

Masakhane-Afrisenti at SemEval-2023 Task 12: Sentiment Analysis using Afro-centric Language Models and Adapters for Low-resource African Languages

Israel Abebe Azime^{1,†}, Sana Sabah Al-Azzawi^{2,†}, Atnafu Lambebo Tonja^{3,†},
Iyanuoluwa Shode⁴, Jesujoba Alabi¹, Ayodele Awokoya⁵, Mardiyah Oduwole,
Tosin Adewumi², Samuel Fanijo⁷, Awosan Oyinkansola, Oreen Yousuf

[∇]Masakhane NLP, ¹Saarland University, Germany, ²Luleå University of Technology, Sweden,

³Instituto Politécnico Nacional, Mexico, ⁴Montclair State University, USA,

⁵University of Ibadan, Nigeria, ⁶Brandeis University, USA, ⁷Iowa State University, USA.

Abstract

In this paper, we describe our submission for the AfriSenti-SemEval Shared Task 12 of SemEval-2023. The task aims to perform monolingual sentiment classification (sub-task A) for 12 African languages, multilingual sentiment classification (sub-task B), and zero-shot sentiment classification (task C). For sub-task A, we conducted experiments using classical machine learning classifiers, Afro-centric language models, and language-specific models. For task B, we fine-tuned multilingual pre-trained language models that support many of the languages in the task. For task C, we made use of a parameter-efficient Adapter approach that leverages monolingual texts in the target language for effective zero-shot transfer. Our findings suggest that using pre-trained Afro-centric language models improves performance for low-resource African languages. We also ran experiments using adapters for zero-shot tasks, and the results suggest that we can obtain promising results by using adapters with limited resources.

1 Introduction

Low-resourced languages receive less attention in natural language processing (NLP) research because they lack the quality of datasets necessary for training, evaluation, and model implementation. However, the increasing abundance of social media platforms on the internet is changing the game for these languages, giving them more visibility and accessibility. This was not always the case, as a significant amount of NLP research focused on sentiment analysis and other tasks geared toward specific languages having a high online presence, thus resulting in the creation of techniques and models tailored to their needs (Yimam et al., 2020; Muhammad et al., 2022).

The rise of region-specific competitions such as the AfriSenti-SemEval competition has led to ef-

forts to curate quality datasets sourced from the internet, create new techniques to maximize the use of these datasets for NLP tasks and investigate the adequacy of existing NLP techniques to cater to the linguistic needs of low-resourced languages. The AfriSenti-SemEval Shared Task 12 provides a corpus of Twitter datasets written in 14 African languages for sentiment analysis tasks (Fröbe et al., 2023). This shared task presents a unique opportunity to advance sentiment analysis development in local languages and help bridge the digital divide in this area. The hope is that this contribution will assist future research and developments in this field.

This paper details our submission for the Afrisenti SemEval-2023 Task 12 (Muhammad et al., 2023b), which investigates the effectiveness of pre-trained Afro-centric language models and adapters for sentiment classification tasks. Our work explores the potential of these models and adapters to enhance sentiment analysis performance in the context of African languages and cultures. Our codes and all project materials are publicly available ¹.

2 Background

Sentiment analysis, a crucial task in natural language processing employs machine learning techniques to identify emotions in text, thus having numerous practical applications in areas such as public health, business, governance, psychology, and more Muhammad et al. (2022). The benefits of sentiment analysis are diverse with its effect felt on almost every aspect of human endeavor (Shode et al., 2022).

While low-resource languages have been neglected in sentiment analysis research, previous works by Yimam et al. (2020), Shode et al. (2022) and Muhammad et al. (2022) have created differ-

¹<https://github.com/SanaNGU/Sentiment-Analysis-for-African-Languages>

[†]These authors contributed equally to this work.

ent sentiment corpora for Nigerian and Ethiopian languages. With text classification used for benchmarking most pre-trained language models, such as [Ogueji et al. \(2021\)](#) and [Alabi et al. \(2022\)](#), their effectiveness on the downstream task of sentiment classification is yet to be properly explored.

The dataset introduced by [Muhammad et al. \(2023a\)](#) in SemEval 2023 Task 12, the first Afro-centric SemEval shared task, includes 14 African languages divided into three sub-tasks: language-specific, multilingual, and zero-shot competitions. As part of our contribution, our team participated in all sub-tasks of the shared task, explored the effectiveness of several pre-trained networks trained on African languages, and examined the use of parameter-efficient approaches like adapters ([Pfeiffer et al., 2020b](#)) for a zero-shot classification task.

3 System Overview

3.1 Development Phase

In the development phase of the competition, we tested several algorithms. First, we worked with the unigram count of words and Tf-idf normalized word count features. These features use very simplified word-counting strategies. Using these features, we tested Multinomial Naive Bayes, Multi-Layer Perception, and XGB classifiers.

The next phase of our experiment focuses on pre-trained language models. We worked on AfroXLMR (small, base) by [Alabi et al. \(2022\)](#), LaBSE by [Feng et al. \(2020\)](#), AfriBERTa-Large by [Ogueji et al. \(2021\)](#), AfroLM by [Dossou et al. \(2022\)](#), Bernice by [DeLucia et al. \(2022\)](#), QARiB by [Abdelali et al. \(2021\)](#), and other language and task-specific models, such as [Barbieri et al. \(2021\)](#), [DarijaBERT²](#) and [twitter-xlm-roberta-base-sentiment³](#). We further experimented with the text-to-text work done by [Jude Ogundepo et al. \(2022\)](#), and on adapters, as proposed by [Pfeiffer et al. \(2020a\)](#). We found some unusual predictions when experimenting with mT5-based ([Xue et al., 2021](#)) Afro-centric models, which was observed in previous work as well ([Adewumi et al., 2022a,b](#)).

During the development phase of both language-specific Task A and multilingual Task B, we trained the previously mentioned models. We also performed data cleaning by removing '@user' tags and eliminating punctuation and emojis. However,

we observed no significant improvement after this cleaning process, so we conducted further experiments without cleaning the dataset and used it in its original state.

For Task C, we investigated the linguistic similarities between Tigrinya, Oromo, and other languages. To aid in this exploration, we drew from the previous work of [Woldeyohannis and Mesheha \(2018\)](#), who examined the similarities between Amharic and Tigrinya. Leveraging the Amharic dataset we had, we also worked on translating it to Tigrinya using Meta's No Language Left Behind (NLLB) ([Costa-jussà et al., 2022](#)) MT model that promises to deliver high-quality translations across 200 languages.

When working with the Oromo language, we were initially uncertain about its linguistic similarities with other languages. To address this, we experimented with using a language spoken in the same region and other languages within the same family, such as Hausa. After conducting our experiments, we found that training models on Hausa data and then using these same models to predict outcomes in Oromo produced better results than with other languages.

For the zero-shot cross-lingual tasks, we used adapters ([Pfeiffer et al., 2020a](#)) while following the two-step procedure as proposed by [Pfeiffer et al. \(2020b\)](#). The procedure is as follows: (i) trained language-specific adapters using monolingual data, and (ii) trained task adapters using the task-specific dataset. For the first step, we trained language-specific adapters for Tigrinya and Oromo using a monolingual dataset and AfroXLMR-base ([Alabi et al., 2022](#)) as our base model. Since the two languages are similar, we used the Amharic training dataset for the task adapter for the Tigrinya zero-shot task, as described in the previous paragraph. Hausa and Swahili datasets were used to train the task adapters for the Oromo zero-shot task as was not a similar language to Oromo in the given training dataset.

We used language adapters for Amharic, Hausa, and Swahili from [Adelani \(2022\)](#). We trained the Amharic task adapter by using the Amharic language adapter and an Amharic sentiment dataset. After training the Amharic task adapter we replaced the Amharic language adapter with the Tigrinya language adapter and evaluated the zero-shot performance for Tigrinya. Similarly, for Oromo, we trained Swahili and Hausa task adapters using

²<https://huggingface.co/SI2M-Lab/DarijaBERT>

³<https://huggingface.co/cardiffnlp/twitter-xlm-roberta-base-sentiment>

Swahili and Hausa language adapters, respectively. After training task adapters, we replaced both language adapters with Oromo and tested their performance for the Oromo zero-shot task.

3.2 Test Phase

In the competition’s Task A testing phase, we selected the top three models from the ones we trained during the development phase. After identifying the better-performing models for each language, we submitted our results based on those models. For our fourth submission, we used a voting-based ensemble approach with the previous three models. Finally, for our last submission, we used the best-performing multilingual model trained on all available training and validation data to predict a specific language.

In multilingual Task B, we identified the four best-performing models and submitted four entries using them. Based on the models’ observed performance, we employed a voting ensemble approach with the top three models for our final submission. For Task C, we submitted five predictions based on the model’s performance during the development phase.

4 Experimental Setup

4.1 Dataset and Evaluation Metrics

During the computation and development phases, we were given only training data with labels and development data without labels. The training dataset was originally from the AfriSenti dataset, which is a corpus of 14 African languages scraped from Twitter for sentiment analysis tasks. To perform experiments repeatedly on the dataset, we created our own training and evaluation data for development, treating the provided development data as a test set. This enabled us to test a list of models we had planned to experiment with. Due to the computationally intensive nature of training, we filtered the best models based on the created data. The promising models based on this data were used in the test phase. We utilized a weighted F1 score as the evaluation metric for our models.

As shown in Figure 1, the dataset across all languages is imbalanced. This problem is more pronounced when the language datasets are joined to form multilingual data. To resolve this, during our development phase, we attempted label-based and language-based balancing, which involved performing repetitive sampling on data with low frequency.

However, these experiments did not yield any significant improvement, therefore we did not include them in our final training for submission.

4.2 Training

To conduct experiments using pre-trained BERT models, we utilized the Huggingface⁴ and PyTorch library⁵. Our implementation of the sentiment analysis models was based on code from the AfriSenti-SemEval GitHub repository⁶. We made modifications to the code to enable repeated training for all models and performed data processing specific to the provided data. Additionally, we explored the use of text-to-text models from the Afriteva GitHub repository⁷ and employed adapter-hub implementations⁸ for our adapter-related experiments.

We used existing hyper-parameters for most of our experiments, except for the pre-trained language models which we trained for 10 epochs during the testing phase. We selected the weights from the best epoch with the highest evaluation F1 score for submission.

For our final submissions, we combined as much data as possible for training. In task A, we used language-specific training and gold label data, while in task B, we utilized all available training and development datasets. Unfortunately, our experiment aimed at incorporating language-specific tags was unsuccessful, as the dataset did not provide the necessary language tags.

5 Results

5.1 Task A

Our final submission for each of the monolingual datasets was the result of LaBSE multilingual model on each language as described in 3.2. Although this model did not give us the best F1 score for all languages, it worked well for Amharic and Xitsonga - resulting in our team ranking 6th out of 29, and 8th out of 31, respectively, on the leaderboard. DziriBERT produced the best F1 score for Moroccan Arabic/Darija and Algerian Arabic, thus ranking our team 14th out of 32, and 19th out of 30 participants, respectively.

Afro-XLMR-base gave rise to our best result for Hausa and Swahili languages with our team

⁴<https://huggingface.co/>

⁵<https://pytorch.org/>

⁶<https://github.com/afrisenti-semeval/afrisent-semeval-2023>

⁷<https://github.com/castorini/afriteva>

⁸<https://github.com/adaptor-hub/adaptor-transformers>




Figure 1: Data distribution among the languages.

Submission	Language											
	amh	ary	hau	ibo	yor	twi	pcm	arq	pt-mz	swa	kin	tso
1	60.60 ¹	61.27²	77.97 ³	77.26 ⁴	73.45 ³	63.76 ³	66.11 ³	70.34 ⁶	67.63 ³	58.80 ³	65.80 ⁵	50.48 ⁷
2	51.49 ³	57.29 ⁸	77.57 ¹	77.26 ⁸	65.54 ⁸	63.56 ⁸	67.18 ⁴	70.34 ⁸	71.26 ³	57.88 ⁷	67.28 ⁴	53.00 ¹
3	53.93 ⁴	58.71 ⁹	78.53⁴	78.42¹	73.85¹	64.62¹	67.18 ¹⁰	72.86²	71.57⁸	59.95⁴	67.31³	52.00 ⁵
4 (Ensemble)	60.60	61.07	77.57	77.26	65.54	63.56	67.18	70.34	67.63	58.80	65.80	52.00
5 (LaBSE on All Data)	68.85	55.50	73.12	73.75	68.98	62.83	66.88	64.35	70.26	60.24	64.14	54.33

Table 1: Test phase comparison of F1-scores (%) for Task A sentiment analysis on 12 African languages using various pre-trained models. Models include (1) AfriBerta-large, (2) Dziribert, (3) LaBSE, (4) AfroXLMR-base, (5) AfroXLMR-small, (6) DarijaBERT, (7) AfroLM, (8) twitter-xlm-roberta-base-sentiment, (9) QARiB , and (10) Bernice.

ranking 31st out of 35 and 16th out of 30 participants. AfriBerta-Large produced the best F1 score for Igbo, Yoruba, and Twi languages resulting in our team ranking 26th out of 32, 23rd out of 33, and 21st out of 31 for each language, respectively. The ensemble approach on LaBSE, Afro-XLMR-base, and Bernice model produced the best prediction for the Nigerian Pidgin language, resulting in our team ranking 16th out of 32. Twitter-XLM-Roberta and LaBSE models predicted the best F1 scores for Mozambican Portuguese and Kinyarwanda, resulting in our team ranking 13th out of 30, and 20th out of 34, respectively. The overview of the F1 scores for each of the models we considered can be found in Table 1.

5.2 Task B

Our best model was an ensembled AfroXLMR-base, LaBSE multilingual, and twitter-xml-roberta-base-sentiment models for the multilingual sentiment classification task. This ensemble model put our team in the 9th position out of 33. The models considered and their F1 scores for this task are

described in table 2.

S.No.	Model	F1-Score (%)
1	AfroXLMR-base	67.69
2	AfroLM	62.45
3	LaBSE_multilingual	68.28
4	twitter-xlm-roberta-base-sentiment	66.07
5	Ensamble of Top 3	70.34

Table 2: Test phase comparison of F1-scores (%) for Task B

5.3 Task C

For the zero-shot classification on Tigrinya (Table 3), our last submitted model was an ensembled AfroXLMR trained on Amharic, multilingual AfroXLMR, and multilingual AfriBerta. This placed our team 18th out of 28. However, our best model was multilingual AfroXLMR which produced an F1 score of 61.48%, as opposed to our submitted ensembled model which produced an F1 score of 57.99%. For the same task on the Oromo language, our ensembled AfroXMLR trained on the multilingual dataset, AfriBerta trained on the

multilingual dataset and the adapter model produced the best F1 score during our training process. This resulted in our team ranking 10th out of 29.

S.No.	Model	F1-Score (%)
1	Adapter	32.95
2	AfroXLMR (Amharic all+ tig dev)	57.99
3	AfroXLMR (Multilingual + tig dev)	61.48
4	AfriBerta (Multilingual + tig dev)	54.74
5	Top 3 Ensemble	57.99

Table 3: Test phase comparison of F1-scores (%) for Task C-Tigrinya

S.No.	Model	F1-Score (%)
1	AfroXLMR Multilingual data + orm dev)	36.89
2	AfriBerta (Multilingual data + orm dev)	42.09
3	Adapter	34.08
4	Top 3 Ensemble	42.09

Table 4: Test phase comparison of F1-scores (%) for Task C-Oromo

6 Conclusion

In this paper, we presented our submission for the AfriSenti-SemEval Shared Task 12 of SemEval-2023, focusing on the classification of sentiment for monolingual, multilingual, and zero-shot settings for low-resource African languages. We explored Afro-centric, language-specific, and general pre-trained language models for fine-tuning. Based on the test result, for monolingual sentiment classification, Afro-centric language models showed better performance for most of the languages. However, language-specific pre-trained language models perform better than Afro-centric language models for Algerian Arabic and Moroccan Darija. For multilingual sentiment classification, Afro-centric language models show promising results. For the zero-shot task, we see that using adapters shows promising results for related languages while Afro-centric languages show better performance for unrelated languages.

We would like to further our research on adapters for low-resource languages as it shows promising results in the zero-shot setting for related languages.

References

Ahmed Abdelali, Sabit Hassan, Hamdy Mubarak, Kareem Darwish, and Younes Samih. 2021. [Pre-training bert on arabic tweets: Practical considerations.](#)

David Ifeoluwa Adelani. 2022. [Adapters for African languages – based on AfroXLMR.](#)

Tosin Adewumi, Lama Alkhaled, Hamam Mokayed, Foteini Liwicki, and Marcus Liwicki. 2022a. [ML_LTU at SemEval-2022 task 4: T5 towards identifying patronizing and condescending language.](#) In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 473–478, Seattle, United States. Association for Computational Linguistics.

Tosin Adewumi, Sana Sabah Sabry, Nosheen Abid, Foteini Liwicki, and Marcus Liwicki. 2022b. [T5 for hate speech, augmented data and ensemble.](#) *arXiv preprint arXiv:2210.05480.*

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.

Francesco Barbieri, Luis Espinosa Anke, and José Camacho-Collados. 2021. [XLM-T: A multilingual language model toolkit for twitter.](#) *CoRR*, abs/2104.12250.

Marta R Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, et al. 2022. [No language left behind: Scaling human-centered machine translation.](#) *arXiv preprint arXiv:2207.04672.*

Alexandra DeLucia, Shijie Wu, Aaron Mueller, Carlos Aguirre, Philip Resnik, and Mark Dredze. 2022. [Bernice: A multilingual pre-trained encoder for twitter.](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6191–6205.

Bonaventure F. P. Dossou, Atnafu Lambebo Tonja, Oreen Yousuf, Salomey Osei, Abigail Oppong, Iyanuoluwa Shode, Oluwabusayo Olufunke Awoyomi, and Chris Chinenye Emezue. 2022. [Afrolm: A self-active learning-based multilingual pretrained language model for 23 african languages.](#)

Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Ariavazhagan, and Wei Wang. 2020. [Language-agnostic bert sentence embedding.](#)

Maik Fröbe, Tim Gollub, Matthias Hagen, and Martin Potthast. 2023. [SemEval-2023 Task 5: Clickbait Spoiling.](#) In *17th International Workshop on Semantic Evaluation (SemEval-2023)*.

Ogunayo Jude Ogundepo, Akintunde Oladipo, Mofetoluwa Adeyemi, Kelechi Ogueji, and Jimmy Lin. 2022. [AfriTeVA: Extending ?small data? pretraining approaches to sequence-to-sequence models.](#) In

- Proceedings of the Third Workshop on Deep Learning for Low-Resource Natural Language Processing*, pages 126–135, Hybrid. Association for Computational Linguistics.
- 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.
- Shamsuddeen Hassan Muhammad, David Ifeoluwa Adelani, Sebastian Ruder, Ibrahim Sa'id Ahmad, Idris Abdulmumin, Bello Shehu Bello, Monojit Choudhury, Chris Chinenye Emezue, Saheed Salahudeen Abdullahi, Anuoluwapo Aremu, Alípio Jorge, and Pavel Brazdil. 2022. [NaijaSenti: A Nigerian Twitter sentiment corpus for multilingual sentiment analysis](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 590–602, Marseille, France. European Language Resources Association.
- 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.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020a. Adapterhub: A framework for adapting transformers. *arXiv preprint arXiv:2007.07779*.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020b. Mad-x: An adapter-based framework for multi-task cross-lingual transfer. *arXiv preprint arXiv:2005.00052*.
- Iyanuoluwa Shode, David Ifeoluwa Adelani, and Anna Feldman. 2022. [YOSM: A NEW YORUBA SENTIMENT CORPUS FOR MOVIE REVIEWS](#). In *3rd Workshop on African Natural Language Processing*.
- Michael Melese Woldeyohannis and Million Meshesha. 2018. Experimenting statistical machine translation for ethiopic semitic languages: The case of amharic-tigrigna. In *Information and Communication Technology for Development for Africa*, pages 140–149, Cham. Springer International Publishing.
- 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](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Seid Muhie Yimam, Hizkiel Mitiku Alemayehu, Abinew Ayele, and Chris Biemann. 2020. Exploring amharic sentiment analysis from social media texts: Building annotation tools and classification models. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1048–1060.