IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.021 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 13, Issue 4, April 2025

DOI: 10.17148/IJIREEICE.2025.13405

AI-Based Fake News Detection System

Mr.Arun Prasanth S[1], Dr. K. Banuroopa, MCA., M.Phil., Ph.D.[2]

Student, Department of Information Technology, Dr. N. G. P. Arts and Science College,

Coimbatore, Tamil Nadu, India[1]

Professor, Department of Information Technology, Dr. N. G. P. Arts and Science College,

Coimbatore, Tamil Nadu, India[2]

Abstract: The rapid spread of false information in the digital age poses a serious threat to society, influencing how people think and make decisions. As a result, identifying fake news has become essential to ensuring the reliability of information found online. Traditional fact-checking methods often rely on slow, labor-intensive manual processes that are increasingly ineffective given the volume and speed of misinformation. This has led to growing interest in machine learning-based solutions for automating fake news detection. In this study, we propose a fake news classification model that uses Logistic Regression for classification and TF-IDF (Term Frequency-Inverse Document Frequency) vectorization for feature extraction, helping to distinguish between real and fake news articles more efficiently.However, many existing fake news detection systems face significant challenges. Traditional models often struggle to adapt to evolving misinformation patterns, leading to outdated or inaccurate results. Additionally, class imbalances in datasets — where one type of news (real or fake) heavily outweighs the other — can create biased predictions. Feature extraction techniques commonly used in older models also fail to capture the deeper, semantic meaning of text, resulting in subpar classification performance. Moreover, many models lack the ability to generalize across diverse datasets, which limits their effectiveness in real-world applications. These challenges highlight the need for a more reliable and adaptable system for fake news detection.

To address these issues, we propose an enhanced machine learning-based detection system. Our approach incorporates Logistic Regression alongside TF-IDF vectorization for effective feature extraction. We also introduce stratified train-test splitting to maintain class distribution during training and use RandomOverSampler to combat class imbalances by generating synthetic samples for underrepresented classes. To thoroughly evaluate performance, we measure accuracy, precision, recall, and visualize results using a confusion matrix, providing a clearer picture of how well the model performs.In addition to these core techniques, our system introduces several novel improvements. We implement automatic dataset validation to identify and handle missing or imbalanced labels, ensuring data is ready for training without manual intervention. If one class is significantly underrepresented or missing altogether, our model performs class augmentation, generating synthetic data to restore balance. We also introduce an interactive user prediction feature, allowing users to input custom news articles for real-time classification. This interactive component enhances the model's practicality, making it a valuable tool for everyday use. These improvements collectively enhance model reliability, resulting in a more robust, accurate, and adaptable fake news detection system capable of keeping up with the ever-changing landscape of misinformation.

INTRODUCTION

The rapid spread of false information in today's digital world is a serious threat to society, influencing how people think and make decisions. Being able to identify fake news has become essential to ensuring the credibility of information online. Traditional methods of verifying news often rely on time-consuming, manual fact-checking, which can be inefficient. To address this, machine learning has emerged as a promising approach for automating fake news detection.

This study presents a fake news classification model that uses logistic regression combined with TF-IDF (Term Frequency-Inverse Document Frequency) vectorization to differentiate between real and fake news articles.

However, existing fake news detection systems face several challenges. Many traditional models struggle to adapt to evolving misinformation patterns and handle class imbalances in datasets, which often leads to biased predictions. Additionally, conventional feature extraction methods fail to capture the deeper meaning of the text, resulting in less accurate classifications. Another major issue is the lack of generalization across different datasets, making these models less effective when applied to real-world scenarios. These limitations highlight the need for a more reliable and adaptable fake news detection system.

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering
Impact Factor 8.021 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 13, Issue 4, April 2025
DOI: 10.17148/IJIREEICE.2025.13405

Our research proposes an improved machine learning-based system that tackles these issues. It leverages Logistic Regression for classification, TF-IDF for feature extraction, stratified train-test splitting to preserve class distribution during training, and RandomOverSampler to handle class imbalances. To evaluate the model's performance, we use accuracy, precision, recall, and confusion matrix visualizations.

1. Introduction: Why Fake News Detection Matters

In today's digital world, information travels faster than ever — and unfortunately, so does misinformation. Fake news has become a major issue, influencing public opinion, creating confusion, and even fueling social and political conflicts. Traditional methods of fact-checking just can't keep up with the speed and scale of misinformation. That's where Artificial Intelligence (AI) steps in, offering a faster, more scalable solution. Let's dive into how AI is tackling fake news and what researchers are doing to improve these systems.

2. Understanding Fake News and Its Impact

Fake news isn't just about made-up stories — it's a mix of half-truths, misleading content, and fabricated information designed to look real. It can range from clickbait headlines to deepfakes. The impact? People believe false information, trust in legitimate media weakens, and harmful narratives spread like wildfire. AI-based systems aim to stop this spread before it causes real damage.

3. How AI Detects Fake News

3.1 Machine Learning: Teaching Computers to Spot Lies

- Machine learning (ML) trains computers to recognize patterns — like how we learn from experience. By feeding large amounts of labeled data into algorithms, AI learns what fake and real news looks like.
- Supervised Learning: This method uses labeled data (e.g., "this is fake, this is real") to help the system learn. Common algorithms include Support Vector Machines (SVM) and Random Forest.
- Unsupervised Learning: When data isn't labeled, AI looks for unusual patterns or clusters that might indicate misinformation.
- Semi-Supervised Learning: A mix of the two, combining a small amount of labeled data with a large set of unlabeled content.

3.2 Natural Language Processing (NLP): Understanding the Words

- AI doesn't just look at words — it tries to understand them. NLP helps AI analyze the tone, context, and structure of the text. For example:
- TF-IDF (Term Frequency-Inverse Document Frequency): This measures how important a word is to a document compared to the whole dataset.
- Word Embeddings (Word2Vec, GloVe): These tools help AI learn the meaning behind words and phrases.
- Advanced models like BERT and GPT: These cutting-edge algorithms analyze context, sentiment, and even writing style — helping detect subtle manipulation.

3.3 Deep Learning: Going Beyond Words

- Deep learning takes things further. It mimics the human brain with layers of "neurons" that learn from data.
- CNNs (Convolutional Neural Networks): Great at spotting patterns in images and headlines — useful for detecting visual clickbait.
- RNNs (Recurrent Neural Networks) & LSTMs (Long Short-Term Memory): These models excel at understanding sequences, like the flow of words in an article.
- Graph Neural Networks (GNNs): These are more advanced and analyze social media interactions — like who shares what — to uncover coordinated misinformation campaigns.

3.4 Hybrid Models: The Best of Both Worlds

Many researchers combine multiple techniques for better results. For instance, a hybrid model might use CNNs for feature extraction and LSTMs to understand the article's overall meaning. This way, the system gets a more complete picture.

4. Datasets: Where Does AI Learn From?

- For AI to learn, it needs good data. Some widely used datasets include:
- LIAR: Contains short statements labeled as true, false, or somewhere in between.

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.021 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 13, Issue 4, April 2025

DOI: 10.17148/IJIREEICE.2025.13405

- FakeNewsNet: Includes both content and social engagement data — helpful for understanding how fake news spreads.

BuzzFeed News & PolitiFact: Real-world datasets with verified articles.
These datasets are crucial for training and testing fake news detection systems. However, they also reveal a big challenge: fake news evolves constantly, so models need regular updates.

5. Challenges: Why It's Not So Easy

- Detecting fake news sounds great — but it's no walk in the park. Here's why:
- Imbalanced Data: There's more real news than fake, which can make AI biased toward believing everything is true.
- Ever-Evolving Fake News: Misinformation creators get smarter and change their tactics.
- Multimodal Content: Fake news isn't just text — it's also images, videos, and even deepfake audio. Current models often struggle to handle this mix.
- Explainability: Many AI models work like a black box — they give answers without explaining why, which makes people skeptical of the results.
- Adversarial Attacks: Bad actors can manipulate content in clever ways to trick AI systems.

METHODOLOGY

1. Data Collection:
The first step involves gathering a large, diverse dataset from reliable sources like trusted news websites, social media platforms, and public datasets such as FakeNewsNet and LIAR. This ensures the system learns from both real and fake news examples.

2. Data Preprocessing:
Next, the collected data is cleaned and organized. This includes breaking the text into smaller parts (tokenization), simplifying words to their base forms (stemming and lemmatization), and removing common, unimportant words like "the" or "is" (stopwords). These steps help the system focus on meaningful content.

3. Feature Extraction:
To make the data usable for the model, the text is transformed into numbers using advanced Natural Language Processing (NLP) techniques. This includes methods like TF-IDF (Term Frequency-Inverse Document Frequency), which highlights important words, and word embeddings like Word2Vec and GloVe, which capture the context and meaning of words.

4. Model Selection:
Various machine learning models are tested to see which works best. Traditional models like Logistic Regression, Decision Trees, and Random Forests are evaluated alongside more advanced deep learning models like LSTM (Long Short-Term Memory) and BERT (Bidirectional Encoder Representations from Transformers), known for their high performance in understanding language.

5. Training and Evaluation:
The models are trained on the prepared data, learning to distinguish between real and fake news. Their performance is then tested using key metrics like accuracy (how often the model is correct), precision (how many flagged articles are truly fake), recall (how many fake articles the model successfully catches), and F1-score (a balanced measure of precision and recall). This helps ensure the system is both effective and reliable.

EXISTING SYSTEM

Right now, detecting fake news is mostly a manual job handled by journalists, fact-checkers, or even the public reporting suspicious content. While these efforts are valuable, they're often slow, influenced by personal bias, and simply can't keep up with how fast misinformation spreads — especially on social media. Traditional methods also lack automation, making it harder to quickly verify the authenticity of articles, posts, or headlines. This leads to delays and inefficiencies, allowing false information to spread before it can be stopped.

DOI: 10.17148/IJIREEICE.2025.13405

Proposed System:

The proposed solution is an AI-powered Fake News Detection System designed to tackle these challenges head-on. Using advanced machine learning and natural language processing (NLP), the system will analyze content in real-time — from news articles and headlines to social media posts — to determine whether it's credible. It will evaluate factors like language patterns, source reliability, and overall content authenticity.

Key features include:

A simple, user-friendly interface where anyone can upload and verify content.
Real-time analysis that delivers a confidence score, showing how likely the content is to be fake.
A dashboard to track fake news trends and monitor how misinformation is spreading.
Automated alerts that notify media outlets and fact-checkers about emerging fake news topics, enabling faster responses.
This upgraded system aims to make fake news detection faster, more accurate, and more reliable — ultimately helping to limit the harmful effects of misinformation on society.

ADVANTAGES AND THE NEW SYSTEM

- Automated Analysis: With machine learning models at its core, the system rapidly identifies fake news, cutting down on manual effort.
- Improved Accuracy: It leverages advanced Natural Language Processing (NLP) algorithms to assess text patterns, sentiment, and overall content credibility, ensuring more reliable results.
- Real-time Monitoring: The system continuously tracks trending stories to spot misinformation as it emerges.
- User-Friendly Interface: Designed with simplicity in mind, it’s accessible to both experts and everyday users without sacrificing functionality.

Enhanced Reporting: Users get detailed insights, complete with confidence scores, helping them make informed decisions.

Technology Behind the System

- This project integrates several cutting-edge technologies:
- Python: A versatile programming language that powers data processing and model development.
- TensorFlow/Keras: These machine learning frameworks help build, train, and fine-tune deep learning models for high performance.
- Natural Language Processing (NLP): Libraries like SpaCy and NLTK enable the system to analyze text data efficiently.
- Flask/Django: These web frameworks make it easy to create an interactive, user-friendly interface.
- BeautifulSoup & Scrapy: Essential tools for web scraping, allowing the system to gather and analyze live news content in real time.

The result? A powerful, responsive, and easy-to-use solution designed to fight fake news faster and more accurately than ever before.

CONCLUSION

In conclusion, the AI Fake News Detection System effectively blends advanced technology with the growing need for accurate information, providing a proactive approach to tackle misinformation. While there’s room for improvement — such as enhancing the model with more training data and refining its performance — the system already establishes a solid foundation for promoting reliable and trustworthy content. Looking ahead, future upgrades could focus on adding multi-language support, improving detection of manipulated visual content, and expanding data sources to ensure even broader and more comprehensive coverage.

REFERENCES

[1]. Academic Journal: Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. Science, 359(6380), 1146-1151. https://doi.org/10.1126/science.aap9559
[2]. Book: Chesney, R., & Citron, D. (2019). Deepfakes and the new disinformation war: The coming age of post-truth geopolitics. Foreign Affairs, 98(1), 147-155.

DOI: 10.17148/IJIREEICE.2025.13405

[3]. Conference Paper: Shu, K., Sliva, A., Wang, S., Tang, J., & Liu, H. (2017). Fake news detection on social media: A data mining perspective. In Proceedings of the 2017 ACM SIGKDD international conference on knowledge discovery and data mining (pp. 22-36).

[4]. Website: Marr, B. (2023). How artificial intelligence is fighting fake news. Forbes. Retrieved from https://www.forbes.com

[5]. Report: World Economic Forum. (2018). Creative Disruption: The impact of emerging technologies on the creative economy. Retrieved from https://www.weforum.org

[6]. Thesis/Dissertation: Gupta, R. (2021). AI-based fake news detection: Challenges and opportunities (Master's thesis, University of Toronto).

[7]. Government or Institutional Report: European Commission. (2020). Tackling online disinformation: A European approach. Retrieved from https://ec.europa.eu

[8]. News Article: Hern, A. (2024). AI is learning to detect fake news — but can it keep up? The Guardian. Retrieved from https://www.theguardian.com.