

Linguistic Feature Extraction and Classification for Online Fake News Detection

Keethi Varman

Research Scholar, Department of Information Technology, Jawaharlal Nehru Technological University, Hyderabad, India

Abstract: The proliferation of deepfakes, misinformation, disinformation, and broader post-truth narratives—collectively characterized as fake news—has generated significant concern regarding the influence of digital platforms on public discourse and democratic stability. The widespread diffusion of manipulated multimedia content and algorithmically amplified false narratives across social media ecosystems has intensified challenges related to information credibility, media literacy, and institutional trust. This project addresses the growing threat of digital deception by developing a computational framework for automated fake news classification using advanced Natural Language Processing (NLP) methodologies. The proposed approach leverages linguistic feature extraction, semantic embeddings, sentiment analysis, contextual modeling, and transformer-based deep learning architectures to differentiate between authentic and fabricated news content. Supervised machine learning models are trained on benchmark datasets to capture syntactic irregularities, rhetorical patterns, source credibility signals, and discourse-level inconsistencies commonly associated with deceptive narratives. In parallel, the study explores the integration of Blockchain and other Distributed Ledger Technologies (DLTs) as complementary mechanisms to enhance data provenance, traceability, and content integrity within digital information ecosystems. By utilizing cryptographic hashing, decentralized consensus protocols, and immutable ledger recording, blockchain-based systems can establish verifiable timestamps, authenticate content origins, and prevent unauthorized post-publication alterations. Such architectures enable transparent audit trails and peer-to-peer verification processes, thereby strengthening trust in news dissemination networks. The research provides a systematic review of existing DLT-based interventions for combating misinformation, including decentralized content verification platforms, reputation-based trust scoring systems, and tokenized incentive models for fact-checking communities. Additionally, it identifies critical open challenges such as scalability limitations, interoperability constraints, privacy-preserving identity management, adversarial attacks on NLP models, and the ethical implications of automated content moderation. By synthesizing advances in NLP-driven content analysis and blockchain-enabled integrity assurance, this study proposes a hybrid resilience framework aimed at mitigating cyber-enabled misinformation threats. The findings offer strategic recommendations for future interdisciplinary research, emphasizing robust model generalization, explainable AI techniques, cross-platform collaboration, and regulatory alignment. Collectively, these efforts contribute toward enhancing information reliability, reinforcing democratic resilience, and fostering trust in contemporary digital media ecosystems.

Keywords: Fake News, Natural Language Processing (NLP), Digital Deception, Distributed Ledger Technologies (DLTs), Immutable Records, Data Verification, Resilience.

I. INTRODUCTION

This project addresses the rising threat of fake news, emphasizing its impact on public perception, decision-making, and elections. Fake news often originates from social media and features hyperbolic and non-attributed content. It has led to social harm, including physical violence and reputational damage. The paper discusses a classifier designed to detect fake news using linguistic analysis. It explores how this system can evolve into an influence mining tool. The proliferation of fake

news via bots and deceptive imagery increases its reach. The study aims to enhance information integrity and media literacy. In today's digital age, the way we consume and interact with information has changed dramatically. With the rise of the Internet and social media platforms, news spreads faster than ever before. While this connectivity has brought many benefits, it has also opened the door to serious challenges—chief among them being the spread of fake news.

Fake news refers to deliberately misleading or false

information presented as legitimate news. It includes deepfakes, misinformation, and disinformation, and it plays a powerful role in shaping public opinion. These deceptive practices threaten the foundation of modern democracies by distorting facts, undermining trust in institutions, and fueling social division.

The age we live in is often referred to as the post-truth era, where emotion and personal beliefs seem to matter more than objective facts. This growing trend of digital deception raises urgent questions about how we can safeguard truth and trust in our online ecosystems.

This project addresses that very challenge. It explores how modern technologies—particularly Natural Language Processing (NLP) and Blockchain—can work together to identify and combat fake news. On one hand, NLP allows computers to understand, interpret, and even generate human language. With the help of NLP models, we can analyze large volumes of text to detect patterns that indicate whether a news article is credible or deceptive. These tools offer powerful capabilities in classifying fake news based on linguistic features, sentiment, and context.

On the other hand, Blockchain and Distributed Ledger Technologies (DLTs) offer a different kind of solution. They provide a decentralized, transparent, and secure method of recording information. By ensuring that data is immutable and verifiable, Blockchain can help track the origin of information, verify its authenticity, and prevent tampering.

Together, these two technologies form a promising defense against digital deception. This project provides a comprehensive overview of how NLP and DLTs are currently being used to combat fake news. It also explores real-world applications and highlights important case studies. In doing so, it identifies key challenges, such as scalability, ethical concerns, and the technical limitations of current systems.

Finally, the study puts forward recommendations for future research. It emphasizes the need for collaboration between technologists, policymakers, and media organizations to build more resilient and trustworthy digital platforms.

In a time when facts are under fire, and online manipulation is rampant, building tools to defend the truth is not just important—it's essential.

II. RELATED WORK

Researchers have extensively explored fake news detection

using Natural Language Processing (NLP). Early studies focused on linguistic features such as hyperbole, emotional tone, and lack of source attribution. Machine learning algorithms like Naive Bayes, SVM, Random Forest, and Logistic Regression have been used with text features like TF-IDF and n-grams.

Deep learning models such as RNNs, LSTMs, and CNNs provide enhanced performance by capturing deeper contextual meaning. Pre-trained embeddings like Word2Vec, GloVe, and BERT have further improved semantic understanding in fake news classification.

Several benchmark datasets, including LIAR, FakeNewsNet, and ISOT, have enabled the development and evaluation of fake news detection systems. Some research also integrates social context by analyzing user behavior and engagement patterns on platforms like Twitter and Facebook.

Rubin et al. (2015) Focused on deception detection using rhetorical structure theory (RST) and linguistic cues in news articles. They identified stylistic differences between fake and legitimate news.

Potthast et al. (2017) Proposed a style-based fake news detection model using writing style analysis and shallow linguistic features like POS tags, readability scores, and syntactic patterns.

Wang (2017) Developed the LIAR dataset and evaluated multiple machine learning models on short political statements, using meta-data and linguistic features.

Pérez-Rosas et al. (2018) Built a fake news dataset and applied traditional ML algorithms with manually extracted linguistic features. They found fake news tends to use simpler, more emotional language.

Shu et al. (2019) Proposed a framework called FakeNews Tracker that used content, social context, and temporal patterns for improved detection accuracy.

Zhou & Zafarani (2020) Emphasized early detection and real-time analysis of fake news using both content-based and context-based approaches.

Kaliyar et al. (2020) Introduced a deep learning model using LSTM architecture to capture long-term dependencies in text, improving classification performance.

Vaswani et al. (2017) Though not directly focused on fake

news, their work on the Transformer architecture led to models like BERT, widely adopted for NLP tasks including fake news detection.

Singhania et al. (2017) Proposed Hybrid models that combine TF-IDF-based features and CNN architectures to detect misinformation with high accuracy.

Ahmed et al. (2018) Investigated the role of user comments and social engagements as additional features in detecting deceptive content.

III. PROPOSED SYSTEM

The proposed system for fake news article detection using Natural Language Processing (NLP) employs multiple modern techniques. A text-based classifier is created to distinguish between fake and real news articles. The system is trained using a labeled dataset of news articles, each marked as fake or real. It also supports categorization into four different classes for more refined classification.

To evaluate the model, standard performance metrics like accuracy, precision, recall, and F1-score are used. Additionally, Natural Language Inference (NLI) models are integrated for enhanced detection. The process includes detailed data collection, preprocessing, interpretation, and result presentation in a structured format.

The system strikes a balance between analysis time and classification quality using advanced NLP models. It leverages cutting-edge models such as BERT, GloVe, and ELMo, which are known for their powerful text representation capabilities.

The proposed system for detecting fake news articles leverages both regular machine learning and deep learning models combined with Natural Language Processing (NLP) techniques. It is designed to be automated, analyzing the text content of articles to detect misinformation. Toolkits such as TextBlob, NLTK, and SciPy are employed to build a novel Bayesian classifier that relies on quoted attribution as a key signal.

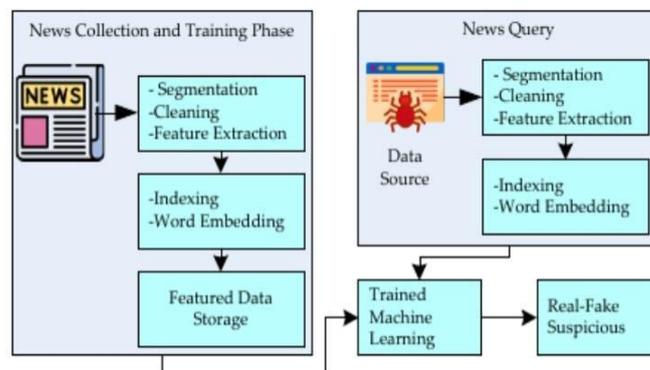


Figure 1: NLP framework for fake news detection

This approach aims to surpass traditional methods in accuracy and efficiency. The model development involves steps like tokenization, regular expressions (regex), and text preprocessing, including punctuation removal and word segmentation. The system thus transforms raw news text into structured data ready for classification.

By combining state-of-the-art NLP tools and structured machine learning pipelines, the proposed system offers a scalable and robust solution to the fake news crisis. The ultimate goal is to create a reliable detector that supports better information integrity across platforms.

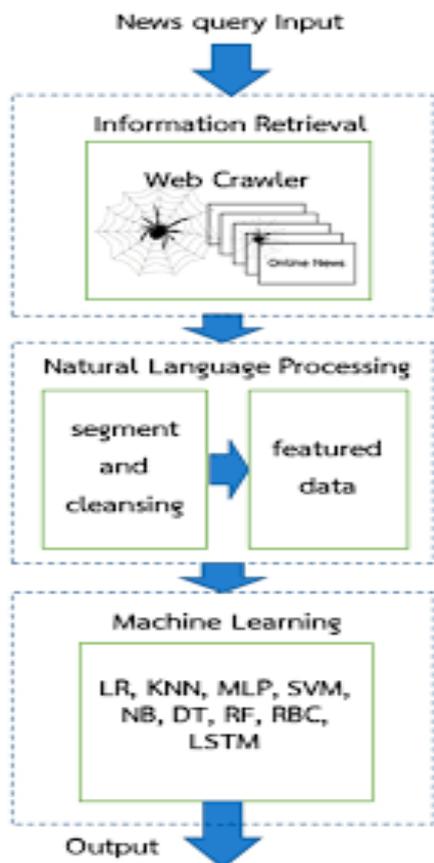
The system follows a multi-stage pipeline that begins with data preprocessing and ends with classification using advanced algorithms.

Initially, a dataset of labeled news articles is collected, where each article is tagged as either fake or real. The text data undergoes preprocessing, including removal of punctuation, stop words, and special characters, followed by tokenization, lemmatization, and vectorization. This prepares the text for analysis by converting it into a machine-readable format.

For feature extraction, word embeddings such as TF-IDF, GloVe, and Word2Vec are used. These help in capturing the semantic meaning of words within the news content. Advanced NLP models like BERT and ELMo are also considered for deeper language understanding.

The system leverages both traditional classifiers (like Logistic Regression, Naive Bayes, and SVM) and modern deep learning models (like LSTM, CNN, and Transformers) to train the classifier. The classification performance is evaluated using standard metrics such as accuracy, precision, recall, and F1-

score.



The use of Natural Language Inference (NLI) models is also explored to understand logical relationships between claims and evidence. Furthermore, quoted attribution analysis is applied using probabilistic frameworks like Bayesian models to enhance reliability.

Overall, the proposed system offers a robust and scalable solution to automatically detect fake news based on the linguistic patterns and semantics of the article text, contributing to the fight against misinformation in digital media.

V. CONCLUSION

This paper presented the results of a study that produced a limited fake news detection system. The work presented herein is novel in this topic domain in that it demonstrates the results of a full-spectrum research project that started with qualitative observations and resulted in a working quantitative model. The work presented in this paper is also promising, because it demonstrates a relatively effective level of machine learning classification for large fake news documents with only one extraction feature. Finally, additional research and work to identify and build additional fake news classification grammars is ongoing and should yield a more refined classification scheme for both fake news and direct quotes.

REFERENCES

- [1] J. Strömbäck, Y. Tsfati, H. Boomgaarden, et al., “News media trust and its impact on media use: Toward a framework for future research,” *Annals of the International Communication Association*, vol. 44, pp. 139-156, 2020.
- [2] E. Mitchelstein and P. J. Boczkowski, “Online news consumption research: An assessment of past work and an agenda for the future,” *New Media & Society*, vol. 12, pp. 1085-1102, 2010.

Citation of this Article:

Keethi Varman. (2025). Linguistic Feature Extraction and Classification for Online Fake News Detection. *Current Journal of Engineering and Science Research*. 2(7), 7-10. Article DOI: <https://doi.org/10.47001/CJESR/2025.207002>

*** End of the Article ***