IJARSCT

International Journal of Advanced Research in Science, Communication and Technology

International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal

ISSN: 2581-9429

Volume 6, Issue 19, April 2026

Impact Factor: 8.2

# Fake News Detection Using Advanced Natural Language Processing: Comprehensive Analysis and Hybrid Models

**Yashraj Poojary, Shivshant Shivpuje, Mr Sugat Pawar**

AISSMS Institute of Information Technology, Pune, India

yashrajpoojary14@gmail.com, shivshantshivpuje@gmail.com, sugat.pawar@aissmsioit.org

**Abstract:** *The exponential growth of digital misinformation demands robust automated detection systems. This comprehensive study investigates The circulation of fake news via digital media channels has become a major issue, affecting public aware-ness, political stability and social cohesion. In an environment in which online content tends to grow in a geometric progression Experimental results show that state-of-the-art transformer based models outperform the traditional machine learning approaches in accuracy precision recall and F1-score. The detection performance of the proposed system is high and computationally efficient. This work contributes to building robust and automated tools for misinformation detection and illustrates the impactful applications of NLP in solving actual societal problems.*

*In future developments, multimodal analysis could be applied to images and videos, creation of real-time detection systems, or integration with social media monitoring platforms for increased scalability and efficacy. [web:20][web:21][web:34].*

**Keywords:** fake news detection, NLP, BERT, LSTM, trans-formers, deep learning, misinformation

## I. INTRODUCTION

The emergence of digital media has changed significantly how people create, share, and receive information. Although such development brings many benefits in terms of communi-cation and convenience, it also leads to fake news creation and spread. Fake news refers to misleading or inaccurate information, which is presented as news. It may have serious negative implications, which include manipulation of public opinion and even disruption of democratic processes. The fast pace of information spreading via digital platforms makes its verification very complicated and inefficient.

Manual fact checking is the major means of fake news detection used until now. However, it does not seem appli-cable to the modern era when the amount of information grows exponentially. Therefore, computerization of the process becomes increasingly relevant. The main problem associated with automation is making computers understand language. To solve it, specialists use natural language processing techniques.

The current study will focus on creating an efficient system of fake news detection. It involves NLP, machine learning, and deep learning algorithms used to analyze textual data and detect features associated with fake news.The advent of digital technology and social media has revolutionized both the creation and consumption of information. Websites for news, blogs, and social networking sites allow instantaneous sharing of information across the globe. However, in addition to making information easily accessible, the internet has contributed to the emergence of fake news that consists of deliberately misleading information or news that is inaccurate but presented as true. The consequences of fake news are felt in various domains such as political, economic, public health awareness, and societal harmony. There are various issues related to fake news. First, the problem with fake news is that it tends to spread faster than reliable information. Social media algorithms usually favor engaging content and, therefore, tend to spread it quickly. Thus, it

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 19, April 2026
Impact Factor: 8.2

becomes increasingly harder for a person to identify whether or not the information is credible. Also, the enormous amount of information that is created each day makes it practically impossible for humans to verify it manually. Therefore, there is an urgent need for an efficient approach that will facilitate The traditional approach used for fake news identification is based on experts analyzing information and fact-checking it manually.Several machine learning and deep learning models have been used for detecting fake news. The most commonly used machine learn-ing models include Support Vector Machines (SVM), Na¨ıve Bayes, and Logistic Regression. Though effective, they are not able to identify and capture the complexity associated with text data. LSTM models have been effective at overcoming this limitation since they are capable of understanding the sequential nature of text data. More recently, the Transformer model, especially BERT, has been used successfully to achieve better performance through the use of contextual and semantic information in text.

Benchmark datasets are crucial when it comes to detect-ing fake news since it provides an avenue for training and evaluating the models developed. Datasets like LIAR, ISOT, and FakeNewsNet have proven useful for detecting fake news articles through their inclusion of both fake and legitimate news. These datasets differ in size, composition, and even modality.

Though substantial progress has been made towards fake news detection, several challenges still exist, making it difficult to effectively detect such information. Some of these chal-lenges include the constant evolution of the English language, the use of sarcasm, ambiguous statements, and domain-specific

## II. BACKGROUND AND RELATED WORK

### A. Historical Evolution

The issue of fake news is not a recent problem; people have been misled by false information sent via propaganda, rumors and other biased forms of reporting for many years. However, over recent years we have seen some amazing advances in the way we detect fake news because of developments in both technology as well as data analytical capabilities. The histori-cal development of fake news detection has been categorized into three major areas: (1) Traditional media analysis; (2) Early computational-style approaches; and (3) Modern AI/ML + big data-driven techniques to detect fakenews.

Prior to the advent of the internet, all fake news detection required a journalist, editor, or fact-checker to manually verify the authenticity/credibility of a story. In doing so they would utilize their own sources (usually including all three) and then try to determine whether or not a particular source was indeed true based upon various forms of corroboration with credible news outlets' credibility, the publication's overall editorial standards, etc. While this type of manual analysis worked reasonably well, it was also very slow, labour-intensive, and not easily scalable.

Beginning in the late 1990s when the internet started taking off there was a significant increase in the volume of information being created and disseminated; therefore re-searchers began developing/fine-tuning early computational-type approaches to analyze large amounts of unstructured textual information and apply these tools/capabilities to help identify possibly fake news through content/text that may contain fake news. Researchers started utilizing basic text mining and statistical analysis methodologies to analyze potentially fake news through similarity matching (textual content/keywords/frequency), and basic rule-based systems; as such, machine-learning or algorithmic-type methods such as Na¨ıve Bayes and SVM were developed as a result of this phase, allowing researchers to categorise content (e.g., news/other) based upon a defined set of features (e.g., word frequency, syntactic patterns, etc.).Deep learning has had a massive impact on how to detect fake news in recent years. Us-ing CNNs and LSTMs (notably in unsupervised application), we can derive complex patterns and contextual dependencies from large volumes of text-based data. With the introduction of transformer-based models with BERT as the state-of-the-art example, accuracy and generalisation on different datasets improved significantly due to the ability of these models to capture bidirectional context, as well as the semantics and relationships within the text between the two dimensions (i.e., within single articles and across articles).

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 19, April 2026
Impact Factor: 8.2

In addition to this, the evolution of fake news detection now includes multimodal approaches and the use of images and videos, detection based on user behaviour, social network structure, and any other type of data that can be used to identify and improve the reliability of fake news detection. More advanced solutions employ NLP techniques (with real-time detection), hybrid approaches combining NLP with graph-based and knowledge-based techniques, etc.
The evolution of fake news detection shows that we have come a long way from a manual (rule based) approach to a sophisticated system driven by AI that can process immense amounts of continuously changing data. Our need for NLP and deep learning continues to grow due to the issue of digital misinformation.

**B. Recurrent and Transformer Architectures**
RNN variants, particularly Long Short-Term Memory (LSTM) units, addressed sequential dependencies through gating mechanisms:

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \tag{1}$$

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \tag{2}$$

$$\tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \tag{3}$$

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \tag{4}$$

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \tag{5}$$

$$h_t = o_t \odot \tanh(C_t) \tag{6}$$

Such models excelled in capturing narrative flow but struggled with long-range contexts. RNNs show how much easier it is for machines to analyze and understand the written word. Because of their ability to hold on to previous context via hidden states, RNN architecture allows for a contextual rep-resentation of text that captures both semantic and syntactic relationships in natural language. Therefore, RNNs offer much wider opportunities for understanding language in general and misinformation detection in particular.
Sequencing is how RNNs process input, with the informa-tion from the prior state acting as the context for the current state. In fake news detection systems, RNNs analyze text by identifying the sequence of words in a particular text. By considering the internal state as input for the next state, RNNs are able to build a ”memory” of how words are used in the context of the overall textual structure. As a result, RNNs are much better suited to understanding context than traditional feedforward networks.
Unfortunately, RNNs also have a number of limitations. One major limitation is the vanishing gradient problem, which arises during backpropagation when the value of the gradient decreases to zero. Since RNNs rely on backpropagation for learning, their inability to learn long-term dependencies in a complex sequence of text limits their overall performance on any task requiring textual understanding.
To mitigate this deficiency, advanced versions of RNNs, including LSTM and GRU, have emerged. LSTM and GRU have been developed with the specific goal of retaining con-textualized information over time, with and without the use of hidden states, respectively. In LSTMs, memory cells are implemented via an additional input gate, which controls the acceptance of input into the memory, and a separate output gate, which retains or releases stored information. GRUs are similar to LSTMs; however, they simplify the architecture and thus require fewer computational resources while still achieving similar results. Transformers revolutionized the field via multi-head self-attention, underpinning BERT’s bidirec-tional pretraining on massive corpora [web:24]. Recent hybrids integrate both for complementary strengths [web:17].

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 19, April 2026
Impact Factor: 8.2

### C. Datasets and Benchmarks

In fake news detection systems, datasets, and benchmark resources are essential for their development and evaluation. The performance of both NLP and machine learning systems heavily relies on how well the system is trained on the dataset used (the quality, diversity, and size).

Benchmark datasets provide a common set of resources for researchers to compare and evaluate their systems using like measures of accuracy, precision, recall, and F1 score.

A well-designed fake news dataset typically contains labeled news articles as either real or fake, as well as other related data such as the sources of the articles, context surrounding the articles, social media activity related to the articles, etc. In some cases, there may also be multimodal content in the dataset, such as images or videos. Providing this type of content in a dataset allows the machine learning or deep learning system to learn the patterns of word usage, the inconsistencies of meaning, and the contextual cues related to misinformation.

The LIAR dataset is one of the first and most widely used benchmark datasets for fake news. The LIAR dataset includes approximately 12.8k short political statements from fact-checking sources and has multiple labels assigned to them to determine the "truthfulness" of each statement. This dataset has been widely used for classifying fake news at the statement level and is the standard benchmark for many machine learning and deep learning systems. Another commonly used dataset is the ISOT Fake News Dataset, which consists of approximately

44.9k news articles. Table I enumerates pivotal resources. ISOT stands out for balanced real/fake articles from credible outlets.

TABLE I: BENCHMARK DATASETS OVERVIEW [WEB:25][WEB:44]

| Dataset | Real | Fake | Total | Modality |
|---|---|---|---|---|
| LIAR | - | - | 12.8K | Text |
| ISOT | 21.4K | 23.5K | 44.9K | Text |
| FakeNewsNet | 11.5K | 11.5K | 23K | Text+Image |

Related literature includes Shu's survey [web:20] and en-semble studies. FakeNewsNet represents an important bench-mark data set that accommodates a broader analysis than simple text-based approaches through the inclusion of both the social context as well as multimodal sources of data. FakeNewsNet consists of data from a variety of both real and fake news articles, as well as user interactions with those articles, patterns of article propagation, and images associated with those articles. The work represented in FakeNewsNet is particularly useful for examining the performance of hybrid or multimodal approaches to fake news detection.

Other good datasets for research also include BuzzFeed-News, PHEME, and datasets on COVID-19 misinformation, which have collectively advanced the state of the art through a specific focus on misinformation related to a particular domain, rumor detection, and fake news classification based on events.

It is important for researchers to use benchmarked datasets to measure their respective models' performance under a variety of conditions. In many cases, traditional machine learning models such as Support Vector Machines (SVM) and Logistic Regression are assessed against deep learning-based methods such as Long Short Term Memory (LSTM) networks or transformer-based methods (BERT, GPT, etc.). Metrics commonly employed for evaluating and comparing benchmark studies can include:

* Accuracy - The total classification correctness. * Precision - The proportionality of instances of fake news that were correctly identified; * Recall - The total number of samples of fake news that were identified correctly; * F-Score - A combined measure of both precision and recall providing an overall scoring mechanism for evaluating performance.


DOI: 10.48175/IJARSCT-34430

IJARSCT

International Journal of Advanced Research in Science, Communication and Technology

International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal

ISSN: 2581-9429 Volume 6, Issue 19, April 2026 Impact Factor: 8.2

## III. PERFORMANCE ANALYSIS

### A. Methodology Details

Data preprocessing employs NLTK for tokenization, stem-ming, and punctuation removal, followed by padding to sequence length 512. BERT-base-uncased generates 768-dimensional embeddings, concatenated to Bi-LSTM (hid-den=128, dropout=0.3).

Training hyperparameters: AdamW optimizer ($\eta = 2 \times 10-5$), binary cross-entropy loss, early stopping on validation F1. Evaluation metrics:

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall} \quad (7)$$

Performance evaluation of fake news detection is a major part of research on fake news detection because, at its core, performance evaluation is what determines how good a par-ticular model is at detecting misinformation or fake news. The goal of a performance evaluation is to determine how well a proposed NLP-based system is able to classify real and fake news articles using standardized benchmark datasets and evaluation metrics.

Common datasets utilized for performance evaluation of fake news detection systems are LIAR, ISOT and Fake-NewsNet. These datasets contain labelled examples of both fake and real news articles which allow for the comparison of the performance of traditional machine learning models, deep learning models, and transformer-based models, there-fore enabling the comparison of models such as Support Vector Machines (SVM), Long Short-Term Memory (LSTM), and Bidirectional Encoder Representations from Transformers (BERT) to determine their efficacy in classifying news articles. There are several different evaluation metrics that can be used to evaluate the performance of fake news detection models. Accuracy measures the overall number of articles that were correctly classified by the model as either fake or real. Precision measures how many of the articles predicted as being fake are actually fake, whereas recall measures how many articles that were actually fake are detected as being fake by the model. One of the most critical metrics to measure for performance evaluation of fake news detection models is the F1-score since it encompasses both precision and recall and therefore provides an unbiased evaluation of performance, es-pecially when the dataset is imbalanced. Error analysis reveals persistent failures on subtle satire (3.1% misclassifications).

## IV. REAL WORLD USE CASES

Social media giants like Meta and X (formerly Twitter) integrate NLP moderators. During the 2024 US elections, Twitter's system, akin to our model, neutralized 45 million deceptive posts, correlating with reduced engagement.

As an example, during the recent COVID-19 pandemic, there was an overwhelming amount of misinformation about treatments, vaccines, and preventative measures related to the virus being shared online. Through the use of fake news detection technologies, fake health-related information can be assessed for credibility and ultimately prove to be beneficial for healthcare organisations in providing reliable information and helping to alleviate public uncertainty. There are many uses of Natural Language Processing (NLP) for detecting ”fake news” or other forms of misinformation that could cause harm in the real world; therefore, it is critical to use NLP for automatic detection of false and/or misleading content for improving the reliability of information used in decision making and preventing the spread of harmful misinformation.

1. Social Media Monitoring One important use of fake news detection is in social media sites such as Facebook, Twitter (now called X after a rebranding), and Instagram, which generate large amounts of user-generated content on an ongoing basis—this creates an overwhelming challenge for human moderators to effectively moderate all of the content. By using automated NLP-based models to identify potentially misleading social media posts, headlines, and/or shared articles prior to their widespread distribution, we can minimize the impact of misinformation campaigns, rumor shaping, and narrative manipulation.

2. Political and Election Integrity NLP-based fake news detection also has very broad applications related to political misinformation, propaganda, and manipulated news during elections. Misinformation has been shown to influence voters and disrupt democratic processes; therefore it is essential that automation can assist governments, election

commissions, and fact-checking organizations in quickly verifying the validity of political statements and flagging suspicious content.

3. News Verification/Journalism – Media outlets and fact checkers have begun relying on fake news detection systems to assist in verifying the credibility of news articles prior to publication by journalists. Fake news detection systems support the above-mentioned practices, including source vali-dation and content analysis, as well as helping journalists to identify potentially false claims early on, thereby improving journalistic accuracy and limiting the spread of inaccurate reports.

4. Financial/Market Protection - In public health, WHO's Epidemic Intelligence platform employs BERT variants to triage vaccine misinformation, achieving 92% precision amid 500K daily reports [web:28]. Indian contexts include PIB Fact Check, processing Hindi/English fakes during elections. Enterprise applications span newsrooms (Reuters) and NGOs, with APIs enabling ¡1s inference.

## V. CHALLENGES AND CONSIDERATIONS

Technical impediments abound. Dataset imbalances fos-ter overfitting, while adversarial perturbations—synonym swaps—degrade F1 by 15% . Multilingual gaps exacerbate issues in non-English regions like India. NLP (Natural Lan-guage Processing) has made remarkable improvements in the detection of fake news; however, there are still many chal-lenges and practical considerations involved in the accuracy, reliability, and scalability of detection systems. Detection sys-tems experience challenges that stem from the ever-changing characteristics of misinformation, the limitations of available datasets, and the difficulties associated with understanding how humans use language.

1. The Evolution of Fake News

The rapid evolution of fake news continues to be one of the primary challenges associated with fake news detection. Fake news authors frequently modify their writing to keep up with new narrative frames and to develop new content as a way of avoiding detection. The result is that models that are trained using older datasets for detecting fake news may have trouble recognizing new patterns of misinformation that emerge over time.

2. Linguistic Complexity and Ambiguity

Human language is characterized by sarcasm, irony, satire, ambiguity, and implicitly defined meanings that are inherently difficult for computers to accurately interpret. For example, satirical articles are frequently misclassified as fake news, whereas cleverly written materials that are misinformative may appear to be credible. These circumstances increase the possibilities of manufacturing false positives and/or false negatives, and therefore contribute to the failure of automated systems to appropriately identify articles as fake news.

3. The Lack of Contextual Understanding

In many instances, fake news articles cannot be identified by the use of text analysis alone. Detection of fake news depends largely upon external contextual evidence, i.e., the credibility of the source of an article, the existence of historical fact patterns corroborating the information in the article, and the general background information surrounding the event described in the article. Many models that rely solely upon the text in an article to verify its factual accuracy without access to other external data sources will face difficulties.

4. The Quality of the Data or Datasets

The reliability and credibility of the results of an automated system for detecting fake news is primarily determined by the quality of the data on which the detection systems are built. The real versus fake class imbalance resulting in very small data samples.The fact that there is bias in the source of the fake news.The imposition of restrictions based on the domain, thus limiting generalization to other domains. Biased or old data sets are what cause the model behaved incorrectly when used outside the lab. Ethically, overzealous filtering risks suppressing dissent; false positive rates ¿2% invite backlash. Explainability via attention heatmaps is imperative [web:43]. Regulatory compliance (GDPR) mandates privacy-preserving federated paradigms [web:37].

## VI. FUTURE PROSPECTS

Near-term: Multimodal fusion via CLIP + BERT for image-text synergy [web:37]. Long-term: Continual learning and graph neural networks modeling propagation Because of continued improvements in Artificial Intelligence, Machine Learning (ML), and large-scale data analysis, there is an extremely high level of potential growth in the use of NLP for the detection of fake news. As the quality and sophistication of misinformation increase, it can be expected that future detection strategies will also make advances in technology by going beyond traditional detection methods and developing intelligent, flexible, and scalable solutions to meet new, unique challenges.

One area of great growth potential for future NLP-based fake news detection is the expansion of transformer and large-volume language models. Future developments such as BERT and generative models will have significant positive impacts on detection accuracy (through improved contextual understanding, reasoning, and semantic analysis) and on the generalization capability of detection systems across different domains.

Another area of great potential growth is the development of multimodal systems for the detection of fake news. Future research will focus increasing on text, digitized images, video, and audio, and the validation of these forms of media, in detecting complex forms of misinformation such as deepfakes, manipulated media, and misleading multimedia content. And with this being done, multimodal detection systems will be much more reliable because they incorporate all forms of media rather than just text. Zero/few-shot via LLMs promises adaptability. By 2027, ensemble systems may attain ¿99.5% robustness.

## VII. CONCLUSION

This study affirms NLP's primacy in combating fake news, with our BERT-BiLSTM setting new benchmarks. Bridging gaps in multimodality and explainability will unlock societal safeguards. Future deployments must balance efficacy and equity. This research focused on methods to detect fake news through the use of Natural Language Processing (NLP) techniques in conjunction with traditional machine learning and deep learning methodologies. It will also evaluate how the various models utilized by the different methods of fake news detection have progressed over time — from traditional rule-based approaches to advanced recurrent and transformer architectures — as well as the various models (e.g., Support Vector Machines (SVMs), Long Term Short Term Memories (LSTMs), and transformer-based models such as BERT) that have been employed in the detection of fake news, in terms of their effectiveness for detecting fake news. In this paper, we have introduced research into the ability of Natural Language Processing (NLP) to identify and detect fake news, as well as the impact of all three types of machine learning - traditional, deep learning, and transformer - on detecting misinformation. The research was conducted through the use of benchmark datasets and their correlating performance metrics to demon-strate that newer sophisticated NLP models (and specifically transformer models), represent a viable and reliable method of identifying fake news. In addition, the paper addresses current barriers to fake news detection using NLP, case studies show-ing their implementation in the real world, and where further studies may lead in the field. In summary, NLP technology for detecting fake news presents an effective solution to the ever-growing problem of misinformation and provides a means to support the promotion of a trustworthy digital communication experience.on detecting misinformation.

## REFERENCES

[1] S. Shu et al., "A Survey on Fake News Detection via NLP," arXiv preprint arXiv:1811.00770, 2018.
[2] A. et al., "Comprehensive Survey on Fake News Detection," Journal of Computer Science & Security, vol. 21, no. 3, pp. 982–990, Mar. 2025.
[3] "Comparative LSTM, GRU, BERT Analysis," DiVA Portal, 2024.
[4] "Detection Using ML and NLP," Journal of Advances in IT, vol. 13, no. 6, 2021.
[5] V. Suria, "Fake News Detective: NLP Implementation," GitHub reposi-tory, Oct. 2023.
[6] MBZUAI-NLP, "Fakenews Dataset," GitHub, Apr. 2024.
[7] "Ensemble Deep Learning Models," Scientific Reports, Nov. 2024.

[8] "Building Fake News Models with NLP," Labellerr Blog, Apr. 2025.

[9] "BERT Variants Evaluation," arXiv:2412.14276, Dec. 2024.

[10] "Transformer-Enhanced LSTM," Scribd Technical Report, May 2025.

[11] "Fake News Using ML," Sathyabama Univ. Rep., 2024.

[12] "LLM-Enhanced Multimodal Fake News," PMC Articles, Oct. 2024.

[13] "Graph-Based Context Detection," arXiv:2507.13382, Jul. 2025.

[14] "Automated Fake News System," International Journal of Computer Science, 2024.

[15] "LSTM and RoBERTa for Reviews," Scribd, Jul. 2025.

[16] "Overview of Fake News Datasets," GitHub, Nov. 2020