

Yankari: Monolingual Yoruba Dataset

Maro Akpobi

African Center for Language Preservation

maro@acflp.org

Abstract

This paper presents Yankari, a large-scale monolingual dataset for the Yoruba language, aimed at addressing the critical gap in Natural Language Processing (NLP) resources for this important West African language. Despite being spoken by over 30 million people, Yoruba has been severely underrepresented in NLP research and applications. We detail our methodology for creating this dataset, which includes careful source selection, automated quality control, and rigorous data cleaning processes. The Yankari dataset comprises 51,407 documents from 13 diverse sources, totaling over 30 million tokens. Our approach focuses on ethical data collection practices, avoiding problematic sources and addressing issues prevalent in existing datasets. We provide thorough automated evaluations of the dataset, demonstrating its quality compared to existing resources. The Yankari dataset represents a significant advancement in Yoruba language resources, providing a foundation for developing more accurate NLP models, supporting comparative linguistic studies, and contributing to the digital accessibility of the Yoruba language.

1 Introduction

Natural Language Processing (NLP) has made tremendous strides in recent years, yet these advancements have primarily benefited high-resource languages, leaving many African languages, including Yoruba, underrepresented in NLP research and applications. This paper introduces Yankari, a large-scale, high-quality monolingual dataset for Yoruba, a language spoken by over 30 million people in West Africa. Despite its significant speaker population, Yoruba has long suffered from a lack of comprehensive, ethically-sourced language resources suitable for modern NLP tasks.

Yankari addresses this critical gap by providing a carefully curated corpus of 51,407 documents from 13 diverse sources, totaling over 30 million tokens. Our methodology prioritizes ethical data collection, rigorous quality control, and the preservation of linguistic authenticity. By avoiding problematic sources such as religious texts and machine-translated content, Yankari offers a more balanced and representative sample of contemporary Yoruba language use.

This paper details our data collection and processing pipeline, discusses the challenges encountered in creating resources for low-resource languages, and provides a thorough analysis of the dataset’s composition and potential biases. We also offer insights into the ethical considerations surrounding the creation and use of such resources. Through Yankari, we aim to facilitate the development of more accurate and culturally appropriate NLP models for Yoruba, contribute to the preservation of linguistic diversity in the digital age, and provide a replicable approach for creating high-quality datasets for other low-resource languages.

To our knowledge, Yankari represents the first large-scale, non-religious domain monolingual resource created specifically for Yoruba. This work not only provides a valuable asset for Yoruba NLP but also offers a replicable approach for developing similar datasets for other low-resource languages.

Source	Docs	Type	Domain
yo.wikipedia.org	16 809	Encyclopedia	General
alaroye.org	10 535	News	Current Affairs
www.bbc.com	8 252	News	Current Affairs
www.awikonko.com.ng	5 438	Blog	Culture
yoruba.von.gov.ng	2 542	News	Current Affairs
sportsinyoruba.wordpress.com	2 328	Blog	Sports
www.asejere.net	2 079	Blog	Entertainment
asa.ooduarere.com	1 744	Blog	Culture
radionigeriaibadan.gov.ng	824	News	Current Affairs
iroyinowuro.com.ng	603	News	Current Affairs
oroyoruba.blogspot.com	139	Blog	Culture
yo.globalvoices.org	81	News	Current Affairs
edeyorubarewa.com	33	Blog	Fashion
Total	51 407		

Table 1: Data sources for the Yankari dataset.

2 Related Works

This section examines existing resources for Yoruba NLP, highlighting their limitations and the need for a comprehensive, ethically-sourced Yoruba dataset.

2.1 Monolingual Yoruba Datasets

2.1.1 Yorùbá Text C3

[Alabi et al. \(2020\)](#) introduced Yorùbá Text C3, compiled from various web sources including the Bible, JW300 ([Agić and Vulić, 2019](#)), books, news articles, and Wikipedia. While broad in scope, this dataset is heavily skewed towards religious content, particularly Christianity. This bias significantly limits its utility for NLG tasks requiring balanced and diverse text. Moreover, the inclusion of JW300 data raises serious ethical and legal concerns. [Hutchinson \(2024\)](#) points out that the Jehovah’s Witnesses have explicitly prohibited the use of their data in NLP research, making the continued use of JW300 not just ethically questionable but potentially illegal.

2.1.2 MENYO-20k

[Adelani et al. \(2021\)](#) developed MENYO-20k, a multi-domain English-Yoruba corpus primarily for machine translation tasks. While it offers more diverse content, its relatively small size of 20,100 sentences and focus on translation rather than monolingual text generation limit its applicability for large-scale NLG tasks.

2.2 Multilingual Datasets Including Yoruba

2.2.1 Wura Dataset

The Wura dataset, developed by [Oladipo et al. \(2023\)](#), is a multilingual dataset containing approximately 68,000 Yoruba documents. It integrates content from JW300 and Wikipedia, inheriting similar biases and ethical issues as Yorùbá Text C3. Our manual inspection of the Wura dataset revealed critical quality issues not previously reported, including formatting errors and inappropriate content.

2.2.2 Large-Scale Web-Crawled Corpora

Multilingual datasets such as mC4, OSCAR, and the Afriberta-Corpus also include Yoruba content. However, these datasets often suffer from noise, poor formatting, and limited source diversity. [Ogueji et al. \(2021\)](#) used the Afriberta-Corpus, which primarily sources data from the BBC News and Common Crawl, resulting in a dataset lacking the domain diversity necessary for robust NLG.

[Nguyen et al. \(2023\)](#) introduced CulturaX, covering 167 languages. However, its Yoruba subset showed an alarmingly high duplication rate of 24.48% and contained machine-translated pages, false positives in language detection, and a significant amount of religious text.

2.3 Large-Scale Multilingual Efforts for African Languages

The Cheetah project by [Adebara et al. \(2024\)](#) focuses on natural language generation for 517 African languages, including Yoruba. This ambitious project utilizes existing corpora from various sources, including OSCAR ([Ortiz Suárez et al.](#),

2020), CC-100 (Conneau et al., 2020), Afriberta-Corpus (Ogueji et al., 2021), and mC4 (Xue et al., 2021).

While Cheetah represents a significant step towards improving NLP capabilities for African languages, it faces challenges common to large-scale multilingual efforts. These include potential issues with data quality, bias towards certain domains or text types, and the inclusion of machine-translated content.

2.4 Ethical Considerations in Using Religious Texts

The ethical implications of using religious texts in NLP, particularly for low-resource languages, are profound and often overlooked. Hutchinson (2024) challenges the NLP community’s casual approach to sacred texts, arguing that the prevalence of Christian texts in datasets for low-resource languages reflects a legacy of colonialism and missionary work. This creates an ethical dilemma where NLP technologies risk becoming unwitting agents of cultural imperialism and religious proselytism.

3 Motivation for Yankari

Given these severe limitations and ethical concerns in existing resources, there is an urgent need for a high-quality, diverse, and extensive monolingual Yoruba dataset that does not rely on restricted or problematic sources. This is where **Yankari** comes in, directly addressing the gaps and ethical issues left by previous datasets.

Yankari aims to provide a large-scale, ethically sourced corpus that represents a wide range of Yoruba language use. By avoiding the pitfalls of previous datasets, such as over-reliance on religious texts or machine-translated content, Yankari seeks to offer a more balanced and authentic representation of the Yoruba language. This approach aligns with recent calls in the NLP community for more thoughtful and ethical dataset creation, particularly for low-resource languages.

4 Methodology

This section details our approach to creating the Yankari dataset, including data collection, processing, quality assurance steps, and corpus analysis.

4.1 Data Collection

Our data collection process focused on gathering content from diverse, high-quality sources to

ensure a representative sample of contemporary Yoruba language use. We carefully selected 13 sources, including news outlets, blogs, educational websites, and Wikipedia. Table 1 provides an overview of these sources and their contributions to the dataset.

4.2 Analysis of Existing Datasets

To inform our data collection and curation process, we conducted a detailed analysis of the Wura dataset (Oladipo et al., 2023). Our investigation revealed several critical issues:

- High repetition: 18.01% of the dataset contains the word 'asteroidi' (asteroid), indicating a significant bias towards astronomical content.
- Duplication: After removing duplicates and cleaning, only 17,103 unique entries remained, representing just 45% of the original dataset.
- Quality issues: We found formatting errors, inappropriate content, and entries in non-Yoruba languages.

These findings highlight the pressing need for more stringent data curation practices in low-resource language datasets. The recent study by Hernandez et al. (2022) on scaling laws and the interpretability of learning from repeated data offers valuable insights that influenced our methodology. Their extensive research reveals that data repetition can severely hinder model performance, particularly by disrupting the balance between memorization and generalization. Additionally, repeated data can obstruct the development of "induction heads," which are vital for in-context learning in large language models. Most importantly, their work underscores the pivotal role that high-quality training data plays in the effectiveness of language models. These insights significantly shaped our careful approach to developing Yankari, underscoring the critical need for rigorous data curation, quality control, and the preservation of diversity in our dataset.

4.3 Data Processing Pipeline

Our data processing pipeline consisted of several key steps:

4.3.1 HTML Parsing and Text Extraction

We implemented a robust HTML parsing system using BeautifulSoup to extract clean text from web pages. This process involved:

- Removing all script and style elements
- Extracting text from relevant HTML tags (e.g., `<p>`, `<h1>`, `<h2>`, etc.)
- Preserving the document structure by maintaining paragraph boundaries

It is important to note that, beyond basic structural parsing, no specific sentence-level filtering based on the presence or absence of terminal punctuation was applied at this stage, as the primary goal was to capture full textual content from the identified relevant HTML elements.

4.3.2 Deduplication

We employed a two-step deduplication process:

1. Exact matching at the document level to remove duplicate web pages
2. Near-duplicate detection at the paragraph level using MinHash and Locality-Sensitive Hashing (LSH) techniques

4.4 Corpus Statistics

Our final Yankari dataset consists of:

- Total number of documents: 51,407
- Total number of tokens: 30,438,702
- Average tokens per document: 592.11

4.5 Domain Analysis

The distribution of web domains in our corpus reflects the diverse sources we targeted. The top 5 domains by number of documents are:

1. yo.wikipedia.org: 32.70%
2. alaroye.org: 20.49%
3. www.bbc.com: 16.05%
4. www.awikonko.com.ng: 10.58%
5. yoruba.von.gov.ng: 4.94%

This distribution ensures a balance between encyclopedic content, news, and cultural discussions, providing a comprehensive representation of written Yoruba across various domains.

4.6 Ethical Considerations and Excluded Content

Throughout our data collection and processing, we prioritized ethical considerations:

- We explicitly removed data from restricted sources, such as those with terms prohibiting use in NLP research.
- We filtered out suspected machine-translated content to maintain linguistic authenticity. This was primarily a manual process conducted by the author, a native Yoruba speaker, during data spot-checking and review. Documents exhibiting unnatural phrasing, common translation artifacts, or content inconsistent with known source characteristics were excluded.
- We removed inappropriate or offensive material to ensure the dataset’s suitability for a wide range of applications. This was also a manual review process performed by the author. The guidelines focused on excluding hate speech, explicit adult content, and other materials generally considered unsuitable for a public research dataset.
- We respected copyright and intellectual property rights. Content was sourced primarily from publicly accessible websites. For sources like Wikipedia, explicit open licenses (e.g., CC-BY-SA) were followed. For other public news and blog content without explicit licenses for redistribution, we operated under the principles of fair use for non-commercial research, providing clear attribution via source URLs. Direct contact for explicit redistribution permission was not feasible for every source due to the scale of collection; this is a recognized challenge in web-scale data gathering and is noted as a limitation.

In line with our commitment to transparency, we acknowledge that our content filtering process may have introduced certain biases:

- The removal of very short texts may have disproportionately affected certain types of content.
- Our focus on standard Yoruba may have led to the underrepresentation of regional dialects or colloquial expressions.

- Excluding content with non-Yoruba characters might have removed some culturally relevant content involving code-switching or borrowings.

4.7 Output Format

The final dataset is stored in JSONL format, with each line containing a separate JSON document with the following fields:

- **text**: The main content of the document in Yoruba.
- **url**: The original URL from which the content was sourced.
- **source**: A code indicating the source of the document.

Here is a sample entry:

```
{
  "text": "O ma se o! Ijamba oko
    ofurufu gba emi eeyan marun-
    un...",
  "url": "https://www.awikonko.
    com.ng/2024/03/o-ma-se-o-
    ijamba-ok-ofurufu-gba-emi.
    html",
  "source": "ACFLP"
}
```

4.8 Quality Assurance

To ensure the highest quality of our dataset:

- We involved native Yoruba speakers in the data cleaning and validation process.
- We conducted regular spot checks throughout the data processing pipeline.
- We performed a final manual review of a randomly selected subset of the data to verify its quality and authenticity.

4.9 Limitations and Potential Biases

We acknowledge the following limitations and potential biases in our dataset:

- Internet Bias
- Written Language Bias
- Source Bias
- Temporal Bias

- Standardization Bias

These limitations highlight areas for future work and expansion of the Yankari dataset.

5 Limitations and Ethical Considerations

While the Yankari dataset represents a significant contribution to Yoruba language resources for NLP, it is important to acknowledge its limitations and the ethical considerations that arise from its creation and potential use.

5.1 Representation Bias

- **Spoken Language**: The dataset does not include samples of spoken Yoruba.
- **Informal Variants**: Internet sources may favor more formal language use.
- **Demographic Skew**: Internet access and content creation are not uniformly distributed across all Yoruba-speaking demographics.

5.2 Diacritization Challenges

- **Tonal Ambiguity**: Incorrect or missing diacritical marks can cause ambiguity.
- **Standardization Issues**: The lack of a universally adopted standard for Yoruba orthography may result in inconsistencies.

6 Conclusion

The Yankari dataset represents a significant step forward in addressing the resource gap for Yoruba in Natural Language Processing. By providing a large-scale, high-quality, and ethically sourced corpus, we have laid a foundation for advancing NLP research and applications in this important West African language. Our rigorous methodology, which prioritizes data quality, diversity, and ethical considerations, sets a new standard for the development of language resources for low-resource languages.

The creation of Yankari highlights several critical challenges in developing NLP resources for languages like Yoruba, including the scarcity of diverse, high-quality online content, the complexities of automated processing for languages with limited existing NLP tools, and the ethical considerations surrounding data collection and potential misuse. By transparently discussing these challenges and our approaches to addressing them, we hope to contribute to the broader conversation on responsible

AI development for diverse languages and cultures. Future work will focus on rigorously evaluating Yankari’s utility in various downstream NLP tasks, such as language modeling, machine translation, and text classification, including comparative performance analyses against other available Yoruba corpora. Such evaluations will further quantify the benefits of Yankari’s curated nature and diverse domain coverage.

The dataset is available on Hugging Face: <https://huggingface.co/datasets/acflp/YANKARI>

Acknowledgments

This work was performed independently. The author received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors for the research described in this manuscript.

References

- Ife Adebara, AbdelRahim Elmadany, and Muhammad Abdul-Mageed. 2024. [Cheetah: Natural language generation for 517 african languages](#). *Preprint*, arXiv:2401.01053.
- David Adelani et al. 2021. Menyo-20k: A multi-domain english - yoruba corpus for machine translation and domain adaptation. In *Proceedings of the Fourth Workshop on Technologies for MT of Low Resource Languages*.
- Željko Agić and Ivan Vulić. 2019. [JW300: A wide-coverage parallel corpus for low-resource languages](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3204–3210, Florence, Italy. Association for Computational Linguistics.
- Jesujoba Alabi et al. 2020. Massive vs. curated embeddings for low-resourced languages: The case of yorùbá and twi. In *Proceedings of the 12th Language Resources and Evaluation Conference*.
- 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.
- Danny Hernandez, Tom Brown, Tom Conerly, Nova DasSarma, Dawn Drain, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Tom Henighan, Tristan Hume, Scott Johnston, Ben Mann, Chris Olah, Catherine Olsson, Dario Amodei, Nicholas Joseph, Jared Kaplan, and Sam McCandlish. 2022. [Scaling laws and interpretability of learning from repeated data](#). *Preprint*, arXiv:2205.10487.
- Ben Hutchinson. 2024. Modeling the sacred: Considerations when using religious texts in natural language processing. *arXiv preprint arXiv:2404.14740*.
- Duc Minh Nguyen et al. 2023. Culturax: A cleaned, enormous, and multilingual dataset for large language models in 167 languages. *arXiv preprint arXiv:2309.09400*.
- Kelechi Ogueji et al. 2021. Small data? no problem! exploring the viability of pretrained multilingual language models for low-resourced languages. *arXiv preprint arXiv:2011.03823*.
- Adewale Oladipo et al. 2023. Wura: A multilingual dataset of african languages. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*.
- Pedro Javier Ortiz Suárez, Laurent Romary, and Benoît Sagot. 2020. A monolingual approach to contextualized word embeddings for mid-resource languages. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1703–1714.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*.