Misinformation Detection in Political News using BERT Model

Halyna Padalko[1,2], Vasyl Chomko[2] and Dmytro Chumachenko[1,2]

[1] National Aerospace University "Kharkiv Aviation Institute", Chkalow str., 17, Kharkiv, 61070, Ukraine
[2] University of Waterloo, 200 University Ave W, Waterloo, N2L 3G5, Canada

Abstract

In the digital age, the rapid dissemination of news via social media platforms has given rise to a significant challenge – the proliferation of fake news. This phenomenon sows confusion among the public and threatens the foundations of informed discourse and democracy. Addressing this issue, the present study explores the application of the Bidirectional Encoder Representations from Transformers (BERT) model for detecting fake news. With its deep learning architecture and contextual understanding of language, BERT offers a promising framework for this purpose. The model was trained and validated on a dataset distinguishing between 'Fake' and 'True' news, achieving an accuracy of 79.88% and an area under the receiver operating characteristic (ROC) curve (AUC) of 0.87. These metrics underscore the model's proficiency in classifying news articles correctly. The study's results demonstrate the potential of BERT in the realm of fake news detection, providing a novel tool for social media platforms to combat misinformation. Despite its efficacy, the research also highlights the model's limitations, including fluctuations in validation accuracy and a tendency to misclassify true news as fake, indicating areas for future improvement. This study contributes to the ongoing efforts to ensure the integrity of news content and offers a foundation for subsequent research in information authenticity.

Keywords

Misinformation, machine learning, deep learning, fake news, BERT

1. Introduction

The proliferation of misinformation, particularly in the political domain, significantly challenges the integrity of democratic processes and informed public discourse [1]. The advent of social media and online platforms has exponentially amplified the spread and impact of misinformation [2]. This has necessitated the development of robust and efficient computational tools to identify and mitigate the effects of false information.

The urgency of this issue is exemplified in the context of recent global political instabilities, most notably the Russian full-scale invasion of Ukraine [3]. This war has been marked not only by physical confrontation but also by an information war, where misinformation has been used as a strategic tool to influence public opinion, sow discord, and manipulate narratives [4]. The rapid dissemination of false information regarding the conflict, its origins, and ongoing developments has seriously threatened international peace and stability. In such a scenario, swiftly and accurately distinguishing between factual information and propaganda becomes paramount.

This conflict illustrates the broader implications of misinformation in international relations and national security [5]. The use of misinformation in the context of the Russian war in Ukraine has had far-reaching consequences, affecting not just the combatant nations but also international alliances, economic stability, and the global information ecosystem [6]. In such

ProfIT AI 2023: 3rd International Workshop of IT-professionals on Artificial Intelligence (ProfIT AI 2023), November 20–22, 2023, Waterloo, Canada
galinapadalko95@gmail.com (H. Padalko); vchomko@uwaterloo.ca (V. Chomko); dichumachenko@gmail.com (D. Chumachenko)
0000-0001-6014-1065 (H. Padalko); 0009-0009-4419-6651 (V. Chomko); 0000-0003-2623-3294 (D. Chumachenko)

situations, false narratives can escalate tensions, perpetuate hostilities, and hinder diplomatic efforts. Consequently, there is a pressing need for advanced tools capable of efficiently identifying and countering such misinformation.

This paper presents a novel approach to addressing this challenge by leveraging the Bidirectional Encoder Representations from Transformers (BERT) model [7]. BERT, a groundbreaking development in natural language processing (NLP), offers a sophisticated mechanism for understanding context in text. Its deep learning architecture, pre-trained on a vast corpus of text, is adept at grasping the nuances of language, making it an ideal candidate for misinformation detection in political news.

The impetus for this research stems from the critical need to uphold the veracity of information in the political sphere. Misinformation in political news can skew public perception, manipulate electoral outcomes, and erode trust in democratic institutions. As such, developing effective tools to counteract misinformation is not only a technical challenge but also a civic imperative.

In this study, we explore the application of the BERT model to the specific task of detecting misinformation in political news articles. We hypothesize that BERT's advanced language comprehension capabilities can be harnessed to discern between factual reporting and misinformation. By training the model on a dataset comprising both authentic news sources and known purveyors of misinformation, we aim to create a system capable of making this distinction with high accuracy.

By advancing the application of BERT in political news, this research contributes to the broader effort to combat misinformation and uphold the integrity of public discourse in the digital age.

2. Current research analysis

The landscape of misinformation detection, especially in political news, has been a dynamic field of study, marked by diverse methodologies and evolving technologies. Current research in this domain primarily focuses on leveraging advanced computational techniques, including machine learning and natural language processing, to develop systems capable of identifying and classifying news content concerning its authenticity. This section critically analyzes contemporary scholarly works that have contributed significantly to the field. These studies have explored various aspects of fake news detection, ranging from developing algorithmic models to examining linguistic and semantic patterns characteristic of misinformation. The subsequent analysis of these papers will delve into their methodologies, findings, and the implications of their contributions to combating misinformation in the digital era.

The paper [8] addresses the crucial issue of misinformation spread, including fake news, propaganda, and conspiracy theories, which pose a serious threat to society by potentially altering beliefs, behaviors, and policies. The authors propose a novel theoretical model to explain the psychological factors underlying the spread of misinformation and effective strategies for its reduction. This model integrates insights from previous research, encompassing various psychological aspects such as partisan bias, analytic thinking, and the need for chaos. The model is structured around several paths: the increase of belief in misinformation leading to increased sharing, the direct increase of sharing even without belief enhancement, and the role of psychological risk factors in increasing exposure to, belief in, and sharing of misinformation. This comprehensive approach incorporates findings from personality psychology, cognitive psychology, political psychology, and political science, providing a multifaceted understanding of the misinformation phenomenon and suggesting interventions for different stakeholders, including users, media outlets, online platforms, policymakers, and institutions.

The study [9], conducted as a two-wave panel survey in Chile, focuses on the intersection of psychological, social, and political factors in spreading fake news. Utilizing OLS regression in a lagged dependent variable model, the research analyzes how personal traits like conspiracy theories, trust in others, education, gender, social media use, and political views influence fake news exposure, belief, and sharing. Key findings indicate a significant relationship between exposure to and belief in fake news and its subsequent sharing. At the same time, frequent social media use surprisingly correlates with less belief in fake news. The study's limitations include a high dropout rate between survey waves and a lack of differentiation among fake news topics. This research contributes to understanding the multifaceted nature of misinformation spread, highlighting the need for more comprehensive studies in this field.

The research [10] focuses on detecting fake news during the COVID-19 infodemic using a novel source-based method. This approach analyzes the community of news propagators, including posters and re-tweeters, connected through follower-following relationships on Twitter. The study combines complex network measures and user profile features in a machine learning framework to classify tweets. An extensive comparative analysis using eleven machine learning and two deep learning models revealed that combining hybrid features significantly outperforms individual network or user features. The most effective models were the Ensemble's boosting model CATBoost and the deep learning model RNN, achieving an AUC score of 98%. The methodology and results highlight the efficacy of incorporating both network dynamics and user profile characteristics in identifying misinformation in a global health crisis context.

The study [11] presents a comprehensive approach to identifying fake news using linguistic analysis. Focusing on political content on social media, the research incorporates 26 significant linguistic features selected based on the Pearson correlation coefficient. These features include complexity, readability index, psycholinguistic, and stylometric features. The study utilizes three feature extraction techniques – term frequency-inverse document frequency (tf-idf), count vectorizer (CV), and hash-vectorizer (HV) – to process the data. Machine learning models were then applied to four different datasets, achieving high accuracy rates: 93.8% for the Random Political dataset, 90% for the Buzzfeed dataset, 86.9% for the Mc_Intire dataset, and 90.8% using the Reuter dataset. The framework's effectiveness is underscored by its improved performance compared to existing state-of-the-art methods. Future work aims to expand the range of linguistic features and datasets, exploring real-time fake news detection methods and various architectural designs.

The study [12] examines the counterintuitive trend of popular and reputedly reliable news sources receiving the most flags on social media. The research introduces a 'bipolar' model based on the assumption that user polarization significantly affects the flagging of news items, incorporating factors such as echo chambers, confirmation bias, and platform-induced polarization. The model predicts that moderate and truthful news sources are more likely to be flagged due to their wider dissemination and consequent exposure to polarized user groups. In contrast, polarized and potentially untruthful news items are less likely to be shared across different network segments and, thus, less frequently flagged. The study's findings suggest that neutral news items are predominantly flagged in highly polarized environments, indicating a potential flaw in the current flagging mechanisms used by social media platforms to identify misinformation.

The current landscape of misinformation detection reveals a complex interplay between user behavior, linguistic features, and machine-learning techniques. The studies analyzed range from exploring the intricacies of user polarization in flagging mechanisms to the utilization of linguistic features in fake news detection, each contributing unique insights into the multifaceted nature of misinformation. These insights underscore the need for sophisticated, adaptive models to navigate the nuanced terrain of fake news. This leads us to our next research phase: applying the BERT model. BERT's advanced capability in understanding the context of language, combined with its deep learning architecture, positions it as a promising solution to address the challenges highlighted by these studies. By integrating the learnings from current research and leveraging BERT's robust processing power, we aim to develop a more effective and nuanced approach to identifying and combating misinformation in the digital realm.

3. Materials and methods

3.1. BERT model

The BERT model represents a significant advancement in NLP, particularly in fake news detection. Developed by researchers at Google, BERT's primary innovation lies in its ability to understand the context of words in a sentence bidirectionally rather than the traditional unidirectional or sequential approach [7]. This contextual awareness is crucial in comprehending the subtleties and complexities of language, which is essential in accurately detecting fake news.

BERT operates on the principle of transformers, which are models that process words about all other words in a sentence instead of one one-by-one in order. This allows for a more holistic understanding of the sentence structure and meaning. The model is pre-trained on a large corpus of text, enabling it to learn various language patterns and nuances. This pre-training includes tasks like predicting missing words in a sentence, which helps the model grasp contextual relationships between words.

For fake news detection, BERT can be fine-tuned with news-specific datasets. This fine-tuning involves training the model on a dataset comprising both legitimate and fake news, allowing it to learn the characteristics and patterns distinguishing authentic news from misinformation. BERT's strength lies in its ability to understand subtle cues and language variations often indicative of fake news, such as exaggerated claims, inconsistent information, or sensational language.

The application of BERT in fake news detection involves several steps. First, the textual content of news articles is input into the model. BERT then processes this content, considering the context of each word and sentence. The model generates representations of the text that capture both its linguistic properties and the learned patterns of fake and legitimate news. Finally, these representations are used to classify the news as either fake or genuine.

One of the key advantages of BERT in this context is its capacity for transfer learning. Having been pre-trained on a vast array of text, it can effectively adapt to news articles' specific language and style, even with a relatively small amount of fine-tuning data. This makes it highly efficient and accurate in identifying fake news, even in limited training data scenarios.

However, it is important to note some limitations of BERT in this application. While highly effective at understanding language, BERT requires significant computational resources, which can be a constraint in some environments. Additionally, the model's performance can be influenced by the quality and representativeness of the training data. Biased or unbalanced training datasets can lead to less accurate classifications.

In the learning process of our BERT-based classifier, the model is trained over multiple epochs using the AdamW optimizer and a linear learning rate scheduler. Training involves backpropagation to adjust weights based on the loss calculated from the model's predictions compared to actual labels. The model's performance is periodically evaluated on a validation dataset, assessing accuracy and generating classification reports. This training and evaluation loop iteratively improves the model's ability to classify text as true or false.

3.2. Data

For the analysis, we have used the dataset "Fake-Real News" available on Kaggle [13]. The dataset is a pivotal resource for research in misinformation detection, particularly tailored for training and evaluating machine learning models in discerning authentic and deceptive news content. This dataset is instrumental for practitioners and researchers in natural language processing, emphasizing the detection of fake news.

Structured in a tabular format, the dataset comprises two distinct files, one encompassing fake news articles and the other containing real news pieces. Each file is meticulously organized with columns representing various attributes of the news articles, including the title, text body, and subject matter. The title column encapsulates the headline of each article, a critical aspect as headlines are often designed to captivate attention and may embody elements of sensationalism.

The text column provides the complete content of each article, presenting a comprehensive source for linguistic analysis to gauge the context, style, and detailed narrative, which are key in determining the article's authenticity. Additionally, the subject column categorizes the article into various domains, such as politics or world events, offering insights into the potential influence of the subject matter on the authenticity of the news.

The dataset is characterized by its voluminous collection of articles, ensuring a rich and diverse pool for analysis. This breadth of subjects enhances the dataset's versatility, making it applicable across different news domains. Its primary utility is aiding the development and testing of machine learning models for fake news detection. It supports a range of analytical approaches, from text classification to sentiment analysis and linguistic pattern recognition.

However, users of this dataset should be cognizant of its sourcing and the timeframe of the news articles to maintain the model's relevance and adaptability to current trends in news production. It is also crucial to evaluate the dataset for potential biases or imbalances that could impact the effectiveness and generalizability of models developed using this resource.

4. Results

In this research, a suite of performance metrics is employed to evaluate the efficacy of the machine learning model in classifying news articles as 'Fake' or 'True'. These metrics include accuracy, precision, recall, F1-score, and the area under the receiver operating characteristic (ROC) curve (AUC). Accuracy measures the proportion of total correct predictions (both true positives and true negatives) to the overall dataset, providing a general sense of the model's overall performance. Precision assesses the model's exactness by measuring the ratio of true positives to the sum of true positives and false positives, reflecting the model's ability to return relevant results. Recall, or sensitivity, evaluates the model's completeness by calculating the ratio of true positives to the sum of true positives and false negatives, indicating the model's capability to find all relevant instances. The F1-score is the harmonic mean of precision and recall, offering a single measure for the balance between them, which is especially useful in the context of imbalanced datasets. Lastly, the AUC represents the degree to which the model can distinguish between the classes across all thresholds, with a higher AUC indicating better model performance. Together, these metrics provide a comprehensive picture of the model's classification abilities, highlighting its strengths and areas for improvement.

The results of the fake news classifications are presented in Table 1.

Table 1
Classification report

	Precision	Recall	F1-score	Support
0	0,82	0,83	0,83	733
1	0,77	0,75	0,76	554
Accuracy			0,80	1287
Macro AVG	0,80	0,79	0,79	1287
Macro AVG	0,80	0,80	0,80	1287

These results indicate a relatively strong performance of the model in distinguishing between the two classes, with a slightly better performance in identifying the real news class than the fake news class. The overall accuracy and macro-average scores demonstrate the model's robustness in handling this classification task.

Figure 1 shows the model's epoch loss, which illustrates a machine learning model's decline in training loss across 15 epochs. As the epochs progress, the line sharply descends, showing a rapid reduction in training loss, which suggests that the model is quickly learning from the training data. The steepest decline occurs between epochs 1 and 4, after which the curve begins to flatten, indicating a slower rate of improvement. By epoch 5, the loss has reduced significantly and continues to decrease at a marginal rate, plateauing as it approaches epoch 15. This flattening of the curve suggests that the model is reaching a point of convergence, where additional training yields little to no improvement in reducing the loss, indicating that the model may have reached its optimal performance on the training dataset. Overall, the graph demonstrates a successful training process where the model's performance, in terms of its ability to minimize the loss function, improves substantially and consistently throughout the training epochs.

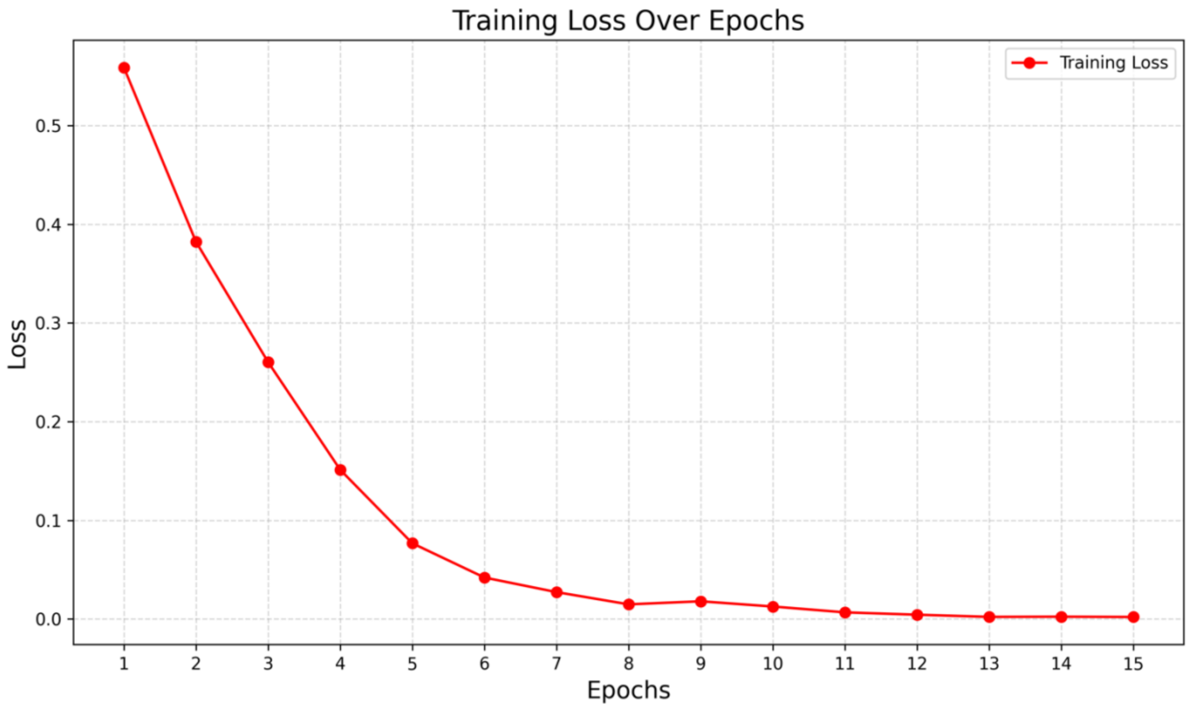

Figure 1: Training loss over epochs

Figure 2 displays a Precision-Recall Curve, a graphical representation commonly used to evaluate the performance of a binary classifier. As recall increases, the precision decreases, suggesting that as the model identifies a higher proportion of positive instances (true positives), it also starts to include more false positives, thus reducing precision. The area under the curve (AUC) is annotated as 0.82, indicating a high level of overall model performance.

Figure 3 shows a Receiver Operating Characteristic (ROC) Curve, which is a graphical plot that illustrates the diagnostic ability of a binary classifier system as its discrimination threshold is varied. The ROC curve shows that the proposed model provides a good separability measure between the two classes. The area under the ROC curve is annotated as 0.87, which quantifies the overall ability of the classifier to discriminate between the positive and negative classes.

Figure 4 shows a graph that tracks a model's change in validation accuracy across 15 epochs. The graph shows significant variability in accuracy from one epoch to the next. The accuracy peaks at certain epochs (near epochs 1, 5, and 11), reaching highs just above 0.805. At the same time, at other points, it dips sharply, with the most notable dip occurring at epoch 9, where accuracy drops below 0.795. The overall trend does not show a consistent improvement or decline in accuracy as training progresses.

Figure 5 shows the confusion matrix, which suggests that the model is quite effective at detecting fake news, as seen by the high TP rate for 'Fake'. However, it also shows room for improvement, especially in reducing the number of FN, where fake news is missed, and FP, where true news is incorrectly flagged as fake. The model's precision, recall, and overall accuracy can be fine-tuned to balance the performance across both classes better.

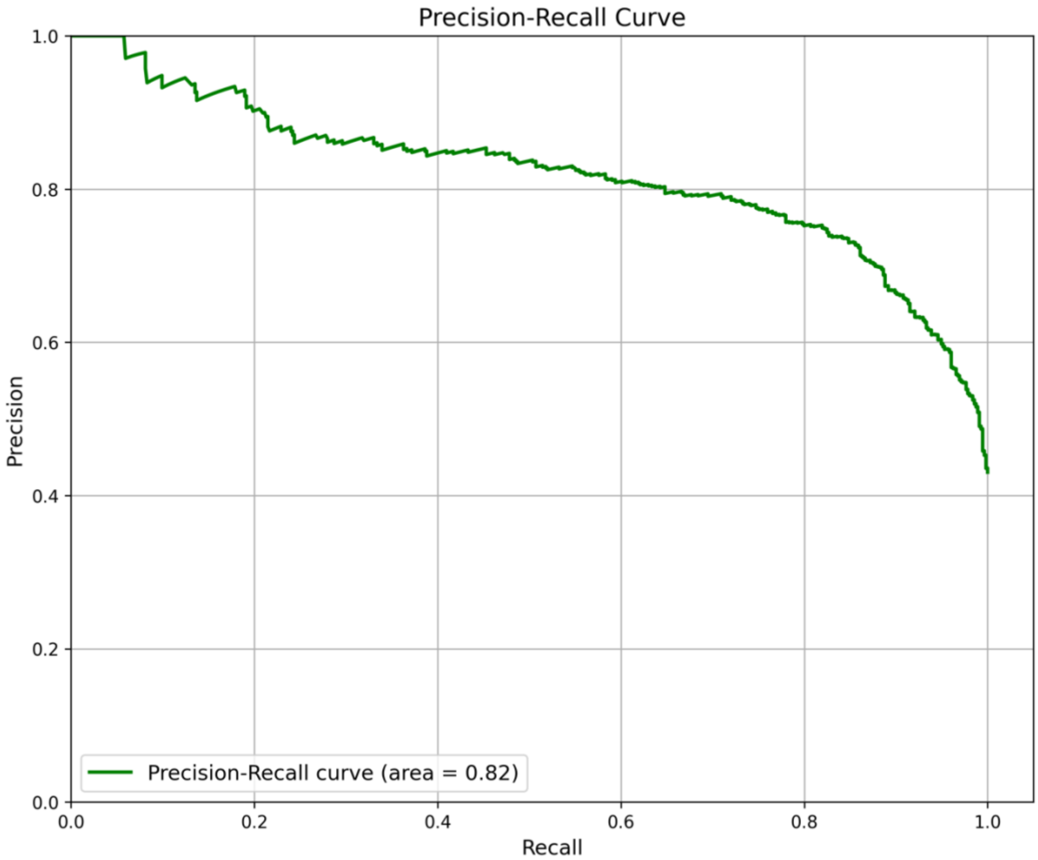

Figure 2: Precision-Recall Curve

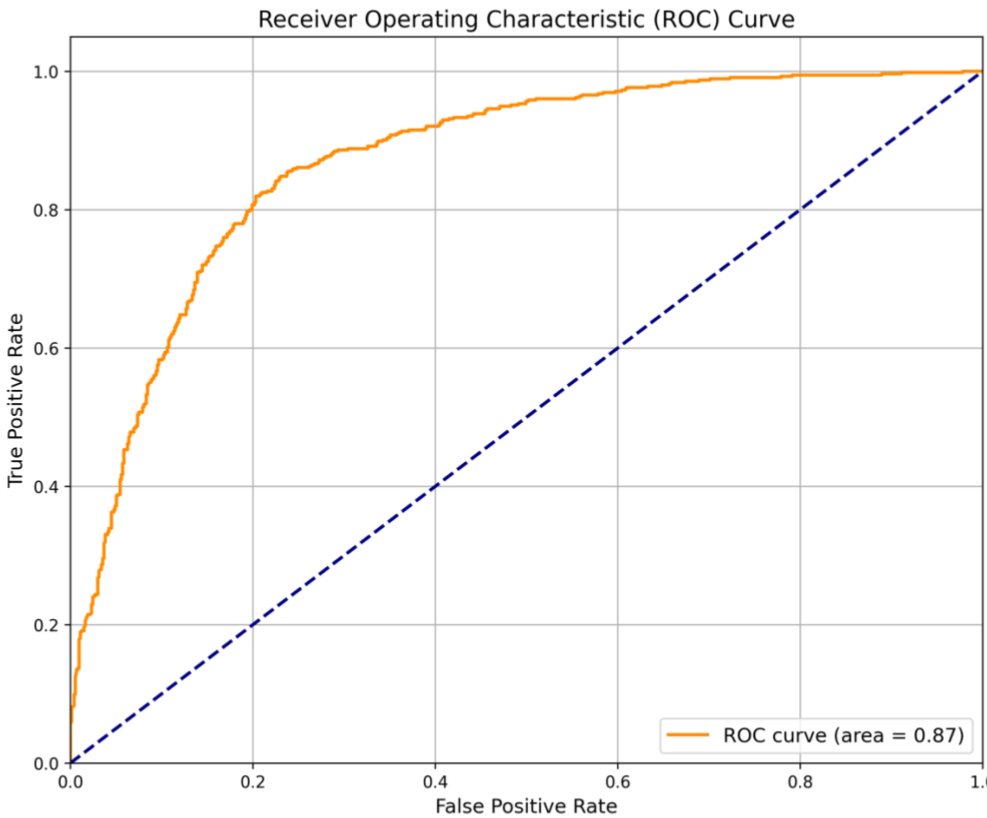

Figure 3: ROC Curve

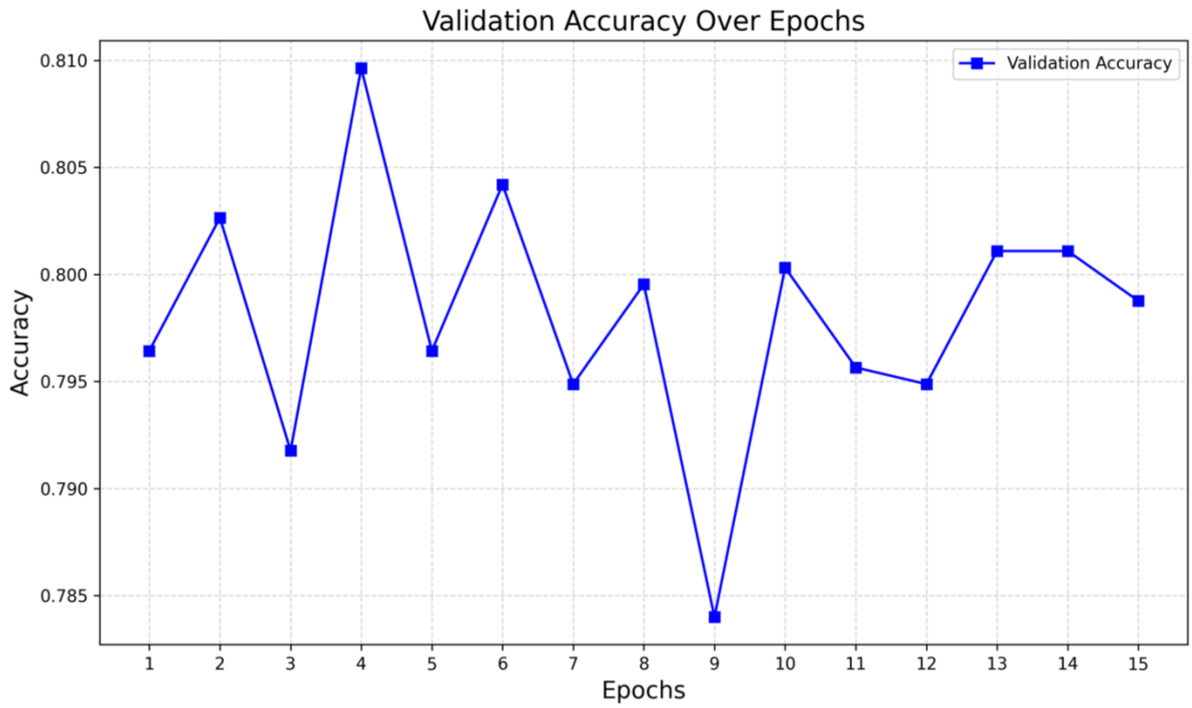

Figure 4: Validation accuracy over epochs

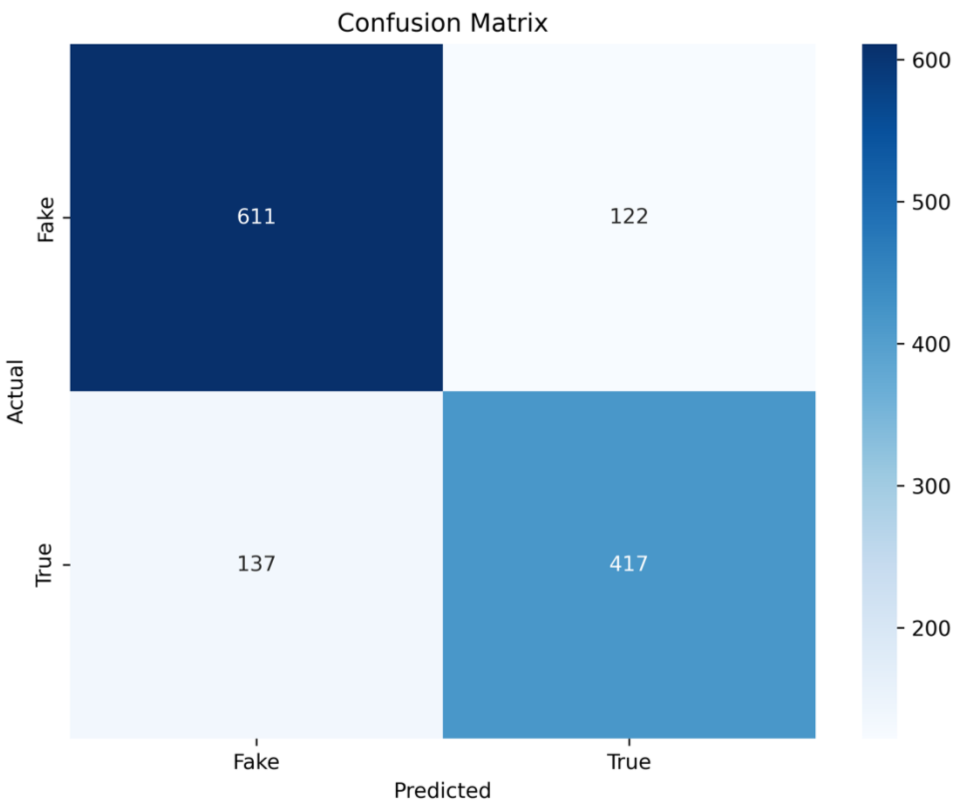

Figure 5: Confusion matrix

5. Discussion

Fake news has emerged as one of contemporary society's most pervasive and insidious challenges, fostering misinformation that percolates through the digital ecosystem with detrimental consequences. Its relevance extends beyond mere academic discourse, touching the very fabric of democracy, influencing public opinion, and shaping political landscapes across the globe [14]. The dissemination of false information has been implicated in swaying election outcomes, inciting social unrest, and undermining public trust in institutions. Consequently, developing effective tools to detect and mitigate the spread of fake news is not only a technological imperative but also a societal one, demanding urgent and concerted efforts.

In this milieu, the intersection of NLP and machine learning presents a fertile ground for innovation. The advent of models such as BERT, which can process the nuances of human language with remarkable depth, has opened new avenues for identifying fake news [15]. The significance of this research lies in its potential to bolster the algorithms underpinning social media platform information curation, thereby enhancing the quality and reliability of shared content. As social media becomes increasingly ingrained in the daily lives of billions, the responsibility to curate content responsibly becomes paramount. Thus, the practical applications of this research resonate with the urgent need to safeguard the information landscape against the proliferation of falsehoods.

Furthermore, the endeavor to refine fake news detection models also carries profound implications for the field of NLP. It propels the domain towards more sophisticated models that can grasp not just the semantic strings of language but also its subtle cues and implicit meanings, which are often exploited in the craft of misinformation. This research contributes to the body of knowledge in NLP. It sets the stage for further explorations that could redefine the boundaries of what artificial intelligence can achieve in the service of truth.

The research explored fake news detection using advanced machine learning techniques, yielding insightful results that warrant a nuanced discussion. The validation accuracy attained was 79.88%, a respectable figure suggesting that the model has a strong potential for identifying the veracity of news articles. This accuracy provides a baseline for the model's general performance but does not elucidate the intricacies of its predictive power across different classes of news.

Further dissecting the results, we observed a precision of 82% for the 'Fake' class and 77% for the 'True' class. These figures indicate a commendable level of exactness, particularly in identifying fake news. While slightly lower, the precision for the 'True' class is still substantial, though it hints at a more conservative approach by the model in classifying news as true. The recall rates of 83% for 'Fake' and 75% for 'True' reveal that the model is more proficient in detecting fake news articles than in catching all genuine news articles, corroborating the insights drawn from precision metrics.

The F1-score, an amalgamation of precision and recall, stands at 83% for 'Fake' and 76% for 'True', indicating a well-balanced performance, especially for the 'Fake' news. This balance is pivotal in scenarios where the relevance of the result (precision) and the ability to identify all relevant items (recall) are crucial.

The ROC curve's AUC of 0.87 reinforces the model's discriminative capacity. An AUC closer to 1 implies that the model has a high true positive rate relative to the false positive rate, showcasing its capability to distinguish between 'Fake' and 'True' news. Such a high AUC indicates a robust model across various thresholds, providing flexibility in model deployment based on the specific needs and trade-offs in real-world scenarios.

However, the confusion matrix provides an additional layer of insight, revealing a discrepancy in the model's ability to classify 'Fake' versus 'True' news. While the true positive rate for 'Fake' news is high (611 out of 733), the model also misclassified a considerable number of 'True' articles as 'Fake' (137 out of 554), which could have significant implications in the context of news dissemination and consumption.

The performance metrics and the confusion matrix suggest an effective model. The results indicate a commendable ability to flag fake news, which is the primary objective in the context of the spread of misinformation. Nevertheless, misclassifying true news as fake is a non-trivial issue that could undermine the credibility of legitimate news sources if the model were deployed in a real-world setting.

Future work could address these limitations by exploring ensemble methods to stabilize predictions across epochs, investigating alternative architectures or feature sets to improve classification, especially for the 'True' class, and implementing techniques to mitigate potential biases in the training data.

6. Conclusions

The quest to fortify the veracity of information in our increasingly digital world has never been more critical, and this research has made strides in utilizing machine learning to discern the genuine from the fraudulent in news dissemination. Our investigation deployed a BERT model, renowned for its contextual language processing capabilities, to tackle the pervasive challenge of fake news.

Scientifically, the research illuminated the efficacy of BERT in a domain that is as dynamic as it is crucial: detecting misinformation. We presented a comprehensive analysis of the model's performance, leveraging various metrics demonstrating a robust capability to classify news articles accurately. The application of BERT for fake news detection is a testament to the model's versatility and advanced understanding of language nuances, a significant leap from traditional vector space models.

From a practical standpoint, this research offers a tangible advancement in the tools available for media platforms and news consumers. With an overall validation accuracy of nearly 80% and an AUC of 0.87, the model is a potent instrument for critically examining news authenticity. The precision and recall metrics further delineate its practical utility in minimizing the spread of fake news. The model's ability to effectively distinguish between 'Fake' and 'True' classes bears significant implications for social media platforms and other news aggregators seeking to preserve the integrity of their content.

Future work will stabilize the model's performance and enhance its discernment abilities, potentially through ensemble methods or alternative training techniques.

This research contributes a novel approach to fake news detection, merging state-of-the-art computational linguistics with machine learning. It stands as a beacon for future endeavors aiming to safeguard information integrity and enrich the reliability of news consumed by the public.

References

[1] M. Mosleh and D. G. Rand, “Measuring exposure to misinformation from political elites on Twitter,” *Nature Communications*, vol. 13, no. 1, Nov. 2022, doi: 10.1038/s41467-022-34769-6.

[2] S. Chen, L. Xiao, and A. Kumar, “Spread of misinformation on social media: What contributes to it and how to combat it,” *Computers in Human Behavior*, vol. 141, no. 141, p. 107643, Apr. 2023, doi: 10.1016/j.chb.2022.107643.

[3] V. Tolz and S. Hutchings, “Truth with a Z: disinformation, war in Ukraine, and Russia’s contradictory discourse of imperial identity,” *Post-Soviet Affairs*, vol. 39, no. 5, pp. 1–19, Apr. 2023, doi: 10.1080/1060586x.2023.2202581.

[4] S. Blank, “Russia, China, and Information War against Ukraine,” *The Journal of East Asian Affairs*, vol. 35, no. 2, pp. 39–72, 2022.

[5] J. Mandić and D. Klarić, “Case Study of the Russian Disinformation Campaign During the War in Ukraine – Propaganda Narratives, Goals and Impacts,” *National security and the future*, vol. 24, no. 2, pp. 97–139, Jul. 2023, doi: 10.37458/nstf.24.2.5.

[6] S.-D. Bachmann, D. Putter, and G. Duczynski, “Hybrid warfare and disinformation: A Ukraine war perspective,” *Global Policy*, vol. 14, no. 5, Aug. 2023, doi: 10.1111/1758-5899.13257.

[7] J. Devlin, M.-W. Chang, K. Lee, K. Google, and A. Language, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” 2019.

[8] J. J. Van Bavel, E. A. Harris, P. Pärnamets, S. Rathje, K. C. Doell, and J. A. Tucker, “Political Psychology in the Digital (mis)Information age: A Model of News Belief and Sharing,” *Social Issues and Policy Review*, vol. 15, no. 1, pp. 84–113, Jan. 2021, doi: 10.1111/sipr.12077.

[9] D. Halpern, S. Valenzuela, J. Katz, and J. P. Miranda, “From Belief in Conspiracy Theories to Trust in Others: Which Factors Influence Exposure, Believing and Sharing Fake News,” *Social Computing and Social Media. Design, Human Behavior and Analytics*, vol. 11578, pp. 217–232, 2019, doi: 10.1007/978-3-030-21902-4_16.

[10] K. A. Qureshi, R. A. S. Malick, M. Sabih, and H. Cherifi, “Complex Network and Source Inspired COVID-19 Fake News Classification on Twitter,” *IEEE Access*, vol. 9, pp. 139636–139656, 2021, doi: 10.1109/access.2021.3119404.

[11] S. Garg and D. Kumar Sharma, “Linguistic Features based Framework for Automatic Fake News Detection,” *Computers & Industrial Engineering*, vol. 172, p. 108432, Jul. 2022, doi: 10.1016/j.cie.2022.108432.

[12] M. Coscia and L. Rossi, “Distortions of political bias in crowdsourced misinformation flagging,” *Journal of The Royal Society Interface*, vol. 17, no. 167, p. 20200020, Jun. 2020, doi: 10.1098/rsif.2020.0020.

[13] K. Yadav, “Fake-Real News,” *www.kaggle.com*, 2020. https://www.kaggle.com/datasets/techykajal/fakereal-news (accessed Oct. 13, 2023).

[14] K. Ognyanova, D. Lazer, R. E. Robertson, and C. Wilson, “Misinformation in action: Fake news exposure is linked to lower trust in media, higher trust in government when your side is in power,” *Harvard Kennedy School Misinformation Review*, vol. 1, no. 4, pp. 1–19, Jun. 2020, doi: 10.37016/mr-2020-024.

[15] A. Herasymova, D. Chumachenko, and H. Padalko, “Development of intelligent information technology of computer processing of pedagogical tests open tasks based on machine learning approach,” *CEUR Workshop Proceedings*, vol. 2631, pp. 121–131, 2020.