here as one that is *native* to a particular region, rather than introduced from elsewhere, in contrast to a *foreign* language, which may nonetheless serve as the mother tongue for certain local populations (Walsh, 2005). A *national* language refers to a language spoken broadly across an entire country. A *regional* language is spoken by a significant number of people within a specific area but does not have widespread national usage. An *educational* language is officially designated as a medium of instruction within the educational system.

that most languages lack NLP datasets beyond language identification.

Language policies play a crucial role in determining resource availability. Adebara and Abdul-Mageed (2022); Hershovich et al. (2022) emphasized the sociolinguistic factors affecting NLP development, noting that languages with official status receive more institutional support, fostering greater digital representation. Bird (2020) further critiqued the Western-centric approach in speech and language technologies, calling for AI models that align with Indigenous linguistic contexts. Model performance in African languages remains uneven, largely due to the dominance of high-resource languages in multilingual training. Conneau et al. (2018) introduced XLM-R, demonstrating the limitations of cross-lingual transfer learning for underrepresented languages. Meta AI’s No Language Left Behind (NLLB) project (Costa-jussà et al., 2022) attempted to improve multilingual models for low-resource languages, but performance remained inconsistent, particularly in generation tasks. Despite these advances, African NLP still requires increased investment in dataset curation, inclusive policy frameworks, and model development tailored to the linguistic diversity of the continent. We now introduce SAHARA.

## 4 Sahara Benchmark

Our objective is to create a comprehensive benchmark for African NLP that enables (i) analysis of existing language resources and (ii) the assessment of language models and tracking of the progress of African NLP. To achieve this goal, we develop SAHARA, adhering to several key design principles that we will now elucidate. SAHARA establishes a comprehensive and adaptable benchmark for African NLP using publicly available datasets. Figure shows coverage of SAHARA.

**Wide and Diverse Coverage.** SAHARA is the most comprehensive benchmark for African NLP, covering a vast range of languages. We achieve this by collecting high-quality, publicly available datasets from as many African languages as possible, ensuring easy accessibility for researchers evaluating their models. As a result, SAHARA supports 517 languages across various tasks, making it the most extensive and representative benchmark for African NLP. It encompasses languages from 50 out of 54 African countries and includes data

written in five different scripts: *Arabic*, *Coptic*, *Ethiopic*, *Latin*, and *Vai*, spanning five language families across the continent.

**Tasks and Task Clusters.** SAHARA is designed to support a broad spectrum of NLP tasks, which we systematically organize into coherent task clusters. The inclusion of well-defined, challenging clusters enables a more meaningful assessment of a model’s capabilities and allows researchers to evaluate performance on specific clusters. As shown in Table 1, SAHARA comprises four clusters: *text classification*, *text generation*, *multiple-choice*, *comprehensive and reasoning (MCCR)*, and *token-level classification*. For further details on the datasets used for each task, please refer to Appendix A.

**Modular Public Leaderboard.** We implement a *user-friendly leaderboard* for model evaluation on SAHARA, hosted on Hugging Face (HF) Spaces. Evaluated models must be publicly available in HF repositories to ensure seamless integration. When submitting a model, users are required to provide detailed *metadata*—including whether the model is pretrained, further pretrained, or instructed—and then submit an evaluation request. The submission enters a processing queue, and once evaluation is complete, the results are automatically displayed on the leaderboard. Our team will actively maintain the leaderboard by continuously integrating new publicly available tasks and datasets to keep it *up-to-date*. SAHARA thus serves as a valuable tool for assessing model performance in African NLP, accelerating the development of high-performance models, and contributing to the field’s growth. In addition to evaluating models on the entire benchmark, users have the flexibility to assess performance on specific task clusters, encouraging the development of specialized models tailored to address the unique challenges of African NLP.

## 5 Experiments and Results

### 5.1 Experimental Setup

To comprehensively assess the progress of African NLP, we systematically evaluate a range of widely used instruction-based LLMs in a few-shot setting on the SAHARA benchmark. We include both open and closed LLMs, such as Phi-3.5 and Phi-4 (Abdin et al., 2024), Gemma family (Team et al., 2024) models including Gemma-2 (2, 9, 27B) and Gemma-3 (4, 12, 27B), Llama family models (Dubey et al., 2024) including Llama-3.1 (8,



Cluster	Task	Identifier	#AL	#DS
Classification	Cross-Lingual Natural Language Inference (Adelani et al., 2024b)	xnli	16	1
	Language Identification (Adebara et al., 2022)	langid	517	1
	News Classification (Azime and Mohammed, 2021; Niyongabo et al., 2020; David, 2020a)	news	4	3
	Sentiment Analysis (Diallo et al., 2021; Oyewusi et al., 2020; Shode et al., 2022a)	sentiment	3	3
	Topic Classification (Hedderich et al., 2020)	topic	2	1
Generation	Machine Translation (Adelani et al., 2022a; Reid et al., 2021a; Ogueji and Ahia, 2019; Akera et al., 2022)	mt	29	5
	Paraphrase (Scherrer, 2020)	paraphrase	4	1
	Summarization (Hasan et al., 2021; Adebara et al., 2024b)	summary	10	2
	Title Generation (Hasan et al., 2021; Adebara et al., 2024b)	title-gen	10	2
MCCR	General Knowledge (Adelani et al., 2024b)	mmlu	16	1
	Mathematical Word Problems (Adelani et al., 2024b)	mgsm	16	1
	Reading Comprehension (Bandarkar et al., 2024a)	reading-comp	25	1
	Context-based Question Answering (Clark et al., 2020)	squad-qa	1	1
Tokens	NER (Adelani et al., 2021a, 2022c; Eiselen, 2016; Alabi et al., 2020a; Pan et al., 2017)	ner	27	5
	Phrase Chunking (Eiselen, 2016)	chunking	8	1
	POS Tagging (Onyenwe et al., 2018a, 2019)	pos	1	1
<b>Total</b>	<b>16</b>		<b>517</b>	<b>30</b>

Table 1: Descriptive statistics of languages and dataset diversity in SAHARA across different tasks and task-clusters. **#AL**: African Languages, **#DS**: datasets, **MCCR**: multiple-choice, comprehensive and reasoning task-cluster.

70B), Llama-3.2 (1, 3B), and Llama-3.3 (70B), Aya models (8, 23B) (Aryabumi et al., 2024), and Command family models (Cohere et al., 2025) including Command-R-Plus-R7B, Command-R-Plus (104B) and Command-A (111B), as well as proprietary models such as GPT-4.1<sup>3</sup> (Achiam et al., 2023) and Claude-4-Sonnet<sup>4</sup> (Anthropic, 2024). We further analyze performance by categorizing models according to *size*: small language models (SLM, <8B parameters) and large language models (LLM, ≥8B parameters), allowing investigation of the relationship between scale and effectiveness. For robust comparison, we aggregate evaluation datasets according to task clusters described in Appendix A. From each task, we randomly sample 1,000 data points for few-shot testing, employing well-curated, concise prompts—examples of which appear in Appendix C1.

## 5.2 Results and Discussion

We evaluate 24 LMs across four task clusters, maintaining a consistent assessment approach while varying evaluation metrics. Depending on the task, we employ *Exact Match*,  $F_1$ , *Accuracy*, *spBLEU*<sup>1K</sup> (Elmadany et al., 2024), and *RougeL* scores.

As shown in Table 2, in the *classification tasks*, the closed models Claude-4-Sonnet and GPT-4.1 consistently lead in performance, followed by larger LLMs such as Gemma3 (27B) and CommandR-Plus-104B. In the *language identification task*, most models struggle to grasp the task effectively, with only the larger LLMs demonstrat-

ing some understanding, suggesting that model size is crucial for achieving higher accuracy across various multilingual tasks. In the *generation task cluster*, larger models perform as expected, with GPT-4.1 and Llama-3.1 (70B) achieving the best results. There is a noticeable performance increase across all models in the paraphrase task, while the summarization task exhibits a more balanced performance among all models. Considering their size, Phi-3.5 and Gemma3 (4B) perform well within this task cluster. In the *MCCR tasks*, Command-A (111B) performs competitively against the closed Claude-4-Sonnet model. In the MGSM task, which involves solving mathematical word problems, most open-source models perform poorly, while Claude-4-Sonnet excels, achieving an exact match score of 22.34%. In token-level tasks, GPT-4.1 and Command-A (111B) consistently deliver strong performances, whereas the Aya model family struggle to grasp the tasks.

Results demonstrate a nuanced landscape of performance across multilingual and African LLMs, with notable variations depending on parameter size and task type. In few-shot settings, the closed model Claude-4-Sonnet achieves the highest overall average score (40.82), outperforming all open models across most task clusters. Among open models, Command-A (111B) and Gemma3 (27B) also exhibit strong performance, particularly in downstream tasks such as reasoning tasks, where Gemma2 (27B) reaches top score of 82.49 on Squad-QA and Command-A achieves 61.91 on MMLU task. For smaller parameter regimes, Gemma-2 (2B) and Phi-4 (3.8B) demonstrate the best performance on their model size. Furthermore,

<sup>3</sup>GPT-41.1 version: gpt-4.1-2025-04-14.

<sup>4</sup>Claude-4-Sonnet version: claude-4-sonnet-20250514

Cluster	Identifier	Metric	Shots	SLM						LLM																Closed LLM	
				S <sub>1</sub>	S <sub>2</sub>	S <sub>3</sub>	S <sub>4</sub>	S <sub>5</sub>	S <sub>6</sub>	L <sub>1</sub>	L <sub>2</sub>	L <sub>3</sub>	L <sub>4</sub>	L <sub>5</sub>	L <sub>6</sub>	L <sub>7</sub>	L <sub>8</sub>	L <sub>9</sub>	L <sub>10</sub>	L <sub>11</sub>	L <sub>12</sub>	L <sub>13</sub>	L <sub>14</sub>	L <sub>15</sub>	L <sub>16</sub>	CL <sub>1</sub>	CL <sub>2</sub>
Classification	xlni	Acc.	5	33.6	33.7	36.7	41.18	40.4	46.68	48.08	55.06	34.3	40.69	37.09	61.09	57.86	<u>65.58</u>	44.2	38.19	51.89	48.38	52.09	35.87	46.77	59.67	69	<b>69.6</b>
	lid	F <sub>1</sub>	10	0	0	0.87	0.26	0.85	0.77	2.82	3.6	0	1.49	2.27	1.95	<b>4.61</b>	4	1.81	3.34	3.26	3.07	3.3	3.05	3.32	4.57	3.73	4.48
	news	F <sub>1</sub>	3	0.86	0.35	0.15	17.54	0	6.64	21.94	0	2.06	24.06	0.12	25.83	7.83	<u>35.56</u>	28.03	0.61	35.26	33.31	34.81	3.45	30.46	28.66	<b>40.09</b>	32.19
	sentiment	F <sub>1</sub>	5	38.51	24.75	30.45	11.05	28.77	1.99	36.09	42.47	38.4	28.36	36.07	34.56	20.55	46.35	35.86	25.94	24.51	45.82	35.89	19.98	37.07	<u>47.14</u>	55.67	<b>57.52</b>
	topic	F <sub>1</sub>	3	8.73	31.62	20.3	19.79	12.46	17.15	48.11	39.35	34.95	28.55	42.12	44.3	31.8	<u>70.73</u>	36.47	43.74	64.89	56.48	48.73	52.09	50.96	53.16	67.93	<b>76.57</b>
	Avg.			16.34	18.08	17.69	17.96	16.50	14.65	31.41	28.10	21.94	24.63	23.53	33.55	24.53	<u>44.44</u>	29.27	22.36	35.96	37.41	34.96	22.89	33.72	38.64	47.28	<b>48.07</b>
Generation	mt_eng2xx	spBleu <sup>1k</sup>	5	1.79	3.63	3.45	3.25	2.12	3.4	7.07	7.61	1.72	4.37	2.79	6.24	6.91	4.56	3.73	3.29	8.84	8.7	<u>9.36</u>	1.4	6.06	8.22	12.18	<b>12.36</b>
	mt_fra2xx	spBleu <sup>1k</sup>	5	0.79	0.36	0.74	1.05	1.25	1.32	0.73	0.95	0.7	1.19	0.66	2.03	1.67	2.05	1.15	0.82	1.93	1.87	1.77	1.05	1.04	2.3	3.41	<b>3.42</b>
	mt_xx2xx	spBleu <sup>1k</sup>	5	0.4	0.45	0.42	0.42	0.36	0.18	0.43	0.63	0.72	0.55	0.46	0.71	0.9	<u>1.39</u>	0.55	0.17	1.28	1.05	1.37	0.67	0.38	0.85	<b>5.15</b>	4.44
	paraphrase	spBleu <sup>1k</sup>	5	23.2	28.88	31.15	19.12	25.93	17.85	28.73	25.83	31.68	31.33	23.66	17.26	21.34	15.37	32.13	30.68	26.86	19.64	27.79	19.23	<b>36.06</b>	26.06	23.2	20.35
	summary	rougeL	2	10.03	3.24	4.37	10.07	0.91	13.34	15.1	10.43	1.45	12.74	5.92	16.16	15.73	17.13	13.27	4.19	16.92	<u>17.19</u>	16.4	15.46	14.88	16.09	5.99	16.3
	title	spBleu <sup>1k</sup>	2	2.14	2.81	0.25	2.3	0.03	5.78	7.48	5.78	3	5.94	1.58	7.75	7.51	8.65	6.83	3.2	<u>11.07</u>	10.34	8.78	7.32	6.69	8.65	<b>13.62</b>	9.48
MCCR	mmu	Acc.	5	24.09	28.5	28.68	27.29	28.9	30.99	31.88	40.9	26.4	44.51	22.1	42.09	44.9	60.71	51.5	24	56.51	59.3	50.42	59.81	29.8	<b>61.91</b>	81.6	58.6
	msgm	ExactM	5	2.6	3.21	0	1.2	3.39	0.7	2.9	2.8	3.3	2.7	1.5	2.3	5.3	0.6	4.4	3	10.1	7.51	7.6	4.5	4.39	<u>11.1</u>	<b>45.9</b>	29.5
	belebele	Acc.	5	27.71	14.71	25.2	29.31	31.31	24.69	22.7	26.93	25.09	28.24	26.33	25.63	31.04	32.34	33.64	26.61	30.4	<u>35.92</u>	33.72	30.83	20.99	29.12	36.6	<b>39.7</b>
	squad_qa	F <sub>1</sub>	5	57.45	70.89	67.32	43.13	71.31	74.97	78.04	80.91	52.66	59.48	63.41	73.78	<u>82.49</u>	79.16	69.25	64.64	77.65	76.35	66.25	53.92	74.64	80.05	78.02	76.11
	Avg.			27.96	29.33	30.30	25.23	33.73	32.84	33.88	37.89	26.86	33.73	28.34	35.95	40.93	43.20	39.70	29.56	43.67	44.77	39.50	37.27	32.46	<u>45.55</u>	<b>60.53</b>	50.98
	Avg.			27.96	29.33	30.30	25.23	33.73	32.84	33.88	37.89	26.86	33.73	28.34	35.95	40.93	43.20	39.70	29.56	43.67	44.77	39.50	37.27	32.46	<u>45.55</u>	<b>60.53</b>	50.98
Tokens	ner	F <sub>1</sub>	5	0	0.78	1.76	2.99	0.13	3.65	2.26	7.49	2.35	3.96	0.2	6.03	13.14	11.4	8.51	1.52	6.86	10.54	8.64	6.04	4.3	<u>16.88</u>	<b>35.06</b>	27.04
	phrase	F <sub>1</sub>	5	14.79	28.5	29.44	11.52	24.91	27.73	15.95	29.35	29.42	16.55	30.87	13.59	30.01	23.97	26.34	23.09	25.55	23.13	16.78	25.13	29.44	<u>31.05</u>	<b>36.69</b>	36.45
	pos	F <sub>1</sub>	5	8.83	6.73	10.56	8.04	10.3	10.46	9.13	13.2	9.81	11.68	12.52	13.86	19	13.99	18.36	11.4	14.13	15.04	11.52	20.76	13.71	<u>27.54</u>	<b>62.84</b>	38.67
	Avg.			7.87	12.00	13.92	7.52	11.78	13.95	9.11	16.68	13.86	10.73	14.53	11.16	20.72	16.45	17.74	12.00	15.51	16.24	12.31	17.31	15.82	<u>25.16</u>	<b>44.86</b>	34.05
Overall				14.64	16.49	17.16	14.19	16.78	17.10	21.08	22.80	17.30	19.61	18.06	22.25	23.80	28.07	24.08	17.75	26.57	27.05	24.42	21.25	23.21	<u>29.93</u>	<b>40.82</b>	36.04

Table 2: Few-shot evaluation across different task clusters/tasks. The highest score for each individual task is in **bold green**, while the best average score across each task cluster (and **overall** score, calculated as average of task cluster averages) is in **bold orange**. Underline: refers to the best score across the open source/weights LLMs. **Shots**: number of shots. **SLMs**: **S<sub>1</sub>** Llama3.2 (1B), **S<sub>2</sub>** Llama3.2 (3B), **S<sub>3</sub>** Gemma2 (2B), **S<sub>4</sub>** Phi-3.5 (3.8B), **S<sub>5</sub>** Phi-4 (3.8B), and **S<sub>6</sub>** Gemma3 (4B). **LLMs**: **L<sub>1</sub>** Llama3.1 (8B), **L<sub>2</sub>** Gemma2 (9B), **L<sub>3</sub>** Aya (8B), **L<sub>4</sub>** Babel (9B), **L<sub>5</sub>** Command-R-Plus-R7B (8B), **L<sub>6</sub>** Gemma3 (12B), **L<sub>7</sub>** Gemma2 (27B), **L<sub>8</sub>** Gemma3 (27B), **L<sub>9</sub>** DeepSeek-R1-Distill-Qwen (32B), **L<sub>10</sub>** Aya (35B), **L<sub>11</sub>** Llama3.1 (70B), **L<sub>12</sub>** Llama3.3 (70B), **L<sub>13</sub>** DeepSeek-R1-Distill-Qwen (70B), **L<sub>14</sub>** Babel (83B), **L<sub>15</sub>** Command-R-Plus (104B), **L<sub>16</sub>** Command-A (111B). **Closed LLMs**: **CL<sub>1</sub>** Claude-4-Sonnet, **CL<sub>2</sub>** GPT-4.1. Table E4 is a full-page version of this table.

Gemma2 (9B, 12B) show consistent mid-to-high performance across various clusters, but do not surpass the larger models in aggregate. Importantly, translating *into* English is consistently easier for the models than translating *from* French, likely reflecting disparities in training data distribution and language representation on the target side. These findings highlight not only the dominance of closed LLMs and very large open models in multilingual tasks, but also the critical influence of resource availability and model scale on task-specific outcomes.

### Which LM should I use for reasoning tasks?

In a limited resource setting, for mathematical QA tasks that fall in the MCCR cluster, Phi-4 ranks top among SLMs for 16 African Languages. In the LLMs category, Command-A (111B) has better performance on a higher number of African languages. For reading comprehension tasks (belebele), Phi-4 is most suitable in constrained compute resource settings while in scenarios without limits in compute resources, Llama 3.3-70B gives better performance.

**Which African language(s) do LMs perform best on?** The models achieve their highest performance on a small set of well-resourced African

languages—chiefly Hausa, Swahili, Yorùbá, and Afrikaans—across diverse NLP tasks. Both open-source models like Command-A (111B) and Aya-35B and proprietary models such as Claude-4-Sonnet consistently excel on these languages, reflecting their ample labeled and parallel data. This trend spans simple classification and NER tasks to complex reasoning benchmarks (e.g., MMLU, XLNI) and machine translation, underscoring that data availability, rather than linguistic typology, is the primary driver of model effectiveness.

Models explicitly designed with African linguistic diversity, such as Aya-35B, also achieve strong results on Southern Sotho and Lingala—languages that have benefited from recent, community-led data initiatives. Swahili, in particular, ranks among the top languages for Llama3 models, likely due to its standardized orthography, regular morphology, and extensive bilingual corpora. Afrikaans and Swahili similarly perform well within Gemma2 and Command-R-Plus (104B), supported by their inclusion in massive multilingual datasets (e.g., mC4, OSCAR, WMT) and, in Afrikaans’s case, established government-supported digital media.

Overall, our findings reveal that LM performance correlates strongly with the volume and diversity of available data. Underperformance on

low-resource African languages is not a function of inherent complexity but of underrepresentation in training data. Bridging this gap will require concerted efforts to collect, annotate, and standardize datasets that capture the dialectal, sociolinguistic, and orthographic richness of all African languages—an essential step toward improved model equity and real-world utility across the continent’s full linguistic spectrum.

## 6 Data-Performance-Policy Analysis

To understand the relationship between policy, data availability, and performance, we carry out a number of analyses inspired by SAHARA data and empirical results on it.

**Data Availability and Diversity.** As shown in Figure E3, among the 517 languages in our benchmark, only 45 have more than one dataset, while the majority are represented solely by language identification data. Amharic leads with 11 unique datasets across all clusters, followed by Yorùbá with ten and Hausa with nine. A total of six languages have only one dataset beyond language identification, while the remaining 36 have between two and eight datasets each. *It is important to note that the disparity in dataset availability is closely linked to language policy.* The 45 languages with multiple datasets generally hold official status or are spoken by a significant majority in their respective countries. Government policies that designate certain languages as official often yield greater institutional support, driving increased documentation, education, and media presence. For example, Amharic’s official status in Ethiopia contributes to its relatively rich dataset availability, while Yorùbá and Hausa, both widely used in Nigeria for regional governance and media, similarly benefit from enhanced linguistic resources. Conversely, languages lacking official recognition or widespread institutional support receive minimal investment in digital and linguistic infrastructure. The dominance of a few languages within policy frameworks perpetuates data scarcity for many Indigenous African languages, thereby limiting their inclusion in NLP advancements and reinforcing digital inequalities.

It is noteworthy that availability of NLP data is not necessarily determined by the number of speakers and that other variables including language prestige, literacy rates, and economic incentives can be involved. This becomes evident when comparing speaker populations with actual

data availability. For example, Hausa (73 million speakers), Swahili (80 million), and Zulu (28 million) are categorized as *hopefuls*—languages with some available labeled datasets. In contrast, Naija (Nigerian Pidgin)—despite having approximately 153 million speakers—is classified as a *left-behind* language (Joshi et al., 2020), largely due to decades of stigmatization and negative public attitudes (Oyebola and Ugwuanyi, 2023), which have significantly contributed to the scarcity of accessible and annotated linguistic resources. Notably, several non-African languages with significantly smaller speaker populations—such as Catalan (5 million), Finnish (10 million), and Swedish (10.5 million)—are regarded as high-resource, primarily as a result of sustained efforts in documentation, digitization, and technological integration. This discrepancy underscores a broader systemic issue in language representation: *it is not just linguistic demographics that shape digital inclusion, but also factors such as historical language prestige, policy decisions, digital infrastructure, and the extent of prior documentation.*

These intersecting factors ultimately determine whether a language is integrated into AI systems or remains marginalized. Addressing this imbalance requires a shift in policy and practice to prioritize equitable digital representation. This includes targeted investments in multilingual data collection, increased funding for research on historically excluded and low-resource languages, and the intentional integration of these languages into education systems, media platforms, and digital technologies. Such efforts must also consider the sociopolitical and infrastructural barriers that shape language inclusion in AI, ensuring that underrepresented languages are not only documented but actively supported within the broader digital ecosystem.

**Data Quality.** Most of the downstream task datasets are limited to simple tasks that require merely an atomic class label at the word, sentence, or document level (e.g., sentiment analysis, named entity recognition, topic classification). Furthermore, many of these datasets are translations of existing resources or content from high-resource languages like English and French, and they do not always reflect authentic language usage within the respective communities. For instance, AfriXLNI, AfriMMLU, and AfriMGSM (Adelani et al., 2024b) were created by translating XLNI, MMLU, and MGSM into African languages. In

these datasets, the frequent use of borrowed words and phonologized variants of foreign terms (Adebara and Abdul-Mageed, 2022) highlights the challenges posed by cross-lingual concept gaps. When translations are applied to classification tasks, misalignments in the attached labels can occur because labels are not always equivalent across languages. Some labels may not exist in the target language or may have restricted usage. For example, while English possesses a productive set of adjectives, languages like Yorùbá have a more limited adjective inventory that is context-dependent, meaning that an English adjective may be translated as a noun or verb in Yorùbá rather than as an adjective.

**Model Performance** Here, we analyze performance of our best model, Claude-4-Sonnet. As shown in Figure 2, the model faces varying degrees of difficulty when handling African languages. Figure 2(a) demonstrates that for generation tasks, even the strongest models struggle to produce high-quality text in African languages, with most scores falling below 15 BLEU/ROUGE. In contrast, Figure 2(b) shows that while many languages achieve relatively strong classification performance (over 80% accuracy), a substantial portion still falls into the 60%–80% range, underscoring a clear performance gap. Finally, Figures 2(c) and 2(d) reveal that translation into English is notably easier than translation out of English or French into African languages, where BLEU scores drop from an average of around 19 down to nine or 11, respectively. This finding suggests that even the best model struggles more with producing fluent, faithful output in African languages than with understanding them—underscoring the urgent need for better data and model development tailored to these languages.

Apart from our empirical analyses, the limited availability and diversity of datasets for African languages in NLP benchmarks has significant implications for developing robust language technologies. *Limited Model Performance and Generalization* – Models trained on a single dataset per language risk overfitting to specific domains or styles, reducing their ability to generalize across different contexts, which is especially problematic for African languages with high dialectal variation. *Bias Toward Well-Resourced Languages* – Languages with more datasets naturally perform better in multilingual models, reinforcing disparities and leading to poor performance for underrepresented

African languages. *Restricted Downstream Applications* – Many NLP applications, such as machine translation and question answering, require diverse datasets; if a language is represented only by language identification data, it cannot support the development of more complex language technologies. *Challenges in Transfer Learning* – Transfer learning techniques, like multilingual fine-tuning, rely on cross-lingual similarities and sufficient data; when African languages suffer from sparse training data, they are unable to effectively leverage knowledge transfer from resource-rich languages. We also analyze the performance of Command-A in Figure E2.

## 7 Recommendations


Policy determines how languages are used across various domains, influencing their visibility and viability in digital spaces. Currently, the limited online presence of many Indigenous African languages restricts the availability of training data, which is crucial for building effective NLP models. We provide the following recommendations:

**Policy matters.** Governments and policymakers should enact national language policies that mandate the inclusion of Indigenous and underrepresented languages across digital infrastructure, education, and media. Countries with multilingual populations should adopt policies that support the systematic documentation, transcription, and digitization of African languages to enhance their presence in AI applications. Such efforts will enrich the diversity of available data, paving the way for the development of more robust and inclusive models.


**Expand data collection and annotation.** Given that only 45 of the 517 languages in the benchmark have more than one dataset based on our collection, it is essential to increase the availability of high-quality data for underrepresented African languages across a broad range of downstream tasks. Moreover, prioritizing community-driven data annotation efforts will be crucial in ensuring both linguistic diversity and accuracy, especially for languages that lack institutional support.

**Develop contextually relevant and culturally accurate datasets.** Many existing African language datasets are derived from translations of high-resource languages, which can introduce misalignments and fail to capture authentic linguistic expressions. To improve NLP performance, datasets should be created using native speakers







(a) Generation tasks.



(b) Classification tasks.



(c) MCCR tasks.



(d) MT tasks grouped.

Figure 2: Distribution of best model performance (Claude-4-Sonnet) on different task clusters across African languages in downstream data.

and sourced from real-world interactions, literature, and media in African languages. This will help preserve linguistic authenticity and avoid the overreliance on loanwords or structures that do not naturally occur in the language. Specialized corpora should also be developed to capture dialectal variations and culturally specific language use.

## 8 Conclusion

In this work, we set out to unravel how current African language policies, by dictating data availability and accessibility, directly influence model performance across the continent’s diverse linguistic spectrum. To achieve this, we conducted a comprehensive empirical evaluation of leading large language models (LLMs) on SAHARA, an extensive benchmark assembled primarily from existing

datasets. SAHARA consists of 30 publicly accessible datasets, covering 16 classification, generation, MCCR, and token level tasks. We also introduced the first publicly available leaderboard for African languages. Our evaluations revealed critical gaps in task performance and language coverage, underscoring pronounced disparities rooted in policy-induced data inequities. Moreover, our findings highlight the importance of integrating theoretical insights on language policies with empirical evaluations to inform forward-thinking policy reforms and promote inclusive data practices. We believe that SAHARA will serve as essential benchmarking standards, catalyzing future research and development efforts aimed at bridging the digital divide and fostering linguistic diversity in artificial intelligence for African communities.

## 9 Limitations

We identify the following limitations for our work:

1. **Data and Data Availability Bias.** Our benchmark relies on publicly available datasets, which inherently reflect existing biases in language documentation and institutional support. Many Indigenous African languages remain severely underrepresented due to historical policies that prioritize dominant languages. As a result, our findings may not fully capture the linguistic diversity of the continent.
2. **Task Scope.** The evaluation primarily focuses on tasks where data is available, such as language identification and text classification. However, the scarcity of datasets for complex NLP tasks (e.g., machine translation, question answering) limits our ability to assess model performance in broader applications. Future efforts should address this gap by developing more comprehensive datasets across a wider range of tasks.
3. **Model Performance Constraints.** While we analyze the impact of data scarcity on model effectiveness, our study does not explore architectural modifications or fine-tuning strategies that could improve performance for low-resource languages. Further research is needed to examine whether adaptive techniques, such as multilingual pretraining with targeted data augmentation, could mitigate these challenges.
4. **Policy Implementation Challenges.** Our recommendations for policy reforms and inclusive data practices require significant institutional commitment and resources. While we outline actionable steps, their adoption depends on the willingness of governments, research institutions, and technology developers to prioritize African language inclusion in NLP.

Despite these limitations, our work underscores the urgency of addressing data inequities and advocating for policies that foster greater linguistic diversity in AI development.

## 10 Ethical Considerations

We make the following ethics-related statements about our work:

1. Our research aims to investigate the progress of NLP for African languages, addressing historical language policies that have contributed to data inequities. By improving representation in NLP, we align with broader efforts to foster linguistic diversity in AI, ensuring that African languages are not sidelined in technological advancements.
2. The datasets used in our benchmark are derived from publicly available sources, yet their existence is shaped by historical and contemporary policies that prioritize certain languages over others. We acknowledge that the digital presence of African languages is not merely a technical issue but a policy-driven reality that influences which languages receive institutional support and investment.
3. Although we do not develop models in this work, our findings highlight the impact of policy-induced data disparities on model performance. Addressing these challenges requires policy interventions that support multilingual data collection, equitable resource allocation, and ethical considerations in dataset creation.
4. Proper attribution to dataset authors is not just an academic necessity but a policy imperative for transparency and recognition. To mitigate concerns about data ownership and ethical usage, we provide a publicly accessible reference file citing all datasets included in our benchmark. We strongly encourage researchers and policymakers to acknowledge these sources, reinforcing the importance of ethical and inclusive data practices.

## Acknowledgments

We acknowledge support from Canada Research Chairs (CRC), CLEAR Global funding from the Gates Foundation, the Natural Sciences and Engineering Research Council of Canada (NSERC; RGPIN-2018-04267), the Social Sciences and Humanities Research Council of Canada (SSHRC; 895-2020-1004; 895-2021-1008), Canadian Foundation for Innovation (CFI; 37771), Digital Research Alliance of Canada,<sup>5</sup> and UBC ARC-Sockeye.<sup>6</sup> We also thank Clémentine Fourier and

<sup>5</sup><https://alliancecan.ca>

<sup>6</sup><https://arc.ubc.ca/ubc-arc-sockeye>

Nouamane Tazi from HuggingFace for providing feedback on this work. The findings and conclusions contained within this work those of the authors and do not necessarily reflect positions or policies of any supporters.

## References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Qin Cai, Vishrav Chaudhary, Dong Chen, Dongdong Chen, Weizhu Chen, Yen-Chun Chen, Yi-Ling Chen, Hao Cheng, Parul Chopra, Xiyang Dai, Matthew Dixon, Ronen Eldan, Victor Fragoso, Jianfeng Gao, Mei Gao, Min Gao, Amit Garg, Allie Del Giorno, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Wenxiang Hu, Jamie Huynh, Dan Iter, Sam Ade Jacobs, Mojan Javaheripi, Xin Jin, Nikos Karampatziakis, Piero Kauffmann, Mahoud Khademi, Dongwoo Kim, Young Jin Kim, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Yunsheng Li, Chen Liang, Lars Liden, Xihui Lin, Zeqi Lin, Ce Liu, Liyuan Liu, Mengchen Liu, Weishung Liu, Xiaodong Liu, Chong Luo, Piyush Madan, Ali Mahmoudzadeh, David Majercak, Matt Mazzola, Caio César Teodoro Mendes, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Liliang Ren, Gustavo de Rosa, Corby Rosset, Sambudha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacrose, Shital Shah, Ning Shang, Hiteshi Sharma, Yelong Shen, Swadheen Shukla, Xia Song, Masahiro Tanaka, Andrea Tupini, Praneetha Vaddamanu, Chunyu Wang, Guanhua Wang, Lijuan Wang, Shuohang Wang, Xin Wang, Yu Wang, Rachel Ward, Wen Wen, Philipp Witte, Haiping Wu, Xiaoxia Wu, Michael Wyatt, Bin Xiao, Can Xu, Jiahang Xu, Weijian Xu, Jilong Xue, Sonali Yadav, Fan Yang, Jianwei Yang, Yifan Yang, Ziyi Yang, Donghan Yu, Lu Yuan, Chenruidong Zhang, Cyril Zhang, Jianwen Zhang, Li Lyna Zhang, Yi Zhang, Yue Zhang, Yunan Zhang, and Xiren Zhou. 2024. [Phi-3 technical report: A highly capable language model locally on your phone](#).
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Ife Adebara and Muhammad Abdul-Mageed. 2022. [Towards afrocentric NLP for African languages: Where we are and where we can go](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3814–3841, Dublin, Ireland. Association for Computational Linguistics.
- Ife Adebara, AbdelRahim Elmadany, and Muhammad Abdul-Mageed. 2024a. [Cheetah: Natural language generation for 517 African languages](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12798–12823, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Ife Adebara, AbdelRahim Elmadany, and Muhammad Abdul-Mageed. 2024b. [Cheetah: Natural language generation for 517 african languages](#).
- Ife Adebara, AbdelRahim Elmadany, Muhammad Abdul-Mageed, and Alcides Alcoba Inciarte. 2023. [SERENGETI: Massively multilingual language models for Africa](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1498–1537, Toronto, Canada. Association for Computational Linguistics.
- Ife Adebara, AbdelRahim Elmadany, Muhammad Abdul-Mageed, and Alcides Alcoba Inciarte. 2022. [Afrolid: A neural language identification tool for african languages](#).
- D. Adelani, Jade Abbott, Graham Neubig, Daniel D’Souza, Julia Kreutzer, Constantine Lignos, Chester Palen-Michel, Happy Buzaaba, Shruti Rijhwani, Sebastian Ruder, Stephen Mayhew, Israel Abebe Azime, S. Muhammad, Chris C. Emezue, Joyce Nakatumba-Nabende, Perez Ogayo, Anuoluwapo Aremu, Catherine Gitau, Derguene Mbaye, J. Alabi, Seid Muhie Yimam, Tajuddeen R. Gwadabe, Ignatius Ezeani, Rubungo Andre Niyongabo, Jonathan Mukiibi, V. Otiende, Iroko Orife, Davis David, Samba Ngom, Tosin P. Adewumi, Paul Rayson, Mofetoluwa Adeyemi, Gerald Muriuki, Emmanuel Anebi, C. Chukwuneke, N. Odu, Eric Peter Wairagala, S. Oyerinde, Clemencia Siro, Tobius Saul Bateesa, Temilola Oloyede, Yvonne Wambui, Victor Akinode, Deborah Nabagereka, Maurice Katusiime, Ayodele Awokoya, Mouhamadane Mboup, D. Gebreyohannes, Henok Tilaye, Kelechi Nwaike, Degaga Wolde, Abdoulaye Faye, Blessing Sibanda, Orevaoghene Ahia, Bonaventure F. P. Dossou, Kelechi Ogueji, Thierno Ibrahima Diop, A. Diallo, Adewale Akinfaderin, T. Marengereke, and Salomey Osei. 2021a. Masakhaner: Named entity recognition for african languages. *ArXiv*, abs/2103.11811.
- David Adelani, Jesujoba Alabi, Angela Fan, Julia Kreutzer, Xiaoyu Shen, Machel Reid, Dana Ruiter, Dietrich Klakow, Peter Nabende, Ernie Chang, Tajuddeen Gwadabe, Freshia Sackey, Bonaventure F. P. Dossou, Chris Emezue, Colin Leong, Michael Beukman, Shamsuddeen Muhammad, Guyo Jarso, Oreen Yousuf, Andre Niyongabo Rubungo, Gilles Hacheme, Eric Peter Wairagala, Muhammad Umair Nasir, Benjamin Ajibade, Tunde Ajayi, Yvonne Gitau, Jade Abbott, Mohamed Ahmed, Millicent Ochieng, Anuoluwapo Aremu, Perez Ogayo, Jonathan Mukiibi, Fatoumata Ouoba Kabore, Godson Kalipe, Derguene

- Mbaye, Allahsera Auguste Tapo, Victoire Memdjokam Koagne, Edwin Munkoh-Buabeng, Valencia Wagner, Idris Abdulmumin, Ayodele Awokoya, Happy Buzaaba, Blessing Sibanda, Andiswa Bukula, and Sam Manthalu. 2022a. [A few thousand translations go a long way! leveraging pre-trained models for African news translation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3053–3070, Seattle, United States. Association for Computational Linguistics.
- David Adelani, Hannah Liu, Xiaoyu Shen, Nikita Vassilyev, Jesujoba Alabi, Yanke Mao, Haonan Gao, and En-Shiun Lee. 2024a. [SIB-200: A simple, inclusive, and big evaluation dataset for topic classification in 200+ languages and dialects](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 226–245, St. Julian’s, Malta. Association for Computational Linguistics.
- David Adelani, Graham Neubig, Sebastian Ruder, Shruti Rijhwani, Michael Beukman, Chester Palen-Michel, Constantine Lignos, Jesujoba Alabi, Shamsuddeen Muhammad, Peter Nabende, Cheikh M. Bamba Dione, Andiswa Bukula, Rooweither Mabuya, Bonaventure F. P. Dossou, Blessing Sibanda, Happy Buzaaba, Jonathan Mukiibi, Godson Kalipe, Derguene Mbaye, Amelia Taylor, Fatoumata Kabore, Chris Chinenye Emezue, Anuoluwapo Aremu, Perez Ogayo, Catherine Gitau, Edwin Munkoh-Buabeng, Victoire Memdjokam Koagne, Allahsera Auguste Tapo, Tebogo Macucwa, Vukosi Marivate, Mboning Tchiaze Elvis, Tajuddeen Gwadabe, Tosin Adewumi, Orevaoghene Ahia, and Joyce Nakatumba-Nabende. 2022b. [MasakhaNER 2.0: Africa-centric transfer learning for named entity recognition](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4488–4508, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- David Adelani, Dana Ruiter, Jesujoba Alabi, Damilola Adebajo, Adesina Ayeni, Mofe Adeyemi, Ayodele Esther Awokoya, and Cristina España-Bonet. 2021b. [The effect of domain and diacritics in Yoruba-English neural machine translation](#). In *Proceedings of Machine Translation Summit XVIII: Research Track*, pages 61–75, Virtual. Association for Machine Translation in the Americas.
- David Ifeoluwa Adelani, Jade Abbott, Graham Neubig, Daniel D’souza, Julia Kreutzer, Constantine Lignos, Chester Palen-Michel, Happy Buzaaba, Shruti Rijhwani, Sebastian Ruder, Stephen Mayhew, Israel Abebe Azime, Shamsuddeen H. Muhammad, Chris Chinenye Emezue, Joyce Nakatumba-Nabende, Perez Ogayo, Aremu Anuoluwapo, Catherine Gitau, Derguene Mbaye, Jesujoba Alabi, Seid Muhie Yimam, Tajuddeen Rabi Gwadabe, Ignatius Ezeani, Rubungo Andre Niyongabo, Jonathan Mukiibi, Verah Otiende, Iroko Orife, Davis David, Samba Ngom, Tosin Adewumi, Paul Rayson, Mofetoluwa Adeyemi, Gerald Muriuki, Emmanuel Anebi, Chiamaka Chukwuneke, Nkiruka Odu, Eric Peter Wairagala, Samuel Oyerinde, Clemencia Siro, Tobius Saul Bateesa, Temilola Oloyede, Yvonne Wambui, Victor Akinode, Deborah Nabagereka, Maurice Katusiime, Ayodele Awokoya, Mouhamadane MBOUP, Dibora Gebreyohannes, Henok Tilaye, Kelechi Nwaike, Degaga Wolde, Abdoulaye Faye, Blessing Sibanda, Orevaoghene Ahia, Bonaventure F. P. Dossou, Kelechi Ogueji, Thierno Ibrahima DIOP, Abdoulaye Diallo, Adewale Akinfaderin, Tendai Marengereke, and Salomey Osei. 2021c. [MasakhaNER: Named entity recognition for African languages](#). *Transactions of the Association for Computational Linguistics*, 9:1116–1131.
- David Ifeoluwa Adelani, Marek Masiak, Israel Abebe Azime, Jesujoba Alabi, Atnafu Lambebo Tonja, Christine Mwase, Odunayo Ogundepo, Bonaventure F. P. Dossou, Akintunde Oladipo, Doreen Nixdorf, Chris Chinenye Emezue, Sana Al-azzawi, Blessing Sibanda, Davis David, Lolwethu Ndoela, Jonathan Mukiibi, Tunde Ajayi, Tatiana Moteu, Brian Odhiambo, Abraham Owodunni, Nnaemeka Obiefuna, Muhidin Mohamed, Shamsuddeen Hassan Muhammad, Teshome Mulugeta Ababu, Saheed Abdullahi Salahudeen, Mesay Gemeda Yigezu, Tajuddeen Gwadabe, Idris Abdulmumin, Mahlet Taye, Oluwabusayo Awoyomi, Iyanuoluwa Shode, Tolulope Adelani, Habiba Abdulganiyu, Abdul-Hakeem Omatayo, Adetola Adeeko, Abeeb Afolabi, Anuoluwapo Aremu, Olanrewaju Samuel, Clemencia Siro, Wangari Kimotho, Onyekachi Ogbu, Chinedu Mbonu, Chiamaka Chukwuneke, Samuel Fanijo, Jessica Ojo, Oyinkansola Awosan, Tadesse Kebede, Toadoun Sari Sakayo, Pamela Nyatsine, Freedom Sidume, Oreen Yousuf, Mardiyyah Odwole, Kanda Tshinu, Ussen Kimanuka, Thina Diko, Siyanda Nxakama, Sinodos Nigusse, Abdulmejid Johar, Shafie Mohamed, Fuad Mire Hassan, Moges Ahmed Mehamed, Evrard Ngabire, Jules Jules, Ivan Ssenkungu, and Pontus Stenetorp. 2023. [MasakhaNEWS: News topic classification for African languages](#). In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 144–159, Nusa Dua, Bali. Association for Computational Linguistics.
- David Ifeoluwa Adelani, Graham Neubig, Sebastian Ruder, Shruti Rijhwani, Michael Beukman, Chester Palen-Michel, Constantine Lignos, Jesujoba Oluwadara Alabi, Shamsuddeen Hassan Muhammad, Peter Nabende, Cheikh M. Bamba Dione, Andiswa Bukula, Rooweither Mabuya, Bonaventure F. P. Dossou, Blessing K. Sibanda, Happy Buzaaba, Jonathan Mukiibi, Godson Kalipe, Derguene Mbaye, Amelia Taylor, Fatoumata Kabore, Chris C. Emezue, Anuoluwapo Aremu, Perez Ogayo, Catherine W. Gitau, Edwin Munkoh-Buabeng, Victoire Memdjokam Koagne, Allahsera Auguste Tapo,



- Tebogo Macucwa, Vukosi Marivate, Elvis Mboning, Tajuddeen R. Gwadabe, Tosin P. Adewumi, Orevaoghene Ahia, Joyce Nakatumba-Nabende, Neo L. Mokono, Ignatius M Ezeani, Chiamaka Ijeoma Chukwuneke, Mofetoluwa Adeyemi, Gilles Hacheme, Idris Abdulmumin, Odunayo Ogundepo, Oreen Yousuf, Tatiana Moteu Ngoli, and Dietrich Klakow. 2022c. Masakhaner 2.0: Africa-centric transfer learning for named entity recognition. *ArXiv*, abs/2210.12391.
- 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 Chukwuneke, Happy Buzaaba, Blessing Sibanda, Godson Kalipe, Jonathan Mukiibi, Salomon Kabongo, Foutse Yuehgo, Mmasibidi Setaka, Lolwethu Ndolela, Nkiruka Odu, Rooweither Mabuya, Shamsuddeen Hassan Muhammad, Salomey Osei, Sokhar Samb, Tadesse Kebede Guge, and Pontus Stenertorp. 2024b. [Irokobench: A new benchmark for african languages in the age of large language models](#).
- Benjamin Akera, Jonathan Mukiibi, Lydia Sanyu Nagayi, Claire Babirye, Isaac Owomugisha, Solomon Nsumba, Joyce Nakatumba-Nabende, Engineer Bainomugisha, Ernest Mwebaze, and John Quinn. 2022. [Machine translation for african languages: Community creation of datasets and models in uganda](#). In *3rd Workshop on African Natural Language Processing*.
- Jesujoba Alabi, Kwabena Amponsah-Kaakyire, David Adelani, and Cristina España-Bonet. 2020a. [Massive vs. curated embeddings for low-resourced languages: the case of Yorùbá and Twi](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2754–2762, Marseille, France. European Language Resources Association.
- Jesujoba Alabi, Kwabena Amponsah-Kaakyire, David Adelani, and Cristina España-Bonet. 2020b. [Massive vs. curated embeddings for low-resourced languages: The case of Yorùbá and Twi](#). In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 2754–2762.
- Jesujoba O. Alabi, David Ifeoluwa Adelani, Marius Mosbach, and Dietrich Klakow. 2022. [Adapting pre-trained language models to African languages via multilingual adaptive fine-tuning](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4336–4349, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Saminu Mohammad Aliyu, Gregory Maksha Wajiga, Muhammad Murtala, Shamsuddeen Hassan Muhammad, Idris Abdulmumin, and Ibrahim Said Ahmad. 2022. [Herdphobia: A dataset for hate speech against fulani in nigeria](#).
- Anthropic. 2024. [The claude 3 model family: Opus, sonnet, haiku anthropic](#). Technical report, Anthropic.
- Viraat Aryabumi, John Dang, Dwarak Talupuru, Saurabh Dash, David Cairuz, Hangyu Lin, Bharat Venkitesh, Madeline Smith, Jon Ander Campos, Yi Chern Tan, Kelly Marchisio, Max Bartolo, Sebastian Ruder, Acyr Locatelli, Julia Kreutzer, Nick Frosst, Aidan Gomez, Phil Blunsom, Marzieh Fadaee, Ahmet Üstün, and Sara Hooker. 2024. [Aya 23: Open weight releases to further multilingual progress](#).
- Israel Abebe Azime and Nebil Mohammed. 2021. [An amharic news text classification dataset](#).
- Lucas Bandarkar, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer, and Madian Khabsa. 2024a. [The belebele benchmark: a parallel reading comprehension dataset in 122 language variants](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 749–775, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Lucas Bandarkar, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer, and Madian Khabsa. 2024b. [The belebele benchmark: a parallel reading comprehension dataset in 122 language variants](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, page 749–775. Association for Computational Linguistics.
- Abhik Bhattacharjee, Tahmid Hasan, Wasi Uddin Ahmad, and Rifat Shahriyar. 2023. [BanglaNLG and BanglaT5: Benchmarks and resources for evaluating low-resource natural language generation in Bangla](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 726–735, Dubrovnik, Croatia. Association for Computational Linguistics.
- Steven Bird. 2020. [Decolonising speech and language technology](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3504–3519, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Jiaao Chen, Derek Tam, Colin Raffel, Mohit Bansal, and Diyi Yang. 2023. [An Empirical Survey of Data Augmentation for Limited Data Learning in NLP](#). *Transactions of the Association for Computational Linguistics*, 11:191–211.
- Wei-Rui Chen, Ife Adebare, Khai Doan, Qisheng Liao, and Muhammad Abdul-Mageed. 2024. [Fumbling in babel: An investigation into ChatGPT’s language identification ability](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 4387–4413, Mexico City, Mexico. Association for Computational Linguistics.
- Victor Ngu Cheo, Esther Phubon Chie, and Yollande Menguie. 2023. [An evaluation of media use of indigenous languages in cameroon](#). *The International Journal of African Language and Media Studies (IJALMS)*, 30:1–20.

- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. [TyDi QA: A benchmark for information-seeking question answering in typologically diverse languages](#). *Transactions of the Association for Computational Linguistics*, 8:454–470.
- Team Cohere, Arash Ahmadian, Marwan Ahmed, Jay Alamar, Milad Alizadeh, Yazeed Alnumay, Sophia Althammer, Arkady Arkhangorodsky, Viraat Aryabumi, Dennis Aumiller, et al. 2025. Command a: An enterprise-ready large language model. *arXiv preprint arXiv:2504.00698*.
- 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.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. Xnli: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Marta R Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, et al. 2022. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*.
- Davis David. 2020a. [Swahili : News classification dataset](#).
- Davis David. 2020b. [Swahili : News classification dataset](#). The news version contains both train and test sets.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Mountaga Diallo, Chayma Fourati, and Hatem Hadad. 2021. [Bambara language dataset for sentiment analysis](#).
- Cheikh M. Bamba Dione, David Adelani, Peter Nabende, Jesujoba Alabi, Thapelo Sindane, Happy Buzaaba, Shamsuddeen Hassan Muhammad, Chris Chinenye Emezue, Perez Ogayo, Anuoluwapo Aremu, Catherine Gitau, Derguene Mbaye, Jonathan Mukiibi, Blessing Sibanda, Bonaventure F. P. Dossou, Andiswa Bukula, Rooweither Mabuya, Allahsera Auguste Tapo, Edwin Munkoh-Buabeng, victoire Memdjokam Koagne, Fatoumata Ouoba Kabore, Amelia Taylor, Godson Kalipe, Tebogo Macucwa, Vukosi Marivate, Tajuddeen Gwadabe, Mboning Tchiazze Elvis, Ikechukwu Onyenwe, Gratien Atindogbe, Tolulope Adelani, Idris Akinade, Olanrewaju Samuel, Marien Nahimana, Th'eogene Musabeyezu, Emile Niyomutabazi, Ester Chimbenga, Kudzai Gotosa, Patrick Mizha, Apelete Agbollo, Seydou Traore, Chinedu Uchechukwu, Aliyu Yusuf, Muhammad Abdullahi, and Dietrich Klakow. 2023a. MasakhaPOS: Part-of-speech tagging for typologically diverse african languages.
- Cheikh M. Bamba Dione, David Ifeoluwa Adelani, Peter Nabende, Jesujoba Alabi, Thapelo Sindane, Happy Buzaaba, Shamsuddeen Hassan Muhammad, Chris Chinenye Emezue, Perez Ogayo, Anuoluwapo Aremu, Catherine Gitau, Derguene Mbaye, Jonathan Mukiibi, Blessing Sibanda, Bonaventure F. P. Dossou, Andiswa Bukula, Rooweither Mabuya, Allahsera Auguste Tapo, Edwin Munkoh-Buabeng, victoire Memdjokam Koagne, Fatoumata Ouoba Kabore, Amelia Taylor, Godson Kalipe, Tebogo Macucwa, Vukosi Marivate, Tajuddeen Gwadabe, Mboning Tchiazze Elvis, Ikechukwu Onyenwe, Gratien Atindogbe, Tolulope Adelani, Idris Akinade, Olanrewaju Samuel, Marien Nahimana, Théogène Musabeyezu, Emile Niyomutabazi, Ester Chimbenga, Kudzai Gotosa, Patrick Mizha, Apelete Agbollo, Seydou Traore, Chinedu Uchechukwu, Aliyu Yusuf, Muhammad Abdullahi, and Dietrich Klakow. 2023b. [MasakhaPOS: Part-of-speech tagging for typologically diverse African languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10883–10900, Toronto, Canada. Association for Computational Linguistics.
- Long Doan, Linh The Nguyen, Nguyen Luong Tran, Thai Hoang, and Dat Quoc Nguyen. 2021. [PhoMT: A high-quality and large-scale benchmark dataset for Vietnamese-English machine translation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4495–4503, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Al-lonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Gregoire Mi-

alon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shao-liang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gouget, Virginie Do, Vish Voleti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyan Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi

Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel, Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind Thattai, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Ibrahim Damla, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khadkelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal Chawla, Kushal Lakhotia, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Maria Tsimpoukelli, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong, Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Maheswari, Russ Howes, Ruty Rinott, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang

- Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vitor Albiero, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaofang Wang, Xiaojian Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. 2024. [The llama 3 herd of models](#).
- David M Eberhard, F Simons Gary, and Charles D Fenning (eds). 2021. [Ethnologue: Languages of the world. Twenty-fourth edition](#), Dallas, Texas: SIL International.
- Roald Eiselen. 2016. [Government domain named entity recognition for South African languages](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3344–3348, Portorož, Slovenia. European Language Resources Association (ELRA).
- AbdelRahim Elmadany, Ife Adebara, and Muhammad Abdul-Mageed. 2024. [Toucan: Many-to-many translation for 150 African language pairs](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 13189–13206, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- AbdelRahim Elmadany, ElMoatez Billah Nagoudi, and Muhammad Abdul-Mageed. 2023. [ORCA: A challenging benchmark for Arabic language understanding](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9559–9586, Toronto, Canada. Association for Computational Linguistics.
- Danny S Foster. 2021. *Language of Instruction in Rural Tanzania: A Critical Analysis of Parents' Discursive Practices and Valued Linguistic Capabilities*. Ph.D. thesis, University of Bristol.
- Sebastian Gehrmann, Tosin Adewumi, Karmanya Aggarwal, Pawan Sasanka Ammanamanchi, Anuoluwapo Aremu, Antoine Bosselut, Khyathi Raghavi Chandu, Miruna-Adriana Clinciu, Dipanjan Das, Kaustubh Dhole, Wanyu Du, Esin Durmus, Ondřej Dušek, Chris Chinenye Emezue, Varun Gangal, Cristina Garbacea, Tatsunori Hashimoto, Yufang Hou, Yacine Jernite, Harsh Jhamtani, Yangfeng Ji, Shailza Jolly, Mihir Kale, Dhruv Kumar, Faisal Ladhak, Aman Madaan, Mounica Maddela, Khyati Mahajan, Saad Mahamood, Bodhisattwa Prasad Majumder, Pedro Henrique Martins, Angelina McMillan-Major, Simon Mille, Emiel van Miltenburg, Moin Nadeem, Shashi Narayan, Vitaly Nikolaev, Andre Niyongabo Rubungo, Salomey Osei, Ankur Parikh, Laura Perez-Beltrachini, Niranjana Ramesh Rao, Vikas Raunak, Juan Diego Rodriguez, Sashank Santhanam, João Sedoc, Thibault Sellam, Samira Shaikh, Anastasia Shimorina, Marco Antonio Sobrevilla Cabeza, Hendrik Strobelt, Nishant Subramani, Wei Xu, Diyi Yang, Akhila Yerukola, and Jiawei Zhou. 2021. [The GEM benchmark: Natural language generation, its evaluation and metrics](#). In *Proceedings of the 1st Workshop on Natural Language Generation, Evaluation, and Metrics (GEM 2021)*, pages 96–120, Online. Association for Computational Linguistics.
- Jian Guan, Zhuoer Feng, Yamei Chen, Ruilin He, Xiaoxi Mao, Changjie Fan, and Minlie Huang. 2022. [LOT: A Story-Centric Benchmark for Evaluating Chinese Long Text Understanding and Generation](#). *Transactions of the Association for Computational Linguistics*, 10:434–451.
- Tri Wahyu Guntara, Alham Fikri Aji, and Radityo Eko Prasajo. 2020. [Benchmarking multidomain English-Indonesian machine translation](#). In *Proceedings of the 13th Workshop on Building and Using Comparable Corpora*, pages 35–43, Marseille, France. European Language Resources Association.
- Jiyeon Ham, Yo Joong Choe, Kyubyong Park, Ilji Choi, and Hyungjoon Soh. 2020. [KorNLI and KorSTS: New Benchmark Datasets for Korean Natural Language Understanding](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 422–430.
- Tahmid Hasan, Abhik Bhattacharjee, Md. Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin Kang, M. Sohel Rahman, and Rifat Shahriyar. 2021. [XLsum: Large-scale multilingual abstractive summarization for 44 languages](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4693–4703, Online. Association for Computational Linguistics.
- Michael A. Hedderich, David Adelani, Dawei Zhu, Jesujoba Alabi, Udia Markus, and Dietrich Klakow. 2020. [Transfer learning and distant supervision for multilingual transformer models: A study on African languages](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2580–2591, Online. Association for Computational Linguistics.
- Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piñeras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. 2022. [Challenges and strategies in cross-cultural NLP](#). In *Proceedings of the 60th Annual*



- Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6997–7013, Dublin, Ireland. Association for Computational Linguistics.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization](#).
- Comfort Eseohen Ilevbare, Jesujoba O. Alabi, David Ifeoluwa Adelani, Firdous Damilola Bakare, Oluwatoyin Bunmi Abiola, and Oluwaseyi Adesina Adeyemo. 2024. [Ekohate: Abusive language and hate speech detection for code-switched political discussions on nigerian twitter](#).
- 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.
- Ogunayo Jude Ogundepo, Akintunde Oladipo, Mofetoluwa Adeyemi, Kelechi Ogueji, and Jimmy Lin. 2022. [AfriTeVA: Extending ?small data? pretraining approaches to sequence-to-sequence models](#). In *Proceedings of the Third Workshop on Deep Learning for Low-Resource Natural Language Processing*, pages 126–135, Hybrid. Association for Computational Linguistics.
- Kentaro Kurihara, Daisuke Kawahara, and Tomohide Shibata. 2022. [Jglue: Japanese general language understanding evaluation](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2957–2966.
- Hang Le, Loïc Vial, Jibril Frej, Vincent Segonne, Maximin Coavoux, Benjamin Lecouteux, Alexandre Alauzen, Benoit Crabbé, Laurent Besacier, and Didier Schwab. 2020. [FlauBERT: Unsupervised Language Model Pre-training for French](#). In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 2479–2490.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- Tshepang Molale and Phillip Mpofu. 2023. (Dis)continuities of african language radio on social media: The case of south africa’s motswedding fm and radio zimbabwe. In *African Language Media*, 1 edition, page 14. Routledge.
- Shamsuddeen Muhammad, Idris Abdulmumin, Abinew Ayele, Nedjma Ousidhoum, David Adelani, Seid Yimam, Ibrahim Ahmad, Meriem Beloucif, Saif Mohammad, Sebastian Ruder, Oumaima Hourrane, Ali-pio Jorge, Pavel Brazdil, Felermimo Ali, Davis David, Salomey Osei, Bello Shehu-Bello, Falalu Lawan, Tajuddeen Gwadabe, Samuel Rutunda, Tadesse Belay, Wendimu Messelle, Hailu Balcha, Sisay Chala, Hagos Gebremichael, Bernard Opoku, and Stephen Arthur. 2023a. [AfriSenti: A Twitter sentiment analysis benchmark for African languages](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13968–13981, Singapore. Association for Computational Linguistics.
- Shamsuddeen Hassan Muhammad, Idris Abdulmumin, Seid Muhie Yimam, David Ifeoluwa Adelani, Ibrahim Sa’id Ahmad, Nedjma Ousidhoum, Abinew Ali Ayele, Saif M. Mohammad, Meriem Beloucif, and Sebastian Ruder. 2023b. [SemEval-2023 Task 12: Sentiment Analysis for African Languages \(AfriSenti-SemEval\)](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*. Association for Computational Linguistics.
- Shamsuddeen Hassan Muhammad, David Ifeoluwa Adelani, Sebastian Ruder, Ibrahim Sa’id Ahmad, Idris Abdulmumin, Bello Shehu Bello, Monojit Choudhury, Chris Chinenye Emezue, Saheed Salahudeen Abdullahi, Anuoluwapo Aremu, Alípio orge, and Pavel Brazdil. 2022. [NaijaSenti: A Nigerian Twitter sentiment corpus for multilingual sentiment analysis](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 590–602, Marseille, France. European Language Resources Association.
- Mary Myers and Nicola Harford. 2020. [Local radio stations in africa: Sustainability or pragmatic viability?](#)
- El Moatez Billah Nagoudi, AbdelRahim Elmadany, Ahmed El-Shangiti, and Muhammad Abdul-Mageed. 2023. [Dolphin: A challenging and diverse benchmark for Arabic NLG](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1404–1422, Singapore. Association for Computational Linguistics.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence RNNs and beyond](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- National Communications Authority. 2016. [List of authorised vhf-fm radio stations in ghana](#). Third Quarter.
- Rubungo Andre Niyongabo, Qu Hong, Julia Kreutzer, and Li Huang. 2020. [KINNEWS and KIRNEWS: Benchmarking cross-lingual text classification for](#)

- Kinyarwanda and Kirundi. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5507–5521, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Kelechi Ogueji and Orevaoghene Ahia. 2019. [Pidgin-unmt: Unsupervised neural machine translation from west african pidgin to english](#).
- 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.
- Akintunde Oladipo, Mofetoluwa Adeyemi, Orevaoghene Ahia, Abraham Owodunni, Odunayo Ogundepo, David Adelani, and Jimmy Lin. 2023. [Better quality pre-training data and t5 models for African languages](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 158–168, Singapore. Association for Computational Linguistics.
- Ikechukwu E Onyenwe, Mark Hepple, Uchechukwu Chinedu, and Ignatius Ezeani. 2018a. [A basic language resource kit implementation for the IgboNLP project](#). *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 17(2).
- Ikechukwu E Onyenwe, Mark Hepple, Uchechukwu Chinedu, and Ignatius Ezeani. 2018b. [A basic language resource kit implementation for the igbonlp project](#). *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 17(2).
- Ikechukwu E. Onyenwe, Mark Hepple, Uchechukwu Chinedu, and Ignatius Ezeani. 2019. [Toward an effective Igbo part-of-speech tagger](#). *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 18(4).
- Sylvester Olubolu Orimaye, Saadat M. Alhashmi, and Siew Eu-gene. 2012. [Sentiment analysis amidst ambiguities in youtube comments on yoruba language \(nollywood\) movies](#). In *Proceedings of the 21st International Conference on World Wide Web, WWW '12 Companion*, page 583–584, New York, NY, USA. Association for Computing Machinery.
- Adama Ouane and Christine Glanz. 2010. [How and why Africa should invest in African languages and multilingual education: An evidence-and practice-based policy advocacy brief](#). UNESCO Institute for Lifelong Learning.
- Folajimi Oyebola and Kingsley Ugwuanyi. 2023. [Attitudes of nigerians towards bbc pidgin: A preliminary study](#). *Language Matters*, 54(1):78–101.
- Wuraola Fisayo Oyewusi, Olubayo Adekanmbi, and Olalekan Akinsande. 2020. [Semantic enrichment of nigerian pidgin english for contextual sentiment classification](#).
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. [Cross-lingual name tagging and linking for 282 languages](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958, Vancouver, Canada. Association for Computational Linguistics.
- Malin Petzell. 2012. The linguistic situation in Tanzania. *Moderna språk*, 106(1):136–144.
- Afshin Rahimi, Yuan Li, and Trevor Cohn. 2019. [Massively multilingual transfer for NER](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 151–164, Florence, Italy. Association for Computational Linguistics.
- Machel Reid, Junjie Hu, Graham Neubig, and Yutaka Matsuo. 2021a. [AfroMT: Pretraining strategies and reproducible benchmarks for translation of 8 African languages](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1306–1320, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Machel Reid, Junjie Hu, Graham Neubig, and Yutaka Matsuo. 2021b. [AfroMT: Pretraining strategies and reproducible benchmarks for translation of 8 african languages](#). In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Punta Cana, Dominican Republic.
- E. ResCue and G. E. Agbozo. 2021. [Creating translated interfaces: The representations of african languages and cultures in digital media](#). In *Rethinking Language Use in Digital Africa: Technology and Communication in Sub-Saharan Africa*, pages 51–72. Multilingual Matters & Channel View Publications.
- Yves Scherrer. 2020. [TaPaCo: A corpus of sentential paraphrases for 73 languages](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 6868–6873, Marseille, France. European Language Resources Association.
- Iyanuoluwa Shode, David Ifeoluwa Adelani, and Anna Feldman. 2022a. [yosm: A new yoruba sentiment corpus for movie reviews](#).
- Iyanuoluwa Shode, David Ifeoluwa Adelani, and Anna Feldman. 2022b. [YOSM: A new Yorùbá Sentiment Corpus for Movie Reviews](#). *AfricaNLP 2022 @ICLR*.
- Oloruntola Sunday, Ayo Yusuff, Simon Godwin Iretomiwa, Vincent Adakole Obia, and Samuel Ejimwunmi. 2018. Use of indigenous languages for social media communication: The nigerian experience. In *African Language Digital Media and Communication*, 1 edition, page 15. Routledge.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter

Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchison, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshv, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem Mohamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Oscar Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshtir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Cogan, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. 2024. [Gemma 2: Improving open language models at a practical size](#).

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

Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#).

Michael Walsh. 2005. Will indigenous languages survive? *Annu. Rev. Anthropol.*, 34(1):293–315.

Bryan Wilie, Karissa Vincentio, Genta Indra Winata, Samuel Cahyawijaya, Xiaohong Li, Zhi Yuan Lim, Sidik Soleman, Rahmad Mahendra, Pascale Fung, Syafri Bahar, et al. 2020. [Indonlu: Benchmark and resources for evaluating indonesian natural language understanding](#). *arXiv preprint arXiv:2009.05387*.

Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong Yu, et al. 2020. [Clue: A chinese language understanding evaluation benchmark](#). *arXiv preprint arXiv:2004.05986*.

Yuan Yao, Qingxiu Dong, Jian Guan, Boxi Cao, Zhengyan Zhang, Chaojun Xiao, Xiaozhi Wang, Fanchao Qi, Junwei Bao, Jinran Nie, et al. 2021. [CUGE: A Chinese Language Understanding and Generation Evaluation Benchmark](#). *arXiv e-prints*, pages arXiv–2112.

Seid Muhie Yimam, Hizkiel Mitiku Alemayehu, Abinew Ayele, and Chris Biemann. 2020. Exploring Amharic sentiment analysis from social media texts: Building annotation tools and classification models. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1048–1060, Barcelona, Spain (Online).



# Appendices

## A Sahara Task Clusters

### A.1 Text Classification

#### Cross-Lingual Natural Language Inference.

This task involves the classification of 2 given sentences—a premise and a hypothesis into entailment, neutral, or contradiction semantic relation (Adelani et al., 2024b). Adelani et al. (2024b) introduced AfriXNLI, a benchmark for evaluating 15 African Languages on Natural Language Inference (NLI). The languages were translated from the English portion of XNLI (Conneau et al., 2018).

**Language Identification.** We include language identification datasets across 518 African languages (Adebara et al., 2022). The dataset includes data from multiple domains including education, government, health care, news, and religion. Additional statistics of the data is available in Table D3. We include language identification tasks to investigate (1.) the capabilities of language models in correctly generating the appropriate language names of an input text and (2.) Correctly generate an additional sentence in a specific language. We include these two approaches because many languages have multiple names (Chen et al., 2024) and a LM’s inability to appropriately name a language does not implicate a lack of information about the language.

**News Classification.** This task involves automatically categorizing news articles into predefined categories or topics based on their content. The goal is to leverage machine learning and NLP techniques to understand and organize large volumes of news data, making it easier for users to access relevant information. News classification plays a crucial role in information retrieval, content recommendation, and other applications that involve organizing and categorizing textual data. The news classification cluster consists of data from four languages including Amharic (Azime and Mohammed, 2021), Kinyarwanda (Niyongabo et al., 2020), Kirundi (Niyongabo et al., 2020), and Swahili (David, 2020a,b). Table D3 provides additional details of the data in this cluster.

**Sentiment Analysis.** Sentiment analysis is crucial in gaining insights into public opinion, customer feedback, and user sentiments across various platforms. The sentiment analysis cluster consists

of the Bambara Sentiment dataset (Diallo et al., 2021), YOSM—a Sentiment Corpus for Movie Reviews in Yorùbá (Shode et al., 2022b), and the Nigerian Pidgin sentiment dataset (Oyewusi et al., 2020), respectively. Further details about the statistical composition of this data is available in Table D3.

**Topic Classification.** The primary goal for this task is to automatically categorize the content of the text based on its theme. Topic classification is widely used in various applications, such as document organization, content recommendation, and information retrieval. We include topic classification datasets for Yorùbá and Hausa (Hedderich et al., 2020).

### A.2 Text Generation

**Machine Translation.** A robust benchmark for MT tasks is essential. SAHARA, in its quest to effectively assess machine translation for African languages, selectively incorporates datasets containing diverse African languages into its benchmark. Specifically, SAHARA includes datasets such as Pidgin-UNMT<sup>7</sup> (Ogueji and Ahia, 2019), AfroMT<sup>8</sup> (Reid et al., 2021a), Lafand-MT (Adelani et al., 2022a), SALT<sup>9</sup> (Akeru et al., 2022), and HornMT, thus encompassing a total of 45 language pairs.

**Paraphrase.** Paraphrasing is a core task in NLG that revolves around the comprehension and generation of text that conveys the same meaning while possibly exhibiting differences in structure, phrasing, or style (Chen et al., 2023). SAHARA incorporates the TaPaCo dataset (Scherrer, 2020), which is publicly accessible and covers a total of 73 languages, including four of African origin. These African languages consist of Kirundi, Amazigh, a macro-language Berber, and Afrikaans.

**Summarization.** Summarization is the task of creating a condensed version of a text that retains its core ideas and essential information (Nallapati et al., 2016). In the context of SAHARA, our focus primarily centers on a subset of the abstractive summarization dataset, XL-SUM (Hasan et al., 2021), which includes several African languages: Amharic, Hausa, Igbo, Kirundi, Oromo, Pidgin,

<sup>7</sup>PidginUNMT GitHub Link

<sup>8</sup>AfroMT GitHub Link

<sup>9</sup>SALT GitHub Link



Somali, Swahili, Tigrinya, and Yoruba. Additionally, we have extended our dataset collection efforts by scraping non-overlapping summarization datasets from BBC and Voice of Africa (VoE) websites, specifically targeting three African languages: Hausa, Ndebele, and Swahili.

**Title Generation.** The task of title generation in NLG involves generating a concise and meaningful title or headline for a given article. This task shares its dataset source with summarization, drawing from XL-SUM, and also includes dataset extensions from BBC and VoE with focus on enabling zero-shot title generation.

### A.3 Multiple-Choice, Comprehensive and Reasoning

This task cluster majorly addresses *Question Answering* type questions. The QA tasks in this cluster are more mathematical-reasoning and reading-comprehension-oriented.

**Mathematical Word Problems.** (Adelani et al., 2024b) introduced AfriMGSM, a QA dataset that contains human-written grade school mathematical word problems. Their dataset covers 15 African Languages, and the translation of each question to English for each language for in-language evaluation of LLMs. In this task, given a word problem, the LLM is expected to return a numerical value as output.

**General Knowledge QA.** (Adelani et al., 2024b) introduced AfriMMLU, a multi-choice QA dataset covering 5 areas of knowledge in 15 African Languages. The knowledge areas are high-school geography, high-school microeconomics, elementary mathematics, international law, and global facts. In this task given a question in any of the aforementioned areas, and 4 options, the LLM is expected to choose the right option.

**Reading Comprehension QA.** In this task, the LLMs are evaluated based on their ability to understand information in a given article. (Bandarkar et al., 2024b) introduced a parallel reading comprehension dataset covering 122 languages, 23 of which are native to Africa. The dataset comprises of short paragraphs accompanied by multi-choice Questions, whose answers can be inferred from accompanying paragraphs.

**Context-based Question-Answering.** QA represents another vital facet of NLG, focusing on a model’s capability to provide answers based on integrated knowledge. SAHARA relies on the gold passage dataset from the multilingual TYDIA<sup>10</sup> QA resource (Clark et al., 2020). This dataset offers concise questions with precisely one corresponding answer, enabling rigorous evaluation of QA performance.

### A.4 Tokens-level Classification

**Named Entity Recognition.** We incorporate NER datasets across 30 languages. We use Distance Supervision NER (DS NER) Data (Hedderich et al., 2020), MasakhaNER (Adelani et al., 2022b), WikiAnn (Rahimi et al., 2019), Yoruba-Twi NER (Alabi et al., 2020b), and multiple NER data from SADIaR. Additional details about the datasets are available in Table D3.

**Phrase Chunking.** This task involves identifying and grouping together consecutive words or tokens into meaningful syntactic units, known as phrases. These phrases can include noun phrases, verb phrases, prepositional phrases, and other grammatical structures. The goal of phrase chunking is to analyze the grammatical structure of a sentence and extract higher-level syntactic units for better understanding. The phrase chunking cluster consists of data for ten Indigenous languages of South Africa (see Table D3). The data has annotations for noun, verb, adjective, adverbial, and prepositional phrase chunks. Words not belonging to these phrase types are labelled with the tag *O*.

**Part of Speech Tagging.** This classification task involves assigning grammatical categories or labels (such as noun, verb, adjective) to each word in a sentence. The primary goal of POS tagging is to analyze the syntactic structure of a sentence and understand the role of each word within it. POS tagging is crucial for various NLP applications, as it provides a foundation for more advanced linguistic analysis. We include datasets for Igbo taken from IgboNLP (Onyenwe et al., 2019).

## B Benchmark

Table B1 shows the comparison between SAHARA with other benchmarks.

<sup>10</sup>TyDiQA GitHub Link


Category	Benchmark	Reference	Lang/Total	Tasks
Multilingual	FLoRES200	(Costa-jussà et al., 2022)	52/200	1
	GEM <sub>v1</sub>	(Gehrmann et al., 2021)	10/52	5
	GEM <sub>v2</sub>	(Gehrmann et al., 2021)	10/52	9
	NLLB M.D.	(Costa-jussà et al., 2022)	2/8	1
	NLLB S.D.	(Costa-jussà et al., 2022)	2/8	1
	SIB-200	(Adelani et al., 2024a)	46/200	1
	Toxicity200	(Costa-jussà et al., 2022)	50/200	1
	XGLUE	(Liang et al., 2020)	1/19	3
Language-specific	XTremE	(Hu et al., 2020)	2/40	4
	BanglaNLP <sub>v1</sub>	(Bhattacharjee et al., 2023)	1	6
	CLUE	(Xu et al., 2020)	1	6
	CUGE	(Yao et al., 2021)	1	8
	Dolphin	(Nagoudi et al., 2023)	1	13
	FLUE	(Le et al., 2020)	1	5
	IndoNLU	(Wilie et al., 2020)	1	5
	IndoNLG	(Guntara et al., 2020)	1	6
	JGLUE	(Kurihara et al., 2022)	1	3
	KorNLU	(Ham et al., 2020)	1	2
	LOT	(Guan et al., 2022)	1	4
	ORCA	(Elmadany et al., 2023)	1	7
African	PhoMT	(Doan et al., 2021)	1	1
	AfriSenti	(Muhammad et al., 2023a)	14	1
	AfroMT	(Reid et al., 2021b)	8/8	1
	AfroNLG	(Adebara et al., 2024b)	517/517	6
	AfroNLU	(Adebara et al., 2023)	28/28	7
	Horn-MT	—	3/3	1
	Iroko-Bench	(Adelani et al., 2024b)	16/16	3
	Mafand-MT	(Adelani et al., 2022a)	17/17	1
	Menyo-20k	(Adelani et al., 2021b)	1/1	1
	Naija-senti	(Muhammad et al., 2022)	4/4	1
	 SAHARA	ours	517/517	16

Table B1: A Comparison of SAHARA with other Benchmarks. **M.D:** Multilingual domain, **S.D:** Seed Data.

## C Few-shot prompt example

Figure C1 shows an example for MT task.

## D Dataset Description

Table D2 describes all datasets included in SAHARA.

## E Data-Performance-Policy Analysis

Compelete the translation of the last sentence from Tigrinya text to Somali

Text: እቲ ብዮርክ ኣሰናዳኢ ኩነት ዝተዳለወ ምርኢት ' 'ኣዲስ ኮንሰርት' ' ብዝብል ሰያሜ እዩ ክቐርብ።

Translation: Waxaa soo abaabulay qabanqaabiyaha xaflada Yorka bandhig faneedkan cinwaankeedu waa "Concert Addis".

Text: ነዋሰ መቐይር ውዕኢት ከመዝግቡ ናይዝኸለሉ ምህታት ኣካልን ኣብ ቀረባ እዋን ዝተረኸበን እዚ ምህቡ ኣብ ሓደ ዙርያ ጉያ ካብ ፍርቂ ሰከንድ ከባብ ሓደ ሰከንድ ምምሕያሽ ከምዕኦ ዝኸለል ተባሂሉ ይለመን።

Translation: Waxa ay ahayd tii ugu dambaysay ee bandhigyo cajiib ah oo lagu sameeyay tignoolajiyada ciyaarta beddeleysa, taas oo lagu tiriyo in ay horumariso waxqabadka inta u dhaxaysa nus ilbiriqsi iyo ilbiriqsi.

Text: ዓለማዊ ውድብ ጥዕና ዝርገኤ ፖለዮ ኣህጉራዊ ህጹጽ ኩነታት ጥዕና ህገዚ ምኒኡ ገለጹ።

Translation: Ururka Caafimaadka Adduunka ayaa sheegaya in faafitaanka cudurka dabayshu ay tahay xaalad caafimaad oo degdeg ah oo caalami ah.

Text: በዓል ሂድ ኣል ኣድሃ ኣብ ከተማታት ጁማ፣ ሃዋሳን ደሴን ኣብ ደገን ኣብ ውሽጢ ዝሃን ብዝተኻየዱ ሃይማኖታዊ ስርዓታት ተኸቢሩ።

Translation: Magaalooyinka Jimma, Hawassa, iyo Dessie ayaa sido kale si wanagsan loogu ciiday munaasabada Ciidul Adxaa oo isugu jirta dhacdooyin lagu qabtay gudaha iyo dibadba.

Text: መገዲ ኣየር ማሊዬዥያ ኣብዚ ዓመት ብሓደጋ ኸልተ ነፈርቱ ድሕሪ ምስኣኑ መሓውራዊ ምምዕርራይ ንምግባር 6000 ሰራሕተኛታት ከንከይ ከምዝተለመ ኣፍሊጡ።

Translation: Shirkada dayaaradaha ee Malaysiya ayaa ku dhawaaqday inay shaqada kajoojinayaan shaqaale dhan 6000 taas oo ay sababtay dib u habayn kadib markii ay sannadkan laba diyaaradood ku waayeen musiibooyin.

Text: ቀዳማይ ሚኒስተር ፓኪስታን፣ ራጃ በርዊዝ ኣሸረፍ፣ ኣንጻር ፕረዚደንት ኣሲፍ ዓሊ በርዳሪ ናይ ብልሽውና መርመራታ ንምኽፋት ብሰንኪ ምፍሻሉ ኣብ ትድሚ ላዕሊዋይ ቤት ፍርድ ፓኪስታን ቀረቡ።

Translation:

Figure C1: Few-shot prompting example for the machine translation task. The prompt provides five examples of the specific task for context and alignment.


Dataset	Brief Description	African Language(s)
MasakhaNER (Adelani et al., 2021a), (Adelani et al., 2022c)	Curated for Named Entity Recognition for several African Languages	amh, bam, bbj, ewe, fon, hau, ibo, kin, lug, Luo, mos, nya, pcm, sna, swa, tsn, twi, wol, xho, yor, zul
MasakhaPOS (Dione et al., 2023a)	Curated from News sources and publicly available datasets for Part-Of-Speech Tagging	bam, bbj, ewe, fon, hau, ibo, kin, lug, Luo, mos, nya, pcm, sna, swa, tsn, twi, wol, xho, yor, zul
MAFAND-MT (Adelani et al., 2022a)	African news corpus for Machine Translation	bam, bbj, ewe, fon, hau, ibo, lug, Luo, mos, swa, tsn, twi, wol, xho, yor, zul
XL-Sum Hasan et al. (2021)	Curated for Summarization and Title Generation	amh, ibo, orm, run, swa, yor, hau, pcm, som, tir
YorubaSenti (Orimaye et al., 2012)	Curated from comments on movies on YouTube	yor
TaPaCo (Scherrer, 2020)	Developed to provide example sentences and translations for particular linguistic constructions and words	afr, ber, run, zgh
TyDi QA (Clark et al., 2020)	Curated for Information-Seeking QA tasks	swa
HornMT	Parallel MT dataset curated for languages in Horn of Africa	aaf, amh, orm, som, tir
yoruba-twi-ner (Alabi et al., 2020a)	Curated for evaluating embeddings for low-resource languages	yor, twi
KINNEWS and KIRNEWS (Niyongabo et al., 2020)	Curated for multi-class classification of news articles	kin, run
AmharicNews (Azime and Mohammed, 2021)	Curated for text classification from news articles	amh
SwahiliNews v0.2 (David, 2020a)	Curated for text classification from news articles	swa
Bambara (Diallo et al., 2021)	Curated for Sentiment Analysis through web crawling	bam
IgboNLP (Onyenwe et al., 2018b)	Contains text corpus, POS tagset and POS-tagged subcorpus	ibo
yosm (Shode et al., 2022a)	Curated from movie reviews for Sentiment Analysis	yor
pidgin-tweet (Oyewusi et al., 2020)	Curated for SA from tweets	pcm
WikiAnn (Pan et al., 2017), (Rahimi et al., 2019)		afr, amh, ibo, mlg, kin, som, swa, yor
HausaTopic, YorubaTopic (Hedderich et al., 2020)	Curated for Topic Classification	hau, yor
SaDiLAR (Eiselen, 2016)	Curated for NER of 10 South African Languages	afr, nbl, nso, sot, ssw, tsn, tso, ven, xho, zul
AfriSenti (Muhammad et al., 2023a)	Curated for African tweet sentiment analysis	amh, ALG-ara, hau, kin, ibo, MOR-ara, MOZ-por, pcm, form, swa, tir, twi, tso, yor
Afro-MT (Reid et al., 2021a)	Benchmark for MT of 8 African Languages	afr, bem, lin, run, sot, swa, xho, zul
PidginUNMT (Ogueji et al., 2021)		pcm
SALT (Akeru et al., 2022)	Curated for Parallel MT of 5 Ugandan Languages	ach, lgg, lug, nyn, teo
Cheetah (Adebara et al., 2024b)	Summarization test set	hau, nde, swa
Cheetah (Adebara et al., 2024b)	Title generation test set	amh, gag (zero-shot), hau, ibo, pcm, som, swa, tir, yor, kin (zero-shot), afr, mlg (zero-shot), orm, nde (zero-shot), sna (zero-shot)
NCHLT-NER (Eiselen, 2016)	Named Entity Recognition	afr, nbl, nso, sot, ssw, tsn, tso, ven, xho, zul
BeleBele (Bandarkar et al., 2024b)	Reading Comprehension MCQA	afr amh, bam, fuv, gaz, hau, ibo, kea, kin, lin, lug, Luo, nso, nya, sna, som, sot, ssw, swh, tir, tsn, tso, wol, xho, yor, zul
IrokoBench (Adelani et al., 2024b)	MGSM, MMLU, XNLI	amh, ewe, hau, ibo, kin, lin, lug, orm, sna, sot, swa, tir, wol, xho, yor, zul

Table D2: A description of all datasets included in *Sahara*. **ALG**: Algerian, **MOR**: Moroccan, **MOZ**: Mozambican.




Task	Source	Language	Train/ Dev/ Test
MT	(Adelani et al., 2022a)	eng-hau	5.9K/ 1.3K/ 1.5K
		eng-ibo	6.9K/ 1.5K/ 1.4K
		eng-lug	4.1K/ 1.5K/ 1.5K
		eng-pcm	4.8K/ 1.5K/ 1.6K
		eng-swa	30.8K/ 1.8K/ 1.8K
		eng-tsn	2.1K/ 1.3K/ 1.5K
		eng-twi	3.3K/ 1.3K/ 1.5K
		eng-yor	6.6K/ 1.5K/ 1.6K
		eng-zul	3.5K/ 1.5K/ 1.0K
		fra-bam	3.0K/ 1.5K/ 1.5K
		fra-bbj	2.2K/ 1.1K/ 1.4K
		fra-ewe	2.0K/ 1.4K/ 1.6K
		fra-fon	2.6K/ 1.2K/ 1.6K
		fra-mos	2.5K/ 1.5K/ 1.6K
		fra-wol	3.4K/ 1.5K/ 1.5K
	(Reid et al., 2021a)	eng-afr	25.8K/ 3.2K/ 3.2K
		eng-bem	12.0K/ 1.5K/ 1.5K
		eng-lin	17.7K/ 2.2K/ 2.2K
		eng-run	12.5K/ 1.6K/ 1.6K
		eng-sot	28.8K/ 3.6K/ 3.6K
		eng-swa	28.1K/ 3.5K/ 3.5K
		eng-xho	26.1K/ 3.3K/ 3.3K
		eng-zul	29.1K/ 3.6K/ 3.6K
	(Ogueji and Ahia, 2019)	eng-pcm	1.7K/ 0.2K/ 0.2K
	(Akeru et al., 2022)	All-pairs <sup>1</sup>	20.0K/ 2.5K/ 2.5K
	HornMT	All-pairs <sup>2</sup>	1.7K/ 0.2K/ 0.2K
NER	(Adelani et al., 2021a)	See Table D2	41.2K/ 5.1K/ 5.1K
	(Adelani et al., 2022c)	See Table D2	41.2K/ 5.1K/ 5.1K
	(Eiselen, 2016)	See Table D2	1.7M/ 219.7K/ 215.6K
	(Alabi et al., 2020a)	See Table D2	20.2K/ 2.8K/ 5.5K
	(Pan et al., 2017)	See Table D2	9.2K/ 9.2K/ 9.4K
News	(Azime and Mohammed, 2021)	amh	41.2K/ 5.1K/ 5.1K
	(Niyongabo et al., 2020)	kin	15.3K/ 1.7K/ 4.3K
		run	3.3K/ 0.4K/ 0.9K
	(David, 2020a)	swa	4.4K/ 0.6K/ 0.6K
Paraphrase	(Scherrer, 2020)	amh, ber, kab, run ber kab	22.4K/ 2.8K/ 2.8K 17.6K/ 2.2K/ 2.2K 4.4K/ 0.6K/ 0.6K
Phrase Chunking	(Eiselen, 2016)	See Table D2	107.5K/ 13.0K/ 13.4K
POS Tagging	(Onyenwe et al., 2018a, 2019)	ibo	756.8K/ 94.7K/ 95.0K
Question Answering	(Clark et al., 2020)	swa	49.9K/ 0.5K/ n/a
Sentiment Analysis	(Diallo et al., 2021)	bam	2.4K/ 0.3K/ 0.3K
	(Oyewusi et al., 2020)	pcm	11.2K/ 1.4K/ 1.4K
	(Shode et al., 2022a)	yor	0.8K/ 0.2K/ 0.5K
Summarization	(Hasan et al., 2021)	amh	5.8K/ 0.7K/ 0.7K
		ibo	4.2K/ 0.5K/ 0.5K
		orm	6.1K/ 0.8K/ 0.8K
		Rundi	5.7K/ 0.7K/ 0.7K
		swa	7.9K/ 1.0K/ 1.0K
		yor	6.4K/ 0.8K/ 0.8K
		hau	6.4K/ 0.8K/ 0.8K
		pcm	9.2K/ 1.2K/ 1.2K
		som	6.0K/ 0.7K/ 0.7K
		Tigrinya	5.5K/ 0.7K/ 0.7K
	(Adebara et al., 2024b)	*†	- / - / 0.4K
Title Generation	(Hasan et al., 2021)	amh	5.8K/ 0.7K/ 0.7K
		ibo	4.2K/ 0.5K/ 0.5K
		orm	6.1K/ 0.8K/ 0.8K
		run	5.7K/ 0.7K/ 0.7K
		swa	7.9K/ 1.0K/ 1.0K
		yor	6.4K/ 0.8K/ 0.8K
		hau	6.4K/ 0.8K/ 0.8K
		pcm	9.2K/ 1.2K/ 1.2K
		som	6.0K/ 0.7K/ 0.7K
		Tigrinya	5.5K/ 0.7K/ 0.7K
	(Adebara et al., 2024b)	*	- / - / 5.9K
Topic Classification	(Hedderich et al., 2020)	hau	2.0K/ 0.3K/ 0.6K
		yor	1.3K/ 0.2K/ 0.4K


Table D3: Statistics of the data in our benchmark. *All – pairs*<sup>1</sup> each have the same size of data. They include ach-eng, ach-lgg, ach-lug, ach-nyn, ach-teo, ach-teo, eng-lgg, eng-lug, eng-nyn, eng-teo, lgg-teo, lug-lgg, lug-teo, nyn-lgg, nyn-lug, and nyn-teo. *All – pairs*<sup>2</sup> are all possible language combinations of aaf, amh, orm, som, tir, eng. \*† is a summarization test set including ‘hau’, ‘nde’ (zero-shot), and ‘swa’. \* is a title generation test set across 15 languages: ‘amh’, ‘gag’ (zero-shot), ‘hau’, ‘ibo’, ‘pcm’, ‘som’, ‘swa’, ‘tir’, ‘yor’, ‘kin’ (zero-shot), ‘afr’, ‘mlg’ (zero-shot), ‘orm’, ‘nde’ (zero-shot), ‘sna’ (zero-shot).



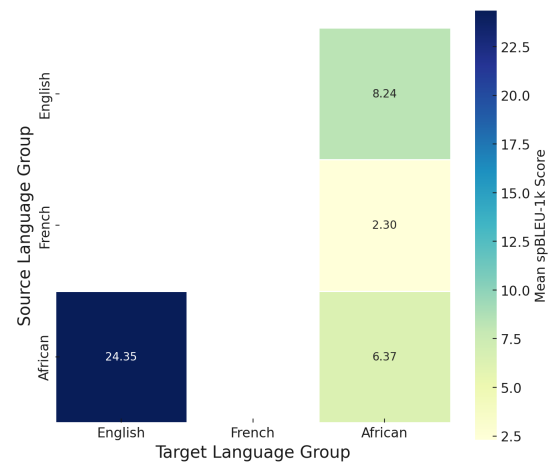
(a) Generation tasks.



(b) Classification tasks.



(c) MCCR tasks.



(d) MT tasks grouped.

Figure E2: Distribution of best open weight model performance (Command-A) on different task clusters across African languages in downstream data.






Figure E3: Distribution of languages across different clusters. Each bar represents the number of datasets a language appears in, categorized into the clusters, with some languages appearing in all four clusters while others are only present in one or two.



(a) Command-A performers.



(b) Claude-4-Sonnet performance.

Figure E4: Performance of Command-A and Claude-4-Sonnet on token-level task cluster across African languages in downstream data.

Cluster	Identifier	Metric	Shots	SLM						LLM																Closed LLM	
				S <sub>1</sub>	S <sub>2</sub>	S <sub>3</sub>	S <sub>4</sub>	S <sub>5</sub>	S <sub>6</sub>	L <sub>1</sub>	L <sub>2</sub>	L <sub>3</sub>	L <sub>4</sub>	L <sub>5</sub>	L <sub>6</sub>	L <sub>7</sub>	L <sub>8</sub>	L <sub>9</sub>	L <sub>10</sub>	L <sub>11</sub>	L <sub>12</sub>	L <sub>13</sub>	L <sub>14</sub>	L <sub>15</sub>	L <sub>16</sub>	CL <sub>1</sub>	CL <sub>2</sub>
Classification	xlnli	Acc.	5	33.6	33.7	36.7	41.18	40.4	46.68	48.08	55.06	34.3	40.69	37.09	61.09	57.86	65.58	44.2	38.19	51.89	48.38	52.09	35.87	46.77	59.67	69	69.6
	lid	F <sub>1</sub>	10	0	0	0.87	0.26	0.85	0.77	2.82	3.6	0	1.49	2.27	1.95	4.61	4	1.81	3.34	3.26	3.07	3.3	3.05	3.32	4.57	3.73	4.48
	news	F <sub>1</sub>	3	0.86	0.35	0.15	17.54	0	6.64	21.94	0	2.06	24.06	0.12	25.83	7.83	35.56	28.03	0.61	35.26	33.31	34.81	3.45	30.46	28.66	40.09	32.19
	sentiment	F <sub>1</sub>	5	38.51	24.75	30.45	11.05	28.77	1.99	36.09	42.47	38.4	28.36	36.07	34.56	20.55	46.35	35.86	25.94	24.51	45.82	35.89	19.98	37.07	57.67	57.52	
	topic	F <sub>1</sub>	3	8.73	31.62	20.3	19.79	12.46	17.15	48.11	39.35	34.95	28.55	42.12	44.3	31.8	70.73	36.47	43.74	64.89	56.48	48.73	52.09	50.96	47.16	65.93	76.57
	Avg.			16.34	18.08	17.69	17.96	16.50	14.65	31.41	28.10	21.94	24.63	23.53	33.55	24.53	44.44	29.27	22.36	35.96	37.41	34.96	22.89	33.72	38.64	47.28	48.07
Generation	mt_eng2xx	spBleu <sup>1k</sup>	5	1.79	3.63	3.45	3.25	2.12	3.4	7.07	7.61	1.72	4.37	2.79	6.24	6.91	4.56	3.73	3.29	8.84	8.7	9.36	1.4	6.06	8.22	12.18	12.36
	mt_fra2xx	spBleu <sup>1k</sup>	5	0.79	0.36	0.74	1.05	1.25	1.32	0.73	0.95	0.7	1.19	0.66	2.03	1.67	2.05	1.15	0.82	1.93	1.87	1.77	1.05	1.04	2.3	3.41	3.42
	mt_xx2xx	spBleu <sup>1k</sup>	5	0.4	0.45	0.42	0.42	0.36	0.18	0.43	0.63	0.72	0.55	0.46	0.71	0.9	1.39	0.55	0.17	1.28	1.05	1.37	0.67	0.38	0.85	5.15	4.44
	paraphrase	spBleu <sup>1k</sup>	5	23.2	28.88	31.15	19.12	25.93	17.85	28.73	25.83	31.68	31.33	23.66	17.26	21.34	15.37	32.13	30.68	26.86	19.64	27.79	19.23	36.06	26.06	23.2	20.35
	summary	rougeL	2	10.03	3.24	4.37	10.07	0.91	13.34	15.1	10.43	1.45	12.74	5.92	16.16	15.73	17.13	13.27	4.19	16.92	17.19	16.4	15.46	14.88	16.09	5.99	16.3
	title	spBleu <sup>1k</sup>	2	2.14	2.81	0.25	2.3	0.03	5.78	7.48	5.78	3	5.94	1.58	7.75	7.51	8.65	6.83	3.2	11.07	10.34	8.78	7.32	6.69	8.65	13.62	9.48
MCRR	Avg.			6.39	6.56	6.73	6.04	5.10	6.98	9.92	8.54	6.55	9.35	5.85	8.36	9.01	8.19	9.61	7.06	11.15	9.80	10.91	7.52	10.85	10.36	10.59	11.06
	mmu	Acc.	5	24.09	28.5	28.68	27.29	28.9	30.99	31.88	40.9	26.4	44.51	22.1	42.09	44.9	60.71	51.5	24	56.51	59.3	50.42	59.81	29.8	61.91	81.6	58.6
	msgm	ExactM	5	2.6	3.21	0	1.2	3.39	0.7	2.9	2.8	3.3	2.7	1.5	2.3	5.3	0.6	4.4	3	10.1	7.51	7.6	4.5	4.39	11.1	45.9	29.5
	belebele	Acc.	5	27.71	14.71	25.2	29.31	31.31	24.69	22.7	26.93	25.09	28.24	26.33	25.63	31.04	32.34	33.64	26.61	30.4	35.92	33.72	30.83	20.99	29.12	36.6	39.7
	squad_qa	F <sub>1</sub>	5	57.45	70.89	67.32	43.13	71.31	74.97	78.04	80.91	52.66	59.48	63.41	73.78	82.49	79.16	69.25	64.64	77.65	76.35	66.25	53.92	74.64	80.05	78.02	76.11
	Avg.			27.96	29.33	30.30	25.23	33.73	32.84	33.88	37.89	26.86	33.73	28.34	35.95	40.93	43.20	39.70	29.56	43.67	44.77	39.50	37.27	32.46	45.55	60.53	50.98
Tokens	ner	F <sub>1</sub>	5	0	0.78	1.76	2.99	0.13	3.65	2.26	7.49	2.35	3.96	0.2	6.03	13.14	11.4	8.51	1.52	6.86	10.54	8.64	6.04	4.3	16.88	35.06	27.04
	phrase	F <sub>1</sub>	5	14.79	28.5	29.44	11.52	24.91	27.73	15.95	29.35	29.42	16.55	30.87	13.59	30.01	23.97	26.34	23.09	25.55	23.13	16.78	25.13	29.44	31.05	36.69	36.45
	pos	F <sub>1</sub>	5	8.83	6.73	10.56	8.04	10.3	10.46	9.13	13.2	9.81	11.68	12.52	13.86	19	13.99	18.36	11.4	14.13	15.04	11.52	20.76	13.71	27.54	62.84	38.67
	Avg.			7.87	12.00	13.92	7.52	11.78	13.95	9.11	16.68	13.86	10.73	14.53	11.16	20.72	16.45	17.74	12.00	15.51	16.24	12.31	17.31	15.82	25.16	44.86	34.05
Overall				14.64	16.49	17.16	14.19	16.78	17.10	21.08	22.80	17.30	19.61	18.06	22.25	23.80	28.07	24.08	17.75	26.57	27.05	24.42	21.25	23.21	29.93	40.82	36.04

Table E4: Few-shot evaluation across different task clusters/tasks. The highest score for each individual task is in **bold green**, while the best average score across each task cluster (and **overall** score, calculated as average of task cluster averages) is in **bold orange**. Underline: refers to the best score across the open source/weights LLMs. **Shots**: number of shots. **SLMs**: S<sub>1</sub> Llama3.2 (1B), S<sub>2</sub> Llama3.2 (1B), S<sub>3</sub> Llama3.2 (3B), S<sub>4</sub> Gemma2 (2B), S<sub>5</sub> Phi-3.5 (3.8B), S<sub>6</sub> Gemma3 (4B). **LLMs**: L<sub>1</sub> Llama3.1 (8B), L<sub>2</sub> Gemma2 (9B), L<sub>3</sub> Aya (8B), L<sub>4</sub> Babel (9B), L<sub>5</sub> Command-R-Plus-R7B (8B), L<sub>6</sub> Gemma3 (12B), L<sub>7</sub> Gemma2 (27B), L<sub>8</sub> Gemma3 (27B), L<sub>9</sub> DeepSeek-R1-Distill-Qwen (32B), L<sub>10</sub> Aya (35B), L<sub>11</sub> Llama3.1 (70B), L<sub>12</sub> Llama3.3 (70B), L<sub>13</sub> DeepSeek-R1-Distill-Qwen (70B), L<sub>14</sub> Babel (83B), L<sub>15</sub> Command-R-Plus (104B), L<sub>16</sub> Command-A (111B). **Closed LLMs**: CL<sub>1</sub> Claude-4-Sonnet, CL<sub>2</sub> GPT-4.1.