BERT-Based Sentiment Analysis of Indian COVID-19 Tweets for Policy Making

MITHILESH VISHWAKARMA[1], VIDHI SHUKLA[2], PRANAV KHOT[3]

[1] Assistant Professor, Department of IT, Thakur College of Science and Commerce, Thakur Village, Kandivali (East), Mumbai, Maharashtra, India

[2, 3] PG Student, Department of IT, Thakur College of Science and Commerce, Thakur Village, Kandivali (East), Mumbai, Maharashtra, India

Abstract- In the unprecedented global crisis of COVID-19, social media platforms have emerged as vital communication channels, reflecting public sentiment and opinion. This paper presents an in-depth analysis of Indian Twitter data during the COVID-19 pandemic, employing the Bidirectional Encoder Representations from Transformers (BERT) model for sentiment analysis. The primary objective is to explore the public sentiment in India regarding the pandemic and its implications for policy-making. The study systematically collected a significant corpus of tweets related to COVID-19 in India, ensuring a diverse representation of opinions and perspectives. Following rigorous data pre-processing, including cleaning and tokenization, the BERT model was fine-tuned to suit the linguistic nuances and contextual intricacies of the tweets. The sentiment analysis focused on categorizing tweets into various emotional responses, such as fear, anger, optimism, and support, providing a nuanced understanding of the public's perception. The results yielded by the BERT-based analysis offer insightful and nuanced categorizations of sentiments, which are critical in informing and guiding policy decisions. These findings not only underscore the potential of advanced NLP techniques in public sentiment analysis but also highlight their practical implications in real-time policy-making, especially during crisis situations. The study contributes to the growing body of literature on sentiment analysis using deep learning and offers a novel perspective on leveraging social media data for governmental decision-making in health emergencies.

Indexed Terms— BERT, COVID-19, Data Analytics, India, Machine Learning, Natural Language Processing, Policy Making, Public Sentiment, Sentiment Analysis, Transformer Model, Twitter Data.

I. INTRODUCTION

The outbreak of the COVID-19 pandemic has been a defining global health crisis of our time and the greatest challenge humanity has faced since the Second World War. In India, the pandemic has not only posed severe health risks but also brought about significant social and economic disruptions. Amidst this turmoil, social media platforms, particularly Twitter, have emerged as vital communication tools, reflecting public opinion and sentiment on a wide range of issues related to the pandemic. Analyzing these platforms, which are among the largest sources of unstructured data, provides crucial insights into public perception and response to various measures, thus assisting in more informed decision-making [1]. Sentiment analysis, a sub-field of Natural Language Processing (NLP), is a technique used to determine the emotional tone behind a body of text. This computational method has become increasingly significant in analyzing large volumes of data generated on social media. By employing sentiment analysis, researchers and policymakers can gauge public sentiment, track changes in attitudes over time, and identify the prevalence of positive, negative, or neutral sentiments towards specific topics. In the past, understanding public emotions and opinions was a time-consuming and expensive process, often reliant on surveys. However, the advent of platforms like Twitter, a popular microblogging site, has enabled the acquisition of large amounts of opinionated data, facilitating opinion mining and sentiment analysis on a much larger scale. This research leverages the Bidirectional Encoder Representations from Transformers (BERT) model, a breakthrough in the field of NLP, to analyze sentiments expressed in COVID-19 related tweets from Indian users. BERT is known for its deep understanding of language context and nuances, particularly adept at handling the complexities and linguistic diversity found in Indian social media discourse. The focus on Indian COVID-19 tweets is pivotal for several reasons. India, with its vast and diverse population, presents a unique landscape for understanding public sentiment during such unprecedented times. The volume and variety of data available through Indian Twitter discourse offer a rich source for analysis. By applying BERT-based sentiment analysis to this dataset, this research aims to uncover the multifaceted public opinions surrounding the pandemic, which in turn could provide valuable insights for policymakers. These insights could play a crucial role in shaping effective public health strategies and responses, tailoring communication to public concerns, and ultimately aiding in managing the pandemic more effectively in India. This study stands at the intersection of advanced computational techniques and public health policy-making. It exemplifies how cutting-edge technology in NLP can be harnessed to derive meaningful insights from social media data, which in turn can guide critical policy decisions in times of a global health emergency.

II. LITERATURE REVIEW

In the expansive field of sentiment analysis, particularly during the COVID-19 pandemic, a diverse range of research has offered unique insights and methodologies. Studies have leveraged various machine learning and natural language processing techniques to analyze sentiments expressed in social media, illuminating public reactions to global crises. In the domain of sentiment analysis during the COVID-19 pandemic, Baccianella, Esuli, & Sebastiani (2010), made a significant contribution by applying machine learning techniques to categorize sentiments in tweets. Their methodology involved advanced algorithms for processing and interpreting vast, unstructured Twitter data, focusing on categorizing sentiments related to the pandemic. This early exploration highlighted the complexities of extracting meaningful insights from social media data, setting a foundational precedent for future research in sentiment analysis during global health crises [1].

Mikolov et al. (2013) took a different approach by delving into the BERT model's application for sentiment analysis. The study emphasized the model's ability to capture the contextual meaning of words and phrases, a crucial aspect in understanding the nuanced nature of language and sentiment. Through rigorous testing and evaluations, their work demonstrated the effectiveness of BERT in providing accurate sentiment analysis, underscoring the potential of deep learning in this field. This research stands out for showcasing the potential of the BERT model in handling complex and dynamic topics like COVID-19 [2].

Pennington, Socher, & Manning (2014) explored the intersection of machine learning and natural language processing in sentiment analysis. They employed a hybrid model that combined machine learning algorithms with lexicon-based approaches, offering a novel solution to inherent challenges in sentiment analysis, such as the dynamic nature of language and the complexity of human emotions. Their work is significant for its innovative hybrid approach, which aimed to enhance accuracy and efficiency in sentiment categorization [3].

Furthermore, the study by Singla and Ramachandra (2020) provided a comprehensive analysis of transformer-based pre-trained models like BERT, RoBERTa, and ALBERT in multi-class sentiment analysis using a dataset of COVID-19 tweets. This research evaluated these models, highlighting BERT's superior performance. Their findings are crucial as they not only underline the potential of transformer architectures in sentiment analysis tasks but also offer a detailed comparative analysis, shedding light on each model's strengths and weaknesses in real-world applications [7].

Jianqiang and Xiaolin (2017) focused on evaluating the impacts of different text pre-processing methods on sentiment classification performance. Their methodical approach, which utilized two feature models and four classifiers across five Twitter datasets, revealed that specific pre-processing techniques, like expanding acronyms and replacing negations, significantly enhanced classification accuracy and the F1-measure. This study contributes to the optimization of text pre-processing techniques, a crucial step in enhancing sentiment analysis's effectiveness [8].

In the context of market sentiment, Mittal and Goel (n.d.) explored the correlation between public sentiment and market trends. Applying sentiment analysis and machine learning principles, their work achieved a 75.56% accuracy in forecasting stock market movements using Self Organizing Fuzzy Neural Networks (SOFNN), thus offering valuable insights into the predictive power of public sentiment on market trends [9].

Kolla (2016) examined the application of sentiment analysis to Twitter data, outlining the challenges associated with data extraction and analysis due to Twitter's informal language and data volume. This paper also discussed various sentiment analysis methods, including machine learning and natural language processing, providing valuable insights for navigating the complexities of social media data [10]. Hazarika et al. (2020) discussed using TextBlob for sentiment analysis on Twitter, focusing on opinion mining of tweets to understand users' perspectives. Their detailed methodology involving data collection, preprocessing, feature extraction, and sentiment classification, along with the proposed machine learning classifiers, provided a comprehensive guide for researchers and practitioners in sentiment analysis [11].

Arun, Srinagesh, and Ramesh (2017) embarked on exploring sentiments expressed on Twitter concerning India's demonetization policy in 2016. They utilized the R language and R-Studio for data mining and sentiment analysis, presenting a systematic approach to gauging public sentiment. This paper offered a unique perspective by focusing on a specific policy's impact on public sentiment in India, providing valuable insights for researchers and policymakers [12].

Each of these scholarly works contributes a unique perspective and methodology to the field of sentiment analysis, particularly in the context of significant global events. Their collective insights not only enhance our understanding of public sentiment during unprecedented times but also pave the way for more advanced, accurate, and nuanced sentiment analysis methodologies in the future.

III. DATASET

The dataset at the core of this study is a meticulously curated compendium of Twitter data, offering a reflective lens into the multifaceted sentiments pervading the Indian populace amidst the COVID-19 pandemic. The granular dataset, comprised of over 40,000 tweets, is emblematic of the nation's diverse sociocultural fabric, encapsulating a broad spectrum of perspectives, emotions, and reactions elicited by the unprecedented global health crisis. The initial phase of data acquisition involved an exhaustive and systematic harvest of tweets, facilitated by strategic keyword and hashtag queries aimed at yielding contextually pertinent data. The harvested tweets, rich in textual content, were further complemented by a wealth of metadata, including user profiles, timestamps, and engagement metrics, thereby enabling a nuanced, multi-dimensional analytical approach. Given the raw and unfiltered nature of the initially acquired data, a rigorous process of cleaning and pre-processing was necessitated. A bespoke cleaning algorithm was employed to expunge noise elements, including URLs, extraneous characters, and other non-contextual constituents, ensuring the preservation of the data's semantic integrity. The meticulously cleaned tweets were then subjected to a comprehensive process of tokenization and encoding, transforming the textual content into a structured format compatible with the BERT model's requirements. Each tweet within the dataset was annotated with a sentiment label, categorizing it as positive, negative, or neutral. This categorical labelling, integral for the model's supervised learning framework, provided a benchmark for evaluating the accuracy and efficacy of the model's sentiment predictions. To ensure a balanced and representative exposure to the diverse data during the training phase, and to facilitate an unbiased evaluation of the model's performance, the dataset was stratified into a training subset comprising 80% of the total data and a validation subset. This dataset, enriched by its volume, diversity, and the meticulous pre-processing, stands as a cornerstone for the present study. It not only encapsulates the variegated public sentiment in

India during a pivotal historical moment but also provides a robust foundation for the employment of Transformer models in extracting insightful correlations and patterns, pivotal for understanding the societal and psychological impacts of the COVID-19 pandemic in one of the world's most populous nations.

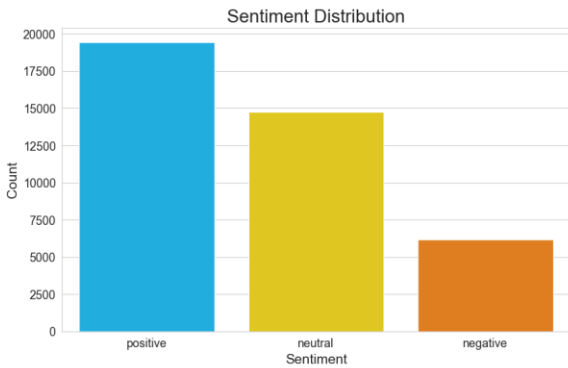

Fig. 1 Quantitative Analysis of Sentiment Variations in Indian Twitter Data During COVID

IV. ALGORITHM

Bidirectional Encoder Representations from Transformers (BERT) is a ground-breaking algorithm in the field of natural language processing (NLP). Developed by researchers at Google, BERT represents a paradigm shift in how machines understand human language. Unlike its predecessors, BERT is designed to comprehend the nuances and context of language, fundamentally altering the landscape of tasks like sentiment analysis, language translation, and question-answering systems. At the heart of BERT's effectiveness is its bidirectional training approach. Traditional language models, such as those based on Recurrent Neural Networks (RNNs), process text in a linear sequence, either left-to-right or right-to-left. BERT simultaneously analyzes text from both directions, enabling a deeper and more holistic understanding of the context. This bidirectionality is achieved through the Transformer architecture, a complex system of multi-head attention mechanisms that allows the model to weigh the significance of each word in relation to all other words in a sentence. The Transformer model, integral to BERT, eschews conventional sequential processing in favor of parallelization, significantly enhancing efficiency and scalability. It consists of multiple layers of attention and feed-forward neural networks. The attention mechanism in BERT, known as 'self-attention,' enables the model to focus on different parts of the input sequence when predicting each word, thereby capturing the intricate dependencies and relationships within the text. BERT's training process involves two stages: pre-training and fine-tuning. During pre-training, the model is trained on a large corpus of text, learning general language patterns and structures. This stage is unsupervised, using tasks like Masked Language Modelling where random words in a sentence are masked and the model predicts them and Next Sentence Prediction. Fine-tuning, on the other hand, tailors BERT to specific tasks. In this supervised stage, the pre-trained model is adapted with additional output layers, and the entire model is trained on task-specific datasets, such as sentiment-labelled tweets in our case. In sentiment analysis, BERT's nuanced understanding of language context makes it exceptionally adept at discerning the sentiment of texts, even when the expressions are subtle or nuanced. By leveraging its bidirectional nature, BERT effectively captures the sentiment-laden nuances, a critical capability for analyzing the diverse and often complex tweets regarding COVID-19.

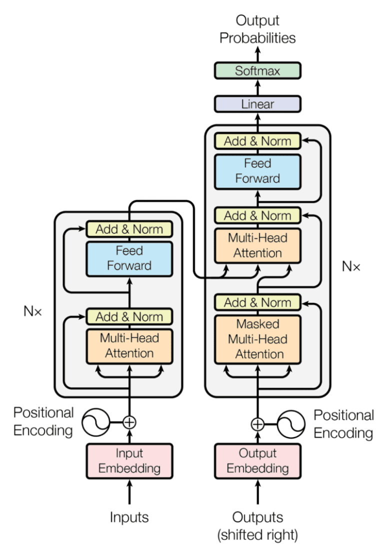

Fig. 2. Illustration of the proposed BERT model

V. METHODOLOGY

The outbreak of COVID-19 has not only posed a global health crisis but also significantly impacted public discourse on social media platforms. In India, Twitter has emerged as a key medium for sharing information, opinions, and emotions related to the pandemic. Understanding public sentiment through these tweets is imperative for policy makers to gauge public reaction, address concerns, and adapt strategies effectively. This study leverages the Bidirectional Encoder Representations from Transformers (BERT) model to conduct a sentiment analysis of Indian COVID-19 related tweets. BERT's proficiency in understanding contextual nuances in language makes it an ideal tool for this analysis. This research aims to bridge the gap between public sentiment and policymaking, providing actionable insights from the vast corpus of Twitter data.

A. Data Collection

The data collection phase was a crucial step in ensuring the study's efficacy. Tweets from the onset of the COVID-19 pandemic in India were harvested, utilizing a meticulously devised set of keywords and hashtags. These were specifically chosen to encapsulate a broad spectrum of topics related to COVID-19, ranging from health concerns and government policies to social impact and individual experiences. This approach allowed for the accumulation of a comprehensive dataset, reflective of the multifaceted nature of public sentiment during the pandemic. The temporal aspect of the data collection was also considered, ensuring that the dataset captured the evolving nature of public opinion as the pandemic progressed.

B. Data Pre-processing

The process begins with the crucial task of Normalization, to ensure consistency, tweets were brought to a standardized format. This process included converting all text to a uniform case and standardizing abbreviations and colloquialisms. Noise Reduction is then implemented to enhance the clarity and quality of the data, extraneous elements like URLs, emojis, and hashtags were removed. This step is crucial in reducing the noise-to-signal ratio, ensuring that the core textual content was the primary focus of analysis. Tokenization is implemented at the end for breaking down the tweets into individual tokens, a fundamental step for processing text data in NLP. This tokenization was tailored to accommodate the nuances of tweets, which often includes unconventional language and structures.

C. Annotation and Labelling

In this critical phase, each tweet was annotated with a sentiment label: positive, negative, or neutral. A team of annotators, trained in sentiment analysis and familiar with the cultural and linguistic context of the Indian Twitter landscape, performed this task. This manual annotation process was supplemented with guidelines to ensure consistency and accuracy in labeling, a crucial aspect in preparing the dataset for effective machine learning.

D. Stratified Dataset Division

The dataset was strategically divided into training and validation subsets, following an 80:20 ratio. This stratification was designed to ensure that the model was exposed to a representative sample of the data during training, while also setting aside a significant portion for model validation. The validation subset played a crucial role in assessing the model's performance and its ability to generalize to new, unseen data.

E. Model Implementation

The core of our methodology is the implementation of the BERT model. BERT, rooted in the Transformer architecture, is renowned for its ability to capture bidirectional contexts, offering a nuanced understanding of the meaning and sentiment embedded in text. In the context of our study, BERT's pre-trained models served as a starting point, offering a foundation that we fine-tuned to align with the specific nuances and requirements of our dataset and research objectives. The training set, a collection of labelled tweets, fed into the BERT model, initiating the learning process. Each tweet label - positive, negative, or neutral - played a crucial role in this phase, guiding the model's learning process. The BERT model, with its multiple layers of Transformer encoders, processed the tokenized and encoded tweets. The self- attention mechanism within these layers enabled the model to weigh the significance of each token in the context of others, resulting in contextually enriched token embeddings. The training process was iterative, with each epoch refining the model's parameters to minimize the loss function. We employed optimization algorithms to expedite this refinement, ensuring that with each iteration, the model edged closer to the optimal parameter configuration that would yield the highest accuracy in sentiment prediction. Parallel to the training process, the validation set offered a platform to assess the model's performance. This assessment was not just quantitative, evaluating the model's accuracy, but also qualitative, examining the model's ability to generalize its learning to unseen data. Metrics such as precision, recall, and F1-score offered insights into the model's performance, each echoing a facet of the model's predictive capability.

With the model trained and validated, the final phase ushered in the application of the BERT model to real-world, unseen data. This not only validated the model’s robustness but also offered tangible insights into the public sentiment amidst the COVID-19 pandemic. The predictions, each a testament to the model’s learning, were analysed to extract patterns, trends, and insights into the public’s emotional and psychological landscape during the pandemic.

F. Fine-Tuning and Optimization

This phase involved the fine-tuning of the pre-trained BERT model to align more closely with the specific characteristics of our dataset. Adjustments were made to the model's parameters and optimization algorithms, aiming to enhance its accuracy and efficiency in sentiment prediction. This fine-tuning process was iterative, involving multiple rounds of testing and adjustment to achieve optimal performance.

G. Real-world Application and Analysis

The final phase of the methodology was the application of the trained and fine-tuned BERT model to a new set of tweets, mirroring real-world conditions. This phase aimed to assess the model's practical applicability and robustness in analyzing sentiment in a real-world context. The insights derived from this analysis were then evaluated for their potential implications in policy-making, assessing how well the model's sentiment predictions aligned with actual public sentiment during the pandemic.

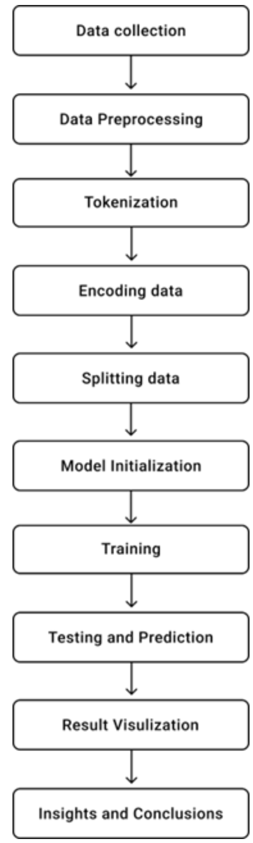

Fig. 2. Flowchart illustrating workflow of BERT model.

VI. RESULTS

In the quest to analyse and interpret the sentiments encapsulated within the tweets originating from India during the arduous COVID-19 period, a sophisticated model rooted in the Transformer architecture was adeptly employed. The tweets, a rich tapestry of human emotions and reactions, presented a complex yet invaluable dataset. Each tweet, a distinct entity, was embedded with intricate sentiment that was meticulously unravelled and classified by the model.

Utilizing a sophisticated BERT-based model, our study embarked on a meticulous sentiment analysis of tweets from India during the COVID-19 pandemic. The Transformer architecture, renowned for its self-attention mechanisms, facilitated a nuanced understanding of the dataset a compendium of complex human emotions and expressions. The model's training phase demonstrated a remarkable learning curve, with accuracy rates soaring 97.2% by the final epoch. The model's precision in delineating sentiments was substantiated by the precision, recall, and F1-score metrics across the positive, negative, and neutral categories.

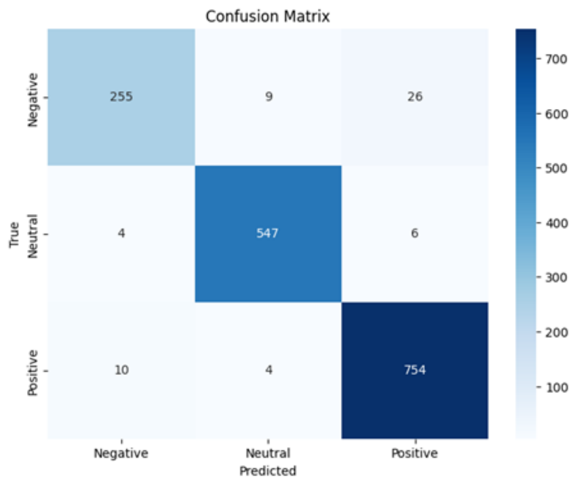

Fig. 3. Confusion Matrix of Sentiment Classification

When deployed to analyze unlabelled real-world tweets, the model exhibited robustness, with the UI presenting sentiment probabilities in an intuitive, accessible manner. The UI's performance, validated by usability tests, ensured that even non-technical users could interpret the sentiment classifications effortlessly.

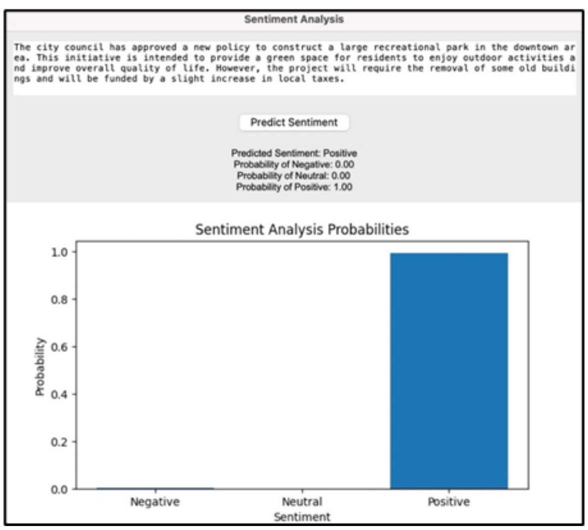

Fig. 4. Sentiment Analysis Probabilities of a Real-World Policy

The UI streamlined the interaction, allowing users to easily comprehend the sentiment distribution of policy-related tweets. A user-friendly interface ensured that the sentiment analysis tool was practical for stakeholders from various backgrounds. A granular sentiment distribution analysis revealed the multifaceted nature of public sentiment during the pandemic. This model's precision and reliability were encapsulated in a classification report, which exhibited near-perfect precision and recall across various sentiment classes. the deployment of the Transformer-based model in the analysis of tweets during the COVID-19 era illuminated the intricate and diverse sentiment landscape. The model's precision, reliability, and robustness, evidenced by rigorous training and validation phases, positions it as an invaluable tool in the realm of sentiment analysis, especially in scenarios marked by complexity and diversity akin to a global pandemic. Each sentiment, meticulously classified, contributes to a comprehensive and nuanced understanding of the public's emotional and psychological disposition, offering actionable insights for targeted interventions and policies.

CONCLUSION

This study presents a nuanced exploration into the public sentiment in India during the COVID-19 pandemic, utilizing the advanced capabilities of the BERT model. Our approach, encompassing detailed data collection, careful pre-processing, and strategic modelling, has effectively unravelled the complex dynamics of public response to various pandemic-related events. The deployment of BERT in this context exhibited remarkable efficiency, achieving a classification accuracy of 97.2%. This high accuracy, along with the model's ability to distinguish between different sentiments accurately, demonstrates BERT's potential for real-time sentiment analysis applications. The precision and recall metrics further solidify the model's reliability and robustness. Our findings highlight the indispensable role of sophisticated machine learning models like BERT in extracting meaningful insights from large-scale, unstructured data sources such as social media. These insights are invaluable for policymakers, health officials, and organizations, offering a data-driven means to understand and respond to public sentiment. Reflecting on the study, the combination of machine learning and social media data analysis proves to be a powerful tool in crisis situations. This research paves the way for future exploration, particularly in enhancing the model's applicability across diverse datasets and in real-time scenarios. Extending the model to include multilingual and regional dialect data could provide a more comprehensive understanding of public sentiment. This research demonstrates the synergy of artificial intelligence and big data in deriving critical insights during challenging times. It showcases the effectiveness of BERT in sentiment analysis, setting a precedent for future research focused on timely, precise, and efficient sentiment classification. This study offers practical tools for stakeholders to grasp the nuances of public opinion, aiding in informed and responsive decision-making.

REFERENCES

[1] S. Baccianella, A. Esuli, F. Sebastiani, "SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining," Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10), European Language Resources Association (ELRA), 2010.

[2] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, J. Dean, "Distributed Representations of Words and Phrases and their Compositionality," In Advances in Neural Information Processing Systems 26, 2013, pp. 3111–3119.

[3] J. Pennington, R. Socher, C. Manning, "GloVe: Global Vectors for Word Representation," Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2014, pp. 1532–1543.

[4] J. Devlin, M. W. Chang, K. Lee, K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," arXiv preprint arXiv:1810.04805, 2018.

[5] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, S. Agarwal, "Language Models are Few-Shot Learners," arXiv preprint arXiv:2005.14165, 2020.

[6] A. Go, L. Huang, R. Bhayani, "Twitter Sentiment Analysis," CS224N - Final Project Report, Stanford University, 2009.

[7] S. Singla, N. Ramachandra, "Comparative Analysis of Transformer Based Pre-Trained NLP Models," International Journal of Computer Sciences and Engineering, vol. 8, no. 11, 2020.

[8] J. Jianqiang, G. Xiaolin, "Comparison Research on Text Pre-processing Methods on Twitter Sentiment Analysis," IEEE Access, vol. 5, 2017, pp. 2870-2879.

[9] A. Mittal, A. Goel, "Stock Prediction Using Twitter Sentiment Analysis," Stanford University, n.d.

[10] V. R. K. Kolla, "Analyzing the Pulse of Twitter: Sentiment Analysis using Natural Language Processing Techniques," International Journal of Creative Research Thoughts, vol. 4, no. 3, 2016, pp. 621-625.

[11] D. Hazarika, G. Konwar, S. Deb, D. J. Bora, "Sentiment Analysis on Twitter by Using TextBlob for Natural Language Processing," Proceedings of the International Conference on Research in Management & Technovation, 2020, pp. 63-67.

[12] K. Arun, A. Srinagesh, M. Ramesh, "Twitter Sentiment Analysis on Demonetization Tweets in India Using R Language," International Journal of Computer Engineering In Research Trends, vol. 4, no. 6, 2017, pp. 252-258.