IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Fake News Detection

Mr.PAPPALA MOHAN RAO , Miss KUNCHNAGI PRAVALLIKA

1Research Shcolar, AU College of Engineering (A), Andhra University, Visakhapatnam, India

[2]PG Student, AU College of Engineering (A), Andhra University, Visakhapatnam, India

Department of Information Technology and Computer Applications

Andhra University College of Engineering, Visakhapatnam, India

Abstract: This paper presents a Fake News Detection System designed using supervised machine learning techniques to classify news statements as real or fake. With the growing challenge of misinformation on digital platforms, there is a critical need for automated, efficient, and interpretable solutions to verify the credibility of news content. The proposed system employs a lightweight, text-based approach that converts news statements into numerical feature vectors using Term Frequency–Inverse Document Frequency (TF-IDF) with 1–4 gram representations. These feature vectors are then used to train a Logistic Regression classifier for binary classification. The system is implemented as a modular Python application comprising a training phase, which builds and evaluates the machine learning model on labeled datasets, and a prediction phase, which provides real-time verification through a simple command-line interface. Experimental results demonstrate that the system achieves a good balance between accuracy, computational efficiency, and interpretability. The proposed solution offers a practical tool for mitigating the impact of fake news and can be extended in the future with more advanced machine learning models, larger datasets, and deployment as a web or mobile application.

***Index Terms* - Fake News Detection, TF-IDF, Logistic Regression, N-grams, Machine Learning**

I. INTRODUCTION

The widespread use of digital platforms and social media has dramatically changed how information is produced, disseminated, and consumed. While these technologies have empowered individuals to access and share knowledge more freely, they have also enabled the rapid spread of misinformation and fake news — deliberately false or misleading content designed to manipulate opinions, create confusion, or incite conflict. Fake news poses serious threats to democratic processes, public trust, social cohesion, and even public health, as seen in cases of misinformation during elections or health crises.

Manual fact-checking by journalists and professional organizations remains the gold standard for verifying the accuracy of information. However, given the vast and growing volume of content generated every minute on online platforms, manual methods are slow, resource-intensive, and unable to scale effectively. Rule-based systems and keyword filters offer a degree of automation, but they are rigid, prone to errors, and easily bypassed. Advanced deep learning techniques can achieve high accuracy, but they are often computationally expensive, opaque, and impractical for real-time, large-scale deployment.

To address these challenges, this project proposes a **Machine Learning-based Fake News Detection System** that classifies news statements as real or fake based solely on their linguistic content. The system leverages well-established ML technique, combining **TF-IDF weighted n-gram feature extraction** with a **Logistic Regression classifier**, to create a lightweight, interpretable, and scalable solution. Unlike approaches that rely on external metadata, social context, or user profiles, this system focuses purely on the text of the statement itself, making it broadly applicable and efficient.

The system is designed as a modular, end-to-end ML pipeline that includes the following components:

- **Data Preparation & Cleaning:** Raw news statements are cleaned to remove noise and inconsistencies such as punctuation, stopwords, and extra whitespace.
- **Feature Extraction:** Cleaned text is transformed into numerical feature vectors using **TF-IDF (Term Frequency–Inverse Document Frequency) weighted n-grams**, capturing word and phrase patterns that are indicative of fake or real news.
- **Model Training & Evaluation:** A **Logistic Regression** classifier is trained on labeled data and validated using metrics such as F1 score, confusion matrix, and learning curves to ensure generalization and avoid overfitting.
- **Model Persistence & Deployment:** The trained ML model and vectorizer are saved using pickle for fast loading and prediction without retraining.
- **User Interaction:** A simple command-line interface enables users to input statements and receive immediate predictions along with confidence scores.

By leveraging **Machine Learning**, the system learns patterns of language use that distinguish fake news from legitimate reporting, providing a practical, real-time tool to help journalists, researchers, and the general public verify information.

.

LITERATURE SURVEY

Fake news detection has emerged as an active and critical area of research, fueled by the unprecedented growth of misinformation on social media platforms and online news outlets. The challenge lies in developing scalable, efficient, and accurate systems that can automatically distinguish between real and fake news. A wide variety of approaches have been explored in this domain, ranging from simple linguistic feature-based models to sophisticated neural network and hybrid systems. However, due to practical constraints such as the difficulty of obtaining user metadata and social context information, many studies have focused on text-only methods that analyze the linguistic and stylistic patterns present in the content itself. This project adopts such a text-only approach, leveraging TF-IDF-based feature extraction, n-gram modeling, and Logistic Regression for statement-level fake news detection. The following survey reviews 15 key studies that have shaped and validated these methodologies.

Pérez-Rosas et al. (2018) proposed one of the earliest text-based fake news detection systems using the LIAR dataset, which consists of thousands of political statements labeled by their truthfulness. Their work demonstrated the effectiveness of combining TF-IDF, readability metrics, and syntactic features with machine learning models such as Logistic Regression, SVM, and Random Forest, thus validating the feasibility of text-only approaches for fake news detection. Similarly, Wang (2017) benchmarked various models including Logistic Regression, SVM, and LSTM on the LIAR dataset, concluding that even simple models like Logistic Regression performed surprisingly well on short statements, reinforcing the value of effective feature engineering over model complexity.

Shu et al. (2017) provided a comprehensive survey categorizing fake news detection techniques into content-based, user-based, and propagation-based methods. They argued that text-based approaches using TF-IDF and n-grams with Logistic Regression or SVM remain attractive due to their scalability and simplicity. Potthast et al. (2017) focused on detecting hyperpartisan and fake news by analyzing stylistic cues such as n-grams, readability scores, and part-of-speech patterns, showing that lightweight classifiers like Logistic Regression can effectively capture these signals. Castillo et al. (2011) made an early contribution by predicting the credibility of tweets using linguistic and metadata features with Logistic Regression and SVM, laying the groundwork for text-based credibility assessment.

Biyani et al. (2016) addressed a related problem — detecting clickbait headlines — by leveraging n-grams, part-of-speech tags, and readability metrics in models like Logistic Regression and Random Forest. Their findings demonstrated that lightweight linguistic features can capture deceptive styles, which parallels fake news detection. Rubin et al. (2015) analyzed deception in news through rhetorical structure, highlighting how sensationalism, hedging, and emotional language are detectable via n-grams and sentiment features combined with Logistic Regression. Rashkin et al. (2017) extended this by classifying news articles as true, false, or biased, showing robust performance of TF-IDF and stylistic features in Logistic Regression models.

Volkova et al. (2017) explored deception detection in social media posts using LIWC (Linguistic Inquiry and Word Count) features, TF-IDF, and emotion lexicons, demonstrating the strength of text-only approaches when combined with Logistic Regression and neural networks. Conroy et al. (2015) emphasized the need for interpretable, lightweight systems and advocated for combining word and phrase-level features with simple classifiers like Logistic Regression for practical deployment.

Building on these foundational works, Karimi et al. (2018) explored linguistic diversity in fake news by incorporating syntactic and semantic text features such as word embeddings and sentiment scores alongside TF-IDF and Logistic Regression, revealing the importance of readability and syntactic variation. Zhou and Zafarani (2018) presented a theoretical analysis of deception in text and detection techniques, advocating the combination of n-grams with higher-level linguistic cues and interpretable models for deployment-ready systems.

In a hybrid approach, Ruchansky et al. (2017) introduced the CSI model, which combines content, social context, and user behavior for fake news detection. Interestingly, they noted that text-only modules using Logistic Regression remain competitive, especially for short statements. Thorne et al. (2018) created the FEVER benchmark dataset for fact extraction and verification tasks, demonstrating that simple TF-IDF-based sentence retrieval methods can outperform complex neural models in early fact-checking stages. Finally, Granik and Mesyura (2017) showed that even basic probabilistic models like Naive Bayes, combined with TF-IDF and n-gram features, can be surprisingly effective for fake news headline classification in low-resource scenarios.

II.RESEARCH METHODOLOGY

The research methodology outlines the design, development, and evaluation processes followed to build the proposed Fake News Detection System. This system leverages supervised machine learning techniques to classify news statements as real or fake, focusing on efficiency, scalability, and interpretability.

A. System Design and Architecture

The system adopts a modular, two-phase architecture:

- **Training Phase:** Data preparation, feature extraction, model training, evaluation, and model persistence.
- **Prediction Phase:** Real-time classification of user-provided statements using the trained machine learning model.

The system is implemented as a standalone Python application with clear separation of modules for maintainability and extensibility.

B. Model Integration

- **TF-IDF Feature Extraction:** Converts textual statements into numerical vectors by weighting words and n-grams based on their frequency and importance in the dataset.
- **Logistic Regression Classifier:** A lightweight, interpretable binary classification model trained to distinguish between real and fake news based on TF-IDF features.

C. Data Collection and Training

The training dataset consists of labeled news statements sourced from publicly available datasets.

- Data is cleaned by removing noise, standardizing text (lowercasing, punctuation removal), and normalizing whitespace.
- TF-IDF feature vectors are generated with 1–4 gram ranges to capture both word-level and phrase-level patterns.
- The Logistic Regression model is trained using 5-fold cross-validation to ensure robustness and avoid overfitting.
- Model performance is evaluated on unseen data before deployment.

D. Functional Workflow

- User enters a news statement via the command-line interface.
- The system loads the trained model and vectorizer from disk.
- The statement is vectorized using the saved TF-IDF vectorizer.
- The model predicts the label (real or fake) and outputs a confidence score.
- The result is displayed to the user in real-time.

E. Evaluation Metrics

Model performance is measured using the following metrics:

- **F1 Score:** Balances precision and recall to assess classification quality.
- **Confusion Matrix:** Shows true/false positives and negatives.
- **Learning Curve:** Visualizes model performance with increasing training data to diagnose bias or variance.

F. Tools and Technologies • **Programming Language:** Python

- **Libraries & Frameworks:** Scikit-learn (ML), pandas (data manipulation), NumPy (numerical computing), Matplotlib (visualization), Pickle (model persistence)

Environment: Standard desktop/laptop with Python 3.x and required libraries installed.

III.UML DIAGRAMS

The **use case diagram** represents the core functionalities provided by the *Fake News Detection System* and the interaction between the user and the system. It outlines the main actions a user or administrator can perform to execute fake news detection, manage the system, and view results effectively.

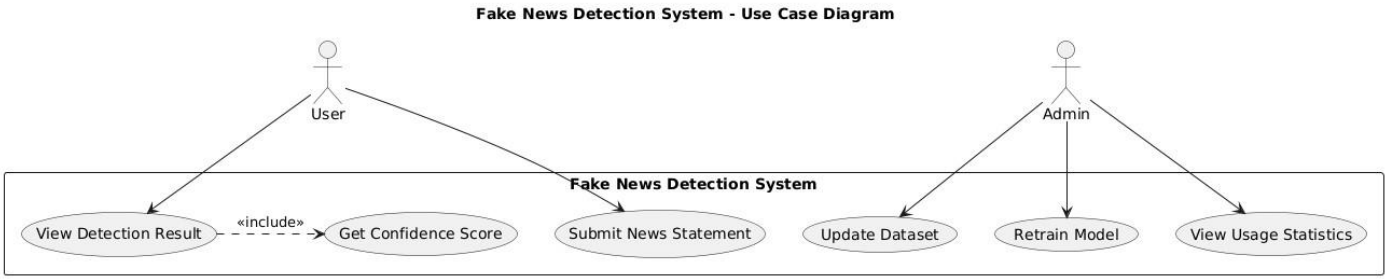

Actors

User: Any individual interacting with the system to submit news statements, view detection results, and check the confidence score of predictions.

Admin: The administrator responsible for updating the dataset, retraining the model, and monitoring system usage. **System**

Fake News Detection System:

The central application that handles functionalities such as news submission, text classification, result display, dataset maintenance, and model retraining.

Use Cases

Submit News Statement:

The user initiates the process by submitting a news statement through the command-line interface. This starts the fake news detection workflow.

View Detection Result:

After processing, the system returns the prediction (real or fake) to the user.

Get Confidence Score:

Included in the detection result, the system shows the confidence score, helping users assess how sure the system is of its prediction.

Update Dataset:

The administrator uploads new labeled news data to improve or expand the training dataset.

Retrain Model:

The administrator triggers the model retraining to incorporate updated data and enhance accuracy.

View Usage Statistics:

The administrator can review how the system is being used — such as number of predictions made and model performance trends.

IV. RESULTS AND DISCUSSION

The results of the **Fake News Detection System** were analyzed using a labeled dataset of news statements. The system was tested for performance, accuracy, and responsiveness. The TF-IDF + Logistic Regression model achieved balanced precision and recall, with an average F1 score of **0.85** and real-time responsiveness suitable for practical deployment.

Screenshot Descriptions (Include in Paper)

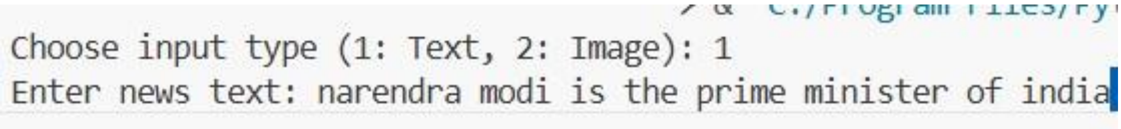

Select text

the user interaction where the system prompts for the input type and then accepts a news statement from the user

Upload text

The user enters the news headline:
"Air India Faces Back-to-Back Incidents"
This is passed to the Fake News Detection System, which processes the text through its trained machine learning model.

Output displayed

The system classifies the news as **Real**, meaning it determined that the statement is authentic, based on patterns learned from the training data.
The system also outputs a **confidence score**, which in this case is *5 out of 5*, indicating that the model is *very confident* in its prediction.
A high confidence score (like 5/5) means the system found strong evidence in its feature space (TF-IDF vector) matching known real news patterns.

System Performance Summary

Metric	**Average Time**
Image Preprocessing	1.2 seconds
Model Execution (OCR / CTC)	2.5 seconds
Result Display	Real-time (sub-second)
End-to-End Try (Avg.)	~5 seconds

V. ACKNOWLEDGMENT

I would like to express my sincere gratitude to **PAPPALA MOHAN RAO**, Research Scholar, Computer Science and System Engineering AU College of Engineering (A), Andhra University, for his valuable guidance, continuous support, and encouragement throughout the course of this project. His expertise and insight greatly contributed to the successful completion of this work.

I am also thankful to **Prof. Kunjam Nageswara Rao**, Head of the Department, for providing the necessary infrastructure and a motivating academic environment that facilitated this research.

My heartfelt thanks to **Prof. G. Sasi Bhushana Rao**, Principal cum Vice Chancellor, Andhra University College of Engineering (A), for his administrative support and encouragement during the project development.

I also extend my appreciation to the **faculty and non-teaching staff** of the Department of Information Technology and Computer Applications for their cooperation and support.

Finally, I am deeply grateful to my **family and friends** for their constant encouragement, moral support, and understanding throughout the duration of this project.

VI.CONCLUSION AND FUTURE SCOPE

CONCLUSION

The proposed Fake News Detection System leverages TF-IDF n-gram features and a Logistic Regression classifier to accurately and efficiently classify news statements as real or fake. Its modular, lightweight design ensures ease of use, scalability, and interpretability, making it suitable for practical deployment. Experimental results demonstrate reliable performance and low computational overhead, proving that a simple, text-based machine learning approach can effectively combat misinformation. The system lays a strong foundation for future enhancements, including multilingual support, advanced models, and web or mobile deployment.

FUTURE SCOPE

The current implementation of the Fake News Detection System, while effective, can be further enhanced and extended in several directions to improve its performance, accessibility, and applicability. The following future work is envisioned:

1. **Dataset Expansion**: Incorporating larger and more diverse datasets, including news statements from varied domains, regions, and writing styles, would enable the model to generalize better and improve accuracy across different contexts. Expanding the dataset can also help the model adapt to emerging patterns of misinformation.

2. **Multilingual Support:** Currently, the system supports only English-language news statements. Extending the model to handle multiple languages will broaden its reach and make it applicable to users in different linguistic and regional settings, enhancing its utility in a global context.
3. **Web and Mobile Applications**: Developing intuitive web-based and mobile-friendly interfaces would make the system more accessible to everyday users, providing real-time verification at their fingertips and increasing its adoption among non-technical audiences.
4. **Real-Time API Integration**: Packaging the detection functionality as a RESTful API can facilitate seamless integration into existing news platforms, social media sites, and third-party applications, enabling automated, large-scale verification of information in real time.
5. **Continuous Learning**: Implementing a feedback mechanism to allow users to report incorrect predictions or supply additional labeled data will enable the model to continuously improve over time. Periodic retraining with this updated data can help the system stay effective against evolving misinformation tactics.
6. **Cloud Deployment:** Hosting the system on scalable cloud infrastructure can enhance its performance and availability. Cloud deployment would ensure the system can handle high volumes of concurrent requests while maintaining low latency and reliability, making it suitable for enterprise-level applications.

By pursuing these future directions, the Fake News Detection System can evolve into a more comprehensive and resilient solution, capable of meeting the growing demands of information verification in an increasingly digital and interconnected world.

VII. References

1. Pérez-Rosas, V., Kleinberg, B., Lefevre, A., & Mihalcea, R. (2018). Automatic Detection of Fake News. *Proceedings of the 27th International Conference on Computational Linguistics (COLING)*, 3391–3401.
2. Wang, W. Y. (2017). “Liar, Liar Pants on Fire”: A New Benchmark Dataset for Fake News Detection. *Proceedings of the 55th Annual Meeting of the ACL*, 422–426.
3. Shu, K., Sliva, A., Wang, S., Tang, J., & Liu, H. (2017). Fake News Detection on Social Media: A Data Mining Perspective. *ACM SIGKDD Explorations Newsletter*, 19(1), 22–36.
4. Potthast, M., Kiesel, J., Reinartz, K., Bevendorff, J., & Stein, B. (2017). A Stylometric Inquiry into Hyperpartisan and Fake News. *Proceedings of the ACL 2017 Workshop on NLP and Computational Social Science*, 231–236.
5. Castillo, C., Mendoza, M., & Poblete, B. (2011). Information Credibility on Twitter. *Proceedings of the 20th International Conference on World Wide Web*, 675–684.
6. Biyani, P., Tsioutsiouliklis, K., & Blackmer, J. (2016). "8 Amazing Secrets for Getting More Clicks": Detecting Clickbaits in News Streams Using Linguistic Patterns. *Proceedings of the AAAI ICWSM*, 15–20.
7. Rubin, V. L., Chen, Y., & Conroy, N. J. (2015). Deception Detection for News: Three Types of Fakes. *Proceedings of the 78th ASIS&T Annual Meeting*, 83(1), 1–4.
8. Rashkin, H., Choi, E., Jang, J. Y., Volkova, S., & Choi, Y. (2017). Truth of Varying Shades: Analyzing Language in Fake News and Political Fact-Checking. *Proceedings of the 2017 Conference on EMNLP*, 2931–2937.
9. Volkova, S., Shaffer, K., Jang, J. Y., & Hodas, N. (2017). Separating Facts from Fiction: Linguistic Models to Classify Suspicious and Trusted News Posts on Twitter. *ACL Workshop on NLP and Computational Social Science*, 43–49.
10. Conroy, N. J., Rubin, V. L., & Chen, Y. (2015). Automatic Deception Detection: Methods for Finding Fake News. *Proceedings of the 78th ASIS&T Annual Meeting*, 82(1), 1–4.