



Model Card

Lelapa-X-Sentiment (isiZulu)

Model Details	2
Intended use	2
Primary intended uses	2
Primary intended users	3
Out-of-scope use cases	3
Factors	3
Relevant factors	3
Evaluation factors	3
Metrics	4
Model performance measures	4
Evaluation data	4
Datasets	4
Motivation	5
Preprocessing	5
Training data	5
Quantitative analyses	5
Unitary results	6
Intersectional result	6
Ethical considerations	6
Caveats and Recommendations	6



Model Details

Basic information about the model: Review section 4.1 of the [model cards paper](#).

Organization	Lelapa AI
Product	Vulavula
Model date	7 November 2023
Feature	Sentiment Analysis
Lang	isiZulu
Domain	General(Social media)
Model Name	Lelapa-X-Sentiment (isiZulu)
Model version	1.0.0
Model Type	Fine-Tuned Proprietary Model

Information about training algorithms, parameters, fairness constraints or other applied approaches, and features: Proprietary Fine-tuning of a Base Model on Text Data

License: Proprietary

Contact: info@lelapa.ai

Intended use

Use cases envisioned during development: Review section 4.2 of the [model cards paper](#).

Primary intended uses

Intended use is governed by the language and domain of the model. The model is intended to be used for the isiZulu sentiment analysis task. The model is trained on the Twitter (X) dataset, and it will be used to analyze sentiment in social media.



Primary intended users

The sentiment analysis model can be used for :

- Customer Feedback Analysis
- Social Media Monitoring
- Market Research and Analysis
- Political Campaigns and Public Opinion
- Content Recommendation Systems

Out-of-scope use cases

All languages and domains outside of sentiment analysis for isiZulu.

Factors

Factors could include demographic or phenotypic groups, environmental conditions, technical attributes, or others listed in Section 4.3: Review section 4.3 of the [model cards paper](#).

Relevant factors

Groups:

- The annotators are recruited by 3rd party company, and their level of understanding of the task is one of the relevant factors. There is no record of the demographic information about the annotators.
- We acknowledge that sentiment analysis is a subjective task and, therefore, our data can still suffer from the label bias that most datasets suffer from.

Evaluation factors

- In our development setting (training and evaluation), we used the factors described above with additional synthetic arrangements to improve the robustness of the model relative to real-world factors.



Metrics

The appropriate metrics to feature in a model card depend on the model being tested. For example, classification systems in which the primary output is a class label differ significantly from systems whose primary output is a score. In all cases, the reported metrics should be determined based on the model's structure and intended use: Review section 4.4 of the [model cards paper](#).

Model performance measures

The model is evaluated using the F1-score and human evaluation: The models' performances are measured by both automatic metrics and human evaluation. As an automatic metric, the F1 score is a measure used in statistics and machine learning to evaluate the accuracy of a binary classification model. It considers both the precision and the recall of the test to compute the score. Precision is the number of true positive results divided by the number of all positive results, including those not identified correctly, and recall is the number of true positive results divided by the number of all samples that should have been identified as positive.. [Read more](#).

F1 score: Testing on sentiment test set in isiZulu

Decision thresholds

No decision thresholds have been specified

Evaluation data

All referenced datasets would ideally point to any set of documents that provide visibility into the source and composition of the dataset. Evaluation datasets should include datasets that are publicly available for third-party use. These could be existing datasets or new ones provided alongside the model card analyses to enable further benchmarking.

Review section 4.5 of the [model cards paper](#).

Datasets

- Proprietary sentiment analysis dataset for isiZulu



Annotation Process

Three native speakers annotate the dataset

Motivation

We are interested in curating a sentiment analysis dataset for isiZulu because there is no publicly available dataset.

Preprocessing

We did the following pre-processing to determine the gold label.

Three-way agreement: Similar to the majority vote approach, if all three annotators agree on a label, we consider the agreed sentiment class to be the gold standard.

Three-way disagreement: When all annotators disagree on a label, we discard the tweet.

Two-way partial disagreement: If two of the annotators agree on a label, and the third annotator has a partial disagreement. For example, if two annotators classify a tweet as POS (or NEG), and the other annotator classifies it as a non-contradicting class such as NEU, we consider the POS (or NEG) classification to be the gold standard.

Two-way disagreement: If two of the annotators agree on a label, and the third annotator has a total disagreement. For example, if two annotators identify a tweet as POS and another as NEG or vice versa, the majority vote is not the final class (in this case, POS). To resolve such subjective disagreement, independent annotators review the disagreement and assign a final label.

Training data

Review section 4.6 of the [model cards paper](#).

Refer to the datasheet provided



Quantitative analyses

Quantitative analyses should be disaggregated, that is, broken down by the chosen factors. Quantitative analyses should provide the results of evaluating the model according to the chosen metrics, providing confidence interval values when possible.

Review section 4.7 of the [model cards paper](#).

Unitary results

Models	F1 score
Lelapa-X-Sentiment	0.6180

Intersectional result

In progress

Ethical considerations

This section is intended to demonstrate the ethical considerations that went into model development, surfacing ethical challenges and solutions to stakeholders. The ethical analysis does not always lead to precise solutions, but the process of ethical contemplation is worthwhile to inform on responsible practices and next steps in future work: Review section 4.8 of the [model cards paper](#).

Tweets were anonymized by replacing all @mentions with @user and removing all URLs.

Caveats and recommendations

This section should list additional concerns that were not covered in the previous sections.



Review section 4.9 of the [model cards paper](#).

Additional caveats are outlined extensively in our Terms and Conditions.