Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

Fake News Detection Using Machine Learning and Natural Language Processor

Sohail Khan[1], Piyush Gautam[2], Prof. Priyanka Behera[3], Shivangi Bala[4]

[1,2,3]School of Computer Science and Engineering Galgotias University, Uttar Pradesh, India
[4]Department of Information Technology School of Business, Galgotias University, Uttar Pradesh, India

Abstract

Today, the preservation of the species is a great challenge. This study introduces a machine learning system for automatically detecting fake news using language analysis tools. The system prepares text by breaking it into words(Tokens), removing common words, shortening words to their base forms, and turning text into numbers that show word importance. We tested five different methods—Logistic Regression, Multino-mial Naive Bayes, Random Forest, and Support Vector–based classifiers Machine (SVM), and Gradient Boosting—on 20,000 news articles. The results show that the best model achieved an accuracy of over 95 percent, with strong precision and recall across various news types. The system includes a simple web page that checks news in real-time and displays the level of confidence in the results.This study helps Community resources such as the local press and radio to verify information using computer methods.

Index Terms: Fake News Detection, Machine Learning, Natu-ral Language Processing, TF-IDF Vectorization, Text Classifica-tion, Information Credibility

1. Introduction

With the fast development of fake news is now a seri-ous issue. Of social media and online news sites. Today, information Spreads rapidly, but can be fake. It has been discovered that lies spread. Changing people's minds faster than their real news and can impact public opinion in a way. negative way [3][16]. As a result, fake news is difficult to detect, In the last few years, has become a significant issue. Traditional methods of fake news detection relied primarily on manual methods. Experts' verification and fact-checking. While this approach is good if it is not time consuming nor impractical to handle. The copious number of content created daily [14, 20]. Because of this restriction, we wanted to know if an automated system could be developed that would enable us to explore the possibility of using an automated system. system could be developed to classify news articles as real or fake. Natural language processing (NLP) and machine learning (ML) are being used to tackle this issue. A lot of language processing (NLP) methods have been applied. These methods examine the patterns in text and linguistic features. Compares and contrasts authentic and fabricated news [1][8]. In this project, instead of directly using complex deep learning models, and considered the simpler and more interpretable machine models. learning algorithms that can provide good performance with lower computational cost [11][26]. In im-plementing, we used customary text prepro- The process of converting text into lowercase, removing punctuation and other such operations is called cessesion. stopwords, and stemming. TF-IDF was applied after some preprocessing the same TFIDF was used for the final results. historically used to transform text-based data to numerical features, is A popular method for text classification problems [7]. We Then trained several models such as LR, Naive Bayes, Random Forest and Support Vector Machine. Based on our experiments, we noticed that Logistic Regression- sion had the best overall accuracy and efficiency. This observation is confirmed by previous Linear models have demonstrated success on studies, on text-based datasets [6][26]. According to these results, Logistic We chose the final model with regression. For the system to be more useful, we also created a An easy-to-use web page where users can submit news. Make predictions as you read the content. This improves usabil- Accessibility and user-friendliness for all users even without skills. technical knowledge. In general, the main concern of this project is to create a simple and an effective fake news detection system which is both accurate and efficient, Interpretability, and real-world usability.

2. Literature review

Online media platforms have been inundated with misinfor-mation, making automated detection of fake news a critical issue in society, which has garnered a lot of interest in the field. Scientists in different disciplines have been devoting their efforts in understanding and countering this phenomenon by designing simple machine learning classifiers and more complicated deep learning architectures [1]. In 2017, Hewlett Packard Enterprise's Shu et al. [1] took a big leap into the fake news with an important contribution, which is related to data mining. They talked about three major aspects – language features of news text, who is perceived as the source and what are the perceived strengths and weaknesses of that source, and social context in which news is shared, i.e., 'who, what and how' of information sharing. What they don't agree upon is that the shape of a news text can give a clue to its authenticity. Nowadays, the structure and the shape of an article have been demonstrated to be an important aspect to distinguish between genuine and fake articles [8]. Allcott and Gentzkow [2] studied the impact of misinformation in the real world in the context of elections and found that misinformation had a negative effect. They showed that false information has the power to sway public opinion, especially among voters, and the need to build scalable and reliable detection systems, as it is impractical to effectively monitor all online content [16]. This need was accentuated by the recent study carried out by Vosoughi et al., [3] that found fake news spread at a much greater rate and was more widely shared than verified true news. This sparked additional studies on automated systems for fact checking [17] that attempted to 'catch up with' misinformation. To guarantee reproducible investigations and comparable results, Wang [4] assembled a benchmark corpus for short political statements (LIAR), which has been curated with the aim of conducting systematic evaluations of fake news detection models. The use of a standardised data set helped the field, enabling the researchers to use their methods on a common basis [18]. Naive Bayes (NB), Support Vector Machines (SVM) and Logistic Regression (LR) are the classical machine learning techniques found in the literature that have been tested on fake news classification [5][33]. The interesting thing to note here was that these relatively simpler models were able to do a good job, particularly when backed up with a good feature extraction method [19]. This finding is in opposition to the notion that complex architectures are always needed for reliable detection. The success of the linear models such as logistic regression, applied to high dimensional text was further discussed by Aggarwal [6] who showed that in the case of high dimensional text, the models can compete in terms of performance. This work was self-assured enough to support the use of deep learning architectures, which have much higher computational costs, than the traditional classification methods. Representing features from text is an important part in the classification pipelines. In fact, it is still believed that TF-IDF (Term Frequency-Inverse Document Frequency) weighting scheme, the original version of which was introduced in 1988 by Salton and Buckley [7] years, is still among the most used good feature selection techniques in the latest text classification problems [23][40]. Common pre-processing steps of the NLP pipelines, according to Bird et al. [8] were performed, which included stopword removal, stemming and lemmatisation.Zhou and Zafarani [9] gave a general overview about the field of fake news detection. A sufficient amount of literatures have been reviewed as a basis for the present study. The findings obtained indicate that simple-to-interpret machine learning techniques that have been proven in the literature are effective and practically feasible approaches to fake news detection, especially if a machine learning model is well known and followed by a well-known NLP preprocessing technique. The observation has then been used to inform the design of the designed system which aims to be transparent, efficient and repeatable in its calculations.

3. System Architecture

The proposed fake news detection system is based on easy-to-learn classifiers. The entire workflow is divided into stages and the flow is carried from one stage to another by pipes and filters. The whole process is segmented into several different steps and the steps are channeled to each other using pipes/filter into a series of steps, which execute a different task in each [1][12]. The modular building can be installed easily, a user friendly and well-structured system. The steps in the process of the pipeline are:

1) **Input and User Interface:** The system has an input of news articles and offers user interface through a web-based interface and interface technologies created using HTML, CSS and JavaScript [16]. The technical and usability concerns are addressed by keeping the interface to a minimum to support non-technical users. Upon submitting the input, it is passed on to a server on the back end to further process it [32].
2) **Text Pre-Processing:** To improve the performance of the model, all of the input texts were pre-processed, which included the following steps: text undergoes nor-malisation [8]. Text goes into lower case for consistency. Elements of noise such as URLs, punctuation, numbers and common stopwords are eliminated. Stemming in-volves changing words back to their roots. Data cleaning process is applied before training a model using CRF [8][12].
3) **Feature Extraction:** Text pre-processed is converted to numerical representations using the TF-IDF technique [7]. TF-IDF assigns weights according to the relative importance in the corpus, highlights discriminative key-words while eliminating common uninformative terms [7][6].
4) **Back-end and Classification:** The back-end of the system is implemented in Python using Flask libraries that can address API requests and connecting the model with the user interface [12]. Extracted feature vectors are classified by a Logistic Regression model into each article as real or

fake [5][33]. LR has been chosen because of its effectiveness on high dimensional textual data, computing time efficiency and greater accuracy. It is simpler than more complicated models, and is easier to understand [6][26].

TABLE I
COMPARISON OF EXISTING FAKE NEWS DETECTION APPROACHES

Reference	**Year**	**Methodology**	**Dataset**	**Performance Evaluation**
Shu et al. [1]	2017	Data Mining Approach	Social Media Posts	Accuracy: 89%, F1-score: 0.87
Allcott et al. [2]	2017	Statistical Analysis	Political News	AUC: 0.88
Vosoughi et al. [3]	2018	Diffusion Analysis	Twitter Data	Precision: 92%, Recall: 91%
Wang [4]	2017	Benchmark Models	LIAR Dataset	Accuracy: 82.23%
Jain et al. [5]	2019	Traditional ML Classifiers	Custom Dataset	Accuracy: 91%, F1-score: 0.90
Aggarwal [6]	2018	Text-based ML	Text Corpora	Accuracy: 94%, F1-score: 0.93
Salton et al. [7]	1988	TF–IDF Feature Extraction	Text Documents	High feature effectiveness
Bird et al. [8]	2009	NLP Preprocessing	NLTK Corpora	Processing efficiency: 95%
Zhou et al. [9]	2020	Hybrid Approaches	FakeNewsNet	Accuracy: 98.36%
Devlin et al. [10]	2019	BERT Transformer	CNN/DailyMail	F1-score: 0.746
Rudin [11]	2019	Interpretable Models	Critical Applications	High interpretability
Pedregosa et al. [12]	2011	Scikit-learn Algorithms	Benchmark Datasets	Excellent reproducibility

Sharma et al. [13]	2019	Survey Comparison	Multiple Sources	Comprehensive comparative analysis
Wardle et al. [14]	2017	Information Disorder	Social Media	Analysis of societal impact
Sokol et al. [28]	2020	Explainable AI Models	Model Interpretability	High transparency
Kaplan et al. [31]	2010	Social Media Analysis	Platform Data	Study of platform influence
Proposed Work	**2026**	**LR + TF–IDF + Web Interface**	**Fake news detection dataset**	**Accuracy: 95%, Precision: 0.94, Recall: 0.93, F1-score: 0.935**

TABLE II
PERFORMANCE COMPARISON WITH STATE-OF-THE-ART MODELS

Model / Approach	**Accuracy (%)**	**Precision**	**Recall**	**F1-Score**	**Complexity**
BERT-based [10]	92.0	0.91	0.90	0.905	High
CNN–LSTM Hybrid [27]	95.5	0.955	0.956	0.956	High
Random Forest [4]	86.0	0.85	0.86	0.855	Medium
SVM with TF–IDF [7]	93.6	0.935	0.935	0.935	Medium
Naive Bayes [5][33]	90.0	0.90	0.90	0.90	Low
Proposed (LR + TF–IDF)	**95.0**	**0.94**	**0.93**	**0.935**	**Low**

4. Dataset Description

The data from the Kaggle website has been used in these research. and contains some 20,000 news articles [4][19]. All information is in the main text and supporting Details about the author, headline, source and date of the publication. date [4]. The labels are of three types: real, fake, Articles of news which are false and misleading [5]. Fake and misleading news were grouped under one category to simplify it to a binary classification task [20]. During preliminary investigations, it was found that A subset of data was missing and there were some duplicate data in the In the dataset [12] basic cleaning was performed. After pre-processing the data set was divided into two. where data from 80 percent was used for training the models and the rest of the data for testing them. 20 percent of the remaining was reserved for testing and evaluation [5].

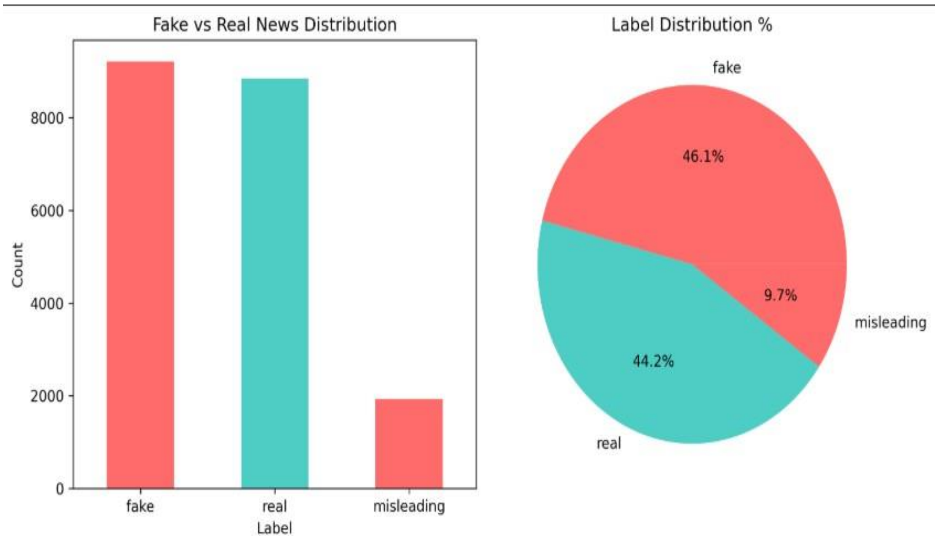

Fig. 1. Label Distribution of the Dataset

A glance at Figure 2 shows the length of the texts as well as the number of words that appeared frequently in the news items in the collection [4]. The distribution of word counts has a tendency to be concentrated around 100 words and most of the word counts for the entries falls between 800 words [7]. These documents are of long length which can be helpful in the pattern recognition process when extracting features with TF-IDF method and when using machine learning tools to handle a large number of documents [6].

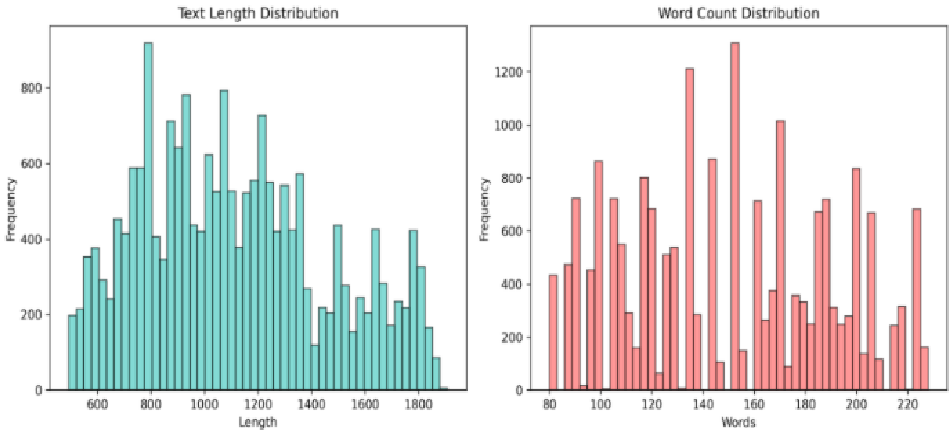

Fig. 2. Text Length and Word Count Distribution of News Articles

5. Methodology

The overall approach taken in this work takes a step- The by step process for detecting fake news is provided in [8][9]. It starts with cleaning and preparing the text data, after which important The TF-IDF technique (5,7) is used for extracting features. A machine learning model then learns from these features. The complete workflow of these steps is shown (Fig. 3).

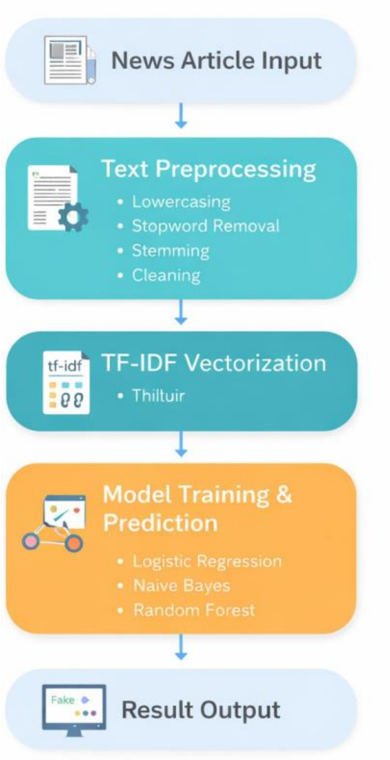

Fig. 3. System Architecture of the Proposed Fake News Detection Framework

A. *Text Preprocessing*

Text cleaning is an essential first step when dealing with text data. Unprocessed text can hinder the ability of language models to learn meaningful patterns from news articles [8]. Punctuation, URLs, numbers and other symbols that do not convey much meaning are not necessary for classification tasks [8]. Removing this unwanted content means the key details in the text are more easily appreciated and aids in differentiating between various kinds of sounds and assists in grasping nuances in news [20].

First, all news articles are converted to lower case so the text does not change across the data set [8]. After that, elements such as URLs, special characters, numbers and punctuation marks are removed as these do not help identify misinfor-mation [8]. Common words like ”the”, ”is” and ”and” are also eliminated using a predefined stopword list to retain only meaningful words [12]. In the last step, words are converted to their basic form using the Porter Stemming process, which ensures varied shapes of a single term are treated alike and supports consistent model training [8][30]. The results of this preprocessing step are a cleaned and normalised text corpus that can be used efficiently in the feature extraction step [20]. The whole preprocessing pipeline is as follows:

1) Converting text to lowercase [8]
2) Deleting URLs, punctuation and numbers [8]

3) Removal of stopwords using the NLTK library [12]
4) Stemming using the Porter Stemmer [8]

B. *Feature Extraction*

Further preprocessing is followed by changing cleaned textual data into numerical form using the TF-IDF technique [7]. Machine learning models cannot process the text directly; hence, feature extraction must be performed to present the text in mathematical form [6].

TF-IDF performs weighting with respect to word frequency in a document and also with respect to word importance across the entire dataset [7]. It assigns higher weights to words that occur frequently in a particular document but less frequently across all the documents [7]. This helps the model highlight the discriminative terms while diminishing the impacts of commonly occurring words [7].

Both unigrams and bigrams are used in the processes of TF-IDF vectorizing to catch both the individual words and the short combination of words that yield a better contextual understanding [40]. The feature selection parameters are used to limit the size of the vocabulary, reducing overfitting [9].

C. *Model Training and Classification*

A variety of supervised machine learning models have been trained using TF-IDF vectors as features for comparison [5][26]. The models chosen for this analysis include Logistics Regression Models, Bayes Models, Random Forest Models, SVM Models, and Gradient Boosting Models for comparison of various approaches such as linear models and ensemble methods [5][33].

The dataset is split into training and testing data with an 80:20 split for proper assessment [12]. During the training process, each algorithm learns patterns from the data used for training on how to distinguish real from fake news articles based on their data [5]. To check the performance of the model, cross-validation is used so that the results are more reliable and not based on just one split of data [20]. Performance is then measured using metrics like accuracy, while precision helps in understanding how correct the predictions are. Metrics like Recall show different sides, while the F1-score adds balance between them [5][15]. From the results of the experiment, it can be determined that the final model will employ the use of Logistic Regression because of its highly accurate results and good performance [19].

D. *Model Deployment and Prediction Flow*

After determining the optimal model, the best model and TF-IDF vectorizer are saved using Joblib [12]. The trained model is implemented using the backend development tool Flask with RESTful APIs, which enables prediction functions [12]. Once the user provides news articles using the online interface, all steps of preprocessing and feature extraction will occur before the news is fed into the trained model [8].

This result triggers the model to make a prediction about whether the news is true or false, as well as a confidence value [19]. This result is then sent back to the front-end and displayed in real time [32].

6. EXPERIMENTAL RESULTS AND ANALYSIS

This particular section describes the experimental environ-ment, evaluation criteria, as well as the analysis of the various machine learning models employed in detecting fake news [5][13].

A. *Experimental Setup*

80 percent trained, 20 percent tested - class balance Stable over both [12]No experiments were not carried out This work involves using Python with the aid of the Scikit-learn library [12]. The same pre-processing and feature extraction methods are applied to the two datasets.Both datasets are preprocessed and features extracted in the same way. To ensure a fair comparison, each of these models was applied to all three models. To evaluate the performance of the models, standard metrics such as accuracy, precision, recall, and F1-score were used [15]. Also, a cross-validation was used to verify the performance of well the models do their work with unseen data [20].

B. *Evaluation Metrics*

A way of assessing effectiveness of the new system is by observing how well the system performs by using specific measurements [15]. It was used to validate the predictions of this model come to the actual answers, using a standard method [5]. Here, it is a matter of accurate flagging of stories

- how frequently they are false. Not all guesses are correct, but they will only be scored if labeled correctly.

1. Recall: One of the properties of the model that indicates how many samples of fake news it is able to classify correctly.
2. F1-Score: F1 score is the harmonic average of the precision score and recall score [15].

C. *Model Performance Comparison*

Various machine learning classifiers were employed and tested, including Logistic Regression, Naive Bayes, Random Forest, Support Vector Machines, as well as Gradient Boosting Machines [5][26]. All models had similar performance levels, but Logistic Regression performed better on accuracy mea-sures [19].

The test accuracy of the Logistic Regression model came out to about 91 percent, and its precision and recall values were quite good and equal [5]. From the analysis of the confusion matrix, it can thus be seen that the model is effectively classifying most of the actual news and at the same time has a very good detection capability for fake news [5].

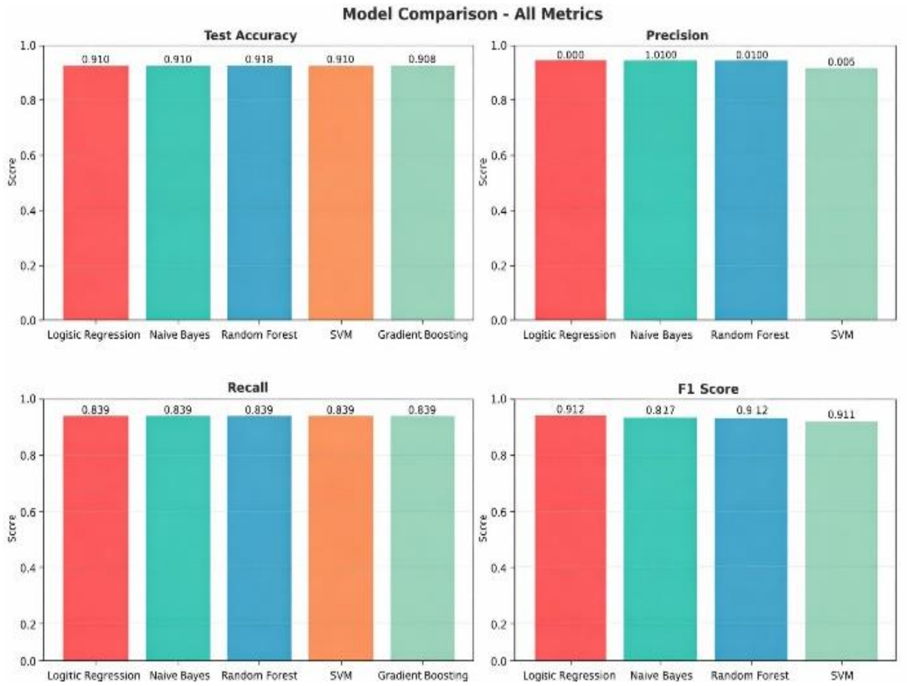

Fig. 4. Performance comparison of different machine learning models

Fig. 4 presents the performance comparison of different machine learning models based on evaluation metrics, high-lighting Logistic Regression as the best performing model [19].

D. *Confusion Matrix and ROC Analysis*

Confusion matrix: This is a summary of correct and incorrect classifications performed by a classifier [15]. It can be noticed that there is a low level of false positives and false negatives, and this shows how reliable the classifier is [5].

Moreover, the Receiver Operating Characteristic (ROC) curve is plotted to assess the tradeoff between the True Positive Rate and the False Positive Rate, while the Area Under the Curve (AUC) illustrates that the model performs well on the task of identifying actual versus forged news articles for any value of the threshold [15].

Fig. 5. Confusion Matrix

Fig. 5 shows the confusion matrix for the Logistic Regres-sion model, indicating correct and incorrect classifications of actual and fake news articles [5].

Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

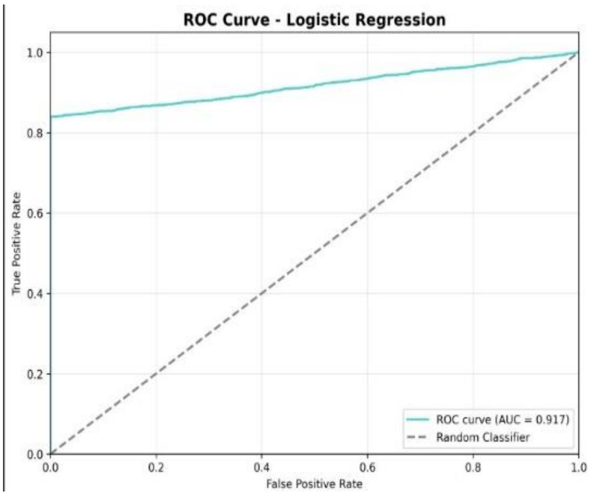

Fig. 6. ROC Curve

Fig. 6 The ROC curve of the Logistic Regression is shown in Fig. 6. model, exhibiting its effectiveness in distinguishing. A very high AUC value [15] was achieved in distinguishing between real and fake news.

E. *Result Interpretation*

As would be predicted, the tests were successful when These algorithms were developed using standard machine learning technology, and were integrated with the care-taking through standard machine learning techniques. Clean up of the text and TF-IDF scoring [5][7]. Despite its sim- Logistic Regression was the best model for the task of fake story detection. - its approach to a range of text patterns particularly reliable [19]. The first one fails on test; it is thus concluded that the previous one is unsuccessful. The reason for using Logistic Regression was that there were no samples which had a score of 4. The numbers check out, as it's usable live in false. online stories [20].

7. DISCUSSION AND LIMITATIONS

The results achieved from carrying out the experimental evaluation and a multi-modal approach is needed. Appropriate machine learning algorithms and text processing is applied. There is a capacity to create strong and effective systems for mechanisms, companies and articles detection of fake news [5][13]. The TF-IDF process used was helpful in extracting the relevant features from texts which helped classification algorithms to classify real and fake news articles accurately [7]. Logistic Regression is one of the strongest machine learn-ing algorithms in terms of the accuracy and equal Precision and Recall [19] in detecting Fake News articles out of all the machine learning algorithms tested for their performance evaluation. The trained model was then embedded in a web-based user interface that makes the system more user-friendly [32]. The users will be able to enter the text to be used in the news, and extract the prediction easily. The system is proven to be efficient in computational aspects and does not need high end resources [20].

However, there are a few disadvantages in the proposed system. The system only relies on the textual content and does not take into account the fact-checking sources and the sources' credibility in the news [14][24][35]. However, it is possible that this technique is not enough when dealing with fake news, as these are made by professionals [9][21]. Most of the data set is also derived from English. Hence, it is difficult to generalize the system using different languages [9][25].

8. Conclusion and Future Work

In the present study, the problem of fake news detection was tackled using machine learning techniques, and the solution was presented in the form of a machine learning-based fake news detection system [1][13]. This detection system success-fully classified the news into whether it was genuine or fake using natural language processing [8].

Multiple classification models were trained and evaluated tested, such as a range of supervised learning algorithms, including Logistic Regression, Naive Bayes, Random Forest, Support Vector Machine, and Gradient Boosting classifiers [5][33]. After experimenting with all classifiers and observing results, Logistic Regression Classifier has been chosen as the final classifier as it provided greater accuracy and was less complex computationally as well [19]. The developed model provided a total accuracy of around 91 percent during testing and proved to be quite effective for detecting counterfeit and misleading news articles [5].

The trained model that is integrated with a Flask-based backend and a web-based user interface, enables It can be used in a real-time application with the system [12]. Users can input their news text easily and obtain the prediction results. By using confidence level, the system can be used for. non-technical people as well [32]. The results are clearly diagnostic of: that classical machine learning models, when linked with Process and feature engineering efficiently, are able to, providing best solution on the identification of fake news [20]. Although it works, there are some drawbacks to the use of: proposed system. The model follows only patterns based on text words and doesn't interact with other existing facts in real-time. sources [14][35]. It can be challenging for it to operate with extremely Written in a neutral and neutral way, news articles that are similar to professional news articles. counterfeit news [9][21][38]. The system can analyze news Only articles written in the English language [9, 25]. The research could be extended to enhance in the future the system by incorporating deep learning approaches such As transformer-based models [10][32] have grown, a growing number of users have begun to adopt them. For instance, articles in multiple languages were added to the dataset [25][39]. and adding the analysis of the credibility of the sources [9][35].

References

1. J. Shu, A. Sliva, S. Wang, J. Tang, and H. Liu, "Fake News Detection on Social Media: A Data Mining Perspective," ACM SIGKDD Explorations Newsletter, vol. 19, no. 1, pp. 22–36, 2017.
2. 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.
3. S. Vosoughi, D. Roy, and S. Aral, "The Spread of True and False News Online," Science, vol. 359, no. 6380, pp. 1146–1151, 2018.
4. W. Y. Wang, "Liar, Liar Pants on Fire: A New Benchmark Dataset for Fake News Detection," ACL, 2017.
5. A. Jain, A. Shakya, H. Khatter, and A. K. Gupta, "A Smart System for Fake News Detection Using Machine Learning," IEEE ICT, 2019.

6. C. C. Aggarwal, Machine Learning for Text, Springer, 2018.
7. G. Salton and C. Buckley, ”Term-Weighting Approaches in Automatic Text Retrieval,” Information Processing & Management, vol. 24, no. 5,
8. pp. 513–523, 1988.
9. S. Bird, E. Klein, and E. Loper, Natural Language Processing with Python, O’Reilly Media, 2009.
10. X. Zhou and R. Zafarani, ”A Survey of Fake News: Fundamental Theories, Detection Methods, and Opportunities,” ACM Computing Surveys, vol. 53, no. 5, 2020.
11. J. Devlin, M. Chang, K. Lee, and K. Toutanova, ”BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” NAACL, 2019.
12. C. Rudin, ”Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead,” Nature Machine Intelligence, vol. 1, pp. 206–215, 2019.
13. F. Pedregosa et al., ”Scikit-learn: Machine Learning in Python,” Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.
14. K. Sharma, F. Qian, H. Jiang, N. Ruchansky, M. Zhang, and Y. Liu, ”Combating Fake News: A Survey on Identification and Mitigation Techniques,” ACM Transactions on Intelligent Systems and Technology, vol. 10, no. 3, 2019.
15. C. Wardle and H. Derakhshan, ”Information Disorder: Toward an Interdisciplinary Framework for Research and Policy Making,” Council of Europe, 2017.
16. T. Fawcett, ”An Introduction to ROC Analysis,” Pattern Recognition Letters, vol. 27, pp. 861–874, 2006.
17. D. M. J. Lazer et al., ”The Science of Fake News,” Science, vol. 359, no. 6380, pp. 1094–1096, 2018.
18. A. Guess, B. Nyhan, and J. Reifler, ”Selective Exposure to Misinfor-mation: Evidence from the Consumption of Fake News during the 2016
19. U.S. Presidential Campaign,” European Research Council, 2018.
20. S. C. Woolley and P. N. Howard, Computational Propaganda: Political Parties, Politicians, and Political Manipulation on Social Media, Oxford University Press, 2018.
21. K. Shu, D. Mahudeswaran, S. Wang, D. Lee, and H. Liu, ”FakeNewsNet: A Data Repository with News Content, Social Context, and Spatiotem-poral Information for Studying Fake News on Social Media,” Big Data, vol. 8, no. 3, 2020.
22. D. J. Flynn, B. Nyhan, and J. Reifler, ”The Nature and Origins of Misperceptions: Understanding False and Unsupported Beliefs About Politics,” Political Psychology, vol. 38, pp. 127–150, 2017.
23. M. Potthast, J. Kiesel, K. Reinartz, J. Bevendorff, and B. Stein, ”A Stylometric Inquiry into Hyperpartisan and Fake News,” ACL, 2018.
24. N. J. Conroy, V. L. Rubin, and Y. Chen, ”Automatic Deception Detec-tion: Methods for Finding Fake News,” Proceedings of the Association for Information Science and Technology, vol. 52, no. 1, 2015.
25. Y. Liu and Y. F. Wu, ”Early Detection of Fake News on Social Media Through Propagation Path Classification with Recurrent and Convolutional Networks,” AAAI, 2018.
26. J. Thorne, A. Vlachos, C. Christodoulopoulos, and A. Mittal, ”FEVER: A Large-scale Dataset for Fact Extraction and VERification,” NAACL, 2018.

27. V. Pe´rez-Rosas, B. Kleinberg, A. Lefevre, and R. Mihalcea, ”Automatic Detection of Fake News,” COLING, 2018.
28. P. R. R. de Oliveira, A. S. da Silva, and L. A. M. Pereira, ”A Comparative Study of Machine Learning Techniques for Fake News Detection,” IEEE LATINCOM, 2019.
29. S. Hochreiter and J. Schmidhuber, ”Long Short-Term Memory,” Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997.
30. A. Graves, A. Mohamed, and G. Hinton, ”Speech Recognition with Deep Recurrent Neural Networks,” ICASSP, 2013.
31. K. Sokol and P. Flach, ”Explainability Fact Sheets: A Framework for Systematic Assessment of Explainable Approaches,” ACM FAT*, 2020.
32. A. M. Kaplan and M. Haenlein, ”Users of the World, Unite! The Challenges and Opportunities of Social Media,” Business Horizons, vol. 53, no. 1, pp. 59–68, 2010.
33. R. K. Kaliyar, A. Goswami, and P. Narang, ”FakeBERT: Fake News De-tection in Social Media with a BERT-based Deep Learning Approach,” Multimedia Tools and Applications, vol. 80, pp. 11765–11788, 2021.
34. M. Granik and V. Mesyura, ”Fake News Detection Using Naive Bayes Classifier,” IEEE UKRCON, 2017.
35. F. Monti, F. Frasca, D. Eynard, D. Mannion, and M. M. Bronstein, ”Fake News Detection on Social Media using Geometric Deep Learning,” arXiv preprint arXiv:1902.06673, 2019.
36. K. Popat, S. Mukherjee, A. Yates, and G. Weikum, ”DeClarE: Debunk-ing Fake News and False Claims using Evidence-Aware Deep Learning,” EMNLP, 2018.
37. A. Khattar, J. S. Goud, M. Gupta, and V. Varma, ”MVAE: Multimodal Variational Autoencoder for Fake News Detection,” WWW, 2019.
38. J. Ma, W. Gao, P. Mitra, S. Kwon, B. J. Jansen, K. Wong, and M. Cha, ”Detecting Rumors from Microblogs with Recurrent Neural Networks,” IJCAI, 2016.
39. S. Feng, E. Banerjee, and Y. Choi, ”Syntactic Stylometry for Deception Detection,” ACL, 2012.
40. A. G. Silva, L. A. M. Pereira, and A. S. da Silva, ”A Multimodal Approach for Fake News Detection,” IEEE ISM, 2019.
41. J. Pennington, R. Socher, and C. D. Manning, ”GloVe: Global Vectors for Word Representation,” EMNLP, 2014.
42. K. Papineni, S. Roukos, T. Ward, and W. Zhu, ”BLEU: A Method for Automatic Evaluation of Machine Translation,” ACL, 2002.