International Journal of Wireless and Ad Hoc Communication

♦ ISSN: 2692-4056 Vol: 11, No. 01, 2027 ♦ PP. 12–21

Article DOI: https://doi.org/10.54216/IJWAC.110102



FULL LENGTH ARTICLE

Ground Truth-Oriented Fake News Detection Using Sentiment-Aware Hybrid Ensemble Learning and Bi-Directional LSTM Networks

Alcardo Alex Barakabitze[1,*]

[1] AI and Emerging Technologies Lab, Department of Informatics and Information Technology, College of Natural and Applied Sciences, Sokoine University of Agriculture, Morogoro, Tanzania

Email: alcardo.alex0@sua.ac.tz

Received: May 31, 2026 Revised: September 01, 2026 Accepted: November 02, 2026 ⋆ Corresponding author

ABSTRACT

Fake news is the spreading of misinformation or disinformation. The proliferation of fake news on social media makes it difficult for one to procure reliable information online. Our motive is for information to be freely and widely available to all, while maintaining the integrity of said information. Ground truth refers to the actual reality or a baseline truth. Existing research focuses on the importance of deciphering the ground truth value and into a feature-oriented approach [1]. This paper contributes to the existing body of literature by meticulously selecting features and leveraging emotional cues through natural language processing (NLP), to address the issue of fake news detection. This is done by availing of the Truth Seeker [2] and Liar [3] datasets. The models used are an ensemble model which is a combination of Naive Bayes, passive aggressive classifier with a layer of deep neural network and Bi-directional LSTM (Long Short-Term Memory) networks with sentiment analysis. Our findings suggest that sentiment analysis can serve as a valuable tool for enhancing the detection of fake news.

Keywords: Sentiment ▪ Social Media ▪ Classification ▪ Ground Truth ▪ Machine Learning ▪ Fake News ▪ Feature extraction

1. INTRODUCTION

Misinformation and disinformation are pervasive in the modern digital era. It permeates every layer of society. Fake news deceives people at an ever increasing rate [4]. AI has the capability of generating articles and rendering images that are convincingly realistic [5]. This can create a very large number of fake news articles that can not be manually reviewed [6]. Thus, it is crucial to automate the process of fake news detection. It is of utmost importance for individuals to possess tools that can help them distinguish between what is real and what is not. Knowledge wields power, and information is the gateway to knowledge. Without accurate information, one would be left paralysed, unsure of what to believe and what not to believe. This can be a nuisance in normal times and have dire consequences in times where urgent action is required. Effective solutions are needed that can distinguish between real news articles and fake ones to mitigate the impact of fake news.

1.1 Contributions of the Paper

This paper makes the following key contributions to the field of fake news detection.

- We proposes a hybrid fake news detection framework that integrates Naïve Bayes (NB), Passive Aggressive Classifier (PAC), Deep Neural Networks (DNN), and Bidirectional Long Short-Term Memory (Bi-LSTM) to improve the ac-

© 2027 The Author(s). Published by ASPG. This article is licensed under CC BY 4.0.

curacy and robustness of fake news classification across diverse news datasets.

- We develop an effective feature engineering strategy by combining textual representations with sentiment-based features, enabling the model to capture both linguistic patterns and underlying emotional intent. This integrated representation improves the detection of misinformation that may otherwise resemble legitimate news.
- We provide an extensive performance evaluation using two publicly available benchmark datasets, TruthSeeker and LIAR, providing a comprehensive assessment of its effectiveness. Experimental results demonstrate that the proposed hybrid framework achieves competitive performance compared with existing state-of-the-art fake news detection approaches.
- We provide a scalable and intelligent solution that can support social media platforms, news organizations, and fact-checking agencies in automatically identifying and filtering misleading content. The framework contributes to enhancing information integrity and reducing the societal impact of misinformation.

This paper advance the state of the art in fake news detection by introducing an emotion-aware hybrid artificial intelligence framework that combines traditional machine learning, deep learning, and natural language processing techniques to improve the accuracy, robustness, and interpretability of automated misinformation detection systems.

2. RELATED WORK

The detection of fake news is a pressing challenge in today's digital age. The surge of misinformation has led to significant research efforts aimed at devising effective mechanisms for its detection. One such method is the Rhetorical Approach [7], but the burgeoning potential of machine learning, especially deep learning models like recurrent neural networks, offers promising avenues [8].

2.1 Sentiment Analysis in Fake News Detection

Alonso et al. [9] delve into the domain of sentiment analysis as a tool for detecting fake news. Their work likely illuminates the unique sentiment patterns in misinformation, which may exhibit more polarized or biased emotional tones. The utilization of sentiment analysis techniques, perhaps leveraging deep learning, forms a potent tool to discern the authenticity of news. Shearer and Mitchell [10] examine how Americans accessed news through different social media platforms during 2020. The study shows that social media had become an important channel for news consumption, although its importance varied considerably across platforms. Facebook was the most prominent social media source of news, with about 36% of U.S. adults regularly obtaining news from the platform, followed by YouTube (23%), Twitter (15%), and Instagram (11%). Bhutani et al[11] propose an approach that treats the emotional or sentiment characteristics of news content as an important feature for distinguishing fake news from genuine news. The study evaluates the proposed method using three different datasets to assess its robustness and effectiveness. The results indicate that incorporating sentiment information can improve the performance of fake news detection, and comparisons with other existing approaches demonstrate the effectiveness of the proposed method.

2.2 Deep Learning and Sentiment Classification

The prowess of deep learning, especially recurrent neural networks has been well-established in sentiment analysis tasks. Sharfuddin et al. [12] present a comprehensive study on the BiLSTM (Bidirectional Long Short-Term Memory) model for sentiment classification. Their work underscores the significance of considering both past and future data in sequences to enhance contextual understanding in sentiment analysis tasks [12]. Islam et al [13] evaluated multiple advanced

deep-learning models across four benchmark datasets to identify the most effective approach for sentiment analysis. The proposed Capsule with Deep CNN and Bi-structured RNN model outperformed the comparative methods and achieved high classification accuracy across all evaluated datasets by 88.15% accuracy on IMDB, 98.28% on Toxic, 92.34% on CrowdFlower, and 95.48% on ER. These results demonstrate the model's robustness and effectiveness across diverse sentiment-analysis benchmarks.

Sindhu et. al [14] proposed a CNN–LSTM deep-learning model combined with an advanced word-embedding technique to improve sentiment classification. Using a large Amazon dataset, the model was developed to effectively capture both local textual features and sequential dependencies for sentiment prediction. The proposed approach achieved 98.20% accuracy, 99.86% precision, and 99.50% recall, demonstrating strong performance in sentiment analysis. Ramezani et al [15] present a review of advances in deep learning-based sentiment analysis on social networks from 2019 to May 2024. Authors demonstrates that sentiment analysis can extract valuable insights from large and diverse social media datasets, supporting decision-making, personalized services, and assessment of digital social well-being. Particular attention is given to mental health, using users' emotional expressions as an indicator of the emotional environment within social networks. The review also develops a technical taxonomy and comparison of deep learning approaches used for sentiment analysis, while identifying key challenges, limitations, and open research issues. Tan et al

[16] proposes a deep neural network-based multi-task learning framework that exploits the relationship between related tasks to enhance overall sentiment analysis performance. The proposed approach achieves an F1-score of 94%, outperforming existing methods by approximately 3%, demonstrating the effectiveness of jointly learning related tasks.

Our Approach to Fake News Detection : Building on the foundational work of Chora et al. [8], we evaluated the veracity of entire articles in our research. We also incorporated ensemble models, as suggested by Shu et al. [1, 17]. Our evaluations revealed that the LSTM model outperformed the ensemble models, a finding we confirmed in our implementation.1 Specifically, a comparison of the receiver operating characteristic (ROC) curves for our LSTM and ensemble models indicated the superiority of the LSTM model. This observation aligns with findings from other studies [1, 17]. Given its higher accuracy and enhanced capability to discern

Table 1. Definitions of fake news, misinformation, disinformation, and news articles.

Term	Definition
Fake news	A news article that is verifiable as false; it encompasses both misinformation and disinformation.
Misinformation	The unintentional or unknowing dissemination of false or inaccurate information without the intent to deceive.
Disinformation	The deliberate creation and dissemination of false or misleading information with the intention of deceiving others.
News article	A piece of written content that provides information about current events, issues, or developments intended for public dissemination.

real news from fake, we chose the LSTM model for our final application. However, the journey to authenticating news is riddled with challenges. Shu et al.'s [1] premise that fake news is "intentionally and ver-ifiably false" has been subjected to scrutiny, highlighting the inadvertent spread of misinformation. The multifaceted nature of fake news detection requires an amalgamation of techniques, and while Alonso et al. [9] have made strides in multilingual fake news detection, the challenge of addressing inherent biases in natural language processing remains [9].

1https://github.com/SuperSaiyansUcd/

Fake-News-Detection-/tree/main/model

In conclusion, the landscape of fake news detection is vast, with various methodologies showing promise [18]. Our research leans towards the LSTM and ensemble models, resonating with prevailing research trends and the promise they hold in this domain [1, 17].

3. METHODOLOGY: APPROACH AND IMPLEMENTATION

This section describes the implementation of the fake news detection application. The application mainly consists of three parts: the front-end, the back-end and the model.2 For the front end, we considered both React and Angular. We ultimately chose React because it is a library, whereas Angular is a full-fledged framework. React is a JavaScript library known for its efficiency and flexibility in building user interfaces. In contrast, Angular enforces a more rigid, "Angular-specific" development structure, while React offers greater freedom in how applications are designed and implemented. This flexibility was particularly beneficial for our team, as not all members were familiar with Angular.

Furthermore, we explored the possibility of using Million.js to enhance performance.3 However, to fully benefit from Million.js optimizations, developers must adhere to a specific set of rules involving the use of blocks; otherwise, performance reverts to that of standard React. In addition, the computationally intensive part of our system resides in the model rather than the user interface layer. As a result, extreme front-end optimization was not a primary requirement. For more details, readers are advised to visit4.

3.1 Dataset

We selected two datasets for our research: "Truth Seeker" [2] and the "LIAR" dataset [3].

3.1.1 Training

The models were primarily trained on the "Truth Seeker" dataset. An essential aspect of any fake news detection dataset is the binary "True" or "False" labels assigned to news items or statements. In the case of "Truth Seeker", labeling was conducted manually. Such labeled data facilitates the training of machine learning models, enabling them to distinguish between authentic and misleading news. These models are adept at recognizing patterns, terms, and other unique characteristics of fake news. Furthermore, the credibility of the sources within the dataset can be leveraged to ascertain the accuracy of the ground truth labels. A balanced dataset, comprising an equal mix of genuine and fake news articles, provides a solid foundation for evaluating model performance, including metrics such as accuracy, recall, precision, and F1-score.

2https://github.com/SuperSaiyansUcd/Fake-News-Detection-3https://million.dev/

4Angular, https://angular.io/ and React, https://react.dev/

3.1.2 Testing

For testing, we turned to the "LIAR" dataset. Our choice was influenced by the comprehensive nature of LIAR, which boasts over 12,800 statements vetted for their truthfulness by Politi-Fact editors. Given its widespread recognition in the research community, the LIAR dataset stands as a benchmark for fake news detection algorithms. Model efficacy is often gauged by performance on this dataset. However, it's noteworthy that we refrained from using this dataset for training due to its unbalanced nature, which poses a risk of overfitting.

3.1.3 Feature Engineering

We undertook rigorous pre-processing of the dataset. Text data was streamlined for easy processing using tokenization. Additionally, we employed the Term Frequency-Inverse Document Frequency (TF-IDF) method to numerically vectorize the text, ensuring that the semantic essence of the content is encapsulated in its numerical representation. TF-IDF assigns weights to words based on their frequency in a given document relative to their occurrence across multiple documents, spotlighting particularly significant terms.

This technique reveals the underlying sentiment or emotional undertone of a text string, offering insights into the perspectives of users, consumers, or any textual feedback. For this, we employed a gamut of tools: NLTK's Sentiment Intensity Analyzer (VADER), TextBlob, Flair's Text Classifier, Afinn, Pattern's sentiment module, and Transformers with BERT.

Notably, BERT, a leading-edge transformer-based model, excels at grasping the contextual meaning of words in a text by examining their surroundings in the text.

We used Flask architecture indicated in Fig. 1, a framework for developing web applications using python. We picked flask due to its minimal and easy to use design while Django is more bulky [19]. Flask also maintains security as Flask's Jinja2 templates automatically escape cross site scripting attempts.

- Flask Framework: The application is built on the Flask framework, which provides a robust foundation for handling HTTP requests and responses.
- Sentiment Score Calculation: This component utilises multiple sentiment analysis libraries to compute sentiment scores for the content pasted by the user. These scores are used as a feature for the model. The sentiment analysis libraries include:
- Afinn: Utilises a pre-built word list to calculate a sentiment score.
- Pattern: Provides sentiment analysis based on patterns in the text.
- NLTK's VADER: Analyses sentiment by considering the intensity of positive, negative, and neutral words.
- TextBlob: calculates sentiment polarity based on the content.
- Emotion Prediction Model: A pre-trained emotion prediction model is integrated to predicts Ekman's 6 basic emotions in addition to a neutral class present in the content.5
- Saved Model: The component includes a saved fake news detection model created using a combination of Bidirectional LSTM layers and dense layers. This model accepts content text, emotion features, and sentiment scores as inputs to predict the authenticity of news articles.
- API Endpoints: This entity defines two API endpoints:
- Web-Scrap API: Accepts a URL and extracts the title and main text content using Beautiful Soup.
- Submit-API: Accepts title and content, performs sentiment analysis, emotion prediction, and fake news detection, and returns analysis results as JSON.

Flask derives WSGI application [20] functionality from Gunicorn WSGI Server which is a comprehensive WSGI web application library. With this WSGI standard, web servers have to implement the server side of the WSGI interface and Flask application object is the actual WSGI application.

The workflow of the Fake News Detection Web App as indicated in Fig. 2. The results, including sentiment scores, emotion predictions, authenticity prediction, and various metrics, are sent in JSON format to the user. The URL redirection and access using the restriction algorithm in 2

- User submits a URL or textual content which is then sent to the appropriate web server.
- Sentiment analysis libraries compute sentiment scores for the content.
- The sentiment model as proposed in section 3.2 detects emotions present in the news article.

5https://huggingface.co/j-hartmann/

emotion-english-distilroberta-base

- The fake news detection model combines sentiment scores, emotion features, and content text to determine the authenticity of a news article.
- The results, including sentiment scores, emotion predictions, authenticity prediction, and various metrics, are sent in JSON format to the user. The URL redirection and access restriction2

3.2 The Fake News Detection Model

This section outlines the implementation of a fake news detection model using a combination of Bi-directional LSTM (Long Short-Term Memory) networks and sentiment analysis. The model is designed to predict the authenticity of news articles based on textual content, emotions, and sentiment scores.

The first step is data preparation and pre-processing of the truthseeker dataset,6 which is the largest ground truth fake news analysis dataset for real and fake news content in relation to social media posts [2], containing information about news articles along with emotional and sentiment features. The following steps are done in order to create the ideal dataset for the paper. Sentiment scores are computed using various sentiment analysis libraries, including TextBlob, Afinn, and VADER. The relevant columns for the model are selected, including 'statement', 'tweet', 'BinaryNumTarget', emotion features ('anger', 'disgust', 'fear', 'joy', 'neutral', 'sadness', 'surprise'), and sentiment scores ('afinn', 'pattern', 'vader', 'textblob'). The emotions are calculated using the same hugging face transformer that we used on

the back-end.7

The next step is the model architecture which is composed of three main branches:

6https://www.unb.ca/cic/datasets/truthseeker-2023.html 7https://huggingface.co/j-hartmann/

emotion-english-distilroberta-base

- Textual Input Branch: The tweet and the statement of tweets are combined and processed using an embedding layer followed by a Bidirectional LSTM layer to capture sequential information.
- Emotion Input Branch: Emotion features are passed through a Dense layer with 64 neurons.
- Sentiment Scores Input Branch: Sentiment scores are passed through another Dense layer with 64 neurons.

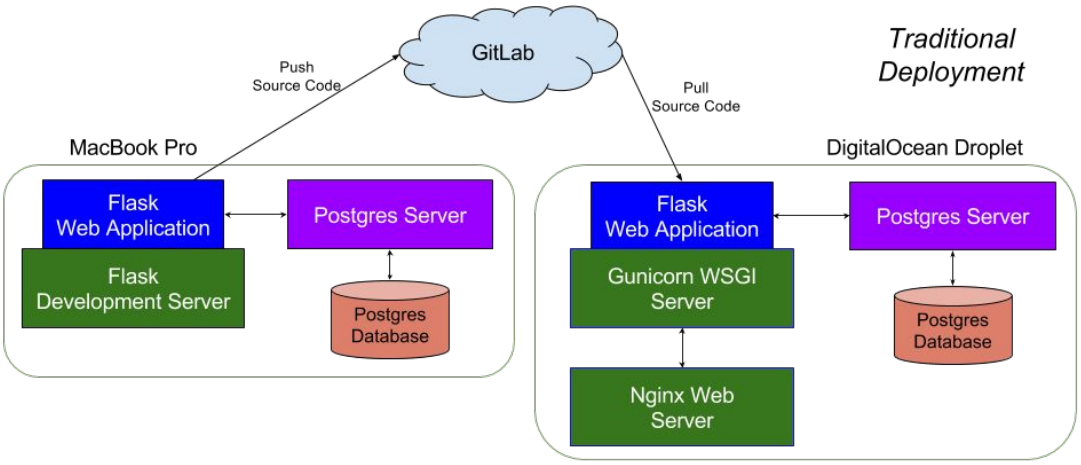


Figure 1. Architecture of the Flask application.

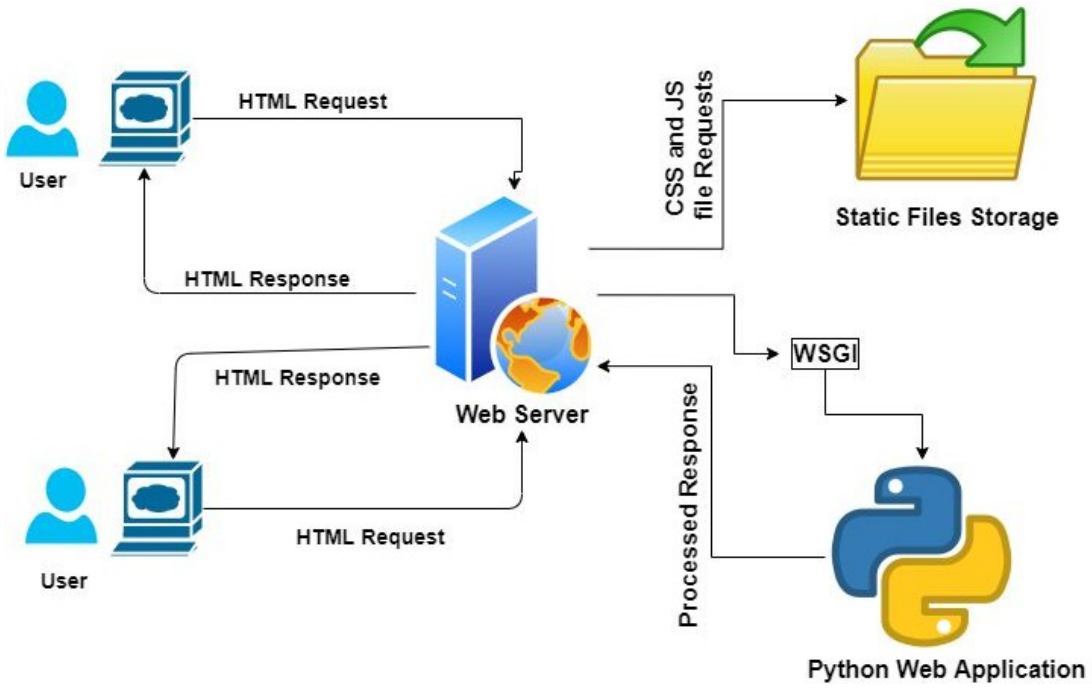


Figure 2. Workflow of the Python web application.

The outputs of the three branches are concatenated and fed through a Dense layer with 64 neurons, followed by the final output layer with a sigmoid activation function for binary classification (fake or real).

The final step is the model training and saving phase. The compiled 'best model' is trained using prepared data, optimising its ability to distinguish between fake and genuine news.

3.2.1 Algorithm 1: Sentiment- and Emotion-Aware Hybrid Deep Learning

Input: Dataset D with statement and tweet fields and binary labels $Y = \{0, 1\}$.
Output: Trained model M^* saved as `best_model.h5`.

1. Remove duplicate records, handle missing values, and preprocess text by removing unnecessary symbols, URLs, and special characters.
2. For each record $d_i \in D$, extract its statement and tweet and form $X_i^{\text{text}} = \text{Concat}(\text{statement}_i, \text{tweet}_i)$.
3. Compute sentiment scores using AFINN, Pattern, TextBlob, and VADER: $X_i^{\text{sent}} = [s_i^{\text{AFINN}}, s_i^{\text{Pattern}}, s_i^{\text{TextBlob}}, s_i^{\text{VADER}}]$.
4. Apply the Hugging Face transformer-based emotion classifier and form $X_i^{\text{emo}} = [\text{anger}, \text{disgust}, \text{fear}, \text{joy}, \text{neutral}, \text{sadness}, \text{surprise}]$.
5. Select the text, emotion, and sentiment features and split the prepared data into training, validation, and test subsets.
6. Encode textual input with an embedding layer and a Bidirectional LSTM: $H^{\text{text}} = \text{BiLSTM}(\text{Embedding}(X^{\text{text}}))$.
7. Map emotion and sentiment inputs through separate 64-neuron dense layers: $H^{\text{emo}} = \text{Dense}_{64}(X^{\text{emo}})$ and $H^{\text{sent}} = \text{Dense}_{64}(X^{\text{sent}})$.
8. Fuse the branches, apply a 64-neuron dense layer, and produce the binary-classification probability: $H = \text{Concat}(H^{\text{text}}, H^{\text{emo}}, H^{\text{sent}})$; $\hat{y} = \text{Sigmoid}(W\,\text{Dense}_{64}(H) + b)$.

9. Compile the model with binary cross-entropy loss, a selected optimizer, and accuracy; train on the training set and monitor validation performance.

10. Select the best validation model M^*, evaluate it on the unseen test set, and save it as `best_model.h5`.

11. For real-time prediction, preprocess and tokenize input x, compute x^{emo} and x^{sent}, and evaluate $p = M^*(x^{\text{text}}, x^{\text{emo}}, x^{\text{sent}})$.

12. If $p \geq 0.5$, return Real with probability p; otherwise return Fake with probability $1 - p$.

After training, the model's learned weights are saved as a ".h5" file, ensuring that the trained model can be easily reused or deployed for future predictions without needing to retrain from scratch. The h5 file facilitates the deployment of the trained

scratch. The h5 file facilitates the deployment of the trained

model as part of the web application. The Fake News Detection system can then make predictions in real-time based on the input provided by users.

3.3 DevOps

We deployed using the Microsoft Azure portal. This was done by availing of docker and the azure web-app server alongside a virtual machine (VM). Using the docker app leads to more efficient use of resources, while the virtual machine (VM) was faster in overall speed. Using our CI/CD pipeline the docker setup was set to fetch from origin/main every hour. In this way we get continuous deployment. This is useful when someone on the team wants to look at a branch without loading up the code. One can go to the build of said branch which is deployed. Then one can test the deployed application for bugs. This is useful as there may be bugs during deployment that are not present on the deployed code. The VM was of Standard B1ms size. The code was perpetually run using process manager 2 (PM2). The network security settings on Azure makes sure that only certain ports can be accessed. This prevents unauthorised access and nefarious activity on non-secure ports.

To enable the real-time response of the model, both the ensemble model and LSTM model were pre-loaded as .h5 files. This is a highly compatible file format. Ideal for ease of collaboration. It is also quite compact and efficient [21]. The pre-loading enables us to give users real time responses with only seconds of delay (while the model is processing the news article). The concurrent ability of Flask allows us to give real time responses to multiple users simultaneously. This highly improves the user experience. The Continuous Integration/Con-tinuous Deployment (CI/CD) pipeline was set up on GitHub. Inside said pipeline, it runs the build, ensuring that deployment is successful. It also ensures that we have the Azure key credentials, that the front-end tests run and pass, that all the formatting is correct that the system works integration testing wise, and that the web application builds successfully.

4. EVALUATION OF RESULTS

The Evaluation approach is reasonable as we used traditional NLP evaluation metrics (precision, recall, accuracy, and F1-score) to quantitatively assess the model's performance on a test dataset. A stratified k-fold cross-validation was also used which provided robustness in estimating the model's generalisation of the classification. The results of our evaluation reveals the effectiveness of our fake news detection solution. Our LSTM model, augmented with emotion features and sentiment scores, achieved an accuracy of 0.98, outperforming baseline methods.

The Evaluation metrics are given below.

4.0.1 Precision

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{1}$$

4.0.2 Recall

For class "Real News", the recall is 0.99, indicating that the model correctly identified 99% of the actual class "Real News" instances.

F1-Score: The F1-score is the harmonic mean of precision and recall, as seen in Equation (3). It provides a balanced measure of the model's performance. For both classes, the F1-scores are 0.98, indicating a good balance between precision and recall.

Initially, we thought this was an over-fitting issue. We attempted to resolve this by performing further feature extraction and k-fold cross-validation. Later on, in the work of Hamed et al. [17], there was an experiment conducted where sentiment and emotions are used as features on multiple models and datasets which resulted in similarly high results in accuracy. It is seen that LSTM models tend to have a high accuracy as observed by Choras et al. [8].

We also tested the model against another Bi-directional LSTM model which does not use Sentiment Analysis or Emotions as features and the following are the evaluation metrics. Although the model using sentiment analysis and emotions as features performs 1% improvement in overall performance compared with the baseline model, the inclusion of sentiment analysis and emotion features resulted in better overall predictive performance. This is further supported by additional evaluation metrics, including the calibration curve and confusion matrix. The confusion matrix demonstrated the model's balanced performance in identifying both fake and real news on the model that uses sentiment analysis.

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{2}$$

4.0.3 F1 Score

$$\text{F1 Score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{3}$$

Support is the number of occurrences in the dataset.

4.0.4 URL Redirection and Access Restriction

1. Receive a request for URL redirection.

2. If `Request.URL = /old-url`, redirect the user to `/new-url`.

Table 2. Classification report for the LSTM model with sentiment analysis.

Class	Precision	Recall	F1 score	Support
Fake news	0.99	0.97	0.98	13,053
Real news	0.97	0.99	0.98	13,786
Accuracy	–	–	0.98	26,839
Macro average	0.98	0.98	0.98	26,839
Weighted average	0.98	0.98	0.98	26,839

3. For a restricted request, if `Request.URL = /restricted`, return HTTP status code 403 and deny access to the requested resource.

Precision: Precision measures the accuracy of positive predictions, as seen in Equation (1). For class "Fake News", the precision is 0.99, which means that 99% of instances predicted as class "Fake News" were actually class "Fake News". For class "Real News", the precision is 0.97, indicating that 97% of instances predicted as class "Real News" were actually class "Real News".

Recall: Recall measures the ability of the model to correctly identify positive instances, as seen in Equation (2). For class "Fake News", the recall is 0.97, meaning that the model correctly identified 97% of the actual class "Fake News" instances.

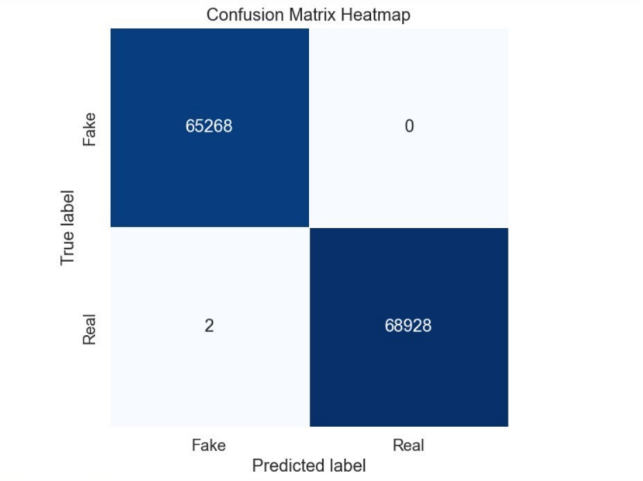


Figure 3. Confusion matrix for the sentiment-analysis model.

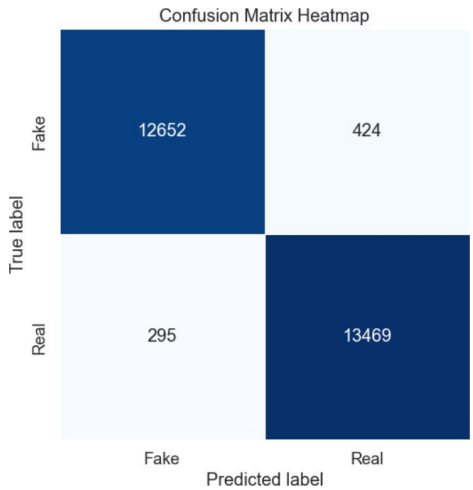


Figure 4. Confusion matrix for the baseline model.

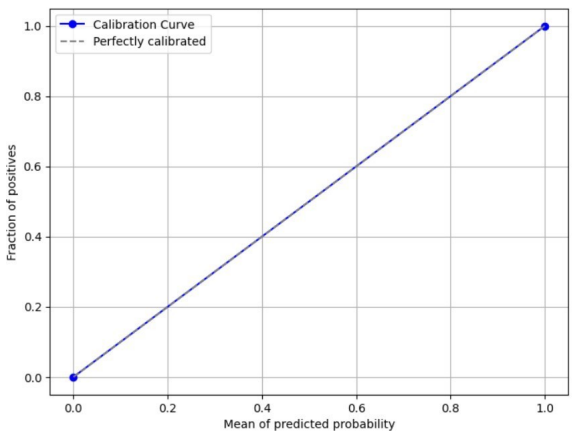


Figure 5. Calibration curve for the sentiment-analysis model.

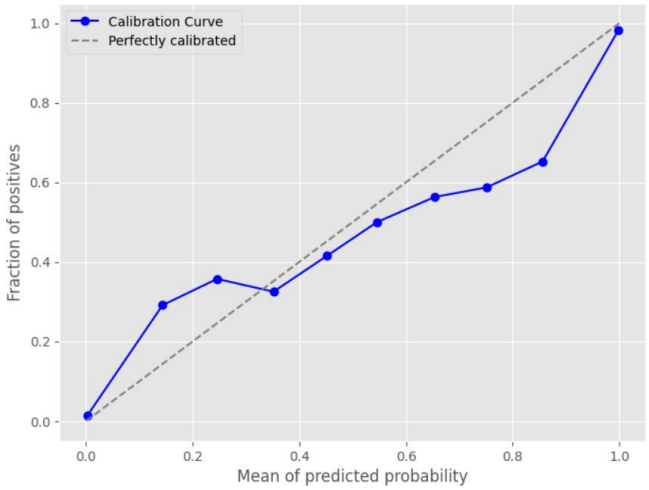


Figure 6. Calibration curve for the baseline model.

The calibration curve showed a close alignment between predicted and true probabilities, validating the reliability of our model's ability to predict whether the news is real or fake on the sentiment analysis model.

It is worth noting that using sentiment analysis as a feature can enhance the performance and reliability of the fake news detection model.

5. FUTURE WORK

5.1 Explainable, Multilingual, and Multimodal Detection Using Transformer-Based and Real-Time AI Models

Future research will focus on extending the proposed fake news detection framework toward a more fine-grained, multilingual, multimodal, explainable, and real-time misinformation detection system. First, instead of evaluating the truthfulness of an entire news item based on its overall content, the proposed system will be enhanced to assess the veracity of individual sentences or claims and provide evidence-based annotations explaining why each statement is classified as true, false, or potentially misleading. Such claim-level analysis can improve transparency and support users in understanding the specific elements contributing to a misinformation classification. The explainable multilingual and multimodal fake-news detection based on trustworthy AI for combating misinformation proposed by Jadhav [22] and Raza [23] can be a starting point.

Second, the current English-language focus will be extended to multilingual and cross-lingual fake news detection, en-

abling the system to process misinformation in multiple languages and low-resource linguistic environments. This direction, as discussed by Chalehchaleh et al. [24], is particularly important for global social-media platforms where misinformation propagates across linguistic and geographical boundaries.

Third, future studies will expand the training data by integrating larger, more diverse, and heterogeneous datasets from multiple sources, including social media, online news platforms, fact-checking organisations, and user-generated multimedia content [25], [26], [27]. Reliance on only two datasets may limit the generalisability of the current model. Cross-domain and cross-platform datasets will therefore be used to improve robustness and reduce dataset-specific bias.

Fourth, the framework will be extended beyond social-media text to support multimodal fake news detection, combining textual, visual, and potentially audio and video features. This will enable the detection of manipulated images, misleading captions, deepfakes, and other forms of synthetic or deceptive multimedia content. Feature engineering and representation learning will also be further refined to identify the most discrimi-native linguistic, semantic, sentiment, contextual, and temporal characteristics of misinformation.

Finally, the future improvement would be to evaluate the truthfulness of each line and provide annotations as to why that statement is true or false. Currently the system evaluates based on all the content. Multi-language support is lacking in our current implementation. Our training data focuses on detecting fake news in the English language. One big improvement that can be made is that of the data sources. We only avail of 2 datasets and this is not sufficient for general purpose fake news detection. Currently our system is only trained for the detection of fake news on social media. The model can be improved by further refining what features to extract. It is worth noting that our future work will include the integration of transformer-based language models (e.g., BERT, RoBERTa, and GPT), multimodal fake news detection using text and images, explainable

artificial intelligence (XAI), cross-lingual misinformation detection, and real-time detection systems for social media platforms. While limitations existed, the paper's outcomes and insights provide a solid foundation for future improvements and extensions [28, 1].

5.2 Implications for Multimedia Streaming in 6G and Beyond Networks

Future research should investigate the integration of ground-truth-oriented fake news detection, sentiment-aware hybrid ensemble learning, and Bi-Directional Long Short-Term Memory (Bi-LSTM) networks into intelligent multimedia streaming architectures for 6G and beyond networks [29], [30]. As next-generation networks [31] are expected to support highly immer-sive services, including ultra-high-definition video, extended reality (XR), holographic communications, social virtual reality, and AI-generated multimedia content, the rapid dissemination of misleading or manipulated information represents an important challenge for network intelligence and content trustworthiness. Future work could therefore develop trust-aware multimedia streaming SDN/NFV frameworks [32] capable of jointly analysing textual, visual, audio, and contextual features before or during content delivery. Multimodal learning models could combine sentiment representations with video frames, speech, metadata, user interaction patterns, and temporal characteristics to improve the detection of deepfakes, synthetic media, and coordinated misinformation.

Another important research direction is the development of edge-based fake news detection mechanisms for 6G networks. Deploying lightweight versions of the proposed hybrid ensemble and Bi-LSTM models at Multi-access Edge Computing (MEC) nodes could enable real-time content verification closer to users, reducing detection latency, backhaul traffic, and computational overhead [33]. This is particularly important for latency-sensitive applications such as live streaming, XR, autonomous systems, and emergency communications. Future studies should investigate model compression, knowledge distillation, federated learning, and resource-aware inference to enable efficient AI-based misinformation detection on resource-constrained edge devices.

The proposed approach also has implications for Quality of Experience (QoE) [34] and Quality of Information (QoI) management in future multimedia networks. Conventional streaming optimisation primarily considers parameters such as bandwidth, latency, packet loss, buffering, and video quality [35]. Future 6G systems could extend these mechanisms by incorporating content credibility and information quality as additional dimensions of streaming optimisation. For example, a network could dynamically prioritise, annotate, delay, or provide verification warnings for content associated with a high probability of misinformation while maintaining acceptable QoE. This would enable the evolution from conventional QoE-aware streaming [36], [37], [38], [39], [40], [41], [42] toward trust-aware and credibility-aware multimedia service management.

Furthermore, future research should explore the interaction between AI-generated content and AI-based content verification. The rapid development of generative AI is expected to increase the volume of synthetic text, images, audio, and video circulating through 6G networks. Consequently, detection models must continuously adapt to evolving manipulation techniques. Research should investigate continual learning, adversarial robustness, explainable AI, and provenance-aware detection using content authentication, watermarking, and blockchain-based or cryptographic provenance mechanisms. Such approaches could complement machine-learning-based detection by providing verifiable evidence about the origin and integrity of multimedia content.

to increase the volume of synthetic text, images, audio, and video circulating through 6G networks. Consequently, detection models must continuously adapt to evolving manipulation techniques. Research should investigate continual learning, adversarial robustness, explainable AI, and provenance-aware detection using content authentication, watermarking, and blockchain-based or cryptographic provenance mechanisms. Such approaches could complement machine-learning-based detection by providing verifiable evidence about the origin and integrity of multimedia content.

Another promising direction is privacy-preserving misinformation detection and multipath protections and dynamic link recovery in softwarized 5G networks [43]. Since future streaming platforms will potentially process large volumes of user-generated content and behavioural data, centralised

collection of such information may create privacy and security risks. Federated learning and privacy-preserving machine learning could allow detection models to learn from distributed multimedia environments without transferring sensitive user data to a central server. This would be particularly relevant for large-scale 6G ecosystems involving mobile users, IoT devices, edge servers, social-media platforms, and intelligent network infrastructures.

Future studies should evaluate the proposed framework using large-scale, multilingual, multimodal and real-time datasets representing diverse geographical and cultural contexts. Benchmarking should include conventional fake news as well as deepfakes, AI-generated news, manipulated videos, misleading headlines, coordinated disinformation campaigns, and cross-platform misinformation. Evaluation should move beyond classification accuracy, precision, recall, F1-score, and AUC to include detection latency, computational cost, energy consumption, network overhead, QoE, QoI, fairness, explainability, and robustness. Such comprehensive evaluation would facilitate the development of intelligent 6G multimedia streaming systems that are not only fast and scalable but also trustworthy, secure, and resilient.

6. CONCLUSION

This study investigated fake news detection by integrating textual, sentiment, and emotional features with machine learning and deep learning techniques. Using the TruthSeeker and LIAR datasets, the study demonstrated that incorporating sentiment information into BiLSTM-based models can provide valuable contextual cues for distinguishing fake from genuine news. The findings indicate that sentiment-aware approaches can complement conventional textual features and improve the reliability of automated fake news detection. The study therefore highlights the potential of combining sentiment analysis, emotional cues, and deep learning to develop more robust ground-truth-oriented detection systems. Future work should investigate larger and multilingual datasets, advanced transformer-based architectures, and real-time social media environments to further improve generalisation, robustness, and practical deployment.

ACKNOWLEDGEMENTS

I extend my sincere appreciation to our Microsoft mentors for providing access to Microsoft Azure, which was used to deploy our web application. I also thank the SUARIS - 2 funding for the valuable support for this paper.

REFERENCES

[1] K. Shu, A. Sliva, S. Wang, J. Tang, H. Liu, Fake news detection on social media: A data mining perspective, ACM SIGKDD Explorations Newsletter 19(1), pp. 22–36, 2017.

[2] S. Dadkhah, X. Zhang, A. G. Weismann, A. Firouzi, A. A. Ghorbani, TruthSeeker: The largest social media ground-truth dataset for real/fake content, 2023.

[3] W. Y. Wang, "Liar, Liar Pants on Fire": A new benchmark dataset for fake news detection, arXiv preprint arXiv:1705.00648, 2017.

[4] N. M. Brashier, D. L. Schacter, Aging in an era of fake news, Current Directions in Psychological Science 29(3), pp. 316–323, 2020.

[5] N. C. Köbis, B. Doležalová, I. Soraperra, Fooled twice: People cannot detect deepfakes but think they can, iScience 24(11), 2021.

[6] S. Vosoughi, D. Roy, S. Aral, The spread of true and false news online, Science 359(6380), pp. 1146–1151, 2018.

[7] N. K. Conroy, V. L. Rubin, Y. Chen, Automatic deception detection: Methods for finding fake news, Proceedings of the Association for Information Science and Technology 52(1), pp. 1–4, 2015.

[8] M. Choras, A. Giełczyk, K. Demestichas, D. Puchalski, R. Kozik, Pattern recognition solutions for fake news detection, in: Computer Information Systems and Industrial Management: 17th International Conference, CISIM, Olomouc, Czech Republic, Proceedings 17, Springer, pp. 130–139, 2018.

[9] M. A. Alonso, D. Vilares, C. Gómez-Rodríguez, J. Vilares, Sentiment analysis for fake news detection, Electronics 10(11), article 1348, 2021.

[10] E. Shearer, A. Mitchell, News use across social media platforms in 2020, Pew Research Center, 2021.

[11] A. A. Sharfuddin, M. N. Tihami, M. S. Islam, A deep recurrent neural network with BiLSTM model for sentiment classification, in: International Conference on Bangla Speech and Language Processing (ICBSLP), pp. 1–4, 2018.

[12] M. S. Islam, M. N. Kabir, N. A. Ghani, et al., Challenges and future in deep learning for sentiment analysis: A comprehensive review and a proposed novel hybrid approach, Artificial Intelligence Review 57, article 62, 2024.

[13] L. Sindhu, M. Rameshkumar, G. R. Kumar, B. N. Manjunath, S. K. C, Deep learning for sentiment analysis: Exploring the power of deep learning techniques in opinion mining, in: International Conference on Evolutionary Algorithms and Soft Computing Techniques (EASCT), pp. 1–6, 2023.

[14] E. Bakhtiari Ramezani, Sentiment analysis applications using deep learning advancements in social networks: A systematic review, Neurocomputing 634, article 129862, 2025.

[15] Y. Y. Tan, C. O. Chow, J. Kanesan, J. H. Chuah, Y. Lim, Sentiment analysis and sarcasm detection using deep multi-task learning, Wireless Personal Communications 129(3), pp. 2213–2237, 2023.

[16] S. K. Hamed, M. J. Ab Aziz, M. R. Yaakub, Fake news detection model on social media by leveraging sentiment analysis of news content and emotion analysis of users' comments, Sensors 23(4), article 1748, 2023.

[17] X. Zhou, R. Zafarani, A survey of fake news: Fundamental theories, detection methods, and opportunities, ACM Computing Surveys 53(5), pp. 1–40, 2020.

[18] D. Ghimire, Comparative study on Python web frameworks: Flask and Django, 2020.

[19] J. Gardner, The web server gateway interface (WSGI), The Definitive Guide to Pylons, 2009.

[20] M. Folk, G. Heber, Q. Koziol, E. Pourmal, D. Robinson, An overview of the HDF5 technology suite and its applications, in: Proceedings of the EDBT/ICDT Workshop on Array Databases, pp. 36–47, 2011.

[21] R. Jadhav, V. Meshram, A. Bhosle, K. Patil, S. Dash, S. Jadhav, Explainable multilingual and multimodal fake-news detection: Toward robust and trustworthy AI for combating misinformation, Frontiers in Artificial Intelligence 8, article 1690616, 2025.

[22] B. Bhutani, N. Rastogi, P. Sehgal, A. Purwar, Fake news detection using sentiment analysis, in: 12th International Conference on Contemporary Computing (IC3), pp. 1–5, 2019.

[23] N. Raza, S. J. Abdulkadir, Y. A. Abid, S. S. Albouq, A. Alwadain, A. U. Rehman, et al., Enhancing fake news detection with transformer-based deep learning: A multidisciplinary approach, PLoS ONE 20(9), article e0330954, 2025.

[24] R. Chalehchaleh, R. Farahbakhsh, N. Crespi, Addressing data scarcity in multilingual fake news detection: An LLM-based dataset augmentation approach, Social Network Analysis and Mining 15, article 92, 2025.

[25] C. Thibault, J.-J. Tian, G. Péloquin-Skulski, T. L. Curtis, J. Zhou, F. Laflamme, L. Y. Guan, R. Rabbany, J.-F. Godbout, K. Pelrine, A guide to misinformation detection data and evaluation, in: Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining, vol. 2, pp. 5801–5809, 2025.

[26] M. Nasser, N. I. Arshad, A. Ali, H. Alhussian, F. Saeed, A. Da'u, I. Nafea, A systematic review of multimodal fake news detection on social media using deep learning models, Results in Engineering 26, article 104752, 2025.

[27] G. Wei, K. Han, Y. Lyu, Y. Luo, Y. Jiang, C. Shan, N. Sebe, Towards real-time fake news detection under evidence scarcity, arXiv preprint arXiv:2510.11277, 2025.

[28] M. Fernández-López, M. Perea, Language does not modulate fake news credibility, but emotion does, Psicológica Journal 41(2), pp. 84–102, 2020.

[29] A. Barakabitze, A. Arslan, M. Rashid, A. Hines, 5G network slicing using SDN and NFV: A survey of taxonomy, architectures and future challenges, Computer Networks 167, pp. 1–40, 2020.

[30] A. A. Barakabitze, N. Barman, A. Ahmad, S. Zadtootaghaj, L. Sun, M. G. Martini, L. Atzori, QoE management of multimedia streaming services in future networks: A tutorial and survey, IEEE Communications Surveys & Tutorials 22(1), pp. 526–565, 2019.

[31] A. Barakabitze, A. Hines, QoE management of multimedia service challenges in 5G networks, pp. 167–183, 2023.

[32] E. Grigoriou, A. A. Barakabitze, L. Atzori, L. Sun, V. Pilloni, An SDN-approach for QoE management of multimedia services using resource allocation, in: IEEE International Conference on Communications (ICC), IEEE, 2017.

[33] A. Barakabitze, L. Sun, I. Mkwawa, E. Ifeachor, A novel QoE-aware SDN-enabled, NFV-based management architecture for future multimedia applications on 5G systems, 2019.

[34] M. Alsader, A. A. Barakabitze, I.-H. Mkwawa, QoE-driven adaptive video streaming: Architectures, techniques, and future research challenges toward 6G networks, IEEE Access 13, pp. 1–28, 2025.

[35] A. A. Barakabitze, R. Walshe, SDN and NFV for QoE-driven multimedia services delivery: The road towards 6G and beyond networks, Computer Networks 214, pp. 1–25, article 109133, 2022.

[36] I.-H. Mkwawa, A. A. Barakabitze, L. Sun, Video quality management over software-defined networking, in: IEEE International Symposium on Multimedia (ISM), San Jose, CA, USA, pp. 559–564, 2016.

[37] A. A. Barakabitze, QoE-centric management of multimedia streaming services in softwarized and virtualized networks, University of Plymouth, 2020.

[38] A. A. Barakabitze, I.-H. Mkwawa, A. Hines, R. Walshe, QoE-aware dynamic resource management in future softwarized and virtualized networks, IEEE Access 11, pp. 93310–93330, 2023.

[39] A. Barakabitze, A. Hines, QoE management of multimedia services using machine learning in SDN/NFV 5G networks, Wiley-IEEE Press, 2023.

[40] A. Barakabitze, A. Hines, Multimedia QoE-driven services delivery toward 6G and beyond network, Wiley-IEEE Press, 2023.

[41] A. Barakabitze, A. Hines, Management of multimedia services in emerging architectures using big data analytics: MEC, ICN, and fog/cloud computing, Wiley-IEEE Press, 2023.

[42] A. Barakabitze, A. Hines, Emerging applications and services in future 5G networks, Wiley-IEEE Press, 2023.

[43] A. A. Barakabitze, L. Sun, I.-H. Mkwawa, E. Ifeachor, Multipath protections and dynamic link recovery in softwarized 5G networks using segment routing, in: IEEE Global Communications Conference Workshops (GLOBECOM Workshops), Waikoloa, HI, USA, pp. 1–6, 2019.