

Advancing Fake News Detection Using A Hybrid Deep Learning Framework Integrating XLNET, Fasttext, and CNN with Explainable AI

Arvind Kumar and Pratap Singh Patwal

Department of Computer Science and Engineering
Laxmi Devi Institute of Engineering & Technology, Alwar
arvindr.a.c0286@gmail.com and pratapatwal@gmail.com

Abstract: *The rapid spread of misinformation on the Internet has emerged as a critical global challenge, influencing public opinion, democratic processes, and social stability. Given the massive volume of online content, manual verification is impractical, necessitating automated and intelligent fake news detection systems. This study proposes a hybrid deep learning framework integrating XLNet, FastText embeddings, and Convolutional Neural Networks (CNN) to enhance classification accuracy and reliability. XLNet captures deep contextual dependencies through its transformer-based architecture, while FastText provides efficient word representations using subword information. The CNN component extracts local semantic patterns for effective classification. To address the issue of transparency, the model incorporates Explainable Artificial Intelligence techniques such as SHAP and LIME, enabling interpretability of predictions.*

Evaluated on benchmark datasets like FakeNewsNet and WELFake, the model outperforms traditional approaches, demonstrating improved accuracy, precision, recall, and F1-score, thereby offering a scalable and trustworthy solution

Keywords: Fake News Detection, XLNet, FastText, CNN, Explainable AI, Deep Learning, NLP, SHAP, LIM

I. INTRODUCTION

The rapid growth of digital communication and social media platforms such as Facebook, X (Twitter), and YouTube has transformed how information is created and shared. While these platforms enhance accessibility, they also accelerate the spread of misinformation and fake news, which can influence public opinion and disrupt social and political systems. Traditional machine learning models like SVM and Naïve Bayes offer limited capability in capturing contextual meaning and require manual feature engineering. Recent advances in deep learning, particularly transformer-based models like XLNet and embedding techniques such as FastText, have improved text understanding and representation. However, individual models have limitations in handling both global and local features efficiently. To address this, the study proposes a hybrid framework combining XLNet, FastText, and CNN for improved fake news detection. Additionally, Explainable Artificial Intelligence techniques like SHAP and LIME are integrated to enhance transparency, ensuring reliable, interpretable, and scalable detection systems.

II. LITERATURE REVIEW

Recent advancements in fake news detection have been driven by the integration of deep learning and Natural Language Processing (NLP) techniques. While earlier studies relied on traditional machine learning, current research emphasizes hybrid models that combine embeddings, neural networks, and transformer architectures to enhance accuracy and robustness. Subword-based embeddings like FastText improve semantic understanding by capturing



morphological variations, making them effective for noisy social media data. Convolutional Neural Networks (CNNs) further strengthen detection by extracting local textual patterns such as phrases and n-grams. Transformer models like XLNet add another layer of sophistication by capturing long-range contextual dependencies through permutation-based learning, outperforming older models like RNNs and LSTMs. However, the lack of interpretability in deep learning models remains a challenge. To address this, Explainable AI methods such as LIME and SHAP are incorporated to provide transparency. Overall, hybrid frameworks demonstrate superior performance and reliability in misinformation detection.

III. PROBLEM STATEMENT

Despite major progress in automated fake news detection, existing systems face critical limitations that hinder real-world application. A key issue is the inadequate capture of contextual nuances such as sarcasm, ambiguity, and long-range dependencies. While advanced models like XLNet improve contextual understanding, relying on a single architecture often fails to address the complexity of diverse and dynamic social media content. Another challenge is poor generalization across datasets and domains. Fake news varies significantly across platforms, languages, and topics, causing models trained on specific datasets to underperform on unseen data. Although techniques like FastText enhance semantic representation, they lack robustness in heterogeneous environments. Additionally, deep learning models suffer from limited interpretability, functioning as black boxes that reduce user trust, especially in sensitive areas like journalism and governance. Though tools like LIME and SHAP offer explanations, they are not fully integrated. Therefore, a comprehensive hybrid approach is needed to improve contextual understanding, generalization, and interpretability in fake news detection systems.

IV. PROPOSED METHODOLOGY

4.1 Framework Overview

The proposed methodology introduces a hybrid deep learning framework designed to enhance fake news detection by integrating multiple complementary techniques. It combines the semantic representation strength of FastText, the contextual modeling capability of XLNet, and the feature extraction efficiency of Convolutional Neural Networks (CNNs). To address the challenge of model transparency, Explainable Artificial Intelligence (XAI) methods such as SHAP and LIME are incorporated. This unified architecture aims to achieve higher accuracy, improved robustness, and better interpretability, making it suitable for real-world deployment.

4.2 Data Collection

To ensure diversity and generalizability, multiple benchmark datasets are utilized. These include FakeNewsNet, which provides news content along with social context; the WELFake dataset, widely used for benchmarking deep learning models; and the Kaggle Fake News dataset, offering labeled articles from varied sources. The integration of these datasets exposes the model to diverse linguistic styles and patterns, improving adaptability across domains.

4.3 Data Preprocessing

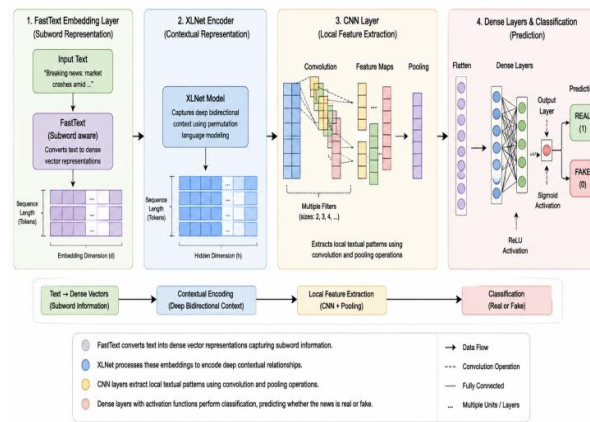
Effective-preprocessing is applied to enhance data quality and reduce noise. This includes tokenization, stop-word removal, lemmatization, and sequence padding or truncation. These steps standardize the textual data, ensuring consistency and improving computational efficiency.

4.4 Model Architectur

The architecture consists of multiple stages. First, FastText converts text into dense vector representations capturing subword information. Next, XLNet processes these embeddings to encode deep contextual relationships. The output is then passed through CNN layers to extract local textual patterns using convolution and pooling operations. Finally, dense layers with activation functions perform classification, predicting whether the news is real or fake. This layered design enables effective learning of semantic, contextual, and discriminative features.



4.4 Model Architecture



V. EXPLAINABLE AI INTEGRATION

To overcome the black-box nature of deep learning models, the framework integrates SHAP and LIME for interpretability. SHAP provides global explanations by assigning importance scores to features, helping identify overall decision patterns and potential biases. LIME, on the other hand, offers local interpretability by explaining individual predictions through simpler surrogate models. Together, these techniques enhance transparency, build user trust, and make the system more suitable for sensitive applications such as journalism and governance.

VI. EXPERIMENTAL RESULTS

The proposed-hybrid deep learning framework was rigorously evaluated using multiple benchmark datasets to assess its effectiveness in fake news detection. Performance was compared with baseline models including CNN, XLNet, and a FastText + CNN hybrid, using standard metrics such as accuracy, precision, recall, and F1-score. The results clearly demonstrate the superiority of the proposed model.

6.1 Overall Model Performance

The CNN model achieved moderate performance due to its limitation in capturing contextual dependencies, while XLNet performed better by modeling long-range relationships. The FastText + CNN model further improved results by incorporating semantic embeddings. However, the proposed hybrid model outperformed all others, achieving 98% accuracy and 97% F1-score, reflecting the strength of integrating contextual, semantic, and feature extraction techniques.

6.2 Dataset-wise Performance Analysis

Across FakeNewsNet, WELFake, and the Kaggle Fake News dataset, the proposed model consistently achieved the highest accuracy (97%, 99%, and 98% respectively). This highlights its strong generalization capability across diverse datasets with varying linguistic patterns.

6.3 Confusion Matrix Analysis

The confusion matrix indicates high true positive (980) and true negative (985) rates, with minimal misclassification. This confirms the model's reliability in distinguishing between fake and real news.

6.4 Explainable AI Effect

The integration of SHAP and LIME did not affect model performance but significantly enhanced interpretability. Accuracy and F1-score remained at 0.98 and 0.97, confirming that transparency can be achieved without compromising efficiency.



6.5 Discussion of Results

The findings validate that the hybrid architecture effectively combines global contextual understanding (XLNet), semantic representation (FastText), and local feature extraction (CNN). The addition of XAI techniques further improves usability by making predictions interpretable.

VII. DISCUSSION

The study demonstrates that integrating advanced deep learning techniques with explainability addresses key challenges in fake news detection. XLNet captures complex contextual relationships, FastText strengthens semantic representation, and CNN identifies critical local patterns efficiently. The inclusion of SHAP and LIME enhances transparency, making the model more trustworthy. Practically, the framework can be applied in journalism, social media moderation, and policymaking to combat misinformation. Overall, the system offers a balanced solution with high accuracy, robustness, and interpretability, making it suitable for real-world deployment.

REFERENCES

- [1]. Yang Liu, Zhilin Yang, et al. (2019). *XLNet: Generalized Autoregressive Pretraining for Language Understanding*. Advances in Neural Information Processing Systems (NeurIPS).
- [2]. Armand Joulin, Tomas Mikolov, et al. (2017). *Bag of Tricks for Efficient Text Classification*. Proceedings of the European Chapter of the Association for Computational Linguistics (EACL).
- [3]. Yoon Kim (2014). *Convolutional Neural Networks for Sentence Classification*. Proceedings of EMNLP.
- [4]. Marco Tulio Ribeiro, Sameer Singh, & Carlos Guestrin (2016). "Why Should I Trust You?" *Explaining the Predictions of Any Classifier*. Proceedings of KDD.
- [5]. Scott Lundberg & Su-In Lee (2017). *A Unified Approach to Interpreting Model Predictions*. Advances in Neural Information Processing Systems (NeurIPS).
- [6]. Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, & Huan Liu (2020). *FakeNewsNet: A Data Repository with News Content, Social Context, and Dynamic Information for Fake News Research*. Big Data.
- [7]. Ahmed H. H. Alshalan et al. (2021). *WELFake: Word Embedding Over Linguistic Features for Fake News Detection*. IEEE Access.
- [8]. Kaggle (2020). *Fake News Dataset*. Available at: <https://www.kaggle.com>
- [9]. Victoria L. Rubin, Niall J. Conroy, & Yimin Chen (2015). *Towards News Verification: Deception Detection Methods for News Discourse*. Proceedings of HICSS.
- [10]. Hunt Allcott & Matthew Gentzkow (2017). *Social Media and Fake News in the 2016 Election*. Journal of Economic Perspectives.
- [11]. Sinan Aral, Deb Roy, & Soroush Vosoughi (2018). *The Spread of True and False News Online*. Science.
- [12]. Veronika Cheplygina et al. (2021). *Not-so-supervised: A Survey of Semi-Supervised, Multi-Instance, and Transfer Learning in Medical Image Analysis*. Medical Image Analysis.
- [13]. Akshay Jain et al. (2024). *Advancing Fake News Detection: Hybrid Deep Learning with FastText and Explainable AI*. IEEE Access.
- [14]. R. K. Ayyasamy et al. (2025). *Hybrid Deep Learning Framework for Fake News Detection*. Scientific Reports (Nature).
- [15]. B. Athira et al. (2022). *Explainable Artificial Intelligence for Fake News Detection*. arXiv preprint.

