

Mafoko: Structuring and Building Open Multilingual Terminologies for South African NLP

Vukosi Marivate^{1,2,3}, Isheanesu Dzingirai¹, Fiskani Banda¹, Richard Lastrucci¹, Thapelo Sindane¹, Keabetswe Madumo¹, Kayode Olaleye¹, Abiodun Modupe¹, Unarine Netshifhefhe¹, Herkulaas Combrink^{4,5}, Mohlatlego Nakeng¹, Matome Ledwaba¹

¹Data Science for Social Impact, Dept. of Computer Science, University of Pretoria,

²AfriDSAI, University of Pretoria, ³Lelapa AI,

⁴Economics and Management Sciences, University of the Free State,

⁵Interdisciplinary Centre for Digital Futures, University of the Free State

Correspondence: vukosi.marivate@cs.up.ac.za

Abstract

The critical lack of structured terminological data for South Africa’s official languages hampers progress in multilingual NLP, despite the existence of numerous government and academic terminology lists. These valuable assets remain fragmented and locked in non-machine-readable formats, rendering them unusable for computational research and development. *Mafoko* addresses this challenge by systematically aggregating, cleaning, and standardising these scattered resources into open, interoperable datasets. We introduce the foundational *Mafoko* dataset, released under the equitable, Africa-centered NOODL framework. To demonstrate its immediate utility, we integrate the terminology into a Retrieval-Augmented Generation (RAG) pipeline. Experiments show substantial improvements in the accuracy and domain-specific consistency of English-to-Tshivenda machine translation for large language models. *Mafoko* provides a scalable foundation for developing robust and equitable NLP technologies, ensuring South Africa’s rich linguistic diversity is represented in the digital age.

1 Introduction

The advancement of Natural Language Processing (NLP) is fundamentally tied to the availability of high-quality language resources. However, the vast majority of the world’s languages, including the 12 official languages of South Africa, remain critically under-resourced in this regard (Joshi et al., 2020). This scarcity creates a significant bottleneck for technological development and linguistic preservation. While substantial government and academic initiatives in South Africa have produced multilingual terminology lists over the years (Taljard, 2015), these valuable assets remain largely fragmented, locked in non-machine-readable formats like PDFs (sometimes available as scanned PDFs or flattened text without reusable structure), and

lack the standardised structure required for modern computational applications.

To bridge this critical gap, we introduce *Mafoko*¹ the South African curated Terminology, Lexicon, and Glossary Project. The mission of *Mafoko* is not to create new terminology from scratch, but to systematically aggregate, digitise, and standardise these scattered, publicly-funded terminological assets. By transforming them into interoperable, machine-readable formats, we unlock their potential for a new wave of linguistic and computational applications.

This paper presents the foundational work and initial release of the *Mafoko* project. Our primary contributions are threefold: First, we release the first version of the *Mafoko* dataset, a structured, multilingual terminology resource covering key domains for South African languages. Second, we release this dataset under the novel, Africa-centered Nwulite Obodo Open Data License (NOODL) to ensure equitable data governance and local benefit-sharing (Okorie and Omino, 2024). Third, we demonstrate the dataset’s immediate practical value by integrating it into a Retrieval-Augmented Generation (RAG) pipeline, which yields substantial improvements in machine translation accuracy and consistency for an English-to-Tshivenda language pair.

Ultimately, *Mafoko* provides both a practical resource and a scalable framework for fostering robust NLP and language technologies that reflect the rich linguistic diversity of South Africa.

2 Motivation

South Africa’s official languages, with the exception of English and to a lesser extent Afrikaans, remain critically under-resourced in the digital domain (Joshi et al., 2020). Despite significant investment from state institutions—including the Depart-

¹Mafoko is a Setswana (TSN) word that means words.



ment of Sports, Arts and Culture (DSAC), the Pan South African Language Board (PanSALB), and Statistics South Africa (StatsSA), in creating terminologies for crucial domains, these valuable assets are largely unusable for modern NLP. The primary barriers are both technical and legal: resources are frequently published as static, non-machine-readable documents ([Handbook](#)) and often lack the clear, permissive licensing required for computational reuse and research. This systemic inaccessibility hinders technological development and undermines efforts to achieve linguistic equity in South Africa's digital sphere. We acknowledge the work of the South African Centre for Digital Language Resources (SADILAR) which has worked to enable creation, preserve some of these resources or provide ways to keep track of them (for example through the *LwimiLinks*² project). *LwimiLinks* still faces the challenge that the data linked or available is mostly in PDF.

South African universities are also key actors in this landscape, developing linguistic resources in response to national policies that mandate the use of indigenous languages in higher education ([of Arts and Culture, 2003](#); [of Higher Education and Training, 2020](#)). Language units at institutions like the University of KwaZulu-Natal, the University of Pretoria, and North-West University have produced valuable discipline-specific glossaries and corpora. However, these academic contributions often suffer from the same fate as government resources: they remain siloed within institutional repositories, lacking the standardisation and interoperability required for broad integration into NLP and AI ecosystems. The need for interventions like the Universities South Africa Community of Practice for African Languages (COPAL) highlights this persistent fragmentation, which a systematic project like *Mafoko* is designed to address.

The core motivation for *Mafoko* is to unlock the potential of these dormant linguistic assets. By systematically digitising ([Taljard et al., 2022](#)), structuring, and releasing these resources under equitable licenses ([Okorie and Omino, 2024](#); [Rajab et al., 2025](#)) that adhere to FAIR principles ([Wilkinson et al., 2016](#)), we can directly enhance AI and NLP capabilities for South Africa's indigenous languages. Properly structured terminologies can be ingested to fine-tune large language models

²*LwimiLinks* is a place for multilingual terminology and other useful language resources. <https://lwimilinks.sadilar.org>

(LLMs), improve machine translation, and power a new generation of inclusive technologies like AI-driven spell checkers and voice assistants. This, in turn, empowers linguists, educators, and innovators to build culturally relevant, domain-specific applications, from healthcare diagnostics in isiZulu to financial literacy tools in Setswana. The transition to open, machine-readable resources is therefore a critical step towards ensuring that South Africa's languages not only survive but thrive in the digital era, fulfilling the multilingual promise of both government and higher education policies.

3 Methodology

The methodology for *Mafoko* is centered on the curation, standardisation, and dissemination of existing linguistic resources, rather than the creation of new terminology from scratch. Our approach systematically aggregates terminologies from disparate sources to enhance their accessibility and utility for linguistic research, education, and computational applications.

3.1 Source Identification

The initial phase involved identifying and collating terminological resources created and archived by South African universities, government departments, and research institutions. Universities, often as part of their language policy implementation, develop such resources, though many are in non-machine-readable formats like PDF. We engaged with the DSAC to assess their portfolio of commissioned terminology projects. Furthermore, the extensive terminology repositories maintained by *Statistics South Africa (StatsSA)*³ and other parastatal bodies were identified as primary data sources.

3.2 Domain Coverage

The scope of *Mafoko* is intentionally domain-agnostic, allowing for the inclusion of terminology lists from a wide array of fields. For instance, the DSAC lists encompass domains such as Information and Communication Technology (ICT), Mathematics, Finance, Health Sciences, and Parliamentary Procedure. The StatsSA collection provides comprehensive multilingual terminology for statistics. Similarly, the Open Educational Resource Term Bank (OERTB) project focused on developing African language terminologies for higher education across multiple disciplines ([University of](#)

³<https://www.statssa.gov.za>

Pretoria, 2019; Taljard, 2015). This broad coverage ensures the dataset’s utility across diverse research and application contexts. We include the University of Pretoria multilingual glossary, the UNISA Multilingual Robotics Glossary: South African Languages Version⁴, Multilingual Linguistic Terminology Project (SAMPLT)⁵ (Griesel and Mojapelo, 2022) and the AI Terminologies in African Languages (in African Languages Contributors, 2025).

3.3 Challenges in Data Acquisition

A primary challenge was overcoming the fragmented and often inaccessible nature of the source data. This included navigating licensing constraints, which were often unclear or restrictive, and dealing with access limitations, such as portals that only permit single-term queries. The heterogeneity of data formats, ranging from scanned PDFs to structured spreadsheets—required significant and bespoke pre-processing efforts. These hurdles are emblematic of the broader challenges in language resource development for African languages (Taljard et al., 2022). Even within a single source like DSAC, we observed inconsistencies in formatting, such as the representation of part-of-speech, across different terminology lists.

3.4 Data Curation and Structuring

The curation pipeline began with automated data extraction. Since much of the source material was in PDF format, we developed a modular extraction pipeline using Python-based tools. The pipeline required custom adaptations for each document’s unique structure, as illustrated by the formatting differences between the DSAC (Figure 1a) and StatsSA (Figure 1b) sources. For some resources, such as the StatsSA list, we were fortunate to be privately provided with a spreadsheet version, which greatly simplified the initial processing.

Automated extraction was followed by extensive manual post-processing to ensure the dataset’s quality and utility. A dedicated team member performed detailed cleaning to correct extraction errors, remove artefacts like page headers and garbled characters, and reconstruct table structures to maintain one-to-one alignment between source and target terms. To preserve the authenticity of the original resources, orthographic (use of hyphens, inconsistent capitalisation, or accent marks in lexical entries) and formatting variations from the source

⁴<https://ir.unisa.ac.za/handle/10500/30440>

⁵<https://hdl.handle.net/20.500.12185/669>

documents were retained. Where multiple translations existed for a single term, all variants were included to enable the study of lexical variation, synonymy, and regional differences. The statistics of the datasets are available in Table 1.

Each record was enriched with provenance metadata, including the originating institution, publication date, and contributor information where available. To ensure interoperability, all languages were standardised using ISO 639-3 codes. The data is currently released in CSV and JSON formats, with a TermBase eXchange (TBX) version planned for future releases. This foundational dataset (v0) is designed for iterative improvement, with future work planned to incorporate part-of-speech tags, semantic domain classification, and TEI-compliant lexicographic structuring (Burnard et al., 2014).

3.5 Data Release and Availability

The dataset is openly licensed (see Section 3.6) and accessible on multiple platforms, including GitHub⁶, Zenodo⁷ (Marivate et al., 2025), and HuggingFace collection⁸, to align with FAIR data principles (Wilkinson et al., 2016). We are working with SADILAR to contribute to their *LwimiLinks* initiative (Mafoko is already listed on there), but we also hope to have a mirror of Mafoko available via SADILAR’s digital mirrors. We plan to provide both bulk download options and API access. A feedback and validation interface is also under development to enable community-driven refinement by linguists, translators, and other stakeholders. This approach supports a virtuous cycle of continuous improvement, ensuring the resource remains relevant and accurate over time.

3.6 Licensing under NOODL

Standard open licenses, while promoting reuse, often fail to address the power asymmetries and historical contexts inherent in community-generated data. To ensure equitable governance, *Mafoko* adopts the *Nwulite Obodo Open Data License* (NOODL), an African-centered framework designed to protect local agency and mandate fair benefit-sharing (Okorie and Omino, 2024).

In contrast to generic licenses, NOODL differentiates access based on user context, mandates reinvestment from commercial use by entities out-

⁶<https://github.com/dsfsi/za-mafoko>

⁷<https://doi.org/10.5281/zenodo.1748499>

⁸<https://huggingface.co/collections/dsfsi/za-mafoko>

4
5

(a) DSAC HIV Terminology snippet
(b) StatsSA Multilingual Terminology snippet

Figure 1: Formatting differences across different terminology lists.

side developing regions, and includes provisions to reinforce community control. For *Mafoko*, this means:

1. South African and other African researchers gain open access with minimal barriers.
2. Community contributors are credited, and downstream use must return value to the originators.
3. Commercial use by external entities requires negotiated terms, correcting historical data-flow asymmetries.

NOODL enables researchers to develop and share with the common agenda, to both promote innovation as well grow the available data for under resourced languages.

4 Terminology Applications

The *Mafoko* datasets are not merely curated artifacts; they are critical assets for both computational evaluation and linguistic inquiry. Their primary applications fall into two key areas: providing a much-needed benchmark for multilingual NLP models and enabling deep analysis of language in a multilingual context.

4.1 A Benchmark for Multilingual NLP Evaluation

A significant bottleneck in developing NLP for African languages is the lack of standardized, domain-specific evaluation benchmarks (Adelani et al., 2023). *Mafoko* directly addresses this gap by providing gold-standard terminologies that can be used to rigorously assess model performance.

One primary use case is in evaluating the cross-lingual consistency of machine translation systems.

Given an English term, a model can be prompted to produce translations in various South African languages. These machine-generated outputs can then be quantitatively compared against the ground-truth terms in the dataset, revealing a model’s ability to handle domain-specific vocabulary.

Furthermore, the aligned multilingual terms enable the evaluation of multilingual word embeddings (Ruder et al., 2019; Upadhyay et al., 2016). By performing cluster analysis or measuring the cosine similarity of term-pairs (Almeida and Xexéo, 2019), researchers can probe how well semantically equivalent concepts are co-located within a shared vector space (Glavaš et al., 2019). This provides crucial insights into the quality of cross-lingual representations, particularly for low-resource languages. By serving as an evaluation resource, *Mafoko* supports the kind of participatory and community-centered benchmarking necessary for building truly useful technologies (Nekoto et al., 2020).

4.2 A Resource for Linguistic and Sociolinguistic Inquiry

Beyond computational applications, the dataset offers rich opportunities for linguistic and lexicographic research. It serves as a valuable corpus for studying the dynamics of language contact, standardisation, and change in South Africa, mirroring the kind of corpus-driven analysis that has been foundational to modern lexicography for African languages (Prinsloo and De Schryver, 2001).

Since the dataset preserves multiple translations for many terms, it facilitates the study of lexical variation (Freixa, 2022), synonymy, and dialectal preferences. Linguists can use this data to investigate term-formation strategies across languages, examining whether translations are neologisms,

Table 1: Overview of the datasets aggregated in the initial release of *Mafoko* (v0).

Source	Primary domains/Categories	Do-	Languages	Entries
DSAC (Combined)	Multiple (Finance, Health, ICT, Law, Mathematics, Arts, Science, Elections)	11	15,554	
AI Terminologies in African Languages Dataset	AI terminologies	4	85	
OERTB	Higher Education Terminology	11	5,744	
StatsSA	Official Statistics (Demography, Economics, Labour, Health, Geography)	11	1,160	
Unisa Robotics*	Robotics Terminology	11	100	
Unisa Multilingual	Provide definitions for linguistic terms in nine languages	9	778	
UP Glossary	Academic Terminology	3	1,768	
Total Entries				25,189

* Unisa Robotics glossary is re-released under CC-BY-NC-SA (not compatible with NOODL) but formatted like all other *Mafoko* releases.

calques, semantic extensions, or borrowings. Such analysis can reveal deeper cognitive or conceptual distinctions between languages.

Moreover, the data provides a unique lens for sociolinguistic inquiry into language planning and policy. It captures the outcomes of official terminology development efforts, allowing researchers to analyze the tensions between top-down standardization and organic, community-level usage. This makes the dataset an essential resource for scholars studying the politics of language and curriculum design in multilingual societies, a challenge common across the African continent (Heugh and Stroud, 2019).

5 Improving Translations with Terminology lists and RAG

To demonstrate the practical value of the *Mafoko* terminologies, we conducted experiments to assess their impact on improving machine translation quality for a low-resource language pair. Despite advances in LLMs, their performance often degrades when translating domain-specific or rare terms, especially for languages with limited high-quality

parallel corpora like South Africa’s (Zhong et al., 2024). This can lead to critical misinterpretations, such as confusing the term *register* in a mathematics context (*ridzhisitara*) versus an electoral one (*redzhistara*) in Tshivenda.

Our experiment investigates whether a Retrieval-Augmented Generation (RAG) pipeline, enriched with our curated terminology, can mitigate these issues. The overall pipeline is visualized in Figure 2.

5.1 Task and Models

We evaluated English-to-Tshivenda translation in two distinct domains: Mathematics and Election, using terminology lists from the DSAC. The datasets that were used for building the RAG terminology were English–Tshivenda Election and Mathematics (Grade R to 6) domain datasets. The election terminology dataset consisted of 576 English–Tshivenda term pairs and their associated context information that defines how a term is used. The mathematics terminology dataset consisted of 966 English–Tshivenda term pairs and their associated context information. The data extraction from PDF files, cleaning and aligning process was auto-

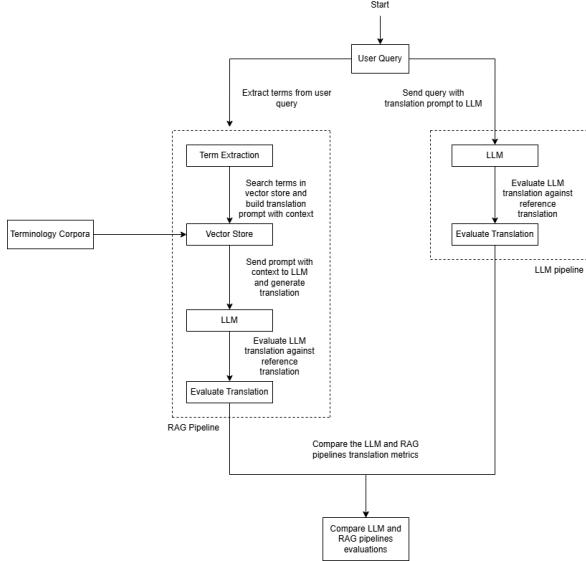


Figure 2: Overview of RAG and LLM Pipelines

mated through a Python script. We used two large language models to assess the impact of our RAG approach: the high-performance GPT-4o-mini and the open-source LLaMA3-8B.

5.2 Experimental Conditions

We tested each model under three conditions to isolate the effect of the RAG pipeline:

1. **No RAG (Baseline):** The LLM was prompted to perform direct translation without any additional context.
2. **RAG with Semantic Terms:** Key terms (nouns, verbs, adverbs) were extracted from the source text using spaCy’s `en_core_web_sm` model. These terms were used to retrieve relevant entries from the *Mafoko* vector store to augment the LLM’s prompt.
3. **RAG with Rare Terms:** Terms were selected from the source text based on their low frequency in general English corpora (Reuters (Lewis, 1997), Inaugural Speeches(Ahrens, 2021)) using the `wordfreq` library. This strategy focuses the retrieval on the most challenging, domain-specific vocabulary.

In both RAG conditions, the retrieved translations and definitions were appended to the prompt, providing in-context examples to guide the LLM.

5.3 Evaluation Metrics

We evaluated translation quality using 20 elections and 20 mathematics test set pairs using standard automatic metrics: BLEU for n-gram precision, and both chrF and chrF++ for character n-gram recall. Higher scores indicate better translation quality. The test sets were generated using: the 2025/2026 annual teaching plans for grade 4 mathematics for mathematics; the 2016 local government elections post proclamation leaflet and former President Kgalema Motlanthe’S 2009 State of the Nation Address for elections. Results are presented in Table 3 and example sentences used for evaluation are shown in Table 2.

Domain	Example Sentences (English and Tshivenda)
Election	English: Local government is in your hands! Tshivenda: Muvhuso wapo u zwanani zwavho!
Mathematics	English: The difference between numbers. Tshivenda: Phambano i re vhukati ha nomboro.

Table 2: Example English and Tshivenda Sentences by Domain

5.4 Results and Analysis

5.4.1 Quantitative Results

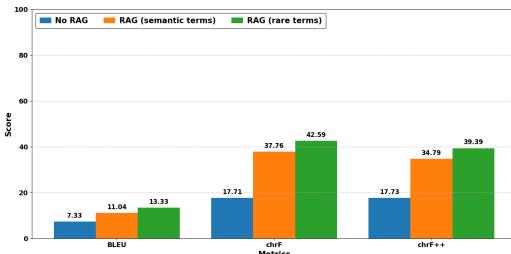
As shown in Table 3, the inclusion of a RAG pipeline with *Mafoko* terminologies leads to substantial improvements in translation quality across all metrics for both models and domains.

For GPT-4o-mini, the gains are significant. In the Mathematics domain, BLEU score improves from 7.33 to **13.33** and chrF++ score from 17.73 to **39.39** using the rare-term RAG. In the Election domain, the semantic-term RAG yields the best results, increasing the BLEU score from 5.87 to **10.41** and chrF++ from 26.99 to **36.85**.

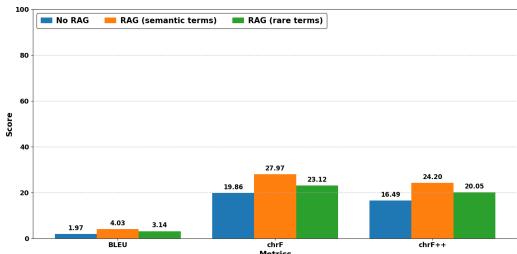
For LLaMA3-8B, while its overall performance is lower than GPT-4o-mini’s, it also benefits greatly from RAG. In both domains, the semantic-term RAG provides the best results. For the Election domain, BLEU improves from 1.97 to **4.03** and chrF++ from 16.49 to **24.20**. The performance gains are visualized in Figure 3a and 3b.

Model	Setup	Mathematics			Election		
		BLEU	chrF	chrF++	BLEU	chrF	chrF++
GPT-4o-mini	No RAG	7.33	17.71	17.73	5.87	29.21	26.99
	RAG (semantic terms)	12.44	41.32	38.95	10.41	40.35	36.85
	RAG (rare terms)	13.33	42.59	39.39	9.73	39.88	35.42
LLaMA3-8B	No RAG	2.28	12.43	10.60	1.97	19.86	16.49
	RAG (semantic terms)	4.54	22.66	20.29	4.03	27.97	24.20
	RAG (rare terms)	3.72	20.01	18.56	3.52	26.31	22.89

Table 3: Translation performance of GPT-4o-mini and LLaMA3-8B across different setups with and without RAG. Best scores per model and domain are in bold.



(a) GPT-4o-mini English to Tshivenda Mathematics



(b) LLaMA3-8B English to Tshivenda Election

Figure 3: Translation performance metrics comparison of GPT-4o-mini and LLaMA3-8B models on English to Tshivenda Mathematics and Election datasets.

5.4.2 Analysis

Although the results are low due to the use of early baseline models, the lack of fine-tuning, and a small test set size of 20 samples per domain, the results nonetheless strongly indicate that providing in-context, domain-specific terminology via RAG is a highly effective method for improving LLM translation performance for low-resource languages with the results almost doubling. The fact that this holds true for both a state-of-the-art proprietary model and a smaller open-source model underscores the robustness of this approach.

Interestingly, the optimal RAG strategy differed between the models. For LLaMA3-8B, retrieving based on semantic terms was consistently better. This suggests the model benefits from guidance on a broader range of vocabulary. For the more capable GPT-4o-mini, the rare-term strategy proved superior in the highly specialized Mathematics domain. This may indicate that the model already possesses a strong grasp of common semantic terms, and its performance is most improved by providing context for the most niche and infrequent vocabulary. The overall performance gap between the

two models likely reflects differences in their pre-training data and inherent capabilities for handling low-resource languages.

5.5 Discussion

These promising results with Tshivenda open several avenues for future work. First, this evaluation framework should be extended to the other official South African languages to confirm the generalizability of our findings. Second, it would be valuable to investigate why the rare-term RAG strategy was particularly effective for GPT-4o-mini and whether this pattern holds across other domains and models.

6 Future Directions and Call for Open Data

The expansion of *Mafoko* depends on the continued identification and integration of scattered terminological resources. As shown in Table 4, numerous valuable glossaries exist across South African institutions, but their accessibility varies dramatically, from openly licensed datasets like Unisa’s robotics glossary to web portals from Stellenbosch University that prohibit bulk download.

Table 4: Examples of Additional Terminological Resources for Integration.

Resource Name	Institution/Body	Accessibility Status
Termbank ¹	UKZN	Offline (as of 31/07/2025)
Full OERTB ²	UP	Offline (as of 31/07/2025)
LwimiLinks ³	SADILAR	Accessible (links to other org resources + PDFs)
Trilingual Wine Industry Dictionary ⁴	SA Wine Industry	Accessible (No clear license for reuse)
Multilingual Glossaries ⁵	Nelson Mandela Uni.	Accessible (PDFs)
Mechanical Engineering and Statistics and Economics Glossaries ⁶	UCT	Accessible (PDF)
Trilingual Terminology Web ⁷	Stellenbosch Uni.	Accessible (Web search, no download)
Statistical Terms Glossary ⁸	Stellenbosch Uni.	Accessible (PDF)
BAQONDE Resources (Polokelo) ⁹	Multiple South African Universities	Multiple formats (PDF, XLS) without clear licensing for reuse
CPUT Multilingual Glossaries ¹⁰	CPUT	Unreliable Access

¹ <https://ukzntermbank.ukzn.ac.za> may have been replaced by ZuluLex

² <http://oertb.tlterm.com/> ³ <https://lwimilinks.sadilar.org>

⁴ https://www.sawis.co.za/dictionary/Dictionary_Eng.pdf ⁵ <https://glossaries.mandela.ac.za> ⁶ <https://ched.uct.ac.za/multilingualism-education-project/projects/multilingual-glossaries-project>

⁷ <https://www1.sun.ac.za/languagecentre-terminologies/> ⁸ https://languagecentre.sun.ac.za/wp-content/uploads/2021/01/Stats_Eng_Afr_fin.pdf

⁹ <https://baqonde.usal.es/polokelo/> ¹⁰ <http://mlg.cput.ac.za>

A significant challenge is the ephemeral nature of digital resources. The impending offline status of both the UKZN Termbank and the Full UP OERTB, for which we fortunately have a partial backup, highlights the critical threat of digital decay and the urgent need for proactive data preservation. Our ability to access The South African Trilingual Wine Industry Dictionary is great, but it does not have a clear license for reuse.

This landscape illustrates the vital need for a structured, centralized effort like *Mafoko*. We acknowledge the institutional constraints that may lead to restrictive access, such as the need to track usage for funding reports. However, we advocate for a collective shift towards open, machine-readable formats under clear, permissive licenses. This not only aligns with Findable, Accessible, Interoperable, and Reusable (FAIR) principles but also empowers researchers, language practitioners, and developers by providing the legal and technical

clarity needed to innovate. Ensuring that South Africa's indigenous languages thrive in the digital age requires a concerted effort to make these foundational resources openly and sustainably available.

7 Conclusion

This paper introduced *Mafoko*, a project that directly confronts the critical scarcity of structured, machine-readable terminologies for South Africa's official languages. By systematically aggregating, cleaning, and standardizing fragmented resources from government and academic sources, we have created a foundational, open-access dataset. Our adoption of the Africa-centered NOODL license further ensures that these resources are used in a manner that is equitable and benefits their communities of origin.

We have demonstrated the immediate, practical value of this structured terminology through RAG

experiments, which yielded substantial improvements in English-to-Tshivenda machine translation accuracy and consistency. This result validates our core premise: that well-curated, accessible terminologies are not merely an academic exercise but are essential for enhancing the performance of language technologies for low-resource languages. Ultimately, *Mafoko* serves as both a valuable new resource and a call to action, providing a scalable foundation for developing more inclusive and capable NLP technologies that reflect the rich linguistic diversity of South Africa and the African continent.

8 Limitations

While *Mafoko* successfully structures existing terminological resources into more accessible formats, several limitations frame the scope of this work and offer avenues for future research.

Firstly, the comprehensiveness of our dataset is inherently constrained by the availability and accessibility of source materials. As noted, many valuable terminology and glossary datasets across South Africa’s language ecosystem remain difficult to incorporate. This inaccessibility stems not only from resources being unpublished or locked in scanned formats but also from digital decay, where resources like the UP OERTB become permanently offline, or are placed behind restrictive web portals that prevent bulk download. The sustainability of digital language resources in the African context is a significant challenge that affects projects like ours (Taljard et al., 2022).

Secondly, the machine translation experiments, while promising, serve primarily as a proof of concept to demonstrate utility. Our evaluation was limited to English-to-Tshivenda translation in two specific domains. A more exhaustive evaluation is needed to assess the impact of *Mafoko* across all 11 official languages and on a wider array of NLP tasks, such as named entity recognition (NER) or cross-lingual information retrieval. Future work should benchmark performance on diverse tasks and languages to fully understand the resource’s capabilities and constraints, following community-driven evaluation standards (Nekoto et al., 2020).

Finally, our adoption of the NOODL license, while principled, may present practical hurdles. As a novel, Africa-centered data governance framework, it may face adoption challenges from institutions or researchers accustomed to more globally recognized licenses like Creative Commons.

Educating potential users on its equitable benefit-sharing model is crucial but requires a dedicated effort beyond the scope of this initial project note. The complexities of data governance and licensing for low-resource languages remain a critical area for further exploration (Okorie and Omino, 2024; Rajab et al., 2025).

Acknowledgements

The authors gratefully acknowledge the support of the ABSA Chair of Data Science and the Data Science for Social Impact (DSFSI) Lab at the University of Pretoria. This work was supported by UK International Development and the International Development Research Centre (IDRC), Ottawa, Canada, under the AI4D Africa Program. DSFSI also acknowledges gifts from NVIDIA, Google.org, OpenAI, and Meta.

References

David Ifeoluwa Adelani, Marek Masiak, Israel Abebe Azime, Jesujoba Alabi, Atnafu Lambebo Tonja, Christine Mwase, Odunayo Ogundepo, Bonaventure FP Dossou, Akintunde Oladipo, Doreen Nixdorf, and 1 others. 2023. Masakanews: News topic classification for african languages. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 144–159.

Kathleen Ahrens. 2021. NLTK Inaugural Address Corpus: U.S. Presidential Inaugural Addresses (1789–2021). Originally compiled from C-SPAN sources by Kathleen Ahrens. Available via the NLTK corpora collection.

Felipe Almeida and Geraldo Xexéo. 2019. Word embeddings: A survey. *arXiv preprint arXiv:1901.09069*.

Lou Burnard, Syd Bauman, and 1 others. 2014. *TEI P5: Guidelines for Electronic Text Encoding and Interchange*. Text Encoding Initiative Consortium.

Judit Freixa. 2022. Causes of terminological variation. In *Theoretical Perspectives on Terminology*, pages 399–420. John Benjamins Publishing Company.

Goran Glavaš, Robert Litschko, Sebastian Ruder, and Ivan Vulić. 2019. How to (properly) evaluate cross-lingual word embeddings: On strong baselines, comparative analyses, and some misconceptions. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 710–721, Florence, Italy. Association for Computational Linguistics.

Marissa Griesel and Mampaka Mojapelo. 2022. *Multilingual linguistic terminology*. SADiLaR Language Resource Repository, License: Creative Commons Attribution 4.0 license.

Open Data Handbook. [Machine readable](#).

Kathleen Heugh and Christopher Stroud. 2019. *Multilingualism in South African Education: A Southern Perspective*, page 216–238. Studies in English Language. Cambridge University Press.

AI Terminologies in African Languages Contributors. 2025. Ai terminologies in african languages dataset. <https://github.com/google-research-datasets/ssa-ai-terminologies/>. Version 2025-07-14. Licensed under CC BY-SA 4.0.

Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The state and fate of linguistic diversity and inclusion in the nlp world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293.

David D. Lewis. 1997. Reuters-21578, distribution 1.0. <http://kdd.ics.uci.edu/databases/reuters21578/reuters21578.html>. ApteMod version, widely used benchmark corpus for text categorization.

Vukosi Marivate, Isheanesu Dzingirai, Fiskani Ella Banda, Lastrucci Richard, THAPELO SINDANE, Keabetswe Madumo, Kayode Olaleye, Abiodun Modupe, Unarine Leo Netshifhefhe, Herkulaas MvE Combrink, Mohlatlego Nakeng, and Matome Brilliant Ledwaba. 2025. *Mafoko companion dataset for mafoko open multilingual terminologies paper*.

Wilhelmina Nekoto, Vukosi Marivate, Tshinondiwa Matsila, Timi Fasubaa, Taiwo Fagbohungbe, Solomon Oluwole Akinola, Shamsuddeen Muhammad, Salomon Kabongo Kabenamualu, Salomey Osei, Freshia Sackey, Rubungo Andre Niyongabo, Ricky Macharm, Perez Ogayo, Orevaoghene Ahia, Musie Meressa Berhe, Mofetoluwa Adeyemi, Masabata Mokgesi-Selinga, Lawrence Okegbemi, Laura Martinus, and 28 others. 2020. *Participatory research for low-resourced machine translation: A case study in African languages*. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2144–2160, Online. Association for Computational Linguistics.

Department of Arts and Culture. 2003. [National language policy framework](#). Accessed: 2025-04-19.

Department of Higher Education and Training. 2020. Language policy framework for public higher education institutions. <https://www.dhet.gov.za/SiteAssets/Policy%20Frameworks/LanguagePolicyFramework.pdf>. Government Gazette No. 43860, 30 October 2020. Effective from 1 January 2022.

C. Okorie and M. Omino. 2024. Licensing African Datasets. [Online]. Available: <https://www.licensingafricandatasets.com> [Accessed: Jun. 30, 2025].

Daniël Jacobus Prinsloo and Gilles-Maurice De Schryver. 2001. Corpus applications for the african languages, with special reference to research, teaching, learning and software. *Southern African Linguistics and Applied Language Studies*, 19(1-2):111–131.

Jenalea Rajab, Anuoluwapo Aremu, Everlyn Asiko Chimoto, Dale Dunbar, Graham Morrissey, Fadel Thior, Luandrie Potgieter, Jessico Ojo, Atnafu Lambebo Tonja, Maushami Chetty, and 1 others. 2025. The esethu framework: Reimagining sustainable dataset governance and curation for low-resource languages. *arXiv preprint arXiv:2502.15916*.

Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2019. A survey of cross-lingual word embedding models. *Journal of Artificial Intelligence Research*, 65:569–631.

Elsabé Taljard. 2015. Collocations and grammatical patterns in a multilingual online term bank. *Lexikos*, 25:387–402.

Elsabé Taljard, Danie Prinsloo, and Michelle Goosen. 2022. Creating electronic resources for african languages through digitisation: a technical report. *Journal of the Digital Humanities Association of Southern Africa*, 4(01).

University of Pretoria. 2019. [Open educational resource term bank \(oertb\)](#). Accessed: 2025-07-25.

Shyam Upadhyay, Manaal Faruqui, Chris Dyer, and Dan Roth. 2016. Cross-lingual models of word embeddings: An empirical comparison. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1661–1670.

Mark D Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E Bourne, and 1 others. 2016. The fair guiding principles for scientific data management and stewardship. *Scientific data*, 3(1):1–9.

Tianyang Zhong, Zhenyuan Yang, Zhengliang Liu, Ruidong Zhang, Yiheng Liu, Haiyang Sun, Yi Pan, Yawei Li, Yifan Zhou, Hanqi Jiang, Junhao Chen, and Tianming Liu. 2024. *Opportunities and challenges of large language models for low-resource languages in humanities research*. *Preprint*, arXiv:2412.04497. ArXiv preprint.