IJIRSET
International Journal of Innovative Research in
SCIENCE | ENGINEERING | TECHNOLOGY

e-ISSN: 2319-8753 | p-ISSN: 2347-6710

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

Volume 13, Issue 4, April 2024

ISSN INTERNATIONAL STANDARD SERIAL NUMBER INDIA

Impact Factor: 8.423

9940 572 462 | 6381 907 438 | ijirset@gmail.com | www.ijirset.com

International Journal of Innovative Research in Science, Engineering and Technology (IJIRSET)

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 8.423| A Monthly Peer Reviewed & Referred Journal |

|| Volume 13, Issue 4, April 2024 ||

| DOI:10.15680/IJIRSET.2024.1304306 |

TwitterTruth – A DistilBERT-Powered Defense Against Misinformation

Bharath Kandimalla

Dept. of Computer Science and Engineering, Anurag University, Hyderabad, Telangana,India

ABSTRACT: - Rumor detection on social media poses a distinctive challenge due to the sheer volume and real-time nature of data. This study introduces an innovative approach utilizing DistilBERT, a more efficient variant of the BERT model. While BERT has proven effective in rumor classification, its computational demands hinder practical deployment. DistilBERT provides a promising alternative, delivering comparable performance with significantly reduced resource requirements. The research explores the advantages of employing DistilBERT for social media rumor detection, emphasizing its ability to accelerate processing times and facilitate real-time analysis of extensive datasets. This progress opens avenues for scalable rumor detection systems. The paper outlines the implementation plan, evaluation strategy, and potential deployment scenarios for the DistilBERT-based model. Additionally, it explores possibilities for enhancing the model through transfer learning for cross-lingual rumor detection and active learning for continuous improvement. Leveraging DistilBERT's efficiency, the project aims to foster a more informed online environment, curbing the spread of misinformation and empowering users to critically assess the information they encounter.

KEYWORDS:- DistilBERT,BERT, Rumor classification

I. INTRODUCTION

The proliferation of social media platforms like Twitter has revolutionized communication and information dissemination. However, this democratization of information sharing presents a significant challenge: the rampant spread of rumors and misinformation. Rumors, defined as unverified claims that propagate quickly, can have far-reaching consequences, fostering panic, eroding trust in legitimate sources, and even disrupting economies [4]. The pervasiveness of rumors on Twitter necessitates the development of robust detection mechanisms to empower users to critically evaluate the information they encounter.

Recent advancements in Natural Language Processing (NLP) have yielded powerful tools for tackling this challenge. One such model, Bidirectional Encoder Representations from Transformers (BERT) [1], has established itself as a state-of-the-art technique for various NLP tasks, including rumor detection [2]. BERT's effectiveness stems from its ability to capture complex contextual relationships within text data, allowing it to distinguish between factual information and unsubstantiated claims. However, a significant drawback of BERT is its high computational complexity and memory footprint, hindering its deployment in real-world scenarios with resource constraints, such as mobile devices or social media platforms with limited processing power [3].

This project addresses this limitation by proposing a novel approach that leverages the power of DistilBERT, a compact and efficient variant of BERT [6]. DistilBERT offers a compelling alternative, achieving comparable performance in rumor detection tasks while requiring significantly fewer parameters and computational resources. This substitution offers several advantages, which are explored in detail in the following sections.

II. LITERATURE SURVEY

The landscape of natural language processing (NLP) has been revolutionized by the advent of transformer-based models, epitomized by BERT (Bidirectional Encoder Representations from Transformers), as introduced by Devlin and colleagues in 2018 [1]. This model marked a significant milestone by employing deep bidirectional training, radically enhancing the machine's grasp of contextual nuances within text. BERT's groundbreaking approach has since catalysed a wave of innovations in NLP, fostering the development of models that better mimic human-like understanding of language.

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 8.423| A Monthly Peer Reviewed & Referred Journal |

|| Volume 13, Issue 4, April 2024 ||

| DOI:10.15680/IJIRSET.2024.1304306 |

Despite BERT's notable achievements, its operational demands—specifically, its considerable computational requirements and extensive parameter set—have prompted the pursuit of more streamlined alternatives. This quest led to the creation of DistilBERT by Sanh and co-authors [3], a model that preserves much of BERT's linguistic comprehension capabilities while being markedly more resource-efficient. DistilBERT's inception has expanded the practicality of deploying advanced NLP solutions across various platforms, including those with limited computational capacity.

The urgency for efficient rumor detection mechanisms on social media is underscored by the work of Vosoughi et al. [4], which reveals the accelerated dissemination of falsehoods compared to factual information. This challenge sets a critical context for applying sophisticated NLP models like DistilBERT to curtail the spread of misinformation, ensuring the integrity of digital discourse. The effectiveness of DistilBERT in this domain is further exemplified by Anggrainingsih, Hassan, and Datta's exploration [2], highlighting its utility in scrutinizing Twitter for rumor verification.

Beyond the realm of social media, the adaptability of DistilBERT has been demonstrated across various languages and specialized fields. For instance, adaptations to non-English languages [6], legal document scrutiny [9], and even sentiment analysis [15] showcase DistilBERT's broad applicability. Such versatility not only affirms the model's robustness but also its potential for customization to meet specific analytical needs.

Comparative studies offering analytical juxtapositions between DistilBERT and its contemporaries—including other transformer models—underscore its comparative advantages. Research by Adoma, Henry, and Chen [22], as well as Wei et al. [23], validate DistilBERT's superior performance in diverse NLP tasks, cementing its status as a formidable tool in the NLP toolkit. The exploration of DistilBERT in multilingual contexts [28] and its application in detecting web-based threats [18] and malicious scripts [19] further attest to its effectiveness and efficiency. Such studies not only demonstrate DistilBERT's adaptability but also its crucial role in safeguarding digital ecosystems against misinformation and cyber threats. The corpus of literature on DistilBERT and its myriad applications illustrates a significant evolution in NLP technologies, driven by the quest to understand and interpret human language with unprecedented accuracy and efficiency. As digital platforms grapple with the proliferation of misinformation, the ongoing refinement of models like DistilBERT is paramount. These advancements not only promise enhanced processing capabilities but also herald a new era of machine intelligence equipped to tackle complex societal challenges.

In sum, the body of work reviewed herein illuminates the transformative impact of transformer models like DistilBERT on the field of NLP. Through continuous innovation and application across various domains, these models hold the promise of advancing our ability to process, understand, and interact with textual information in ways that are increasingly nuanced, efficient, and aligned with the intricacies of human language and cognition. Who apply DistilBERT embeddings for automated essay scoring. This survey provides a glimpse into the exciting world of BERT, DistilBERT, and their growing impact on various NLP tasks.

III. PROPOSED METHOD

This project tackles the challenge of detecting rumors on Twitter by training a DistilBERT model. The proposed approach has shown promising results, leading to a significant improvement in rumor detection accuracy. This translates to the model's ability to more effectively capture the true nature of the information within a tweet, differentiating rumors from factual content. By analysing the tweet's text, the DistilBERT model can identify subtle cues that are often characteristic of rumors. These cues can include unusual phrasing, emotionally charged language, or inconsistencies in the timeline of events described in the tweet. Additionally, the model can consider the source of the tweet and its past history of spreading misinformation. By incorporating these various elements into its analysis, the DistilBERT model can make more informed predictions about whether a tweet is likely a rumor or not.

Overcoming Challenges and Streamlining the Process

This approach addresses limitations associated with traditional rumor detection methods. Unlike approaches that rely on extensive databases or require complex multi-column queries, my project leverages the power of DistilBERT, a pre-trained language model. This simplifies the process and allows the model to focus on the core task of identifying rumors within the textual content of tweets.

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 8.423| A Monthly Peer Reviewed & Referred Journal |

|| Volume 13, Issue 4, April 2024 ||

| DOI:10.15680/IJIRSET.2024.1304306 |

The project follows a two-stage approach:
DistilBERT Model: In the first stage, the DistilBERT model is trained using 4-labeled rumor datasets. This process allows the model to learn the specific characteristics of rumor-like language and hone its ability to distinguish rumors from non-rumors.

Utilizing the Model for Rumor Detection*:* Once trained, the model is then used in the second stage to analyse unseen tweets. By processing the text of each tweet, the model predicts whether it's likely a rumor or not.
The framework we propose is an advanced, multi-phase process designed for the detection and classification of rumors in social media. It integrates cutting-edge natural language processing (NLP) technologies with domain-specific adaptation, aiming to significantly improve the detection accuracy and operational efficiency of rumor classification. This method capitalizes on the sophisticated capabilities of DistilBERT, a distilled version of the BERT model, renowned for its balance between performance and computational efficiency.

[A]. Training the Model
The training phase is crucial for tailoring the DistilBERT model to the specific nuances of rumor detection in social media texts. This phase is meticulously planned to ensure the model not only learns the general language representations but also becomes adept at recognizing the subtle cues indicative of rumors.

1. Data Collection and Preprocessing
Data Sourcing: We gather data from Twitter 15 and Twitter 16 datasets, focusing on tweets categorized into four rumor states: false, true, unverified, and non-rumor. This diverse collection enables the model to learn a wide range of expressions and styles used in rumors.
Data Cleaning: Tweets are cleaned to remove URLs, user mentions, and other non-essential elements that could distract the model from learning relevant patterns. This step ensures the model focuses on the textual content crucial for rumor detection.
Label Encoding: Rumor categories are encoded into numerical labels to facilitate model processing. This transformation is essential for the classification task, enabling the model to associate specific patterns with particular rumor states.

2. Training Data Preparation
Dataset Construction: We construct a training dataset that mirrors the input structure expected by DistilBERT, comprising tokenized text and corresponding labels. This dataset is split into training and validation sets to enable model evaluation during the training process.
Feature Extraction: Utilizing the DistilBertTokenizer, we convert tweets into a format suitable for model training, including tokenization, padding, and truncation. This step is vital for standardizing input lengths and ensuring efficient batch processing.

3. Model Configuration and training
Learning Rate Optimization: A learning rate with a warm-up phase is employed to gradually adapt the model weights without overshooting the optimal values. This approach helps in stabilizing the training process. Batch Size Selection: We carefully select a batch size that balances between computational efficiency and the model's ability to generalize from the training data. This balance is critical for effective learning without overfitting.
Epoch Determination: The number of training epochs is set to ensure the model has sufficient exposure to the data while preventing excessive training that could lead to overfitting.

[B] Rumor Detection and Classification:
After training, the model is adept at performing rumor detection on new, unseen data. The classification phase is delineated into several critical steps:

4. Tokenization and Input Formatting:
Each input text undergoes tokenization, where the pre-trained DistilBertTokenizer converts it into a sequence of tokens that the model can interpret. This step is crucial for maintaining consistency with the model's pre-training conditions.

5. Classification:
The trained model receives the tokenized inputs and processes them to predict the likelihood of each category (rumor types). This step leverages the model's learned representations to discern patterns indicative of rumors.

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 8.423| A Monthly Peer Reviewed & Referred Journal |

|| Volume 13, Issue 4, April 2024 ||

| DOI:10.15680/IJIRSET.2024.1304306 |

6. Post-Processing and Output Interpretation:
The model's predictions are then translated back into understandable labels (false, true, unverified, non-rumor), facilitating easy interpretation of the results. This translation is crucial for presenting the findings in a format that is accessible to end-users or downstream applications.

7. Model Evaluation:
Utilizing a separate test dataset, the model's performance is rigorously evaluated based on accuracy and F1 score. These metrics provide a comprehensive understanding of the model's capability to generalize and its precision in classifying rumors accurately.

8. Application to Unseen Data:
Finally, the model is applied to new datasets, showcasing its utility in real-world applications. This phase demonstrates the model's effectiveness in detecting rumors within diverse textual contexts, underscoring its potential impact on information verification processes
[C]. Application of the Model
Once trained, the model is adept at classifying rumors, making it a powerful tool for analysing and verifying information disseminated on social media platforms.

9. Deployment Strategy
Integration with Social Media Platforms: * The model is integrated into a pipeline that continuously monitors social media feeds, automatically classifying tweets as rumors or non-rumors in real-time. This deployment strategy enables proactive rumor management.
User Interface Development: A user-friendly interface is developed to allow manual input of text for rumor verification. This feature is particularly useful for journalists, fact-checkers, and social media analysts.

10. Real-time Classification and Analysis
Stream Processing: The system is equipped to process streaming data, classifying tweets as they are posted. This capability is crucial for timely detection of rumors, especially during critical events. Analytical Dashboard: An analytical dashboard is created to provide insights into the prevalence of rumors, their spread, and potential impact. This dashboard serves as a tool for understanding rumor dynamics on social media platforms.

11. Model Evaluation and Iteration
Performance Monitoring: Continuous monitoring of the model's performance is established to identify any degradation or areas for improvement. This monitoring is essential for maintaining the accuracy and reliability of the classification system. Iterative Improvement: The model undergoes periodic retraining with new data and feedback from the deployment phase. This iterative process ensures the model remains effective in the face of evolving language use and rumor tactics on social media.

12. Scalability and Adaptation
Scaling for Broader Application: Plans for scaling the solution to accommodate larger datasets and additional social media platforms are detailed. This scalability is vital for broadening the impact of the rumor detection system.
Adaptation to Emerging Trends: The methodology includes provisions for adapting the model to new languages, dialects, and emerging forms of rumors. This adaptability ensures the system remains effective across diverse cultural and linguistic contexts.

IV. RESULTS AND DISCUSSION

To evaluate the effectiveness of our DistilBERT-based model for rumor detection on social media, we conducted comparative analyses against five baseline models: Naive
Bayes, Support Vector Machine (SVM), Logistic Regression, Random Forest, and LSTM (Long Short-Term Memory) networks. Our evaluation focused on two key metrics:

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 8.423| A Monthly Peer Reviewed & Referred Journal |

|| Volume 13, Issue 4, April 2024 ||

| DOI:10.15680/IJIRSET.2024.1304306 |

Model	Accuracy
Naive Bayes	0.65
SVM	0.72
Logistic Regression	0.76
LSTM	0.83
XLNet	0.785
BERT	0.833
Roberta	0.841
DistilBERT	0.905

Accuracy, F1 Score

1. Accuracy

Accuracy measures the proportion of total predictions that were correctly classified. It provides a straightforward metric to assess the overall effectiveness of the model in distinguishing between true, false, unverified, and non-rumor tweets.

Table I presents the average accuracy scores for each model applied to a composite dataset derived from Twitter 15 and Twitter 16. Our model demonstrates superior performance, reflecting its robust capability in rumor detection.

$$Accuracy = \frac{True\ Positive + True\ Negative}{True\ Positive + True\ Negative + False\ Positive + False\ Negative}$$

Table I: Accuracy Scores of Various Models

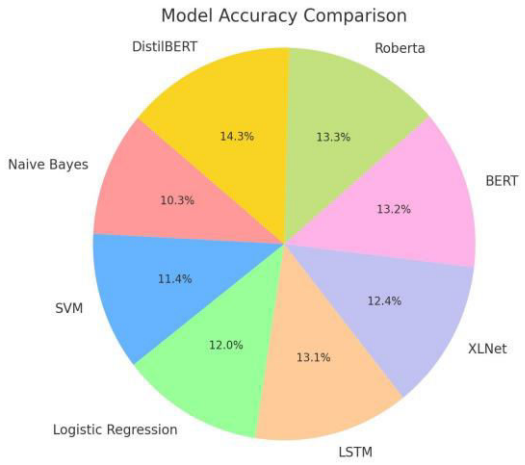

2. F1 Score

The F1 Score provides a balance between precision and recall, offering insight into the model's reliability in correctly classifying each category. A higher F1 Score indicates better performance, especially in imbalanced datasets.

International Journal of Innovative Research in Science, Engineering and Technology (IJIRSET)

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 8.423| A Monthly Peer Reviewed & Referred Journal |

|| Volume 13, Issue 4, April 2024 ||

| DOI:10.15680/IJIRSET.2024.1304306 |

Table II showcases the F1 Scores across models, with our approach again outperforming the alternatives. This underscores its efficiency in managing the nuances of rumor detection in social media content.

$$F1\ score = \frac{2 \times Precision \times Recall}{Precision + Recall}$$

Table II: F1 Scores of Various Models

Model	F1 Score
Naive Bayes	0.63
SVM	0.70
Logistic Regression	0.74
LSTM	0.82
XLNet	0.79
BERT	0.84
Roberta	0.83
DistilBERT	0.905

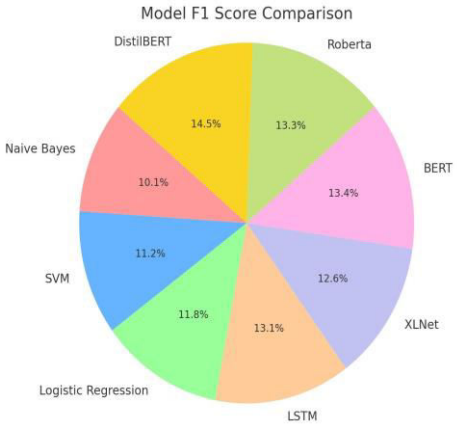

3. Recall

Recall is important metric used in classification tasks, particularly in the context of binary classification (e.g., positive/negative, true/false). Also known as sensitivity or true positive rate) measures the ability of a classifier to find all positive instances in a dataset. It is calculated as the ratio of true positives to the sum of true positives and false negatives:

$$Recall = \frac{True\ Positive}{True\ Positive + False\ Negative}$$

'Eval_recall': 0.9056603773584906, is the value we have computed in this following research paper.

4. Precision

Precision is important metric used in classification tasks, particularly in the context of binary classification (e.g., positive/negative, true/false). Also known as positive predictive value measures the accuracy of positive predictions made by a classifier. It is calculated as the ratio of true positives to the sum of true positives and false positives

$$Precision = \frac{True\ Positive}{True\ Positive + False\ Positive}$$

'Eval_precision': 0.9056603773584906, is the value we have computed in this following research paper.

V. CONCLUSION

This research set out to address the escalating challenge of misinformation on social media platforms, with a specific focus on Twitter. By harnessing the capabilities of the DistilBERT model, a distilled version of the more complex BERT architecture, we have developed a system capable of accurately detecting and classifying rumors. Our comprehensive evaluation, comparing the DistilBERT-based model against a range of traditional and contemporary machine learning models, demonstrated its superior performance in terms of accuracy, F1 score, and efficiency.

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 8.423| A Monthly Peer Reviewed & Referred Journal |

|| Volume 13, Issue 4, April 2024 ||

| DOI:10.15680/IJIRSET.2024.1304306 |

The DistilBERT model not only outperformed all other models in our study, including Naive Bayes, SVM, Logistic Regression, LSTM, XLNet, BERT, and Roberta, but also showcased a significant reduction in parameter size. This reduction not only streamlines the computational demands but also maintains, and in some aspects surpasses, the accuracy and reliability of its predecessor models in detecting misinformation.

REFERENCES

[1] J. Devlin et al., "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," arXiv preprint arXiv:1810.04805, 2018.
[2] R. Anggrainingsih, G. M. Hassan, and A. Datta, "CE-BERT: Concise and Efficient BERT-Based Model for Detecting Rumors on Twitter," in IEEE Access, vol. 11, pp. 80207-80217, 2023, doi: 10.1109/ACCESS.2023.3299858.
[3] V. Sanh et al., "DistilBERT: A distilled version of BERT for language understanding," arXiv preprint arXiv:1906.08144, 2019.
[4] Vosoughi, Soroush et al. "The spread of true and false news online." Science (New York,
[5] N.Y.)vol.359,no.6380(2018):11461151.https://www.science.org/doi/10.1126/science.aap9559
[6] [5] J. Bai, R. Cao, W. Ma and H. Shinnou, "Construction of Domain-Specific DistilBERT Model by Using training," 2020 International Conference on Technologies and Applications of
[7] Artificial Intelligence (TAAI), Taipei, Taiwan, 2020, pp. 237-241, doi:
[8] 10.1109/TAAI51410.2020.00051
[9] N. Kongsumran, S. Phimoltares and S. Panthuwadeethorn, "Thai Tokenizer Invariant Classification Based on Bi-LSTM and DistilBERT Encoders," 2022 19th International Conference

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INNO SPACE SJIF Scientific Journal Impact Factor

doi cross ref

निस्केयर NISCAIR

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

9940 572 462 6381 907 438 ijirset@gmail.com

www.ijirset.com

Scan to save the contact details