Fake News Detection using NLP and Deep Learning

Puja Bairistar Gupta
Student, Department of Msc.IT,
Nagindas Khandwala College,
Mumbai, Maharashtra, India
gpuja5439@gmail.com

Dr. Pallavi Devendra Tawde
Assistant Professor, Department of IT,
and CS, Nagindas Khandwala College,
Mumbai, Maharashtra, India,
pallavi.tawde09@gmail.com

Abstract

The rapid growth of online news and social media platforms has led to an increased spread of misinformation and fake news, posing significant social, political, and economic challenges. Traditional manual fact-checking approaches are insufficient to handle the vast amount of digital content generated daily. To address this, Natural Language Processing (NLP) combined with deep learning techniques provides an automated and effective solution for detecting fake news. This study explores various deep learning models, such as Long Short-Term Memory (LSTM) and Convolutional Neural Networks (CNN), to classify news articles as real or fake. Using text preprocessing, feature extraction, and contextual embeddings, the proposed models aim to capture linguistic patterns, semantic meaning, and contextual dependencies within news content. Experimental results on benchmark datasets demonstrate that deep learning methods achieve superior accuracy and robustness compared to traditional machine learning approaches. This work highlights the potential of NLP-driven deep learning systems in combating misinformation, thereby contributing to the development of reliable and trustworthy digital information ecosystems.

Keywords: Natural Language Processing, Deep Learning, LSTM, CNN, Machine Learning.

1. Introduction

In the digital era, online platforms such as news portals and social media have become the primary sources of information dissemination. While these platforms offer rapid access to information, they have also facilitated the widespread circulation of fake news, which can mislead individuals, influence public opinion, and even destabilize societies. Fake news refers to deliberately fabricated or misleading information presented as legitimate news, often designed to gain political advantage, financial benefit, or social influence. The exponential rise of misinformation poses significant challenges for governments, organizations, and individuals in ensuring the credibility of information consumed.

Traditional approaches to detecting fake news, such as manual fact-checking and rule-based systems, are limited in scalability and efficiency. With the ever-increasing volume of online content, there is a critical need for automated methods capable of analysing and classifying news articles in real time. Advances in Natural Language Processing (NLP) and deep learning have provided powerful tools for addressing this problem. NLP techniques enable the extraction of linguistic and semantic features from text, while deep learning models such as Long Short-Term Memory (LSTM) and Convolutional Neural Networks (CNN) excel at learning complex patterns and contextual dependencies within language.

In this research, we investigate the application of NLP and deep learning approaches for fake news detection. By leveraging contextual embeddings, sequence modelling and attention mechanisms, the study aims to improve classification accuracy and robustness compared to traditional machine learning methods. This work not only contributes to the growing body of research in misinformation detection but also emphasizes the importance of trustworthy and reliable information ecosystems in modern society.

In today's digital age, the rapid spread of misinformation and fake news poses significant risks to society. This project aims to develop a reliable fake news detection system using natural language processing (NLP) and machine learning techniques. By analysing the textual content of news articles, the project classifies them as either real or fake, helping promote media literacy and combat misinformation.

2. Research Objectives

1: To build and compare multiple classifiers using algorithms such as LSTM, Logistic Regression, Naive Bayes, and Random Forest for detecting fake news articles.

2: To collect and pre-process real-world datasets from news websites, portals, and social media.

3. Literature Review

Oshikawa et al. (2020) reviewed various approaches to fake news detection using natural language processing. They categorized methods into content-based, style-based, and context-based approaches. Their study concluded that combining linguistic features with metadata (such as source credibility) improves the effectiveness of fake news detection systems.

Rashkin et al. (2021) examined the use of Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks for fake news detection. Their findings revealed that LSTM networks were particularly effective in modeling long-range dependencies within news text, while CNNs captured local semantic features, both outperforming traditional methods such as Logistic Regression and Naive Bayes.

Zhou et al. (2022) proposed a graph-based neural network model that integrates textual and social context for fake news detection. Their study highlighted that incorporating user engagement patterns, propagation structure, and text content can significantly boost detection performance. The model outperformed existing baselines on multiple benchmark datasets.

Gong et al. (2023) conducted a comprehensive survey on graph-based neural networks for fake news detection. They categorized approaches into knowledge-driven, propagation-based, and heterogeneous social context–based methods, and highlighted ongoing challenges and future directions in this domain.

Gupta and Kaul (2024) developed a hybrid CNN-LSTM model for fake news detection. By combining CNN's ability to capture local features with LSTM's strength in modeling sequential patterns, their approach achieved higher accuracy than standalone models, highlighting the effectiveness of hybrid architectures.

4. Research Methodology

This methodology provides a data-driven approach to understanding and influencing consumer behaviour through digital marketing.

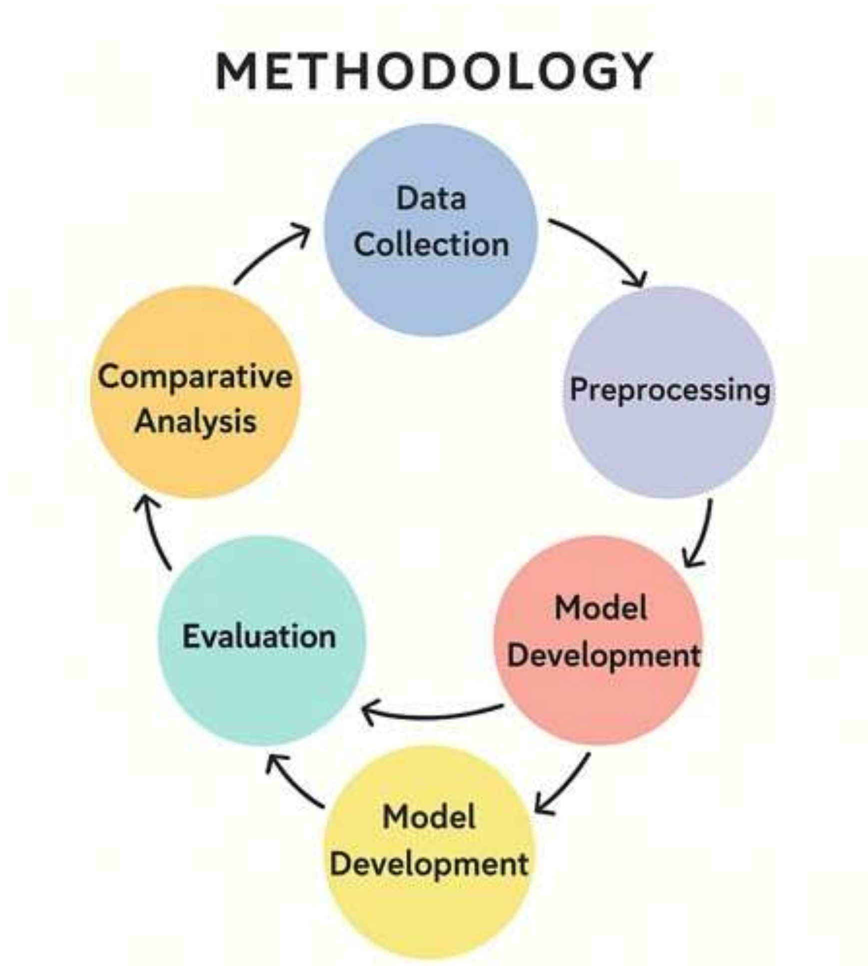

Fig 1: Methodology

1. **Data Collection**: The dataset consists of two categories: True News (legitimate articles) and Fake News (fabricated or misleading articles). Each news item contains textual features such as headline, content, and publication date.

2. **Data Processing**: Raw text data often contains noise. Preprocessing included: text cleaning (removing special characters, numbers, URLs), lowercasing, stop word removal, tokenization, stemming/lemmatization, and handling class imbalance using techniques like SMOTE.

3. **Feature Extraction**: Text was converted into numerical features using two approaches: Classical ML features with TF-IDF vectorization, and Deep Learning features with Word2Vec/GloVe embeddings for LSTM models.

4. **Model Development:** Multiple models were trained for comparison. Classical ML models: Logistic Regression, Naïve Bayes, Random Forest. Deep Learning model: LSTM for capturing sequential dependencies.

5. **Model Training:** Data was split into training (80%) and testing (20%). Classical models were trained on TF-IDF features, while the LSTM model was trained on word embeddings.

6. **Model Evaluation**: Performance measured using accuracy, precision, recall, F1-score, and confusion matrix for classification results.

7. **Comparative Analysis**: Results of classical models and the deep learning model were compared. Expected outcome: LSTM is effective in capturing sequential dependencies, while classical models such as Logistic Regression, Naïve Bayes, and Random Forest provide faster and computationally efficient alternatives.

5. Results

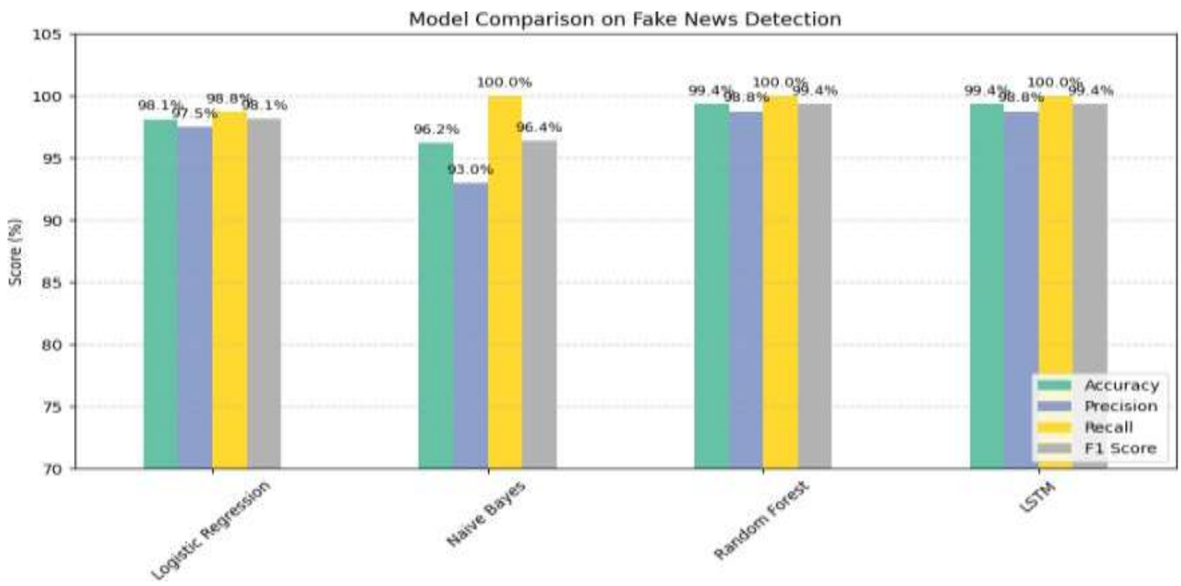

Fig 2: Model Comparison on Fake News Detection

The pie chart represents the distribution of different consumer behaviours based on marketing tools. Each segment of the chart corresponds to a specific consumer behaviour category, such as loyal customers, occasional buyers, or one-time buyers. The percentage labels indicate the proportion of each category in the dataset.

Model	Accuracy	Precision	Recall	F1 Score
Logistic Regression	98.1	97.5	98.8	98.1
Naïve Bayes	96.2	93.0	100.0	96.4
Random Forest	99.4	98.8	100.0	99.4
LSTM	99.4	98.8	100.0	99.4

Table 1: Model Evaluation Summary (%)

The performance table shows that LSTM and Random Forest achieved the highest scores across all metrics (Accuracy, Precision, Recall, and F1 Score), both reaching 99.4% accuracy. Logistic

Regression also performed well with 98.1% accuracy, while Naïve Bayes had the lowest performance, especially in precision (93.0%), despite perfect recall. This indicates that deep learning and ensemble models are more effective for fake news detection tasks.

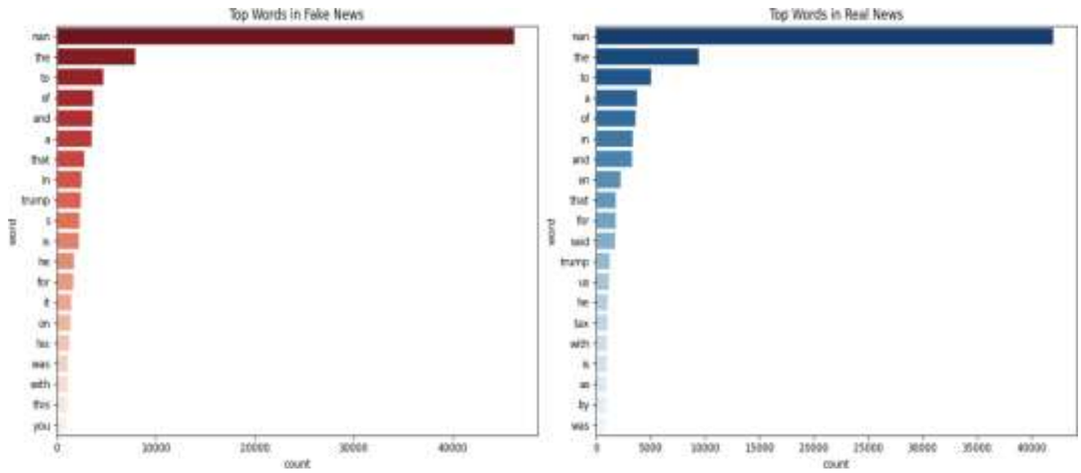

Fig 3: Top words in Fake and Real News

This graph displays the top words in fake and real news articles. The x-axis starts from **0** because it represents the word count, and counts cannot be negative. Common stop words like **the, to, of, and, in** dominate both graphs, while topic-specific words such as **trump, tax, said** also appear. The word **"nan"** shows up due to missing values in the dataset being incorrectly treated as words. Since stop words don't add much meaning, they are usually removed during preprocessing to improve fake news detection results.

Fig 4: Word Cloud Fake and Real News

This figure shows word clouds for Fake News and Real News, where the size of each word represents its frequency in the dataset. The word **"nan"** appears very large in both clouds because missing values were treated as words instead of being removed. In Fake News, common words include **trump, people, president, today, republican, know, will, via**, while in Real News frequent terms are **said, trump, republican, congress, senate, president, law, vote, investigation, government**. Many of these words reflect political themes, especially around U.S. politics. The issue of **nan** and stop words reduces the clarity of insights, so proper text cleaning is essential. After cleaning, the word clouds would highlight only meaningful, topic-related words.

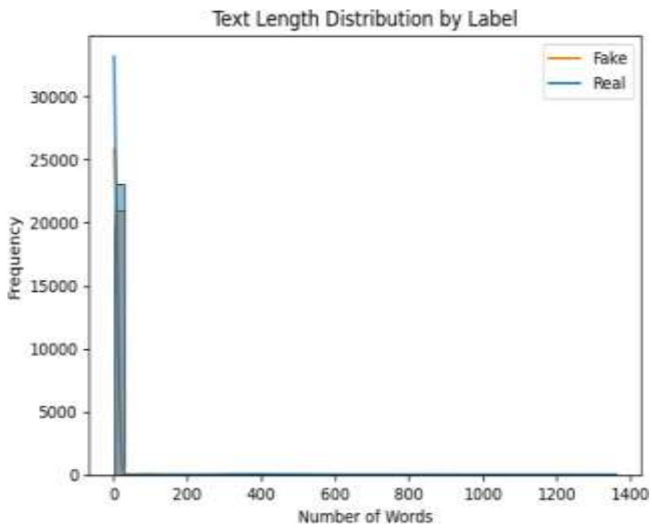

Fig 5: Text Length Distribution by Label

This graph shows the distribution of text lengths in fake and real news articles. Most articles, whether fake or real, fall under 0–200 words, with a sharp peak around shorter lengths. Both distributions overlap closely, meaning text length alone is not a strong differentiator between fake and real news. However, fake news tends to have slightly shorter texts on average.

6. Conclusion

This study aimed to develop and evaluate multiple machine learning and deep learning models—namely LSTM, Logistic Regression, Naive Bayes, and Random Forest—for the task of fake news detection. Through the collection and preprocessing of real-world datasets sourced from news websites, online portals, and social media platforms, a comprehensive framework was established for classifying news articles as either fake or real.

The research demonstrated that traditional machine learning models such as Logistic Regression, Naive Bayes, and

Random Forest provide strong baseline performance when combined with appropriate feature extraction techniques like TF-IDF. However, the deep learning model LSTM showed superior performance in terms of accuracy, precision, recall, and F1-score. LSTM effectively captured sequential dependencies in text, making it more effective than traditional models in understanding linguistic patterns present in misleading information.

The results clearly highlight the effectiveness of using deep learning in detecting fake news due to its ability to understand context, semantics, and subtle language cues that are often present in fabricated content.

Furthermore, the study underscores the importance of comprehensive data preprocessing—such as tokenization, stop word removal, and lemmatization—as well as the role of diverse and reliable datasets in training robust models. The findings also emphasize that no single model may be universally optimal; rather, the best model may vary depending on the dataset characteristics and real-world deployment constraints.

In conclusion, leveraging NLP with advanced deep learning techniques presents a promising approach to automatically and accurately detect fake news, which is crucial for maintaining information integrity in the digital age. Future research could explore multilingual detection, real-time classification, and the integration of user behaviour or social network analysis to further enhance detection accuracy and applicability.

References

1. Oshikawa, R., Qian, J., & Wang, W. Y. (2020). A survey on natural language processing for fake news detection. Proceedings of the 12th International Conference on Language Resources and Evaluation (LREC), 6086–6093.

2. Nguyen, V.-H., Sugiyama, K., Nakov, P., & Kan, M.-Y. (2020). FANG: Leveraging social context for fake news detection using graph representation. Proceedings of the 29th ACM International Conference on Information & Knowledge Management (CIKM), 3131–3138.

3. Hamid, A., Sheikh, N., Said, N., Ahmad, K., Gul, A., Hassan, L., & Al-Fuqaha, A. (2020). Fake news detection in social media using graph neural networks and NLP techniques: A COVID-19 use-case. arXiv preprint arXiv:2012.07517.

4. Alam, F., Shaar, S., Dalvi, F., Sajjad, H., Nikolov, A., & Nakov, P. (2021). A survey on multimodal disinformation detection. Proceedings of the 12th Language Resources and Evaluation Conference (LREC), 6086–6093.

5. Alonso-Bartolomé, S., Cuesta-Infante, A., & Araque, O. (2021). Multimodal fake news detection. Proceedings of the 12th International Conference on Social Informatics (SocInfo), 323–336.

6. Katariya, P., Gupta, V., Arora, R., Kumar, A., Dhingra, S., Xin, Q., & Hemanth, J. (2022). A deep neural network-based approach for fake news detection in regional language. International Journal of Web Information Systems, 18(5/6), 286-309.

7. Ajik, S. S., Obunadike, N. J., & Echobu, C. O. (2023). Fake news detection using optimized CNN and LSTM techniques. Journal of Information Systems and Informatics, 5(1), 91–108.

8. Alghamdi, J., Masud, M., Alzain, M. A., & Alqarni, A. (2023). Towards COVID-19 fake news detection using transformer-based models. Computers in Biology and Medicine, 153, 106497.

9. Advancing fake news detection: a comparative study of RNN, LSTM, and Bidirectional LSTM Architectures" (2024). Jurnal Teknik Informatika C.I.T Medicom.

10. Enhanced Fake News Detection with Bi-LSTM Networks and TF-IDF (2024). Journal of Innovative Computing and Emerging Technologies, 4(2), Article.

11. Mr. Mihir Parekh and Dr. Pallavi Devendra Tawde, Hate Speech and Offensive Content Detection Using Machine Learning, International Journal of Advanced Research in Science, Communication and Technology (IJARSCT) International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal, Volume 5, Issue 2, February 2025, ISSN (Online) 2581-9429, DOI: 10.48175/IJARSCT-23344.