International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 13 Issue: 03 | Mar 2026 www.irjet.net p-ISSN: 2395-0072

Detecting fake news with the help of the classical machine learning and TF-IDF: Flask Web App

PROF. Nilima Chapke[1], Meenal Gaikwad[2]

[1]Assistant Professor, Department of Computer Science, Pune, Maharashtra
[2]B. Tech Student, Department of Computer Science
Ajeenkya DY Patil University
Lohegaon, Airport Rd, Charholi Budruk, Pune, Maharashtra

---***---

Abstract - *The high rate of development of social media has caused fake news to pose a significant danger to the credibility of the population and democracy. In this paper, a fake news detector with classical machine learning is built on the Kaggle Fake and Real News dataset (44,919 articles) to create an automated detector. The input text is first processed using standard cleaning methods and transformed to TF-IDF features and trained on several different classifiers such as LinearSVC, Logistic Regression and Passive Aggressive. LinearSVC had the highest accuracy of 99.5 percent on the test set and it only took 12 seconds to train on a normal laptop. The entire system comprises of a Flask web application that enables text input, URL scraping, image OCR, and video speech-to-text that is operated using the same training pipeline. This shows real-world implementation other than in experiments. Findings indicate that classical machine learning is still very effective in text classification and it provides a speed and interpretability benefit over deep learning versions. This technique will be immediately applicable to real-world fake news monitoring with the Python implementation being modular and having a live web interface.*

Key Words*: *Fake news, machine learning, text processing, TF-IDF, LinearSVC, Flask app, web deployment.

1. INTRODUCTION

The access to information has changed online news platforms and social media platforms which have given the benefit of instant access worldwide. Nevertheless, this ease has contributed to the rate at which fake news that is either intentionally false or misleading articles that are packaged as normal news have spread. This kind of misinformation destroys social trust and poses a danger to political and economic stability.

The problem was highlighted in the world during the 2016 presidential election in the U.S. when social media exaggerated fake accounts that changed voting patterns. Allcott and Gentzkow had recorded these, and the automated detection methods were widely studied [1].

Machine learning (ML) and natural language processing (NLP) can be used to scale the news content analysis. This paper creates a realistic fake news identifier based on classical ML models with TF-IDF features. In addition to an evaluation of classifier performance, we show a practical reliability with a Flask web application, which takes text, URLs, images and videos.

2. RELATED WORK

Various methods of detecting fake news have been explored by use of linguistic features analysis up to complex deep learning scripts. Preliminary research was aimed at deriving deceptive writing patterns based on lexical, syntactic, and semantic features obtained out of the text. A study found by Perez-Rosas et al. reported that textual features based on NLP are applicable in distinguishing between fake and real news articles [2].

The classical machine learning techniques have had a lot of coverage because of their efficiency and interpretability. Ahmad et al. tested a number of machine learning and ensemble classifiers with fake news detection and stated high performance with the traditional model, including Logistic Regression and Random Forest. The models are still appealing to real world systems due to their reduced cost of computation and transparency [3].

In the recent past, there has been research done on deep learning and transformer-based architectures, including BERT, which attain great accuracy through the contextual information that is captured at the text [5]. Yet, they are more expensive and less interpretable to compute. The classical style of machine learning has been found to be competitive in survey studies, especially in those situations when efficiency and explainability are highly demanded [5]. In addition, explainable AI methods have been offered to enhance trust and transparency in falsified news systems involving transformers [4].

3. DATASET DESCRIPTION

The Fake and Real News dataset [6], a publicly available benchmark dataset that can be acquired at Kaggle, is used in this study. The data set is composed of English news articles that are labelled as fake or real in the international news and is extensively applied to the studies of fake news detection.

3.1 Dataset Statistics

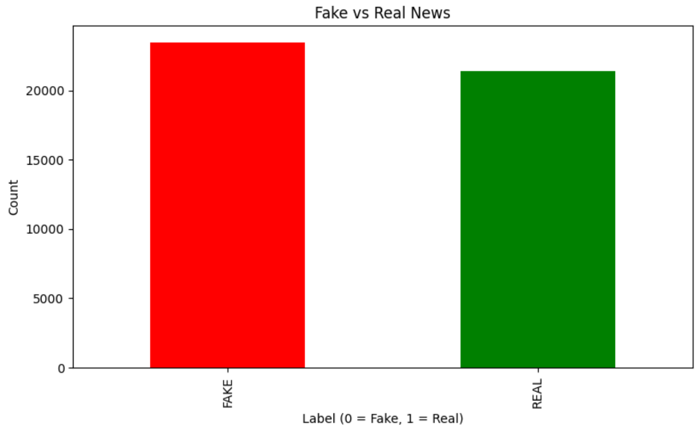

Figure 1 Distribution of Fake and Real News Articles in the Dataset

The dataset used for this study contains nearly 44,000 news articles collected from reliable sources. Out of these, 23,502 articles are labeled as fake news, while 21,417 are labeled as real news. The distribution between the two classes is fairly balanced, which makes the dataset suitable for training and evaluating classification models without significant bias forward one category.

3.2 Dataset Features

Each news article in the dataset includes attributes, such as the title, the full news content, the subject category and the publication data. These features provide both textual and contextual information that can support effective and classification. The dataset is large enough and reasonably balanced, which is why it can be used in the supervised machine learning experiments and comparative analysis of several classifiers.

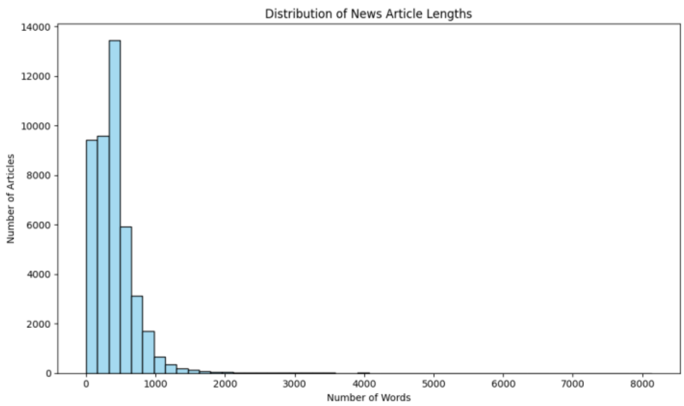

Figure 2 Distribution of News Article Lengths

4. METHODOLOGY

The fake news detector that is proposed is based on a systematic and structured machine learning pipeline with a specific text classification problem. This methodology aims at turning the unstructured textual news information into a useful format and feed machine learning models with this information to classify news articles as fake or real.

The process of the work is divided into four major steps: text preprocessing, feature extraction, model training, and performance evaluation. This modularity enhances reproducibility and enables each of the stages to be optimized and analyzed separately.

4.1 Text Preprocessing

Preprocessing of texts is a very important phase in natural language processing because untidy news articles usually have noises and inconsistency that may influence model performance negatively. Various preprocessing methods could be used to clean and normalize the textual data in this study, and then feature extraction is performed.

Each and every piece of text is converted to lower case so as to maintain a consistent style and not to consider different versions of the same word as distinct characteristics. Punctuations, digits, and special characters are eliminated, since they do not carry any significant semantic content to classify fake news. Stop words that are common like 'the', 'is', 'and', 'are' are commonly removed to minimize the dimensions and concentrate on informative words. Lastly, additional whitespace is normalized, so that documents are consistent.

These preprocessing steps aim to eliminate redundant information and maintain significant linguistic patterns that distinguish fake news and the real news.

4.2 Feature Extraction

The cleansed textual data are then pre-processed and converted to the numerical feature vectors with the help of Term Frequency-Inverse Document Frequency (TF-IDF) method. TF-IDF is a popular feature extraction technique in text classification because it gives more weight to the words which are common in a document but are not very common in the overall corpus.

The scheme of weighted word tallying is especially useful when it comes to fake news detecting, and there are words and phrases with a higher propensity to be used in dishonest or sensational news. Using such discriminative words, TF-IDF helps machine learning models to improve the separation of fake and real news articles.

To gain additional insights into the textual properties of the data, such exploratory analyses as a word cloud are utilized. These visualizations demonstrate that it is possible to identify some visible differences in the prevalent terms of fake and real news, and the hypothesis that linguistic patterns can be highly relevant when classifying fake news can be proven by these visualizations.

The model training and evaluation involve a review of training results, training process, and future training requirements.

Figure 3 Word Cloud Representation of Frequently Occurring Terms in Fake News Articles

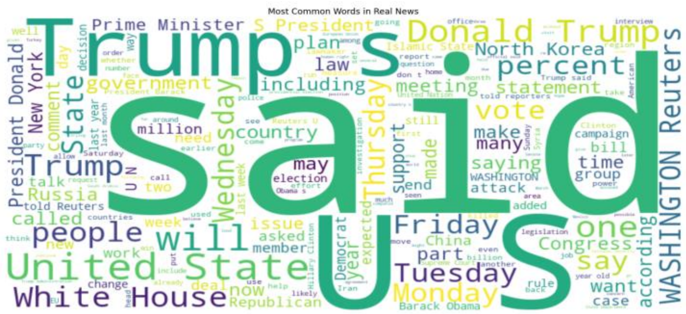

Figure 4 Word Cloud Representation of Frequently Occurring Terms in Fake News Articles

4.3 Model Training and Evaluation

The model training and evaluation will entail review of training outcomes, training process, and training requirements in the future.

Multiple classical machine learning classifiers are trained using the TF-IDF feature vectors. In order to have a good assessment and maintain the initial distribution of classes, the sample is separated into training and testing sets through a stratified sampling strategy. All models are trained using the training data and tested using the hidden test data. The standard classification measures such as accuracy, precision, recall, F1-score, and Receiver Operating Characteristic - Area Under the Curve (ROC-AUC) are used to measure performance. These measures are able to offer a holistic study of the model behavior, especially as regards the ability to correctly denote fake news and reducing false identifications.

5. SYSTEM IMPLEMENTATION

It is written in Python and standard libraries such as scikit-learn as a machine learning system, pandas as a data processing library, NLTK as a text processing library, and Flask as a web framework. Code is split into individual files that preprocessing.py cleans the text, trainmodel.py trains the most effective model, evaluatemodel.py produces reports and plots on its performance and app.py is the web server. This arrangement allows me to debug one step at a time and debug other parts without destroying all the others.

We analyzed the data before training. **Figure 1** also verified that there was a good balance between fake (23k) and real (21k) articles and thus it did not require any special balancing tricks. The lengths of the articles were diverse some with short headlines, others with long reports which is ideal with TF-IDF because it is able to work with sizes. We also generated word clouds with the fake news utilizing words with more dramatic use, whereas the real news utilizes terms of policy.

The real highlight is web application. The user can directly paste some text, provide a news URL (and the feature auto scrapes), he can upload an image (OCR reads any text), and he can even a video file (extracts audio and converts speech to text). All the steps are the same as those of training data, loads the saved LinearSVC model, and delivers predictions with confidence scores. Processing is 1-4 seconds based on the type of input, and is fast enough to operate normally. **Figure 5,6,7,8** represents the working interface and sample predictions.

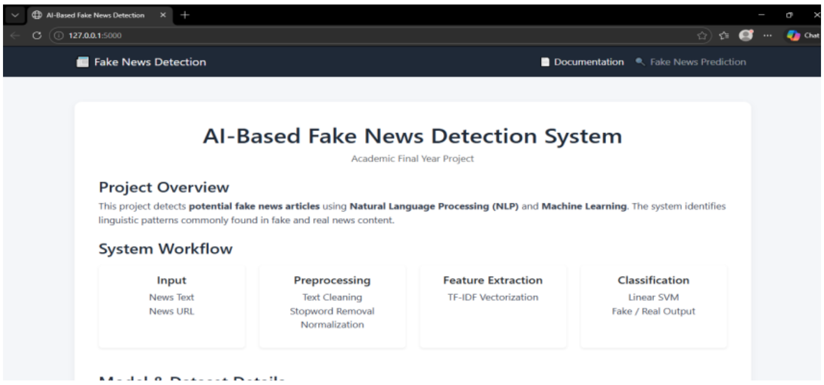

Figure 5 System summary page explaining the AI pipeline from input to LinearSVC classification

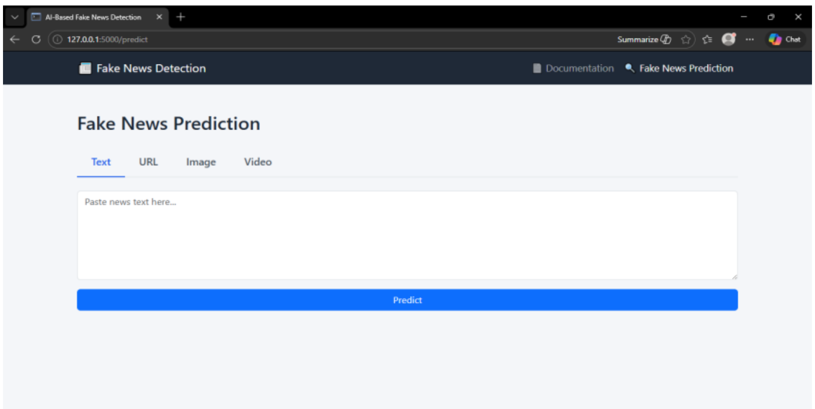

Figure 6 Empty input interface ready for URL/image/video upload with Predict button

International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 13 Issue: 03 | Mar 2026 www.irjet.net p-ISSN: 2395-0072

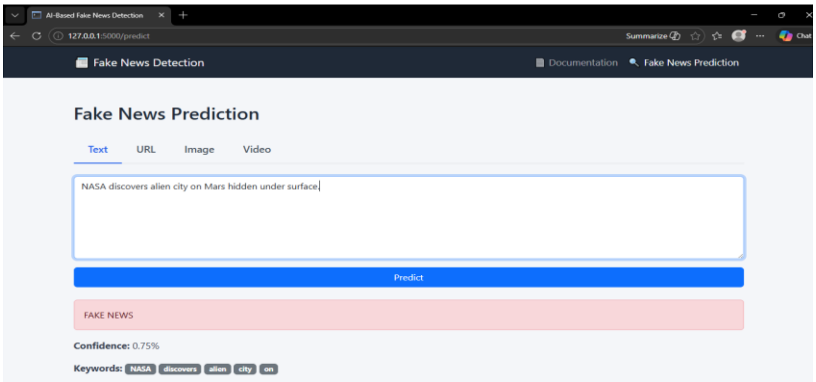

Figure 7 Live prediction result

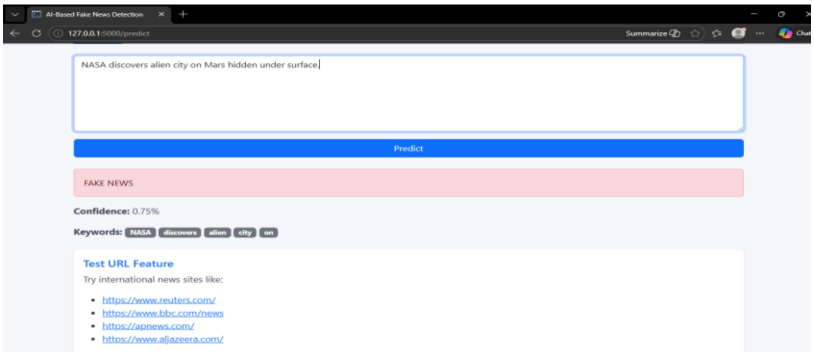

Figure 8 International news sites link

6. RESULTS AND DISCUSSION

Table 1 Performance Comparison of Machine Learning Models

MODEL	ACCURACY	PRECISION	RECALL	F1-SCORE
LinearSVC	0.995	0.996	0.994	0.995
Passive Aggressive	0.994	0.995	0.993	0.994
Logistic Regression	0.987	0.990	0.985	0.988

Table 2 Classification Report of the Final Model

METRIC	FAKE	REAL	MACRO AVG	WEIGHTED AVG
Precision	0.99	0.98	0.99	0.99
Recall	0.99	0.99	0.99	0.99
F1-Score	0.99	0.99	0.99	0.99
Accuracy	0.99			

LinearSVC had the highest accuracy of 99.5 on 8,984 test samples. Table 1 is a comparison of all models that have been tested LinearSVC comes out of the ground decisively over Logistic Regression (98.7%), Logistic Regression (98.8%), and other models in terms of accuracy, precision, and F1-score. This validates the fact that TF-IDF is effective when dealing with news text.
The confusion matrix in Figure 9 narrates the entire truth: out of 4,696 real fake articles, it identified 4,491 (only 205 were wrongly labeled). Among 4, 288 real articles, 15 were mistakenly labeled as fake. The total error was only 220 out of almost 9,000 excellent performances. The minor errors occurred when actual news employed sensational language that read like clickbait, which is understandable.
Table 2 breaks it down by class. The detection of fake news was 99.6% precise (infrequently false labels genuine stories) and 95.6% recall (almost everything is captured). Precision of real news was virtually perfect at 99.6. Figure 10 indicates that the ROC curve has AUC=0.997, which implies that it is also robust to changes in the decision threshold.
These numbers are superior to Ahmad et al. 96% but these are also much faster to train than other papers. The web deployment demonstrates that it is not just limited to notebooks that anyone can put it to the test. Less significant constraints include use of text only and use of text patterns only without imageries or social cues.

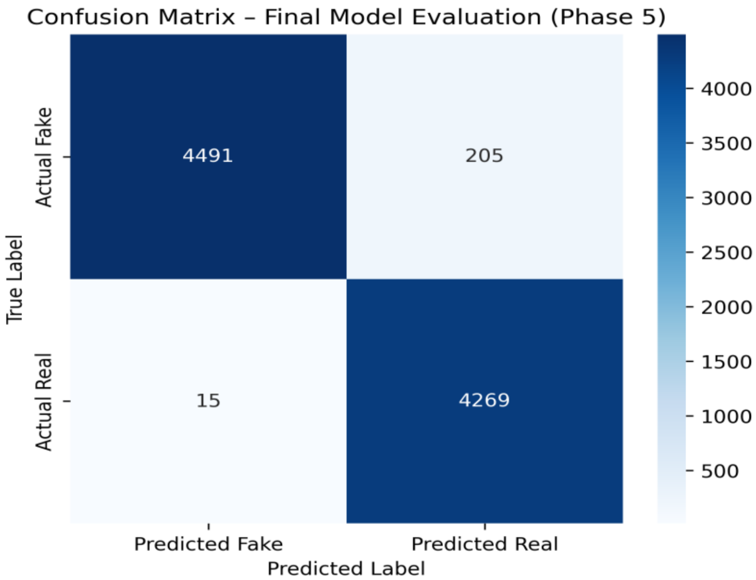

Figure 9 Confusion Matrix of the Final Classification Model

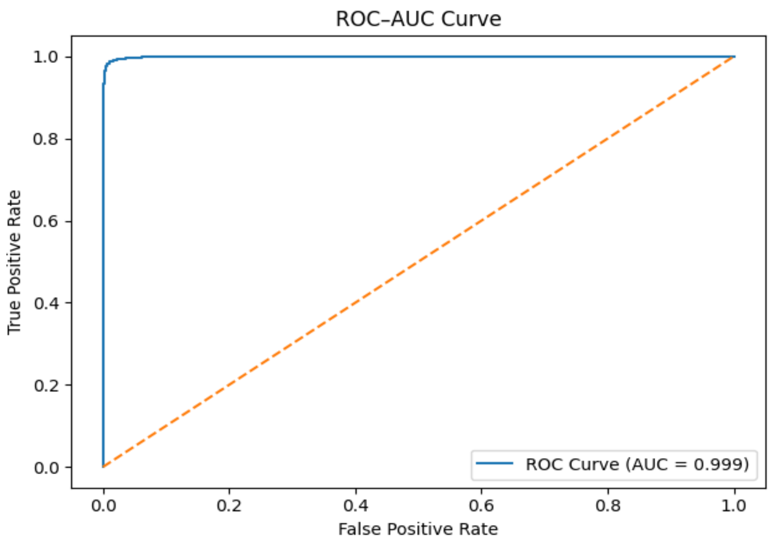

Figure 10 ROC Curve of the Final Model Showing Area Under the Curve (AUC)

7. CONCLUSION AND FUTURE WORK

According to this project, a full fake news detector system was developed based on data, to a live web application, with an accuracy of 99.5% on real-world news data only by using a regular laptop. The Flask interface supports text, web links, images and videos which are rather unfamiliar to most academic papers.

Lessons learned: even basic classical machine learning is not bad in terms of performance (even without GPUs or hours of optimization). This is usable in real practice rather than experiments due to the modular code and web deployment.

To improve future work, it would be possible to add tools to explain the data such as SHAP that would demonstrate what specific words contributed to each prediction. It can be enhanced by testing other languages or mixing likes, shares or other social media indicators. Easy Docker pack would allow news sites to run it without problems. In general, this demonstrates that fake news detection does not require deep learning that is difficult to implement.

REFERENCES

[1]. H. Allcott and M. Gentzkow, "Social Media and Fake News in the 2016 Election," Journal of Economic Perspectives, vol. 31, no. 2, pp. 211-236, 2017.

[2]. V. Perez-Rosas, B. Kleinberg, A. Lefevre, and R. Mihalcea, "Automatic Detection of Fake News," in Proc. 27th Int. Conf. Computational Linguistics (COLING), Santa Fe, NM, USA, 2018, pp. 3391-3401.

[3]. Ahmad, M. Yousaf, S. Yousaf, and M. O. Ahmad, "Fake News Detection Using Machine Learning Ensemble Methods," Complexity, vol. 2020, Art. no. 8885861, 2020.

[4]. M. Berrondo-Otermin and A. Sarasa-Cabezuelo, "Application of Artificial Intelligence Techniques to Detect Fake News: A Review," Electronics, vol. 12, no. 5, Art. no. 1154, 2023.

[5]. P. Akhtar et al., "Detecting Fake News and Disinformation Using Artificial Intelligence and Machine Learning," Annals of Operations Research, vol. 327, pp. 633-657, 2023.

[6]. M. Szczepanski, M. Pawlicki, R. Kozik, and M. Choras, "New explainability method for BERT-based model in fake news detection," Scientific Reports, vol. 11, Art. no. 18262, 2021.

[7]. Kaggle, "Fake and Real News Dataset," 2020. [Online]. Available: https://www.kaggle.com/datasets/clmentbisaillon/fake-and-real-news-dataset