




International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 14, Issue 11, November 2025

AI-Powered Fake News Detection System using Natural Language Processing (NLP) and Deep Learning

Dr. Pankaj Agarkar¹, Tushar Zalte², Farhan Shikalgar³, Umesh Kardile⁴

Department of Computer Engineering, Ajeenkya DY Patil School of Engineering, Lohegaon, Pune, India¹⁻⁴

ABSTRACT: Fake news has grown to be a serious social issue that influences social trust and public opinion. The rapid spread of false information via online media platforms makes it difficult to discern between trustworthy and dishonest content. To solve this issue, a lot of work has been done on machine learning and natural language processing (NLP) techniques to develop automated detection systems. These techniques use statistical, contextual, and linguistic features to categorize news stories as either authentic or synthetic. This book presents a variety of machine learning techniques, including more traditional classifiers such as Logistic Regression (LR), Naïve Bayes (NB), Support Vector Machines (SVM), and Decision Trees, as well as deeper, more complex models. Furthermore, the hybrid models that integrate natural language processing (NLP) and ensemble learning are highlighted for their potential to increase classification reliability and accuracy. Textual data is arranged into machine learning inputs using feature extraction methods such as word embeddings, N-grams, and Term Frequency-Inverse Document Frequency (TF-IDF). In the battle against misinformation, network-based methods and sentiment analysis are also investigated as alternatives to pattern recognition. The results offer compelling proof that integrating deep learning methods with statistics improves the efficacy and accuracy of false news detection. By illuminating the strengths and weaknesses of different methods, this study lays the foundation for future research aimed at developing reliable and scalable systems for detecting fake news.

KEYWORDS: Fake News Detection, Natural Language Processing (NLP), Machine Learning, Deep Learning, Text Classification, Feature Extraction, Sentiment Analysis, Transformer Models, TF-IDF.

I. INTRODUCTION

The circulation of misleading information has become one of the most critical challenges of the digital era. Unlike traditional media, where content passes through editorial checks, modern platforms such as social media, instant messaging, and online blogs allow unrestricted sharing of information. This unregulated flow of content influences public perception, impacts election outcomes, spreads false health advice, and even provokes social unrest. Major global events, including the 2016 U.S. Presidential Election and the COVID-19 pandemic, demonstrated how disinformation can often overshadow official sources in visibility and engagement.

Conventional news verification methods that rely on human fact-checkers, journalists, and regulatory agencies are too slow to cope with the speed and scale of online content generation. This creates an urgent demand for automated solutions that can analyze large volumes of text in real time and reliably identify misinformation. Natural Language Processing (NLP) and Machine Learning (ML) have emerged as key technologies for this purpose. These systems can detect linguistic cues such as exaggerated headlines, emotional tone, inconsistent semantics, and stylistic deviations from genuine reporting. Coupled with NLP techniques like tokenization, lemmatization, sentiment analysis, and parsing, they form the foundation for effective fake news detection.

Advances in feature representation, such as word embeddings (Word2Vec, GloVe) and vectorization methods (TF-IDF, n-grams), combined with deep learning approaches (CNNs, RNNs, LSTMs), have further enhanced performance by capturing contextual and structural patterns in text. More recently, transformer-based architectures such as BERT, RoBERTa, and GPT have revolutionized this field by providing pre-trained models capable of understanding fine-grained semantic dependencies. These models significantly outperform traditional ML techniques due to their deep contextual learning and transfer learning capabilities, making detection systems more accurate and scalable.

This paper compares classical classifiers (Logistic Regression, Naïve Bayes, Support Vector Machines), deep learning models, and transformer architectures, while also examining hybrid approaches that combine statistical and neural techniques. The study evaluates these methods on benchmark datasets in terms of accuracy, scalability, and adaptability to real-world challenges. Special emphasis is given to the threats posed by AI-generated content and multimodal fake news, as well as the importance of Explainable AI (XAI) for building transparent and trustworthy systems. Ultimately, detecting misinformation requires an interdisciplinary approach that balances linguistic analysis, computational innovation, and ethical responsibility. By leveraging NLP and ML, this work contributes to developing intelligent tools that strengthen digital literacy, preserve information integrity, and support an informed society.

II. RESEARCH METHODOLOGY

This section outlines the methodology adopted for developing a fake news detection system using Machine Learning (ML) and Natural Language Processing (NLP). The process follows a structured pipeline, beginning with dataset acquisition and progressing through preprocessing, feature extraction, model training, evaluation, and future recommendations.

2.1 Dataset Collection

The dataset employed in this research was sourced from Kaggle and contains a balanced mix of legitimate and fabricated news articles. Each entry is accompanied by a label that indicates whether the news is real or fake, enabling supervised learning. The dataset encompasses a wide range of domains to ensure diversity, fairness, and better generalization of models. Using a publicly available and curated dataset supports reproducibility and facilitates unbiased model evaluation.

2.2 Text Preprocessing

To ensure consistency and reliability, the raw textual data underwent preprocessing. Tokenization was applied to split text into meaningful units (words or phrases). Stop words, redundant symbols, and irrelevant characters were removed. Stemming and lemmatization were used to reduce words to their root forms, thereby improving consistency across the dataset. Normalization steps, including converting text to lowercase and removing punctuation, helped minimize noise and refine the input for downstream tasks.

2.3 Feature Engineering

Following preprocessing, numerical representations of text were generated to serve as inputs for ML models. Techniques such as Count Vectorization, n-gram modeling, and Term Frequency-Inverse Document Frequency (TF-IDF) were applied to capture word frequency and contextual patterns.

Additionally, feature selection methods like Chi-Square Test and Mutual Information were used to reduce dimensionality and emphasize the most relevant attributes. These steps enhanced computational efficiency and improved classification performance.

2.4 Model Development and Training

The extracted features were utilized to train multiple machine learning and deep learning models. Classical ML algorithms such as Logistic Regression, Support Vector Machines (SVM), and Random Forest were compared against advanced neural models, including Convolutional Neural Networks (CNNs) and Long Short-Term Memory networks (LSTMs). Transformers such as BERT and RoBERTa were also incorporated to capture deep contextual dependencies. Each model was trained on a split dataset with cross-validation to ensure reliability and robustness.

2.5 Evaluation and Future Scope

The performance of the models was evaluated using standard metrics: accuracy, precision, recall, and F1-score. Accuracy measured the overall correctness of predictions, precision highlighted the proportion of correctly identified fake news, and recall measured the ability to capture all fake news instances. The F1-score provided a balanced measure in cases of class imbalance, which is common in real-world scenarios.

While the models showed strong detection capabilities, there is still scope for improvement. Future research should focus on:

- **Real-time Fact-Checking Systems:** Capable of detecting misinformation as it spreads across platforms.
- **Explainable AI (XAI):** To improve transparency, accountability, and trust in high-stakes applications.
- **Self-Supervised Learning:** To reduce dependency on large labeled datasets by enabling models to learn representations from raw, unlabeled text.

These innovations will support the development of more scalable, adaptive, and reliable fake news detection systems, ensuring resilience against evolving disinformation strategies.




Figure1: Structured Methodology for Automated Fake News Classification

III. LITERATURE REVIEW

The identification of fake news has emerged as a prominent research area, with Machine Learning (ML) and Natural Language Processing (NLP) techniques increasingly used to improve detection accuracy [14,15]. Studies focusing on linguistic and semantic features—such as sentence structure, vocabulary richness, and textual coherence—have revealed that fabricated news articles often employ exaggerated wording, demonstrate reduced lexical variety, and present extreme readability levels [16]. Sentiment and emotion-based analysis further indicate that fake news frequently relies on emotionally charged language to manipulate readers' responses [17].

Stylometric approaches, which analyze writing patterns and syntactic style, have also shown promising results in distinguishing authentic content from deceptive reporting [18]. While traditional ML models such as Naïve Bayes, Support Vector Machines (SVM), and Decision Trees have been widely applied in this field, the advent of deep learning architectures—including Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTMs), Convolutional Neural Networks (CNNs), and hybrid models—has led to considerable improvements in accuracy by effectively capturing contextual and sequential relationships within text data [19].

Transformer-based models like BERT, GPT, and T5 leverage self-attention mechanisms to achieve state-of-the-art performance in understanding contextual semantics [20]. Moreover, ensemble and hybrid methods, which integrate statistical learners with deep neural models, have further enhanced robustness and generalization [21]. In addition, innovations in real-time detection increasingly incorporate user behavior analytics, network propagation patterns, and Explainable AI (XAI) frameworks to provide greater transparency and accountability [22].


Despite these advances, fake news detection remains a challenging problem due to the evolving tactics of misinformation, the rise of multimodal content (combining text, images, and videos), and concerns regarding bias in algorithms [23]. Looking ahead, research must focus on building adaptive, scalable, and ethically responsible detection systems capable of addressing the complex and rapidly changing nature of online disinformation.

3.1 Linguistic and Semantic Analysis

One of the earliest and most widely adopted methods for detecting fake news involves **linguistic and semantic analysis**, which examines textual properties such as sentence construction, vocabulary usage, and overall coherence of meaning [24–26]. Research findings highlight several distinctive linguistic patterns commonly present in deceptive content:

- **Exaggerated or sensational wording:** Fabricated articles often rely on emotionally persuasive terms, hyperbolic statements, and misleading titles designed to maximize reader engagement and provoke strong reactions.
- **Reduced lexical diversity:** Fake news tends to recycle particular words and phrases repeatedly, reflecting limited variety in expression and signaling a lack of richness in vocabulary.
- **Irregular sentence complexity and readability:** Deceptive texts may display abnormally simplistic or, conversely, unnecessarily complex sentence structures, deviating from the balanced and standardized style usually observed in professional journalism.

These linguistic and semantic indicators form the foundation for many detection systems, providing a crucial layer of analysis before integrating more advanced machine learning or deep learning models.



3.2 Advanced Techniques for Fake News Detection

Fake news detection has evolved through techniques such as sentiment and emotion analysis, which have shown that fake news tends to incorporate emotionally charged or subjective language, as opposed to the neutral tone of authentic news[27]. Emotion classifiers that identify fear, anger, or surprise assist in identifying deceptive content. Stylometric analysis, such as n-gram patterns, grammar, and readability scores, also assists in detecting fake news[28]. Machine learning algorithms such as Naïve Bayes, SVMs, Decision Trees, and Random Forests categorize content using linguistic features but demand heavy feature engineering. In order to combat increasingly complex disinformation, scientists are integrating deep learning and NLP for higher accuracy.




Figure 2: Distribution of Core NLP and ML Techniques for Fake News Detection

3.3 Advanced Deep Learning Techniques for Fake News Detection

With the rise of misinformation, identifying fake news is becoming more crucial with deep learning. RNNs, LSTMs, and CNNs learn textual patterns and contextual relationships, improving classification accuracy[29]. Hybrid models that employ a combination of CNNs and LSTMs, and transformer-based models like BERT, GPT, T5, and XLNet offer state-of-the-art results with deep language and contextual understanding. Researchers also use ensemble methods and combine statistical models like SVM and Naïve

Bayes with deep learning to provide improved accuracy and explainability[30]. Knowledge-based systems and external fact-checkers assist in making it reliable. User behavior analysis and network-based tracking are components of real-time detection now. Explainable AI (XAI) is also used to build trust. Some of the challenges are the novel misinformation tactics, the rise of multimodal content, bias, and scalability[31]. Future work seeks to create adaptive, ethical, and robust AI systems using techniques like self-supervised, reinforcement, and federated learning.




Figure 3: Proportion of Deep Learning Strategies for Fake News Classification

IV. ALGORITHM

5.1 BERT-based Fake News Detection

Input: Dataset $D = \{(x_i, y_i)\}_{i=1}^N$, pre-trained BERT weights, hyperparameters: epochs E , batch size B , learning rate η , max length L

Output: Trained classifier f_θ that maps input text to label $\{\text{Real}, \text{Fake}\}$

1. Data Preparation

- 1.1 Load dataset D and inspect balance of labels.
- 1.2 Clean text: lowercase, remove HTML tags, punctuation, and normalize whitespace.
- 1.3 Optionally apply lemmatization/stemming and remove stopwords (for analysis only).
- 1.4 Split D into train (80%) and test (20%) sets; create validation split from train (e.g., 10% of train).

2. Tokenization & Encoding

- 2.1 Initialize BERT tokenizer with max sequence length L .
- 2.2 For each x_i : compute $\text{input_ids}_i, \text{attention_mask}_i = \text{Tokenizer}(x_i, \text{max_length}=L, \text{truncation}=\text{True}, \text{padding}=\text{'max_length'})$.

3. Model Setup

- 3.1 Load pre-trained BERT encoder and attach a classification head: $\text{Dense}(\text{hidden_size} \rightarrow 1) + \text{Sigmoid}$ (or Softmax for multiclass).
- 3.2 Initialize optimizer (AdamW) with learning rate η and weight decay; set scheduler (linear warmup + decay).

4. Training Loop

For epoch in $1..E$:

For each mini-batch of size B from training set:

- a) Forward pass: $\text{logits} = f_\theta(\text{input_ids}, \text{attention_mask})$
- b) Compute loss L_{batch} using binary cross-entropy between logits and labels.
- c) Backpropagate gradients: $\nabla_\theta L_{\text{batch}}$
- d) Update weights θ using optimizer step.

End for

Evaluate on validation set; save model checkpoint if validation metric (F1) improves.

5. Evaluation

- 5.1 Load best checkpoint θ^* .
- 5.2 Compute predictions on test set; derive metrics: Accuracy, Precision, Recall, F1-score, Confusion Matrix.
- 5.3 Perform error analysis; compute explainability maps (attention weights or Integrated Gradients) for sample predictions.

6. Deployment

- 6.1 Export tokenizer and θ^* ; build inference API (Flask/FastAPI) that accepts raw text and returns label + confidence + explanation.

V. RESULT AND DISCUSSION

The performance of a fake news detection system depends on several factors, including the choice of machine learning algorithms, the representativeness of the dataset, the effectiveness of feature extraction, and the strategies used for hyperparameter tuning [32]. In this study, a broad range of models were examined, spanning from traditional classifiers to advanced deep learning architectures and transformer-based approaches [33].

Among conventional methods, Logistic Regression (LR) and Naïve Bayes (NB) remain the simplest and most efficient classifiers for text categorization tasks. Their ability to process high-dimensional data with minimal computational cost makes them useful as baseline models and suitable for low-resource environments [34]. Logistic Regression works particularly well when the data is linearly separable, while Naïve Bayes, despite its assumption of feature

independence, has shown strong performance in text classification due to natural word distribution patterns [35]. Support Vector Machines (SVMs) are also effective, especially with sparse input representations such as TF-IDF or n-grams. With proper kernel optimization, SVMs generally outperform simpler models, though they require higher computational resources and careful tuning for large-scale datasets [36].

Tree-based learners, such as Decision Trees and Random Forests, provide interpretability and are capable of learning non-linear feature interactions. While they can capture hierarchical decision rules, they often face overfitting challenges when the dataset is imbalanced or feature engineering is limited. Ensemble methods like Random Forests and XGBoost mitigate these issues by combining multiple learners, thereby improving robustness and reducing bias [37].

Deep learning models have significantly advanced fake news detection because of their ability to capture semantic and syntactic nuances in text. Convolutional Neural Networks (CNNs) excel at learning local patterns, while Recurrent Neural Networks (RNNs) and Long Short-Term Memory networks (LSTMs) capture sequential dependencies, making them effective at modeling context and narrative flow [38]. However, these models demand large labeled datasets and substantial computational resources.

Recent progress in transformer-based architectures—such as BERT, RoBERTa, and GPT—has set new benchmarks in natural language understanding. These models utilize self-attention mechanisms to better capture contextual relationships, offering higher accuracy and F1-scores compared to both traditional classifiers and earlier deep learning models [39] [40]. Despite their superior performance, transformers come with trade-offs: they require powerful GPU infrastructure, consume extensive memory, and often face transparency issues, raising concerns about explainability and fairness in sensitive applications.

To balance performance and interpretability, hybrid and ensemble approaches have gained popularity. For example, combining an SVM with an LSTM, followed by a logistic regression meta-classifier, allows the system to exploit both linear and sequential dependencies. Similarly, bagging and boosting methods such as AdaBoost and XGBoost often outperform standalone models by reducing bias and variance [41].

In real-world scenarios, integrating metadata such as article source, author credibility, and user engagement metrics further enhances detection accuracy. Studies show that blending content-based features with propagation signals (how news spreads across social networks) strengthens system reliability [42].

In summary, while traditional ML models are fast, interpretable, and resource-efficient, they fall short of the contextual understanding achieved by deep learning. Transformers, though computationally expensive, deliver state-of-the-art results and continue to drive advancements in fake news detection [43]. Ultimately, the selection of a model should align with the application requirements—whether prioritizing scalability, transparency, computational feasibility, or real-time processing.

VI. CONCLUSION

The detection of fake news continues to be a critical challenge in the digital information era, as misinformation influences social perceptions, political narratives, and public decision-making. This paper explored multiple approaches based on Machine Learning (ML) and Natural Language Processing (NLP) to classify news articles as genuine or deceptive. The study highlighted that hybrid models, which integrate statistical learning with deep learning techniques, offer significant improvements in detection performance.

Traditional ML algorithms such as Logistic Regression, Naïve Bayes, and Support Vector Machines have shown efficiency in identifying deceptive content using linguistic and statistical features. However, deep learning models, particularly transformer-based architectures like BERT and RoBERTa, demonstrated superior accuracy by effectively capturing contextual and semantic dependencies within text. Preprocessing strategies such as TF-IDF, n-grams, and word embeddings enhanced feature representation, while sentiment analysis and propagation-based methods provided additional insight into emotional and relational patterns present in misinformation.

Despite significant advancements, challenges remain in achieving fully reliable detection. Issues such as adversarial attacks, constantly evolving disinformation tactics, and the need for real-time detection require ongoing research. Future work should emphasize adaptability, multi-modal data integration (combining text, images, and videos), and the use of Explainable AI (XAI) to ensure transparency, accountability, and trustworthiness of detection systems.

In conclusion, the synergy of NLP and ML provides a strong foundation for combating misinformation. Continued innovation in model design, data handling, and interpretability is essential to build scalable and effective systems. The outcomes of this research contribute to the growing global effort to reduce the harmful effects of misinformation and support the creation of a more informed and resilient digital society.

REFERENCES

1. Alghamdi, J., Lin, Y., & Luo, S. (2023). *Towards COVID-19 fake news detection using transformer-based models. Knowledge-Based Systems*, 274, 110642. The paper achieves an impressive F1-score of 98% using CT-BERT with BiGRU layers.
2. Jabar, B. A., Seline, S., Bintang, Victoria, C. J., & Arifin, R. N. (2022). *COVID-19 Fake News Detection With Pre-trained Transformer Models. Ultimatics: Jurnal Teknik Informatika*, 14(2), 51–56. It compares BERT, RoBERTa, DeBERTa, and Electra, finding Electra achieves ~97.8% accuracy.
3. Narayanan, M. B., Ramesh, A. K., Gayathri, K. S., & Shahina, A. (2023). *Fake news detection using a deep learning transformer-based encoder-decoder architecture. Journal of Intelligent and Fuzzy Systems*, 45(5), 8001–8013.
4. Raja, E., Soni, B. S., Borgohain, S. K., & Lalrempuii, C. (2023). *Dravidian Fake News Detection with Gradient Accumulation based Transformer Model. In Proceedings of ICON 2023*, pages 466–471.
5. Yadav, A., Gaba, S., Khan, H., Budhiraja, I., Singh, A., & Singh, K. K. (2022). *ETMA: Efficient Transformer Based Multilevel Attention framework for Multimodal Fake News Detection.*
6. Choi, E., & Kim, J.-K. (2024). *TT-BLIP: Enhancing Fake News Detection Using BLIP and Tri-Transformer.* ArXiv preprint — introduces a multimodal model combining text and image features using BLIP and a tri-transformer.
7. Huang, T., Xu, Z., Yu, P., Yi, J., & Xu, X. (2025). *A Hybrid Transformer Model for Fake News Detection: Leveraging Bayesian Optimization and Bidirectional Recurrent Unit.* ArXiv preprint — achieves up to 99.73% test accuracy by combining transformer, BiGRU, and Bayesian optimization.
8. Saadi, A., Belhadeh, H., Guessas, A., & Hafirassou, O. (2025). *Enhancing Fake News Detection with Transformer Models and Summarization. Engineering, Technology & Applied Science Research*, 15(3), 23253–23259.
9. Rout, J., Mishra, M., & Saikia, M. J. (2025). *Towards Reliable Fake News Detection: Enhanced Attention-Based Transformer Model. Journal of Cybersecurity and Privacy*, 5(3), 43.
10. Papageorgiou, E., Varlamis, I., & Chronis, C. (2025). *Harnessing Large Language Models and Deep Neural Networks for Fake News Detection. Information*, 16(4), 297.
11. Buddhadev, M., & Parekh, V. (2025). *Fake News Detection: Benchmarking Machine Learning and Deep Learning Approaches. ESP Journal of Engineering & Technology Advancements*, 5(2), 39–46.
12. Anonymous (2025). *An Unsupervised Fake News Detection Framework Based on Structural Contrastive Learning. Cybersecurity*, 8:41. Proposes a no-annotation model combining propagation structure and contrastive learning.
13. Tummala, M., Malkhed, M. E., & Vekkot, S. (2025). *Hybrid Semantic and Contextual Analysis for Multilingual Fake News Detection Using Deep Learning. In ICDT 2025.* Utilizes Hindi, Tamil, and English datasets achieving ~97%, 84%, and 93% accuracy respectively.
14. Chen, L., Zhang, X., & Wang, Y. (2023). *Transformer-driven fake news detection across multiple languages. Journal of Information Security and Applications*, 75, 103573.
15. Das, R., & Mehta, S. (2024). *A cross-domain study on misinformation detection using hybrid transformer models. Expert Systems with Applications*, 233, 120947.
16. Li, H., Kumar, A., & Zhao, Q. (2023). *Fake news detection with contextual embeddings and attention networks. Neural Computing and Applications*, 35(12), 8563–8578.
17. Singh, A., Choudhury, P., & Verma, R. (2024). *Leveraging BERT and BiLSTM for enhanced fake news classification. Procedia Computer Science*, 225, 190–198.
18. Wu, Y., Gao, S., & Lin, J. (2023). *Transformer-based cross-lingual misinformation detection framework. Information Sciences*, 639, 119020.

18. Patel, N., Sharma, K., & Gupta, V. (2024). Comparative evaluation of transformer and graph neural networks for detecting fake news. *Applied Intelligence*, 54(3), 2681–2694.
19. Roy, M., & Chakraborty, S. (2023). Hybrid semantic attention models for COVID-19 misinformation detection. *Social Network Analysis and Mining*, 13(2), 88.
20. Zhang, T., & Lee, H. (2024). Transformer models enhanced with contrastive learning for fake news identification. *Journal of Big Data*, 11(5), 215.
21. Ahmed, R., & Khan, M. (2025). Ensemble transformer frameworks for reliable fake news detection. *Engineering Applications of Artificial Intelligence*, 135, 109868.
22. Brown, D., & Cooper, J. (2023). Investigating multimodal approaches for detecting misinformation using transformers. *Multimedia Tools and Applications*, 82(7), 9491–9507.
23. Iqbal, T., & Rahman, A. (2024). Cross-lingual fake news detection using multilingual transformers and adversarial training. *Future Generation Computer Systems*, 149, 275–286.
24. Wang, P., & Zhou, L. (2025). Attention-guided deep learning framework for fake news detection in social media. *Knowledge and Information Systems*, 67(1), 145–162.
25. Silva, R., & Costa, F. (2023). Transformer and reinforcement learning integration for misinformation analysis. *International Journal of Data Science and Analytics*, 16(4), 505–518.
26. Meena, R., & Thomas, G. (2024). Fake news detection through hybrid transformer and convolutional architectures. *Neural Processing Letters*, 56(3), 2135–2149.
27. Ali, K., & Hussain, S. (2023). Enhancing fake news detection with deep contextualized embeddings and transformer models. *Transactions on Asian and Low-Resource Language Information Processing*, 22(1), 16.
28. Zhang, M., & Sun, Y. (2025). Transformer-based model fusion for misinformation detection on short texts. *Information Processing & Management*, 62(3), 103127.
29. Prasad, V., & Menon, A. (2023). Fake news detection with attention-enhanced multilingual transformer models. *International Journal of Artificial Intelligence in Education*, 33(6), 1871–1884.
30. Gao, Y., & Li, Z. (2024). Contrastive self-supervised learning for transformer-based misinformation detection. *Pattern Recognition Letters*, 172, 46–54.
31. Banerjee, S., & Mukherjee, A. (2023). Transformer models for fake news detection: An interpretability perspective. *Journal of Computational Social Science*, 6(2), 289–307.
32. Park, J., & Choi, H. (2025). Fake news detection via multimodal transformer with hierarchical attention. *Multimedia Systems*, 31(1), 87–101.
33. Gupta, R., & Das, P. (2024). Detecting misinformation using hybrid transformers and knowledge graphs. *Journal of Web Semantics*, 78, 100778.
34. Luo, J., & Feng, L. (2023). Graph-enhanced transformer models for social media fake news detection. *IEEE Access*, 11, 44521–44533.
35. Kaur, P., & Malhotra, N. (2024). Comparative analysis of transformer variants for multilingual fake news classification. *Language Resources and Evaluation*, 58(3), 911–929.
36. Zhao, W., & Lin, D. (2023). Transformer-based multimodal misinformation detection integrating sentiment and credibility cues. *Information Fusion*, 98, 101832.
37. Oliveira, F., & Martins, J. (2025). A hybrid transformer architecture with attention calibration for fake news detection. *Applied Soft Computing*, 151, 111978.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details