

AfriSpeech-200: Pan-African Accented Speech Dataset for Clinical and General Domain ASR

Tobi Olatunji^{1*†}, Tejumade Afonja^{2,3*}, Aditya Yadavalli^{4*}, Chris Chinenye Emezie^{5,6*},
Sahib Singh^{7*}, Bonaventure F. P. Dossou^{5,6,8,9*}, Joanne Osuchukwu¹, Salomey Osei^{10*},
Atnafu Lambebo Tonja^{8,11,12*}, Naome Etori^{13*}, Clinton Mbataku^{3*}

*Masakhane NLP, ¹Intron Health, ²CISPA Helmholtz Center for Information Security,
³AI Saturdays Lagos, ⁴Karya Inc, ⁵Mila Quebec AI Institute, ⁶Lanfrica, ⁷Ford Motor Company,
⁸Lelapa AI, ⁹McGill University, ¹⁰University of Deutso, ¹¹Instituto Politécnico Nacional, ¹²University
of Colorado Colorado Springs, ¹³University of Minnesota

Abstract

Africa has a very poor doctor-to-patient ratio. At very busy clinics, doctors could see 30+ patients per day—a heavy patient burden compared with developed countries—but productivity tools such as clinical automatic speech recognition (ASR) are lacking for these overworked clinicians. However, clinical ASR is mature, even ubiquitous, in developed nations, and clinician-reported performance of commercial clinical ASR systems is generally satisfactory. Furthermore, the recent performance of general domain ASR is approaching human accuracy. However, several gaps exist. Several publications have highlighted racial bias with speech-to-text algorithms and performance on minority accents lags significantly. To our knowledge, there is no publicly available research or benchmark on accented African clinical ASR, and speech data is non-existent for the majority of African accents. We release AfriSpeech, 200hrs of Pan-African English speech, 67,577 clips from 2,463 unique speakers across 120 indigenous accents from 13 countries for clinical and general domain ASR, a benchmark test set, with publicly available pre-trained models with SOTA performance on the AfriSpeech benchmark.

1 Introduction

The African continent and the nearby islands constitute one-fourth of the land surface of the earth (Lodhi, 1993). Approximately 1.3 billion people live in Africa, which is about 18% of the world's population (Wikipedia contributors, 2023a). Of the estimated 7,000+ languages and dialects in the world, over 3,000 languages are found in Africa

(Wikipedia contributors, 2023b; Heine and Nurse, 2000).

Despite its large and predominantly young population, Africa bears a significant proportion of the global disease burden (de Graft Aikins et al., 2010) with multiple socioeconomic factors contributing to high mortality and morbidity rates (Baingana and Bos, 2006). Healthcare systems are overburdened and underfunded in many African countries (Oleribe et al., 2019; Naicker et al., 2009; Nkomazana et al., 2015), struggling to cope with the increasing demand for services, while at the same time facing significant shortages of trained health workers (WHO, 2022; Ahmat et al., 2022; Naicker et al., 2010; Nkomazana et al., 2015; Kinfu et al., 2009; Etori et al., 2023). A recent study conducted by Ahmat et al. (2022) in 47 African countries shows that the region has a ratio of 1.55 health workers (physicians, nurses, and midwives) per 1000 people—3x less than the WHO-recommended density of 4.45 health workers per 1000 people.

While technology can help mitigate some of these problems, Bukachi and Pakenham-Walsh (2007) and Manyati and Mutsau (2021) aptly show that although Africa has enjoyed massive growth in mobile technology, telecommunication, and internet penetration over the past two decades, healthcare technology lags significantly.

A 2019 systematic review on the use of Automatic Speech Recognition (ASR) for clinical documentation in the US from 1990 to 2018 by Blackley et al. (2019) and other similar studies (Goss et al., 2019; Blackley et al., 2020; Ahlgrim et al., 2016; Vogel et al., 2015) showed that the use of speech recognition led to a 19-92% decrease in mean documentation time, 50.3-100%

[†]Corresponding Author: tobi@intron.io.

decrease in turnaround time, and 17% improvement in documentation quality. However, in the African context, the lack of training datasets for many of the 3000+ languages and accents in the continent remains an obstacle in developing and adopting robust speech recognition systems for the general domain and for clinical ASR in particular (Doubouya et al., 2021; Siminyu et al., 2021; Babirye et al., 2022; Ogayo et al., 2022). While recent efforts have begun to turn this tide for the majority of African languages like Swahili, Kinyarwanda, and Yoruba (Gutkin et al., 2020; Dossou and Emezue, 2021; Olaleye et al., 2022), over a thousand African languages and accents remain excluded from global speech research advancements.

Recent single-digit word error rates (WER) (Chen et al., 2022; Radford et al., 2022; Hsu et al., 2021; Baevski et al., 2020b) in multiple SOTA publications and benchmarks on LibriSpeech (Panayotov et al., 2015), TED-LIUM3 (Hernandez et al., 2018), and other datasets using architectures like Wav2vec2 (Baevski et al., 2020b), Conformer (Gulati et al., 2020), Transducer, and Whisper (Radford et al., 2022) contrast significantly with ASR performance for African accented speech (Gutkin et al., 2020; Dossou and Emezue, 2021) (see Figure 2). We explore whether curating a large pan-African speech corpus might unlock comparable single-digit performance on African accents. We restrict this investigation to accented speech in English because English is the official language for the medical record in most Anglophone African countries, expanding the utility of this work to multiple Anglophone African countries.

Our contributions are as follows:

- We present *AfriSpeech-200*,¹ the first and most diverse open-source pan-African accented English speech corpus for clinical and general domain ASR, providing 200.70 hrs of accented speech, 67,577 speech-transcript pairs in 120 African accents across 13 countries, a benchmark dataset that paves the way for out-of-distribution, few-shot and zero-shot analyses on very-low-resource accents.²

¹<https://huggingface.co/datasets/tobiolatunji/afrispeech-200>.

²AfriSpeech-200 is licensed under a CC BY-NC-SA 4.0 license.

- We present a templating framework to augment existing corpora with native African proper nouns and evaluate multiple SOTA pre-trained models and leading commercial ASR systems on our benchmark dataset. We provide in-depth analysis of selected models to explain their failure modes and offer helpful insights.
- We fine-tune the best-performing open-source models and achieve SOTA performance on the AfriSpeech benchmark dataset (108 African accents) as well as show promising zero-shot performance on very low-resource accents. We provide best models³ as publicly available pre-trained checkpoints.

2 Related Work

With the advent of large multilingual speech datasets (Panayotov et al., 2015; Javed et al., 2022; Chen et al., 2021; Ardila et al., 2020; Valk and Alumäe, 2021), various research groups have proposed large self-supervised speech models such as wav2vec (Schneider et al., 2019), vq-wav2vec (Baevski et al., 2020a), wav2vec 2.0 (Baevski et al., 2020b), HuBERT (Hsu et al., 2021), XLSR (Conneau et al., 2021), and XLS-R (Babu et al., 2022). These models achieved state-of-the-art performance on many downstream tasks such as automatic speech recognition (ASR), automatic speech translation (AST), and language identification. However, most existing systems still perform poorly on accented speech (Javed et al., 2022). Koenecke et al. (2020) further showed that popular commercial ASR systems—like Amazon, Apple, Google, IBM, and Microsoft—exhibit substantial racial disparities in their speech recognition capabilities. Most ASR systems work best for native English speakers and their accuracy plummets dramatically with non-native English speakers (Hassan et al., 2022; Prasad and Jyothi, 2020).

To enhance the performance of accented speech recognition, various methods have been proposed, which can be categorized into modeling and dataset approaches. On the modeling front, there have been efforts such as dialect-aware ASR models (Yadavalli et al., 2022), domain adversarial training (DAT) (Sun et al., 2018), combining DAT

³https://huggingface.co/Seyfelislem/afrispeech_large_A100.

with transfer learning (Chen et al., 2020), using voice conversion (VC) (Zhang et al., 2022), combining VC with speed perturbation (Zhang et al., 2022), and accent pre-training (Acc-PT) (Das et al., 2021). These efforts, however, produced marginal improvements and still exhibit poor generalization capabilities.

Datasets have played a major role in improving ASR performance. The current SOTA in ASR (Radford et al., 2022) demonstrated the superior utility of large supervised datasets. Therefore, to bridge the ASR performance gap for African accented speech, multiple dataset creation efforts (Doubouya et al., 2021; Siminyu et al., 2021; Babirye et al., 2022; Ogayo et al., 2022; Gutkin et al., 2020; Dossou and Emezue, 2021; Afonja et al., 2021; Kamper and Niesler, 2011; Ibejih et al., 2022) have been established. However, many of these datasets are limited in size and diversity. For example, Common Voice (Ardila et al., 2020) contains less than 10 hours of African English speech, Li et al. (2021) evaluate on 50 hrs of African accented English (not released), Sanabria et al. (2023) provide 40 hrs of accented English, less than 20% is African. Kamper and Niesler (2011) and De Wet et al. (2007) are limited to a few South African accents, and Ibejih et al. (2022) include less than 8 hours, while Afonja et al. (2021) include less than 2 hours of accented African English speech. Furthermore, there are no available benchmarks for clinical ASR for African languages, creating a need for evaluation datasets that help identify areas of improvement in this domain.

While previous works have primarily focused on adapting Western accents to African accents, to the best of our knowledge, there has been limited research specifically addressing domain adaptation from a general domain to the clinical domain in the African context. In this regard, our work is the first attempt to bridge this gap and tackle the unique challenges associated with adapting accented African English ASR systems to the clinical domain.

3 AfriSpeech Dataset

We introduce AfriSpeech, a Pan-African accented English speech dataset for clinical and general domain ASR crowd-sourced from 2,463 African speakers, 200.70 hrs with an average audio duration of 10.7 seconds. Speaker, gender, age

Country	Clips	Speakers	Hours
Nigeria	45875	1979	142.40
Kenya	8304	137	20.89
South Africa	7870	223	22.69
Ghana	2018	37	5.16
Botswana	1391	38	3.96
Uganda	1092	26	2.89
Rwanda	469	9	1.47
United States ⁴	219	5	0.53
Turkey ⁵	66	1	0.18
Zimbabwe	63	3	0.18
Malawi	60	1	0.15
Tanzania	51	2	0.18
Lesotho	7	1	0.02

Table 1: Contributions by country showing speakers, number of clips, and speech duration in seconds and hours.

group, and clip domain distributions are shown in Table 2. In the following subsections, we describe the dataset creation process.

3.1 Focus Languages

We conducted an investigation on 120 African accents across 13 countries including the United States and Turkey. These accents originate from languages that belong to five language families, as documented by Eberhard (Eberhard et al., 2019): Afro-Asiatic, Indo-European, Khoe-Kwadi (Hainum), Niger-Congo, and Nilo-Saharan. This selection represents the diverse linguistic landscape across western, eastern, and southern Africa. In Table 1, we provide an overview of the number of clips, speakers, and hours of data per country, with Nigerian accents constituting 67% of the dataset. Since some languages are spoken across several countries (e.g., Swahili, isiZulu, Hausa, and Luganda), accents are not unique to countries.

3.2 Obtaining AfriSpeech Transcripts

Neural network models learn concepts from training data. Where the training data is predominantly Western (e.g., Common Voice [Ardila

⁴Although the self-reported country from the speakers is the United States, their reported accents, namely, Yoruba and Igbo, is mostly spoken in the western part of Africa.

et al., 2019]), the resulting ASR systems fail to capture important pan-African contexts. For example, ASR systems fail woefully at transcribing African names like “Ogochukwu” (Igbo), “Malaika” (Swahili), or “Uwimana” (Rwandan), while excellently transcribing Western names like “Lauren” and “Bryan”—representative of the bias in their training corpora. To solve the problem of scarce African-centric text in the general and clinical domains, we created AfriSpeech using the following strategies.

3.2.1 Finding Available Transcripts

Our first task was to supplement existing large multi-domain corpora with African-centric text. Our first target was **Wikitext-103** (Merity et al., 2016), a collection of over 100 million tokens extracted from the set of verified “good” and “featured” articles on Wikipedia curated by Salesforce. We split this corpus on sentence boundaries and randomly sampled sentences for our transcript corpus. Our next strategy was **web scraping**. We crawled and scraped major African news websites across multiple African countries on topics like politics, entertainment, sports, religion, education, etc. In contrast to Wiki-text, the resulting corpus contained several African names, cities, and highly relevant vocabulary applicable to real-world use cases for downstream ASR. By scraping health-focused websites and health sections of news websites, we were able to get content from the clinical domain, albeit very little.

To increase clinical content representation, we focused on two multi-specialty biomedical datasets: **PubMed** (Wheeler et al., 2007) and the **NCBI disease** (Doğan et al., 2014). We split these corpora on sentence boundaries and randomly sampled sentences for our transcript corpus.

3.2.2 Finding African Entities

We sourced for African-centric entities in two places: First, we leveraged an existing database of over 90,000 African names from the transatlantic slave trade between 1808 and 1863 (Anderson et al., 2013), which increased our coverage of African names, phonemes, and morphemes. We then used Okagbue et al.’s (2017) dataset of 965 Igbo names collected to reflect the dialectal classification of Igbo people and supplemented it with 1,000 more Nigerian names from other cultures

⁵Even though the reported country is Turkey, the reported Zulu accent is mostly spoken in the southern part of Africa.

such as Yoruba, Hausa, Fulani, Tiv, Efik, Ibibio, etc. These names were obtained from freely available textbooks, online baby name websites, oral interviews, published articles, and online forums like Instagram and Twitter. Finally, we obtained a list of African cities from Wikipedia (Wikipedia contributors, 2023c).

3.2.3 AfriSpeech Templates

The web scraping corpus was highly relevant but small. In the larger biomedical and Wikitext datasets, African content was sparse. We, therefore, sought to increase the utility of the curated corpora by creating “Africanized” versions. Several studies have demonstrated the utility of “templates” as an effective way to create richer, more expressive training datasets, especially for Question-Answering and prompt engineering (Pawar and Shrawankar, 2016; Brown et al., 2020; Yao et al., 2022) and named entity recognition (Davody et al., 2022). Inspired by this approach, we augment our dataset by sampling sentences from the corpora described above in addition to template sentences contributed by professional clinicians, hand-crafting a total of 140 template sentences. For each template sentence, we masked proper nouns (first names, last names, organizations, and cities), replacing them with their corresponding NER tags [PER, ORG, LOC]. We then randomly replaced the masked tokens with African-centric entities—African names and cities, derived from section 3.2.2 above, as well as common tropical diseases and medications. Each template sentence was reused 200 times. A random subset was sampled, sent as prompts for recording, and included with this release. Templated sentences represent approximately 30% of this corpus.

3.3 Audio Recording

Collection: Inspired by Common-Voice (Ardila et al., 2019) and SautiDB (Afonja et al., 2021), we developed and deployed a web-based application in Python/Flask (Figure 1) to collect crowd-sourced speech samples. The application also facilitates tracking of completion status, user demographics, reviews, and quality control. The app presents randomly selected sentences (prompts) to the speakers and prompts them to record their voices while reading the text. The speech recordings are persisted as mono-channel, 16-bit wav files, with a 48 kHz sampling rate.




Figure 1: Intron Online Recording platform.

Speaker Gender Ratios - # Clip %	
Female	57.11%
Male	42.41%
Other/Unknown	0.48%
Speaker Age Groups - # Clips	
<18yrs	1,264 (1.87%)
19–25	36,728 (54.35%)
26–40	18,366 (27.18%)
41–55	10,374 (15.35%)
>56yrs	563 (0.83%)
Unknown	282 (0.42%)
Clip Domain - # Clips	
Clinical	41,765 (61.80%)
General	25,812 (38.20%)

Table 2: Dataset statistics.

Post-processing tasks were performed on the audio recordings to remove samples shorter than 2 seconds and longer than 17 seconds. Raw unedited samples are provided as part of this release. Speakers in this dataset have been de-identified. Demographic information available includes gender, age group, accent, and country.

Annotation Instructions Recorder demographics are presented in Table 2. Instructions were provided to crowd-sourced recorders as detailed in Appendix A.2. Notably, the recorders were instructed to read punctuation marks in full and encouraged to use their natural accent.

3.4 Quality Control

Projects: Transcripts were bucketed into projects to separate clinical from general domain prompts. This approach maximized the time value of clinician contributors focusing their efforts more on medical prompts.

Reviewers: We hired a team of human reviewers who up-voted or down-voted clips to indicate quality. Text feedback was also provided to recorders in 30% of cases where negative feedback was indicated. The text feedback contained the reason for the down-vote and was intended to help recorders improve future recording quality.

Guest Clip Review: New recorders were admitted as guests and allowed to record a maximum of 200 clips before quality review. Ten to 30 clips were reviewed per guest and those who passed review were promoted to a ‘Paid’ status.

Paid Clip Review: In the paid category, users were allowed a maximum of 200 clips before a temporary pause for quality check. During the temporary suspension, reviewers randomly reviewed 10% of the speech samples provided and positive, negative, or text feedback was provided. Access was restored if quality remained satisfactory, or users were blacklisted if over 30% of clips reviewed were down-voted.

Delisting Problematic Sentences: Where an audio clip receives a down-vote, the corresponding sentence is released for re-recording by a different user. If a clip recorded for the same sentence receives a second down-vote, the transcript itself is blacklisted.

4 Experiments

4.1 Data

AfriSpeech-200 is a manually reviewed and curated subset, representing 7% of the total AfriSpeech dataset, intended as an initial public release to stimulate research into African clinical and general domain ASR for accents with little or no representation in speech research. Table 1 shows the distribution of clips, unique speakers, and hours by country.

As shown in Table 3, the train, test, and development sets are bucketed such that any given speaker may appear in only one. This ensures that contributors seen at train time are not seen at test time, which would skew the results.

Item	Train	Dev	Test
# Speakers	1466	247	750
# Hours	173.4	8.74	18.77
# Accents	71	45	108
Avg secs/speaker	425.80	127.32	90.08
clips/speaker	39.56	13.08	8.46
speakers/accent	20.65	5.49	6.94
secs/accent	8791.96	698.82	625.55
# general domain	21682	1407	2723
# clinical domain	36318	1824	3623

Table 3: Dataset splits showing speakers, number of clips, and speech duration in Train/Dev/Test splits.

4.2 Benchmarks

We compare SOTA open-source pre-trained ASR models, Whisper (Radford et al., 2022), Wav2vec2 (Baevski et al., 2020b), XLSR (Babu et al., 2022), Hubert (Hsu et al., 2021), WavLM (Chen et al., 2022), Conformer (Gulati et al., 2020), and CRDNN-RNNLM (Ravanelli et al., 2021), with commercial clinical and non-clinical ASR systems. We refer readers to read the respective papers for details on pretraining corpora, model architecture, and hyperparameters. For each model, we compare performance (WER) on Librispeech test-clean partition (Panayotov et al., 2015) with WER on the AfriSpeech dev and test sets. Single-run results are provided.

4.3 Fine-tuning

Based on the benchmark results in Table 4 and GPU memory constraints, 2 top performing open-source model architectures were selected for fine-tuning. Although commercial ASR systems outperformed many open-source models, they are excluded from fine-tuning experiments because their model architectures and underlying pre/post-processing logic are unknown.

Selected Model Architectures

1. wav2vec-large-xlsr-53 (Grosman, 2021): an Encoder-decoder architecture with CNN-based feature extractor, code book, and transformer-based encoder, 378.9M parameters; LR 1e-4.

2. whisper-medium (Radford et al., 2022): a Decoder-only multi-task architecture, 789.9m parameters; LR 2.5e4.

For each model, we fine-tuned with FP16, AdamW (Loshchilov and Hutter, 2017), batch size of 16, for 10 epochs, with a linear learning rate decay to zero after a warmup over the first 10% of iterations. We fine-tune and evaluate on 3 domains: (1) **general** (25,812 clips), (2) **clinical** (41,765 clips), and (3) **both** (67,577 clips). We train on each domain and test across all 3 domains to investigate the effect of out-of-domain data on model performance. XLSR models were trained on a single Tesla T4 GPU with 16GB GPU memory while Whisper and Conformer models were trained on RTX8000 GPU with 48GB GPU memory. Fine-tuning took 24-48 hrs for all domains.

4.4 Model Vocabulary

Most pre-trained models define a limited vocabulary of only Latin alphabets with no numbers or punctuation (Baevski et al., 2020b). In stark contrast, numbers are critical in healthcare, e.g., blood pressure 130/80 mmHg, or Lab results 0.428 mmol/L. Eliminating all numerical references in clinical text is dangerous and counterproductive. Post-processing to convert all numerical values to long form is imperfect so we retain numbers in their original form. For fine-tuning experiments, we define an alphanumeric vocabulary with semantically important punctuations, characters, and symbols commonly used in medical practice (colon, question mark, plus, etc.).

4.5 Evaluation

We report our results as WER on AfriSpeech dev and test sets in addition to domain and accent-specific performance. Results are compared with Librispeech (Panayotov et al., 2015) test set performance. We also report the zero-shot performance of fine-tuned models on unseen accents in the test set.

5 Results and Discussion

5.1 Africa-centric Fine-tuning Improves Robustness

As shown in Table 4, compared with its pre-trained version, xlsr-53 fine-tuned on general domain speech (AfriSpeech-general) yields 53.4% relative

Model	Params	Training/Fine-tuning Corpora	ls-clean	Dev (45 accents)			Test (108 accents)		
				General	Clinical	Both	General	Clinical	Both
Open-Source SOTA Models									
openai/whisper-large	1550M	Multi, 680k hrs	0.167	0.235	0.287	0.261	0.240	0.375	0.306
openai/whisper-medium	769M	Multi, 680k hrs	0.166	0.246	0.300	0.273	0.276	0.392	0.332
openai/whisper-medium-en	769M	Multi, 680k hrs	0.169	0.267	0.315	0.291	0.304	0.414	0.358
openai/whisper-small	244M	Multi, 680k hrs	0.167	0.313	0.372	0.343	0.330	0.455	0.391
openai/whisper-small-en	244M	Multi, 680k hrs	0.167	0.319	0.384	0.352	0.350	0.482	0.414
nvidia/stt-en-conformer-ctc-large	118M	Multi, 10	0.210	0.410	0.486	0.448	–	–	–
nvidia/stt-en-conformer-transducer-large	139M	Multi, 10	0.150	0.408	0.477	0.443	–	–	–
jonatasgrosman/wav2vec2-large-xlsr-53-english	317M	Multi, 3	0.100	0.498	0.561	0.530	0.506	0.650	0.576
jonatasgrosman/wav2vec2-xls-r-1b-english	317M	Multi, 4	0.087	0.502	0.571	0.537	0.521	0.670	0.594
facebook/wav2vec2-large-960h-lv60-self	317M	Single, 2	0.051	0.512	0.587	0.550	0.533	0.694	0.611
facebook/hubert-xlarge-ls960-ft	1B	Single, 1	0.052	0.531	0.610	0.571	0.562	0.725	0.641
patrickvonplaten/wavlm-libri-clean-100h-large	317M	Single, 1	0.091	0.606	0.679	0.643	0.631	0.783	0.705
facebook/wav2vec2-large-960h	317M	Single, 1	0.062	0.610	0.695	0.652	0.641	0.797	0.717
facebook/wav2vec2-large-robust-ft-swbd-300h	317M	Single, 5	0.093	0.689	0.778	0.734	0.733	0.906	0.817
Commercial ASR APIs									
Azure	–	–	–	0.438	0.468	0.453	0.340	0.444	0.391
AWS	–	–	–	0.332	0.437	0.385	0.354	0.536	0.442
GCP	–	–	0.132	0.494	0.565	0.530	0.534	0.624	0.578
Commercial Clinical ASR APIs									
AWS [Medical] (Primary Care)	–	–	–	0.385	0.416	0.400	0.439	0.520	0.478
GCP [Medical]	–	–	–	0.550	0.475	0.512	0.567	0.537	0.552
Ours									
facebook/wav2vec2-large-xlsr-53-english-general	317M	+ AfriSpeech-general	0.253	0.254	0.437	0.347	0.236	0.468	0.349
facebook/wav2vec2-large-xlsr-53-english-clinical	317M	+ AfriSpeech-clinical	0.415	0.437	0.312	0.374	0.424	0.308	0.368
facebook/wav2vec2-large-xlsr-53-english-all	317M	+ AfriSpeech	0.314	0.295	0.308	0.302	0.279	0.308	0.293
openai/whisper-medium-general	769M	+ AfriSpeech-general	0.351	0.205	0.486	0.347	0.186	0.525	0.351
openai/whisper-medium-clinical	769M	+ AfriSpeech-clinical	0.568	0.491	0.264	0.376	0.464	0.266	0.368
openai/whisper-medium-all	769M	+ AfriSpeech	0.418	0.213	0.241	0.227	0.192	0.242	0.216

Table 4: Results showing selected models, number of parameters, number of pre-training/fine-tuning corpora [“Multi” refers to multilingual or multi-task], Librispeech (Panayotov et al., 2015) test clean WER and AfriSpeech dev and test set performance for open-source, commercial ASR models, and fine-tuned models (Ours). Missing values indicate incomplete or failed experiments.

improvement. Xlsr-53 fine-tuned on clinical domain speech (AfriSpeech-clinical) yields 52.6%, and xlsr-53 fine-tuned on the combined domains (AfriSpeech-all) yields 49.1% relative improvement. The trend is similar with pre-trained Whisper-medium, yielding 32.6% relative improvement on the general domain, 32.1% on the clinical domain, and 34.9% when finetuned on combined domains.

5.2 Training Data Bias

In the Open-Source section of Table 4, AfriSpeech dev and test set performance correlates with the number and diversity of pre-training datasets. For example, Wav2vec2 models trained exclusively on Librispeech significantly underperform when compared with those trained on multiple (Baevski et al., 2020b) or multilingual corpora (Babu et al., 2022). Models trained on multilingual or multi-task corpora (Radford et al., 2022; Gulati et al., 2020) learn more useful representations, are more linguistically diverse, are more robust, and generalize better to accented speech.

5.3 Clinical ASR is Sensitive to Model Vocabulary

As mentioned in Section 4.4, most ASR models tend to transcribe numbers in their extended forms, which have a detrimental effect on their WER as shown in Table 4, particularly in the clinical domain where numerical values need to be transcribed accurately (columns 6 & 9). However, ASR models with a larger vocabulary, such as Whisper, Commercial ASR models, and our fine-tuned models, demonstrate superior performance by effectively transcribing numbers in clinical speech and converting them into correct numeric representations.

5.4 Punctuation Prediction is Critical for Clinically Useful ASR

Medical documents typically follow preset sequence and formatting, for example, patient history, general examination, laboratory investigation, etc., separated by new lines, section titles, or semi-colons. Punctuation commands such as “Next line”, “full stop” (.), “query” (?), “comma” (,), “colon” (:) are frequently used

Accent	Samples	OpnSrc Whisper	Commercial			Ours Whisper
			Azure	GCP	AWS	
Niger-Congo						
Ukwuani	119	0.364	0.393	0.677	0.484	0.244
Eggon	100	0.254	0.316	0.616	0.359	0.122
Bini	76	0.830	0.840	0.916	1.061	0.412
Yoruba, hausa	75	0.462	0.367	0.463	0.437	0.133
Ekpeye	70	0.376	0.406	0.582	0.539	0.190
Bajju	61	0.229	0.323	0.428	0.378	0.171
Ikulu	60	0.406	0.388	0.650	0.543	0.195
Jaba	59	0.462	0.475	0.798	0.529	0.268
Ekene	55	0.414	0.350	0.673	0.519	0.192
Agatu	54	0.734	0.725	0.903	0.793	0.387
Ijaw(nembe)	49	0.478	0.529	0.743	0.675	0.275
Delta	48	0.384	0.351	0.724	0.473	0.205
Igarra	45	0.591	0.539	0.839	0.687	0.258
Khana	45	0.539	0.584	0.761	0.785	0.318
Gbagyi	42	0.327	0.461	0.633	0.475	0.195
Jukun	42	0.182	0.234	0.415	0.244	0.122
Brass	39	0.147	0.269	0.357	0.309	0.131
Afro-Asiatic						
Mada	78	0.485	0.560	0.684	0.634	0.236
Mwaghavul	67	0.444	0.513	0.690	0.613	0.235
Angas	58	0.605	0.580	0.862	0.653	0.343

Table 5: Zero shot (OOD) accents. Test set WER on top 20 accents absent from the training set for open-source (OpnSrc), commercial, and fine-tuned ASR models (Ours).

in healthcare dictations to add structure to documents. ASR systems without support for such commands force clinicians to review every line of the ASR transcript to add/revise punctuations and document structure, prolonging documentation time and patient wait time (Sunkara et al., 2020). As a result, commercial clinical ASR systems supporting these commands are preferable and outperform general-purpose models.

5.5 Commercial ASR APIs are Not So Global

The 3 large commercial ASR systems evaluated in this study have global presence. Millions of African Android users have access to Voice typing through the Google keyboard and Microsoft Word users have access to its ASR engine. Table 6 compares the performance of these ASR APIs on majority African accents and we show that despite their global presence, performance lags significantly on some of Africa’s most populous accents like Swahili and Yoruba.

5.6 Domain Adaptation

Pre-trained whisper models performed better on general domain speech (AfriSpeech-general) when compared with the clinical domain, demonstrating the relative domain-driven difference in difficulty despite the robust training data for Whisper models (680k hours, 90 languages). Cross-domain fine-tuning yields significant gains

helping to somewhat bridge this gap. Our results agree with prior work on domain adaptation (Sun et al., 2017; Abdelwahab and Busso, 2015) showing that models trained exclusively on clinical data improve when general domain data is added. Whisper shows 9% relative improvement on the clinical domain with the addition of general domain data. However, this trend is reversed with general domain data. Adding speech from the clinical domain leads to a 3% and 18.2% relative drop for Whisper and xlsr-53, respectively. Domain adaptation is no silver bullet. Care must be taken to apply this approach where benefits outweigh risks.

5.7 Accent-level Performance

Table 6 shows test set performance on the top 23 AfriSpeech accents grouped by their language families. We report the results for open-source, commercial, and fine-tuned ASR models. Fine-tuned models (ours) average relative improvement is 26.7% over the open-source ASR models and 36.5% over the commercial ASR models. For several accents, we observe that the whisper model fine-tuned with our AfriSpeech dataset shows the best overall performance with an average relative improvement of 16.2% across all accents, except in 4 South African languages (Zulu, isiZulu,⁶ Tswana, Afrikaans), Luo, and Kinyarwanda, where the fine-tuned model underperforms compared to the pretrained whisper model and commercial Azure model performs best on Luo accent. Although counter-intuitive, it is possible these accents are highly represented in Whisper pre-training data and require further investigation.

5.8 Zero-Shot Performance

We further explore generalizability to unseen accents, i.e., out-of-distribution (OOD) accents. Table 5 shows the results for the top 20 OOD accents in the test set. We observe an impressive 44.4% relative performance improvement across all OOD accents with our fine-tuned Whisper

⁶We note that both Zulu and isiZulu are the same but they are labeled differently in our dataset. We further discuss this in the Limitations section.

Accent	Country	Test Samples	Train Samples	Open Source		Commercial			Ours, Finetuned	
				xlsr-53	whisper	Azure	GCP	AWS	XLSR	Whisper
Niger-Congo										
Yoruba	[NG]	575	14233	0.576	0.327	0.364	0.581	0.421	0.291	0.218
Swahili	[KE, TZ, UG, ZA]	485	5484	0.448	0.192	0.307	0.436	0.305	0.244	0.181
Igbo	[NG]	319	8068	0.564	0.338	0.393	0.563	0.441	0.273	0.197
Zulu	[TR, LS, ZA]	156	1309	0.471	0.223	0.329	0.477	0.345	0.315	0.237
Setswana	[BW, ZA]	96	1275	0.448	0.208	0.288	0.446	0.300	0.291	0.234
Isizulu	[ZA]	88	779	0.457	0.182	0.254	0.406	0.292	0.265	0.206
Ijaw	[NG]	77	2371	0.608	0.364	0.372	0.671	0.446	0.321	0.238
Luhya	[KE]	69	426	0.538	0.310	0.548	0.489	0.427	0.296	0.245
Twi	[GH]	54	1321	0.504	0.184	0.382	0.510	0.361	0.236	0.177
Idoma	[NG]	53	1767	0.607	0.384	0.424	0.639	0.543	0.294	0.243
Luganda	[KE, UG, BW]	44	529	0.525	0.320	0.362	0.526	0.378	0.381	0.277
Tswana	[BW, ZA]	34	289	0.362	0.184	0.265	0.425	0.267	0.249	0.241
Akan (fante)	[GH]	29	230	0.732	0.418	0.425	0.803	0.604	0.290	0.197
Kikuyu	[KE]	24	163	0.406	0.160	0.275	0.387	0.300	0.221	0.126
Xhosa	[ZA]	17	342	0.498	0.265	0.322	0.332	0.389	0.318	0.237
Sepedi	[ZA]	17	176	0.651	0.373	0.394	0.659	0.458	0.414	0.285
Kiswahili	[KE]	16	811	0.466	0.159	0.389	0.394	0.274	0.173	0.163
Urhobo	[NG]	15	578	0.551	0.378	0.423	0.678	0.423	0.345	0.210
Nembe	[NG]	14	546	0.571	0.352	0.449	0.556	0.449	0.372	0.296
Kinyarwanda	[RW]	14	439	0.495	0.216	0.338	0.527	0.437	0.369	0.311
Afro-Asiatic										
Hausa	[NG]	168	5453	0.627	0.358	0.457	0.633	0.488	0.320	0.243
Indo-European										
Afrikaans	[ZA]	49	1911	0.373	0.142	0.202	0.443	0.209	0.283	0.211
Nilo-Saharan										
Luo	[UG, KE]	12	179	0.411	0.234	0.229	0.343	0.343	0.309	0.234

Table 6: Test set performance per accent for open-source, commercial, and fine-tuned ASR models.

model compared to the baselines and 49.8% average relative improvement over the commercial models (Azure, GCP, AWS). These results demonstrate significant generalizability gains are achievable with better training data diversity.

5.9 Take SOTA LibriSpeech Results with a Grain of Salt

Figure 2 contrasts LibriSpeech and AfriSpeech WER for several models. Many ASR leaderboards rank ASR models based on single-digit LibriSpeech (Panayotov et al., 2015) WER. Pre-trained ASR models, therefore, overfit to LibriSpeech at the expense of robust ASR performance for all people. As seen in Table 4, several models are 3-10x worse on African accented speech with the exception of multi-lingual or multi-task models like Whisper, Conformer, and XLSR.

6 Limitations and Future Work

Limited Clinical Subdomains: Although this dataset includes a variety of clinical text, several specialties are not represented. As a result, ASR performance may vary between clinical specialties.

Read Speech: All audio samples in this release are read based on text prompts. Without appropriate augmentation, ASR Models trained on this dataset may underperform with conversational or spontaneous speech.

North-African Accents are not included in this work. Because of the distinct nature of those accents, performance on sub-Saharan accents may not necessarily generalize to the Northern African Region.

Self-reported Accents: Similar to Common-Voice, recorders self-report their native tongue in free-text making it difficult to map to ISO-3 in all cases. Some users also reported their accents as ‘‘French’’, ‘‘English’’, ‘‘South African English’’, or a combination of accents. Although we attempted to clean and normalize the self-reported languages, this process was by no means perfect. As a result, accent names sometimes overlap (e.g., Zulu and IsiZulu). Further cleanup could be done to consolidate these closely related accents. The dataset release will therefore include a normalized accent field for each sample.

Medical Abbreviations are Inconsistent: Since crowd-sourced recorders had varying levels of familiarity with the prompts, abbreviations like ‘‘Breast CA’’ may be pronounced fully as




Figure 2: WER on LibriSpeech vs AfriSpeech for selected pre-trained models and commercial ASR systems.

“Breast Cancer” or “Breast see-A”. Since abbreviations abound in medical text and WER is not robust to such idiosyncrasies, models with correct predictions, e.g., “Breast Cancer” are sometimes wrongly penalized where the transcript reads “Breast CA”.

Integrating ASR in Healthcare Settings is Challenging: Cloud-based ASR presents some well-known challenges in healthcare. Privacy is a major concern as there is a risk of unauthorized or malicious third-party access to confidential patient information. Furthermore, the perceived higher value of healthcare data among malefactors also heightens security risks for hospitals and ASR vendors. Additionally, Unethical ASR vendors could misuse confidential data for model training and development without proper consent.

7 Ethical Considerations

Clinical ASR models can improve productivity for clinicians, they can also increase documentation errors, especially through incorrect transcription of numbers, fractions, dates, and proper nouns which have legal, safety, and prognostic implications in healthcare. We caution clinicians to use ASR with full discretion and review transcripts carefully before final submission into the medical record. We release AfriSpeech hoping that

it will be beneficial to clinical and non-clinical use cases within and outside Africa, improving ASR performance for accented speech and it may contain biases due to publicly available datasets. We do not have access to reviewers who are native speakers of most of the languages covered in AfriSpeech who can provide a rigorous review of self-reported accents. This hinders our ability to investigate samples from all languages. We hope that future users of the dataset will further investigate AfriSpeech’s utility and quality for their languages.

Acknowledgments

Tobi Olatunji acknowledges Intron Health for providing the dataset and compute resources. Chris Chinenye Emezue acknowledges the support of the Mila - Quebec AI Institute for compute resources.

References

World Health Organization. Chronic staff shortfalls stifle Africa’s health systems: WHO study — afro.who.int. <https://www.afro.who.int/news/chronic-staff-shortfalls-stifle-africas-health-systems-who-study>. [Accessed 15-Oct-2022].

- Mohammed Abdelwahab and Carlos Busso. 2015. Supervised domain adaptation for emotion recognition from speech. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5058–5062. IEEE. <https://doi.org/10.1109/ICASSP.2015.7178934>
- Tejumade Afonja, Oladimeji Mudele, Iro Orife, Kenechi Dukor, Lawrence Francis, Duru Goodness, Oluwafemi Azeez, Ademola Malomo, and Clinton Mbatuku. 2021. Learning nigerian accent embeddings from speech: Preliminary results based on sautidb-naija corpus. *arXiv preprint arXiv:2112.06199*.
- Christoph Ahlgrim, Oliver Maenner, and Manfred W. Baumstark. 2016. Introduction of digital speech recognition in a specialised outpatient department: A case study. *BMC Medical Informatics and Decision Making*, 16(1):1–8. <https://doi.org/10.1186/s12911-016-0374-4>, PubMed: 27756284
- Adam Ahmat, Sunny C. Okoroafor, Isabel Kazanga, James Avoka Asamani, Jean Jacques Salvador Millogo, Mourtala Mahaman Abdou Illou, Kasonde Mwinga, and Jennifer Nyoni. 2022. The health workforce status in the WHO African region: Findings of a cross-sectional study. *BMJ Global Health*, 7(Suppl 1):e008317. <https://doi.org/10.1136/bmjgh-2021-008317>, PubMed: 35675966
- Richard Anderson, Alex Borucki, Daniel Domingues Da Silva, David Eltis, Paul Lachance, Philip Misevich, and Olatunji Ojo. 2013. Using african names to identify the origins of captives in the transatlantic slave trade: Crowd-sourcing and the registers of liberated Africans, 1808–1862. *History in Africa*, 40(1):165–191. <https://doi.org/10.1017/hia.2013.3>
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M. Tyers, and Gregor Weber. 2019. Common voice: A massively-multilingual speech corpus. *arXiv preprint arXiv:1912.06670*.
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M. Tyers, and Gregor Weber. 2020. Common voice: A massively-multilingual speech corpus. In *LREC*.
- Claire Babirye, Joyce Nakatumba-Nabende, Andrew Katumba, Ronald Ogwang, Jeremy Tusubira Francis, Jonathan Mukiibi, Medadi Ssentanda, Lilian D. Wanzare, and Davis David. 2022. Building text and speech datasets for low resourced languages: A case of languages in east Africa. In *3rd Workshop on African Natural Language Processing*.
- Arun Babu, Changhan Wang, Andros Tjandra, Kushal Lakhota, Qiantong Xu, Naman Goyal, Kritika Singh, Patrick von Platen, Yatharth Saraf, Juan Miguel Pino, Alexei Baevski, Alexis Conneau, and Michael Auli. 2022. Xls-r: Self-supervised cross-lingual speech representation learning at scale. In *INTERSPEECH*. <https://doi.org/10.21437/Interspeech.2022-143>
- Alexei Baevski, Steffen Schneider, and Michael Auli. 2020a. vq-wav2vec: Self-supervised learning of discrete speech representations. *ArXiv*, abs/1910.05453.
- Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. 2020b. wav2vec 2.0: A framework for self-supervised learning of speech representations. *ArXiv*, abs/2006.11477.
- Florence K. Baingana and Eduard R. Bos. 2006. Changing patterns of disease and mortality in sub-saharan Africa: An overview. *Disease and Mortality in Sub-Saharan Africa. 2nd edition*.
- Suzanne V. Blackley, Jessica Huynh, Liqin Wang, Zfania Korach, and Li Zhou. 2019. Speech recognition for clinical documentation from 1990 to 2018: A systematic review. *Journal of the American Medical Informatics Association*, 26(4):324–338. <https://doi.org/10.1093/jamia/ocy179>, PubMed: 30753666
- Suzanne V. Blackley, Valerie D. Schubert, Foster R. Goss, Wasim Al Assad, Pamela M. Garabedian, and Li Zhou. 2020. Physician use of speech recognition versus typing in clinical documentation: A controlled observational study. *International Journal of Medical Informatics*, 141:104178. <https://doi.org/10.1016/j.ijmedinf.2020.104178>, PubMed: 32521449
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla

- Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901.
- Frederick Bukachi and Neil Pakenham-Walsh. 2007. Information technology for health in developing countries. *Chest*, 132(5):1624–1630. <https://doi.org/10.1378/chest.07-1760>, PubMed: 17998362
- Guoguo Chen, Shuzhou Chai, Guan-Bo Wang, Jiayu Du, Weiqiang Zhang, Chao Weng, Dan Su, Daniel Povey, Jan Trmal, Junbo Zhang, Mingjie Jin, Sanjeev Khudanpur, Shinji Watanabe, Shuaijiang Zhao, Wei Zou, Xiangang Li, Xuchen Yao, Yongqing Wang, Yujun Wang, Zhao You, and Zhiyong Yan. 2021. Gigaspeech: An evolving, multi-domain asr corpus with 10, 000 hours of transcribed audio. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2021-1965>
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al. 2022. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1505–1518. <https://doi.org/10.1109/JSTSP.2022.3188113>
- Yi-Chen Chen, Zhaojun Yang, Ching feng Yeh, Mahaveer Jain, and Michael L. Seltzer. 2020. Aipnet: Generative adversarial pre-training of accent-invariant networks for end-to-end speech recognition. *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6979–6983. <https://doi.org/10.1109/ICASSP40776.2020.9053098>
- Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdel rahman Mohamed, and Michael Auli. 2021. Unsupervised cross-lingual representation learning for speech recognition. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2021-329>
- Nilaksh Das, S. Bodapati, Monica Sunkara, Sundararajan Srinivasan, and Duen Horng Chau. 2021. Best of both worlds: Robust accented speech recognition with adversarial transfer learning. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2021-1888>
- Ali Davody, David Ifeoluwa Adelani, Thomas Kleinbauer, and Dietrich Klakow. 2022. TOKEN is a MASK: Few-shot named entity recognition with pre-trained language models. In *Text, Speech, and Dialogue - 25th International Conference, TSD 2022, Brno, Czech Republic, September 6–9, 2022, Proceedings*, volume 13502 of *Lecture Notes in Computer Science*, pages 138–150. Springer. https://doi.org/10.1007/978-3-031-16270-1_12
- Febe De Wet, Philippa Louw, and Thomas Niesler. 2007. Human and automatic accent identification of Nguni and Sotho black South African English. *South African Journal of Science*, 103(3):159–164.
- Rezarta Islamaj Doğan, Robert Leaman, and Zhiyong Lu. 2014. Ncbi disease corpus: A resource for disease name recognition and concept normalization. *Journal of Biomedical Informatics*, 47:1–10. <https://doi.org/10.1016/j.jbi.2013.12.006>, PubMed: 24393765
- Bonaventure F. P. Dossou and Chris C. Emezue. 2021. Okwugb\`e: End-to-end speech recognition for Fon and Igbo. *arXiv preprint arXiv:2103.07762*.
- Moussa Doumbouya, Lisa Einstein, and Chris Piech. 2021. Using radio archives for low-resource speech recognition: Towards an intelligent virtual assistant for illiterate users. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14757–14765. <https://doi.org/10.1609/aaai.v35i17.17733>
- David Eberhard, Gary Simons, and Chuck Fennig. 2019. *Ethnologue: Languages of the World, 22nd Edition*.
- Naome Etori, Ebasa Temesgen, and Maria Gini. 2023. What we know so far: Artificial intelligence in African healthcare. *arXiv preprint arXiv:2305.18302*.

- Foster R. Goss, Suzanne V. Blackley, Carlos A. Ortega, Leigh T. Kowalski, Adam B. Landman, Chen-Tan Lin, Marie Meteer, Samantha Bakes, Stephen C. Gradwohl, David W. Bates, and Li Zhou. 2019. A clinician survey of using speech recognition for clinical documentation in the electronic health record. *International Journal of Medical Informatics*, 130:103938. <https://doi.org/10.1016/j.ijmedinf.2019.07.017>, PubMed: 31442847
- Ama de Graft Aikins, Nigel Unwin, Charles Agyemang, Pascale Allotey, Catherine Campbell, and Daniel Arhinful. 2010. Tackling Africa's chronic disease burden: From the local to the global. *Globalization and Health*, 6(1):1–7. <https://doi.org/10.1186/1744-8603-6-5>, PubMed: 20403167
- Jonatas Grosman. 2021. Fine-tuned XLSR-53 large model for speech recognition in English. <https://huggingface.co/jonatasgrosman/wav2vec2-large-xlsr-53-english>
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*. <https://doi.org/10.21437/Interspeech.2020-3015>
- Alexander Gutkin, Isin Demirsahin, Oddur Kjartansson, Clara E. Rivera, and Kólá Túbòsún. 2020. Developing an open-source corpus of Yoruba speech. <https://doi.org/10.21437/Interspeech.2020-1096>
- Muhammad Ahmed Hassan, Asim Rehmat, Muhammad Usman Ghani Khan, Muhammad Haroon Yousaf, and Muhammad Fazal Ijaz. 2022. Improvement in automatic speech recognition of south asian accent using transfer learning of deepspeech2. *Mathematical Problems in Engineering*, 2022. <https://doi.org/10.1155/2022/6825555>
- Bernd Heine and Derek Nurse. 2000. *African Languages: An Introduction*. Cambridge University Press.
- François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia Tomashenko, and Yannick Esteve. 2018. Ted-lium 3: Twice as much data and corpus repartition for experiments on speaker adaptation. In *International Conference on Speech and Computer*, pages 198–208. Springer. https://doi.org/10.1007/978-3-319-99579-3_21
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460. <https://doi.org/10.1109/TASLP.2021.3122291>
- Sharon Ibejeh, Wuraola Fisayo Oyewusi, Olubayo Adekanmbi, and Opeyemi Osakuade. 2022. EDUSTT: In-domain speech recognition for Nigerian accented educational contents in English. In *3rd Workshop on African Natural Language Processing*.
- Tahir Javed, Sumanth Doddapaneni, Abhigyan Raman, Kaushal Santosh Bhogale, Gowtham Ramesh, Anoop Kunchukuttan, Pratyush Kumar, and Mitesh M. Khapra. 2022. Towards building asr systems for the next billion users. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 10813–10821. <https://doi.org/10.1609/aaai.v36i10.21327>
- Herman Kamper and Thomas Niesler. 2011. Multi-accent speech recognition of Afrikaans, black and white varieties of South African English. In *Twelfth Annual Conference of the International Speech Communication Association*. <https://doi.org/10.21437/Interspeech.2011-798>
- Yohannes Kinifu, Mario R. Dal Poz, Hugo Mercer, and David B. Evans. 2009. The health worker shortage in Africa: Are enough physicians and nurses being trained? <https://doi.org/10.2471/BLT.08.051599>, PubMed: 19377719
- Allison Koenecke, Andrew Nam, Emily Lake, Joe Nudell, Minnie Quartey, Zion Mengesha, Connor Toups, John R. Rickford, Dan Jurafsky, and Sharad Goel. 2020. Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences*, 117(14):7684–7689. <https://doi.org/10.1073/pnas.1915768117>, PubMed: 32205437
- Jialu Li, Vimal Manohar, Pooja Chitkara, Andros Tjandra, Michael Picheny, Frank Zhang,

- Xiaohui Zhang, and Yatharth Saraf. 2021. Accent-robust automatic speech recognition using supervised and unsupervised wav2vec embeddings. *arXiv preprint arXiv:2110.03520*.
- Abdulaziz Y. Lodhi. 1993. The language situation in africa today. *Nordic Journal of African Studies*, 2(1):11–11.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Tarisai Kudakwashe Manyati and Morgen Mutsau. 2021. A systematic review of the factors that hinder the scale up of mobile health technologies in antenatal care programmes in sub-saharan Africa. *African Journal of Science, Technology, Innovation and Development*, 13(1):125–131. <https://doi.org/10.1080/20421338.2020.1765479>
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models. <https://doi.org/10.48550/arXiv.1609.07843>
- Saraladevi Naicker, John B. Eastwood, Jacob Plange-Rhule, and Roger C. Tutt. 2010. Shortage of healthcare workers in sub-saharan Africa: A nephrological perspective. *Clinical Nephrology*, 74:S129–33. <https://doi.org/10.5414/CNP74S129>, PubMed: 20979978
- Saraladevi Naicker, Jacob Plange-Rhule, Roger C. Tutt, and John B. Eastwood. 2009. Shortage of healthcare workers in developing countries–Africa. *Ethnicity & Disease*, 19(1):60. PubMed: 19484878
- Oathokwa Nkomazana, Robert Mash, Sheila Shaibu, and Nthabiseng Phaladze. 2015. Stakeholders’ perceptions on shortage of healthcare workers in primary healthcare in Botswana: Focus group discussions. *PloS One*, 10(8):e0135846. <https://doi.org/10.1371/journal.pone.0135846>, PubMed: 26284617
- Perez Ogayo, Graham Neubig, and Alan W. Black. 2022. Building African voices. *arXiv preprint arXiv:2207.00688*. <https://doi.org/10.21437/Interspeech.2022-152>
- Hilary I. Okagbue, Abiodun A. Opanuga, Muminu O. Adamu, Paulinus O. Ugwoke, Emmanuela C. M. Obasi, and Grace A. Eze. 2017. Personal name in Igbo culture: A dataset on randomly selected personal names and their statistical analysis. *Data in Brief*, 15:72–80. <https://doi.org/10.1016/j.dib.2017.08.045>, PubMed: 28971125
- Kayode Olaleye, Dan Oneață, and Herman Kamper. 2022. Yfacc: A Yorùbá speech-image dataset for cross-lingual keyword localisation through visual grounding. *ArXiv*, abs/2210.04600.
- Obinna O. Oleribe, Jenny Momoh, Benjamin S. C. Uzochukwu, Francisco Mbofana, Akin Adebisi, Thomas Barbera, Roger Williams, and Simon D. Taylor-Robinson. 2019. Identifying key challenges facing healthcare systems in Africa and potential solutions. *International Journal of General Medicine*, 12:395. <https://doi.org/10.2147/IJGM.S223882>, PubMed: 31819592
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: An asr corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210. <https://doi.org/10.1109/ICASSP.2015.7178964>
- Komal Pawar and Urmila Shrawankar. 2016. Question systematization using templates. In *3rd International Conference on Computing for Sustainable Global Development*.
- Archiki Prasad and Preethi Jyothi. 2020. How accents confound: Probing for accent information in end-to-end speech recognition systems. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3739–3753, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.345>
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. Robust speech recognition via large-scale weak supervision. *arXiv preprint arXiv:2212.04356*.
- Mirco Ravanelli, Titouan Parcollet, Peter Plantinga, Aku Rouhe, Samuele Cornell, Loren Lugosch, Cem Subakan, Nauman Dawalatabad, Abdelwahab Heba, Jianyuan Zhong, Ju-Chieh Chou, Sung-Lin Yeh, Szu-Wei Fu, Chien-Feng Liao, Elena Rastorgueva, François Grondin, William Aris, Hwidong Na, Yan Gao, Renato

- De Mori, and Yoshua Bengio. 2021. Speech-brain: A general-purpose speech toolkit. *arXiv preprint arXiv:2106.04624*.
- Ramon Sanabria, Nikolay Bogoychev, Nina Markl, Andrea Carmantini, Ondrej Klejch, and Peter Bell. 2023. The Edinburgh International Accents of English corpus: Towards the democratization of English asr. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE. <https://doi.org/10.1109/ICASSP49357.2023.10095057>
- Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. 2019. wav2vec: Unsupervised pre-training for speech recognition. In *Proceedings of Interspeech 2019*, pages 3465–3469. <https://doi.org/10.21437/Interspeech.2019-1873>
- Kathleen Siminyu, Godson Kalipe, Davor Orlic, Jade Abbott, Vukosi Marivate, Sackey Freshia, Prateek Sibal, Bhanu Neupane, David I. Adelani, Amelia Taylor, Jamiil Toure ALI, Kevin Degila, Momboladji Balogoun, Thierno Ibrahima DIOP, Davis David, Chayma Fourati, Hatem Haddad, and Malek Naski. 2021. Ai4d-african language program. *arXiv preprint arXiv:2104.02516*.
- Sining Sun, Ching feng Yeh, Mei-Yuh Hwang, Mari Ostendorf, and Lei Xie. 2018. Domain adversarial training for accented speech recognition. *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4854–4858. <https://doi.org/10.1109/ICASSP.2018.8462663>
- Sining Sun, Binbin Zhang, Lei Xie, and Yanning Zhang. 2017. An unsupervised deep domain adaptation approach for robust speech recognition. *Neurocomputing*, 257:79–87. <https://doi.org/10.1016/j.neucom.2016.11.063>
- Monica Sunkara, Srikanth Ronanki, Kalpit Dixit, Sravan Bodapati, and Katrin Kirchhoff. 2020. Robust prediction of punctuation and true casing for medical asr. In *ACL 2020 Workshop on NLP for Medical Conversations*. <https://doi.org/10.18653/v1/2020.nlpmc-1.8>
- Jörgen Valk and Tanel Alumäe. 2021. Voxlingua-107: A dataset for spoken language recognition. *2021 IEEE Spoken Language Technology Workshop (SLT)*, pages 652–658. <https://doi.org/10.1109/SLT48900.2021.9383459>
- Markus Vogel, Wolfgang Kaisers, Ralf Wassmuth, and Ertan Mayatepek. 2015. Analysis of documentation speed using web-based medical speech recognition technology: Randomized controlled trial. *Journal of Medical Internet Research*, 17(11):e5072. <https://doi.org/10.2196/jmir.5072>, PubMed: 26531850
- David L. Wheeler, Tanya Barrett, Dennis A. Benson, Stephen H. Bryant, Kathi Canese, Vyacheslav Chetvernin, Deanna M. Church, Michael DiCuccio, Ron Edgar, Scott Federhen, et al. 2007. Database resources of the national center for biotechnology information. *Nucleic Acids Research*, 36(suppl_1):D13–D21. <https://doi.org/10.1093/nar/gkm1000>, PubMed: 18045790
- Wikipedia contributors. 2023a. Demographics of africa — Wikipedia, the free encyclopedia. https://en.wikipedia.org/w/index.php?title=Demographics_of_Africa&oldid=1132870977. [Online; accessed 20-January-2023].
- Wikipedia contributors. 2023b. Languages of africa — Wikipedia, the free encyclopedia. https://en.wikipedia.org/w/index.php?title=Languages_of_Africa&oldid=1133594141. [Online; accessed 20-January-2023].
- Wikipedia contributors. 2023c. List of cities in africa by population — Wikipedia, the free encyclopedia. https://en.wikipedia.org/w/index.php?title=List_of_cities_in_Africa_by_population&oldid=1146587606. [Online; accessed 31-March-2023].
- Aditya Yadavalli, Ganesh S. Mirishkar, and Anil Kumar Vuppala. 2022. Multi-task end-to-end model for telugu dialect and speech recognition. In *Interspeech*. <https://doi.org/10.21437/Interspeech.2022-10739>
- Yuan Yao, Bowen Dong, Ao Zhang, Zhengyan Zhang, Ruobing Xie, Zhiyuan Liu, Leyu Lin, Maosong Sun, and Jianyong Wang. 2022.

Prompt tuning for discriminative pre-trained language models. *arXiv preprint arXiv:2205.11166*. <https://doi.org/10.18653/v1/2022.findings-acl.273>

Yuanyuan Zhang, Yixuan Zhang, Bence Halpern, Tanvina Patel, and Odette Scharenborg. 2022. Mitigating bias against non-native accents. In *Proceedings of the Annual Conference of the International Speech Communication Association, Interspeech 2022*, pages 3168–3172. <https://doi.org/10.21437/Interspeech.2022-836>

A Appendix

A.1 Transcript Preprocessing

Date and Time Replacement: Dates are a critical part of clinical documentation as they typically contain several references to dates and times, for example, date of admission, date of discharge, time of death, and so on. Sampled subsets of sentences containing data and time references from the clinical and general domain were randomly replaced with random dates and times in different formats including “10/12/1999”, “10th December, 1999”, “10th Dec, 1999”, “10-12-1999”, “Mon 10 Dec, 1999”, “Monday 10th December, 1999”. Similar timestamp variations were added to our templates.

Cleaning: Final corpus was pre-processed and cleaned by splitting on sentence boundaries, normalizing spaces, removing carriage return characters, removing non-alphanumeric characters except those with important structural or semantic meaning in the clinical domain such as question marks, parenthesis, colon, a hyphen, plus sign, and greater/lesser than sign. We removed transcripts with less than 5 characters and greater than 300 characters.

Privacy and Patient Information: Although the clinical corpora used were already anonymized, we re-examined several sentence samples for inadvertent exposure of patient names. Anonymized datasets with de-identification tokens like [NAME] and [DATE] were replaced with African names and randomly generated dates as described above.

A.2 Annotation Instructions

Recorders were provided with the following instructions:

Accuracy It is very important that the recorded words match the text in the script exactly. If you accidentally deviate from the script, become unsure, or lose track of your thought, please delete and record the prompt again.

Punctuations All punctuations should be pronounced in full, not just observed. That is, when reading a text sample that contains punctuation, you say “comma”, “full stop”, “semi-colon”, “colon”, “slash”, “hyphen”, “question mark”, “exclamation mark”, and so on as appropriate. Brackets should be pronounced as “open bracket” or “close bracket”.

Punctuation Exclusions/Exceptions to the above rule: In measurements or units like “mg/dl”, please say “milligram PER dl” NOT “milligram slash dl”. In situations where “?” is used to represent “query”, please say “query” NOT “question mark”.

Abbreviations Pronounce common short-hand forms (such as r/o, prn, tds, PO, mg, W/O), dates, times, and numbers as you would in a clinical setting. For example, “r/o” should be pronounced “rule out” as usual not “arr slash ohh”. Common Abbreviations SHOULD be pronounced in full. “CT” should be pronounced “see tee” as usual NOT “Computed Tomography”. “CXR” should be pronounced “Chest Xray” as usual NOT “see ex arr” “mmHg” should be pronounced in full as “millimeters of mercury”. Pronounce CA as “Carcinoma” NOT “See Ay”.

Tone Also be sure to use your natural accent. The goal is to build a speech-to-text system that understands African accents. This tool is for us. Be natural.

Speed Do not speak unrealistically fast. While an increased reading speed is recommended, take care to avoid vocal fatigue from rushing through the phrases at lightning speed! This will only result in a lower-quality voice. Record a maximum of 2 hours a day, taking a break every half hour.

A.3 Annotator Management

Consent Recorders signed a Terms of Use agreement and consented to the privacy policy on the recording platform.

Payment Recorders were paid \$5 to \$10 per hour depending on task difficulty and clinical experience. Most recorders considered payment satisfactory compared with task difficulty.

A.4 AfriSpeech Vocabulary

AfriSpeech models use a 50-character vocab including numbers and punctuations and symbols with important semantic roles in healthcare.

“-”, “w”, “a”, “7”, “,”, “0”, “d”, “i”, “:”, “p”, “g”, “u”, “(”, “5”, “1”, “e”, “9”, “j”, “b”, “3”, “s”, “””, “h”, “o”, “+”, “l”, “v”, “y”, “q”, “n”, “2”, “r”, “f”, “m”, “%”, “t”, “/”, “6”, “z”, “?”, “8”, “)”, “x”, “.”, “4”, “c”, “k”, “l”, “[UNK]”, “[PAD]”.