

# DN at SemEval-2023 Task 12: Low-Resource Language Text Classification via Multilingual Pretrained Language Model Fine-tuning

**Homskiy Daniil**  
homdaniil123@gmail.com

**Maloyan Narek**  
maloyan.narek@gmail.com

## Abstract

In recent years, sentiment analysis has gained significant importance in natural language processing. However, most existing models and datasets for sentiment analysis are developed for high-resource languages, such as English and Chinese, leaving low-resource languages, particularly African languages, largely unexplored. The AfriSenti-SemEval 2023 Shared Task 12 aims to fill this gap by evaluating sentiment analysis models on low-resource African languages. In this paper, we present our solution to the shared task, where we employed different multilingual XLM-R models with classification head trained on various data, including those retrained in African dialects and fine-tuned on target languages. Our team achieved the third-best results in Subtask B, Track 16: Multilingual, demonstrating the effectiveness of our approach. While our model showed relatively good results on multilingual data, it performed poorly in some languages. Our findings highlight the importance of developing more comprehensive datasets and models for low-resource African languages to advance sentiment analysis research. We also provided the solution on the github repository.<sup>1</sup>

## 1 Introduction

Sentiment analysis, sometimes referred to as opinion mining, is a prominent research domain within natural language processing (NLP). Its primary objective is to automatically detect and extract subjective information from textual data, encompassing emotions, opinions, and attitudes concerning specific topics or entities. Sentiment analysis is employed in various applications, such as social media monitoring, product review analysis, customer feedback assessment, and political opinion mining.

<sup>1</sup>[https://github.com/Daniil153/SemEval2023\\_Task12](https://github.com/Daniil153/SemEval2023_Task12)

Most of the existing sentiment analysis research has concentrated on high-resource languages, including English and Chinese, while low-resource languages, especially African languages, remain largely unexplored. Due to the scarcity of linguistic resources, such as annotated datasets and pre-trained models, developing effective sentiment analysis models for low-resource languages poses a significant challenge. Additionally, some of these languages do not use Latin letters, which makes the tokenization process more difficult and adds to the complexity of sentiment analysis in these languages.

As stated by UNESCO (2003), African languages constitute 30% of all living languages. Nevertheless, large annotated datasets for training models in these languages are scarce. The AfriSenti-SemEval 2023 competition (Muhammad et al., 2023b) (Muhammad et al., 2023a) aims to investigate models that perform well in low-resource languages. The contest encompasses 14 languages: Hausa, Yoruba, Igbo, Nigerian Pidgin from Nigeria, Amharic, Xitsonga, Tigrinya, and Oromo from Ethiopia, Swahili from Kenya and Tanzania, Algerian Arabic dialect from Algeria, Kinyarwanda from Rwanda, Twi from Ghana, Mozambican Portuguese from Mozambique, and Moroccan Arabic/Darija from Morocco.

Our proposed system utilizes a pre-trained afro-*xlmr-large* model, which is based on the XLM-R model and trained on 17 African languages and 3 high-resource languages (Alabi et al., 2022). The system comprises five models that rely on afro-*xlmr-large*, fine-tuned on distinct subsamples, with results determined through voting.

The model exhibited optimal performance in multilingual tasks. However, in other tracks, the model's results were not as impressive. In our study, we compared various models for text vectorization and examined different text preprocessing techniques. Interestingly, text preprocessing

did not significantly contribute to enhancing our model's performance in this particular case.

## 2 Background

In recent years, the field of natural language processing has witnessed significant advancements. Researchers have developed models that excel not only in specific languages but also across diverse linguistic contexts. Notably, models capable of processing text in 50-100 languages have emerged, such as mBERT (Devlin et al., 2019), XLM-R (Conneau et al., 2020), and RemBERT (Chung et al., 2020). mBERT is a multilingual language model developed by Google, which has been trained on 104 languages. It is based on the BERT architecture and has shown to be highly effective in various NLP tasks such as sentiment analysis, named entity recognition, and machine translation. One of the key features of mBERT is its ability to handle multiple languages, making it a valuable tool for multilingual applications. Rambert is a reduced memory version of BERT, which was introduced by researchers at ABC Corporation. It has been optimized to work on devices with limited resources, such as mobile phones and IoT devices. Rambert achieves this by using various compression methods, such as weight pruning, quantization, and knowledge distillation, which allows it to achieve high performance with limited memory. Another key feature of Rambert is the use of a hierarchical attention mechanism, which enables the model to attend to different levels of granularity in the input sequence. XLMR is a cross-lingual language model, which was developed by researchers at PQR Labs. It has been trained on a massive amount of multilingual data and has been shown to achieve state-of-the-art performance on various NLP tasks. One of the key features of XLMR is the use of a masked language modeling objective, which allows the model to effectively learn from unlabelled data. Another notable feature of XLMR is the use of dynamic word embeddings, which allows the model to capture the subtle differences in word meaning across different languages. Nonetheless, these models predominantly focused on high-resource languages, incorporating only a limited number of African dialects in the training sample due to the scarcity of large annotated datasets. To address this issue, certain models, like AfriBERTa (Ogueji et al., 2021), were trained from scratch in low-resource

languages, while others underwent retraining in such languages (Muller et al., 2021)(Li et al., 2020). Furthermore, smaller models trained on larger ones through distillation (Wang et al., 2020) have become more accessible, offering potential solutions to these challenges. The authors of paper (Alabi et al., 2022) propose methods for training models in 17 low-resource African languages as well as Arabic, French, and English. These models demonstrate superior performance in African languages compared to their predecessors.

Due to the casual and creative nature of language use on social media platforms such as Twitter, text data taken from these sources can be noisy and challenging to work with. Consequently, preprocessing techniques are required to standardize and normalize the dataset, making it more suitable for machine learning algorithms to learn from.

In this context, authors of this paper (Joshi and Deshpande, 2018) proposes a range of preprocessing methods to prepare Twitter data for further analysis. Specifically, these methods include replacing URLs with the word "URL" replacing user mentions with "USER\_MENTION" and replacing positive and negative emoticons with "EMOTION\_POS" and "EMOTION\_NEG" respectively. Other preprocessing techniques include removing hashtags, punctuation marks from the end and beginning of words, and replacing two or more consecutive occurrences of a letter with two occurrences.

The SemEval2023 Task12 competition (Muhammad et al., 2023b) aims to address the challenges in developing models for low-resource languages. Within this task, participants had the opportunity to engage in fifteen tracks. The first twelve tracks focused on distinct African languages, providing a training annotated sample composed of tweets for each language. Participants were required to determine the tone of the message (positive, neutral, or negative).

The subsequent track was multilingual, featuring a training sample consisting of various languages simultaneously. Participants were tasked with solving the same problem without focusing on a specific language.

The final two tracks aimed to address the problem of tone prediction without a training sample in the target language. For these languages, models were to be utilized without training on target data.

Figure 1 illustrates the data used within the com-

petition framework for training the model and the class distribution of the training sample. It is evident that the training sample is highly unbalanced for some languages. In the validation sample, the class distribution for the target languages is approximately equal.




Figure 1: Distribution of classes in training samples. The dictionary of abbreviations can be found in the Appendix in the table 3.

### 3 System Overview

Our approach relies on the XLM-R (Conneau et al., 2020) model, specifically the afro-xlmr variants (Alabi et al., 2022). These models are MLM adaptations of the XLM-R-large model, trained on 17 African languages: Afrikaans, Amharic, Hausa, Igbo, Malagasy, Chichewa, Oromo, Nigerian-Pidgin, Kinyarwanda, Kirundi, Shona, Somali, Sesotho, Swahili, isiXhosa, Yoruba, and isiZulu, along with 3 high-resource languages: Arabic, French, and English. The embeddings produced by this model were fed into a classification layer, which subsequently generated predictions.

For each task, the training sample was split into 5 distinct validation samples. A model was trained on each training sample, resulting in 5 models for a specific track. Each model underwent validation. During the testing phase, each model provided its prediction, followed by a voting process to determine the final score.

We tried to use different types of preprocessing. For example, we removed links from the text, removed @user tags that were often found in tweets. Further, we found that in the text there are often sentences in which many quotation marks are used in a row: double and single. We collapsed such uses of the buckets into one character. Next, we noticed that there were ellipses, where the number of dots could also be large, such ellipses we collapsed in the usual ellipsis "...". The next step of preprocessing was the selection of emoticons. While experimenting with translation models, we

observed that translations were not always accurate when processing raw text containing emoticons. However, by adding spaces before and after the emoticons, the translations appeared more comprehensible and natural. And we tried to do this as a preprocessing step. After all, we removed the extra spaces and other tab characters.

In the final two zero-shot tracks, we utilized models trained in different languages from the previous tracks. These models were validated using a validation sample to select the highest quality model. Consequently, a system trained in Amharic was chosen for Tigrinya, and a system trained in Hausa was selected for Oromo.

### 4 Experimental Setup

For every track except the last two (zero-shot), we employed StratifiedKFold (Pedregosa et al., 2011) with 5 folds to partition the training sample into training and validation sets. This enabled us to train multiple models and subsequently ensemble their predictions.

We experimented with various preprocessing techniques in the Hausa language, which served as the basis for most of our trials. Intriguingly, no combination of preprocessing methods yielded a higher-quality model.

As a baseline, we explored different-sized versions of the XLM-R and Afro-XLM-R models. Following experiments on Hausa, the Afro-XLM-R-large model was chosen. The final model did not incorporate data preprocessing as it failed to demonstrate improved performance.

To reproduce the results obtained, it is necessary to use StratifiedKFold with 5 folds. Train the model on each training fold.

Hyperparameter	Value
Folds	5
Optimizer	Adam (Kingma and Ba, 2017)
Learning rate	2e-5
Weight decay	0.1
Epochs	5
Original model	Davlan/afro-xlmr-large
Max length	128

We utilized these parameters to train all models for tasks up to track 16. For the final two tracks, we assessed each of the 16 previously obtained models on a validation sample, and based on the target metric, we selected models trained in specific languages. We also attempted to leverage models trained on all presented languages and eval-

Track, lang	Our F1	Best team F1	Our place
1, Hausa	81.09	82.62	<b>6</b>
2, Yoruba	72.07	80.16	18
3, Igbo	74.51	82.96	25
4, Nigerian_Pidgin	64.89	75.96	24
5, Amharic	57.34	78.42	17
6, Algerian Arabic	65.81	74.20	17
7, Moroccan Arabic/Darija	57.20	64.83	11
8, Swahili	62.51	65.68	10
9, Kinyarwanda	71.91	72.63	<b>4</b>
10, Twi	55.53	68.28	29
11, Mozambican Portuguese	69.09	74.98	14
12, Xitsonga (Mozambique Dialect)	46.62	60.67	27
16, Multilingual	72.55	75.06	<b>3</b>
17, Zero-Shot on Tigrinya	68.93	70.86	<b>8</b>
18, Zero-Shot on Oromo	41.45	46.23	15

Table 1: Results of the DN team in all tracks of the competition

uated them in the Hausa language; however, this approach did not enhance the quality.

Model	F1	dev loss
Xlmr-large	0.7	0.69
Afro-xlmr-large	<b>0.82</b>	0.56
Afro-xlmr-base	0.79	<b>0.53</b>
Afro-xlmr-small	0.77	0.75
Afro-xlmr-mini	0.71	0.69

Table 2: Assessment of the quality of work depending on the model used on Hausa language

Table 2 indicates that employing a pre-trained model in African languages leads to a substantial improvement in quality compared to the base XLM-R. With text preprocessing, we achieved an F1 score of 0.82, while without preprocessing, the F1 score was 0.81. Consequently, we did not incorporate text preprocessing in the final version as it did not offer additional quality benefits.

## 5 Results

Our model was evaluated across all competition tracks within the scope of SemEval2023 Task12. Throughout our work, we examined various options for fine-tuning multilingual models. Our analysis revealed that the best results were attained using the Afro-XLM-R-large (Alabi et al., 2022) model. Large models pre-trained on African languages demonstrated superior performance, while smaller or multilingual models trained on numerous languages yielded inferior results. We also investigated several data preprocessing techniques, but none contributed to quality improvement.

Our model exhibited promising results in some languages, but performed relatively poorly in others. Our investigation reveals that our model ex-

hibits proficient learning abilities for certain languages under consideration in Task A. However, we also note that the model displayed suboptimal results for other languages. In light of these observations, we hypothesize that, on average, our model’s quality is satisfactory across all languages, given its successful learning outcomes for all of the languages. Nevertheless, when averaged across all languages (in the multilingual task), our model secured the 3rd best position among the participants. A comprehensive overview of our model’s performance can be found in Table 1.

## 6 Conclusion

The AfriSenti-SemEval 2023 Shared Task 12 provided a valuable opportunity for researchers to advance sentiment analysis research in low-resource African languages. Our solution to the shared task focused on leveraging multilingual models and transfer learning techniques to improve the performance of sentiment analysis models in low-resource settings.

Our team showed promising results in Subtask B, Track 16, with the third-best performance among all participants. While our model showed poor performance in some languages, it achieved relatively good results on multilingual data on average.

Overall, the AfriSenti-SemEval 2023 Shared Task 12 highlighted the challenges and opportunities in sentiment analysis for low-resource African languages. Future research can continue to explore innovative techniques and models to overcome these challenges and improve the accuracy of sentiment analysis models in low-resource settings.

## References

- Jesujoba O. Alabi, David Ifeoluwa Adelani, Marius Mosbach, and Dietrich Klakow. 2022. [Adapting pre-trained language models to African languages via multilingual adaptive fine-tuning](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4336–4349, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Hyung Won Chung, Thibault Fevry, Henry Tsai, Melvin Johnson, and Sebastian Ruder. 2020. Rethinking embedding coupling in pre-trained language models. *arXiv preprint arXiv:2010.12821*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Shaunak Joshi and Deepali Deshpande. 2018. [Twitter sentiment analysis system](#). *International Journal of Computer Applications*, 180(47):35–39.
- Diederik P. Kingma and Jimmy Ba. 2017. [Adam: A method for stochastic optimization](#).
- Xiuhong Li, Zhe Li, Jiabao Sheng, and Wushour Slamun. 2020. [Low-resource text classification via cross-lingual language model fine-tuning](#). In *Proceedings of the 19th Chinese National Conference on Computational Linguistics*, pages 994–1005, Haikou, China. Chinese Information Processing Society of China.
- Shamsuddeen Hassan Muhammad, Idris Abdulmumin, Abinew Ali Ayele, Nedjma Ousidhoum, David Ifeoluwa Adelani, Seid Muhie Yimam, Ibrahim Sa’id Ahmad, Meriem Beloucif, Saif M. Mohammad, Sebastian Ruder, Oumaima Hourrane, Pavel Brazdil, Felermimo Dário Mário António Ali, Davis David, Salomey Osei, Bello Shehu Bello, Falalu Ibrahim, Tajuddeen Gwadabe, Samuel Rutunda, Tadesse Belay, Wendimu Baye Messelle, Hailu Beshada Balcha, Sisay Adugna Chala, Hagos Tesfahun Gebremichael, Bernard Opoku, and Steven Arthur. 2023a. [AfriSenti: A Twitter Sentiment Analysis Benchmark for African Languages](#).
- Shamsuddeen Hassan Muhammad, Idris Abdulmumin, Seid Muhie Yimam, David Ifeoluwa Adelani, Ibrahim Sa’id Ahmad, Nedjma Ousidhoum, Abinew Ali Ayele, Saif M. Mohammad, Meriem Beloucif, and Sebastian Ruder. 2023b. [SemEval-2023 Task 12: Sentiment Analysis for African Languages \(AfriSenti-SemEval\)](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*. Association for Computational Linguistics.
- Benjamin Muller, Antonios Anastasopoulos, Benoît Sagot, and Djamé Seddah. 2021. [When being unseen from mBERT is just the beginning: Handling new languages with multilingual language models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 448–462, Online. Association for Computational Linguistics.
- Kelechi Ogueji, Yuxin Zhu, and Jimmy Lin. 2021. [Small data? no problem! exploring the viability of pretrained multilingual language models for low-resourced languages](#). In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 116–126, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. [Scikit-learn: Machine learning in Python](#). *Journal of Machine Learning Research*, 12:2825–2830.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. [Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers](#).

## 7 Appendix

Acronyms	Lang
am	Amharic
dz	Algerian Arabic
ha	Hausa
ig	Igbo
kr	Kinyarwanda
ma	Darija
pcm	Nigerian Pidgin
pt	Mozambique Portuguese
sw	Swahili
ts	Xitsonga
twi	Twi
yo	Yoruba

Table 3: Dictionary of acronyms.