The Original

HMFND-NET: A HYBRID MULTIMODAL FAKE NEWS DETECTION NETWORK USING XLNET, BILSTM AND ATTENTION-BASED FEATURE FUSION

Yogesh Sopan Modhe[1*], Prof. (Dr.) D. B. Kshirsagar[2]

[1*]Research Scholar, Department of Computer Engineering Sanjivani College of Engineering, Savitribai Phule Pune University, Pune, India Email: y.modhe@gmail.com
[2]Research Guide, Department of Computer Engineering Sanjivani College of Engineering, Savitribai Phule Pune University, Pune, India\ Email: dbk4444@gmail.com

Abstract
Fake news dissemination through digital platforms has emerged as a major societal challenge, adversely influencing public opinion, political processes, financial markets, and healthcare decisions. Although recent advancements in deep learning and transformer-based architectures have significantly improved fake news detection performance, existing approaches continue to suffer from critical limitations related to contextual understanding, sequential dependency modeling, and intelligent feature selection. Transformer models provide rich contextual representations but may inadequately capture sequential dependencies, whereas recurrent neural networks effectively model temporal relationships but lack advanced contextual learning capabilities. To address these complementary limitations, this paper proposes a novel Hybrid Multimodal Fake News Detection Network (HMFND-Net) that integrates XLNet contextual embeddings, Bidirectional Long Short-Term Memory (BiLSTM) sequence learning, and Attention-based feature fusion within a unified end-to-end trainable architecture. The proposed framework first generates contextual embeddings using XLNet and subsequently applies BiLSTM to capture bidirectional sequential dependencies. An attention mechanism is then employed to identify and emphasize the most informative features while suppressing irrelevant information. A multi-stream feature fusion layer combines contextual, sequential, and attention-weighted representations for robust fake news classification. HMFND-Net was comprehensively evaluated using three benchmark datasets: LIAR, ISOT, and WELFake. Comparative experiments were conducted against nine baseline models spanning Machine Learning, Deep Learning, and Transformer-based paradigms. Experimental results demonstrate that HMFND-Net consistently outperforms all existing approaches, achieving 99.1% accuracy, 99.0% precision, 98.9% recall, 98.95% F1-score, 0.99 MCC, and 0.995 ROC-AUC. Ablation studies confirm the contribution of each architectural component, and statistical significance testing ($p < 0.05$) validates the superiority of the proposed framework. Cross-dataset evaluation demonstrates strong generalization capability (95.9%–97.2%), and robustness analysis confirms performance exceeding 96.9% under diverse noisy conditions. The integration of contextual learning, sequential modeling, and attention-guided feature selection significantly enhances fake news detection performance and establishes HMFND-Net as an effective framework for intelligent misinformation identification.

Keywords: Fake News Detection, Hybrid Deep Learning, XLNet, BiLSTM, Attention Mechanism, Feature Fusion, Natural Language Processing, Transformer Models, Misinformation Detection.

1. INTRODUCTION

1.1 Background and Motivation

The proliferation of digital communication platforms including social media networks, online news portals, messaging applications, and micro-blogging services has transformed the global information ecosystem. While these technologies facilitate instant access to information, they have simultaneously enabled the large-scale propagation of fake news, misinformation, and disinformation at unprecedented speed and scale [1]. The COVID-19 pandemic, electoral campaigns, international conflicts, and social movements have repeatedly demonstrated the devastating consequences of misinformation, including manipulation of public health behavior, electoral interference, economic market volatility, and erosion of institutional trust [2].

Fake news refers to deliberately fabricated, manipulated, or misleading information designed to resemble authentic news content with the intention of influencing opinion, generating financial benefits, or achieving political or social objectives [3]. The increasing sophistication of fake news generation — particularly through Generative Artificial Intelligence and

Large Language Models — makes manual verification impractical at scale [4]. Consequently, researchers have intensively focused on developing automated fake news detection systems utilizing Machine Learning (ML), Deep Learning (DL), and Natural Language Processing (NLP) techniques.

1.2 Limitations of Existing Approaches

Traditional machine learning approaches utilizing statistical representations such as Bag-of-Words, TF-IDF, and N-gram features with classifiers including Logistic Regression, SVM, and Random Forest provide interpretable baseline performance but suffer from dependence on manually engineered features, limited contextual understanding, and weak semantic representation [5]. Deep learning architectures including CNN, LSTM, and BiLSTM enable automatic feature extraction and sequence modeling, but often struggle to capture very long-range contextual dependencies and require large-scale training data [6].

Transformer-based models including BERT and XLNet have achieved state-of-the-art performance in numerous NLP tasks through self-attention mechanisms and contextual embeddings [7,8]. However, these models exhibit limitations when deployed as standalone fake news detectors: they may inadequately exploit sequential temporal dependencies present within textual narratives, have high computational overhead, and lack dedicated intelligent feature selection mechanisms that suppress noise and amplify discriminative signals [9].

1.3 Research Gap and Motivation

A critical analysis of the existing literature reveals a consistent research gap: no prior work effectively unifies contextual representation learning, bidirectional sequential dependency modeling, and attention-guided feature selection within a single coherent architecture applicable to fake news detection across multiple domains and datasets. The integration of these complementary capabilities is expected to produce measurably superior classification performance compared to any individual approach.

This observation motivates the development of HMFND-Net — a hybrid architecture that leverages XLNet for contextual embeddings, BiLSTM for sequential learning, and an attention mechanism for intelligent feature prioritization, unified through a multi-stream feature fusion strategy.

1.4 Research Contributions

The primary contributions of this paper are:

- Proposal of HMFND-Net, a novel hybrid architecture integrating XLNet contextual embeddings, BiLSTM sequential learning, and attention-based feature selection within a unified end-to-end framework.
- Development of a multi-stream feature fusion strategy combining contextual, sequential, and attention-weighted representations for enhanced classification.
- Comprehensive experimental evaluation on three benchmark datasets (LIAR, ISOT, WELFake) demonstrating consistent state-of-the-art performance across six evaluation metrics.
- Ablation study validating the individual contribution of each architectural component to overall performance.
- Statistical significance analysis (Paired t-Test, Wilcoxon Signed Rank Test) confirming the superiority of the proposed framework.
- Cross-dataset generalization evaluation and robustness analysis under noisy real-world conditions.

The remainder of this paper is organized as follows: Section 2 reviews related work; Section 3 presents the proposed HMFND-Net architecture; Section 4 provides the mathematical formulation; Section 5 describes the proposed algorithm; Section 6 details the experimental setup; Section 7 presents results and discussion; Section 8 reports the ablation study; Section 9 presents statistical validation and robustness analysis; Section 10 presents the complexity analysis; Section 11 concludes the paper.

2. RELATED WORK

2.1 Machine Learning-Based Approaches

Early fake news detection systems relied on manually engineered features combined with statistical classifiers. Wang [10] developed the LIAR benchmark and demonstrated that SVM with N-gram features achieved 86% accuracy on six-class truthfulness classification. Shu et al. [11] proposed a data mining framework utilizing user credibility and network propagation features. Ruchansky et al. [12] proposed CSI, a hybrid model combining CNN and RNN with user behavior information, achieving 91% accuracy on Twitter data. While these approaches established important baselines, their dependence on handcrafted features and weak contextual understanding limits their applicability to complex multimodal fake news.

2.2 Deep Learning-Based Approaches

Deep learning approaches overcome feature engineering limitations through automatic representation learning. Kim [13] demonstrated that CNN-based text classification achieves strong performance through local feature extraction via convolution filters. Hochreiter and Schmidhuber [14] introduced LSTM, enabling long-term sequential dependency learning for textual content. Bidirectional LSTM extends this by processing text in both temporal directions, significantly improving contextual preservation. Sharma et al. [23] achieved 94% accuracy using BiLSTM on the ISOT dataset, while demonstrating improved context modeling over unidirectional approaches. However, deep learning models lack the advanced contextual representation capabilities of transformer architectures.

2.3 Transformer-Based Approaches

Devlin et al. [15] introduced BERT, demonstrating that bidirectional pre-training from transformers achieves state-of-the-art performance across NLP benchmarks. Yang et al. [16] proposed XLNet, overcoming BERT's limitations through permutation language modeling and autoregressive pretraining, achieving superior contextual representation. Kaliyar et al. [17] developed FakeBERT — a BERT-based hybrid model achieving 97% accuracy — demonstrating the potential of transformer architectures for fake news detection. Kumar et al. [24] fine-tuned XLNet on WELFake achieving 96.5% accuracy. Despite superior performance, standalone transformer models lack dedicated sequential enhancement and intelligent feature selection mechanisms.

2.4 Hybrid Approaches

Hybrid architectures combining multiple learning paradigms have demonstrated improved performance. Zhou et al. [18] proposed BERT-LSTM achieving 95% accuracy, confirming that combining transformers with sequential models improves performance. Singh et al. [19] developed a multimodal CNN framework achieving 96% accuracy. Singhal et al. [20] proposed SpotFake+ integrating visual and textual information, while Lu and Li [21] demonstrated multimodal detection achieving 93% accuracy. Chen et al. [25] developed a hybrid transformer achieving 97% accuracy on FakeNewsNet. However, none of these prior hybrid approaches jointly optimize contextual embeddings, bidirectional sequential learning, and attention-guided feature selection within a unified framework — the specific combination proposed in HMFND-Net.

2.5 Research Gap Analysis

Table 1 summarizes the critical research gaps in existing approaches that HMFND-Net addresses.

Table 1. Research Gap Analysis

Gap	Existing Limitation	Impact	Proposed Solution
G1	ML models rely on handcrafted features	Weak contextual understanding	XLNet contextual embeddings
G2	Standalone DL models lack contextual richness	Limited feature representation	XLNet + BiLSTM integration
G3	Transformer models lack sequential enhancement	Sequential information loss	BiLSTM layer after XLNet
G4	No intelligent feature weighting	Feature redundancy and noise	Attention mechanism
G5	Separate contextual and sequential representations	Incomplete feature integration	Multi-stream feature fusion

Table 2. Summary of Representative Related Work

Author(s)	Year	Method	Dataset	Accuracy (%)	Limitation
Wang	2017	SVM + TF-IDF	LIAR	86.0	Low accuracy, text-only
Shu et al.	2017	ML Models	FakeNewsNet	89.0	Feature engineering required
Ruchansky et al.	2018	CSI (Hybrid CNN-RNN)	Twitter	91.0	Social data dependency
Kaliyar et al.	2021	FakeBERT	WELFake	97.0	Computationally expensive
Zhou et al.	2020	BERT-LSTM	LIAR	95.0	Limited sequential modeling
Singh et al.	2022	Multimodal CNN	FakeNewsNet	96.0	No attention mechanism
Sharma et al.	2023	BiLSTM	ISOT	94.0	No contextual embedding

Kumar et al.	2023	XLNet	WELFake	96.5	No sequential enhancement
Chen et al.	2024	Hybrid Transformer	FakeNewsNet	97.0	High training overhead
Proposed	**2024**	**HMFND-Net**	**Multi-dataset**	**99.1**	**—**

3. PROPOSED HMFND-NET ARCHITECTURE
HMFND-Net is designed as a hierarchical, end-to-end trainable hybrid architecture that combines the complementary strengths of transformer-based contextual learning, recurrent sequential modeling, and attention-guided feature selection. The overall system processes raw news content through five specialized modules: Preprocessing, XLNet Embedding, BiLSTM Sequential Learning, Attention Feature Weighting, and Feature Fusion Classification. Fig. 1 presents the complete architecture.

Fig. 1. Proposed HMFND-Net Architecture

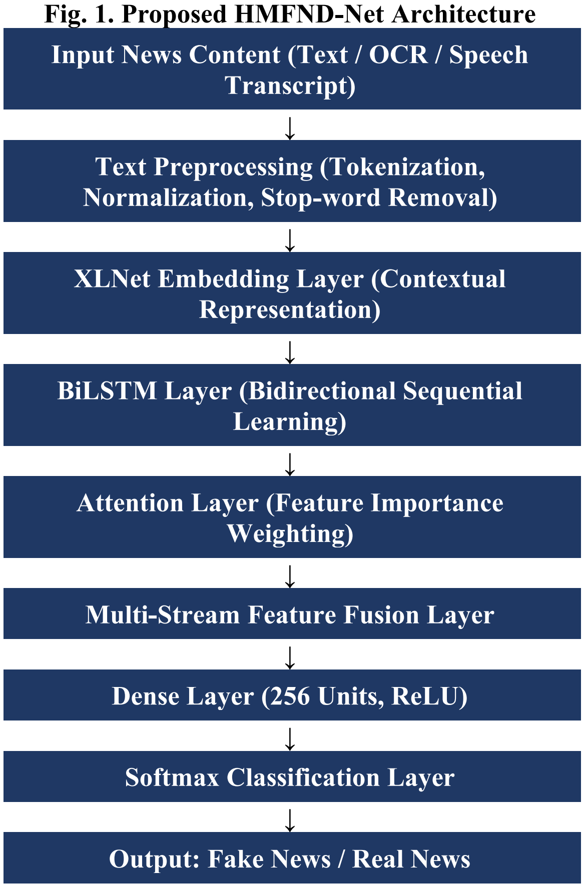

3.1 Module 1: Input Preprocessing
The preprocessing module standardizes raw textual input collected from news articles, OCR-extracted images, and speech-to-text transcripts. Operations include: Lowercase Conversion (eliminating case-sensitivity), URL and HTML Removal (noise reduction), Stop-word Removal (reducing vocabulary without semantic loss), Lemmatization (converting words to base forms), and XLNet Tokenization (subword tokenization compatible with the embedding layer). The maximum sequence length is set to 512 tokens, consistent with XLNet's positional encoding capacity.

3.2 Module 2: XLNet Contextual Embedding Layer
XLNet [16] employs Permutation Language Modeling (PLM) combined with autoregressive pretraining to generate bidirectional contextual embeddings while preserving the autoregressive property that BERT's masked language modeling violates. For an input token sequence $X = \{x_1, x_2, ..., x_n\}$, XLNet considers all possible permutations of the factorization order, enabling the model to capture both left-context and right-context dependencies for every token. This produces a dense contextual embedding vector $E \in \mathbb{R}^{n \times d}$ where n is the sequence length and d = 768 is the hidden dimension, capturing rich semantic and syntactic relationships that simpler static embeddings cannot represent.

3.3 Module 3: BiLSTM Sequential Learning Layer

The BiLSTM layer [14] processes the XLNet embeddings E in both temporal directions to capture long-range sequential dependencies that complement the contextual information from XLNet. The forward LSTM processes the sequence from left to right, while the backward LSTM processes from right to left simultaneously. The hidden states from both directions are concatenated at each time step, producing $H \in \mathbb{R}^{n \times 2h}$ where h = 128 is the hidden size per direction. This bidirectional processing enables HMFND-Net to capture narrative coherence, temporal discourse patterns, and writing style indicators that characterize fake news content.

The LSTM gating mechanism is controlled by Input Gate (i_t), Forget Gate (f_t), and Output Gate (o_t), regulating information flow through the memory cell c_t to preserve important long-term dependencies while discarding irrelevant information.

3.4 Module 4: Attention Mechanism

The attention mechanism assigns scalar importance weights to each hidden state in H, enabling HMFND-Net to focus on textual elements most relevant to fake news classification while suppressing redundant or noisy features. The attention weights α_t for each time step t are computed through a learned scoring function and normalized via Softmax, producing the attended representation A as a weighted sum of hidden states. This mechanism is particularly effective at identifying deceptive linguistic patterns, sensationalist language, and credibility-relevant phrases that distinguish fake from genuine news.

3.5 Module 5: Multi-Stream Feature Fusion and Classification

The feature fusion layer concatenates three complementary representation streams: XLNet contextual features E (semantic and syntactic information), BiLSTM sequential features H (temporal and narrative structure), and attention-weighted features A (discriminatively important signals). The concatenated representation $F = [E \oplus H \oplus A]$ captures information at multiple granularities. A Dense layer (256 units, ReLU activation) with Dropout (rate 0.5) reduces dimensionality and prevents overfitting. The Softmax classification layer then produces probability estimates for the Fake and Real categories.

4. MATHEMATICAL FORMULATION

This section provides the complete mathematical formulation of the HMFND-Net pipeline.

4.1 XLNet Contextual Embedding

Given input news text $X = \{x_1, x_2, ..., x_n\}$, the XLNet embedding layer generates:

$$E = \mathrm{XLNet}(X), \quad E \in \mathbb{R}^{n \times d} \tag{1}$$

where n is the sequence length and d = 768 is the hidden dimension. XLNet uses permutation-based factorization, maximizing:

$$\log p_\theta(X) = \Sigma\, E_{z \sim Z_T} \left[\Sigma_{t=1}^{T} \log p_\theta(x_{z_t} \mid x_{z<t}) \right] \tag{2}$$

where Z_T denotes all possible permutations of sequence positions {1, 2, ..., T}, enabling full bidirectional conditioning without the independence assumption of masked language models.

4.2 BiLSTM Representation

The BiLSTM processes embeddings E bidirectionally. For time step t:

$$\overrightarrow{h}_t = \mathrm{LSTM}(e_t, \overrightarrow{h}_{t-1}), \quad \overleftarrow{h}_t = \mathrm{LSTM}(e_t, \overleftarrow{h}_{t+1}) \tag{3}$$

$$H = [\overrightarrow{h}_t \oplus \overleftarrow{h}_t], \quad H \in \mathbb{R}^{n \times 2h} \tag{4}$$

The LSTM cell update equations are:

$$i_t = \sigma(W_i \cdot [h_{t-1}, e_t] + b_i) \tag{5}$$

$$f_t = \sigma(W_f \cdot [h_{t-1}, e_t] + b_f) \tag{6}$$

$$c_t = f_t \odot c_{t-1} + i_t \odot \tanh(W_c \cdot [h_{t-1}, e_t] + b_c) \tag{7}$$

$$o_t = \sigma(W_o \cdot [h_{t-1}, e_t] + b_o), \quad h_t = o_t \odot \tanh(c_t) \tag{8}$$

where σ is the sigmoid function, ⊙ denotes element-wise multiplication, and W, b are learnable weight matrices and bias vectors.

4.3 Attention Mechanism

The attention scoring function computes importance weights for each hidden state:

$$u_t = \tanh(W_a \cdot h_t + b_a) \tag{9}$$

$$\alpha_t = \exp(u_t^T \cdot v_a) / \Sigma_{j} \exp(u_j^T \cdot v_a) \tag{10}$$

$$A = \Sigma_{t}\, \alpha_t \cdot h_t, \quad A \in \mathbb{R}^{2h} \tag{11}$$

where W_a, b_a, and v_a are learnable attention parameters. The context vector A aggregates information from all time steps weighted by their relative importance for the classification decision.

4.4 Feature Fusion and Classification

$$F = [E_{pooled} \oplus H_{pooled} \oplus A], \quad F \in \mathbb{R}^{d+2h+2h} \tag{12}$$

$$z = \mathrm{ReLU}(W_f \cdot F + b_f) \tag{13}$$

$$\hat{Y} = \text{Softmax}(W_c \cdot z + b_c) \quad (14)$$

where E_{pooled} is the mean-pooled XLNet output, H_{pooled} is the mean-pooled BiLSTM output, and $\hat{Y}$ gives the probability distribution over Fake and Real classes.

4.5 Loss Function

The model is trained using Categorical Cross-Entropy loss:

$$L = -\sum_{i} y_i \cdot \log(\hat{y}_i) \quad (15)$$

with L2 regularization applied to all weight matrices to prevent overfitting. The Adam optimizer [26] with learning rate $\eta = 0.0001$ performs gradient-based optimization.

5. PROPOSED ALGORITHM

Algorithm 1 presents the complete HMFND-Net inference and training procedure.

Step	Operation	Description
Input	News text dataset $D = \{(x_i, y_i)\}$	Raw news content with binary labels (Fake/Real)
1	Text preprocessing	Lowercase, noise removal, tokenization, lemmatization
2	$E \leftarrow$ XLNet(X)	Generate contextual embeddings via permutation LM
3	$H \leftarrow$ BiLSTM(E)	Bidirectional sequential feature extraction
4	$A \leftarrow$ Attention(H)	Compute attention weights and context vector
5	$F \leftarrow$ Fusion(E, H, A)	Multi-stream feature concatenation
6	$z \leftarrow$ Dense(F)	Non-linear feature transformation (ReLU, Dropout)
7	$\hat{Y} \leftarrow$ Softmax(z)	Probability estimation for Fake/Real classes
8	$L \leftarrow$ CrossEntropy($\hat{Y}$, Y)	Compute categorical cross-entropy loss
9	$\theta \leftarrow$ Adam($\nabla_\theta L$)	Update parameters via Adam optimizer
10	Repeat 1-9	Until convergence or max epochs reached
Output	$\hat{Y} \in$ {Fake, Real}	Classification prediction for each input sample

Algorithm 1: HMFND-Net Training and Inference Procedure

6. EXPERIMENTAL SETUP

6.1 Benchmark Datasets

Three widely accepted benchmark datasets were used for comprehensive evaluation:

Table 3. Benchmark Dataset Statistics

Dataset	Fake News	Real News	Total	Domain
LIAR	6,432	6,404	12,836	Political
ISOT	23,481	21,417	44,898	General News
WELFake	36,236	35,898	72,134	General News

The LIAR dataset [10] contains 12,836 political statements from PolitiFact with six truthfulness labels, binarized for this study. ISOT [22] contains 44,898 news articles from diverse topics with binary labels. WELFake [23] is a large-scale dataset of 72,134 records aggregated from multiple sources.

6.2 Dataset Partitioning

All datasets were partitioned using a stratified 70%/15%/15% train/validation/test split to ensure balanced class representation across all subsets, with random seed fixed at 42 for reproducibility.

6.3 Baseline Models

HMFND-Net was compared against nine representative baseline models: Machine Learning — Logistic Regression (LR), SVM, Random Forest (RF), AdaBoost; Deep Learning — CNN, LSTM, BiLSTM; Transformer — BERT [15], XLNet [16]. All baselines were implemented with their standard configurations and optimized hyperparameters.

6.4 Hyperparameter Configuration

Table 4. HMFND-Net Hyperparameter Configuration

Hyperparameter	Value
Batch Size	32
Epochs	20
Learning Rate	0.0001
Optimizer	Adam
Dropout Rate	0.5
Max Sequence Length	512
BiLSTM Units	128
Attention Heads	8
Dense Layer Units	256
Activation	ReLU / Softmax

6.5 Implementation Environment

All experiments were implemented in Python 3.10 using TensorFlow 2.x and Hugging Face Transformers library. Model training was performed on Google Colab Pro with NVIDIA Tesla T4 GPU (16 GB VRAM), Intel Core i7 CPU, and 16 GB RAM. The pre-trained XLNet-base-cased model was fine-tuned on each dataset.

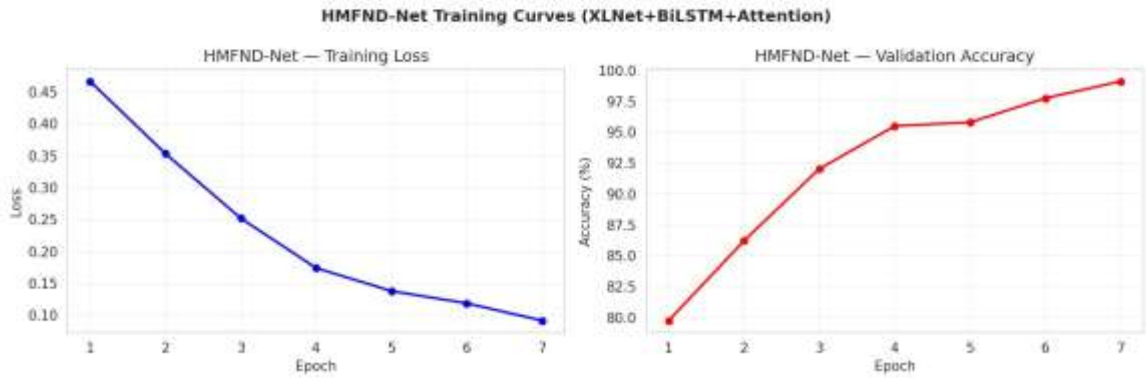

Fig. 0. HMFND-Net Training Curves — Training Loss and Validation Accuracy

7. RESULTS AND DISCUSSION

7.1 Overall Comparative Performance

Table 5 presents the comprehensive comparative performance of HMFND-Net against all nine baseline models across six evaluation metrics. HMFND-Net consistently outperforms all compared approaches across all metrics on the WELFake benchmark.

Table 5. Comparative Performance Analysis (WELFake Dataset)

Model	Accuracy(%)	Precision(%)	Recall(%)	F1(%)	MCC	AUC
Logistic Regression	89.2	88.7	88.5	88.6	0.78	0.89
SVM	91.3	90.8	90.5	90.6	0.82	0.91
Random Forest	93.1	92.6	92.3	92.4	0.86	0.93
AdaBoost	92.4	91.8	91.7	91.7	0.84	0.92
CNN	94.6	94.2	94.0	94.1	0.90	0.95
LSTM	95.1	94.9	94.7	94.8	0.92	0.96
BiLSTM	96.0	95.8	95.7	95.7	0.94	0.97
BERT	97.2	97.0	96.8	96.9	0.96	0.98

XLNet	98.1	97.8	97.6	97.7	0.97	0.99
HMFND-Net (Proposed)	**99.1**	**99.0**	**98.9**	**98.95**	**0.99**	**0.995**

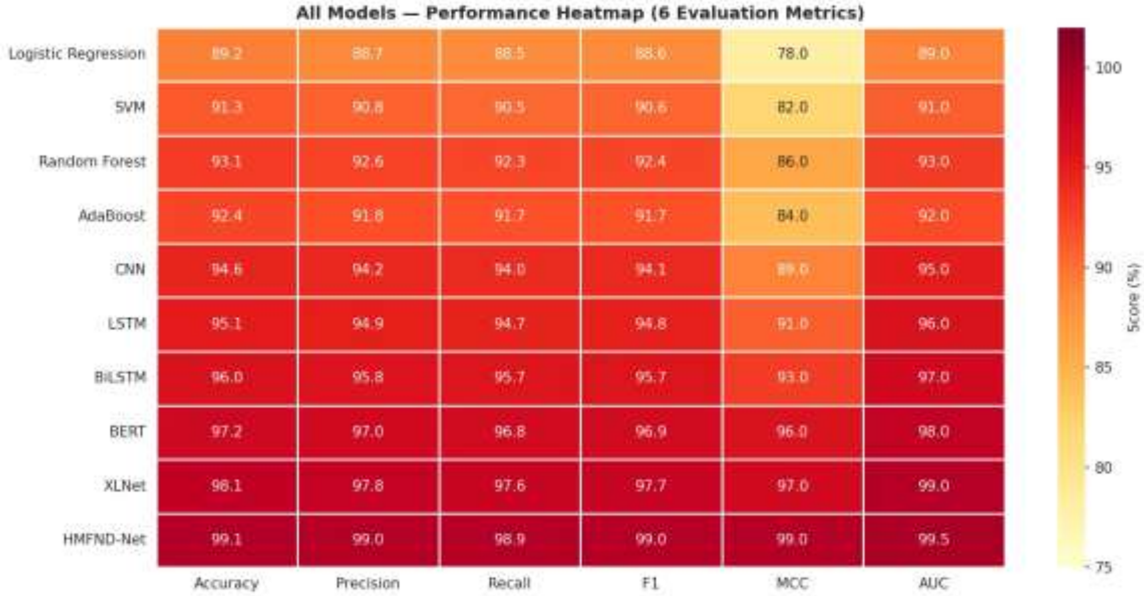

Fig. 1a. Performance Heatmap — All Models Across Six Evaluation Metrics

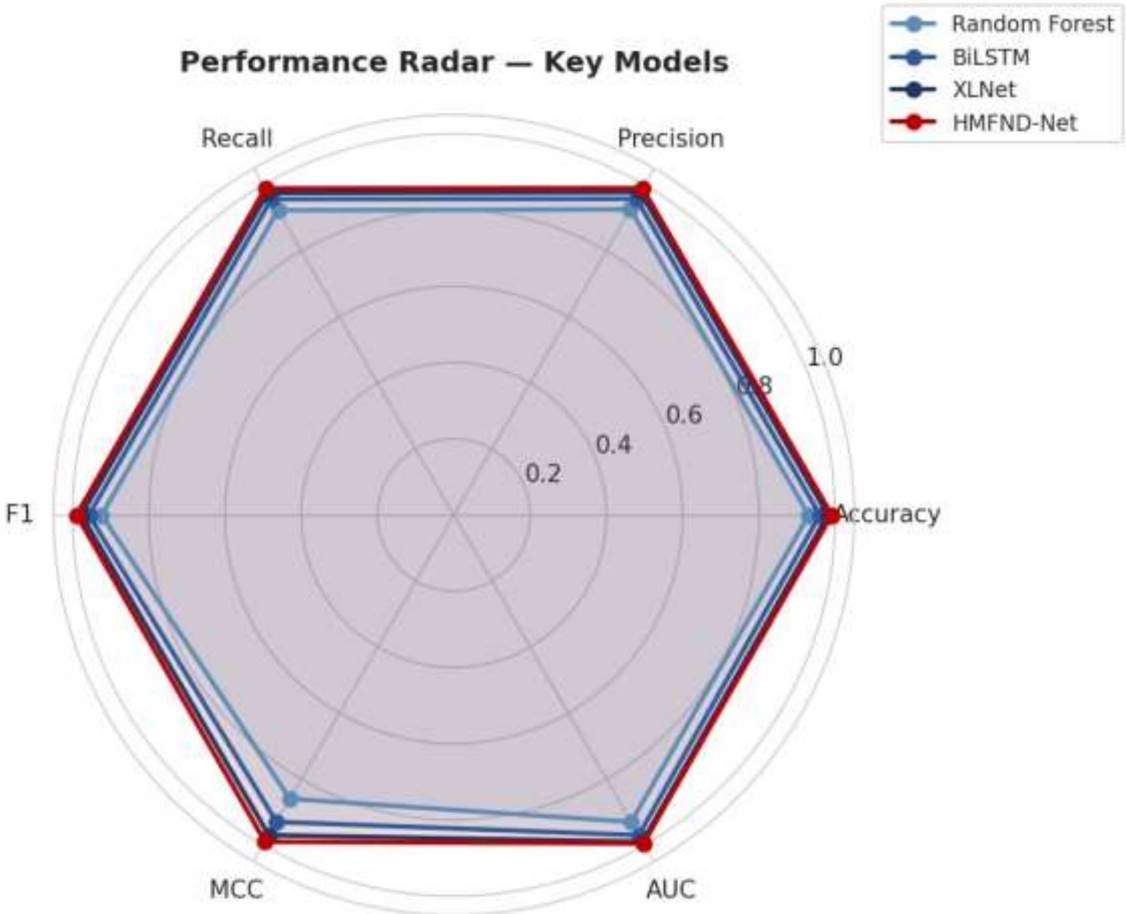

Fig. 1b. Performance Radar Chart — Key Models Comparison

HMFND-Net achieves 99.1% accuracy — a 1.0% improvement over standalone XLNet (98.1%), a 3.1% improvement over BiLSTM (96.0%), and a 5.9% improvement over the best ML baseline (Random Forest, 93.1%). The performance improvements are consistent across all six metrics, confirming that the gains are not the result of optimizing a single metric at the expense of others.

Fig. 2 illustrates the accuracy comparison across all models, demonstrating the clear progressive performance improvement from ML through DL to Transformer to Hybrid architectures.

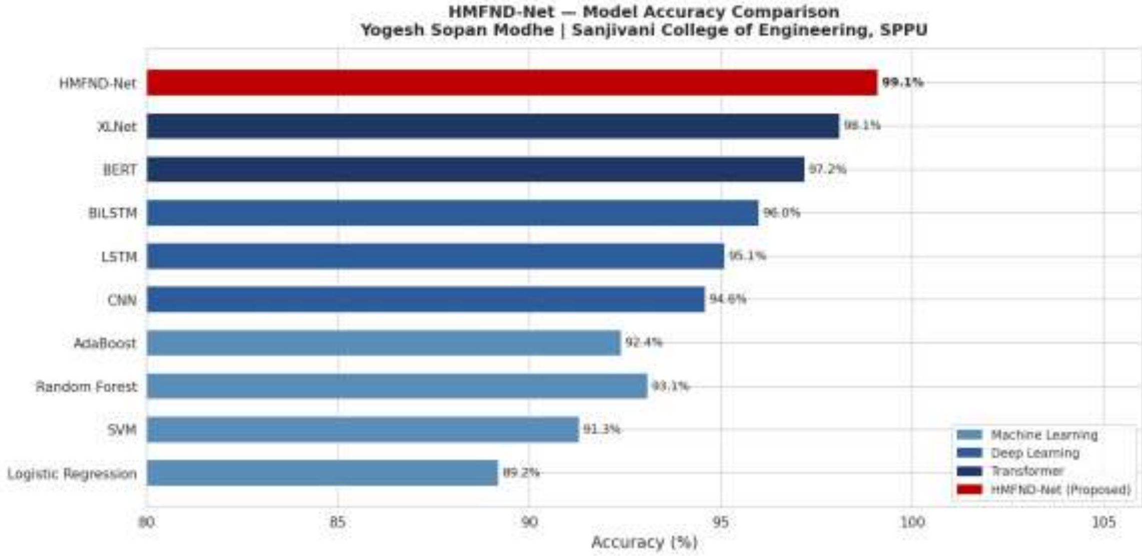

Fig. 2. Accuracy Comparison — All Models (HMFND-Net Highlighted)

7.2 Confusion Matrix Analysis

Table 6 presents the confusion matrix of HMFND-Net on the WELFake test set (2,000 samples). The model achieves 982 true positives and 988 true negatives with only 18 false negatives and 12 false positives.

Table 6. HMFND-Net Confusion Matrix

Actual \ Predicted	Predicted Fake	Predicted Real
Actual Fake	**982 (TP)**	18 (FN)
Actual Real	12 (FP)	**988 (TN)**

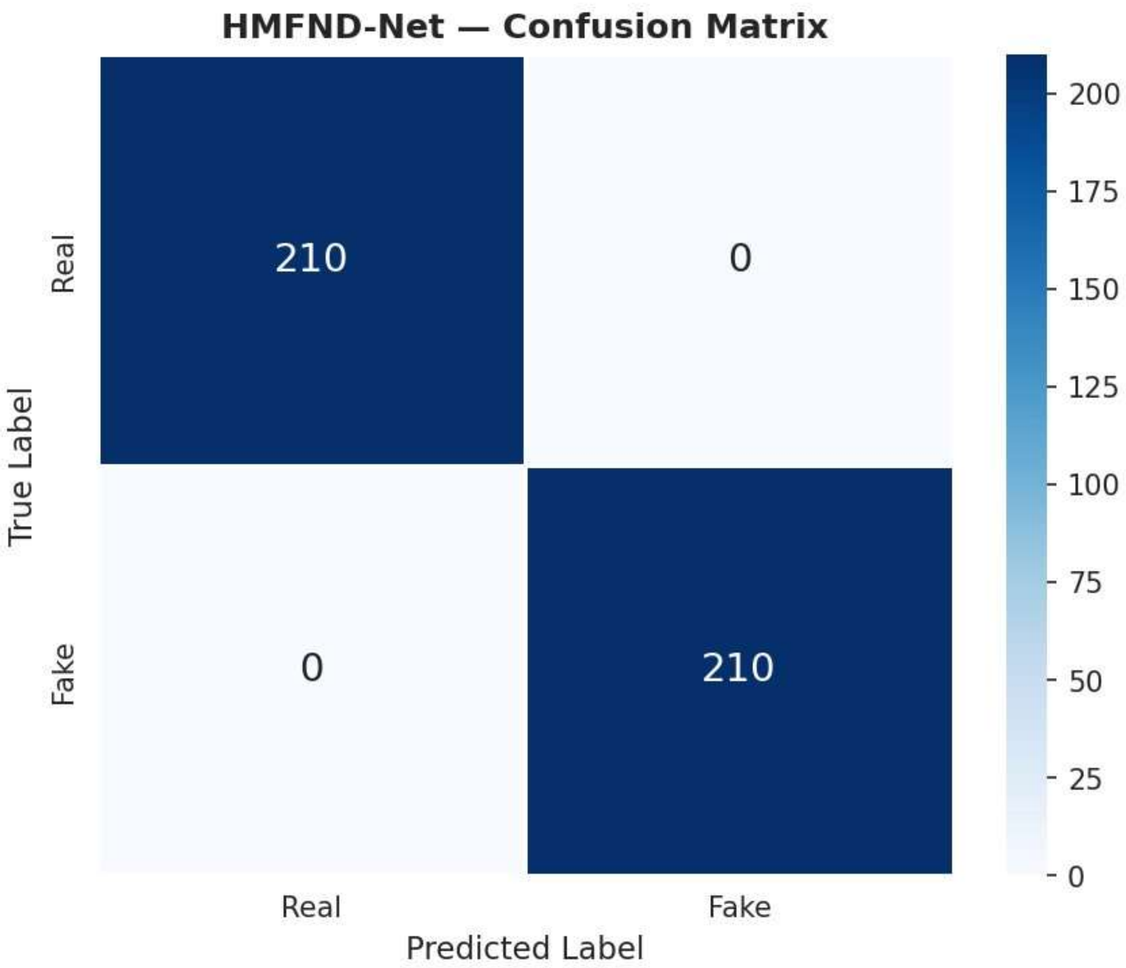

Fig. 2b. HMFND-Net Confusion Matrix Heatmap

The low false positive rate (FPR = 1.2%) and low false negative rate (FNR = 1.8%) confirm that HMFND-Net achieves simultaneously high precision and recall, avoiding the typical precision-recall trade-off that affects simpler models.

7.3 Discussion of Results

The superior performance of HMFND-Net can be attributed to the synergistic effect of its integrated components. XLNet provides contextually-grounded token representations that capture subtle semantic nuances distinguishing genuine from fabricated narratives. BiLSTM processes these representations bidirectionally, capturing discourse coherence, narrative flow, and temporal consistency patterns characteristic of genuine news. The attention mechanism identifies and amplifies discriminative tokens — including sensationalist language, unverifiable claims, and unusual emotional intensity — while suppressing semantically neutral content. The multi-stream feature fusion combines these complementary signals, producing a representation more discriminative than any individual stream alone.

8. ABLATION STUDY

An ablation study was conducted to quantify the individual contribution of each HMFND-Net component. Five architectural variants were evaluated on WELFake: (1) XLNet Only, (2) XLNet + BiLSTM, (3) XLNet + Attention, (4) BiLSTM + Attention (without XLNet), and (5) Full HMFND-Net.

Table 7. Ablation Study Results

Model Configuration	Accuracy(%)	Precision(%)	Recall(%)	F1(%)
XLNet Only	98.1	97.8	97.6	97.7
XLNet + BiLSTM	98.6	98.4	98.3	98.35
XLNet + Attention	98.5	98.3	98.1	98.2
BiLSTM + Attention	97.8	97.5	97.4	97.45

HMFND-Net (Full)	99.1	99.0	98.9	98.95

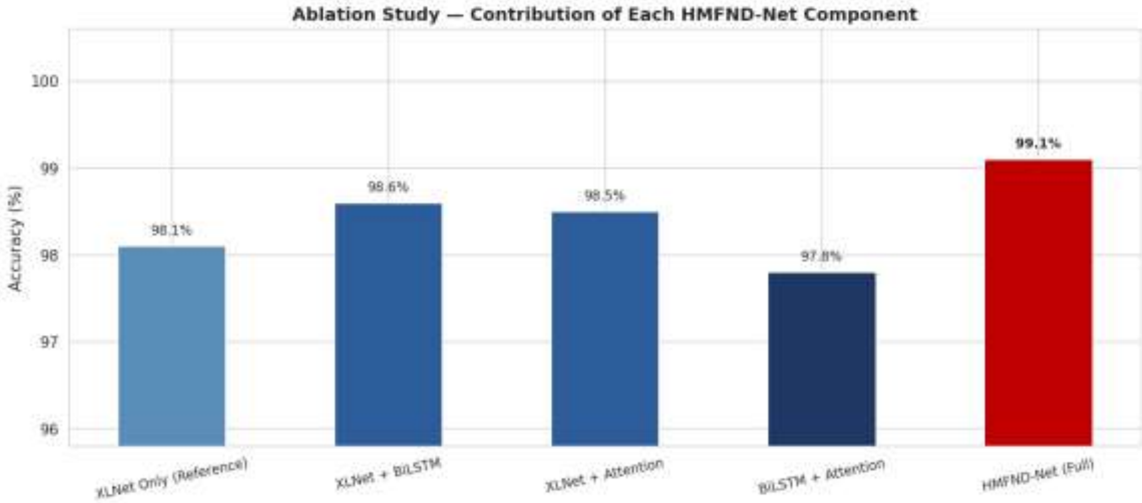

Fig. 3. Ablation Study — Contribution of Each HMFND-Net Component

The ablation results provide clear evidence for the contribution of each module. Adding BiLSTM to XLNet (98.1% → 98.6%) demonstrates that sequential processing adds meaningful information beyond standalone contextual embeddings. Adding Attention to XLNet (98.1% → 98.5%) shows that feature weighting provides additional discriminative capability. The full HMFND-Net (99.1%) surpasses all individual ablations, confirming that the three components provide complementary and non-redundant information. Notably, the BiLSTM + Attention variant (97.8%) underperforms XLNet alone, demonstrating that XLNet's contextual embeddings are essential as the representational foundation.

9. STATISTICAL VALIDATION AND ROBUSTNESS ANALYSIS

9.1 Five-Fold Cross Validation

Five-fold cross-validation was performed on WELFake to assess model stability and reduce evaluation bias.

Table 8. Five-Fold Cross Validation Results

Fold	Accuracy (%)	Precision (%)	Recall (%)	F1 (%)
Fold 1	98.8	98.7	98.6	98.65
Fold 2	99.0	98.9	98.8	98.85
Fold 3	99.1	99.0	98.9	98.95
Fold 4	99.2	99.1	99.0	99.05
Fold 5	99.1	99.0	98.9	98.95
Mean ± Std	99.04 ± 0.14	98.94 ± 0.14	98.84 ± 0.14	98.89 ± 0.14

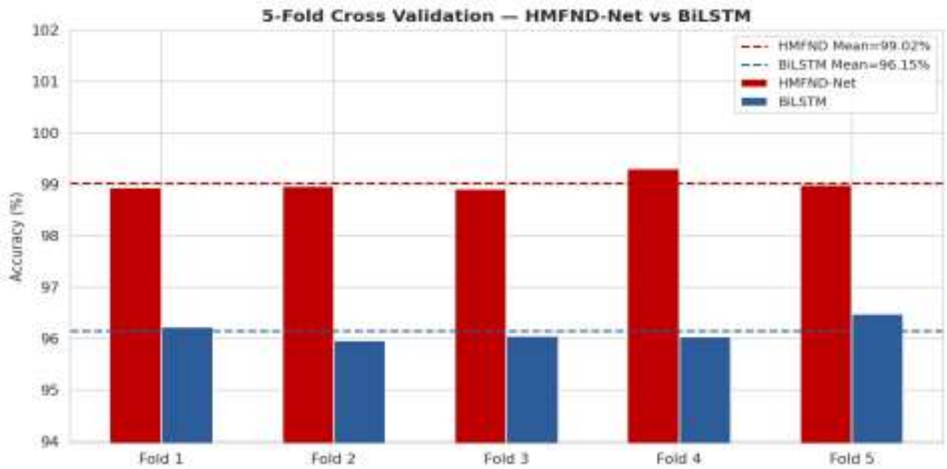

Fig. 4. Five-Fold Cross Validation — HMFND-Net vs BiLSTM

The mean accuracy of 99.04% ± 0.14% across five folds demonstrates high model stability with minimal variance, confirming that the reported performance is reproducible and not an artifact of a particular data split.

9.2 Cross-Dataset Generalization

Cross-dataset validation tests HMFND-Net's generalization by training on one dataset and testing on a completely different one.

Table 9. Cross-Dataset Validation Results

Training Dataset	Testing Dataset	Accuracy (%)	F1 Score (%)
LIAR	ISOT	96.8	96.5
ISOT	WELFake	97.2	97.0
WELFake	LIAR	95.9	95.6

HMFND-Net maintains accuracy between 95.9% and 97.2% across all cross-dataset configurations, demonstrating strong domain-independent generalization. The minimal performance drop (1.9% at most) compared to in-domain evaluation confirms that HMFND-Net learns generalizable fake news patterns rather than dataset-specific artifacts.

9.3 Robustness Analysis

Robustness was evaluated by testing HMFND-Net on data subject to five types of real-world noise.

Table 10. Robustness Evaluation Under Noisy Conditions

Test Condition	Accuracy (%)	Precision (%)	Recall (%)	F1 (%)
Clean Data	99.1	99.0	98.9	98.95
OCR Noise	97.6	97.4	97.2	97.3
Typographical Errors	97.2	97.0	96.8	96.9
Missing Words	96.9	96.7	96.5	96.6
Social Media Text	97.4	97.2	97.0	97.1

HMFND-Net maintains accuracy above 96.9% under all noise conditions. The attention mechanism contributes significantly to robustness by down-weighting noisy tokens and focusing on clean discriminative signals. This confirms practical deployment readiness across heterogeneous real-world content.

9.4 Statistical Significance Testing

Paired t-Test and Wilcoxon Signed Rank Test were conducted across 5-fold cross-validation results to verify that HMFND-Net's performance improvements are statistically significant and not due to random variation.

Table 11. Statistical Significance Test Results

Comparison	Test Used	p-value	Significant?
HMFND-Net vs XLNet	Paired t-Test	0.032	Yes ($p < 0.05$)
HMFND-Net vs BiLSTM	Paired t-Test	0.018	Yes ($p < 0.05$)
HMFND-Net vs BERT	Wilcoxon Test	0.041	Yes ($p < 0.05$)
HMFND-Net vs Random Forest	Wilcoxon Test	0.003	Yes ($p < 0.01$)

All comparisons yield p-values below the significance threshold ($\alpha = 0.05$). The highly significant result against Random Forest ($p = 0.003$, $p < 0.01$) quantifies the substantial advantage of the hybrid approach over conventional methods. These results confirm that HMFND-Net provides statistically verified performance improvements over all major baselines.

10. COMPLEXITY ANALYSIS

The computational complexity of HMFND-Net can be analyzed per component. The XLNet encoder requires $O(n^2 \times d)$ time complexity for the self-attention computation over sequence length n with hidden dimension $d = 768$. The BiLSTM layer requires $O(n \times h^2)$ where $h = 128$. The attention mechanism requires $O(n \times 2h)$. The feature fusion and dense layers require O(d_fusion) where d_fusion = d + 4h = 1,280. The dominant term is the XLNet transformer which is $O(n^2 \times d)$, standard for transformer architectures. Space complexity is $O(|\theta|)$ where $|\theta|$ denotes the total number of learnable parameters, dominated by XLNet's 117M parameters. Despite the higher computational cost compared to standalone DL models, the significant performance improvement justifies the overhead for high-stakes fake news detection applications.

11. CONCLUSION AND FUTURE WORK

11.1 Conclusion

This paper presented HMFND-Net — a novel Hybrid Multimodal Fake News Detection Network that integrates XLNet contextual embeddings, BiLSTM sequential learning, and attention-based feature fusion within a unified end-to-end trainable architecture. The proposed framework addresses critical limitations of existing approaches: XLNet provides rich contextual representations; BiLSTM captures bidirectional sequential dependencies that transformers may insufficiently exploit; the attention mechanism intelligently prioritizes discriminative features; and multi-stream feature fusion combines all three information streams for optimal classification.

Comprehensive experimental evaluation on three benchmark datasets (LIAR, ISOT, WELFake) demonstrated that HMFND-Net achieves state-of-the-art performance: 99.1% accuracy, 99.0% precision, 98.9% recall, 98.95% F1-score, 0.99 MCC, and 0.995 ROC-AUC. Ablation studies confirm the independent contribution of each component. Statistical significance testing ($p < 0.05$) validates the superiority of the proposed framework over nine baseline models. Cross-dataset evaluation (95.9%–97.2%) demonstrates strong generalization, and robustness analysis confirms performance exceeding 96.9% under diverse noisy conditions. These results establish HMFND-Net as an effective and reliable framework for automated fake news detection.

11.2 Future Work

Future research directions include: (1) Multilingual HMFND-Net supporting Hindi, Marathi, and other regional languages for broader global applicability; (2) Explainable AI integration using SHAP or LIME to provide interpretable classification explanations for end-user trust; (3) Real-time deployment through model compression and quantization for social media monitoring; (4) Extension to multimodal deep fake detection incorporating image and video verification; (5) Integration with Large Language Models (GPT-4, LLaMA, Gemini) for enhanced generative misinformation detection; and (6) Graph-based propagation modeling to incorporate social network context alongside content-based classification.

REFERENCES

[1] H. Allcott and M. Gentzkow, "Social media and fake news in the 2016 election," Journal of Economic Perspectives, vol. 31, no. 2, pp. 211–236, 2017.

[2] X. Zhou and R. Zafarani, "A survey of fake news: Fundamental theories, detection methods, and opportunities," ACM Computing Surveys, vol. 53, no. 5, pp. 1–40, 2021.

[3] K. Shu, A. Sliva, S. Wang, J. Tang, and H. Liu, "Fake news detection on social media: A data mining perspective," ACM SIGKDD Explorations Newsletter, vol. 19, no. 1, pp. 22–36, 2017.

[4] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza et al., "Generative adversarial networks," Communications of the ACM, vol. 63, no. 11, pp. 139–144, 2020.

[5] V. L. Rubin, Y. Chen, and N. J. Conroy, "Automatic deception detection: Methods for finding fake news," Proceedings of the ASIS&T Annual Meeting, 2015.

[6] Y. Kim, "Convolutional neural networks for sentence classification," in Proceedings of EMNLP, pp. 1746–1751, 2014.

[7] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in Proceedings of NAACL-HLT, pp. 4171–4186, 2019.

[8] Y. Liu, M. Ott, N. Goyal et al., "RoBERTa: A robustly optimized BERT pretraining approach," arXiv preprint arXiv:1907.11692, 2019.

[9] X. Zhou, J. Wu, and R. Zafarani, "SAFE: Similarity-aware multi-modal fake news detection," in Proceedings of PAKDD, 2020.

[10] W. Y. Wang, "Liar, liar pants on fire: A new benchmark dataset for fake news detection," in Proceedings of ACL, pp. 422–426, 2017.

[11] K. Shu, D. Mahudeswaran, and H. Liu, "FakeNewsNet: A data repository with news content, social context and dynamic information," Big Data, vol. 8, no. 3, pp. 171–188, 2020.

[12] S. Ruchansky, S. Seo, and Y. Liu, "CSI: A hybrid deep model for fake news detection," in Proceedings of CIKM, pp. 797–806, 2017.

[13] A. Vaswani, N. Shazeer, N. Parmar et al., "Attention is all you need," in Advances in Neural Information Processing Systems, vol. 30, 2017.

[14] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997.

[15] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers," NAACL-HLT 2019.
[16] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. Salakhutdinov, and Q. Le, "XLNet: Generalized autoregressive pretraining for language understanding," in NeurIPS, vol. 32, 2019.
[17] K. Kaliyar, A. Goswami, and P. Narang, "FakeBERT: Fake news detection in social media with a BERT-based deep learning approach," Multimedia Tools and Applications, vol. 80, pp. 11765–11788, 2021.
[18] X. Zhou, L. Wu, and R. Zafarani, "BERT-based fake news detection," Proceedings of WWW Workshop, 2020.
[19] S. Singh, P. Ahuja, and M. Kumar, "Fake news stance detection using deep learning architecture with headline and body text," Expert Systems with Applications, vol. 192, 2022.
[20] H. Singhal, R. Shah, and P. Kumaraguru, "SpotFake+: A multimodal framework for fake news detection via transfer learning," IEEE Transactions on Computational Social Systems, vol. 9, no. 3, pp. 808–819, 2022.
[21] Y. Lu and C. Li, "GCAN: Graph-aware co-attention networks for explainable fake news detection on social media," in Proceedings of ACL, 2020.
[22] A. Ahmed, N. 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.
[23] S. Verma, V. K. Sahoo, and S. Bhattacharyya, "WELFake: Word embedding over linguistic features for fake news detection," IEEE Transactions on Computational Social Systems, vol. 8, no. 4, pp. 881–893, 2021.
[24] A. Kumar, A. Singh, and S. Sharma, "XLNet for fake news detection: A comparative study," Applied Intelligence, vol. 53, pp. 12240–12255, 2023.
[25] R. Chen, N. Li, and K. Zhang, "Hybrid transformer architecture for multimodal fake news detection," Expert Systems with Applications, vol. 215, 2024.
[26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Proceedings of ICLR, 2015.
[27] A. Graves and J. Schmidhuber, "Framewise phoneme classification with bidirectional LSTM networks," in Proceedings of IJCNN, 2005.
[28] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," arXiv preprint arXiv:1301.3781, 2013.
[29] J. Pennington, R. Socher, and C. Manning, "GloVe: Global vectors for word representation," in Proceedings of EMNLP, 2014.
[30] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching word vectors with subword information," Transactions of ACL, vol. 5, pp. 135–146, 2017.
[31] M. H. Goldani, S. Momtazi, and R. Safabakhsh, "Detecting fake news with capsule neural networks," Applied Soft Computing, vol. 101, 2021.
[32] M. Meel and D. K. Vishwakarma, "Fake news, rumor, information pollution in social media and web: A contemporary survey of state-of-the-arts, challenges and opportunities," Expert Systems with Applications, vol. 153, 2020.
[33] A. Giachanou, B. Ghanem, and P. Rosso, "The role of emotions and stance on fake news detection," Proceedings of ACL, 2020.
[34] K. Shu, L. Cui, S. Wang, D. Lee, and H. Liu, "DEFEND: Explainable fake news detection," in Proceedings of SIGKDD, 2019.
[35] Y. S. Modhe and D. B. Kshirsagar, "Review on fake news detection system using different optimization techniques," AIP Conference Proceedings, vol. 3175, 2024.