Analyzing Public Opinion and Sentiments on Nigerian Government Policy on Fuel Subsidy Removal

Adedotun O. Ogundare[1][0009-0004-1896-3187] and Stevan Gostojić[2][0000-0002-4458-3565]

[1] University of Novi Sad, Dr. Zorana Djindjica 1, Novi Sad, Serbia
[2] University of Novi Sad, Dr. Zorana Djindjica 1, Novi Sad, Serbia
adedotun@uns.ac.rs

Abstract. In recent time, social media has become a strong tool in the Nigerian political landscape, fostering interaction on governance and promoting active involvement of citizens in policymaking. The most disruptive policy in Nigeria since the Covid 19 era is the decision of the government to remove fuel subsidy in May 2023. Like many other policies of the government, this has been brought under public scrutiny and it has generated a lot of reactions on social media, particularly "X" Many applaud the decision while many argue it is an ill-conceived policy that will bring hardship to the masses. In this study, a substantial amount of data on Nigerians' opinion on fuel subsidy removal is collected from X as a univariate dataset. In the first phase of the study, we classify the data using VADER and got an accuracy of 82.22% after evaluating the performance. In the second phase, we utilize the dataset to train a BERT model for binary classification, with the goal of enabling our model to classify future opinions on the subject as either positive or negative sentiment. With an average accuracy of 84.98% across different variants of BERT, we demonstrate the effectiveness of transformer-based model in identifying sentiments expressed by Nigerians regarding the fuel subsidy removal policy. Thus, the study contributes to Nigerian natural language processing studies.

Keywords: sentiment analysis, VADER, Nigeria fuel subsidy removal, sentiment analysis, BERT.

1 Introduction

1.1 Nigeria

Nigeria is one of the largest oil producers in Africa, and it is also the most populated in the continent with an estimation of 230 million people [1]. Most of its economic activities depend on petroleum products which powers different sectors of the economy, including transportation of agricultural and other consumable products from one part of the country to another. As such, fuel subsidy is connected to livelihood for the poor masses and a hike in petroleum price could easily trigger inflation in price of food items in the market. Making living cost affordable to the masses has been the justification for

the long history of subsidizing fuel prices in Nigeria. Fuel subsidy was intended to cushion the economic hardship on Nigerians by keeping fuel prices artificially low.

However, the fuel subsidy has been a significant burden on the economy of Nigeria as it costs the government billions of dollars annually, limiting investments in other critical sectors such as education, health and infrastructure [2]. The fuel subsidy was also plagued by inefficiencies with several claims of diversion of subsidized fuel to neighboring countries where prices are higher. Fuel subsidy became a controversial subject in Nigeria, and there came a time when the masses, who once resisted its removal, began to advocate for it. The argument for the removal of fuel subsidy began as far back as 2005 and it gained momentum as the year passes. Many economic experts argued that removal of the subsidy would attract investment in the downstream oil sector and benefit the nation in the long run. Nigerians had a general election in February 2023 and all the presidential candidates of the major political had the removal of fuel subsidy on top of their agenda. The administration that governed the nation before the February 2023 election produced a new president made several attempts to remove the fuel subsidy. However, the immediate strains caused by Covid 19 pandemic made it inexpedient to remove subsidy as the masses continues to grapple with economic hardship. In May 2023, the newly inaugurated president announced the removal of the fuel subsidy as part of broader economic reforms to stabilize and diversify the economy. Consequently, the phrase 'subsidy is gone' was amplified by both traditional and new media, becoming widely used in Nigerian discourse

1.2 Sentiment Analysis

Sentiment analysis is a technique of Natural Language Processing (NLP) that is used to estimate an opinion towards an entity by classifying them as positive, negative, or sometimes neutral. For a data driven world of today, sentiment analysis is extremely valuable in providing wider opinion on topics. Past studies have utilized it to extract sentiments from tweets while analyzing the emotions in it. Tweets are a great resource for a high volume of sentiment data. “X” formerly twitter, is typically a social network that inspires people to share thoughts on a variety of topics as tweets. “X” has always attracted people willing to share their thoughts on ongoing topic, brand and business. Therefore, tweets can be used by governments, business or researchers to evaluate public opinions or emotions about a specific issue. In the Nigerian context, social media has recently become a powerful tool for public relations management and “X” has been one of the most popular platforms used by Nigerians. Companies and governments are doing a lot to manage perceptions about their image and brands on “X” and other relevant social media platforms. Such is the potential that social media now has [3].

VADER. Valence Aware Dictionary Sentiment Reasoning (VADER) is a model of text sentiment analysis which is sensitive to both polarity in terms of positivity and negativity of texts, and intensity in terms of the strength of the emotions in the text. VADER differs from the conventional approach of classification algorithms because it can be applied directly to an unlabeled text data such as the datasets extracted from “X” for this research. VADER sentiment analysis uses sentiment scores which is derived by

mapping lexical features of texts to emotional intensities in a dictionary. Sentiment score is obtained by the summation of the intensity of each word in the texts [4].

The VADER sentiment uses lexicons that connects lexical features to intensities of emotion and subsequently allot a sentiment score. The sentiment score is computed by calculating the strength of each word. VADER has been effectively used for social media postings, film reviews, and customer reviews [4]. It is capable of analyzing the meaning of texts that comprises of slangs, expressions, emphasis, punctuation and many more. In this study, the NLTK's VADER analyzer is used to analyze tweets and classify them based on the sentiments each tweet conveys. The tweets are specifically tailored to commentaries and remarks of Nigerian citizens on the fuel subsidy removal and every other subject around it. This research work will be useful for government officials in determining public opinion on the specific policy under context. It is also a useful research reference for the performance of VADER in analyzing sentiments in this specific domain of application.

1.3 Classification

Classification is one of the fundamental tasks in machine learning and deep learning. It involves predicting categorical labels for given inputs. Classification using deep learning has numerous applications across diverse fields. In natural language processing (NLP), sentiment analysis is one of the prominent use cases for text classification. Other use cases include, spam detection, and topic categorization. While classification can be done with traditional machine learning techniques, the adaptability and robustness of deep learning models make them suitable for tackling complex classification problems in various domains [5].

2 Related Works

2.1 VADER

We are in an age of big data where almost everything relies on data. Businesses rely on data to make informed decision as much as government relies on it to make informed policies. Sentiment analysis is useful for tracking and interpreting the thoughts people express on social media. This has become an important consideration in the standard measures for analyzing the effectiveness of marketing on social media. It is a crucial element for business especially in terms of product and brand awareness, promotion and customer loyalty. As noted by [6], data gathering, pre-processing of tweets, extraction and classification, sentiment recognition, and classification are the key steps for performing sentiment analysis. In many literatures, VADER has been used to perform sentiment classification on X DataStream.

[7] explained Parsimonious Rule based Model for social media context in their work. In addition to applying VADER on tweets, they were able to present a comparative study using algorithms like SentiWordNet, and Support Vector Machine (SVM). The results show that VADER outperformed all other algorithms they used in their analysis. However, they were unable to perform aspect-based analysis with VADER.

[8] conducted their study on Verb Based Sentiment Research where they suggested a model for verb-based sentiment that provides a structure for developing sentiment analysis on movie reviews. They made use of some tools in addition to the NLP library. They identified POS tags and used text blob to detect the polarity and sentiment.

Lexicon. Lexicon is defined in the dictionary as the complete set of meaningful units in a language. Polarity lexicons are lexicons that have word listings with initial level of polarities. Polarity lexicon is very important for texts sentiment analysis that relies on computational techniques. ** identifies three ways of building polarity lexicons in his work as interpreting existing lexicons from different languages, extracting polarity lexicons from corpora, and annotating sentiments. Aside VADER, there are several English polarity lexicons available online such as SentiWordNet [9].

English Language. There are numerous works on sentiment analysis performed in English language. VADER, being one of them, combines important lexical features obtained from five generalized rules that embody grammatical and syntactical conventions of human speech. It also retains the advantages of conventional sentiment lexicons such as LIWC.

2.2 Transformer Based Architecture

Transformers have become the backbone of modern NLP models. Unlike RNNs (Recurrent Neural Networks), they use self-attention mechanisms to process entire sequences in parallel. This allows for better handling of long-range dependencies in text [10].

BERT-Based Language Models. Bidirectional Encoder Representations from Transformers (BERT) utilizes the transformer architecture to pre-train deep bidirectional representations, allowing it to understand context from both directions in a sentence. It is useful in understanding and processing text-based data. Language model have also been widely applied in sentiment analysis tasks. Several variations of the BERT model have emerged over the years each trying to either optimize or improve performance of the original BERT model.

The study conducted by [11] demonstrate sentiment analysis on product reviews from an e-commerce site. The study compared the success rate between language models and machine learning models. He found Naïve Bayes to achieve the best performance among the machine learning models with an accuracy of 89.95%, while ELECTRA achieved a result of 92.54%. In a similar research endeavor, [12] evaluated BERT-based models and machine learning models on Turkish data collected through Net Promoter Score (NPS) surveys of customers using banking service. They also found BERT model to achieve the highest score with 91% F1-score. There are several other studies that have compared language models with machine learning models, each arriving at the same conclusion. Our study does not aim to compare performance between machine learning models and language models since it is already widely established in literature that language models are superior to classical machine learning

models. We have trained our models using the data obtained from 'X' after classifying its sentiment with VADER. Consequently, we tested performance on several hyperparameters using different variants of BERT.

DistilBERT. Given the revolutionary effect of BERT in the Natural Language Processing field, DistilBERT emerged as a compressed version of BERT, designed to achieve faster and more efficient language model. It removes some layers of BERT model and make the remaining layers smaller. This helps to achieve a reduction of up to 40% in the size of the model thereby boosting performance and speed. This model has proven to be most adequate for large dataset [13].

RoBERta. Another improved version of BERT language model is the Robustly Optimized BERT Pretraining Approach. It is designed to optimize the training phase of BERT thereby leading to a better language model. Again, it is more ideal for a large dataset and it has achieved good results in several NLP tasks such as text classification and sentiment analysis [14][15][16].

ALBERT. This is a lighter and faster version of BERT which helps to overcome the high costs associated with training a BERT model. It uses a more efficient training method to improve the sharing of parameters and learning efficiency in the BERT model. ALBERT has been applied widely in several applications including sentiment analysis as reported in the works of [16][17]

3 Data

3.1 Data Acquisition

Data needs to be collected in order to perform sentiment analysis. The data used for this analysis are collected from "X". The data is then passed through a series of preprocessing stages to make it more useful for the intended analysis. The first step involves collecting data from twitter on the specific subject relating to the government policy on fuel subsidy. This was done by retrieving X data using Tweepy API. The search query was set from May 29 2023 until August 29th 2023 in order to collect tweets within 3 months of the subsidy removal. Botometer library was used to filter out tweets from accounts with high bot scores to ensure the tweets we analyze are indeed expressed by people. Tweets are further filtered by Nigerian geographic coordinates. At the end of the acquisition process.

3.2 Data Preprocessing

As an essential preprocessing step, a number of text cleaning processes were carried out on the data to prepare it for sentiment analysis classification task. This entails removing special characters, emojis, symbols, and punctuations. URLs are also stripped from the dataset because they do not contribute to sentiment analysis. To enhance the

quality of the data, the "@" symbols which are typical of "X" data were also removed. All stop words were removed from the dataset because they are generic and ubiquitous terms that do not carry any significant meaning. This helps to remove noise and improve the quality of the analysis.

All the texts are converted to lowercase to minimize variations. Stemming and lemmatization were employed to ensure words are in their root forms. This is particularly useful for sentiment analysis as it reduces redundancy.

3.3 Labelling Data

Since the structure of the dataset is univariate, we need to add a label column to it. The Vader lexicon is used alongside the sentiment intensity analyzer to label all the records in the dataset. The object of Sentiment Intensity Analyzer is created to determine the polarity score of each tweet in the dataset. The polarity scores are then appended as a new column to the dataset alongside a new label that categorizes the tweets based on the polarity scores. In this sense, all data with polarity scores less than 0 are classified as negative while those above 0 are classified as positive. The result of our VADER sentiment analysis is presented in the section on results.

4 Results

4.1 Experiment 1: labelling dataset

A subset consisting of 150 records was extracted from the original dataset. After applying VADER sentiment analysis, 92 of the records were classified as negative while 58 were positive. In view of the imbalanced dataset, more subsets were randomly extracted and classified with the goal of achieving a balanced dataset. After shuffling and combining some of the subsets, we achieved a total of 270 tweets. After running a VADER analysis on this set, 140 returned negative sentiment while 130 returned positive sentiment. We validated the VADER sentiment classification by manually reading all 270 tweets. We found that 48 were classified incorrectly. This gives an accuracy score of 82.22% and that becomes our performance measure for the VADER classifier. We corrected the placement of the 19 records that were incorrectly classified. We then proceeded to the second phase of the experiment to train a model that could predict the sentiment in a future tweet.

4.2 Experiment 2: training and validating model

We utilized the BERT-base-uncased model from Hugging Face's transformers library, fine-tuned for binary classification. Using the pre-trained BERT weights at the embedding layer along with a dropout layer to prevent overfitting, we trained our model with the Adam optimizer using a learning rate of 2e-5 and changing hyperparameters until we record the best possible result. The dataset was split into 80% training set, 10% validation set and 10% test set. We evaluated our results and achieved an accuracy of 84.60% and an F1 of 85.21%. We tested for other variants of BERT (ALBERT,

DistillBERT and RoBERTa) and found the best result with ALBERT with an accuracy of 85.40% and F1 of 84.98%. We decided to evaluate our model by testing it with a variety of Language Model Models (LLMs) fine-tuned for sentiment analysis. We directly utilized the sentiment-analysis classifier provided by the Hugging Face transformer for each variant of BERT. Upon comparison, we found the results of the model we trained to be better than the off-the-shelf sentiment classifiers. This shows that a model may achieve better performance on the dataset upon which it is trained, as it captures the nuances and unique characteristics of the specific data unlike the case of a pre-trained model.

Table 1. Results obtained with BERT-based language models

Language Models	Accuracy		F1-score	
	Pre-trained	**Trained**	**Pre-trained**	**Trained**
BERT	83.41	**84.60%**	82.99	**85.21%**
ALBERT	85.21%	**85.40%**	84.98%	**84.98%**
RoBERTa	83.20%	84.00%	83.73%	87.01%
DistilBERT	85.10%	85.93%	84.49%	84.50%

5 Conclusion

5.1 Summary of Findings

This study has two primary aims: to analyze the opinions of Nigerians on fuel subsidy removal and to train a model capable of classifying the sentiments of future opinions on the subject, considering the unique nuances of how Nigerians discuss the subject. We observed, from the dataset, that majority of Nigerians expressed negative sentiments regarding the removal of the fuel subsidy. However, to balance the data used for training, we extracted additional dataset to include more positive sentiments, with the goal of training the model with a balanced dataset. After fine-tuning different hyperparameters on various BERT variants, we achieved an accuracy of 88.00%. This model can then be used to predict future opinions on the subject. Future work will focus on training a model with a larger dataset.

5.2 Discussion

While our study provides insights into the sentiments expressed towards the public policy, it is important to acknowledge its limitation in not considering segmentation by

ethnic groups. Nigeria is a highly diverse country with over 200 distinct extinct group and opinions are often held along ethnic and religious affiliations. Therefore, not considering demographic segmentation may have implications for the generalizability and contextualization of our findings within broader societal dynamics. Moving forward, future research could explore the incorporation of demographic information and segmentation analysis to better contextualize and understand the nuances of public opinion towards the policy.

Acknowledgment

This research has been supported by the Ministry of Science, Technological Development and Innovation (Contract No. 451-03-65/2024-03/200156) and the Faculty of Technical Sciences, University of Novi Sad through project "Scientific and Artistic Research Work of Researchers in Teaching and Associate Positions at the Faculty of Technical Sciences, University of Novi Sad" (No. 01-3394/1).

References

1. Pillah, T. (2023). Population and Housing Census in Nigeria. *Journal of Public Administration, Policy and Governance Research*, 9-21.
2. Nwachukwu, M. U., & Chike, H. (2011). Energy. *Fuel Subsidy in Nigeria: Fact or Fallacy*.
3. Burstein, P. (2003). The Impact of Public Opinion on Public Policy: A Review and an Agenda. *Political Research Quarterly*, 29-40.
4. Elbagir, S., & Yang, J. (2019). Twitter Sentiment Analysis Using Natural Language Toolkit and VADER Sentiment. *Proceedings of the International MultiConference of Engineers and Computer Scientists.* Hong Kong: IMECS.
5. Soofi, A. A., & Arshad, A. (2017). Classification Techniques in Machine Learning: Applications and Issues. *Journal of Basic & Applied Sciences*, 459-465. Retrieved from https://pdfs.semanticscholar.org/2678/e213cec548d278879ceaf01582ee8913cc3f.pdf
6. Shihab, E., & Jing , Y. (2020). Sentiment Analysis on Twitter with Python's Natural Language Toolkit and VADER Sentiment Analyzer. *IAENG Transactions on Engineering Sciences*, 63-80.
7. Huto, C., & Gilbert, E. (2019). A Parsimonious Rule-Based Model for Sentiment analysis of Social Media. *Proceedings of the eighth International AAAI Conference on Weblogs and social media.*
8. Sharada, A., & Krishna, P. P. (2019). VErb Based Sentiment Research. *International Journal of Recent Technology and Engineering (IJRTE)*.
9. Maite, T., Kulian , B., Milan , T., Kimberly, V., & Manfred, S. (2011). Lexicon-Based Methods for Sentiment Analysis. *Computational Linguistics*, 267-307. Retrieved from https://direct.mit.edu/coli/article/37/2/267/2105/Lexicon-Based-Methods-for-Sentiment-Analysis
10. Haber, J., & Poesio, M. (2024). Polysemy- Evidence from Linguistics, Behavioral Science, and Contextualized Language Models. *Computational Linguistics* , 351-417.
11. Engin, D., & Metin, B. (2023). Sentiment Analysis from Turkish News Texts with BERT-Based Language Models and Machine Learning Algorithms. *8th Internationmal Conference on Computer Science and Engineering.*

12. Jacob , D., Ming-Wei, C., Kenton, L., & Kristina, T. (2018). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *Computation and Language.* Cornell University
13. Sanh, V., Debut, L., Cahumond, J., & Wolf, T. (2020). DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *Computation and Language*.
14. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., . . . Stoyanov, V. (2019). RoBERTa: A Robustly Optimized BERT Pretraining Approach. *Computation and Language*.
15. Warstadt, A., Zhang, Y., Li, H.-S., Liu, H., & Bowman, S. R. (2020). Learning Which Features Matter: RoBERTa Acquires a Preference for Linguistic Generalizations (Eventually). *Computation and Language*.
16. A. F., A., M., H., & W., C. (2020). Comparative Analyses of Bert, Roberta, Distilbert and Xlnet for Text-Based Emotion Recognition. *17th International Computer Conference on Wavelet Active Media Technology and Information Processing*, 117-121.
17. Cheng-Han, C., Sung-Feng, H., & Hung-yi, L. (2020). Pretrained Language Model Embryology: The Birth of ALBERT. *Computation and Language*
18. Tripathy, J., Chakkaravarthy, S. S., Satapathy, S., Sahoo, M., & Vaidehi, V. (2022). ALBERT-based fine-tuning model for Cyberbullying Analysis. *Multimedia System*