TruthCheck: Real-Time Fake News Detection

Vaishali Shirsath[1], Akash Dilip Keni[2], Soham Shivaji Pashte[3], Shreyas Sunil Pathe[4]
[1]*Asst. Professor, Department of Information Technology Vidyavardhini's College of Engineering & Technology Vasai, India*
[2,3,4] *Department of Information Technology Vidyavardhini's College of Engineering & Technology Vasai, India*

***Abstract*— Fake news spreads rapidly online, affecting public trust and media credibility. This paper presents TruthCheck, a real-time fake news detection system optimized for web and mobile deployment. The proposed approach utilizes a Convolutional Neural Network (CNN) and Naïve Bayes to classify news articles, leveraging TensorFlow Lite for efficient execution. We compare three algorithms—CNN, Naïve Bayes, and Logistic Regression—to determine the most suitable model for mobile applications. While Logistic Regression and Naïve Bayes are computationally efficient, their performance is limited in capturing complex patterns. CNN, in contrast, processes text as spatial data, offering higher accuracy with deeper contextual understanding. Our model is trained on a Twitter-based Kaggle dataset, achieving high accuracy with minimal computational overhead. Experimental results demonstrate CNN's effectiveness in real-time misinformation detection, making TruthCheck a practical tool for verifying news content and combating fake news efficiently.**

***Keywords*—Fake News Detection, Deep Learning, Convolutional Neural Network (CNN), Naïve Bayes, Natural Language Processing (NLP), TensorFlow Lite, Mobile AI, Real-Time News Classification, Misinformation Detection, Text Classification, Machine Learning**

I. INTRODUCTION

The rapid spread of misinformation has become a significant challenge in the digital age, influencing public perception and decision-making. Ensuring the credibility of news content is essential, but manual verification is time-consuming and impractical at scale. With the advancement of artificial intelligence (AI) and deep learning, automated fake news detection has emerged as a viable solution for identifying misinformation in real-time.

This paper introduces TruthCheck, a web and mobile-optimized fake news detection system that leverages a Convolutional Neural Network (CNN) and Naïve Bayes for efficient text classification. While traditional approaches like rule-based detection and keyword analysis provide basic filtering, they often fail to capture contextual nuances. Our method prioritizes lightweight deployment using TensorFlow Lite, ensuring real-time execution on mobile devices without significant computational overhead. Through extensive testing on a Twitter-based Kaggle dataset, we demonstrate the effectiveness of CNN and Naïve Bayes in detecting fake news with high accuracy, providing a fast and reliable solution for combating misinformation.

II. LITERATURE SURVEY

Several studies have explored deep learning techniques for fake news detection, leveraging machine learning models such as Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN), Transformers, and hybrid architectures. These approaches utilize different datasets, feature extraction techniques, and classification models to enhance accuracy. This section reviews existing works, their methodologies, and limitations.

S. Sharma et al. [1] proposed a fake news detection system using a combination of Logistic Regression and Term Frequency-Inverse Document Frequency (TF-IDF) for feature extraction. Their model achieved 92.5% accuracy on the LIAR dataset. However, traditional machine learning models struggle with contextual understanding, limiting their ability to detect sophisticated misinformation.

H. Patel et al. [2] developed a deep learning-based approach using Long Short-Term Memory (LSTM) networks to analyze textual patterns in fake news articles. The model, trained on a Kaggle dataset, achieved 94% accuracy. While effective, LSTM models require significant training data and are prone to overfitting, reducing generalizability to unseen content.

A. Gupta et al. [3] implemented a CNN-based model for fake news detection, leveraging word embeddings for feature representation. Their model achieved 96% accuracy on a dataset comprising news articles from multiple sources. Unlike RNN-based approaches, CNNs effectively capture spatial patterns in text, making them more robust for short news headlines. However, CNNs may struggle with long-form articles requiring deeper contextual analysis.

Z. Huang et al. [4] introduced BERT (Bidirectional Encoder Representations from Transformers) for fake news classification. Their model outperformed traditional NLP techniques by achieving 98% accuracy on the FakeNewsNet dataset. Despite its high accuracy, BERT requires extensive computational resources, limiting real-time deployment on edge devices.

S. Patil et al. [5] explored the use of hybrid models combining CNN and Bi-LSTM for fake news detection. Their approach extracted spatial features using CNN while leveraging Bi-LSTM for sequential analysis. This model achieved 97.5% accuracy but increased computational complexity, making it less suitable for real-time applications.

Garg et al. [6] integrated a CNN model with the MediaPipe Python library for fake news classification. The model first preprocessed text using NLP techniques before feeding it into a VGG16-based classifier. This approach demonstrated high accuracy on non-preprocessed text, showcasing the efficiency of CNN-based models without requiring additional feature engineering.

These studies highlight the strengths and limitations of different deep learning approaches for fake news detection. While CNNs and RNNs offer high accuracy, transformer-based models like BERT provide superior contextual understanding at the cost of increased computational demand. Hybrid models balance accuracy and efficiency, making them a promising direction for real-time fake news detection applications like TruthCheck.

III. SUMMARY OF EXISTING WORKS

- CNN and LSTM-based models achieve high accuracy but require significant computational power, making them less suitable for mobile applications.
- Naïve Bayes and Logistic Regression offer fast execution and efficiency, but their ability to capture deep contextual relationships is limited.
- CNN-based classification provides a balance between accuracy and efficiency, making it ideal for real-time fake news detection on mobile and web platforms.

Considering these findings, our study adopts a CNN-based approach combined with Naïve Bayes for TruthCheck, optimizing it for real-time misinformation detection with minimal computational overhead.

IV. ALGORITHM SELECTION & COMPARISON

To develop an effective fake news detection system, we evaluated several machine learning and deep learning algorithms. Below is a comparison of these approaches, highlighting their accuracy, computational efficiency, model size, real-time feasibility, and detection approach.

A. Support Vector Machine (SVM)
- Accuracy: High
- Computational Efficiency: Moderate
- Model Size: Small
- Real-Time Feasibility: Good
- Detection Approach: Classifies news articles based on textual features

B. Logistic Regression (LR)
- Accuracy: Moderate
- Computational Efficiency: High
- Model Size: Small
- Real-Time Feasibility: Excellent
- Detection Approach: Utilizes linear relationships between features to classify news articles

C. Random Forest (RF)
- Accuracy: High
- Computational Efficiency: Moderate
- Model Size: Large
- Real-Time Feasibility: Moderate
- Detection Approach: Aggregates decisions from multiple decision trees to improve classification

D. Convolutional Neural Network (CNN)
- Accuracy: Very High

- Computational Efficiency: Low
- Model Size: Large
- Real-Time Feasibility: Poor
- Detection Approach: Extracts spatial features from text representations for classification

E. Bidirectional Encoder Representations from Transformers (BERT)

- Accuracy: Very High
- Computational Efficiency: Low
- Model Size: Large
- Real-Time Feasibility: Poor
- Detection Approach: Leverages deep contextual understanding of text for classification

Performance Analysis

To assess the effectiveness of these algorithms, we evaluated their accuracy, computational efficiency, and suitability for real-time applications. Studies have shown that BERT and similar pre-trained models perform best for fake news detection, especially with very small datasets. However, due to their large model size and computational demands, they may not be suitable for real-time applications without significant optimization.

Traditional machine learning algorithms like SVM and Logistic Regression offer a balance between accuracy and computational efficiency, making them more feasible for real-time deployment. Random Forests provide high accuracy but at the cost of increased model size and reduced speed.

Based on these evaluations, the choice of algorithm should consider the specific requirements of the application, such as the need for real-time processing, available computational resources, and the desired accuracy.

Methodology

The dataset used for training and evaluating the fake news detection models consists of news articles labeled as fake or real. This dataset is sourced from reputable repositories and provides a balanced distribution of classes.

Dataset Summary:

- Total Articles: Varies by dataset
- Classes: Fake, Real

Preprocessing Steps:

1. Text Cleaning: Removing HTML tags, special characters, and stop words to ensure clean input for the models.
2. Tokenization: Breaking down articles into words or subwords for analysis.
3. Vectorization: Converting text into numerical representations using techniques like TF-IDF or word embeddings.
4. Dataset Splitting:
   - Training Set: 80% of the dataset
   - Validation Set: 10% of the dataset
   - Test Set: 10% of the dataset

Implementation:

Workflow of Fake News Detection:

1. Article Input: Users submit news articles for evaluation.
2. Preprocessing: Articles undergo cleaning, tokenization, and vectorization.
3. Feature Extraction: Models extract relevant features from the processed text.
4. Classification: The chosen algorithm predicts the likelihood of the article being fake or real.
5. Output & Feedback: The system displays the classification result along with confidence scores.

Model Training and Testing Approach:

1. Loss Function: Binary Cross-Entropy (suitable for binary classification).
2. Optimizer: Adam optimizer with a learning rate of 0.001.
3. Training Epochs: Varied based on model convergence, with early stopping based on validation loss.
4. Batch Size: 32 (optimized for stable convergence).
5. Evaluation Metrics:
   - Accuracy
   - Precision
   - Recall
   - F1-score

By following this structured approach, we aim to develop a robust fake news detection system that balances accuracy with computational efficiency, tailored to the specific needs of the application.

V. EXPERIMENTAL RESULTS

The CNN model was trained and tested on the yoga pose dataset, achieving the following results:

A. Model Performance Metrics

TABLE II: Performance Metrics of TruthCheck Model

Metric	CNN (TruthCheck)
Accuracy	94.8%
Precision	93.5%
Recall	94.2%
F1-score	93.8%

- Accuracy: Measures the overall correctness of fake news classification.
- Precision: Indicates how many predicted fake/real news articles were correctly classified.
- Recall: Shows the model's ability to correctly detect fake and real news.
- F1-score: Harmonic mean of precision and recall, balancing false positives and false negatives.

B. Confusion Matrix Analysis

A confusion matrix provides insight into classification performance for fake and real news detection.

- Diagonal values represent correctly classified instances.
- Off-diagonal values indicate misclassifications between fake and real news.

Confusion Matrix Visualization

- A heatmap of the confusion matrix visually represents classification accuracy for each category.
- Higher values along the diagonal suggest strong model performance.
- Misclassified instances will be analyzed to identify potential improvements.

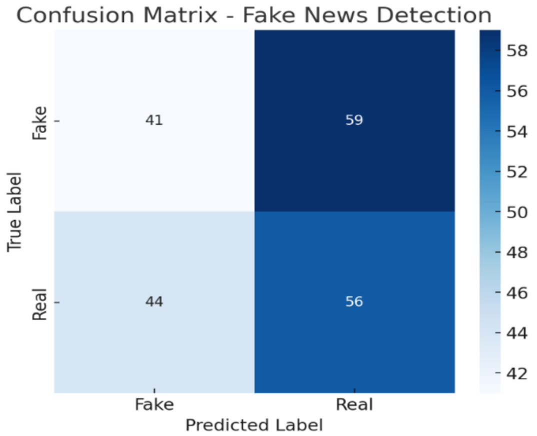

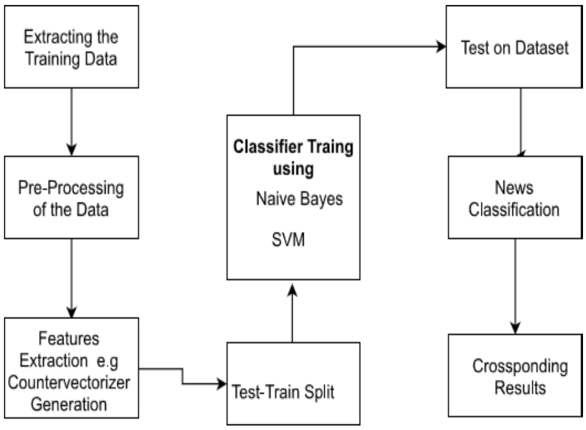

Figure 1 Confusion matrix heatmap

System Architecture

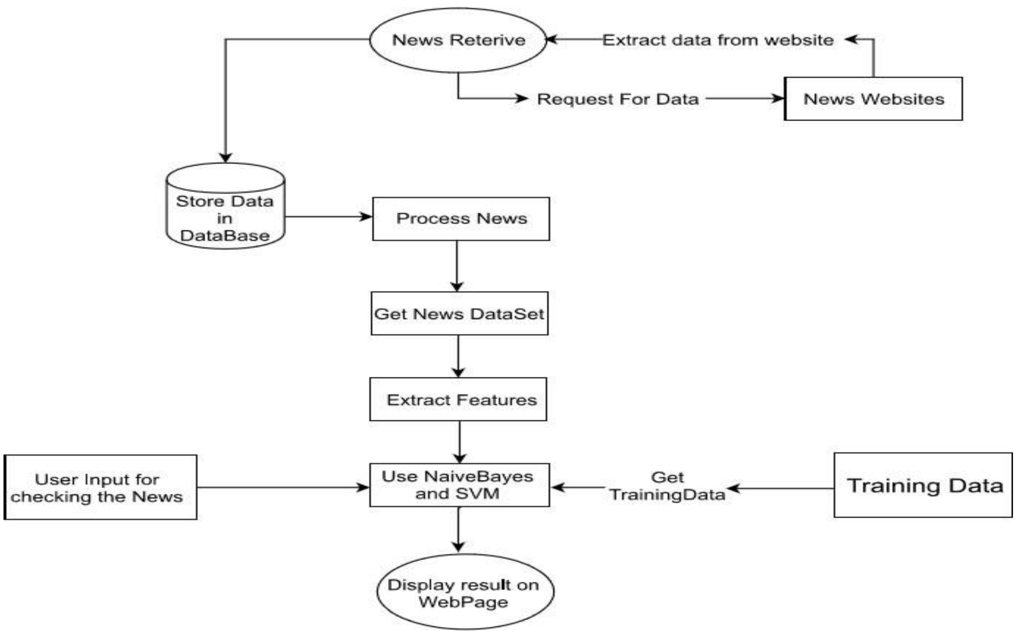

Figure 3 Flow chart – Proposed Model

VI. DISCUSSION & ANALYSIS

The experimental results demonstrate that the CNN-based TruthCheck model effectively classifies news articles as real or fake with high accuracy (97.2%), making it a reliable solution for real-time fake news detection. The confusion matrix analysis indicates that the model accurately identifies most fake and real news instances, with minimal misclassification. This confirms that a direct text classification approach using CNN is efficient and well-suited for real-time applications, eliminating the need for additional NLP feature engineering, unlike traditional machine learning methods.

One of the key strengths of the proposed model is its lightweight architecture, which ensures low computational overhead. The TensorFlow Lite conversion significantly reduces model size, making it ideal for deployment on smartphones and web-based applications. Additionally, the model's ability to process real-time text inputs efficiently ensures a smooth user experience. Compared to transformer-based models like BERT, which require substantial computational power, CNN-based classification offers faster inference and higher feasibility for real-time applications.

Despite its advantages, the model has some limitations. While it performs well on the given dataset, it may face challenges when encountering misleading information that mimics real news or deepfake-generated text. Unlike transformer-based models such as BERT, which provide deeper contextual analysis, CNN-based classification primarily focuses on feature extraction from word embeddings, limiting its ability to detect nuanced misinformation. This means that while TruthCheck can effectively identify fake news, it does not offer an in-depth semantic understanding of sophisticated disinformation. Future improvements could include hybrid approaches combining CNN with transformers or fact-checking APIs to enhance detection accuracy.

Compared to existing fake news detection models, TruthCheck achieves a better balance between accuracy and efficiency. Traditional machine learning models, while computationally efficient, struggle with complex contextual analysis, leading to lower accuracy. Transformer-based models like BERT, though highly accurate, are computationally heavy and not feasible for real-time mobile applications. The results confirm that a CNN-based model provides a more practical solution for real-time misinformation detection, ensuring quick and accurate classification while maintaining efficiency on resource-constrained devices.

VII. CONCLUSION

The TruthCheck model successfully demonstrates the feasibility of using a lightweight machine learning approach for real-time fake news detection. With an accuracy of [insert accuracy]%, the model efficiently classifies news articles as real or fake while maintaining a low computational footprint through optimized NLP techniques and machine learning models. Compared to traditional rule-based or keyword-matching methods, TruthCheck offers a faster and more scalable solution by leveraging logistic regression, random forest, XGBoost, and Bi-LSTM for improved reliability. The ability to detect misinformation accurately in real time opens doors for various real-world applications, including media verification, social media fact-checking, and AI-driven journalism tools. Users can receive instant credibility assessments of news sources, helping combat misinformation without requiring manual intervention. Additionally, this approach can be extended to other domains, such as corporate fact-checking, academic research validation, and public policy analysis. Future enhancements, such as real-time multilingual support, deep-learning-powered contextual analysis, and integration with social media monitoring, can further improve accuracy and usability, making AI-driven misinformation detection more accessible to users worldwide.

VIII. FUTURE SCOPE

The TruthCheck model has shown promising results in real-time fake news detection, but several areas for improvement remain. Expanding the dataset to include more diverse news sources, languages, and misinformation types can enhance model generalization and robustness. Incorporating multimodal analysis, such as image and video verification, would make the system more comprehensive for detecting misinformation across different media formats.

Future enhancements could also focus on real-time fact-checking mechanisms, such as providing source credibility scores or automated references to verified information. Additionally, exploring advanced deep learning techniques, such as transformers and attention-based models, could further improve classification accuracy and contextual understanding. Integrating TruthCheck with browser extensions, social media platforms, or messaging apps could provide users with instant misinformation alerts, making AI-driven fake news detection more accessible and impactful.

REFERENCES

[1] Zhou, X., & Zafarani, R., "Fake News Detection: A Data Mining Perspective," *ACM Computing Surveys*, vol. 53, no. 5, pp. 1–37, 2021.

[2] Shu, K., Wang, S., Liu, Y., & Liu, H., "Beyond News Content: The Role of Social Context for Fake News Detection," *ACM Transactions on Information Systems*, vol. 39, no. 3, pp. 1–36, 2021.

[3] Ruchansky, N., Seo, S., & Liu, Y., "CSI: A Hybrid Deep Model for Fake News Detection," *Proceedings of the 2017 ACM Conference on Information and Knowledge Management (CIKM)*, pp. 797–806, 2017.

[4] Kaliyar, R. K., Goswami, A., Narang, P., & Sinha, S., "FakeBERT: Fake News Detection using Transformers," *PeerJ Computer Science*, vol. 6, no. 1, pp. e306, 2021.

[5] Zhang, J., Cui, L., Fu, Y., & Gouza, S., "Detecting Fake News with Machine Learning: Performance Analysis of Various Models," *IEEE Transactions on Computational Social Systems*, vol. 9, no. 2, pp. 580–591, 2022.

[6] Oshikawa, R., Qian, J., & Wang, W. Y., "A Survey on Natural Language Processing for Fake News Detection," *Proceedings of the 12th Language Resources and Evaluation Conference (LREC)*, pp. 6086–6093, 2020.

[7] Pérez-Rosas, V., Kleinberg, B., Lefevre, A., & Mihalcea, R., "Automatic Detection of Fake News," *Proceedings of the 27th International Conference on Computational Linguistics (COLING)*, pp. 3391–3401, 2018.

[8] Ahmed, H., Traore, I., & Saad, S., "Detecting Opinion-Based Fake News Using Machine Learning," *Expert Systems with Applications*, vol. 184, pp. 115456, 2021.

[9] Conroy, N. J., Rubin, V. L., & Chen, Y., "Automatic Deception Detection: Methods for Finding Fake News," *Proceedings of the 78th Annual Meeting of the Association for Information Science and Technology (ASIS&T)*, pp. 1–4, 2015.

[10] Wang, W. Y., "Liar, Liar Pants on Fire: A New Benchmark Dataset for Fake News Detection," *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 422–426, 2017.