JCTA Journal of Computing Theories and Applications ISSN:3024-9104

Research Article

Fake News Detection Using Bi-LSTM Architecture: A Deep Learning Approach on the ISOT Dataset

Maaruf M. Lawal [1], and Abdulrashid Abdulrauf [2,*]

[1] Department of Computer Science, Ahmadu Bello University Zaria, PMB 1045, Kaduna State, Nigeria; e-mail : mmlawal80@gmail.com
[2] Department of Computer Science, Federal Polytechnic Kaltungo, PMB 09, Gombe State, Nigeria; e-mail : abdulrashid.abdulrauf2022@nda.edu.ng
* Corresponding Author : Abdulrashid Abdulrauf

Abstract: The proliferation of fake news across digital platforms has raised critical concerns about information reliability. A notable example is the viral rumor falsely claiming that the Nigerian Minister of the Federal Capital Territory, Nyesom Wike, had collapsed at an event and was rushed to an undisclosed hospital, an entirely fabricated claim that caused public confusion. While both traditional machine learning and deep learning approaches have been explored for automated fake news detection, many existing models have been limited to topic-specific datasets and often suffer from overfitting, especially on smaller datasets like ISOT. This study addresses these challenges by proposing a standalone Bidirectional Long Short-Term Memory (BiLSTM) model for fake news classification using the ISOT dataset. Our model outperformed the state-of-the-art models, which achieved 96.3% accuracy. In contrast, the proposed BiLSTM model achieved superior results with 98.98% accuracy, 98.22% precision, 99.65% recall, and a 98.93% F1-score. The model demonstrated balanced classification across both fake and real news, exhibiting strong generalization capabilities. However, training and validation performance plots revealed signs of overfitting after epoch 2, suggesting the need for regularization in future work. This study contributes to the growing body of research on fake news detection by demonstrating the efficacy of a focused, sequential deep learning model over more complex architectures, providing a practical, scalable, and robust solution for misinformation detection.

Keywords: BiLSTM; Deep Learning; Detection; Fake News; HOAX Detection; ISOT; Misinformation Detection.

Received: August, 14th 2025
Revised: August, 30th 2025
Accepted: August, 31st 2025
Published: September, 3rd 2025

1. Introduction

The proliferation of information across the Internet and the World Wide Web has raised significant concerns regarding its reliability. In recent years, the alarming spread of rumors and fake news has profoundly impacted social and political discourse [1]. This issue is exacerbated by the substantial amount of time individuals dedicate to social media platforms and online news sources, a trend that has diminished the prominence of traditional information dissemination channels as the Internet becomes increasingly integrated into modern life [1].

Nigeria exemplifies this escalating threat, having witnessed a significant intensification of fake news dissemination over the past decade. Notably, during the 2023 general elections, numerous websites emerged, reaching thousands of users while propagating false information [2]. A striking example includes a widely circulated false claim that Nyesom Wike, the Minister of the Federal Capital Territory (FCT), had collapsed at an event and was hospitalized [3].

Research consistently demonstrates that misinformation spreads at a significantly faster rate than factual information. Studies have shown that tweets containing false information are retweeted six times more quickly than verified content, contributing to widespread fear, panic, and even financial losses [4], [5]. The pervasive nature of this issue is further highlighted by findings that over one-third of trending topics on Chinese microblogging platforms contained fake news [6]. On Twitter, political misinformation is notably more likely to be retweeted and to spread more rapidly than truthful content [7]. Even seemingly reliable sources, such as Wikipedia, have been found to be susceptible to misinformation [8], underscoring the rapid and damaging effects of fake news dissemination across various platforms.

While both traditional machine learning and deep learning approaches have been explored for automated fake news detection [9]–[11], much of this research has been confined to specific types of news, particularly political news, leading to models that are often tailored to topic-specific datasets. Existing comparative studies have also been limited in their scope. For instance, the ISOT dataset, introduced by Ahmed et al. [7] for benchmarking various models, presents challenges for training deep neural networks due to its limited length, which occasionally leads to overfitting. This existing gap in generalizable and robust fake news detection models, particularly with challenging datasets, represents a critical problem. Addressing these limitations, this study proposes a standalone model utilizing BiLSTM for classifying news as fake or real, specifically employing the ISOT dataset to mitigate the challenges previously encountered.

The ISOT dataset was chosen for this study because it is a well-established benchmark dataset that has been widely used in fake news detection research, thereby facilitating reproducibility and meaningful comparison with existing studies. Although this dataset does not directly capture the Nigerian context of fake news, it provides a balanced and structured collection of real and fake news articles suitable for model evaluation. In the Nigerian case, where fake news circulated widely during the 2023 elections, the ISOT dataset provides a suitable foundation for developing and benchmarking models before their extension to more localized datasets. This enables the validation of model robustness in a controlled setting before applying it to real-world, context-specific challenges.

Our choice of BiLSTM over traditional ML models such as Naïve Bayes, Random Forests, or XGBoost is based on the distinct advantages of deep learning for handling textual data [12]. While traditional ML algorithms (e.g., SVM [13], Random Forest [13], [14], Logistic Regression [14], [15], Naïve Bayes [14]–[17], and XGBoost [18]) remain effective when combined with feature engineering, they rely heavily on handcrafted features and lack the ability to capture sequential dependencies within text fully[12]. In contrast, BiLSTM architectures excel at modelling long-range dependencies and sequential patterns [19]–[21], making them particularly suited for fake news detection where contextual meaning and word order significantly influence classification outcomes. Unlike CNNs, which mainly capture local patterns [20]–[22], or basic RNNs prone to vanishing gradient problems, BiLSTMs process information in both forward and backward directions, thereby preserving semantic context and improving text representation. The key contributions of this study are:

- Propose a standalone BiLSTM model emphasizing sequential learning of fake news texts.
- Evaluate the model on the ISOT dataset to ensure robustness and generaliz bility.
- Benchmark results against existing state-of-the-art models.

2. Related Work

Recent years have seen extensive research into automated fake news detection, employing both classical machine learning and deep learning approaches. Potey et al. [23] developed a web-based system for real-time detection of fake news and sentiment on Twitter. Empirical results from 14,000 tweets showed effective classification using TF-IDF and sentiment polarity, with visual representation aiding result interpretation. Their system proved effective in detecting COVID-19 misinformation. Similarly, Kula et al. [24] utilized Google Colab and the Flair NLP library to build a deep learning model, achieving 99.8% accuracy on training data. Their cloud-based empirical implementation emphasized the scalability of DNNs for fake news detection and highlighted Flair's advantages in text feature extraction. Ensemble-based strategies have also shown promise; Akinyemi et al. [25] proposed a stacked ensemble model that combines content and social features, tested on the PHEME dataset. Their empirical evaluation demonstrated a 17.25% increase in accuracy and a 15.78% improvement in sensitivity compared to existing approaches, validating the robustness of the ensemble strategy. While Bauskar et al. [26] developed a hybrid model that combines social and content-based features using NLP methods. Their experiments achieved an average accuracy of 90.62% and an F1 score of 90.33%, underscoring the benefits of incorporating how fake news spreads socially.

Benchmark studies highlight dataset-specific model performance. Khan et al. [27] conducted a benchmark study across three datasets and 19 models, including BERT. Their empirical analysis revealed BERT's superiority on small datasets and LSTM's reliance on long, rich content. Naïve Bayes with N-grams performed comparably to neural models on large datasets. Alghamdi et al. [28] confirmed similar trends by benchmarking various models across four datasets (LIAR, PolitiFact, GossipCop, COVID-19), finding that BERT and RoBERTa ex-celled, particularly on short-text datasets. However, classical ML models outperformed deep learning in certain scenarios, such as GossipCop, demonstrating dataset-dependent performance differences. Ensemble approaches were further validated by Ahmad et al. [29], who demonstrated that ensemble models achieved an average accuracy of 97.67% across four datasets (DS1–DS4) and outperformed single classifiers, with Random Forest reaching 99% accuracy on ISOT. Advances in architectures also emerged: Goldani et al. [30] introduced capsule neural networks for ISOT and LIAR, improving accuracy by up to 7.8%; Girgis et al. [31] tested RNN, GRU, and LSTM on LIAR, with GRU slightly outperforming others; and Sarnovský et al. [32] demonstrated the applicability of CNN, LSTM, and BiLSTM to Slovak COVID-19 misinformation, reaching 98.93% accuracy and a 94% F1 score.

Focusing specifically on the ISOT dataset, several studies highlight its role as a benchmark for evaluating the detection of fake news. Alsuwat and Alsuwat [33] proposed the MM-FND framework, which integrates textual, temporal, and spatial features, achieving 96.3% accuracy and 96.4% F1-score on ISOT, as well as strong results on the LIAR and COVID-19 datasets. Fayaz et al. [34] employed Random Forest with feature selection on ISOT, achieving 97.25% accuracy with all features and up to 97.33% with optimized subsets, outperforming gradient boosting and other baselines. Yu et al. [35] introduced a multi-tier filtering model with consensus layers, obtaining 98.95% accuracy and an F1 of 0.9892 on ISOT, surpassing conventional ML pipelines. Rani and Shokeen [36] combined blockchain and deep learning (FNNet) for secure detection, achieving 98.53% accuracy on ISOT along with high precision and recall. Kansal [37] emphasized linguistic style features in an ensemble model, achieving 92.32% accuracy and an F1-score of 0.9185 on ISOT, underscoring the importance of stylistic cues in deceptive writing.

While the MM-FND framework [33] shows strong performance, particularly on the COVID-19 dataset, the model's accuracy on the ISOT Fake News Dataset (96.3%) indicates room for improvement, especially in refining temporal feature extraction. The integration of Bi-LSTM was limited to capturing temporal patterns within a multi-modal context. However, the model could benefit from a more focused architecture that leverages the sequential dependencies of textual data more effectively. This presents a gap where a dedicated and optimized Bi-LSTM-based model could potentially enhance accuracy and overall performance on the ISOT dataset.

To address this gap, a standalone Bi-LSTM model is proposed, designed to learn temporal patterns in fake news texts on the ISOT dataset in a deep manner. This focused architecture may offer improved accuracy by eliminating feature-level fusion complexities and emphasizing sequential modeling, potentially surpassing the 96.3% accuracy threshold set by the MM-FND framework [32]. The BiLSTM was selected over traditional machine learning approaches such as Naïve Bayes, Random Forest, and XGBoost because these rely heavily on handcrafted or statistical features, which fail to capture the sequential and contextual dependencies present in textual data.

Compared to transformer-based models like BERT, which require large-scale pretraining and high computational costs, BiLSTM strikes a balance between efficiency and performance. Its bidirectional nature allows it to process text from both past and future contexts, effectively capturing nuanced semantic cues and contextual relationships that are vital in distinguishing real from fake news [12]. This makes BiLSTM a simpler yet powerful framework, well-suited for the ISOT dataset, and capable of utilizing sequential characteristics of language without the overhead of hybrid or ensemble architectures.

3. Proposed Method

This section presents the dataset, data preprocessing, feature selection, and models used. Figure 1 gives a clear diagrammatic description of the proposed model, and the steps involved are discussed in subsequent sections.

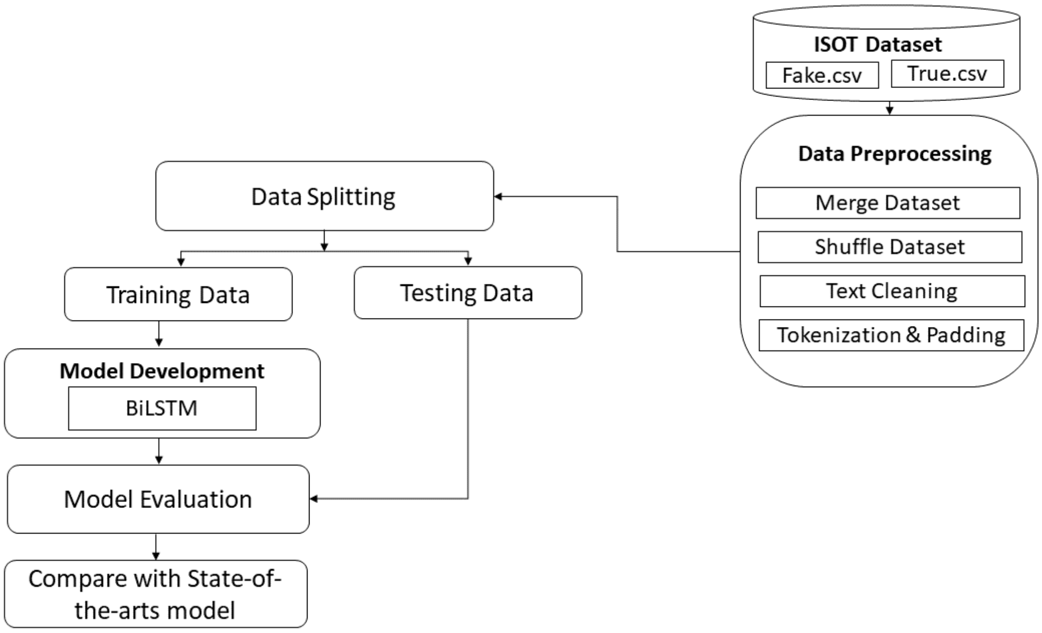

Figure 1. Proposed framework

3.1. Dataset

The dataset considered in this study is the ISOT Fake News Dataset, compiled by Ahmed et al. [7] from real-world platforms, specifically Reuters.com and Kaggle.com. This dataset is widely used as a benchmark in fake news detection research due to its structured format and reliable data sources. Each entry in the dataset exceeds 200 characters in length and includes various metadata such as the article title, text, type, and publication date. The dataset is publicly available on the Kaggle platform. The ISOT dataset consists of two CSV files. The detailed statistics of the dataset are presented in Table 1.

Table 1. Type and size of articles per category for the ISOT dataset provided by Ahmed et al. [7]

News Type	Total Record	Subject	
		Type	Size
Real News	21,417	World-News	10,145
		Politics-News	11,272
Fake News	23,481	Government-News	1,570
		Middle-east	778
		US News	783
		Left-News	4,459
		Politics	6,841
		News	9050

Each record provides essential textual and metadata fields useful for model training and evaluation. Due to its balanced and diverse structure, the ISOT dataset has been extensively adopted in numerous studies as a benchmark for evaluating the performance of fake news detection models. In this study, the original ISOT dataset consisted of 23,481 fake news articles and 21,417 real news articles, totaling 44,898 records. These counts served as the baseline for evaluating whether any reduction occurred during the preprocessing phase.

3.2. Data Preprocessing

The study's data preprocessing involved preparing two datasets of real and fake news for deep learning models. The datasets were first labeled (0 for fake, 1 for real) and merged into a single DataFrame; then, they were shuffled to remove ordering bias. The preprocessing stage combined article titles and body text into a single field, normalized text to lowercase, and removed brackets, URLs, HTML tags, punctuation, newline characters, and numeric tokens. Following preprocessing, the dataset size was reduced from 44,898 to 44,678 articles, representing a data loss of 220 records. The data loss occurred due to the removal of null or empty records during the cleaning process. Figure 2 illustrates a comparison of dataset sizes before and after preprocessing.

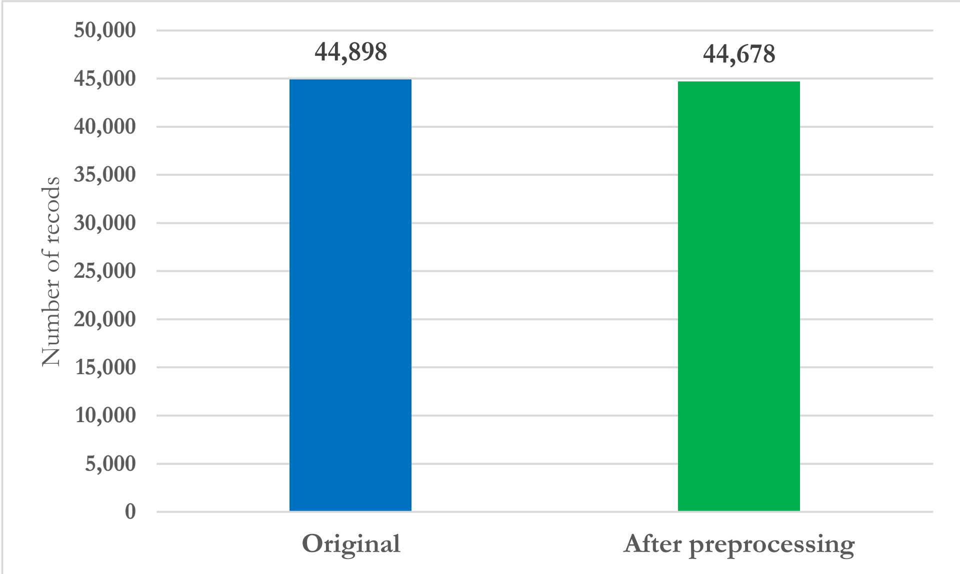

Figure 2. Dataset size before and after preprocessing

3.3. Data Splitting

After preprocessing, the merged dataset of real and fake news was split into training and testing sets using Scikit-learn's train_test_split() function. Tokenized and padded text sequences served as input features (X), while binary labels (0 for fake, 1 for real) were the targets (y). The cleaned dataset was divided into training, validation, and testing subsets following an 80:10:10 ratio. The split sizes are illustrated in Figure 3, which confirms balanced partitioning and sufficient representation across all sets. To ensure reproducibility, the random_state parameter was set to 42, enabling consistent training and test sets across multiple runs.

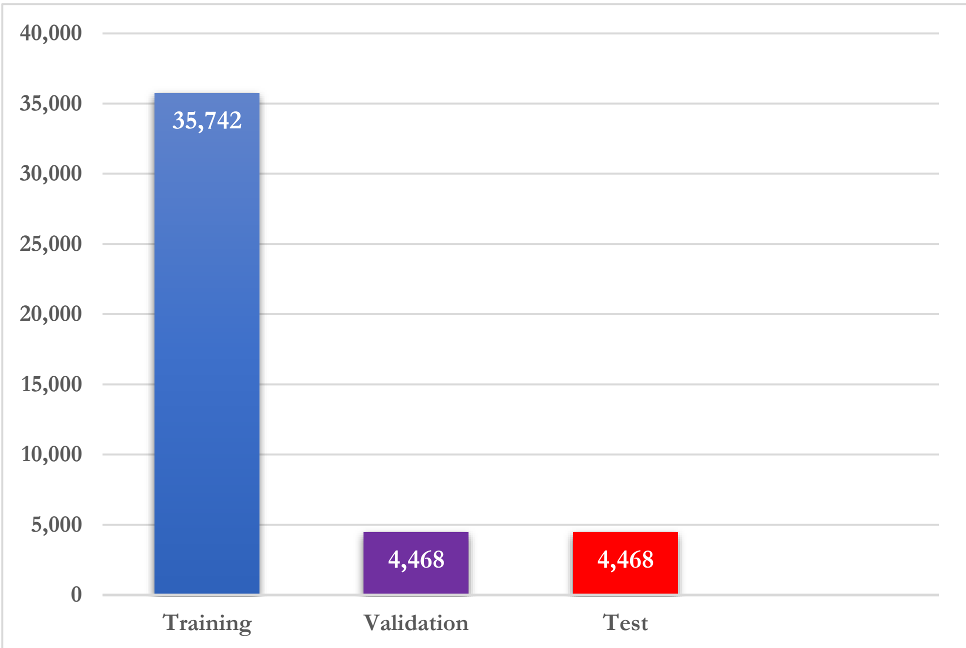

Figure 3. Bar chart of train, validation, and test split sizes.

It is important to note that cross-validation was not employed in this study. This choice is consistent with several state-of-the-art works in fake news detection and text classification that also adopt a fixed train-validation-test split. For instance [33], [34], [37], all report results using predetermined splits rather than cross-validation. Considering the large dataset size and the high computational cost associated with retraining deep learning models multiple times, the fixed 80:10:10 split provides an efficient yet robust evaluation strategy. Moreover, this approach ensures comparability with prior studies reporting results using similar splitting strategies.

3.4. Model Development

To classify news articles as either real or fake, four different deep learning models were developed using the combined dataset from True.csv (real news) and Fake.csv (fake news). After data preprocessing and tokenization, the text sequences were input into these models to learn discriminative features for classification.

The model employs a Bidirectional Long Short-Term Memory (BiLSTM) network, which processes text sequences in both forward and backward directions, allowing it to capture contextual dependencies from both sides of a sentence. This capability is especially valuable in natural language processing tasks, such as fake news detection, where understanding the full context of a sentence can determine its meaning and credibility. Table 2 presents the BiLSTM architecture.

Table 2. Bi-LSTM architecture.

Layer (Type)	Output Shape	Param #
Embedding (embedding_1)	(None, 300, 128)	1,280,000
Bidirectional LSTM (bidirectional_1)	(None, 128)	98,816
Dropout (dropout_2)	(None, 128)	0
Dense (dense_2)	(None, 64)	8,256
Dropout (dropout_3)	(None, 64)	0
Dense (dense_3)	(None, 1)	65

As shown in Table 2, the model begins with an Embedding layer that converts each word into a 128-dimensional vector, thereby capturing its semantic meaning. This is fed into a BiLSTM layer, which processes the sequence bidirectionally to capture context from both past and future words. A 50% Dropout layer is applied to reduce overfitting, followed by a dense layer with 64 ReLU-activated units that extracts higher-level features. Another 30% Dropout is applied before the final dense layer with a sigmoid activation, which outputs the probability that the article is fake (0) or real (1).

The architecture was deliberately designed with relatively low complexity, utilizing only one Embedding layer, a single BiLSTM layer, and two sets of Dropout and dense layers. This choice strikes a balance between model expressiveness and computational efficiency, ensuring faster training, reduced memory requirements, and a lower risk of overfitting. While deeper or more complex models might capture additional nuances, they often demand large datasets and heavy computation, which may not generalize well. By keeping the model compact yet strategically structured, it effectively captures semantic relationships (via Embedding), sequential context (via BiLSTM), and robust feature extraction (via Dense + Dropout).

Table 3. Bi-LSTM training configuration.

Layer (Type)	Output Shape
Loss Function	Binary Cross-Entropy
Optimizer	Adam
Batch Size	128
Epochs	5

As shown in Table 3, the model was compiled with Binary Cross-Entropy loss, which is well-suited for binary classification problems as it compares predicted probabilities against the true class labels. The Adam optimizer was selected for its adaptive learning rate and computational efficiency, which helps the model converge more quickly. A batch size of 128 was employed to strike a balance between training speed and stable gradient updates. Although the model was initially built and tested with a single epoch, the final training was conducted over five epochs, allowing the network to refine its parameters progressively across multiple passes through the dataset. This configuration ensured that the model was computationally efficient, avoided overfitting, and maintained strong generalization on unseen data.

4. Results and Discussion

4.1. Training, Validation Accuracy, and Loss of BiLSTM Model

Figure 4 illustrates the Training and Validation Accuracy (left) and Training and Validation Loss (right) of the BiLSTM model over five epochs. These plots provide a visual assessment of the model's learning behavior and generalization ability.

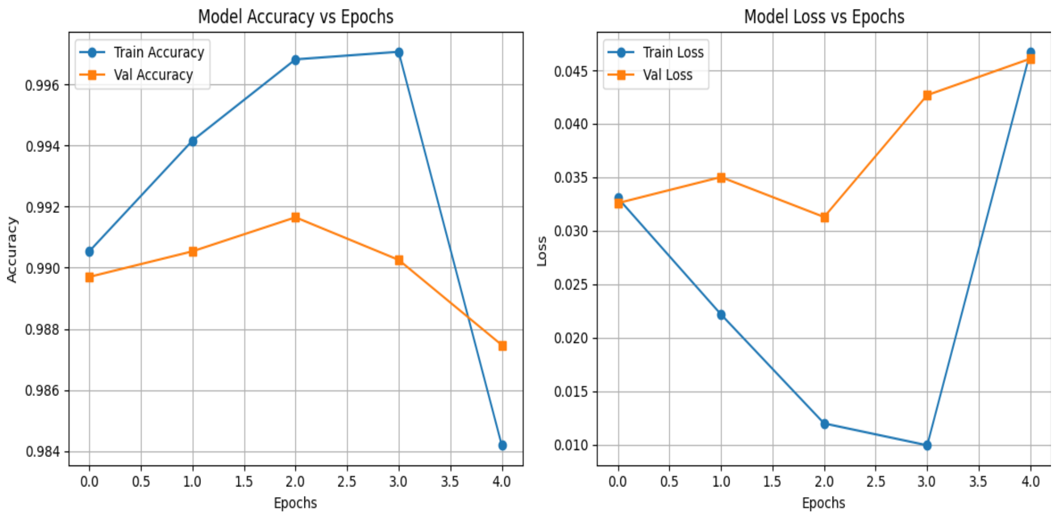

Figure 4. Training, validation accuracy, and loss of BiLSTM

The training and validation performance plots of BiLSTM, as shown in Figure 4, indicate that the model generally learns well over the five training epochs. Training accuracy steadily increases from approximately 98.6% to nearly 99.8%, while training loss consistently decreases, reflecting effective learning on the training data. Validation accuracy follows a similar upward trend until epoch two but drops slightly in epoch three before partially recovering in epoch four. Notably, the validation loss shows a spike at epoch 3, despite continued improvement in training performance, which suggests that the model may have begun to overfit the training data at that point. Overall, the model achieves high accuracy on both the training and validation sets. However, the divergence in loss and accuracy trends around epoch three highlights the need to monitor for overfitting in further training.

Similarly, the Training and Validation Loss plot reinforces this trend. The training loss shows a consistent decrease, reaching its lowest point around epoch three, reflecting the model's increasing fit to the training data. In contrast, the validation loss, after an initial fluctuation, begins to rise steadily after epoch two, peaking at epoch four. This increasing gap between training and validation loss, despite continued improvement on the training set, indicates slight overfitting.

4.2. Testing Performance

As shown in Tables 4 and 5, the BiLSTM model, employed in this study, has demonstrated strong performance on the test set in distinguishing between fake and real news articles. The overall accuracy of the model was 98.74%, indicating that nearly all predictions were correct. While this high accuracy value is promising, it is essential to interpret it in conjunction with other key metrics, namely precision, recall, and F1 score, for a more comprehensive assessment.

Table 4. Performance metrics of BiLSTM on the test set

Metric	Value
Accuracy	0.9898
Precision	0.9822
Recall	0.9965
F1 Score	0.9893

Table 5. Detailed Classification Report of BiLSTM on test set

Class	Precision	Recall	F1-Score	Support
Fake	1.00	0.98	0.99	4710
Real	0.98	1.00	0.99	4270
		Overall Scores		
Accuracy			0.99	8980
Macro Avg	0.99	0.99	0.99	8980
Weighted Avg	0.99	0.99	0.99	8980

The model demonstrated outstanding performance on the test dataset, achieving an overall accuracy of 98.98%, indicating that nearly all instances were correctly classified as either fake or real news. The precision score of 98.22% shows that when the model predicts a news item as fake, it is correct approximately 98% of the time. The recall score of 99.65% highlights the model's ability to identify nearly all actual fake news cases, minimizing false negatives. The F1-score, which balances precision and recall, is exceptionally high at 98.93%, confirming the model's robustness and reliability in detecting fake news.

Looking at the detailed classification report, the model achieved perfect precision (1.00) for fake news and a recall of 0.98, resulting in an F1-score of 0.99. For real news, it scored 0.98 in precision and 1.00 in recall, also yielding an F1-score of 0.99. These metrics suggest that the model is equally effective in classifying both fake and real news. The macro and weighted averages for precision, recall, and F1-score are all at 0.99, further emphasizing the model's high generalization capability and balanced performance across both classes.

4.3. Discussion

The findings of this study demonstrate the superior performance of a Bidirectional Long Short-Term Memory (BiLSTM) model in classifying fake and real news articles using the ISOT dataset. In contrast to prior approaches such as the MM-FND framework by Alsuwat and Alsuwat [33], which relied on multimodal feature fusion and achieved 96.3% accuracy, the proposed BiLSTM achieved 98.98% accuracy, 98.22% precision, 99.65% recall, and an F1-score of 98.93%. These results suggest that the BiLSTM's ability to capture contextual and sequential dependencies in text provides a distinct advantage over models dependent on handcrafted or multimodal features. The particularly high recall indicates the model's effectiveness in minimizing false negatives, a crucial aspect of fake news detection where undetected misinformation poses serious risks.

Table 5. Performance comparison of the proposed BiLSTM model with state-of-the-art approaches on the ISOT dataset

References	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Alsuwat & Alsuwat [33]	96.30	95.80	97.10	96.40
Fayaz et al. [34] (chi^2 best)	97.33	–	–	–
Yu et al. [35]	98.95	–	–	98.92
Rani and Shokeen [36]	98.53	98.57	98.30	–
Kansal [37]	92.32	–	–	91.85
Proposed Model (This Study)	98.98	98.22	99.65	98.93

Beyond numerical improvements, the comparative analysis in Table 5 underscores key insights into why the proposed approach is superior. Prior studies, including those by Kansal [37] and Rani and Shokeen [36], introduced techniques such as POS tagging or blockchain-based ensembles, yet achieved only moderate improvements at the cost of increased complexity. Similarly, ensemble-based frameworks, such as those reported by Yu et al. [35], achieved competitive results but required computationally expensive integrations. By contrast, the BiLSTM achieves state-of-the-art performance with a streamlined pipeline consisting of an embedding layer for semantic representation, a BiLSTM for contextual learning, and lightweight dense layers for classification. This implies that the model's success lies not only in its sequential architecture but also in its design simplicity, which enhances scalability and facilitates real-time deployment on social media platforms. While early signs of overfitting were observed after the second epoch, suggesting the need for regularization strategies, the findings support the hypothesis that a focused sequential architecture, such as BiLSTM, can outperform more complex feature-fusion or ensemble-based models. This proves that simplicity in model design can lead to both higher accuracy and reduced computational cost.

5. Conclusions

This study successfully achieved its objective of proposing and implementing a standalone Bidirectional Long Short-Term Memory (BiLSTM) model for fake news detection using the ISOT dataset. The model demonstrated superior performance, attaining an accuracy of 98.98% along with near-perfect precision, recall, and F1 scores for both fake and real news classifications. These findings highlight the effectiveness of leveraging the sequential nature of textual data, showing that an optimized BiLSTM can outperform more complex, multimodal frameworks while offering a simpler and more practical solution for real-world applications. However, the study has limitations. Although the ISOT dataset provides a substantial number of records, the distribution of articles is uneven across categories. For instance, the majority of real news is concentrated in Politics and World News, while fake news is heavily skewed toward broad categories like News and Politics, with relatively fewer samples in specialized areas such as the Middle East and US News. This imbalance may limit the model's ability to generalize across diverse topics and real-world scenarios where fake news manifests in more varied forms.

Additionally, signs of potential overfitting suggest that the model's generalizability requires further validation. Future work should therefore explore data augmentation or resampling techniques to address class and topical imbalances, and consider hybrid architectures that combine BiLSTM with other advanced models. Future extensions will focus on Nigerian-specific datasets that reflect the linguistic, cultural, and socio-political landscape of the region. The spread of misinformation during Nigeria's 2023 general elections highlights the need for locally adaptable models that can effectively handle code-switching, regional dialects, and culturally nuanced narratives. Building upon the robust baseline established here, Such efforts would enhance robustness, broaden applicability, and further strengthen the model's potential for deployment in high-stakes information ecosystems.

Author Contributions: Conceptualization: A.A. and M.M.L.; Methodology: M.M.L.; Software: A.A.; Validation: A.A. and M.M.L.; Formal analysis: A.A.; Investigation: A.A. and M.M.L.; Resources: A.A.; Data curation: A.A.; Writing—original draft preparation: A.A.; Writing—review and editing: A.A. and M.M.L.; Visualization: A.A.; Supervision: M.M.L.; Project administration: M.M.L.; Funding acquisition: A.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The dataset is publicly available on the Kaggle platform and can be accessed via URL: https://www.kaggle.com/datasets/emineyetm/fake-news-detection-datasets.

Conflicts of Interest: The authors declare no conflict of interest.

References

[1] A. Bondielli and F. Marcelloni, "A survey on fake news and rumour detection techniques," *Inf. Sci. (Ny).*, vol. 497, pp. 38–55, Sep. 2019, doi: 10.1016/j.ins.2019.05.035.

[2] C. Nwonwu, F. Tukur, and O. Alo, "Nigeria elections: Websites use false stories to attract views and ads," *bbc.com*, 2023. https://www.bbc.com/news/world-africa-66647768

[3] O. Sunday, "NewsWike did not slump in Abuja – Lere Olayinka," *Daily Post Nigeria*, 2025. https://dailypost.ng/2025/03/29/wike-did-not-slump-in-abuja-lere-olayinka/

[4] 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.

[5] X. Zhou and R. Zafarani, "Network-based Fake News Detection," *ACM SIGKDD Explor. Newsl.*, vol. 21, no. 2, pp. 48–60, Nov. 2019, doi: 10.1145/3373464.3373473.

[6] K. Langin, "Fake news spreads faster than true news on Twitter—thanks to people, not bots," *Science (80-.).*, Mar. 2018, doi: 10.1126/science.aat5350.

[7] 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.

[8] A. Zubiaga, A. Aker, K. Bontcheva, M. Liakata, and R. Procter, "Detection and Resolution of Rumours in Social Media," *ACM Comput. Surv.*, vol. 51, no. 2, pp. 1–36, Mar. 2019, doi: 10.1145/3161603.

[9] S. Vosoughi, D. Roy, and S. Aral, "The spread of true and false news online," *Science (80-.).*, vol. 359, no. 6380, pp. 1146–1151, Mar. 2018, doi: 10.1126/science.aap9559.

[10] Z. Zhao, P. Resnick, and Q. Mei, "Enquiring Minds," in *Proceedings of the 24th International Conference on World Wide Web*, May 2015, pp. 1395–1405. doi: 10.1145/2736277.2741637.

[11] S. Kumar, R. West, and J. Leskovec, "Disinformation on the Web," in *Proceedings of the 25th International Conference on World Wide Web*, Apr. 2016, pp. 591–602. doi: 10.1145/2872427.2883085.

[12] 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.

[13] L. Ilias and I. Roussaki, "Detecting malicious activity in Twitter using deep learning techniques," *Appl. Soft Comput.*, vol. 107, p. 107360, Aug. 2021, doi: 10.1016/j.asoc.2021.107360.

[14] E. Puraivan, R. Venegas, and F. Riquelme, "An empiric validation of linguistic features in machine learning models for fake news detection," *Data Knowl. Eng.*, vol. 147, p. 102207, Sep. 2023, doi: 10.1016/j.datak.2023.102207.

[15] P. K. Verma, P. Agrawal, I. Amorim, and R. Prodan, "WELFake: Word Embedding Over Linguistic Features for Fake News Detection," *IEEE Trans. Comput. Soc. Syst.*, vol. 8, no. 4, pp. 881–893, Aug. 2021, doi: 10.1109/TCSS.2021.3068519.

[16] M. García Lozano *et al.*, "Veracity assessment of online data," *Decis. Support Syst.*, vol. 129, p. 113132, Feb. 2020, doi: 10.1016/j.dss.2019.113132.

[17] H. A. Santoso, E. H. Rachmawanto, A. Nugraha, A. A. Nugroho, D. R. I. M. Setiadi, and R. S. Basuki, "Hoax classification and sentiment analysis of Indonesian news using Naive Bayes optimization," *TELKOMNIKA (Telecommunication Comput. Electron. Control.*, vol. 18, no. 2, p. 799, Apr. 2020, doi: 10.12928/telkomnika.v18i2.14744.

[18] S. Rezaei, M. Kahani, B. Behkamal, and A. Jalayer, "Early multi-class ensemble-based fake news detection using content features," *Soc. Netw. Anal. Min.*, vol. 13, no. 1, p. 16, Dec. 2022, doi: 10.1007/s13278-022-01019-y.

[19] 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.

[20] I. Segura-Bedmar and S. Alonso-Bartolome, "Multimodal Fake News Detection," *Information*, vol. 13, no. 6, p. 284, Jun. 2022, doi: 10.3390/info13060284.

[21] Y. Tashtoush, B. Alrababah, O. Darwish, M. Maabreh, and N. Alsaedi, "A Deep Learning Framework for Detection of COVID-19 Fake News on Social Media Platforms," *Data*, vol. 7, no. 5, p. 65, May 2022, doi: 10.3390/data7050065.

[22] R. K. Kaliyar, A. Goswami, P. Narang, and S. Sinha, "FNDNet – A deep convolutional neural network for fake news detection," *Cogn. Syst. Res.*, vol. 61, pp. 32–44, Jun. 2020, doi: 10.1016/j.cogsys.2019.12.005.

[23] G. Potey, R. Jadhav, K. Shroff, A. Gore, D. Phalke, and J. Shimpi, "Fake News Detection and Sentiment Analysis in Twitter," *Int. J. Res. Appl. Sci. Eng. Technol.*, vol. 8, no. 4, pp. 1572–1577, 2020.

[24] 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.

[25] B. Akinyemi, O. Adewusi, and A. Oyebade, "An Improved Classification Model for Fake News Detection in Social Media," *Int. J. Inf. Technol. Comput. Sci.*, vol. 12, no. 1, pp. 34–43, Feb. 2020, doi: 10.5815/ijitcs.2020.01.05.

[26] S. Bauskar, V. Badole, P. Jain, and M. Chawla, "Natural Language Processing based Hybrid Model for Detecting Fake News Using Content-Based Features and Social Features," *Int. J. Inf. Eng. Electron. Bus.*, vol. 11, no. 4, pp. 1–10, Jul. 2019, doi: 10.5815/ijieeb.2019.04.01.

[27] J. Y. Khan, M. T. I. Khondaker, S. Afroz, G. Uddin, and A. Iqbal, "A benchmark study of machine learning models for online fake news detection," *Mach. Learn. with Appl.*, vol. 4, p. 100032, Jun. 2021, doi: 10.1016/j.mlwa.2021.100032.

[28] J. Alghamdi, Y. Lin, and S. Luo, "A Comparative Study of Machine Learning and Deep Learning Techniques for Fake News Detection," *Information*, vol. 13, no. 12, p. 576, Dec. 2022, doi: 10.3390/info13120576.

[29] I. Ahmad, M. Yousaf, S. Yousaf, and M. O. Ahmad, "Fake News Detection Using Machine Learning Ensemble Methods," *Complexity*, vol. 2020, pp. 1–11, Oct. 2020, doi: 10.1155/2020/8885861.

[30] M. H. Goldani, S. Momtazi, and R. Safabakhsh, "Detecting fake news with capsule neural networks," *Appl. Soft Comput.*, vol. 101, p. 106991, Mar. 2021, doi: 10.1016/j.asoc.2020.106991.
[31] S. Girgis, E. Amer, and M. Gadallah, "Deep Learning Algorithms for Detecting Fake News in Online Text," in *2018 13th International Conference on Computer Engineering and Systems (ICCES)*, Dec. 2018, pp. 93–97. doi: 10.1109/ICCES.2018.8639198.
[32] M. Sarnovský, V. Maslej-Krešňáková, and K. Ivancová, "Fake News Detection Related to the COVID-19 in Slovak Language Using Deep Learning Methods," *Acta Polytech. Hungarica*, vol. 19, no. 2, pp. 43–57, 2022, doi: 10.12700/APH.19.2.2022.2.3.
[33] E. Alsuwat and H. Alsuwat, "An improved multi-modal framework for fake news detection using NLP and Bi-LSTM," *J. Supercomput.*, vol. 81, no. 1, p. 177, Jan. 2025, doi: 10.1007/s11227-024-06671-z.
[34] M. Fayaz, A. Khan, M. Bilal, and S. U. Khan, "Machine learning for fake news classification with optimal feature selection," *Soft Comput.*, vol. 26, no. 16, pp. 7763–7771, Aug. 2022, doi: 10.1007/s00500-022-06773-x.
[35] C. Chang Yu, I. R. A Hamid, Z. Abdullah, K. Kipli, and H. Amnur, "A Multi-tier Model and Filtering Approach to Detect Fake News Using Machine Learning Algorithms," *JOIV Int. J. Informatics Vis.*, vol. 8, no. 2, p. 643, May 2024, doi: 10.62527/joiv.8.2.2703.
[36] P. Rani and J. Shokeen, "FNNet: a secure ensemble-based approach for fake news detection using blockchain," *J. Supercomput.*, vol. 80, no. 14, pp. 20042–20079, Sep. 2024, doi: 10.1007/s11227-024-06216-4.
[37] A. Kansal, "Fake News Detection Using Pos Tagging and Machine Learning," *J. Appl. Secur. Res.*, vol. 18, no. 2, pp. 164–179, Apr. 2023, doi: 10.1080/19361610.2021.1963605.