

# AI4D -- AFRICAN LANGUAGE DATASET CHALLENGE

Kathleen Siminyu;Sackey Freshia;Jade Abbott;Vukosi Marivate;

;

© 2021, KATHLEEN SIMINYU



This work is licensed under the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/legalcode>), which permits unrestricted use, distribution, and reproduction, provided the original work is properly credited.

Cette œuvre est mise à disposition selon les termes de la licence Creative Commons Attribution (<https://creativecommons.org/licenses/by/4.0/legalcode>), qui permet l'utilisation, la distribution et la reproduction sans restriction, pourvu que le mérite de la création originale soit adéquatement reconnu.

*IDRC Grant/ Subvention du CRDI: 109187-002-Laying the foundations for artificial intelligence for development (AI4D) in Africa*

# AI4D - African Language Dataset Challenge

**Kathleen Siminyu**

Artificial Intelligence for Development  
Africa  
kathleensiminyu@gmail.com

**Sackey Freshia**

Jomo Kenyatta University  
of Agriculture and Technology  
freshiasackey@gmail.com

**Jade Abbott**

Retro Rabbit  
jabbott@retorabbit.co.za

**Vukosi Marivate**

University of Pretoria

vukosi.marivate@cs.up.ac.za

## Abstract

As language and speech technologies become more advanced, the lack of fundamental digital resources for African languages, such as data, spell checkers and Part of Speech taggers, means that the digital divide between these languages and others keeps growing. This work details the organisation of the AI4D - African Language Dataset Challenge<sup>1</sup>, an effort to incentivize the creation, organization and discovery of African language datasets through a competitive challenge. We particularly encouraged the submission of annotated datasets which can be used for training task-specific supervised machine learning models.

## 1 Introduction

Africa has a language diversity of over 2000 languages, many of which are only spoken and not written (Eberhard et al., 2019).

As language technologies advance and more sophisticated tools are built using Artificial Intelligence, the divide between low resource languages and others is likely to get even larger as a common prerequisite of these advanced systems is the existence of a large amount of digital data. African languages, low resource languages, are at risk of being left behind (Joshi et al., 2020).

Data Science and Machine Learning skills are increasingly becoming widespread on the African continent. This can be attributed to the rise of grassroots capacity building efforts through organisations such as Data Science Africa<sup>2</sup>, Data Science Nigeria<sup>3</sup>, Deep Learning Indaba<sup>4</sup>, as well as NLP-specific communities such as Masakhane<sup>5</sup> (Orife et al., 2020). These movements have facilitated a critical mass of individuals with the relevant skills, who speak African languages that can start contributing to the overall body of work that currently exists and begin the work where none does. With this challenge, we sought to engage the African NLP community in the task of dataset creation.

## 2 Methodology

The realisation of this work was wholly driven by the intended outcome, the need for more African language datasets for use in NLP research.

### 2.1 Framing the Challenge

Early framing of this challenge was predicated on the fact that pre-trained language models are producing state-of-the-art NLP results (Devlin et al., 2018; Radford et al., 2019). African NLP would undoubtedly benefit from the creation of such language models, and so the competition was initially envisioned as a language model challenge. Unfortunately, not only do most pre-trained language models require large amounts of monolingual data to train, they require labelled NLP datasets in order to usefully evaluate the models. This motivated the decision to create a challenge focused on data collection, rather than model building.

<sup>1</sup><https://zindi.africa/competitions/ai4d-african-language-dataset-challenge>

<sup>2</sup><http://www.datascienceafrica.org/>

<sup>3</sup><http://datasciencenigeria.ai>

<sup>4</sup><https://deeplearningindaba.com/>

<sup>5</sup><https://www.masakhane.io/>

## 2.2 Securing Donor Funding

With the aim of creating datasets that are openly available, we prepared and circulated a proposal and succeeded in securing donor funding. That being committed, we tailored the challenge to take place in 2 phases. The initial phase focused on data collection and the second phase being a more conventional machine learning (ML) challenge, where the datasets developed in the first phase could be used as evaluation sets for a pre-trained language model challenge.

## 2.3 Hosting the Challenge

Our target audience was researchers, practitioners and enthusiasts from African countries who could create datasets for the languages that they speak. We approached Zindi<sup>6</sup>, an African data science competition platform, to host the challenge in a bid to leverage their existing user base. They have over 12,000 individuals signed up.

## 2.4 Evaluation of Submissions

Unlike a conventional ML challenge that would have an agreed upon automated metric to evaluate and rank submitted models, evaluation of a challenge of this kind cannot be automated. Instead, we put together a panel of judges with experience in NLP who would review the datasets each month. We also indicated in the challenge guidelines that each dataset submission should be accompanied by a datasheet (Geburu et al., 2018) that documents its motivation, composition, collection process, recommended uses, and so on (Example in Appendix A).

Evaluation was done by judges analysing the datasheets and awarding points to each submission based on a scoring rubric. The rubric took into account the following: how representative and balanced the corpus was, the dataset size in terms of tokens and unique tokens, whether it was annotated for a specific downstream task, under-representation of the language in terms of digital data, methodology of the data collection and labelling process, originality of the data collection and labelling process (Reviewer Documentation in Appendix B).

## 3 Results

Language	Tasks	Submissions
Yoruba	MT, Diacritic Verification, Text Classification, NER, misc	7
Kiswahili	Document Classification, misc	6
Igbo	NER, misc	4
Hausa	Sentiment Analysis, Document classification, misc	4
Fongbe	MT, Speech to Text, misc	3
Amharic	Hate speech detection, stop words list, misc	3
Asante Twi	Sentiment Analysis, MT, misc	3
Chichewa	NER, MT	2
Ewe	MT, misc	2
Wolof	ASR	1
Tunizian Arabizi	Sentiment Analysis	1
Kikuyu	misc	1
Kabiwe	MT	1
Oromo	misc	1
Zulu	misc	1

Table 1: Language and Task distribution of submissions.

The challenge ran for a period of five months with a total of 270 people registering on the Zindi platform to participate.

<sup>6</sup><http://zindi.africa>

The data in the submitted datasets was compiled from a wide variety of sources. These were largely digital sources such as news websites, religious texts, Facebook, Twitter and YouTube. This outcome is likely an indication of the ease of access that online sources present, given that the data is already digitised. Other data creation processes included participants convening to carry out manual translation of existing pieces of text. Unique contributions included OCR to digitize printed texts and using the user base of a commercial application to crowd source and validate recordings of phrases and texts common on the platform.

#### **Observations and Lessons Learned:**

- Teams composed of individuals from relevant multi-disciplinary backgrounds, including computer scientists, professional translators and linguists, were able to create and annotate datasets that captured fundamental lexical and semantic nuances of languages.
- The challenge framing allowed for anyone to participate. While useful as an exercise in evaluating the interest in such a challenge, the top evaluated submissions came from teams who had been exposed to NLP research work. Targeting such a challenge to NLP researchers could lead to higher quality submissions in future.
- A portion of submissions contained very few data points. As the aim is to use the datasets for NLP, in future, we'd set explicit minimum requirements with regards to the size of datasets admissible, file formats and require the inclusion of any cleaning or pre-processing code used.
- Since the challenge was evaluated monthly, we often received disparate submissions from the same teams as they managed to obtain more data. Instead, one large dataset built over a couple of months would have been the ideal outcome, so in future we'd select and support teams for a sustained period of time to enable them build sizeable datasets.
- Participants and judges had to rely on their own understanding of what "Representative and Balanced" means in a dataset. In future, more specificity of what "representative" and "balanced" means would enable participants to produce better datasets.

## **4 Future Work**

A large number of opportunities were identified to support future work in African, and low resource, language dataset creation, as follows:

- Research and analysis of the legal implications of obtaining textual, visual and audio data from a variety of online sources, which were noted as a common source among participants. The copyright and intellectual property implications will have to be thoroughly assessed ahead of the publication and further public use of relevant datasets.
- Outlining of best case practise techniques for protecting the identities and privacy of users, in instances where data is obtained from social media/content platforms like Twitter, Facebook and YouTube. Social media sites have been noted as a common data source.
- Recommended techniques for identifying and ascertaining whether data obtained from online sources(news publications,social media and content platforms) contains biased sentiments(sexist, racist) and offensive material(hateful), as well as techniques for removing any biased sentiment and offensive material, if need be, depending on the use of the dataset.

Courtesy of this dataset creation challenge, we have secured further funding to support 5 of the top teams for a 6 month period. During this time, they will further flesh out their datasets with the aim of using these as the basis of future NLP challenges/shared tasks. This project will also be used as a model case to inform evidence-based policy making concerning Artificial Intelligence and we hope that it will be replicated to support the development of data for other low resource languages.

## Acknowledgements

This work has been funded through a partnership between the International Development Research Centre, the Swedish International Development Cooperation Agency, the Knowledge4All Foundation, Zindi Africa and the AI4D-Africa Network. The expert panel that volunteered their time to undertake the difficult qualitative task of dataset assessment was composed of Jade Abbott - Retro Rabbit, John Quinn - Google AI / Makerere University, Kathleen Siminyu - AI4D-Africa, Veselin Stoyanov - Facebook AI and Vukosi Marivate - University of Pretoria.

## References

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- David M Eberhard, Gary F. Simons, and Charles D. Fenning. 2019. *Ethnologue: Languages of the worlds*. twenty-second edition.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumeé III, and Kate Crawford. 2018. Datasheets for datasets. *arXiv preprint arXiv:1803.09010*.
- Pratik M. Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The state and fate of linguistic diversity and inclusion in the nlp world. *ArXiv*, abs/2004.09095.
- Iroro Orife, Julia Kreutzer, Blessing Sibanda, Daniel Whitenack, Kathleen Siminyu, Laura Martinus, Jamiil Toure Ali, Jade Abbott, Vukosi Marivate, Salomon Kabongo, et al. 2020. Masakhane-machine translation for africa. *arXiv*, pages arXiv–2003.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.

---

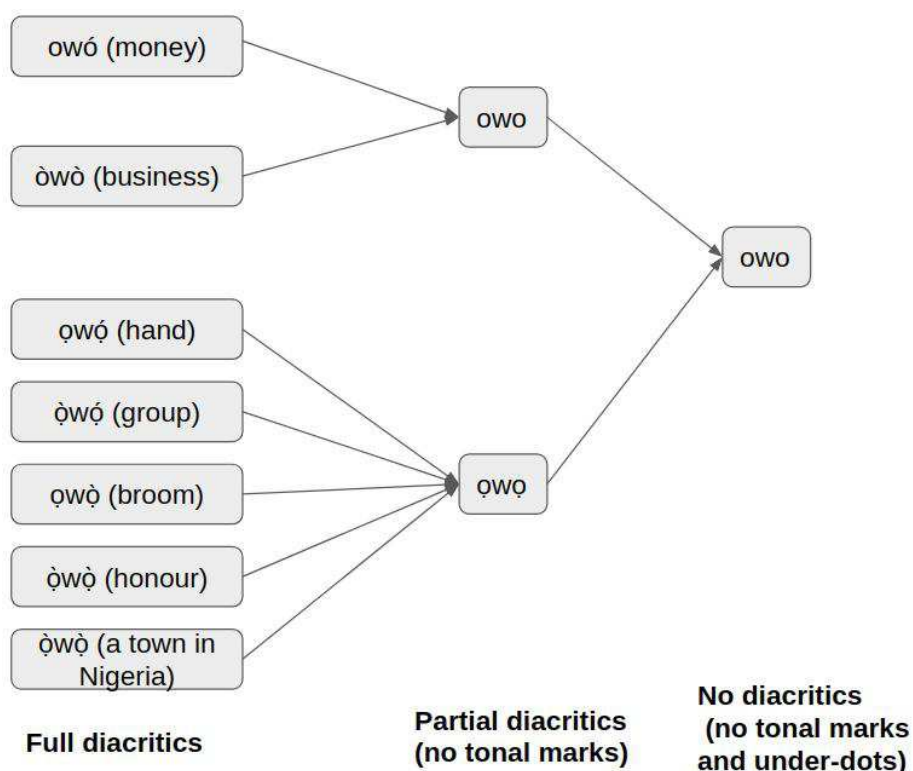
This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

## Multi-domain machine translation dataset for Yoruba Language

### Motivation

**For what purpose was the dataset created? Was there a specific task in mind? Was there a specific gap that needed to be filled? Please provide a description**

The dataset was created to train a machine translation model from English language to Yoruba language. The major contribution of this dataset is that the english-yoruba parallel sentences come from different domain like technology, medicine, science, movie transcripts, book translation and news articles. This is very uncommon for low-resource languages. The most available parallel text corpora are biased to religious text like Bible and JW300. We provide high quality Yoruba corpora with proper diacritics. The diacritics problem of Yoruba is summarized in the diagram below:



**Who created this dataset (e.g., which team, research group) and on behalf of which entity (e.g., company, institution, organization)?**

The dataset was created by Di ( ), Di ( ) and Qi ( ). Di contacted them because they are the major

contributors to the Yoruba language section of the Global Voices (<https://yo.globalvoices.org/> ). Another reason for contacting them is is to verify that all the corpora that we are submitting have the correct diacritics. Improper diacritics affects the performance on down stream NLP tasks (See <https://arxiv.org/abs/1912.02481>)

**Who funded the creation of the dataset? If there is an associated grant, please provide the name of the grantor and the grant name and number**

We did not receive any funding for the creation of this dataset.

## **Composition**

**What do the instances that comprise the dataset represent (e.g., documents, photos, people, countries)?**

The dataset compose of text data from different domains like technology, medicine, science, movie transcripts, book translation and news articles.

**How many instances are there in total (of each type, if appropriate)?**

In total, there are 6047 parallel sentences in all the CSV files from 5 sources:

- (1) 45 news articles (or 119 parallel sentences) from Global Voices.
- (2) 2700 parallel sentences from Yoruba Proverbs
- (3) First 6 chapters (or 862 parallel sentences) from “Out of his mind book”
- (4) 812 parallel sentences from Unsane movie on YouTube
- (5) 549 Multi-domain sentences comprising of technological and scientific terms.

**Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set?**

We took only 6 chapters of the 12 chapters of the Out of his mind book due to not enough time to manually extract the parallel sentences before submission. We hope to complete it if we are successful in the competition

**What data does each instance consist of?**

Each instance consist of sentence id, English sentence and Yoruba translation arranged in a csv format.

## **Collection Process**

## **How was the data associated with each instance acquired?**

- (1) 45 news articles (or parallel 119 sentences) crawled from Global Voices. (<https://yo.globalvoices.org/>) and (<https://globalvoices.org/>). For each Yoruba news article, we manually search for the English equivalence.
- (2) 2700 parallel sentences from Yoruba Proverbs were crawled from [https://twitter.com/yoruba\\_proverbs](https://twitter.com/yoruba_proverbs)
- (3) First 6 chapters (or 862 parallel sentences) from “Out of his mind book”. The entire book was translated by Dr. [Oluwatobi Oluwalaniran](#) for her Bachelors thesis, and has been edited by the supervisor. The quality is excellent.
- (4) 812 parallel sentences from Unsane movie on YouTube was translated by Dr. [Oluwatobi Oluwalaniran](#). This is what she does for living.
- (5) 549 Multi-domain sentences comprising of technological and scientific terms were translated by Dr. [Oluwatobi Oluwalaniran](#)

## **Preprocessing/cleaning/labeling**

The main preprocessing done was to align the english and Yoruba sentences into a CSV format. Some were aligned with scripts and others manually depending on the difficulty of the corpora.

## **Uses**

### **Has the dataset been used for any tasks already? If so, please provide a description**

The dataset has not been used for machine translation. We will probably be the first.

### **What (other) tasks could the dataset be used for?**

The dataset can also be used for training language models, automatic diacritic restoration.

### **Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses?**

No risk that we are aware of

## **Distribution**

### **Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created?**

The dataset will be publicly made available after the competition

### **How will the dataset will be distributed (e.g., tarball on website, API, GitHub)?**

By Github or Google Drive. The link will be available after the competition.

### **When will the dataset be distributed?**



Hopefully, January 2020.

**Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)?**

Creative Commons license.

**Have any third parties imposed IP-based or other restrictions on the data associated with the instances?**

No

**Do any export controls or other regulatory restrictions apply to the dataset or to individual instances?**

No

## **Maintenance**

**Who is supporting/hosting/maintaining the dataset?**

Digital Language Acquisition

**How can the owner/curator/manager of the dataset be contacted (e.g., email address)?**

[dla@uwaterloo.ca](mailto:dla@uwaterloo.ca) or [dla@uwaterloo.ca](mailto:dla@uwaterloo.ca) gmail.com

**Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)? If so, please describe how often, by whom, and how updates will be communicated to users (e.g., mailing list, GitHub)?**

The dataset will be updated when we get more parallel corpora.

## REVIEWER DOCUMENTATION

### **Representative and Balanced (30%) - 12 points**

- 12-18 points: The corpus is very representative and cuts across several different genres eg literary fiction, blogs, newspaper articles, spoken speech etc, and balanced with various genres being equally represented
- 6-12 points: A little representative and is comprised of two or more genres and somewhat balanced
- 1-6 points: Not representative and derived from one genre

### **(Automated) Number of tokens and number of unique tokens (25%) - 10 points**

- 1-5 points: Large number of tokens and large number of unique tokens
- 5-10 points: Large number of either tokens or unique tokens

### **Annotated for a specific downstream task (15%) - 6 points**

- 3-6 points: Specific enough downstream task to evaluate future language models on
- 1-3 points: Not tailored for specific enough downstream task

### **Is the language underrepresented on the Internet and in terms of available digital data? (15%) - 6 points**

- 3-6 points: Very underrepresented language
- 1-3 points: A little underrepresented

### **Methodology of data collection and labelling process (7.5%) - 3 points**

- Is the collection and labelling methodology sound, relevant and well documented?

### **Originality of data collection and labelling process (7.5%) - 3 points**

- How original and creative is the data collection and labelling process? Eg. Including data from JW300 as a source versus translating poetry from a book that had previously never been digitised.