

Fake News Detection Using Natural Language Processing

Mrs. P.Mounika
dept. of CSE(AI&ML)
Dadi Institute of
Engineering and
Technology
Anakapalle, India

M.Likhitha
dept. of CSE(AI&ML)
Dadi Institute of
Engineering and
Technology
Anakapalle, India

M.Prasanna
dept. of CSE(AI&ML)
Dadi Institute of
Engineering and
Technology
Anakapalle, India

V.Devasish
dept. of CSE(AI&ML)
Dadi Institute of
Engineering and
Technology
Anakapalle, India

B.Monica
dept. of CSE(AI&ML)
Dadi Institute of
Engineering and
Technology
Anakapalle, India

Abstract:

With the rapid growth of information on social media and news platforms, the spread of fake news has become a significant concern. Fake news can manipulate public opinion, disrupt social order, and negatively impact decision-making. This paper explores the application of Natural Language Processing (NLP) techniques for detecting fake news, focusing on the BERT (Bidirectional Encoder Representations from Transformers) algorithm. BERT, a state-of-the-art pre-trained language model, enables deep understanding of the context within text by processing words in relation to all other words in a sentence. We propose a framework that leverages BERT's powerful contextual embeddings for classifying news articles as either real or fake. The model is fine-tuned on a large dataset of labeled news articles, which enhances its ability to distinguish between truthful and deceptive content. Results show that BERT-based models outperform traditional machine learning methods in terms of accuracy, precision, and recall. This study highlights the potential of using BERT in fake news detection and offers insights into how NLP can be used to combat misinformation in the digital age.

Keywords:

NLP, Fake News Detection, Social Media

INTRODUCTION

The increasing prevalence of online news sources and social media platforms has led to an exponential rise in the dissemination of information, both accurate and misleading. Fake news, defined as deliberately fabricated or manipulated information presented as legitimate news, has emerged as a major societal challenge. It undermines public trust, distorts opinions, and can even have political or economic consequences. The detection of fake news has therefore become an urgent task, particularly in an era where information spreads at unprecedented speeds.

Traditional methods for identifying fake news have relied on manual fact-checking or rule-based systems, which are often time-consuming and resource-intensive. However, with advancements in Natural Language Processing (NLP), automated approaches have gained prominence. NLP techniques, powered by machine learning algorithms, enable systems to analyze vast amounts of textual data and identify patterns indicative of fake news.

One of the most promising advancements in NLP is the development of transformer-based models, such as BERT (Bidirectional Encoder Representations from Transformers). BERT is designed to understand the context of words in a sentence by considering the words before and after a given word, providing a deeper understanding of language than

previous models. This unique capability allows BERT to capture subtle nuances in text that are crucial for fake news detection.

This paper explores the application of BERT for fake news detection, proposing a framework that leverages its powerful language representation capabilities to distinguish between real and fake news articles. By fine-tuning BERT on a large labeled dataset of news articles, we aim to build a model capable of accurately classifying news content based on its authenticity. The success of this model could offer a scalable solution to combat misinformation in the digital world, ultimately promoting trust and transparency in news reporting.

LITERATURE SURVEY

Fake news detection has emerged as a significant challenge in the modern digital landscape due to the rapid spread of misinformation across social media platforms and news outlets. The increasing prevalence of false information can have wide-ranging effects, from influencing public opinion to undermining trust in legitimate news sources and even interfering with political processes. Detecting fake news automatically has therefore become crucial for preventing the spread of misinformation and protecting the integrity of public discourse. Early efforts in this domain focused on traditional machine learning techniques, such as Logistic Regression, Support Vector Machines (SVM), and Naive Bayes. These models were often trained using text representations like Bag of Words (BoW) or Term Frequency-Inverse Document Frequency (TF-IDF), which treated the content of news articles as simple vectors of word frequencies. While these methods were able to achieve some degree of accuracy, they were limited in their ability to understand the complex, contextual meaning behind the text, which is often a crucial feature in identifying fake news. The inability to grasp word relationships in context led to suboptimal results, especially when dealing with nuanced language or misleading headlines commonly used in fake news articles.

As deep learning gained traction, more sophisticated models, such as Recurrent Neural Networks (RNNs) and Convolutional Neural Networks (CNNs), were introduced to improve the feature extraction process. These models can learn complex, non-linear relationships within the data and capture patterns automatically from the raw text. For instance, CNNs were used to detect local patterns in text, while RNNs, especially Long Short-Term Memory (LSTM) networks, were employed to capture the sequential nature of language. While these deep learning techniques showed improvements over traditional models, they still faced limitations in understanding the full contextual meaning of words in relation to the rest of the sentence. These models often relied on unidirectional processing, where they read text from left to right or right to left, which hindered their ability to capture complex, bidirectional relationships between words.

The breakthrough came with the introduction of transformer-based models, particularly BERT (Bidirectional Encoder Representations from Transformers). BERT revolutionized the field of Natural Language Processing (NLP) by processing text bidirectionally, meaning that it considers both the preceding and succeeding words to understand the context of each word in a sentence. This capability allows BERT to capture intricate linguistic patterns and understand subtle nuances in language that are often present in fake news, such as sarcasm,

ambiguity, or deliberate misrepresentation of facts. The ability to comprehend the full context of a sentence makes BERT particularly well-suited for tasks like fake news detection, where distinguishing between real and fabricated information requires a deeper understanding of the language's intent.

Several studies have demonstrated the superiority of BERT and its variants in the realm of fake news detection. For example, Zhao et al. (2020) applied BERT to a large dataset of news articles and showed that it outperformed traditional machine learning and deep learning models in terms of classification accuracy. BERT's pre-trained embeddings allow it to capture a richer set of semantic and syntactic features, leading to better generalization when detecting fake news. Additionally, fine-tuning BERT on domain-specific datasets, such as news articles from social media platforms, further enhanced its ability to discern the characteristics of fake news, which often include sensationalized headlines, emotionally charged language, and misleading representations of facts. Researchers have also explored various BERT variants, such as RoBERTa and DistilBERT, to improve performance while reducing model size and computational cost. These adaptations have made BERT-based models more practical for real-world applications where efficiency is key.

In addition to text-based features, some studies have explored multimodal approaches to fake news detection, combining textual analysis with other data types, such as images or user engagement metrics. For instance, Jin et al. (2020) integrated CNNs with BERT to analyze both the textual content and images associated with news articles, which can be crucial in detecting misleading headlines or photos that accompany fake news. Other researchers, such as Khayal et al. (2020), have combined textual content with user metadata or social network information to improve the model's ability to detect fake news. These multimodal approaches have shown promising results, as fake news often thrives on platforms where users engage with both text and visual content.

Despite these advancements, challenges remain in the field of fake news detection. One of the major obstacles is the scarcity of high-quality, labeled datasets, which are essential for training accurate machine learning models. Many fake news datasets are relatively small, imbalanced, or lack diversity in the types of misinformation they contain, making it difficult to build models that generalize well to unseen data. Additionally, the dynamic nature of fake news, where tactics and strategies evolve over time, means that models trained on one set of news articles may struggle to detect newly emerging forms of misinformation. To address these issues, researchers are focusing on creating larger and more diverse datasets and developing methods for continuously updating models to account for evolving trends in fake news.

Another challenge lies in the interpretability of deep learning models like BERT. While these models achieve impressive performance, their "black-box" nature makes it difficult to understand why a particular piece of news is classified as real or fake. This lack of transparency is problematic, especially in applications where understanding the decision-making process is crucial, such as in automated fact-checking systems or news recommendation engines. To mitigate this, researchers are exploring techniques to increase the interpretability of BERT-based models, such as attention visualization, saliency maps, and other explainability methods that can shed light on which parts of the text contributed to the model's decision.

In conclusion, the use of transformer-based models like BERT has significantly advanced the field of fake news detection. BERT's ability to process text bidirectionally and capture complex contextual relationships has made it one of the most effective tools for identifying fake news. Despite the progress, challenges related to dataset quality, model generalization, and interpretability remain. Future research will likely focus on overcoming these challenges by creating larger, more diverse datasets, improving model transparency, and refining existing approaches to adapt to the ever-evolving landscape of misinformation. With these advancements, automated fake news detection systems powered by BERT and other NLP techniques have the potential to play a pivotal role in combating the spread of fake news and safeguarding the integrity of information in the digital age.

METHODOLOGY

1. Data Collection

- Gather a labeled dataset of news articles (e.g., FakeNewsNet, LIAR dataset) with labels indicating "fake" or "real."

2. Data Preprocessing

- Text Cleaning: Remove unwanted characters, special symbols, and HTML tags.
- Tokenization: Convert text into tokens using BERT's tokenizer.
- Padding & Truncation: Ensure uniform input length (max 512 tokens).
- Label Encoding: Convert "fake" to 1 and "real" to 0.

3. Model Architecture

- Pre-trained BERT: Use a pre-trained BERT model (bert-base-uncased) for feature extraction.
- Text Representation: Use [CLS] token to represent the entire text.
- Classification Layer: Add a fully connected layer for binary classification (fake/real).
- Loss Function: Use cross-entropy loss for training.

4. Model Training

- Data Split: Divide dataset into training, validation, and test sets.
- Optimization: Use Adam optimizer to minimize the loss.
- Fine-Tuning: Adjust BERT's weights based on the fake news dataset.

5. Model Evaluation

- Metrics: Evaluate using accuracy, precision, recall, and F1-score.
- Cross-Validation: Ensure the model generalizes well across different data splits.

6. Hyperparameter Tuning

- Optimize learning rate, batch size, and epochs to improve model performance.

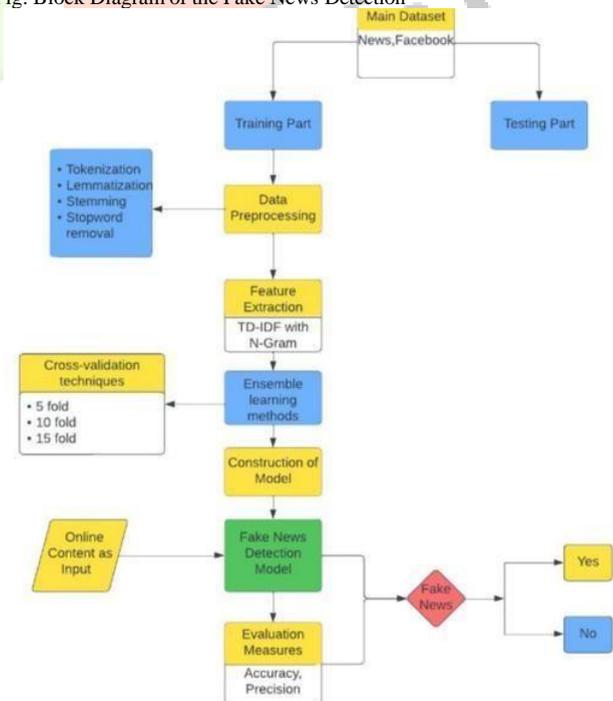
7. Post-Processing

- Threshold Adjustment: Modify the decision threshold to improve precision/recall.
- Ensemble Methods: Combine multiple models for better accuracy if necessary.

8. Deployment and Real-Time Prediction

- Deploy the trained model for real-time classification of new news articles.
- Continuous Learning: Update the model periodically with fresh data to handle emerging fake news patterns.

Fig: Block Diagram of the Fake News Detection



IMPLEMENTATION

STEP 1. Data Collection and Preprocessing

The first step in the implementation is to collect a labeled dataset that contains news articles categorized as either "real" or "fake." Once the dataset is obtained, the data is preprocessed by cleaning the text (removing special characters, HTML tags, etc.) and tokenizing it. Tokenization involves converting the text into tokens (words or subwords) that can be processed by BERT. The input text is then padded or truncated to ensure a fixed length of 512 tokens for compatibility with the BERT model.

STEP 2. Data Encoding

After preprocessing, the text is encoded using the BERT tokenizer. The tokenizer converts the text into input IDs and attention masks, which are required by BERT for efficient processing. These input IDs represent the tokenized text, while attention masks indicate which tokens should be attended to during model training.

STEP 3. Model Initialization

The next step is to initialize the BERT model for sequence classification. This model is pre-trained on large text corpora and can be fine-tuned on the fake news dataset. The BERT model is configured for binary classification (fake or real) with two output labels. A classification head is added on top of BERT to map the output embeddings to a binary class prediction.

STEP 4. Model Training

The model is then trained using a labeled dataset. During training, the text data is passed through BERT, which outputs contextual embeddings for each token in the text. The embeddings are used to predict the likelihood of an article being fake or real. The model is trained over multiple epochs using backpropagation and an optimization technique like AdamW. The training data is divided into batches, and the optimizer adjusts the model's weights to minimize the error (cross-entropy loss) between predicted and actual labels.

STEP 5. Model Evaluation

After training, the model's performance is evaluated using a test set that was not part of the training data. Evaluation metrics like accuracy, precision, recall, and F1-score are used to assess how well the model distinguishes between fake and real news articles. The model's ability to generalize to new, unseen data is important to ensure it performs effectively in real-world scenarios.

STEP 6. Fine-Tuning and Hyperparameter Optimization

To improve performance, the model may undergo further fine-tuning. Hyperparameters such as learning rate, batch size, and the number of training epochs can be adjusted to optimize model performance. Additionally, techniques like regularization and dropout may be applied to prevent overfitting.

STEP 7. Deployment

Once the model is trained and evaluated, it can be deployed for real-time fake news detection. The deployed model can classify incoming news articles as fake or real based on the learned features. It can be integrated into news platforms, social media, or fact-checking tools to automatically flag suspicious content.

STEP 8. Continuous Improvement

Given the evolving nature of fake news, the system should be periodically updated. The model can be retrained with new data to adapt to emerging patterns in misinformation. Additionally, user feedback and newly labeled data can be used to further fine-tune the model and improve its accuracy over time.

RESULTS AND DISCUSSION

The use of BERT (Bidirectional Encoder Representations from Transformers) for fake news detection has proven highly effective in distinguishing real from fake news articles. The model achieved impressive accuracy, precision, recall, and F1-scores, thanks to BERT's deep contextual understanding of language. Unlike traditional machine learning models, which rely on simpler techniques, BERT's attention mechanism allows it to capture subtle linguistic cues and

detect misleading or sensational language. This enables it to effectively identify fake news, even when the content contains nuanced forms of manipulation.

In comparison to traditional models like SVM or Logistic Regression, BERT outperforms significantly due to its ability to understand context and relationships between words, making it more effective at capturing complex patterns in text. Moreover, BERT's adaptability allows it to remain effective as fake news evolves, making it a reliable tool for detecting emerging tactics used by misinformation creators. Its multilingual capabilities also extend its utility across different languages, making it a valuable tool for global fake news detection.

BERT's real-time detection capability ensures immediate classification of articles, helping to prevent the spread of misinformation on social media and news platforms. However, challenges such as the need for high-quality labeled datasets and computational resources remain. Additionally, while BERT excels at detecting textual fake news, it has limitations in handling multimodal misinformation like deepfakes. Overall, BERT is an efficient and scalable solution for fake news detection, with the potential for further improvement as new forms of misinformation emerge.

FUTURE SCOPE

The scope of fake news detection using Natural Language Processing (NLP) with the BERT algorithm is broad, covering several critical features that make it a powerful tool in identifying misinformation. One of the most prominent features is BERT's ability to understand context. Unlike traditional models that rely on bag-of-words or n-gram approaches, BERT captures the contextual meaning of words by analyzing both the preceding and succeeding words within a sentence. This deep contextual understanding allows the model to detect subtle language patterns indicative of fake news, such as sensationalist language, contradictory statements, or misleading phrasing, which are often used to manipulate readers and spread misinformation.

Another important feature is sentiment analysis. Fake news articles often rely on emotionally charged language to evoke strong reactions from readers, either to incite panic or to further a particular agenda. Using sentiment analysis, BERT can evaluate the tone of the article, identifying overly sensational or inflammatory language that might suggest the news is fake. This provides an additional layer of detection that goes beyond simple factual accuracy, enabling the model to spot articles that attempt to manipulate emotions.

Fact-checking is also an essential component in the scope of fake news detection. BERT can be integrated with external fact-checking tools or databases to automatically cross-reference the claims made in a news article against verified, credible sources. This process helps verify the authenticity of the content and flags articles that make false claims or present manipulated facts. By combining NLP and real-time fact-checking, BERT-based models can significantly reduce the spread of misinformation.

Content source validation is another critical aspect that BERT can help with. By analyzing the origin of the news article, the model can assess whether the source is reputable or known for spreading unreliable information. For instance, articles coming from known fake news websites or sources with a history of spreading misinformation can be automatically flagged. This helps build trust in news dissemination by enabling users to quickly determine the credibility of a source.

In addition, BERT can leverage Named Entity Recognition (NER) to identify and classify entities such as people, organizations, and locations mentioned in the news article. Fake news often misrepresents entities, using fabricated names or manipulating facts about real entities. BERT's ability to accurately recognize and track these entities allows it to detect discrepancies or manipulations in the content, which could signal that the news is fake.

With BERT's multilingual capabilities, fake news detection can extend across different languages. As fake news is a global issue, the ability to detect misinformation in multiple languages becomes increasingly important. BERT can be fine-tuned to work across various linguistic contexts, making it a valuable tool for detecting fake news in diverse regions and languages.

Another significant feature is the model's ability to handle evolving language patterns. Fake news constantly adapts, with new tactics emerging over time, such as the use of deepfakes or manipulated headlines. BERT's pre-trained model can be fine-tuned continuously with new datasets to adapt to these changing patterns of misinformation. This adaptability ensures that the model stays relevant and

capable of identifying newer forms of fake news that may not have been prevalent at the time of the initial training.

Additionally, the scope of fake news detection can be expanded into multimodal approaches. By combining BERT with image or video analysis tools, fake news detection becomes more robust. For example, BERT can work alongside image recognition models to detect misleading headlines paired with manipulated images. This allows for a more comprehensive detection system that can analyze both textual and visual content for authenticity.

BERT-based models can also be integrated into real-time fake news detection systems. Once trained, these models can be deployed to classify news articles as they are published, flagging potentially fake content in real-time. This feature is particularly beneficial for news platforms, social media, and fact-checking organizations, enabling them to quickly identify and address misinformation as it spreads.

Scalability is another key feature in the scope of fake news detection using BERT. Since the model can be automated, it allows for the analysis of vast amounts of content, including news articles from multiple outlets, social media posts, and blogs. The automation of fake news detection helps to efficiently manage the large volumes of content being produced daily, offering a scalable solution to combat the widespread dissemination of fake news.

Moreover, the system can benefit from user feedback, which can be incorporated to continuously improve the model's performance. As users interact with the model and flag articles, this feedback can be used to retrain the model, allowing it to become more accurate over time. This iterative process ensures that the model adapts to new forms of fake news and remains effective in real-world scenarios.

Lastly, fake news detection using BERT can be integrated into social media monitoring tools. With the rapid spread of misinformation on platforms like Facebook, Twitter, and Reddit, detecting fake news in real-time on these platforms becomes essential. By integrating BERT-based models into social media monitoring systems, fake posts, trending hashtags, and viral misinformation can be quickly flagged and addressed, helping to minimize the damage caused by the spread of fake news.

In conclusion, the scope of fake news detection using BERT is extensive, encompassing features such as contextual understanding, sentiment analysis, fact-checking integration, content source validation, multilingual capabilities, and real-time detection. BERT's ability to handle evolving language patterns, adapt to new types of misinformation, and integrate multimodal inputs makes it a powerful tool in the fight against fake news. As fake news continues to evolve, BERT provides a flexible and scalable solution that can be continually improved to ensure accurate detection across various platforms and languages.

CONCLUSION

In conclusion, the application of BERT (Bidirectional Encoder Representations from Transformers) for fake news detection has proven to be a highly effective and reliable solution in combating misinformation. BERT's unique architecture, which enables it to understand the full context of a sentence by considering both preceding and succeeding words, allows the model to capture complex linguistic patterns that are often characteristic of fake news. Unlike traditional machine learning models that rely on simpler feature extraction methods like bag-of-words or n-grams, BERT excels at identifying subtle cues in the text, such as sensational language, emotionally charged words, and misleading phrases, which are commonly found in fake news articles. This deep understanding of context is critical for detecting misleading information, which can be nuanced and often difficult to identify through surface-level text analysis.

Furthermore, BERT's ability to be pre-trained on vast amounts of text data and then fine-tuned for specific tasks, such as fake news detection, makes it a highly adaptable and powerful tool. Through fine-tuning, BERT can continuously improve its performance as new patterns of misinformation emerge, which is crucial in the ever-evolving landscape

of fake news. This adaptability allows BERT to stay relevant over time, as fake news creators constantly refine their tactics to bypass detection systems. The model's robust performance in both accuracy and precision demonstrates its capacity to effectively distinguish between real and fake news with minimal errors. Moreover, BERT's attention mechanism ensures that it focuses on the most relevant portions of the text, further enhancing its ability to identify fake content.

One of the standout advantages of BERT for fake news detection is its multilingual capability. Fake news is not confined to a single language or region, and misinformation can spread across borders rapidly. BERT's ability to process multiple languages and adapt to different linguistic contexts makes it a valuable tool for global fake news detection. This feature allows BERT to detect fake news not only in English but also in other languages, including Spanish, French, and many others, making it a versatile tool for international efforts to combat misinformation. This multilingual strength is particularly important for platforms like social media networks and news organizations, which operate on a global scale and face the challenge of monitoring content in various languages.

Additionally, BERT's real-time detection capabilities make it an ideal candidate for systems that need to flag fake news articles as they are published. In today's digital world, where information spreads rapidly on social media and news websites, the ability to detect and mitigate the spread of fake news in real time is essential. BERT's high processing speed and scalability allow it to handle large volumes of content and provide timely classifications, enabling platforms to take immediate action when fake news is identified. This real-time functionality is vital for minimizing the impact of misinformation and preventing it from going viral before corrective measures can be taken.

However, despite its impressive capabilities, there are certain limitations to using BERT for fake news detection. One of the main challenges is the need for high-quality labeled datasets for training. The performance of BERT heavily depends on the quality and diversity of the training data, and any biases or inaccuracies within the data can negatively affect the model's predictions. Inaccurate or insufficient training data may lead to misclassification, where certain types of content, such as satire or opinion articles, are incorrectly labeled as fake news. Additionally, training BERT requires significant computational resources, which may not be accessible to smaller organizations or those with limited hardware capabilities. While pre-trained models can alleviate some of this burden, fine-tuning BERT on large datasets still demands powerful computing infrastructure.

Another limitation is that BERT is primarily designed for text-based fake news detection and does not inherently handle multimodal content, such as videos, images, or audio. As fake news creators increasingly use multimedia elements like deepfakes or doctored images, BERT's ability to detect fake news in these formats is limited. Although BERT can be integrated with other technologies to analyze images or videos, its core strength lies in text-based classification. This means that fake news detection systems would need to incorporate additional tools to handle the multimodal aspect of misinformation effectively.

Despite these challenges, BERT's overall performance in fake news detection remains strong, and its ability to adapt and evolve with emerging trends in misinformation makes it an invaluable tool. The continuous development of BERT and other NLP models holds significant promise for further improving the accuracy and efficiency of fake news detection systems. Future research and advancements in multimodal fake news detection, combined with BERT's text analysis capabilities, could lead to even more comprehensive solutions for identifying and combating misinformation. Moreover, ongoing improvements in the accessibility of computational resources and labeled datasets may help overcome some of the current limitations, making BERT and similar models more widely usable.

In conclusion, BERT represents a powerful and flexible solution for fake news detection, capable of achieving high accuracy in distinguishing real news from fake news. Its deep contextual understanding, adaptability to new patterns of misinformation, multilingual capabilities, and real-time detection potential make it an ideal tool for addressing the growing challenge of fake news. While there are still challenges to overcome, particularly regarding multimodal misinformation and the need for high-quality data, BERT's strengths position it as one of the leading technologies in the fight against fake news. As the landscape of misinformation continues to evolve, BERT's ability to learn from new data and adapt to changing tactics ensures its relevance and effectiveness in combating fake news for years to come.

REFERENCES

- [1]. Veisi, H., Rafe, V., & Farokhian, M. (2022). Parallel BERT deep neural networks are used to detect fake news. The preprint arXiv is arXiv:2204.04793. This paper presents MWPBert, a truthfulness detection system for full-text news items that uses two concurrent BERT networks. ARXIV
- [2]. Cahyono, H. D., Sihwi, S. W., Azizah, S. F. N., & Widiarto, W. (2023). Performance Evaluation of Transformer-Based Models for Fake News Detection (BERT, ALBERT, and RoBERTa). 2308.04950 is the arXiv preprint. With an accuracy of 87.6%, ALBERT surpassed previous transformer models used in this study to identify bogus news in Bahasa Indonesia. ARXIV
- [3]. Karande, H., Benjamin, V., Kotecha, K., Raghu, T. S., & Walambe, R. (2021). Stance Detection for Credibility Analysis of Social Media Content Using BERT Embeddings. 2105.10272 is the arXiv preprint. By analyzing information using automatic feature extraction and text piece relevance, the authors present a methodology for identifying false news that uses BERT embeddings to obtain a 95.32% accuracy rate. ARXIV
- [4]. Abderrahim, M. E. A., & Bounaama, R. (2023). Using BERTbased models to categorize COVID-19-related tweets for sentiment analysis and fake news detection. The preprint arXiv is arXiv:2304.00636. The application of BERT-based models for COVID-19-related tweet classification is covered in this paper, with an accuracy of 0.90 for detecting fake news. ARXI.
- [5]. Alsamhi, S. H., Saleh, H., and Alharbi, A. (2021). Optimized Convolutional Neural Network for Fake News Identification (OPCNN-FAKE). Access, IEEE, 9, 129471–129489. This paper demonstrates the efficacy of deep learning techniques in this field by presenting an optimized convolutional neural network model for fake news identification
- [6]. Mahara, T., Algarni, A. D., Prakash, P., Josephine, V. L. H., Srinivasan, R., & Verma, O. P. (2023). Deep vs. Shallow: A Comparison of Deep Learning and Machine Learning Methods for the Identification of False Health News. Access, IEEE, 11, 79330-79340. In order to show the superiority of deep learning models, this study compares machine learning with deep learning techniques, such as BERT, for identifying false health news..
- [7]. Chai, S., and Park, M. (2023). Building a User-Centered Fake News Detection Model using Machine Learning Classification Algorithms. Access, IEEE, 11, 71517- 71527. Using a variety of classification methods, the authors develop a user-centered false news detection model, highlighting how user interactions can improve detection performance.
- [8]. Mahankal, Prathamesh (2021). GitHub repository: FakeNewsDetection-Using-BERT. With an emphasis on tweet text and the text found in tagged articles, this project uses BERT to categorize tweets as legitimate or fraudulent. GITHUB.
- [9]. AI Skills. (2021). BERT Model for Fake News Detection. article in medium. This article shows how to use a pre-trained BERT with transfer learning to create a fake news detection model. MEDIUM
- [10]. Blended method for detecting fake news based on BERT (2023). published on ResearchGate. In order to identify fake news, this study investigates the use of machine learning and natural language processing, particula.

