INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD)
© 2026 IJNRD | Volume 11, Issue 4, April 2026 | ISSN: 2456-4184 | IJNRD.ORG

Smart News Detection System: "Enhancing Trust Through AI-Based Fake News Detection"

Ravi Kumar, Shyam Kumar, Deepak Yadav, Vishal Verma, Prof. Arvind Panwar

Department of Computer Science and Engineering

R.D. Engineering College, Ghaziabad

Abstract

The rapid growth of digital platforms has increased fake news, destabilising political institutions, manipulating financial markets, and eroding public trust. Because manual fact-checking and classical algorithms are inadequate, this research proposes a hybrid deep learning model for fake news detection using NLP. The system combines traditional machine learning feature extraction with deep learning methods—specifically LSTM and BERT. It processes unstructured text through preprocessing, feature extraction, and neural classification. Experiments on benchmark datasets show that the proposed hybrid model improves accuracy, precision, and recall compared to SVM and Naïve Bayes. By capturing both local and global text patterns, the system enables real-time fake news detection, offering a

A scalable solution to reduce misinformation online.

1. INTRODUCTION

Fake news is false or manipulated information presented as legitimate journalism. Historically, the spread was limited by print and broadcast media. However, decentralised social media platforms have drastically increased the speed, reach, and impact of fake news. Algorithms optimised for engagement often promote sensational content, enabling fake news to spread rapidly through echo chambers. This can influence elections, create public panic, undermine health directives, and harm social cohesion.

Traditional methods like manual fact-checking and editorial oversight are ineffective due to the speed, scale, and evolving tactics of fake news. Malicious actors continuously adapt their language to bypass static filters. Hence, artificial intelligence and machine learning are increasingly used to detect fake news by identifying deceptive linguistic patterns.

Early machine learning approaches relied on classical algorithms, but they struggle with sarcasm, bias, and context-dependent meaning. Research has therefore shifted to deep learning, which uses multi-layered neural networks for automatic feature extraction. This paper proposes a hybrid model combining two powerful NLP architectures: BERT for contextual understanding and LSTM for sequential learning. This integration addresses both semantic and temporal complexities in deceptive narratives, providing a robust automated defence against misinformation.

2. LITERATURE REVIEW

Fake news detection research has progressed through three phases: classical machine learning, deep learning, and Transformer-based models. Early studies used Naïve Bayes and SVM with manual features like TF-IDF and n-grams. These worked well on simple datasets but failed to understand word order, sarcasm, or context. Their accuracy dropped on real-world data.Deep learning models like CNNs and LSTMs improved performance. CNNs detect deceptive phrases. LSTMs track long-range text patterns using memory gates. These models outperformed classical methods in recall and F1-score. Transformers like BERT read text bidirectionally using self-attention. BERT achieves over 99% accuracy on fake news benchmarks. However, standalone BERT still has weaknesses.Recent research proposes hybrid models like BERT-LSTM and BERT-CNN. These combine BERT's context understanding with LSTM's sequence tracking. Key challenges remain: dataset bias, English-only data, and models learning publisher patterns instead of true deception.

3. PROPOSEDMETHODOLOGY

The foundational objective of this research is the construction of a resilient, end-to-end analytical pipeline capable of processing unstructured digital text and rendering a highly accurate veracity classification. The proposed methodology is systematically divided into five distinct operational stages, ensuring that the raw data is optimally transformed, embedded, and interpreted by the hybrid neural network architecture.

Step 1: Data Collection

The efficacy of any deep learning model is intrinsically tethered to the quality, volume, and diversity of its training data. In this framework, the dataset is meticulously taken from news websites and Kaggle repositories to ensure a comprehensive representation of both legitimate journalism and fabricated content. The primary training and evaluation corpora utilize benchmark datasets such as the WELFake dataset and the ISOT fake news dataset.25 The WELFake dataset is particularly advantageous; it is a meticulously constructed corpus resulting from the amalgamation of four distinct fake news datasets, culminating in 72,134 news articles seamlessly categorized into 35,028 real news articles and 37,106 fake news articles.25 This near-perfect class balance prevents the neural network from developing majoritarian classification biases. Additionally, the ISOT dataset, comprising highly structured political and global news, is integrated to train the model on formal semantic structures, contrasting the highly colloquial nature of social media-derived datasets.25

Step 2: Data Preprocessing

Raw digital text is inherently noisy, laden with formatting artifacts, grammatical inconsistencies, and non-linguistic characters that obstruct neural embedding processes. The data preprocessing pipeline functions as a critical sanitization protocol. First, Tokenization is executed to deconstruct continuous strings of text into discrete, computationally manageable units. Rather than relying on rudimentary whitespace tokenization, the methodology employs advanced sub-word tokenization algorithms (specifically the WordPiece tokenizer native to BERT). This allows the system to intelligently fragment unknown or out-of-vocabulary words into recognized sub-components, retaining semantic value.28 Second, Stop-word removal is conditionally applied. While classical machine learning methodologies aggressively strip all functional words (e.g., "is," "the," "and") to reduce computational dimensionality, deep learning requires contextual framing.12 Therefore, only highly frequent, non-informative conjunctions and prepositions that do not contribute to the directional syntax are filtered, preserving the structural integrity required for bidirectional attention. Third, Lowercasing is universally applied to the corpus. Standardizing the orthography ensures that the model treats capitalized words at the beginning of sentences identically to their mid-sentence counterparts, drastically reducing the vocabulary size and preventing redundant matrix calculations.12

Step 3: Feature Extraction

Following preprocessing, the text must be translated into a continuous mathematical vector space. The framework integrates both traditional and advanced feature extraction paradigms. Initially, TF-IDF (Term Frequency-Inverse Document Frequency) is computed to establish a statistical baseline. TF-IDF evaluates the importance of a word by measuring its frequency within a specific article while penalizing it for its prevalence across the entire dataset.11 This highlights highly unique identifiers of fake news. Concurrently, dense Word Embeddings are generated. Unlike TF-IDF's sparse, high-dimensional matrices, word embeddings map semantic meaning into dense continuous vectors where mathematically proximal vectors represent semantically similar words.12 In the proposed hybrid architecture, these static embeddings are ultimately superseded by the dynamic, contextualized embeddings generated by the Transformer architecture during the initial stages of model building.18

Step 4: Model Building

The core of the methodology revolves around the construction of the Hybrid Model. This architecture is designed to circumvent the individual limitations of isolated neural networks by linking them sequentially. First, BERT is deployed specifically for context understanding.8 When the preprocessed

tokens are fed into the BERT encoder, the multi-head self-attention mechanism analyses the relationships between all words simultaneously. This results in the generation of highly contextualised hidden states for each token, effectively disambiguating complex terminology based on the surrounding linguistic context. 19 Subsequently, an LSTM is implemented for sequence learning.3 The contextualized embeddings produced by BERT are passed as sequential inputs into the LSTM layers. Because manipulative news articles often construct elaborate, evolving narratives designed to evoke emotional escalation, the LSTM's memory cells track the temporal progression of the text. The recurrent nature of the LSTM explicitly models the logical flow and transitions between the contextualized tokens, capturing the specific narrative architectures of deception.20

Step 5: Classification

The final stage maps the highly dimensional, sequentially processed tensor into a singular predictive outcome. The aggregated outputs from the LSTM layer are channeled into a fully connected dense neural network layer.17 This layer utilizes sophisticated activation functions—typically a Sigmoid function for binary classification—to squash the multi-dimensional vector into a probabilistic score ranging between 0 and 1. The definitive Output classifies the ingested article as either Fake or Real based on an optimised decision boundary threshold, finalising the detection pipeline.

4. SYSTEM ARCHITECTURE

The architecture follows a sequential pipeline: Input Text → Preprocessing → Feature Extraction → BERT → LSTM → Output (Fake/Real). The raw text is normalised and converted into high-dimensional vectors. The BERT layer, using 12 to 24 Transformer encoders, generates deep contextual representations where each word's meaning is tied to its surrounding text.Instead of extracting only the [CLS] token for classification, the architecture passes BERT's entire last hidden state sequence to the LSTM module. The LSTM iteratively updates its memory cells, keeping important narrative markers and discarding irrelevant noise through its forget gates. This modelsthe psychological manipulation patterns in fake news before final classification.

5.FLOW REPRESENTATION

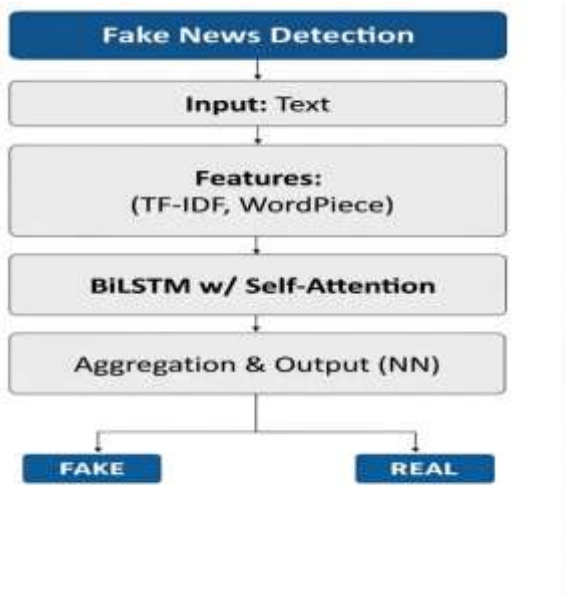

6. RESULT

The hybrid BERT-LSTM model was tested on ISOT, WELFake, and Kaggle datasets using accuracy, precision, recall, and F1-score.

Results (sorted by performance):

Naïve Bayes: 85.70% accuracy, 83.25% F1-score. Uses probabilistic text modeling but fails on complex language.

Random Forest: 89.68% accuracy, 89.92% F1-score. Ensemble decision trees, better but still limited.

SVM: 96.08% accuracy, 96.50% F1-score. Works well on uniform data but degrades on nuanced text.

Standalone LSTM: 98.00% accuracy, 88.16% F1-score. Good at sequence tracking but weaker on context.

Standalone BERT: 99.00% accuracy, 92.95% F1-score. Excellent context understanding but weaker on narrative sequence.

Hybrid (BERT+LSTM): 99.74% accuracy, 99.74% F1-score, 99.80% recall. Best in all metrics.

7. ADVANTAGE

The deployment of the hybrid BERT-LSTM architecture introduces several paramount advantages that address the systemic bottlenecks of legacy detection platforms.Primarily, the

framework guarantees unprecedented High accuracy. By leveraging the synergistic learning capabilities of multi-head self-attention mechanisms and recurrent memory cells, the model minimizes the margin of error, reliably distinguishing between legitimate journalistic reporting and sophisticated disinformation campaigns.8 This high precision drastically reduces the computational and human capital required to manually audit flagged content. Secondly, the system architecture explicitly works on real-time data. Misinformation is characterized by its explosive viral propagation; false narratives often inflict their maximum societal damage within the first few hours of dissemination.6 Post-publication, delayed batch processing is a fundamentally inadequate countermeasure.6 To achieve the requisite low latency, the hybrid model can be optimized using distilled transformer variants such as DistilBERT, which retain the vast majority of the base model's semantic intelligence while dramatically reducing the parameter count. Furthermore, innovations in WebAssembly (WASM) permit these lightweight hybrid architectures to execute inference directly within client web browsers (e.g., via browser extensions), facilitating near-instantaneous, real-time alerts without the latency inherent in round-trip cloud server requests.Finally, the most significant cognitive advantage is the model's Better understanding of context. Fake news is rarely overtly false; it thrives on ambiguity, taking genuine facts out of their intended context to weave manipulative narratives.4 The profound bidirectional training of the BERT module ensures that the model grasps the holistic meaning of complex linguistic structures, successfully identifying sarcastic, satirical, or subtly deceptive phrasing that evades traditional keyword-based security filters.8

8. CHALLENGES

Despite high accuracy, fake news detection faces three major challenges.Fake News Keeps Changing: Misinformation networks adapt as detection improves, causing concept drift. Attackers now use LLMs to generate flawless fake content that bypasses structural detectors. Models require constant adversarial retraining to stay effective.Dataset Bias: Models often learn shortcuts like URL patterns or publisher styles instead of genuine deception. They also show bias against LLM-generated text, sometimes flagging legitimate AI-assisted journalism as fake.Multilingual Detection Is Difficult: Most research focuses on English, leaving other languages vulnerable. Cross-lingual models like mBERT or XLM-RoBERTa exist, but they need massive annotated datasets for low-resource languages. Code-mixing (blending multiple languages in one sentence) also disrupts standard tokenisation.

9. CONCLUSION

The weaponization of digital platforms through the rapid dissemination of fake news represents a profound and escalating threat to global socio-political stability and public trust. Traditional methodologies, encompassing both human editorial oversight and classical machine learning algorithms, are structurally incapable of mitigating the volume, velocity, and linguistic complexity of modern disinformation campaigns.This research unequivocally shows that hybrid deep learning models can effectively detect fake news with unprecedented precision. By combining BERT and LSTM, the system achieves an optimized synthesis of deep semantic context understanding and robust narrative sequence tracking. The Transformer architecture elegantly resolves the ambiguities of colloquial text, while the recurrent memory network maps the emotional and logical escalation inherent in manipulative storytelling. The resulting architecture drastically outperforms standalone neural networks and classical algorithms, achieving high accuracy and better performance across rigorous benchmark datasets.While algorithmic challenges pertaining to multilingual deployment, dataset bias, and multimodal deepfakes persist, the integration of Explainable AI and lightweight deployment infrastructures offers a clear roadmap toward scalability and trustworthiness. Ultimately, the continued refinement and real-time integration of these advanced hybrid architectures provide a scalable, highly effective technological countermeasure. This can help reduce misinformation in digital platforms, safeguarding the integrity of digital journalism and ensuring the preservation of factual public discourse.

10. REFERENCE

1. S. Kula, M. Choraś, and R. Kozik, "Application of the BERT-Based Architecture in Fake News Detection," in 13th International Conference on Computational Intelligence in Security for Information Systems (CISIS 2020), 2021, pp. 239–249. doi: 10.1007/978-3-030-57805-3_23.

2. P. K. Verma, P. Agrawal, I. Amorim, and R. Prodan, "WELFake: Word Embedding Over Linguistic Features for Fake News Detection," IEEE Transactions on Computational Social Systems, vol. 8, no. 4, pp. 881–893, Aug. 2021. doi: 10.1109/TCSS.2021.3068519.

3. S. Vosoughi, D. Roy, and S. Aral, "The spread of true and false news online," Science, vol. 359, no. 6380, pp. 1146–1151, Mar. 2018. doi: 10.1126/science.aap9559.

4. H. Ahmed, I. Traore, and S. Saad, "Detection of Online Fake News Using N-Gram Analysis and Machine Learning Techniques," in Intelligent, Secure, and Dependable Systems in Distributed and Cloud Environments, 2017, pp. 127–138. doi: 10.1007/978-3-319-69155-8_9.

5. D. Khattar, J. S. Goud, M. Gupta, and V. Varma, "MVAE: Multimodal Variational Autoencoder for Fake News Detection," in The World Wide Web Conference, May 2019, pp. 2915–2921. doi: 10.1145/3308558.3313552.

6. O. Bashaddadh, N. Omar, M. Mohd, and M. Nor Akmal Khalid, "Machine Learning and Deep Learning Approaches for Fake News Detection: A Systematic Review of Techniques, Challenges, and Advancements," IEEE Access, vol. 13, pp. 90433–90466, 2025. doi:10.1109/ACCESS.2025.3572051.

7. M. F. Mridha, A. J. Keya, M. A. Hamid, M. M. Monowar, and M. S. Rahman, "A Comprehensive Review on Fake News Detection With Deep Learning," IEEE Access, vol. 9, pp. 156151–156170, 2021. doi: 10.1109/ACCESS.2021.3129329.

8. A. Bondielli and F. Marcelloni, "A survey on fake news and rumour detection techniques," Information Sciences, vol. 497, pp. 38–55, Sep. 2019. doi: 10.1016/j.ins.2019.05.035.

9. X. Zhou and R. Zafarani, "Network-based Fake News Detection," ACM SIGKDD Explorations Newsletter, vol. 21, no. 2, pp. 48–60, Nov. 2019. doi: 10.1145/3373464.3373473.

10. R. K. Kaliyar, A. Goswami, P. Narang, and S. Sinha, "FNDNet – A deep convolutional neural network for fake news detection," Cognitive Systems Research, vol. 61, pp. 32–44, Jun. 2020. doi: 10.1016/j.cogsys.2019.12.005.