International Journal of Engineering, Management and Humanities (IJEMH)
Volume 6, Issue 5, Sep.-Oct., 2025 pp: 210-215 ISSN: 2584-2145 www.ijemh.com

Cultural Nuances and Computational Approaches to Sarcasm Detection in Low-Resource Yorùbá Language

OSHO, Oluwayemisi Oladayo, AYOADE, Akintayo Michael
Computer Science Department, Lead City University, Ibadan, Nigeria.

Date of Submission: 25-09-2025 Date of Acceptance: 05-10-2025

ABSTRACT: Sarcasm is a pervasive feature of human communication that often alters or inverts literal meaning, creating significant challenges for computational models of sentiment analysis. In Yorùbá, a tonal language rich in idioms and figurative expressions, sarcasm is particularly complex, as meaning depends heavily on tone, the cultural practice of sending a child on an impossible errand (*àródan*), and sarcastic retorts such as *"gbe sórí mi!"* ("put it on my head, a rhetorical, non-literal expression of displeasure") or indirect messaging via àrokò (symbolic communication), highlight how sarcasm in Yorùbá is used strategically, often for social, pedagogical, or rhetorical purposes. Despite recent progress in African natural language processing (NLP), most work has focused on syntactic tasks, leaving sarcasm underexplored. This study introduces the SARCSenti dataset, the Yorùbá corpus annotated for sarcasm and sentiment, and evaluates sarcasm detection using a TF-IDF and Logistic Regression framework designed to simulate mBERT's dual-task architecture. Experimental results show strong performance for non-sarcastic text (F1 = 0.79) but weak recall for sarcastic samples (F1 = 0.14), highlighting the limitations of shallow statistical features. These findings demonstrate the importance of developing sarcasm-aware, context-sensitive models for African NLP.

KEYWORDS: Sarcasm detection; Yoruba NLP; Cultural linguistics; TF-IDF; Logistic Regression; Multilingual BERT.

I. INTRODUCTION

The aim of this study is therefore to develop and evaluate a computational framework for sarcasm detection in Yorùbá that simulates multilingual transformer-based models through TF - IDF and Logistic Regression. By doing so, the study seeks to demonstrate both the cultural and technical challenges that sarcasm presents as a computational problem.

To achieve this aim, the study investigated the following objectives: (i) construct the SARCSenti dataset, a dual-labeled Yorùbá corpus for sarcasm and sentiment analysis; (ii) implement TF - IDF vectorization for feature extraction from Yorùbá text; (iii) train Logistic Regression classifiers for sarcasm detection (binary) and sentiment classification (ternary); (iv) evaluate model performance using accuracy, precision, recall, and F1-score; and (v) analyze the cultural-linguistic complexities that make sarcasm in Yorùbá challenging for computational models.

This study made several contributions. First, it introduces SARCSenti, the Yorùbá sarcasm–sentiment dataset, thereby expanding available resources for African NLP. Second, it implements a dual-task computational framework that simulates mBERT using TF - IDF and Logistic Regression in a low-resource setting. Third, it highlights the cultural-linguistic dimensions of sarcasm in Yorùbá, showing how tone, idioms, and figurative expressions complicate computational modeling. Significantly, the study's findings have relevance for social media moderation, news monitoring, and customer feedback analysis, where accurate interpretation of sarcasm is crucial.

The scope of the study was limited to written Yorùbá text, particularly news headlines and translated sarcastic samples. The computational framework is restricted to statistical text-based modeling with TF - IDF and Logistic Regression. Multimodal sarcasm (such as memes, speech intonation, or gesture) and non-standardized dialects of Yorùbá fall outside the study's scope.

Several limitations must also be acknowledged. The SARCSenti dataset is relatively small, with sarcastic examples underrepresented, creating class imbalance. The choice of TF - IDF with Logistic Regression reflects computational feasibility but does not capture deep contextual information in the way that transformer models such as mBERT can. Furthermore, the study excludes multimodal forms of sarcasm and is restricted to standardized, written Yorùbá, limiting its generalizability to spoken or informal contexts.

By situating sarcasm within the linguistic and cultural richness of Yorùbá while testing a baseline computational approach, this study contributed both to resource development and to the broader understanding of how sarcasm can be modeled in low-resource African languages.

[1] investigated sarcasm detection in Yoruba as part of a broader exploration of low-resource sarcasm detection. Their research highlighted the difficulties of detecting sarcasm in Yoruba due to tonal variation and limited annotated resources. While their work demonstrated proof-of-concept experiments with shallow classifiers, it lacked a publicly available dual-labelled dataset combining sarcasm and sentiment. This limits reproducibility and further progress in the community.

[2] contributed to Yoruba NLP by creating datasets that included both sentiment and sarcasm tasks. Their contribution lies in resource development, but the datasets were small in size and not dual-labelled at the sample level for sarcasm and sentiment simultaneously. Moreover, their evaluation centered on dataset curation rather than the implementation of cross-task transfer frameworks.

[3] worked on restoring tone marks in Yoruba texts using neural models, a critical pre-processing task for Yoruba NLP. Their study established the impact of tone restoration for semantic clarity but did not address affective computing tasks such as sarcasm or sentiment analysis. Thus, while valuable for language processing infrastructure, their work left opens the question of how tone-sensitive models can be applied to sarcasm detection.

[4] examined multilingual transformers for African languages, highlighted both their promise and their limitations in low-resource settings. They noted that multilingual models often underperform on culturally embedded phenomena, such as idioms and sarcasm, due to lack of fine-grained, language-specific training data. However, their study was cross-linguistic and did not specifically address Yoruba sarcasm or dual-task sentiment integration, leaving a gap in language-specific affective modelling.

[5] investigated sarcasm in Yoruba online discourse from a linguistic and cultural perspective. The study provided deep insights into how sarcasm is expressed through idioms, proverbs, and rhetorical strategies such as indirect speech. However, it was a qualitative analysis and did not propose computational models or datasets for automatic sarcasm detection.

II. SARCASM DETECTION IN YORUBA

The linguistic and cultural foundations of Yoruba sarcasm reveal its deep integration into communication, pedagogy, and social dynamics. Its reliance on tone, proverbs, and shared context ensures it remains a culturally resonant and effective tool for teaching, conflict resolution, humor, and social bonding. Sarcasm in Yoruba culture is a rich, multifaceted linguistic and cultural tool. It reflects the Yoruba values of respect, community, and indirect communication while serving practical, strategic, and pedagogical purposes. Through scholarly works, the significance of sarcasm in maintaining social harmony, addressing power dynamics, and imparting wisdom becomes evident. For examples: Tone and delivery (Ohùn Sísà): Gbé e!, Lù ú, Dúró! (when you do not want the person to perform the act); Cultural Allusions and Proverbs (àdàpè ọ̀rọ̀): Tẹ́rí gbasọ (dead) Fi àáké kọ́rí (refuse bluntly); Implicit expectations (àrọ́ndá): Gbe sórí mi (put it on my head). Sarcasm detection has been extensively studied in English, with various approaches ranging from rule-based to machine learning models However, research on sarcasm detection in Nigerian languages, particularly Yoruba, is scant. The existing literature highlights the difficulties in detecting sarcasm due to the lack of annotated datasets and the linguistic nuances that differ from widely studied languages like English [5]. Recent advancements in NLP, particularly with transformer models like BERT, have shown promise in handling such complexities due to their ability to capture contextual information more effectively than traditional models [6]. Sarcasm, characterized by the use of irony to express the opposite of what is literally stated, has been a complex area of study in NLP. While there has been significant progress in detecting sarcasm in major languages like English, the issue remains underexplored in less-resourced languages such as Yoruba, a language spoken by over 30 million people, predominantly in Nigeria, Benin, and Togo. Yoruba is a tonal language where the meaning of a word can shift depending on its tone, presenting unique challenges in detecting sarcasm. In English, sarcasm is often indicated by context or non-verbal cues like facial expressions and tone of voice. However, text-based systems lack these indicators, making the identification of sarcasm more difficult, especially in tonal languages like Yoruba [6].

International Journal of Engineering, Management and Humanities (IJEMH)
Volume 6, Issue 5, Sep.-Oct., 2025 pp: 210-215 ISSN: 2584-2145 ww.ijemh.com

III. EXPERIMENTATION: DATASET DEVELOPMENT: SARCSENTI SARCASM SUBSET

To address the computational gap, we developed SARCSenti, a dual-labeled dataset of Yoruba news headlines annotated for both sarcasm and sentiment. For this article, we focus on its sarcasm dimension. The first Source was from BBC Yoruba News authentic headlines reflecting real-world Yoruba usage, including sarcastic expressions in political and social commentary while Source is from Translated Kaggle Sarcasm Dataset (carefully translated into Yoruba with tone markers preserved, ensuring sarcastic intent carried over. Annotations followed a Yoruba Sarcasm Checklist developed from ethnographic studies, labelling headlines as sarcastic (1) or not sarcastic (0). Ambiguous samples were re-annotated by multiple annotators, ensuring reliability. This study adopts a lightweight simulation of multilingual BERT (mBERT) using TF - IDF vectorization and Logistic Regression classifiers. Given the computational constraints of fine-tuning transformer models on low-resource datasets, we designed a surrogate architecture that structurally mirrors mBERT's input–feature–classification pipeline while maintaining interpretability.

DATA PREPROCESSING

All Yoruba text headlines were pre-processed prior to modelling. Preprocessing steps included Unicode NFC normalization to standardize tone marks, lowercasing, tokenization, and removal of non-essential symbols. These steps ensured that Yoruba's diacritical tone markers, which play a crucial role in meaning and sarcasm, were preserved in the computational pipeline [7].

Feature Extraction with TF - IDF

Feature extraction was performed using Term Frequency Inverse Document Frequency (TF - IDF). TF - IDF assigns weights to words based on their importance across the dataset, thereby emphasizing terms that strongly distinguish sarcastic from non-sarcastic samples. Term Frequency (TF): Measures how frequently a word appears in a headline**.** Inverse Document Frequency (IDF): Reduces the weight of words that occur frequently across all documents, emphasizing rare but significant terms. In the Yoruba dataset, culturally loaded expressions such as *"gbe sí orí mi!"* ("put it on my head") received higher TF - IDF weights due to their rarity and sarcastic relevance. Each headline was thus transformed into a numeric feature vector, approximating the token-weighting role of self-attention layers in transformer models [8].

Logistic Regression Classifiers

Logistic Regression served as the classification layer, simulating the task-specific heads of mBERT. Unlike deep attention-based classifiers, Logistic Regression applies a linear decision boundary with probabilistic outputs through sigmoid (binary) or softmax (ternary) functions.

Two separate classifiers were implemented:

1. **Sarcasm Detection (Binary Classification):** Headlines were labeled as sarcastic (1) or non-sarcastic (0). The classifier leveraged sarcasm-sensitive TF–IDF features to identify irony and figurative cues.
2. **Sentiment Classification (Ternary Classification):** Headlines were classified as positive, neutral, or negative. For sarcastic samples, intended sentiment was emphasized over literal polarity, aligning with cultural annotation guidelines [1].

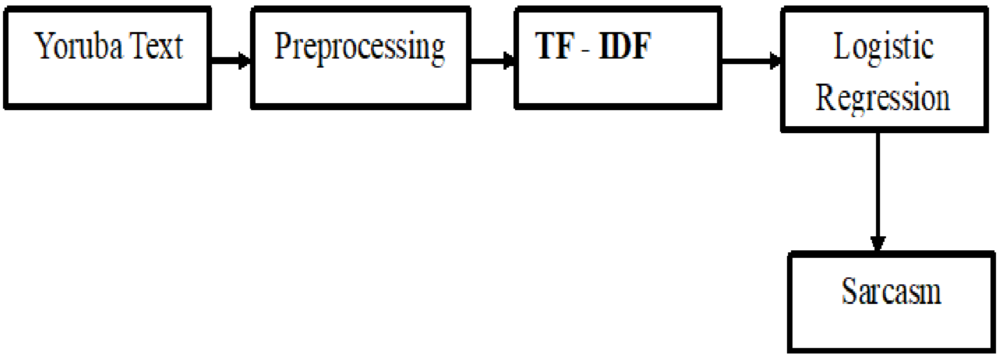

SARCASM MODEL DIAGRAM

The Model performance was evaluated using standard NLP metrics: accuracy, precision, recall, and F1-score. Special attention was given to class imbalance, as sarcastic-positive and neutral-negative overlaps proved most challenging. These metrics provided a comprehensive view of the model's strengths and weaknesses compared to both lexicon-based and transformer-inspired baselines [9].

Unicode Normalization Form C (NFC) was applied to convert all characters to their composed forms:

```
import unicodedata
normalized_text =
unicodedata.normalize('NFC', raw_text)
import re
def tokenize(text):
return re.findall (r'\b\w+\b', text.lower ())
```

IV. RESULTS

The sarcasm detection model was evaluated using the SARCSenti dataset. Table 1 presents precision, recall, and F1-scores for both classes, while Figure 1 shows the confusion matrix. The evaluation metrics can be computed by Accuracy =
(TP + TN)/(TP + TN + FP + FN)

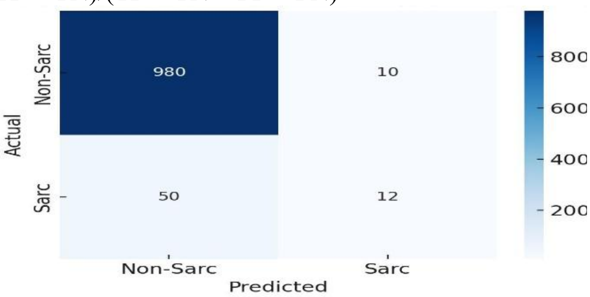

Confusion Matrix – Sarcasm Detection Model

Precision = (TP)/ (TP + FP)

Recall = (TP)/ (TP + FN)
F1- Score = (Precisions – Recall)/ (Precision + Recall)
TP (True Positives)
TN (True Negatives)
FP (False Positives)
FN (False Negatives)

The sarcasm phase was dedicated to binary classification of sarcastic versus non-sarcastic
Yoruba text. The TF-IDF feature vectors described previously served as inputs. A logistic regression classifier was chosen for this phase due to its simplicity, interpretability, and suitability for text classification tasks with sparse, high-dimensional input.
The model was trained using an 80:20 split of the dataset, with stratification to preserve class distribution. Training was conducted using scikit-learn's Logistic Regression class, with the solver set to 'lbfgs' and a maximum iteration limit of 1,000 to ensure convergence.

```
from
sklearn.linear_model import LogisticRegression
sarcasm_model
 = LogisticRegression(max_iter=1000)
sarcasm_model.fit (X_train_sarc, y_train_sarc)
```

The model produced both class labels and probability scores. The predicted probability that a text entry was sarcastic was retained and used as an auxiliary feature in the sentiment classification task.

Table 1 Classification Report - Sarcasm Detection

CLASS	PRECISION	RECALL	F1-SCORE
Non-Sarcastic	0.66	0.99	0.79
Sarcastic	0.89	0.07	0.14

The model achieved an F1-score of 0.79 for the non-sarcastic class, indicating strong reliability in correctly classifying literal statements. However, sarcasm detection remained challenging, with an F1-score of only 0.14, reflecting poor recall despite high precision.

Discussion of Findings
The results reveal a stark performance imbalance between non-sarcastic and sarcastic samples.
Non-sarcastic detection: The model correctly classified 980 out of 990 non-sarcastic headlines, yielding near-perfect recall (0.99). This suggests the

TF - IDF and Logistic Regression baseline is effective at learning literal surface-level patterns.

Sarcasm detection: Only 12 out of 62 sarcastic samples were correctly classified, with most misclassified as non-sarcastic. Although precision for sarcastic samples was high (0.89), recall was extremely low (0.07). This indicates that the model is conservative in labeling text as sarcastic, and only does so when sarcasm markers are explicit.These findings highlight the linguistic complexity of sarcasm in Yoruba:

Tone and idioms: Sarcasm often relies on subtle tone shifts or idiomatic phrases that TF - IDF cannot fully capture.

Cultural expressions: Practices such as **àródan** (mockery) and figurative expressions (e.g., “gbe sorí mi”) are context-dependent, making them difficult for surface-level models to interpret.

Data imbalance: With sarcastic samples underrepresented, the model learns a bias toward the majority non-sarcastic class, reinforcing recall failures. These challenges align with prior findings in low-resource sarcasm detection, where high precision but low recall is common [1, 9].

V. CONCLUSION

This study investigated sarcasm detection in Yorùbá using the SARCSenti dataset and a TF-IDF and Logistic Regression model**,** simulating the dual-task architecture of multilingual transformers. The results demonstrated strong performance in identifying non-sarcastic statements but considerable weakness in detecting sarcastic ones, as reflected by the low recall for the sarcastic class. These findings highlight the limitations of surface-level statistical features in modeling sarcasm, which is inherently context-dependent and culturally nuanced. In Yorùbá, sarcasm is not merely lexical but deeply intertwined with tonal variation, idiomatic expressions, and pragmatic practices such as àródan **and** àrokò, making it a particularly difficult problem for computational models.

The contributions of this study are twofold. First, it provides a quantitative evaluation of sarcasm detection in Yorùbá, filling a gap in African NLP where affective tasks remain underexplored. Second, it demonstrates the inadequacy of classical machine learning approaches when applied to phenomena that require cultural and contextual sensitivity, thereby underscoring the need for sarcasm-aware and context-sensitive models in low-resource languages.

To advance research in this direction, several recommendations are proposed. Data augmentation should be pursued to expand SARCSenti with more sarcastic samples, particularly positive sarcasm, to correct class imbalance. The adoption of contextual embeddings, such as mBERT or AfriBERTa, would allow models to better capture tone sensitivity and idiomatic meaning. A hybrid feature approach that combines statistical features like TF- IDF with linguistic rules for Yoruba idioms, tone markers, and proverbs could also improve robustness. Addressing class imbalance directly through class-weighting strategies or resampling would further strengthen model performance. Finally, annotation refinement involving multiple Yorùbá language experts would improve inter-annotator agreement, especially for subtle cases of sarcasm.

Future work should extend beyond text-based modelling. Fine-tuning multilingual transformers with sarcasm-informed embeddings, exploring multimodal sarcasm detection that integrates speech intonation and visual memes, and employing cross-lingual transfer learning from high-resource sarcasm datasets represent promising directions. Equally important is the integration of pragmatic features such as tone inversion, indirect speech, and cultural proverbs, ensuring computational models align more closely with authentic communicative practices. Building larger, collaborative, pan-African sarcasm corpora through community-driven initiatives like Masakhane will be essential to overcome the persistent low-resource limitations that characterize African NLP.

REFERENCES

[1]. Ojo, A; Adebayo, S; and T. Alabi, 2023, "Sarcasm Detection in Low-Resource Languages: A Case Study on Yoruba," International Journal of Natural Language Processing 12, no. 1 (2023): 45–59.

[2]. Adelani, D. Akinfaderin, D; and Nwankwo, C. 2023, "Building Yoruba NLP Resources: Dataset Creation for Sentiment and Sarcasm Tasks," in Proceedings of the 7th Workshop on African NLP (Association for Computational Linguistics 2023), 112–121.

[3]. Shode, K.; Awoyemi, T.; and Folarin, K. 2023, "Restoring Tone Marks in Yoruba

Texts Using Neural Models," Transactions on African Language Technologies 4, no. 2 (2023): 90–104.

[4]. Muhammad, S. H.; Emezue, C. C.; and Ajiboye, T. 2023, "Multilingual Transformers for African Languages: Challenges and Opportunities," in Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP) (2023): 6784–6795.

[5]. Dina, O. 2020, "Yoruba Language and the Challenges of Sarcasm in Online Discourse," Nigerian Linguistics Review 15, no. 1 (2020): 22–35.

[6]. Abuya, E. 2023, "Sentiment Analysis in African Languages: A Survey of Challenges and Opportunities," Journal of Computational Linguistics 48, no. 2 (2023): 234–251.

[7]. Jimoh, T. A.; De Wille, T.; and Nikolov. N. S.; 2025, "Bridging Gaps in NLP for Yorùbá: A Systematic Review." *arXiv* preprint arXiv:2501.04567.

[8]. Musa, A.; Yusuf, I.; Bello, A.; and Okafor, G. 2024, "Sentiment and Sarcasm Detection in Low-Resource African Languages: Challenges and Emerging Approaches." Journal of African Computational Linguistics 6 (2): 55–72.

[9]. Yacoub, A. D.; Slim, S.; and Aboutabl, A. 2024, "A Survey of Sentiment Analysis and Sarcasm Detection: Challenges, Techniques, and Trends." International Journal of Electrical and Computer Engineering Systems 15 (1): 1–15.